Edge stabilization in the homology of graph braid groups

BYUNG HEE AN
GABRIEL C. DRUMMOND-COLE
BEN KNUDSEN

We introduce a novel type of stabilization map on the configuration spaces of a graph, which increases the number of particles occupying an edge. There is an induced action on homology by the polynomial ring generated by the set of edges, and we show that this homology module is finitely generated. An analogue of classical homological and representation stability for manifolds, this result implies eventual polynomial growth of Betti numbers. We calculate the exact degree of this polynomial, in particular verifying an upper bound conjectured by Ramos. Because the action arises from a family of continuous maps, it lifts to an action at the level of singular chains, which contains strictly more information than the homology level action. We show that the resulting differential graded module is almost never formal over the ring of edges.

20F36, 55R80, 13D40; 05C40

1 Introduction

Configuration spaces of manifolds have numerous applications in algebraic topology and homotopy theory—see Arnold [3], McDuff [22], and Cohen–Lada–May [6] for a few notable examples. More recently, there has been a growing swell of interest in configuration spaces of *graphs*. For a graph Γ—which is to say a finite, 1-dimensional cell complex—we write

$$B_k(\Gamma) = \left\{(x_1, \dots, x_k) \in \Gamma^k : x_i \neq x_j \text{ if } i \neq j\right\}_{/\Sigma_k}$$

for the kth unordered *configuration space* of Γ. First considered by Ghrist [11] from the point of view of robotics and motion planning, these spaces were shown by Abrams [1, Corollary 3.11] to be aspherical for Γ connected and so classify their respective fundamental groups, the graph braid groups of Γ. Much effort has been dedicated to understanding the homological, geometric, and combinatorial properties of these groups by a number of different researchers and groups. A non-exhaustive list includes [2, 4, 8, 9, 12, 13, 16, 18, 19, 20, 21, 25, 26, 27].

Published: XX Xxxember 20XX DOI: 10.2140/gt.20XX.XX.1001

1.1 Stability phenomena

Configuration spaces of different cardinalities relate to one another in a variety of ways, and it is often simpler to study the graded space $B(-) = \bigsqcup_{k\geq 0} B_k(-)$. For example, if M is a manifold with non-empty boundary, there is a *stabilization map*

$$\sigma : B(M) \to B(M)$$

which inserts a new particle near the boundary. This map increases the cardinality k by 1, and, for sufficiently large k, a theorem of McDuff asserts that the induced map $H_i(B_k(M)) \to H_i(B_{k+1}(M))$ is an isomorphism [22, Theorem 1.2]. One says that the configuration spaces of such manifolds exhibit *homological stability*.

One can define such a map on the configuration spaces of a graph, inserting a new particle at the end of a dangling edge, but stability almost never occurs. In order to remedy this defect, we draw inspiration from a related situation in which homological stability fails, namely that of the *ordered* configuration spaces $\mathrm{Conf}_k(M)$. Although the Betti numbers of these spaces do not stabilize, the symmetric group representations occurring in homology do. This phenomenon of *representation stability*, as formalized by Church–Ellenberg–Farb, may be summarized in the statement that the collection $\{H_*(\mathrm{Conf}_k(M))\}_{k\geq 0}$ is a *finitely generated module* over a certain combinatorial category [5, Theorem 6.2.1].

In light of these results, a philosophy has emerged that a good notion of homological stability in a given context should be the property of finite generation with respect to some naturally occurring action. For example, in hindsight, one interprets the theorem of McDuff mentioned above as a statement about the action of the polynomial ring $\mathbb{Z}[\sigma_*]$.

1.2 Edge stabilization

In this paper, we investigate a new stability phenomenon in the homology $H_*(B(\Gamma))$, which takes the form of an action by the polynomial ring generated by the edges of Γ. We prove the following analogue of the homological and representation stability enjoyed by configuration spaces of manifolds.

Theorem 1.1 *Let Γ be a graph with set of edges E. For any $i \geq 0$, the $\mathbb{Z}[E]$-module $H_i(B(\Gamma);\mathbb{Z})$ is finitely generated.*

In fact, since finite generation and presentation are equivalent over Noetherian rings, the module is finitely presented. By a theorem of Hilbert, Theorem 1.1 implies eventual polynomial growth of Betti numbers. Our next result gives the exact degree of this polynomial in terms of a certain connectivity invariant Δ^i_Γ, which is roughly the largest number of connected components of edges obtainable by removing i essential vertices from Γ—see Definition 3.1 for a precise definition. In stating the result, we assume that Γ is not a discrete graph, for which the question is trivial.

Theorem 1.2 *Let $\mathbb{F}$ be a field and Γ a graph with at least one edge. For k sufficiently large, $\dim_{\mathbb{F}} H_i(B_k(\Gamma);\mathbb{F})$ coincides with a polynomial in k of degree $\Delta^i_\Gamma - 1$.*

First introduced and exploited as a purely algebraic phenomenon by Ramos in [25, §3.2] for trees and by the authors of the present paper [2, §4.2] in general, this module structure has a simple topological origin. The *edge stabilization map*

$$\sigma_e : B(\Gamma) \to B(\Gamma)$$

acts by replacing the subconfiguration of particles on e with the collection of averages of consecutive particles or endpoints—see Figure 1 and Section 2.2 below.

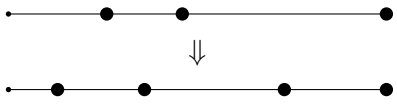

Figure 1: Edge stabilization

1.3 Edge formality

Because the various edge stabilization maps commute on the nose—not merely up to homotopy—the action of $\mathbb{Z}[E]$ on homology arises from a $\mathbb{Z}[E]$-action on singular chains. This chain level structure is surprisingly rich; indeed, we show that it almost always carries strictly more information than the action on homology.

Theorem 1.3 *Let Γ be a graph with set of edges E and R a commutative ring. The singular chain complex of $B(\Gamma)$ with coefficients in R is formal as a differential bigraded $R[E]$-module if and only if Γ is a small graph.*

Here we declare a graph to be *small* if, after smoothing as many bivalent vertices as possible, no vertex has three distinct edges—see Definition 2.3. Small graphs are very

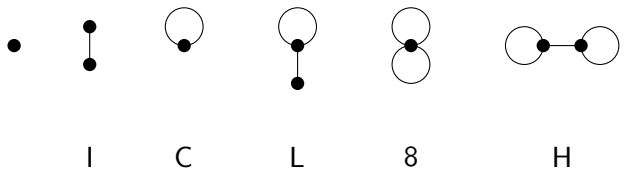

Figure 2: Homeomorphism types of connected small graphs

primitive objects; indeed, each connected component of a small graph is an isolated vertex, an interval I, a cycle C, a lollipop L, a figure-eight 8, or a handcuff H. See Figure 2.

This result is a measure of the inherent complexity of the family of graph braid groups of Γ. It also provides an explanation for this complexity, as all non-formal behavior arises from variations on the following simple example. The configuration space of two points in a star graph with three edges is homotopy equivalent to a circle. The fundamental class of this circle is represented by a cycle where the two points orbit one another by taking turns passing through the central vertex. This cycle is decomposable over the ring of edges, but the class itself is indecomposable, and non-formality follows.

1.4 Techniques and previous work

We prove our theorems by appealing to a small and explicit chain complex $S(\Gamma)$ with an action of $\mathbb{Z}[E]$—see Theorem 2.10—which functorially computes the homology and non-functorially models the singular chains of $B(\Gamma)$, both as $\mathbb{Z}[E]$-modules. In the proof of Theorem 1.2, we exploit a family of spectral sequences computing the homology of this complex—see Remark 2.14 and Section 3. These spectral sequences are an expansion of the "vertex explosion" technique introduced by the authors in [2, Definition 5.14], and we expect them to prove useful in future computations.

The Świątkowski complex $S(\Gamma)$ first arose as the cellular chain complex of a cubical deformation retract of $B(\Gamma)$ [27, §1] and was later derived independently by the authors of this paper from ideas involved in factorization homology and discrete Morse theory [2, §4.2]. We use features of both approaches to $S(\Gamma)$, so we must compare them. This comparison, which appears as Proposition 5.11, is both a key technical step and a unification of perspectives.

For the case of configuration spaces of trees, Theorem 1.2 was proven by direct computation of all Betti numbers by Maciążek and Sawicki [20, Equation (23)] and by

Ramos [25, Theorem A], who also stated the upper bound half of the general case as a conjecture.

Ramos' methods involved a phenomenon similar to edge stabilization. In his work, the action of $\mathbb{Z}[E]$ is purely algebraic, occurring at the level of discrete Morse complexes, and limited to trees. We expect the two actions on homology to coincide.

1.5 Questions

Our work invites the following questions.

(1) Is there an analogue of edge stabilization for *ordered* configuration spaces of graphs?

Lütgehetmann and Recio-Mitter recently constructed a stabilization map on ordered configuration spaces of graphs on an edge near an essential vertex [19, Proposition 5.6] using other information at the vertex. The nature of the optimal algebraic structure organizing these maps remains unclear.

(2) Are there analogues of edge stabilization for higher dimensional cells?

One conceptual explanation for edge stabilization is that the unordered configuration spaces of an edge are all contractible, which guarantees that a certain extension problem is unobstructed. In the context of both of these questions, the corresponding extension problems are obstructed. This fact explains why Lütgehetmann and Recio-Mitter's stabilization involves additional information beyond the choice of an edge. It also indicates that stabilization for higher dimensional cells would likely have to involve a similar choice of further local information or, more radically, some kind of non-local invariants.

(3) What is the significance of the invariant Δ^i_Γ? For example, is its role here connected to its appearance in the "cut polynomials" of right-angled Artin groups studied by Papadima and Suciu [23, §4.2]?

(4) Is there a simple characterization, extending Theorem 1.3 and Proposition 4.9 below, of the subsets $E_0 \subseteq E$ for which formality holds over $\mathbb{Z}[E_0]$?

(5) Our proof of Theorem 3.1 shows that the W-tori of Definition 3.11 account for a positive fraction of the homology of configuration spaces of graphs—see Remark 3.14 for further discussion of this point. Can this estimate be improved and/or extended to other families of generators?

1.6 Linear outline

The paper following the introduction is divided into five sections. In Section 2, we introduce edge stabilization and the Świątkowski complex, and we connect the two via the statement of Theorem 2.10, which implies Theorem 1.1. Section 3 is concerned with the lower and upper bounds on growth necessary to establish Theorem 1.2, and Section 4 assembles the proof of Theorem 1.3. Finally, in Sections 5 and 6, we return to complete the proof of Theorem 2.10, in the process forging a connection between the cellular models of Abrams and Świątkowski.

1.7 Conventions

Bigradings of modules are by *degree* and *weight*, and all are non-negative. The braiding isomorphism for a tensor product of modules has a sign which depends on degree and not on weight: if x and y have degree i and j, the braiding isomorphism takes $x \otimes y$ to $(-1)^{ij} y \otimes x$. We write $[m]$ for the degree shift functor by m and $\{n\}$ for the weight shift functor by n so that the degree i and weight j homogeneous component of $M[m]\{n\}$ is the degree $i-m$ and weight $j-n$ homogeneous component of M. In a differential graded context, differentials preserve weight.

The category $\mathcal{M}\mathrm{od}$ has objects pairs (R, M) where R is a weight-graded commutative unital ring and M a differential bigraded R-module. A morphism in $\mathcal{M}\mathrm{od}$ between (R_1, M_1) and (R_2, M_2) is a pair (f, g), where f is a weight-graded ring morphism from R_1 to R_2 and g is a differential bigraded R_1-module morphism from M_1 to $f^* M_2$ (that is, M_2 with the action $rm := f(r)m$).

Acknowledgements

The first two authors were supported by IBS-R003-D1. The third author was supported by NSF award 1606422. The authors would like to thank Daniel Lütgehetmann, Hyo Won Park, and Eric Ramos for illuminating conversations.

2 Edge stabilization

After establishing terminology and notation regarding graphs, we define the edge stabilization map $\sigma_e : B(\Gamma) \to B(\Gamma)$, where e is an edge of the graph Γ. At the level

of homology, these maps give rise to an action by the polynomial ring generated by the edges of Γ. In Theorem 2.10, we present a small chain complex computing this homology, together with its module structure.

2.1 A category of graphs

A *graph* Γ is a finite 1-dimensional CW complex. Its 0-cells and open 1-cells are its *vertices* and *edges*, and the set of such is denoted $V(\Gamma)$ and $E(\Gamma)$, or simply V and E, respectively. A *half-edge* is an end of an edge, and the set of such is denoted $H(\Gamma)$ or simply H. The vertices of an edge $V(e)$ are the vertices contained in the closure of that edge in Γ. The edges of a vertex $E(v)$ are the edges incident to the vertex v. The half-edges of a vertex $H(v)$ (or of an edge $H(e)$) are the half-edges incident to v (contained in e). For h in H, we write $v(h)$ and $e(h)$ for the corresponding vertex and edge.

The *valence* of a vertex v is the cardinality of $H(v)$, denoted $d(v)$. The vertex v is *isolated* if $d(v) = 0$ and *essential* if $d(v) \geq 3$. We shall sometimes write $V^{\geq 2}$ and V^{ess} for the set of vertices of valence at least 2 and the set of essential vertices, respectively. An edge with a 1-valent vertex is a *tail*. A *self-loop* at a vertex is an edge whose entire boundary is attached at that vertex.

Example 2.1 The cone on $\{1, \dots, n\}$ is a graph S_n with $n+1$ vertices, with one of valence n and n of valence 1. These graphs are called *star graphs* and the cone point the *star vertex*.

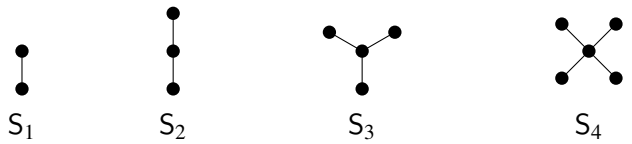

Figure 3: Star graphs

A *parametrization* of a graph Γ is a set of homeomorphisms $D_h : e(h) \to (0, 5)$ for $h \in H$ such that

(1) D_h maps the h end of $e(h)$ to the 0 end of $(0, 5)$, and

(2) if $h_1 \neq h_2 \in H(e)$, then $D_{h_2} = 5 - D_{h_1}$.

Making a choice of parametrization does not affect the homeomorphism type of configuration spaces. We will sometimes implicitly identify an edge of a parametrized graph with the interval $(0,5)$. Up to homotopy, all constructions on parametrized graphs will be independent of the choice of parametrization.

Definition 2.2 Let $f : \Gamma_1 \to \Gamma_2$ be a continuous map between graphs. We say that f is a *graph morphism* if

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

We call a graph morphism a *smoothing* if it is a homeomorphism and a *graph embedding* if it preserves vertices. A graph morphism can be factored into a graph embedding followed by a smoothing. At times, we may refer to the inverse of a smoothing as a *subdivision*, and we caution the reader that these are typically not graph morphisms.

The composite of graph morphisms is a graph morphism, and we obtain in this way a category $\mathcal{G}\mathrm{ph}$. Although the objects of $\mathcal{G}\mathrm{ph}$ are simply finite 1-dimensional CW complexes, not all morphisms are cellular. A *subgraph* is the image of a graph embedding.

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

Since graph morphisms are injective, they induce maps at the level of configuration spaces. Thus, it is natural to view $H_*(B(-))$ as a functor from the category $\mathcal{G}\mathrm{ph}$ to bigraded Abelian groups, where the weight grading records the cardinality of a configuration. As we will see in the following section, there is more structure to be found.

We close this section with the following pair of definitions premised on our notion of a smoothing.

Definition 2.3 Let Γ be a graph.

(1) We say that Γ is *smooth* if every smoothing with domain Γ is an isomorphism.
(2) We say that Γ is *small* if, in any maximal smoothing of Γ, there is no vertex with three distinct edges. Otherwise, Γ is *large*.

Smoothness is almost, but not quite, equivalent to having no bivalent vertices; indeed, the cycle C (see Figure 2) is smooth.

2.2 Topological edge action

We now introduce the promised stabilization.

Definition 2.4 Let Γ be a parametrized graph and e an edge. *Edge stabilization* at e is the map $\sigma_e : B(\Gamma) \to B(\Gamma)$ that preserves partial configurations in the complement of $(2,3) \subseteq e$ and replaces the partial configuration $\{x_1 \leq \cdots \leq x_j\} \subseteq (2,3)$ with

$$\left\{ \frac{2+x_1}{2}, \frac{x_1+x_2}{2}, \ldots, \frac{x_{j-1}+x_j}{2}, \frac{x_j+3}{2} \right\}.$$

It is a direct verification that the map σ_e is continuous and independent of parametrization up to homotopy. See Figure 5 for a depiction of this map.

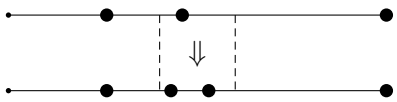

Figure 5: Edge stabilization redux

Remark 2.5 Parametrizing edges by $(0,5)$ and stabilizing in the subinterval $(2,3)$ are choices designed to interface well with the arguments of Section 5. The more obvious choices of parametrizing edges by $(0,1)$ and stabilizing in this entire interval (see Figure 1) lead to maps that differ only up to homotopy.

Remark 2.6 Stabilization by adding points to the tails of a graph has been considered [9, §5], and addition of points near the boundary of a manifold is a well-studied phenomenon—see [22, §4], for example. The existence of stabilization maps at internal edges is new, but see [25, §3.2] for a related algebraic stabilization mechanism in the context of trees.

Passing to homology, we obtain an action of the weight graded ring $\mathbb{Z}[E]$ on $H_*(B(\Gamma))$. This action is natural in the sense that a graph morphism from Γ_1 to Γ_2 induces a commutative diagram

$$\begin{array}{ccc} \mathbb{Z}[E(\Gamma_1)] \otimes H_*(B(\Gamma_1)) & \longrightarrow & \mathbb{Z}[E(\Gamma_2)] \otimes H_*(B(\Gamma_2)) \\ \downarrow & & \downarrow \\ H_*(B(\Gamma_2)) & \longrightarrow & H_*(B(\Gamma_2)). \end{array}$$

In this way, the homology of configuration spaces of graphs lifts to a functor $H_*(B(-)) : \mathcal{G}\mathrm{ph} \to \mathcal{M}\mathrm{od}$.

2.3 The Świątkowski complex

We now present a convenient chain model for $H_*(B(\Gamma))$, thought of as a functor valued in $\mathcal{M}\mathrm{od}$ via edge stabilization.

Definition 2.7 Let Γ be a graph and let R be a commutative ring. For $v \in V$, set $S(v) = \mathbb{Z}\langle \varnothing, v, h \in H(v)\rangle$. The *Świątkowski complex* of Γ (with coefficients in R) is the $R[E]$-module

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

endowed with the bigrading $|\varnothing| = (0,0)$, $|v| = |e| = (0,1)$, and $|h| = (1,1)$, together with the differential ∂ determined by setting $\partial(h) = e(h) - v(h)$.

Typically, the ring R will be $\mathbb{Z}$ or a field. When the coefficient ring is $\mathbb{Z}$, we omit it from the notation.

Remark 2.8 The following geometric heuristic is often helpful in dealing with the Świątkowski complex.

(1) We think of a generator of $S(v)$ as prescribing the local "state" of a configuration near v. Specifically, the generator $\varnothing$ represents the absence of a particle at v, the generator v represents the presence of a stationary particle at v, and the generator $h \in H(v)$ represents an infinitesimal path of a particle exiting v in the direction of h. See Figure 6.

(2) We think of a monomial in $\mathbb{Z}[E]$ as prescribing the state of a configuration on the edges of Γ; for example, the generator $e_1 e_2^2$ represents the presence of a stationary particle on e_1 and two stationary particles on e_2. See Figure 6.

(3) A basis element of the Świątkowski complex is a prescription of a global state, which is to say that we choose a local state from each $S(v)$ and populate the edges with a monomial in E.

(4) The first grading is the natural homological grading, since $h \in H$ is a path and all other generators involve stationary particles, and the second grading is given by counting the number of particles.

(5) The differential takes a path to the difference of its endpoints.

Notation 2.9 We systematically omit all factors of $\varnothing$ and all tensor symbols from the notation, and we regard half-edge generators at different vertices as permutable up to

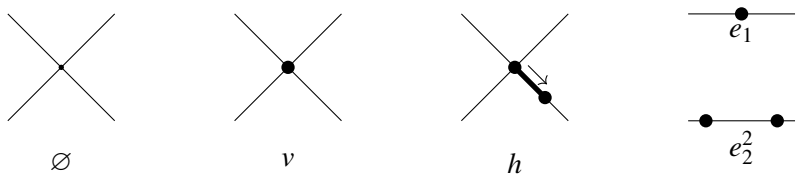

Figure 6: Local states near a vertex v and the state $e_1e_2^2$

sign. Thus, if Γ is the interval with vertices 0 and 1, half-edges h_0 and h_1, and edge e, then we write

$$\begin{aligned} e^3 \otimes \varnothing \otimes \varnothing &= e^3 \\ e^3 \otimes \varnothing \otimes h_2 &= e^3 h_2 \\ e^3 \otimes h_1 \otimes h_2 &= e^3 h_1 h_2 = -e^3 h_2 h_1. \end{aligned}$$

See Figure 7.

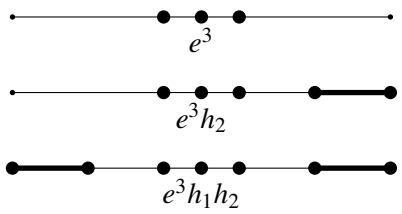

Figure 7: Three global states on the interval

A graph morphism $f : \Gamma_1 \to \Gamma_2$ determines a map $S(f;R) : S(\Gamma_1;R) \to S(\Gamma_2;R)$ as follows [2, §4.2]: an edge is sent to its image under f; for a vertex v such that $f(v)$ is also a vertex, there is an evident map $S(v) \to S(f(v))$; and, for a vertex such that $f(v)$ is an edge, we instead use the map $S(v) \to \mathbb{Z}[f(v)]$ sending $\varnothing$ to 1, v to e, and half-edges to 0. This map respects the bigrading, differential, and module structures, so we obtain a functor $S(-;R) : \mathcal{G}\mathrm{ph} \to \mathcal{M}\mathrm{od}$. Since ∂ is $R[E]$-linear by definition, this module structure descends to homology.

Theorem 2.10 *There is a natural isomorphism*

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

of functors from $\mathcal{G}\mathrm{ph}$ *to* $\mathcal{M}\mathrm{od}$.

At the level of bigraded Abelian groups, this natural isomorphism was established in [2, Theorem 4.5]—see Theorem 5.4 below. To conclude the full statement over $\mathbb{Z}$, we will check that this isomorphism is compatible with the respective $\mathbb{Z}[E]$-actions, a task which we take up in Section 5 below. The general case follows by the universal coefficients theorem.

Since $S(\Gamma)$ is finitely generated over $\mathbb{Z}[E]$ by definition, and since $\mathbb{Z}[E]$ is Noetherian, Theorem 1.1 follows immediately from Theorem 2.10.

We close this section with an introduction of a smaller variant of the Świątkowski complex, which is often more convenient.

Definition 2.11 Let Γ be a graph. For each $v \in V$, let $\widetilde{S}(v) \subseteq S(v)$ be the subspace spanned by $\varnothing$ and the differences $h_{ij} := h_i - h_j$ of half-edges. The *reduced Świątkowski complex* with coefficients in R is

$$\widetilde{S}(\Gamma; R) := R[E] \otimes_{\mathbb{Z}} \bigotimes_{v \in V} \widetilde{S}(v),$$

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

The inclusion $\widetilde{S}(\Gamma; R) \subseteq S(\Gamma; R)$ is an $R[E]$-linear quasi-isomorphism. The $R = \mathbb{Z}$ case is [2, Proposition 4.9], and the general case follows by the universal coefficients theorem. Note that $\widetilde{S}(-; R)$ is functorial for graph morphisms.

2.4 Vertex explosion and star classes

In this section, we review some tools from [2, §5] afforded by the Świątkowski complex, which will play an important role in the proofs of Theorems 1.2 and 1.3. The first of these is an exact sequence which is useful in reducing computations of $H_*(B(\Gamma))$ to computations for simpler graphs.

Definition 2.12 For $v \in V$, we write Γ_v for the *vertex explosion* of Γ at v, which is the graph obtained by

(1) replacing the vertex v with $\{v\} \times H(v)$ and
(2) modifying the attaching maps for half-edges at v by attaching h to (v, h).

There is a graph morphism from Γ_v to Γ which takes each edge to itself, takes the vertex (v, h) to $e(h)$, and takes each other vertex to itself. Defining this morphism requires choices of precisely where in $e(h)$ to send (v, h), but the isotopy class of this graph morphism is unique. See Figure 8.

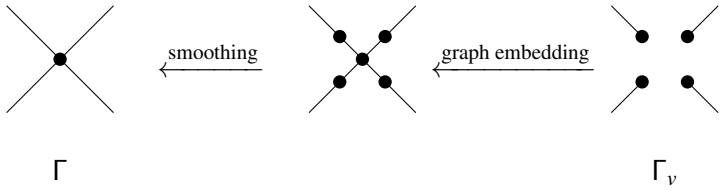

Figure 8: A local picture of vertex explosion along with an intermediate graph which admits a graph morphism from Γ_v and a smoothing to Γ

Proposition 2.13 ([2, Corollary 5.16]) *Fix a half-edge $h_0 \in H(v)$. There is a short exact sequence of differential bigraded $R[E]$-modules*

$$0 \to \widetilde{S}(\Gamma_v; R) \to \widetilde{S}(\Gamma; R) \xrightarrow{\psi} \bigoplus_{h \in H(v)\setminus\{h_0\}} \widetilde{S}(\Gamma_v; R)\,[1]\{1\} \to 0,$$

where ψ sends an element of $\widetilde{S}(\Gamma; R)$, written uniquely as $\beta + \sum_{h \in H(v)\setminus\{h_0\}} (h - h_0)\alpha_h$ with β involving no half-edges incident on v, to $(\alpha_h)_{h \in H(v)\setminus\{h_0\}}$.

The result in [2] is only taken with $\mathbb{Z}$ coefficients. Since $\widetilde{S}(\Gamma_v)$ is degreewise flat, tensoring with R preserves exactness.

In the homology long exact sequence corresponding to this short exact sequence of chain complexes, the connecting homomorphism δ from $\bigoplus H_*(\widetilde{S}(\Gamma_v; R))\{1\} \to H_*(\widetilde{S}(\Gamma_v; R))$ is given by the formula

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

Remark 2.14 This exact sequence is a degenerate example of a spectral sequence interpolating between the homology groups of the configuration spaces of Γ and those of the graph obtained by exploding a specified collection of vertices. This type of spectral sequence will play an important role in Section 3 below.

The second tool is a type of atomic homology class. Recall that the configuration space of two points in the star S_3 is homotopy equivalent to a circle.

Definition 2.15 [2, §5.1] Let $\mathsf{S}_3 \to \Gamma$ be a graph morphism. A *star class* is a class in $H_1(B_2(\Gamma); R)$ which is the image under the induced morphism of the generator α in $H_1(B_2(\Gamma); R) \simeq R$ represented by the chain $a \in S(\mathsf{S}_3; R)$

$$a := e_1(h_2 - h_3) + e_2(h_3 - h_1) + e_3(h_1 - h_2),$$

where h_i is the half-edge of e_i adjacent to the star vertex of S_3.

We refer to the chain a, and to its image, as the *standard representative* of α. See Figure 9. Given m graph morphisms $\mathsf{S}_3 \to \Gamma$ with pairwise disjoint images, we obtain a degree m homology class, called the *external product* of the star classes [2, Definition 5.10], whose standard representative is the tensor product of the standard representatives of the factors.

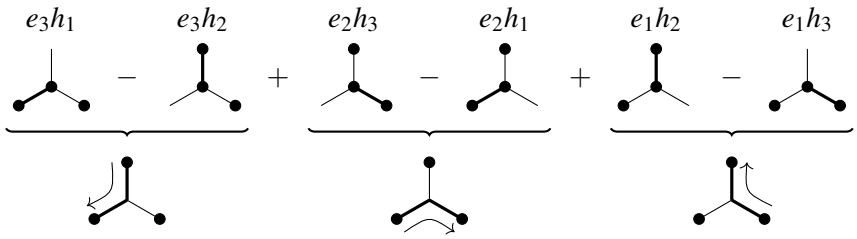

Figure 9: The standard representative of a star class

3 Growth of Betti numbers

Theorem 1.1 implies that the $\mathbb{F}$-Betti numbers of $B_k(\Gamma)$ are eventually given by a polynomial in k for any field $\mathbb{F}$ (see, eg, [7, Theorem 1.11]). In this section, we determine the exact degree of this polynomial.

3.1 Connectivity and growth

We write $D^i_\Gamma(\mathbb{F})$ for the degree of eventual polynomial growth of the $\mathbb{F}$-dimension of $H_i(B_k(\Gamma);\mathbb{F})$ in the weight k. By convention, $D^i_\Gamma(\mathbb{F}) = -\infty$ if $H_i(B_k(\Gamma);\mathbb{F}) = 0$ for $k \gg 0$.

This degree of growth is controlled by a certain elementary connectivity invariant. In order to define this invariant, we introduce the following equivalence relation.

Definition 3.1 Let Γ be a smooth graph. A subset $W \subseteq V^{\geq 2}$ determines an equivalence relation $\sim_W$ on the edges of Γ, where $e \sim_W e'$ if and only if $[e] = [e']$ in $\pi_0(\Gamma \setminus W)$. Writing $E_W = E/\sim_W$ for the set of equivalence classes, we define

$$\Delta^W_\Gamma = |E_W|\,, \qquad\qquad \Delta^i_\Gamma = \max \Delta^W_\Gamma\,.$$

where the maximum is taken over subsets W of cardinality i of $V^{\geq 2}$. We use the convention that $\Delta^i_\Gamma = -\infty$ if $i > |V^{\geq 2}|$. If Γ is not smooth, define Δ^i_Γ as $\Delta^i_{\Gamma'}$, where Γ' is a smooth graph homeomorphic to Γ.

See Figure 10.

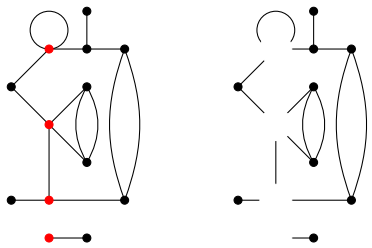

Figure 10: A graph with a choice of W in red, and the seven equivalence classes of edges under $\sim_W$

Remark 3.2 Although we work with a different set of conventions, this invariant was essentially defined by Ramos [25, page 2306], who conjectured Proposition 3.9 below.

Example 3.3 Here are some basic examples of the behavior of the invariant Δ^i_Γ.

(1) If Γ is an isolated vertex, then $\Delta^0_\Gamma = 0$.

(2) If Γ is an isolated edge, then $\Delta^0_\Gamma = 1$.

(3) If Γ is a cycle, then $\Delta^0_\Gamma = \Delta^1_\Gamma = 1$.

In these examples, $\Delta^i_\Gamma = -\infty$ for all other values of i.

(4) For $\Gamma = \Gamma_1 \sqcup \Gamma_2$ a disjoint union of (not necessarily connected) graphs,

$$\Delta^i_\Gamma = \max_{a+b=i} (\Delta^a_{\Gamma_1} + \Delta^b_{\Gamma_2}).$$

We recall the statement of Theorem 1.2.

Degree Theorem *If Γ has at least one edge, then $D^i_\Gamma(\mathbb{F}) = \Delta^i_\Gamma - 1$.*

Remark 3.4 A statement valid for an arbitrary graph is available for the sum

$$\sum_{\ell=0}^{k} \dim_{\mathbb{F}} H_i(B_\ell(\Gamma); \mathbb{F}),$$

which exhibits eventual polynomial growth of degree Δ^i_Γ.

Example 3.5 For $n \geq 3$, let Θ_n denote the Theta-n graph, which is the double cone on n points, as shown:

$\Theta_n :=$ 

It is known that

$$\dim_{\mathbb{F}} H_i(B_k(\Theta_n);\mathbb{F}) = \begin{cases} 1 & i = 0; \\ \binom{n}{2} & i = 1; \\ \binom{n}{2} - 1 + \sum_{j=0}^{2} (1-n)^j \binom{2}{j} \binom{k-j+n-1}{n-1} & i = 2. \end{cases}$$

Calculating the zeroth Betti number is trivial, the first is given in [16, Lemma 3.14], and the second follows from these and knowledge of the Euler characteristic for $B_k(\Theta_n)$, which can be extracted from [10, Theorem 2] and is given explicitly in [2, Corollary 5.11]. These Betti numbers exhibit polynomial growth of degree 0 for $i \leq 1$ and $n-1$ for $i = 2$, and these numbers are one less than respective numbers of components obtained by removing i vertices. Thus, the theorem holds in this example.

Example 3.6 For the complete graph K_4, the explicit formula for the ith Betti number $\dim_{\mathbb{F}} H_i(B_k(\mathsf{K}_4);\mathbb{F})$ in terms of k can be computed directly from [2, §5.6]. Table 1 shows the growth $D^i_{\mathsf{K}_4}(\mathbb{F})$ of the Betti number and the invariant $\Delta^i_{\mathsf{K}_4}$, verifying Theorem 1.2 for K_4.

Definition 3.7 We say that a graph Γ is *normal* if it is connected, smooth, and has an essential vertex.

The bulk of the theorem is contained in the following two results, whose proofs will occupy the rest of the section. The second is essentially Conjecture 4.3 of [25].

Proposition 3.8 *If Γ is normal, then $D^i_\Gamma(\mathbb{F}) \geq \Delta^i_\Gamma - 1$.*

Proposition 3.9 *If Γ is normal, then $D^i_\Gamma(\mathbb{F}) \leq \Delta^i_\Gamma - 1$.*

Assuming these results for the moment, we complete the proof.

Proof of Theorem 1.2 We may assume that Γ is smooth. Since the presence or absence of finitely many isolated vertices in Γ does not change the eventual growth rate, we may further assume that each connected component of Γ contains an edge. Thus,

i	$\dim_{\mathbb{F}} H_i(B_k(\mathsf{K}_4);\mathbb{F})$	$D^i_{\mathsf{K}_4}(\mathbb{F})$	$\Delta^i_{\mathsf{K}_4}$	$(\mathsf{K}_4)_W$
0	1	0	1	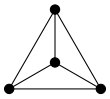
1	4 (for $k>1$)	0	1	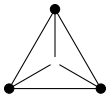
2	$6k-15$ (for $k>2$)	1	2	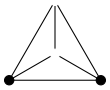
3	$4\binom{k-3}{3}$	3	4	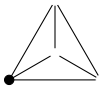
4	$\binom{k-3}{5}$	5	6	
≥ 5	0	$-\infty$	$-\infty$	

Table 1: For the complete graph K_4, a table containing the Betti number $\dim_{\mathbb{F}} H_i(B_k(\mathsf{K}_4);\mathbb{F})$, the growth rate $D^i_{\mathsf{K}_4}(\mathbb{F})$, the invariant $\Delta^i_{\mathsf{K}_4}$, and the corresponding complement $(\mathsf{K}_4)_W$

each connected component of Γ is either normal, an interval, or a cycle. The theorem is known for each of these components; indeed, Propositions 3.8 and 3.9 supply the normal case, and the other cases are classical. This observation forms the base case for an induction on $|\pi_0(\Gamma)|$.

For the inductive step, we make use of "big theta" notation [14, page 19] defined as follows:

$$f(k)=\Theta(g(k)) \iff \exists c, C>0 : c\cdot g(k)\leq f(k)\leq C\cdot g(k) \quad \forall k \gg 0.$$

Consider a decomposition of a smooth graph into two non-empty graphs $\Gamma=\Gamma_1\sqcup\Gamma_2$, and assume that the degree theorem holds for each of Γ_1 and Γ_2. By the Künneth

theorem, we have

$$\begin{aligned}\dim_{\mathbb{F}}(H_i(B_k(\Gamma);\mathbb{F})) &= \sum_{\ell=0}^{k}\sum_{a+b=i}\dim_{\mathbb{F}}(H_a(B_\ell(\Gamma_1);\mathbb{F}))\dim_{\mathbb{F}}(H_b(B_{k-\ell}(\Gamma_2);\mathbb{F}))\\ &= \sum_{a+b=i}\Theta\left(\sum_{\ell=0}^{k}\ell^{\Delta^a_{\Gamma_1}-1}(k-\ell)^{\Delta^b_{\Gamma_2}-1}\right)\\ &= \sum_{a+b=i}\Theta\left(k^{\Delta^a_{\Gamma_1}+\Delta^b_{\Gamma_2}-1}\sum_{\ell=0}^{k}\frac{1}{k}\left(\frac{\ell}{k}\right)^{\Delta^a_{\Gamma_1}-1}\left(1-\frac{\ell}{k}\right)^{\Delta^b_{\Gamma_2}-1}\right)\\ &= \sum_{a+b=i}\Theta\left(k^{\Delta^a_{\Gamma_1}+\Delta^b_{\Gamma_2}-1}\right)\int_0^1 x^{\Delta^a_{\Gamma_1}-1}(1-x)^{\Delta^b_{\Gamma_1}-1}dx,\end{aligned}$$

where the second equality uses our assumption on Γ_1 and Γ_2. By our assumption on the connected components of Γ, both $\Delta^a_{\Gamma_1}$ and $\Delta^b_{\Gamma_2}$ are nonzero for every a and b, so the integral shown is a nonnegative rational number independent of k. The largest exponent of k present in this sum is

$$\max_{a+b=i}\left(\Delta^a_{\Gamma_1}+\Delta^b_{\Gamma_2}-1\right)=\Delta^i_{\Gamma}-1,$$

as desired—see Example 3.3. □

3.2 Lower bound

To prove Proposition 3.8, we will define a special kind of homology class, called a W-torus, depending on a set of vertices W. For a suitable choice of W, we will show that a W-torus exists and generates an $\mathbb{F}[E]$-submodule with polynomial growth of the expected degree.

Throughout this section, we assume that Γ is normal. Given a set W of vertices, we write Γ_W for the graph obtained from Γ by successive explosion of each vertex of W.

Definition 3.10 A subset $W \subseteq V^{\mathrm{ess}}$ is *well-separating* if the open star of each $v \in W$ intersects more than one connected component of Γ_W.

The empty set is vacuously well-separating. We fix a field $\mathbb{F}$ and a subset $W \subseteq V^{\mathrm{ess}}$ for the remainder of this section, and we write $F_\ell H_*(B(\Gamma);\mathbb{F})$ for the filtration on homology induced by filtering $\widetilde{S}(\Gamma;\mathbb{F})$ by the number of half-edge generators at vertices of W.

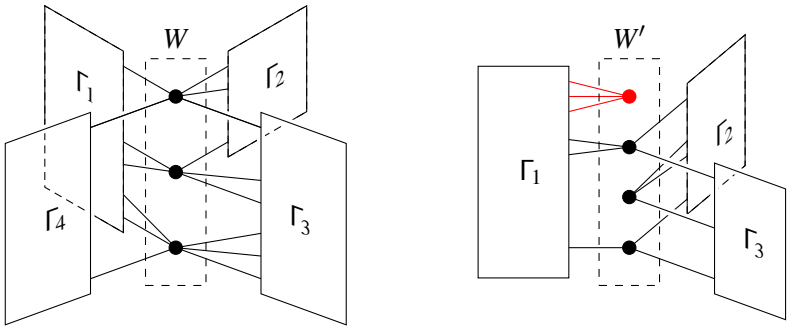

Figure 11: A well-separating subset W and a non-well-separating subset W' (assuming $\Gamma_1 \setminus W'$ is connected)

Definition 3.11 A class $\alpha \in H_{|W|}(B(\Gamma);\mathbb{F})$ is a *W-torus* if α is the external product of classes $\{\alpha_v\}_{v\in W}$, where α_v is a star class at v. We call α_v a *star factor* of α. We further say that α is *rigid* if W is well-separating and α lies in $F_{|W|}H_{|W|}(B(\Gamma);\mathbb{F})$ but not in $F_{|W|-1}H_{|W|}(B(\Gamma);\mathbb{F})$.

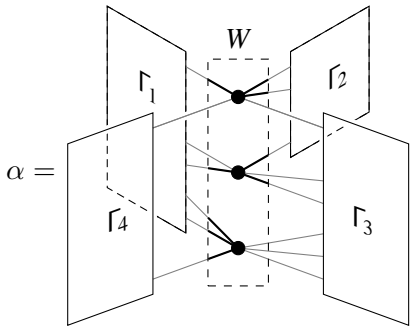

Figure 12: A W-torus which Lemma 3.15 shows to be rigid

Example 3.12 There is a unique $\varnothing$-torus, namely the class of the empty configuration in $H_0(B_0(\Gamma);\mathbb{F})$, and this class is rigid since $F_{-1}H_0(B(\Gamma);\mathbb{F}) = 0$.

Example 3.13 Consider the theta graph Θ_3, which is the double cone on three points with vertices the cone points v and w as depicted in Figure 13. According to the Θ-relation [2, Lemma 5.8(4) and Definition 5.9(4)], any star class α_v at v is equal to some star class α_w at w. We have defined rigidity only for well-separating W because if W is not well-separating in Γ, pulling the Θ-relation back along a smoothing $\Theta_3' \to \Theta_3$ and

then pushing forward along a graph embedding $\Theta_3' \to \Gamma$ shows that no W-torus could ever have a rigidity property.

We can use the same strategy to give examples of non-rigid W-tori in cases where W is well-separating. To wit, consider the graph Θ_3^+ of Figure 13. There is an evident graph morphism ι from Θ_3 to Θ_3+. The vertex set $W = \{v'\}$ is well-separating in Θ_3^+ but the star class $\iota(\alpha_v)$ is not rigid because it is equal to the star class $\iota(\alpha_w)$ by pushing forward the Θ-relation.

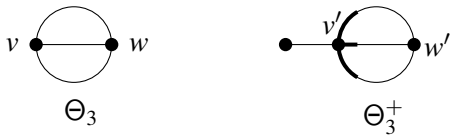

Figure 13: The theta graph Θ_3 and a graph Θ_3^+ with a well-separating vertex set $W = \{v'\}$ and a non-rigid W-torus (indicated in bold)

Remark 3.14 (Universal Generators) The definition and use of W-tori can be viewed as an outgrowth of the study of an important question in configuration spaces of graphs.

This question is to determine a *universal presentation* for the homology of configuration spaces of graphs (or of specific classes of graphs). By a universal presentation we mean

(1) a set of homology classes (*universal generators*) which generate all homology groups of configuration spaces of graphs under pushforward along graph morphisms, edge stabilization, and external products

(2) a set of *universal relations* which generate all relations among universal generators under the same three operations.

For example, the empty configuration is a universal generator for degree 0 homology while star classes and *loop classes* (see [2, Definition 5.3]) are universal generators for degree 1 homology. The Proof of Theorem 1.2 for the special case of trees in [20, Theorem V.3] amounts to showing that star classes are in fact universal generators for all homology groups for trees (it is known that this is not true for general graphs). In later work, the same authors show that the same statement holds for wheel graphs [21, §5.4].

One small step in the study of universal presentations is to ask quantitatively how much of the homology for a given graph or class of graphs is generated by such tori in general. Our proof of Proposition 3.8 implies that for any graph and in any homological degree

the amount of homology generated by such tori makes up a positive proportion of all homology as the weight increases. But this is a very crude estimate and it would be useful to have further information about this in order to know where to look for new universal generators and relations.

Set $R_W = \mathbb{F}[E_W]$ and write π_W for the projection from $\mathbb{F}[E]$ to R_W. If α is a W-torus, then the surjection $\mathbb{F}[E]\langle\alpha\rangle \to \mathbb{F}[E]\cdot\alpha$ factors through π_W, since any two edges in a connected component of Γ_W may be connected by a path of edges disjoint from the standard representative of α.

Lemma 3.15 *Let W be well-separating and let α be a W-torus. The following are equivalent.*

(1) *The W-torus α is rigid.*

(2) *The action of $\mathbb{F}[E]$ induces an isomorphism*

$$R_W\langle\alpha\rangle \xrightarrow{\simeq} \mathbb{F}[E]\cdot\alpha.$$

(3) *The standard representative of every star factor α_v involves edges in at least two distinct connected components of Γ_W.*

The third condition is a little technical but is also easier than the other two to check. In some sense it says that to guarantee rigidity, it is enough to avoid the class of "obvious" pitfalls illustrated in Example 3.13. In particular, it has the following corollary.

Corollary 3.16 *For any well-separating subset W, there is a rigid W-torus.*

Proof Choose a star class α_v at each $v \in W$. Since W is well-separating, the star class α_v may be chosen to satisfy the star factor condition (3). □

Remark 3.17 For an arbitrary set of vertices W, one could instead filter according to a maximal well-separating subset $W_0 \subseteq W$. The resulting generalized rigid W-tori do not exist for arbitrary W, $\mathbb{F}$, and Γ; however, if $\mathbb{F} = \mathbb{Z}/2\mathbb{Z}$ or if Γ is planar, then such classes exist for any W. To construct one, use Corollary 3.16 to generate a rigid W_0-torus and form the external product with a star class at each vertex in $W \setminus W_0$. Either assumption implies that this external product is nonzero, and the proof of rigidity proceeds along similar lines.

In the situation of interest, we can guarantee that W is well-separating.

Lemma 3.18 *If $W \subseteq V^{\mathrm{ess}}$ with $|W| = i$ is such that $\Delta_\Gamma^W = \Delta_\Gamma^i$, then either W is well-separating or $i = \Delta_\Gamma^i = 1$.*

Assuming these lemmas, we can establish the lower bound.

Proof of Proposition 3.8 Let $W \subseteq V^{\mathrm{ess}}$ be a subset with $|W| = i$ such that $\Delta_\Gamma^W = \Delta_\Gamma^i$. Except in the special case $i = \Delta_\Gamma^i = 1$, Lemma 3.18 guarantees that W is well-separating, so Corollary 3.16 supplies a rigid W-torus α. The dimension of $H_i(B_k(\Gamma);\mathbb{F})$ is no less than that of the submodule $\mathbb{F}[E]\cdot\alpha$, which is isomorphic to $R_W\langle\alpha\rangle$ by Lemma 3.15. Suppose α is in $H_i(B_{k_0}(\Gamma);\mathbb{F})$. Then the $\mathbb{F}$-dimension of the space of monomials in R_W of weight $k-k_0$ is exactly $\binom{\Delta_\Gamma^W+(k-k_0)-1}{\Delta_\Gamma^W-1}$ and hence $\mathbb{F}[E]\cdot\alpha$ has polynomial growth of degree $\Delta_\Gamma^W - 1$.

For the special case $i = \Delta_\Gamma^i = 1$, we must verify that $H_1(B_k(\Gamma);\mathbb{F})$ is eventually non-zero, which is well known—see [16, Theorem 3.16], for example. □

It remains to prove Lemmas 3.15 and 3.18. In the proof, we make use of the spectral sequence arising from the filtration $F_\ell H_*(B(\Gamma);\mathbb{F})$ introduced above, which we denote by $E^r_{p,q}$. This spectral sequence is a spectral sequence of $\mathbb{F}[E]$-modules.

Proof of Lemma 3.15 The unique $\varnothing$-torus is rigid, and clearly $\mathbb{F}[E]\cdot[\varnothing] \cong R_\varnothing$. In the remainder of the proof, we take $W \neq \varnothing$.

We prove first that (1) implies (2). Suppose that α is rigid, and write

$$[\alpha] = \alpha + F_{|W|-1}H_{|W|}(B(\Gamma);\mathbb{F}) \in E^\infty_{|W|,0}$$

for the resulting element in the associated graded. Since α is rigid, $[\alpha]$ cannot be zero. We wish to show that the lefthand map in the composite

$$R_W\langle\alpha\rangle \to \mathbb{F}[E]\cdot\alpha \to \mathbb{F}[E]\cdot[\alpha]$$

is an isomorphism. Since this map is surjective, it suffices to show that the composite is injective. We will show that $E^\infty_{|W|,0}$ is a submodule of a free R_W-module $R_W\langle X\rangle$ for some set X. Writing $\{[\alpha]_x\}_{x\in X}$ for the coordinates of $[\alpha]$ in this basis, the kernel of the map in question is $\bigcap_{x\in X}\mathrm{Ann}_{R_W}([\alpha]_x)$. Since $[\alpha]\neq 0$ and R_W has no zero divisors, one of the terms of this intersection is zero, implying the claim.

Our spectral sequence is concentrated in the first quadrant, so we have the containment $E^\infty_{|W|,0} \subseteq E^2_{|W|,0}$, and, since $E^1_{|W|+1,0} = 0$, we have the further containment $E^2_{|W|,0} \subseteq$

$E^1_{|W|,0}$. Both are containments of $\mathbb{F}[E]$-submodules. Since Γ has no isolated vertices, there is a $\mathbb{F}[E]$-linear isomorphism

$$E^1_{|W|,0} \cong H_0(B(\Gamma_W);\mathbb{F}) \otimes \mathbb{F}\langle X\rangle \cong R_W \otimes \mathbb{F}\langle X\rangle,$$

where X is the set of generators of $\widetilde{S}(\Gamma)$ of the form $\bigotimes_{v\in W} h^v$ with h^v a difference of half-edge generators at v. This module is the desired R_W-free module.

We prove that (2) implies (3) by proving the contrapositive. Let α_v be a star factor represented by a graph morphism $\iota : S_3 \to \Gamma$ whose image intersects only one connected component of Γ_W, and extend ι to a graph morphism $\widehat{\iota} : \Theta'_3 \to \Gamma$ with the same property, where Θ'_3 is a subdivision of Θ_3 (in other words, there is a smoothing $f : \Theta'_3 \to \Theta_3$). By pulling back via f the Θ-relation on Θ_3 as seen in Example 3.13, we have the Θ-relation on Θ'_3 as well and so $\alpha_v = \alpha_w$ for some star class α_w at w if we write w for the image of the other essential vertex of Θ'_3 under $\widehat{\iota}$. Replacing α_v by α_w in the external product defining α, we conclude that α has a representative a that does not involve v. Since W is well-separating, there are edges e_i and e_j at v lying in distinct connected components of Γ_W, and appending the generator h_{ij} to a produces an element whose boundary is $(e_i - e_j)\alpha$. Thus, the map $R_W\langle\alpha\rangle \to \mathbb{F}[E]\cdot\alpha$ has nonzero kernel.

Finally, we prove that (3) implies (1). Throughout, products and tensor products are indexed by W. At each $v \in W$, choose a basis for the half-edge generators at v such that α_v is represented by $h^v_{12}(e^v_1 - e^v_3) - h^v_{13}(e^v_1 - e^v_2)$. By our assumption (3), we may assume that e^v_1 and e^v_3 intersect distinct connected components of Γ_W. We claim that $\bigotimes h^v_{12}$ appears with a nonzero $\mathbb{F}[E]$-coefficient in every chain representative of α, implying in particular that $\alpha \notin F_{|W|-1}H_{|W|}(B(\Gamma);\mathbb{F})$, as desired.

We begin by examining the coefficient $\prod(e^v_1 - e^v_3)$ of $\bigotimes h^v_{12}$ in the standard representative a of α. We apply π_W to this coefficient to obtain an element of R_W. By construction, $\pi_W(e^v_1 - e^v_3) \neq 0$ for each v. Choose an ordering of $\pi_0(\Gamma_W)$ and write $\pi_W(e^v_1 - e^v_3) = \pm(\Delta^v_i - \Delta^v_j)$ with $\Delta^v_i < \Delta^v_j$ in the ordering. The leading term $\prod \Delta^v_i$ appears with coefficient ± 1 in the polynomial $\pi_W(\prod e^v_1 - e^v_3)$, so $\prod(e^v_1 - e^v_3) \neq 0$, establishing the claim for the standard representative a.

Suppose now that $a - a' = \partial c$. In our preferred basis for the tensor product, c is necessarily a sum of elementary tensors of $(|W| + 1)$ half-edge generators, some of which are of the form

$$p(E)h_{ij} \otimes \bigotimes h^v_{12} \tag{1}$$

for some half-edge generator h_{ij} at a vertex away from W. Terms not of this form contain strictly fewer than $|W|$ tensor factors of the form h^v_{12}, and the boundary of such

a term is a sum of terms containing strictly fewer than $|W|$ tensor factors of the form h_{12}^v. On the other hand, the boundary of (1) is a sum of terms containing strictly fewer than $|W|$ tensor factors of the form h_{12}^v, together with the term

$$p(E)(e_i - e_j) \bigotimes h_{12}^v.$$

By our assumption on h_{ij}, e_i and e_j lie in the same connected component of Γ_W, so $\pi_W(p(E)(e_i - e_j)) = 0$. Thus, the coefficient of $\bigotimes h_{12}^v$ in a' coincides with the coefficient in a modulo $\ker \pi_W$, so this coefficient is nonzero by the previous paragraph. □

Remark 3.19 The same techniques serve to establish a version of Lemma 3.15 for toric classes with some factors given by loop classes.

Proof of Lemma 3.18 If $|W| = 0$, or if $|W| = 1$ and $\Delta_\Gamma^W > 1$, then the claim is obvious. If $|W| > 1$ and W is not well-separating, then there is a vertex $v \in W$ whose open star intersects only one connected component of Γ_W. We claim that there is a second vertex $v' \in V^{\mathrm{ess}} \setminus W$ lying in this component of Γ_W and sharing an edge e with some vertex of $W \setminus \{v\}$.

Given such a vertex, we replace v with v' in W to obtain a set W' with $\Delta_\Gamma^{W'} \geq \Delta_\Gamma^W + 1$. Indeed, the relations $\sim_W$ and $\sim_{W\setminus\{v\}}$ are identical, but the equivalence class of e under the relation $\sim_W$ is split into at least two distinct equivalence classes under $\sim_{W'}$. Thus, Δ_Γ^W does not achieve the maximum value Δ_Γ^i.

Assume for contradiction that such a v' does not exist, and suppose there is a reduced edge path in Γ beginning at v, terminating at some other vertex of W, and passing through no third vertex of W. Such a path is necessarily contained in the connected component of Γ_W containing the open star of v. Since this path is reduced and Γ has no bivalent vertices, the path passes only through essential vertices. We obtain a contradiction of the non-existence of v' unless all such paths have length one; in other words, every edge at v terminates either at a vertex of W or a vertex of valence 1. Therefore, either the valence of v is 1, or the open star of v intersects more than one component of Γ_W. Either is a contradiction, so we conclude that no edge path exists between v and any other vertex of W. Since $|W| > 1$, it follows that Γ is disconnected, a contradiction. □

3.3 Upper bound

Let Γ be a graph and $\mathbb{F}$ be a field. Given a subset $W \subseteq V^{\mathrm{ess}}$ and a vertex v, we write $E_W(v) \subseteq E_W$ for the set of equivalence classes of edges edges adjacent to $v \in V$. We set

$$\widetilde{S}_W(\Gamma; \mathbb{F}) := \widetilde{S}(\Gamma_W; \mathbb{F}) \otimes_{\mathbb{F}[E]} \mathbb{F}[E_W].$$

The action of $\mathbb{F}[E]$ on $\mathbb{F}[E_W]$ is weight-respecting so the complex $\widetilde{S}_W$ is weight-graded. Since $\widetilde{S}_\varnothing$ is simply the reduced Świątkowski complex over $\mathbb{F}$, Proposition 3.9 is a special case of the following result.

Proposition 3.20 *Let Γ be smooth and contain no isolated vertices. For any $W \subseteq V^{\mathrm{ess}}$, the dimension of $H_i(\widetilde{S}_W(\Gamma;\mathbb{F}))$ is eventually polynomial in the weight of degree at most $\Delta^i_\Gamma - 1$.*

Remark 3.21 Geometrically, the weight k subcomplex of $\widetilde{S}_W(\Gamma)$ corresponds to the space of configurations of k points in Γ which are permitted to collide at vertices in W—see [4, §2] and [24, §1.1].

Note in particular that the behavior of $\widetilde{S}_W(\Gamma)$ is different from that of $\widetilde{S}(\Gamma_W)$. In both cases, there is no local homological information near W, but particles pass freely through vertices in W in the former while in the latter they avoid W. This is easiest to see in the star graph S_n, where the former models a contractible space and the latter a space with n contractible components.

The strategy of the proof of Proposition 3.20 will be to show that the desired growth rate is already achieved on the E^2 page of a certain spectral sequence converging to the homology of $\widetilde{S}_W(\Gamma;\mathbb{F})$. In order to introduce this spectral sequence, we require the following notation. Given a vertex $v \in V \setminus W$, we write S_v for a star graph of maximal valence equipped with a graph morphism $\iota : \mathsf{S}_v \to \Gamma$ sending the star vertex to v. Setting $R_{W,v} = \mathbb{F}[E_W(v)]$, we write $C_{W,v}$ for the chain complex

$$C_{W,v} := \widetilde{S}(\mathsf{S}_v;\mathbb{F}) \otimes_{\mathbb{F}[E(\mathsf{S}_v)]} R_{W,v}$$

of $R_{W,v}$-modules. The chain complex $C_{W,v}$ is finitely generated over $R_{W,v}$ because $\widetilde{S}(\mathsf{S}_v;\mathbb{F})$ is finitely generated over $\mathbb{F}[E(\mathsf{S}_v)]$.

Lemma 3.22 *Let v be an essential vertex of $V \setminus W$ in Γ. There is a convergent, homological, weight-graded spectral sequence of $\mathbb{F}[E_W]$-modules*

$$E^2_{p,q} \cong H_p\left(\widetilde{S}_W(\Gamma_v;\mathbb{F}) \otimes_{R_{W,v}} H_q(C_{W,v})\right) \implies H_{p+q}(\widetilde{S}_W(\Gamma;\mathbb{F})).$$

Proof We write the differential of $\widetilde{S}_W(\Gamma;\mathbb{F})$ as $\partial = \partial_v + \partial_-$, where ∂_v is the sum of the terms of the differential involving half-edges at v. These two operators square to zero individually and commute, giving $\widetilde{S}_W(\Gamma;\mathbb{F})$ the structure of a bicomplex of $\mathbb{F}[E_W]$-modules. The desired spectral sequence is the spectral sequence of this bicomplex with

zeroth differential ∂_v and first differential ∂_-—see Figure 14. Since this bicomplex is concentrated in finitely many bidegrees, the spectral sequence collapses and in particular converges.

Using the decomposition $\widetilde{S}(\Gamma) \cong \widetilde{S}(\Gamma_v) \otimes_{\mathbb{Z}[E(v)]} \widetilde{S}(\mathsf{S}_v)$, we obtain the isomorphism

$$E^0 \cong \left(\widetilde{S}_W(\Gamma_v;\mathbb{F}) \otimes_{R_{W,v}} C_{W,v},\ 1 \otimes \partial\right)$$

of trigraded $\mathbb{F}[E_W]$-modules, where the homological bidegree on the righthand side is the natural bigrading of the tensor product. Since $\widetilde{S}_W(\Gamma_v;\mathbb{F})$ is $\mathbb{F}[E_W]$-free, and hence $R_{W,v}$-free, the Künneth isomorphism gives

$$E^1 \cong \left(\widetilde{S}_W(\Gamma_v;\mathbb{F}) \otimes_{R_{W,v}} H_*(C_{W,v}),\ \partial \otimes 1\right),$$

completing the proof. □

$$\begin{array}{ccccccccc}
E^0_{0,1} & \xleftarrow{d^1} & \cdots & \xleftarrow{d^1} & E^0_{p-1,1} & \xleftarrow{d^1} & E^0_{p,1} & \xleftarrow{d^1} \cdots \xleftarrow{d^1} & E^0_{|V|,1} \\
\downarrow{\scriptstyle d^0} & & & & \downarrow{\scriptstyle d^0} & & \downarrow{\scriptstyle d^0} & & \downarrow{\scriptstyle d^0} \\
E^0_{0,0} & \xleftarrow{d^1} & \cdots & \xleftarrow{d^1} & E^0_{p-1,0} & \xleftarrow{d^1} & E^0_{p,0} & \xleftarrow{d^1} \cdots \xleftarrow{d^1} & E^0_{|V|,0}
\end{array}$$

Figure 14: The potentially nonzero entries of the E^0 page of the spectral sequence for the bicomplex $(\widetilde{S}_W(\Gamma;\mathbb{F}), \partial_v, \partial_-)$

The key to the desired growth estimate is the following technical lemma.

Lemma 3.23 *Let (M, ∂_M) be a differential bigraded $R_{W,v}$-module and N a bigraded $R_{W,v}$-module with N finitely generated and concentrated in strictly positive homological degrees. If $\dim_{\mathbb{F}} H_j(M)$ is eventually polynomial in weight of degree at most d for every $0 \leq j < i$, then $\dim_{\mathbb{F}} H_i(M \otimes_{R_{W,v}} N, \partial_M \otimes 1)$ is as well.*

We assume this result for the moment.

Proof of Proposition 3.20 Let Γ be smooth with no isolated vertices. We proceed by induction on $|V^{\mathrm{ess}} \setminus W|$.

For the base case $W = V^{\mathrm{ess}}$, the complex $\widetilde{S}_W(\Gamma;\mathbb{F})$ is isomorphic to $\mathbb{F}[E_W]$ concentrated in degree 0. In weight k, the dimension of this vector space is the number of ways of putting k indistinguishable balls in E_W distinct bins, which is polynomial of degree $E_W - 1 = \Delta_\Gamma^{|W|} - 1$, as desired.

For the inductive step, suppose the statement has been shown true for $|V^{\mathrm{ess}} \setminus W| < r$ and let $|V^{\mathrm{ess}} \setminus W| = r > 0$. Fixing a vertex v in $V^{\mathrm{ess}} \setminus W$, it will suffice to show that the $E^2_{i,0}$ and $E^2_{i-1,1}$ entries of the spectral sequence of Lemma 3.22 each have eventual polynomial growth of degree at most $\Delta^i_\Gamma - 1$ (since $C_{W,v}$ is concentrated in degrees 0 and 1, these are the only nonzero entries in the appropriate degree on the E^2 page).

We note first that $H_0(C_{W,v})$ is one-dimensional in each weight, since there is a degree 1 chain interpolating between any two edges of S_v. Thus,

$$\begin{aligned} E^2_{i,0} &\cong H_i(\widetilde{S}_W(\Gamma_v;\mathbb{F}) \otimes_{R_{W,v}} \mathbb{F}) \\ &\cong H_i(\widetilde{S}\left(\Gamma_{W\sqcup\{v\}};\mathbb{F}\right) \otimes_{\mathbb{F}[E]} \mathbb{F}[E_W] \otimes_{R_{W,v}} \mathbb{F}) \\ &\cong H_i(\widetilde{S}_{W\sqcup\{v\}}(\Gamma;\mathbb{F})), \end{aligned}$$

and the latter has polynomial growth of degree at most $\Delta^i_\Gamma - 1$ by induction (we have used the isomorphism $\mathbb{F}[E_{W\sqcup\{v\}}] \cong \mathbb{F}[E_W] \otimes_{R_{W,v}} \mathbb{F}$).

Next, we have $E^2_{i-1,1} \cong H_i(M \otimes_{R_{W,v}} N)$, where $M = \widetilde{S}_W(\Gamma_v;\mathbb{F})$ and $N = H_1(C_{W,v})$. As previously noted, $C_{W,v}$ is finitely generated over the Noetherian ring $R_{W,v}$, so N is finitely generated, as well as concentrated in strictly positive degrees by definition. Moreover, M is degreewise free and thus degreewise flat. Since Γ was smooth and contained no isolated vertices, the same is true of Γ_v. Then the inductive hypothesis guarantees that, for $j < i$, the homology group $H_j(M)$ is eventually polynomial in the weight of degree at most

$$\Delta^j_{\Gamma_v} - 1 \leq \Delta^{j+1}_\Gamma - 1 \leq \Delta^i_\Gamma - 1,$$

so Lemma 3.23 implies that $\dim_{\mathbb{F}} E^2_{i-1,1}$ has the same property. □

We conclude with the proof of the technical lemma.

Proof of Lemma 3.23 Up to associated graded, the homology group $H_i(M \otimes_{R_{W,v}} N)$ is the sum of the ith anti-diagonal on the E^∞ page of a first-quadrant (hence convergent) Künneth spectral sequence. Thus, it suffices to bound the growth of the ith antidiagonal of

$$E^2 \cong \mathrm{Tor}_{R_{W,v}}(H_*(M), N).$$

To calculate these Tor groups, we resolve N. Since $R_{W,v}$ is Noetherian of finite global dimension, and since N is finitely generated, there is a finite length resolution P of N by finitely generated projective $R_{W,v}$-modules. By the Quillen–Suslin theorem, P is in fact a free resolution, generated by a trigraded $\mathbb{F}$-vector space $G = \bigoplus G_{a,b}\{k\}$ in which every summand is finite dimensional and almost all vanish. The trigraded

complex $H_*(M) \otimes_{R_{W,v}} P$ computes the desired (weight-graded) Tor groups. Since N is concentrated in strictly positive homological degrees, P is as well, so the ith antidiagonal of this complex is isomorphic to

$$\bigoplus_{j=0}^{i-1} \bigoplus_{a+b=i-j} \bigoplus_{k} H_j(M) \otimes_{\mathbb{F}} G_{a,b}\{k\},$$

which is a finite sum of vector spaces having eventual polynomial growth of degree at most d. □

4 Edge formality

We now undertake the in-depth study of one aspect of the chain level $R[E]$-module structure induced by edge stabilization, namely the question of its formality.

4.1 Edge formality

Recall that a differential bigraded A-module is said to be *formal* if it is connected to its homology by a finite zig-zag of (bigraded) quasi-isomorphisms:

$$(M,d) \xleftarrow{\simeq} (M_1,d_1) \xrightarrow{\simeq} (M_2,d_2) \xleftarrow{\simeq} \cdots \xrightarrow{\simeq} (H(M,d),0).$$

Definition 4.1 We say that a graph Γ is *edge formal* over the commutative ring R if the the singular chain complex of $B(\Gamma)$ with coefficients in R is formal as a differential bigraded $R[E]$-module.

The goal of this section is to prove Theorem 1.3, which gives a complete characterization of edge formal graphs independent of coefficient ring. We first recall the statement of this theorem—see Definition 2.3 for a reminder on terminology.

Formality Theorem *A graph is edge formal if and only if it is small.*

In proving this theorem, we may replace the large and unwieldy complex of singular chains with the smaller complex $S(\Gamma;R)$; indeed, as we will see below in Corollary 5.12, this $R[E]$-module is quasi-isomorphic to the $R[E]$-module of singular chains with coefficients in R. Moreover, we may work interchangeably with the complexes $S(\Gamma;R)$ and $\widetilde{S}(\Gamma;R)$. Since the underlying bigraded $R[E]$-module of $S(\Gamma;R)$ (resp. $\widetilde{S}(\Gamma;R)$) is free, we may use the following criterion, whose proof is a standard argument in homotopical algebra.

Lemma 4.2 *Let A be a weight-graded commutative ring and M a differential bigraded A-module which is projective in each homological degree. Then M is formal if and only if there is a map of differential bigraded A-modules $M \to H_*(M)$ inducing the identity on homology.*

Thus, Theorem 1.3 amounts to characterizing the (non-)existence of such maps.

Remark 4.3 We caution the reader that the notion of an edge formal graph has no direct connection with that of a formal topological space in the sense of rational homotopy theory. For partial results about this other kind of formality, see [15, Theorems 1.2 and 1.3], where a necessary and sufficient criterion is given for $B_4(\Gamma)$ to be formal as a space.

4.2 Small graphs are formal

We begin with a classification of smooth small graphs—see Figure 2.

Lemma 4.4 *If Γ is a small graph that is both connected and smooth, then Γ is isomorphic to an isolated vertex, an interval I, a cycle C, a lollipop L, a figure-eight 8, or a handcuff H.*

Proof Any vertex of a small graph must be either isolated, 1-valent, 2-valent (with either two distinct edges or a self-loop), 3-valent, with one self-loop, or 4-valent, with 2 self-loops. A connected graph with a vertex which is isolated, 2-valent with a self-loop, or 4-valent with 2 self-loops is necessarily an isolated vertex, the cycle graph C, or the figure-eight graph 8, respectively. A graph containing a 2-valent vertex with two distinct edges cannot be smooth. Then in any remaining case, every vertex must be 1-valent or 3-valent with one self-loop. A connected graph all of whose vertices are of these two kinds must have precisely two vertices; then there are three cases, namely the interval I, the lollipop graph L, and the handcuff graph H. □

This gives us most of what we need, since smoothings reflect formality.

Lemma 4.5 *Let $f : \Gamma \to \Gamma'$ be a smoothing. If Γ' is edge formal over R, then so is Γ.*

Proof The hypothesis guarantees the existence of the middle arrow in the diagram

$$S(\Gamma;R) \xrightarrow{S(f;R)} S(\Gamma';R) \dashrightarrow H_*(S(\Gamma';R)) \xleftarrow{H_*(S(f;R))} H_*(S(\Gamma;R))$$

of quasi-isomorphisms of $R[E(\Gamma)]$-modules, where $R[E(\Gamma)]$ acts on the middle two entries by restriction. □

We can also eliminate self-loops by making choices.

Lemma 4.6 *Let Γ be a graph with a self-loop e at the vertex v. Let Γ_- be the graph obtained by replacing e with an edge which is not a loop to a new vertex. Then the reduced Świątkowski complex $\widetilde{S}(\Gamma;R)$ (unnaturally) decomposes as a direct sum of differential graded $R[E]$-modules as*

$$\widetilde{S}(\Gamma_v;R)[1]\{1\} \oplus \widetilde{S}\left(\Gamma_-;R\right) \cong \widetilde{S}(\Gamma;R).$$

In particular, Γ is edge formal over R if and only if both Γ_v and Γ_- are so.

Proof Let h_1 and h_2 be the two half-edges of e. Write all half-edge generators at v as differences $h_{1i} = h_1 - h_i$ We obtain a direct sum decomposition of $\widetilde{S}(\Gamma;R)$ by grouping generators (using this choice of presentation at v) according to whether or not they involve h_{12}. Since h_{12} is closed, this decomposition respects the differential, and it respects the $R[E]$-action by inspection.

Now, there is a graph morphism ι from Γ_- to Γ taking the half-edge of e to h_1, and we obtain a map $(\alpha,\beta) \mapsto h_{12}\alpha + \iota_*(\beta)$. It is an immediate verification that this map realizes the direct sum decomposition. □

Removing a vertex or turning a self-loop to an edge as in this lemma preserves smallness.

Proof of Theorem 1.3, "if" direction By Lemmas 4.5 and 4.6, we may assume that Γ is a smooth graph without self-loops. By Lemma 4.4, Γ is a disjoint union of isolated vertices and intervals. But then $\widetilde{S}(\Gamma;R)$ has no differential, and there is nothing to prove. □

4.3 Paradoxically decomposable cycles

We will derive the "only if" direction of Theorem 1.3 from the existence of certain cycles that we dub *paradoxically decomposable*.

Definition 4.7 Let A be a commutative ring, $I \subseteq A$ an ideal, M a differential graded A-module, and $b \in M$ a cycle.

(1) We say that b is *decomposable* with respect to I if $b \in IM$.

(2) We say that b is *paradoxically decomposable* with respect to I if

  (a) b is decomposable with respect to I in M, and

(b) $[b]$ is *not* decomposable with respect to I in $H_*(M)$.

Lemma 4.8 *Let M be a differential graded A-module that is projective in each degree, $I \subseteq A$ an ideal, and $b \in M$ a cycle. If b is paradoxically decomposable with respect to I, then M is not formal.*

Proof Assume that M is formal. Then, by our hypothesis on M and Lemma 4.2, there is a map $f : M \to H_*(M)$ of A-modules inducing the identity on homology. Since b is decomposable with respect to I, we may write $b = \sum r_i b_i$ with $r_i \in I$, whence $f(b) = \sum r_i f(b_i)$; thus, $f(b)$ is decomposable with respect to I. But b is a cycle and f induces the identity on homology, so $f(b) = [b]$, a contradiction. □

We briefly detour from our main discussion to present the following result, which is of independent interest, particularly from the point of view of [2, Theorem 3.20].

Proposition 4.9 *Let Γ be a graph with tails e_1, e_2, and e_3 in the same connected component, and let R be $\mathbb{Z}$ or $\mathbb{Z}/2\mathbb{Z}$. Then $S(\Gamma; R)$ is not formal over $R[e_1, e_2, e_3]$.*

Proof We will show that the $S(\Gamma; R)$ has a parodoxically decomposable cycle with respect to the ideal I generated by e_1, e_2, and e_3.

Because e_1, e_2, and e_3 lie in the same connected component, there is a subgraph $\Gamma_0 \subseteq \Gamma$ containing these three edges and admitting a smoothing $\Gamma_0 \to \mathsf{S}_3$. We will show that the corresponding star class $\alpha \in H_*(B(\Gamma_0); R) \cong H_*(B(\mathsf{S}_3); R)$ is not decomposable with respect to I, which will imply the claim, since a decomposable representative for this class is given by $\sum p_{ij} e_k$, where p_{ij} is the unique element in $S(\Gamma_0; R)$ with $\partial p_{ij} = e_i - e_j$.

Recall the canonical homomorphism $\sigma : H_1(B_2(\Gamma); R) \to \mathbb{Z}/2\mathbb{Z}$ that records the permutation of the endpoints of a braid. Since α is a star class, $\sigma(\alpha) = 1$. On the other hand, $\sigma(e\gamma) = 0$ for any $e \in \{e_1, e_2, e_3\}$ and $\gamma \in H_1(B_1(\Gamma); R)$ because e is a tail and thus is not part of the simple closed curve representing γ. Then the assumption that α is decomposable with respect to I implies the contradiction $\sigma(\alpha) = 0$. □

We expect this statement is true with arbitrary coefficients.

Next, we will use paradoxically decomposable cycles to establish non-formality in two basic cases, from which we will deduce the statement for a general large graph. For the rest of this section, we shall only consider decomposability with respect to the ideal generated by E.

We write N for the number of essential vertices of Γ. As is well known, the homology of $B_k(\Gamma)$ vanishes above degree N; indeed, the Świątkowski complex itself vanishes in these degrees (see [11, Theorem 3.3] for an earlier argument). Moreover, this top homology is nonzero, since there can be no boundaries in degree N, and since the external product of standard star class representatives at the essential vertices of Γ is a nonzero cycle in the Świątkowski complex.

Lemma 4.10 *Let Γ be a graph with no self-loops all of whose vertices are of valence 1 or 3. If Γ has a vertex of valence 3, then Γ is not edge formal.*

Proof By [2, Proposition 5.25], $H_N(B_k(\Gamma); R)$ is 0 for $k < 2N$ and is one dimensional for $k = 2N$, spanned by a W-torus, where W is the set of trivalent vertices. This class is not decomposable, since there is nothing in lower weight and degree N, but its standard representative is decomposable. The claim now follows from Lemma 4.8. □

For the second result, we require the following preliminary notion:

Definition 4.11 We say that a vertex v is *simple* if there are no self-loops at v and no vertex w with multiple edges between v and w.

Lemma 4.12 *If Γ has a simple essential vertex, then Γ is not edge formal.*

Proof Let v be a simple essential vertex of Γ, and assume without loss of generality that Γ has no vertices of valence 2. Then $H_N(B(\Gamma); R) \neq 0$, so $H_{N-1}(B(\Gamma_v); R) \neq 0$ by Proposition 2.13. Choose a nonzero class β_v in the group $H_{N-1}(B_k(\Gamma_v); R)$ with k chosen minimally so that this group is nonzero, and write β for the external product of β_v and a star class at v. Since β is represented by a nonzero cycle b in top degree, we conclude that $\beta \neq 0$. There are no boundaries in top degree, so the cycle b is necessarily unique, and it is decomposable, since the standard representative of the star class is so. Thus, it suffices to show that β is not decomposable.

Applying Proposition 2.13 at v, we obtain the following piece of the exact sequence

$$0 \to H_N(B_{k+1}(\Gamma); R) \to \bigoplus^{\ell-1} H_{N-1}(B_k(\Gamma_v; R)) \xrightarrow{\delta} H_{N-1}(B_{k+1}(\Gamma_v; R))$$

where $\ell := d(v)$. Since β is a nonzero element of $H_N(B_{k+2}(\Gamma); R)$, the assumption of decomposability is the assumption that β is a nontrivial $R[E]$-linear combination of elements from $H_N(B_{k+1}(\Gamma); R)$. Therefore it will suffice to show that $H_N(B_{k+1}(\Gamma); R)$ vanishes, ie, that the connecting homomorphism δ is injective. Since Γ_v has $N - 1$

essential vertices, this degree is the top homological degree for $B_k(\Gamma_v)$, so homology and cycles coincide. It follows that the kernel of δ consists of tuples of cycles $(b_2, \dots, b_\ell)$ in the weight N summand of $S_{N-1}(\Gamma_v; R)$ such that

$$\sum_{j=2}^{\ell} b_j(e_j - e_1) = 0$$

where e_j is the jth edge incident on v. The e_j are distinct by the simplicity of v. We may assume without loss of generality that $b_2 \neq 0$.

Let $b_{j,1}$ be the chain obtained from b_j by replacing e_ℓ with e_1. Since the differential is $R[E]$-linear, each $b_{j,1}$ is again a cycle. We claim further that $b_{j,1}$ is nonzero if b_j was. To see this, we note that the difference $b_j - b_{j,1}$ is divisible by $(e_\ell - e_1)$, so that, if $b_{j,1} = 0$, we may conclude that b_j is the product of $(e_\ell - e_1)$ and a *cycle* of top degree in weight $k-1$, which is necessarily not a boundary. Since k was chosen to be minimal with respect to the existence of such a cycle, this is a contradiction.

Thus, the sum $\sum_{j=2}^{\ell-1} b_{j,1}(e_j - e_1)$ vanishes modulo $(e_\ell - e_1)$ and so vanishes, since none of the terms contain e_ℓ by construction. Applying this procedure repeatedly, we obtain a nonzero cycle $b_{2,\ell-2}$ with $b_{2,\ell-2}(e_2 - e_1) = 0$, a contradiction. □

4.4 Large graphs are not formal

In order to reduce the case of a general large graph to the cases already considered, we will make use of the following device:

Definition 4.13 Let M_1 be a differential graded A_1-module and M_2 be a differential graded A_2-module. An (A_1, A_2)-*retraction* of M_2 onto M_1 consists of

- a retraction of rings $A_1 \xrightarrow{\iota} A_2 \xrightarrow{\pi} A_1$, and
- a retraction of $\mathbb{Z}$-modules $M_1 \xrightarrow{i} M_2 \xrightarrow{p} M_1$

where i is A_1-linear with respect to the A_1-module structure on M_2 induced by ι and p is A_2-linear with respect to the A_2-module structure on M_1 induced by π.

The relation of this notion to questions of formality is the following:

Lemma 4.14 *Let M_1 and M_2 be differential graded A_1- and A_2-modules, respectively, and suppose that there exists an (A_1, A_2)-retraction of M_2 onto M_1. If M_2 is projective in each degree and formal over A_2, then M_1 is formal over A_1.*

Proof Our hypotheses on M_2 and Lemma 4.2 guarantee the existence of a map $M_2 \to H_*(M_2)$ of differential graded A_2-modules inducing the identity on homology. It follows that the map in homology induced by the composite

$$M_1 \to M_2 \to H_*(M_2) \to H_*(M_1)$$

is $H_*(p) \circ \mathrm{id}_{H_*(M_2)} \circ H_*(i) = \mathrm{id}_{H_*(M_1)}$, so it suffices to check that this map is A_1-linear, which follows from a diagram chase. □

Definition 4.15 Let Γ and Γ' be graphs. An *algebraic retraction* (over R) from Γ' to Γ is a $(R[E(\Gamma)], R[E(\Gamma')])$-retraction of $\widetilde{S}(\Gamma'; R)$ onto $\widetilde{S}(\Gamma; R)$. In the presence of an algebraic retraction, we say that Γ is an *algebraic retract* of Γ' (over R).

The composition of algebraic retractions is an algebraic retraction. An algebraic retraction need not arise from a topological retraction between configuration spaces.

The "only if" direction of Theorem 1.3 is an immediate consequence of the following result in combination with Lemmas 4.10, 4.12, and 4.14 (see Figure 15 for an explanation of unfamiliar terminology).

Lemma 4.16 *If Γ is a smooth large graph, then Γ has one of the following graphs as an algebraic retract:*

(1) *a graph with a simple essential vertex,*
(2) *the theta graph Θ_3, or*
(3) *the A graph* A.

Figure 15: The theta graph and the A graph

The proof of this result relies on the fact that algebraic retractions arise geometrically as the result of *surgeries* on graphs (see Figure 16).

Lemma 4.17 *Let Γ and Γ' be graphs, and suppose that Γ is obtained from Γ' by replacing a connected subgraph Δ, attached to the rest of Γ only at the vertices $v_1 \neq v_2$, with a single edge e_0 between v_1 and v_2, ie,*

$$\Gamma' = (\Gamma \setminus e_0) \coprod_{v_1 \cup v_2} \Delta$$

Then Γ is an algebraic retract of Γ' over any R.

We call this operation a *surgery*. Assuming this result momentarily, we complete the

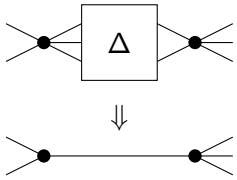

Figure 16: Depiction of a surgery

proof of Theorem 1.3.

Proof of Lemma 4.16 A connected component is an algebraic retract and largeness is a local property so it suffices to consider the connected case. Suppose Γ has a vertex of valence strictly greater than three with a self-loop. The direct sum decomposition of Lemma 4.6 gives an algebraic retraction from Γ to the graph Γ_- obtained by replacing a self-loop with a tail. The graph Γ_- is smooth, connected, and large if Γ was. Suppose Γ has a vertex of valence three with a self-loop. There is a surgery replacing the closed star of the vertex with a tail which preserves smoothness and largeness. Then iteratively we may assume Γ contains no self-loops.

Suppose that Γ has a pair of vertices sharing three edges. We obtain an algebraic retraction onto Θ_3 using Lemma 4.17 with Δ the complement in Γ of two of these (open) edges. See Figure 17(a). Suppose that Γ has a pair of vertices v and w sharing two edges e_1 and e_2. Since Γ is smooth, each of v and w has a third edge. If there is a path of edges disjoint from e_1 and e_2 connecting v and w, then surgery produces an algebraic retraction onto Θ_3. See Figure 17(b). If not, two surgeries produce an algebraic retraction onto A. See Figure 17(c).

Then if there is no algebraic retraction from Γ onto Θ_3 or A, the removal of self-loops constitutes an algebraic retraction onto a connected smooth large graph with a simple essential vertex. □

To prepare for the proof of the surgery lemma, we observe that the Świątkowski complex enjoys a certain non-canonical functoriality for subdivisions. More precisely, the map induced by a smoothing $f : \Gamma_1 \to \Gamma_2$ admits a non-canonical left inverse, which we regard as corresponding to the subdivision f^{-1}. It suffices to construct this morphism in the case of a subdivision of a single edge e with half-edges h_1 and h_2 into two edges e_1 and e_2 meeting at the common vertex v with corresponding half-edges h_{v1} and h_{v2}. We define the map in this case by the assignments $e \mapsto e_1$, $h_1 \mapsto h_1$, and $h_2 \mapsto h_{v1} - h_{v2} + h_2$.

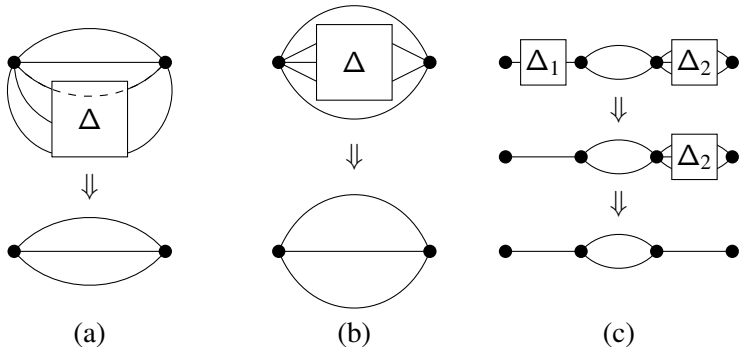

Figure 17: Surgeries for the proof of Lemma 4.16

See Figure 18. This map is a map of differential bigraded $R[E_2]$-modules, where the action of $R[E_2]$ on the target is via the map $R[E_2] \to R[E_1]$ just constructed and hence also non-canonical. The same remarks hold for reduced complexes.

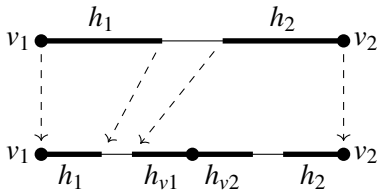

Figure 18: A choice of left inverse for the map induced by a smoothing

Proof of Lemma 4.17 There is a (non-canonical) graph embedding of a subdivision of Γ into Γ', where only the edge e_0 is subdivided. Thus, for any surgery, there is a non-canonical map $\widetilde{S}(\Gamma;R) \to \widetilde{S}(\Gamma';R)$. This non-canonical map admits a left inverse r, defined as follows. Write Γ' as $(\Gamma \setminus e_0) \coprod_{v_1 \cup v_2} \Delta$. Then:

$$r(h) = \begin{cases} h_i & h \in H(\Delta),\ v(h) = v_i \\ 0 & h \in H(\Delta),\ v(h) \neq v_i \\ h & h \notin H(\Delta) \end{cases} \qquad r(e) = \begin{cases} e_0 & e \in E(\Delta) \\ e & e \notin E(\Delta). \end{cases}$$

The map on edges endows $\widetilde{S}(\Gamma;R)$ with a $R[E(\Gamma')]$-module structure, and linearity is a direct verification. By inspection, these maps constitute a retraction both at the level of modules and at the level of rings. □

5 Two isomorphisms

We return to the proof of Theorem 2.10. Along the way, we will encounter two very different routes to the Świątkowski complex. In Section 5.1, we review the method of [2], which has the advantage that it naturally outputs a functor on the category $\mathcal{G}\mathrm{ph}$. On the other hand, there is the cubical deformation retract introduced by Świątkowski [27], which offers a more direct and geometric comparison to $B(\Gamma)$ at the cost of non-functorial choices. Our argument will combine the advantages of these two approaches.

In this section and the next, we write $C^{\mathrm{sing}}(X)$ for the singular chain complex of the topological space X. If X is a CW complex, we denote the cellular chain complex of X by $C(X)$. We will make use of the existence of a zig-zag

$$C^{\mathrm{sing}} \xleftarrow{\sim} \bullet \xrightarrow{\sim} C$$

of quasi-isomorphisms connecting these complexes, which is natural for cellular maps. The specifics of the intermediate object will play no role here, but see [2, Construction 2.22] for one option.

5.1 Functorial model

We now recall some of the work of [2]. The starting point is a cell complex introduced by Abrams [1, Definition 2.1].

Definition 5.1 Let X be a cell complex. The kth unordered *configuration complex* of X is the subspace

$$B_k^{\square}(X) = \left(\bigcup_{\bar{c}_i \cap \bar{c}_j = \varnothing} c_1 \times \cdots \times c_k \right)_{/\Sigma_k} \subseteq B_k(X)$$

where the union is taken over the set of k-tuples of disjoint open cells of X.

The configuration complexes of X are approximations to the configuration spaces of X by cell complexes, and the accuracy of this approximation improves with subdivision. Following [8, §3], we say that a graph Γ is *sufficiently subdivided for k* if, first, every path in Γ between essential vertices passes through at least $k-1$ edges, and, second, every loop passes through at least $k+1$ edges.

Theorem 5.2 ([1, Theorem 2.1]) *Let $\Gamma \to \Gamma'$ be a subdivision with Γ' sufficiently subdivided for k. The inclusion $B_k^{\square}(\Gamma') \to B_k(\Gamma)$ is a deformation retract.*

In order to work functorially and with all k at once, it is more convenient to consider all subdivisions simultaneously. Denoting by $\mathcal{P}$ the set of subdivisions of Γ, viewed as a category under refinement, Theorem 5.2 implies that the natural map

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

is a weak homotopy equivalence [2, Theorem 2.8].

Including arbitrary subdivisions has the further benefit of allowing one to work locally on Γ. In order to leverage this flexibility, we return to the heuristic interpretation of Remark 2.8, in which generators of $S(\Gamma)$ represent "states" in $B(\Gamma)$ obtained by prescribing "local states" near vertices and along edges, which are compatible on half-edges.

More precisely, define $S^{\curlywedge}(\mathsf{S}_n)$ to be the subcomplex of $S(\mathsf{S}_n)$ spanned over $\mathbb{Z}[E]$ by the symbols $\varnothing$, the star vertex and half-edges at the star vertex. For e an edge, viewed as a 1-cell, define $S^{\curlywedge}(e)$ as the subcomplex of $S(e)$ spanned by basis elements with no vertices or half-edges (then $S^{\curlywedge}(e)$ is canonically isomorphic to $\mathbb{Z}[e]$). For these atomic graphs, one can write down a map $C(B^{\square}(\Gamma')) \to S^{\curlywedge}(\Gamma)$ and check by hand that it is a quasi-isomorphism for sufficiently fine subdivisions—see [2, §4.2–4.3].

Now, for Ξ a disjoint union of star graphs and edges, define $S^{\curlywedge}(\Xi)$ to be the tensor product over the connected components of Ξ of the corresponding subcomplexes defined above. Then there is an isomorphism

$$S(\Gamma) \cong S^{\curlywedge}\left(\coprod_{v \in V} \mathsf{S}_{d(v)}\right) \bigotimes_{S^{\curlywedge}((\coprod_{e \in E} \mathsf{I}) \times_E H)} S^{\curlywedge}\left(\coprod_{e \in E} \mathsf{I}\right).$$

We now introduce a device that aids in piecing together these local identifications. Recall that each edge of Γ is identified with $(0,5)$ via its parametrization. We define a map $\pi : \Gamma \to [0,1]$ by setting

$$\pi(t) = \begin{cases} t-1 & t \in [1,2] \subseteq e \\ 1 & t \in [2,3] \subseteq e \\ 4-t & t \in [3,4] \subseteq e \\ 0 & \text{otherwise,} \end{cases}$$

and by sending every vertex of Γ to 0.

Definition 5.3 A *gap* in Γ is a subspace of the form $A = \pi^{-1}(A_0)$, where $A_0 \subseteq [0,1]$ is a nonempty open subset such that

(1) the complement $[0,1] \setminus A_0$ is a (possibly empty) finite union of closed intervals of positive length, and

(2) if $i \in \{0,1\}$ lies in the closure of A_0, then $i \in A_0$.

See Figure 19. We write $\mathcal{G}$ for the poset of gaps.

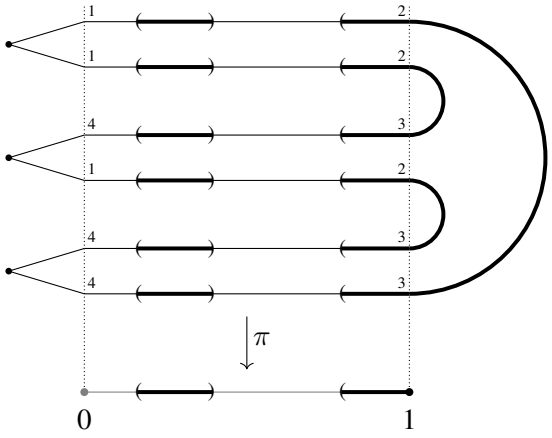

Figure 19: A gap in the complete graph K_3

Since the complement of a gap is a disjoint union of stars and intervals, we have a map $C(B^{\square}(\Gamma' \setminus A)) \to S^{\lambda}\left(\Gamma \setminus A\right)$ whenever A is a gap and Γ' is a subdivision in which A is a union of cells. In this way, we obtain the following zig-zag of quasi-isomorphisms:

$$\begin{array}{ccc}
\operatorname*{colim}_{\Gamma' \in \mathcal{P}} C^{\mathrm{sing}}(B^{\square}(\Gamma')) & \xrightarrow{\ \sim\ } & C^{\mathrm{sing}}(B(\Gamma)) \\
\uparrow\wr & & \\
\bullet & & \\
\downarrow\wr & & \\
\operatorname*{colim}_{\Gamma' \in \mathcal{P}} C(B^{\square}(\Gamma')) & & \\
\uparrow\wr & & \\
\operatorname*{hocolim}_{A \in \mathcal{G}^{op}} \operatorname*{colim}_{\Gamma' \in \mathcal{P}_A} C(B^{\square}(\Gamma' \setminus A)) & & \\
\downarrow\wr & & \\
\operatorname*{hocolim}_{A \in \mathcal{G}^{op}} S^{\lambda}\left(\Gamma \setminus A\right) & \xrightarrow{\ \sim\ } & S(\Gamma).
\end{array}$$

Here, $\mathcal{P}_A \subseteq \mathcal{P}$ contains only those subdivisions of Γ in which A is a union of cells. For details on why these maps are quasi-isomorphisms and why the resulting isomorphism

on homology is functorial, see [2, §4].

Theorem 5.4 ([2, Theorem 4.5]) *There is a natural isomorphism*

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

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

5.2 Cubical model

We now recall the cubical model of $B(\Gamma)$ introduced in [27] and corrected in [4, §2.1]. Given a graph Γ, we write $A(\Gamma)$ for the set of labelings

$$\lambda : E \sqcup V \to \mathbb{Z}_{\geq 0} \sqcup V \sqcup H \sqcup \{\varnothing\}$$

such that $\lambda(e) \in \mathbb{Z}_{\geq 0}$ and $\lambda(v) \in \{\varnothing, v\} \sqcup H(v)$.

Construction 5.5 (Świątkowski) Define a space $UK(\Gamma)$ as the quotient

$$UK(\Gamma) = \coprod_{\lambda \in A(\Gamma)} \{\lambda\} \times [0,1]^{\lambda^{-1}(H)} \Big/ \sim,$$

where the equivalence relation is determined as follows. Fix a half-edge h with edge e and vertex v. Suppose λ, λ_0, and λ_1 are labellings which agree except on v and e, where

$$\begin{array}{lll} \lambda(v) = h & \lambda_0(v) = v & \lambda_1(v) = \varnothing \\ \lambda(e) = n & \lambda_0(e) = n & \lambda_1(e) = n+1. \end{array}$$

Then we glue according to the following identifications for $\epsilon \in \{0,1\}$:

$$\{\lambda_\epsilon\} \times [0,1]^{\lambda_\epsilon^{-1}(H)} \cong \{\lambda\} \times [0,1]^{\lambda^{-1}(H \setminus \{h\})} \times \{\epsilon\}$$

as a subset of $\{\lambda\} \times [0,1]^{\lambda^{-1}(H)}$.

To specify a point in $UK(\Gamma)$, it is enough to give a labeling $\lambda \in A(\Gamma)$ together with numbers $t(h) \in [0,1]$ for each $h \in \lambda^{-1}(H)$.

We endow the set $A(\Gamma)$ with a bigrading by declaring that

$$|\lambda| = \left(|\lambda^{-1}(H)|, \sum_E \lambda(e) + |\lambda^{-1}(V)| + |\lambda^{-1}(H)| \right)$$

and note that $UK(\Gamma)$ splits as a disjoint union of cell complexes $UK_k(\Gamma)$ whose i-cells are those in bigrading (i,k).

Observation 5.6 By inspection, the map

$$C(UK(\Gamma)) \to S(\Gamma) \qquad \lambda \mapsto \left(\prod_E e^{\lambda(e)}\right) \otimes \bigotimes_V \lambda(v)$$

is a bigraded chain isomorphism.

Remark 5.7 To make Observation 5.6 precise, we should take care with orientations. One way to specify an orientation on a cube which is a product of intervals indexed by a set J is to give an order on J up to even permutation. One way to specify the correct sign on a tensor product of odd degree vector spaces indexed by a set J is to give an order on J up to even permutation. Then, with appropriate conventions as to which is the "positive" and which is the "negative" end of an interval, the map of the observation intertwines these conventions.

Remark 5.8 In previous work along these lines [4, 17, 27], neither 2-valent nor 1-valent vertices were considered, but there is no obstruction to this mild generalization.

We now show that the complex $UK(\Gamma)$ is homotopy equivalent to the configuration space of interest. Apart from slight modifications for simplicity and compatibility with our setup, this argument is essentially that of [4, §2.1].[1] In particular, our homotopy equivalence will be homotopic to the one considered in that work.

In order to compare $B(\Gamma)$ and $UK(\Gamma)$, we first deform $B(\Gamma)$ onto the subspace $\widetilde{B}(\Gamma)$ of configurations x with the property that, for each vertex v, at most one coordinate lies in the open star $\mathrm{st}_1(v)$ of radius 1 at v. This deformation is achieved by radial expansion in each star simultaneously. Having taken this intermediate step, we define our comparison map ρ as the composite

$$\rho : B(\Gamma) \to \widetilde{B}(\Gamma) \to UK(\Gamma),$$

where the second map records the presence or absence of particles at vertices, the coordinate of any particle in $\mathrm{st}_1(v) \setminus \{v\} \cong H(v) \times (0,1)$, and the number of particles lying in the subinterval $[1,4]$ of each edge. The details of the deformation are given below in Construction 5.10; for now, we state the following result concerning ρ.

Proposition 5.9 *The map ρ is a homotopy equivalence.*

[1] As observed by Lütgehetmann, the map defined by Świątkowski is not quite a retraction [17, p. 24]. Unfortunately, the replacement retraction constructed by Lütgehetmann is not continuous, as one can see by comparing the formulas for $t_x(s)$ appearing on p. 22 for $k = 2$ and $k = 3$ in a situation in which the first of three particles in an edge approaches the initial endpoint of that edge. Chettih–Lütgehetmann give a continuous retraction.

Construction 5.10 We define the deformation of $B(\Gamma)$ onto $\widetilde{B}(\Gamma)$ as follows. For each vertex v, we have $\overline{\mathrm{st}_2(v)} \setminus \{v\} \cong H(v) \times (0,2]$. On each of these intervals, for a fixed configuration x, we use the homotopy

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

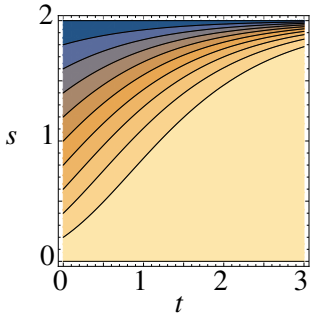

on all points of $x \cap (\{h\} \times (0,2])$ simultaneously. Here, $s \in (0,2]$ and $t \in [0, t(x,v)]$, where $t(x,v)$ is the least t such that the resulting configuration has at most one particle in $\mathrm{st}_1(v)$.[2] Since $t(x,v)$ is continuous in x, and since $t(x,v) = 0$ for $x \in \widetilde{B}(\Gamma)$, this prescription defines a deformation retraction.

Proof of Proposition 5.9 Chettih–Lütgehetmann [4, §2.1] define a cube complex $\widetilde{UK}_k(\Gamma)$ which is a deformation retract of the *ordered* configuration space of k points in Γ. Their cube complex is a Σ_k-cover of the weight k subcomplex of $UK(\Gamma)$. Their deformation retraction is equivariant and so passes to a deformation retraction of $B(\Gamma)$ onto $UK(\Gamma)$. The composite $B(\Gamma) \xrightarrow{\rho} UK(\Gamma) \to B(\Gamma)$ with the inclusion of this quotient deformation retract differs from the identity only in the positions of particles in individual open edges. Therefore, the two maps are homotopic by edgewise straight line homotopies. It follows that ρ is a one-sided homotopy inverse to a homotopy equivalence and hence itself a homotopy equivalence. □

5.3 Comparison of models

We have two isomorphisms $H_*(B(\Gamma)) \cong H_*(S(\Gamma))$. The isomorphism of Theorem 5.4 is an isomorphism of functors on the category $\mathcal{G}\mathrm{ph}$. This naturality is a powerful tool in applications [2, §5]. On the other hand, the isomorphism obtained by combining Proposition 5.9 and Observation 5.6 is more geometric in nature. Fortunately, we need not choose between these virtues.

[2] Explicitly, $t(x,v)$ is 0 if there are fewer than 2 points in $\mathrm{st}_2(v)$ and otherwise $\log(1 - \log_2(s_v))$ where s_v is the distance from v to the second closest point in $\mathrm{st}_2(v)$ in x.

Proposition 5.11 *The diagram of isomorphisms*

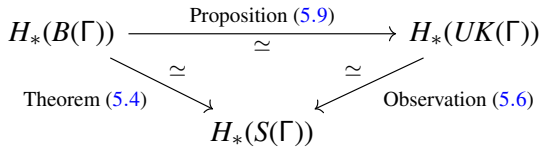

commutes.

The proof of this result will occupy Section 6 below. For now, we will use it to deduce the desired conclusion regarding edge stabilization.

For $\lambda \in A(\Gamma)$, write $e\lambda$ for the labeling that differs from λ only in that $e\lambda(e) = \lambda(e) + 1$ (in particular, $\lambda^{-1}(H) = (e\lambda)^{-1}(H)$). There is a version of edge stabilization at the level of $UK(\Gamma)$ which sends λ to $e\lambda$ and fixes all t_h coordinates. The induced $\mathbb{Z}[E]$-action on $C(UK(\Gamma))$ coincides, through the isomorphism of Observation 5.6, with the canonical action on $S(\Gamma)$.

Proof of Theorem 2.10 It suffices by Proposition 5.11 to show that the map

$$\rho_* : H_*(B(\Gamma)) \to H_*(\widetilde{B}(\Gamma)) \to H_*(UK(\Gamma))$$

is $\mathbb{Z}[E]$-linear. By inspection, we have the commuting diagram

$$\begin{array}{ccc} \widetilde{B}(\Gamma) & \longrightarrow & UK(\Gamma) \\ {\scriptstyle \sigma_e}\downarrow & & \downarrow \\ \widetilde{B}(\Gamma) & \longrightarrow & UK(\Gamma), \end{array}$$

so the second map in this composite is $\mathbb{Z}[E]$-linear. In order to show that the first map is also $\mathbb{Z}[E]$-linear, we note that the inverse, which is induced by the inclusion $\widetilde{B}(\Gamma) \subseteq B(\Gamma)$, is $\mathbb{Z}[E]$-linear. □

We also record the following useful conclusion.

Corollary 5.12 *The differential bigraded $\mathbb{Z}[E]$-modules $C^{\mathrm{sing}}(B(\Gamma))$ and $S(\Gamma)$ are quasi-isomorphic.*

Proof We have the zig-zag of $\mathbb{Z}[E]$-linear quasi-isomorphisms

$$C^{\mathrm{sing}}(B(\Gamma)) \xleftarrow{\sim} C^{\mathrm{sing}}(\widetilde{B}(\Gamma)) \xrightarrow{\sim} C^{\mathrm{sing}}(UK(\Gamma)) \xleftarrow{\sim} \bullet \xrightarrow{\sim} C(UK(\Gamma)) \cong S(\Gamma). \qquad \square$$

Unlike the homology isomorphism, this quasi-isomorphism is not natural, since it relies on a choice of parametrization.

6 Long ends and the proof of Proposition 5.11

In this section, we compare the two isomorphisms $H_*(B(\Gamma)) \cong H_*(S(\Gamma))$. The key observation is that, for a certain class of subdivision $\Gamma \to \Gamma'$, the natural map

$$B_k^{\square}(\Gamma') \subseteq B_k(\Gamma') \cong B_k(\Gamma) \to UK_k(\Gamma)$$

from Abrams' model to Świątkowski's model is cellular. With this observation in hand, Proposition 5.11 follows after checking that this collection of special subdivisions is large enough to support the argument of [2].

6.1 Long ends

In this section, we fix a parametrization of a graph Γ, identifying each edge with $(0,5)$. Given a subdivision $\Gamma \to \Gamma'$, we do *not* independently parametrize Γ'. Rather, we use the homeomorphism underlying the subdivision to identify Γ' with Γ. This allows us to specify the data of the subdivision Γ' up to isomorphism by naming a finite set of points in $(0,5)$ for each edge of Γ.

Definition 6.1 We say that a subdivision $\Gamma \to \Gamma'$ has *long ends* if, for every edge of Γ, the induced subdivision $\{0, t_1, \dots, t_r, 5\}$ of $[0,5]$ has $t_1 = 1$ and $t_r = 4$.

If Γ' has long ends, we identify a half-edge h of Γ with the corresponding 1-cell of length 1 in Γ'. The following observation is the heart of our comparison of the two models in question.

Lemma 6.2 *If $\Gamma \to \Gamma'$ has long ends, then the composite*

$$B_k^{\square}(\Gamma') \subseteq B_k(\Gamma) \xrightarrow{\rho_k} UK_k(\Gamma)$$

is cellular.

Proof A cell of $B_k^{\square}(\Gamma')$ is specified by a function μ from the set of cells of Γ' to $\{0,1\}$ with the following properties:

(1) $\sum_c \mu(c) = k$

(2) if $c_i \cap c_j \neq \varnothing$, then $\mu(c_i) + \mu(c_j) \leq 1$.

The degree of the cell μ is the sum of μ over the 1-cells of Γ'.

By inspection, the composite in question maps the cell μ onto the cell λ_μ with

$$\lambda_\mu(v) = \begin{cases} v & \mu(v) = 1 \\ h & \mu(h) = 1 \\ \varnothing & \text{otherwise.} \end{cases}$$

and $\lambda_\mu(e)$ defined by summing μ over all cells of the induced subdivision on $[1,4]$. Since the degree of λ_μ is at most the degree of μ, this claim implies the lemma. □

Write $\mathcal{L}$ for the category of subdivisions of Γ with long ends. Although this category is filtered, it is not convergent in the sense of [2, Definition 2.6]. Nevertheless, we have the following.

Lemma 6.3 *For every $k \geq 0$, the natural map*

$$\operatorname*{colim}_{\mathcal{L}} C(B_k^{\square}(\Gamma')) \to \operatorname*{colim}_{\mathcal{P}} C(B_k^{\square}(\Gamma'))$$

is a quasi-isomorphism.

Proof Let $\mathcal{P}_k \subseteq \mathcal{P}$ denote the subcategory of subdivisions of Γ that are sufficiently subdivided for k, and set $\mathcal{L}_k = \mathcal{L} \cap \mathcal{P}_k$. Since $\mathcal{L}_k$ is final in $\mathcal{L}$ and $\mathcal{P}_k$ in $\mathcal{P}$, the vertical arrows in the commuting diagram

$$\begin{array}{ccc} \operatorname{colim}_{\mathcal{L}_k} C(B_k^{\square}(\Gamma')) & \longrightarrow & \operatorname{colim}_{\mathcal{P}_k} C(B_k^{\square}(\Gamma')) \\ \downarrow & & \downarrow \\ \operatorname{colim}_{\mathcal{L}} C(B_k^{\square}(\Gamma')) & \longrightarrow & \operatorname{colim}_{\mathcal{P}} C(B_k^{\square}(\Gamma')) \end{array}$$

are isomorphisms, so it suffices to show that the top arrow is a quasi-isomorphism. By sufficient subdivision and Theorem 5.2, every arrow in $\mathcal{P}_k$ induces a quasi-isomorphism on $C(B_k^{\square}(-))$, so the claim follows from the observation that the inclusion of $\mathcal{L}_k$ in $\mathcal{P}_k$ induces a weak homotopy equivalence on nerves, since both nerves are contractible. Indeed, both $\mathcal{L}$ and $\mathcal{P}$ are filtered, hence contractible, and final functors induce weak equivalences on nerves. □

Lemma 6.4 *Let $A = \pi^{-1}(A_0)$ be a gap. If $A_0 \cap \{0,1\} = \varnothing$, then the natural map*

$$\operatorname*{colim}_{\Gamma' \in \mathcal{L} \cap \mathcal{P}_A} C(B^{\square}(\Gamma' \setminus A)) \to \operatorname*{colim}_{\Gamma' \in \mathcal{P}_A} C(B^{\square}(\Gamma' \setminus A))$$

is a quasi-isomorphism.

Proof The category of subdivisions of $\Gamma' \setminus A$ that are sufficiently subdivided for a fixed k is final in the category of subdivisions restricted from $\mathcal{P}$. Because $A_0 \cap \{0,1\} = \varnothing$, the same is true for $\mathcal{L}$, so the claim follows in the manner of Lemma 6.3. □

Lemma 6.5 *Suppose that $\Gamma \to \Gamma'$ has long ends, and let A be a gap in Γ that is a union of cells of Γ'. The following diagram of chain maps commutes:*

$$\begin{array}{ccccc} C(B^{\square}(\Gamma' \setminus A)) & \longrightarrow & C(B^{\square}(\Gamma')) & \xrightarrow{C(\rho)} & C(UK(\Gamma)) \\ \downarrow & & & & \| \wr \\ S^{\curlywedge}\left(\Gamma \setminus A\right) & & \longrightarrow & & S(\Gamma). \end{array}$$

Proof The claim is immediate from the explicit description of the value of ρ on cells given in Lemma 6.2 and the description of the lefthand vertical map given in [2, Definitions 4.13 and 4.15]. □

6.2 Proof of Proposition 5.11

We wish to compare two isomorphisms identifying $H_*(B(\Gamma))$ with $H_*(S(\Gamma))$. The first is induced on homology by the zig-zag of quasi-isomorphisms in the righthand portion of the diagram of Figure 20, in which each of the square subdiagrams commutes. The first

$$\begin{array}{ccccc} \operatorname*{colim}_{\Gamma' \in \mathcal{L}} C^{\mathrm{sing}}(B^{\square}(\Gamma')) & \longrightarrow & \operatorname*{colim}_{\Gamma' \in \mathcal{P}} C^{\mathrm{sing}}(B^{\square}(\Gamma')) & \xrightarrow{\sim} & C^{\mathrm{sing}}(B(\Gamma)) \\ {\scriptstyle (1)}\uparrow & & \uparrow \wr & & \\ \bullet & \longrightarrow & \bullet & & \\ {\scriptstyle (2)}\downarrow & & \downarrow \wr & & \\ \operatorname*{colim}_{\Gamma' \in \mathcal{L}} C(B^{\square}(\Gamma')) & \xrightarrow{(3)} & \operatorname*{colim}_{\Gamma' \in \mathcal{P}} C(B^{\square}(\Gamma')) & & \\ \uparrow & & \uparrow \wr & & \\ \operatorname*{hocolim}_{A \in \mathcal{G}^{op}} \operatorname*{colim}_{\Gamma' \in \mathcal{L} \cap \mathcal{P}_A} C(B^{\square}(\Gamma' \setminus A)) & \xrightarrow{(4)} & \operatorname*{hocolim}_{A \in \mathcal{G}^{op}} \operatorname*{colim}_{\Gamma' \in \mathcal{P}_A} C(B^{\square}(\Gamma' \setminus A)) & & \\ & & \downarrow \wr & & \\ & & \operatorname*{hocolim}_{A \in \mathcal{G}^{op}} S^{\curlywedge}\left(\Gamma \setminus A\right) & \xrightarrow{\sim} & S(\Gamma) \end{array}$$

Figure 20: A commutative diagram of quasi-isomorphisms for the proof of Proposition 5.11

step in the proof is to verify that *all* of the maps in the diagram are quasi-isomorphisms, so that we may replace this zig-zag with the outer zig-zag in the diagram.

(1)–(2) These maps are induced by natural quasi-isomorphisms after taking the colimit over $\mathcal{L}$. Since $\mathcal{L}$ is filtered, the claim follows.

(3) This quasi-isomorphism is supplied by Lemma 6.3.

(4) This quasi-isomorphism follows from Lemma 6.4 and the observation that the poset of gaps satisfying the hypotheses of that lemma is homotopy initial in $\mathcal{G}$ (hence homotopy final in $\mathcal{G}^{op}$).

The remaining arrows are quasi-isomorphisms by two-out-of-three.

The second isomorphism $H_*(B(\Gamma)) \to H_*(S(\Gamma))$ is induced by the map $\rho : B(\Gamma) \to UK(\Gamma)$ of topological spaces, together with the identification of $S(\Gamma)$ with cellular chains on $UK(\Gamma)$. In order to compare this isomorphism with the previous, we note that Lemma 6.2 supplies the dashed fillers in the commuting diagram

$$\begin{array}{ccccc}
C^{\mathrm{sing}}(B^{\square}(\Gamma')) & \longrightarrow & C^{\mathrm{sing}}(B(\Gamma)) & \xrightarrow{\sim} & C^{\mathrm{sing}}(UK(\Gamma)) \\
\wr\uparrow & & & & \uparrow\wr \\
\bullet & & \dashrightarrow & & \bullet \\
\wr\downarrow & & & & \downarrow\wr \\
C(B^{\square}(\Gamma')) & & \dashrightarrow & & C(UK(\Gamma))
\end{array}$$

whenever $\Gamma \to \Gamma'$ has long ends. Passing to the colimit over $\mathcal{L}$, we obtain the diagram

$$\begin{array}{ccccc}
\operatorname*{colim}_{\Gamma'\in\mathcal{L}} C^{\mathrm{sing}}(B^{\square}(\Gamma')) & \xrightarrow{\sim} & C^{\mathrm{sing}}(B(\Gamma)) & \xrightarrow{\sim} & C^{\mathrm{sing}}(UK(\Gamma)) \\
\wr\uparrow & & & & \uparrow\wr \\
\bullet & & \longrightarrow & & \bullet \\
\wr\downarrow & & & & \downarrow\wr \\
\operatorname*{colim}_{\Gamma'\in\mathcal{L}} C(B^{\square}(\Gamma')) & & \longrightarrow & & C(UK(\Gamma)) \\
\wr\uparrow & & & & \| \\
\operatorname*{hocolim}_{A\in\mathcal{G}^{op}} \operatorname*{colim}_{\Gamma'\in\mathcal{L}\cap\mathcal{P}_A} C(B^{\square}(\Gamma'\setminus A)) & & & & \wr\| \\
\wr\downarrow & & & & \| \\
\operatorname*{hocolim}_{A\in\mathcal{G}^{op}} S^{\curlywedge}\left(\Gamma\setminus A\right) & & \xrightarrow{\sim} & & S(\Gamma).
\end{array}$$

The upper portion of the diagram commutes by what has already been said, and the bottom portion of the diagram commutes by Lemma 6.5 and the universal properties of the colimit and the homotopy colimit. We established in the first half of the proof that the

first of the isomorphisms in question is induced on homology by the counterclockwise zig-zag from $C^{\mathrm{sing}}(B(\Gamma))$ to $S(\Gamma)$. Since the clockwise zig-zag induces the second of the isomorphisms, the proof is complete.

References

[1] **A Abrams**, *Configuration spaces of braid groups of graphs*, PhD thesis, UC Berkeley (2000)

[2] **B H An**, **G C Drummond-Cole**, **B Knudsen**, *Subdivisional spaces and graph braid groups*, Doc. Math. (2019)To appear

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

[4] **S Chettih**, **D Lütgehetmann**, *The homology of configuration spaces of trees with loops*, Alg. Geom. Topol. 18 (2018) 2443–2469

[5] **T Church**, **J S Ellenberg**, **B Farb**, *FI-modules and stability for representations of symmetric groups*, Duke Math. J. 164 (2015) 1833–1910

[6] **F Cohen**, **T Lada**, **J P May**, *The homology of iterated loop spaces*, Lecture Notes in Math. 533, Springer (1976)

[7] **D Eisenbud**, *Commutative Algebra With a View Toward Algebraic Geometry*, Springer-Verlag (1999)

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

[9] **D Farley**, **L Sabalka**, *On the cohomology rings of tree braid groups*, J. Pure Appl. Algebra 212 (2008) 53–71

[10] **S Gal**, *Euler characteristic of the configuration space of a complex*, Colloq. Math. 89 (2001) 61–67

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

[12] **J M Harrison**, **J P Keating**, **J M Robbins**, **A Sawicki**, *n-Particle Quantum Statistics on Graphs*, Comm. Math. Phys. 330 (2014) 1293–1326

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

[14] **D E Knuth**, *Big Omicron and Big Omega and Big Theta*, SIGACT News 8 (1976) 18–24

[15] **K H Ko**, **J H La**, **H W Park**, *Graph 4-braid groups and Massey products*, Topology Appl. 197 (2016) 133–153

[16] **K H Ko**, **H W Park**, *Characteristics of Graph Braid Groups*, Discrete Comput. Geom. 48 (2012) 915–963

[17] **D Lütgehetmann**, *Configuration spaces of graphs*, Master's thesis, Freie Universität Berlin (2014)

[18] **D Lütgehetmann**, *Representation stability for configuration spaces of graphs* (2017)ArXiv:1701.03490

[19] **D Lütgehetmann**, **D Recio-Mitter**, *Topological complexity of configuration spaces of fully articulated graphs and banana graphs* (2018)ArXiv:1806.00659

[20] **T Maciążek**, **A Sawicki**, *Homology groups for particles on one-connected graphs*, J. Math. Phys. 58 (2017)

[21] **T Maciążek**, **A Sawicki**, *Non-abelian quantum statistics on graphs* (2018)ArXiv:1806.02846

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

[23] **S Papadima**, **A I Suciu**, *Algebraic invariants for right-angled Artin groups*, Math. Ann. 334 (2006) 533–555

[24] **E Ramos**, *Configuration spaces of graphs with certain permitted collisions* (2017)ArXiv:1703.05535

[25] **E Ramos**, *Stability phenomena in the homology of tree braid groups*, Alg. Geom. Topol. 18 (2018) 2305–2337

[26] **L Sabalka**, *On rigidity and the isomorphism problem for tree braid groups*, Groups Geom. Dyn. 3 (2009) 469–523

[27] **J Świątkowski**, *Estimates for homological dimension of configuration spaces of graphs*, Colloq. Math. 89 (2001) 69–79

Center for Geometry and Physics, Institute for Basic Science (IBS), Pohang 37673, Republic of Korea

Center for Geometry and Physics, Institute for Basic Science (IBS), Pohang 37673, Republic of Korea

Department of Mathematics, Harvard University, Cambridge 02138, USA

anbyhee@ibs.re.kr, gabriel.c.drummond.cole@gmail.com, knudsen@math.harvard.edu

drummondcole.com/gabriel/academic/