

ON THE SECOND HOMOLOGY OF PLANAR GRAPH BRAID GROUPS

BYUNG HEE AN AND BEN KNUDSEN

ABSTRACT. We show that the second homology of the configuration spaces of a planar graph is generated under the operations of embedding, disjoint union, and edge stabilization by three atomic graphs: the cycle graph with one edge, the star graph with three edges, and the theta graph with four edges. We give an example of a non-planar graph for which this statement is false.

1. INTRODUCTION

A fundamental dichotomy, observed since the dawn of the study of configuration spaces, is that of stability and instability [Arn69, McD75, Chu12, CEF15]. Writing $B_k(X)$ for the k th unordered configuration space of the topological space X , a phenomenon is said to be *stable* if it occurs for all sufficiently large k , or perhaps in the limit as $k \rightarrow \infty$. Stable phenomena tend to be structured and calculable, unstable phenomena fleeting and irregular, hence more difficult to grasp.

This paper is concerned with the unstable homology of configuration spaces of graphs, or equivalently of their fundamental groups, the graph braid groups [Abr00]. This investigation is a companion and counterpoint to the recent complete calculation of the stable homology [ADCK21], premised on the action by *edge stabilization* of the polynomial ring generated by the edges of the graph Γ on $H_*(B(\Gamma))$, where $B(\Gamma) := \bigsqcup_{k \geq 0} B_k(\Gamma)$ [ADCK20].

Unstably, little systematic is known beyond the landmark calculation by Ko–Park of the first homology [KP12]. One consequence of this calculation is that $H_1(B(\Gamma))$ is generated under edge stabilization by *loop classes* and *star classes*, geometric generators represented by maps from circles. Disjoint unions of stars and loops then give rise to higher degree classes represented by maps from tori.

As observed independently in [CL18] and [WG17], the space $B_3(\Theta_4)$ has the homotopy type of a surface of genus 3, whose fundamental class cannot be represented by a map from a torus—here, Θ_4 is the suspension of four points. Our main result is that, in the planar case, this *theta class* is the only “exotic” generator in degree 2.

Theorem 1.1. *Let Γ be a planar graph with set of edges E . The $\mathbb{Z}[E]$ -module $H_2(B(\Gamma))$ is generated by toric classes and theta classes.*

Precise descriptions of the action of $\mathbb{Z}[E]$ and of the classes in question can be found in Sections 2 and 3, respectively. As illustrated by Example 6.8, the assumption of planarity cannot be removed.

1.1. Questions. Our work invites the following questions, which we hope to pursue in the future.

- (1) *Universal relations.* Star classes and loop classes are universal generators for the first homology of graph braid groups, and Theorem 1.1 provides universal generators for the second homology in the planar case. In degree 1, a complete set of universal relations is known. What are the relations in degree 2?
 - (2) *Combinatorial bases.* Ko–Park give a basis for $H_1(B(\Gamma))$ in terms of combinatorial invariants of Γ [KP12, Thm. 3.16]. Can a similar basis be given for $H_2(B(\Gamma))$?
 - (3) *Higher degrees.* We conjecture that a direct analogue of Theorem 1.1 holds in all degrees.
- Conjecture 1.2.** *For any planar graph Γ and $i > 0$, the $\mathbb{Z}[E]$ -module $H_i(B(\Gamma))$ is generated by classes arising from disjoint unions of cycle graphs, star graphs, and theta graphs.*
- (4) *Non-planar graphs.* What extra generators are needed in order to remove the assumption of planarity from Theorem 1.1 and Conjecture 1.2?

1.2. Strategy and outline. Writing $M(\Gamma)$ for the submodule spanned by toric classes and theta classes, the theorem is the equality $M(\Gamma) = H_2(B(\Gamma))$. This conclusion being well known for trees, we proceed by induction on the first Betti number of Γ . The tool facilitating this induction is the exact sequence of Proposition 2.6, originally introduced in [ADCK19], which expresses $H_2(B(\Gamma))$ as an extension by a submodule contained in $M(\Gamma)$ by induction. By exactness, proving the theorem becomes a matter of showing that $M(\Gamma)$ surjects onto the quotient by this submodule (Theorem 5.1), which is achieved through consideration of a special generating set for the cokernel (Proposition 5.10). Two inductions reduce the vanishing of this generating set to the triconnected case, which is handled by a combinatorial argument.

In linear order, Section 2 is concerned with background material, Section 3 discusses the geometric classes of interest, Section 4 details decomposition tactics for the purposes of induction, Section 5 provides the key reformulation of Theorem 1.1 and the special generating set for the cokernel, Section 6 is concerned with examples and the first induction, Section 7 deals with the triconnected case, and Section 8 carries out the second induction.

1.3. Acknowledgements. This paper benefited substantially from conversations with Gabriel Drummond-Cole, who declined to be named as coauthor. The authors also thank Roberto Pagaria for a simplifying observation. The first author was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. 2020R1A2C1A0100320). The second author was supported by NSF grant DMS 1906174.

2. PRELIMINARIES

This section presents a brief overview of some necessary background material in the study of the homology of graph braid groups. A more leisurely exposition is available in [ADCK20].

2.1. Conventions on graphs. A graph is a finite CW complex of dimension at most 1, whose 0-cells and open 1-cells are called vertices and edges, respectively. A graph is called a tree if it is contractible and a cycle if it is homeomorphic to S^1 . A half-edge is a point in the preimage of a vertex under the attaching map of a 1-cell; thus, every edge determines two half-edges. In general, sets of vertices, edges, and half-edges are denoted $V(\Gamma)$, $E(\Gamma)$, and $H(\Gamma)$, respectively, but we omit Γ from the notation wherever doing so causes no ambiguity.

A half-edge h has an associated vertex $v(h)$ and an associated edge $e(h)$, and we write $H(v) = \{h \in H : v = v(h)\}$ for the set of half-edges incident on $v \in V$. The degree or valence of v is $d(v) = |H(v)|$. A vertex is essential if its valence is at least 3. An edge is a tail if its closure contains a vertex of valence 1 and a self-loop if its closure contains only one vertex. A multiple edge is a set of edges incident on the same pair of vertices.

A subgraph is a subcomplex of a graph. A graph morphism is a finite composition of isomorphisms onto subcomplexes and inverse subdivisions, which we call smoothings—see Figure 1 and [ADCK20, §2.1].



FIGURE 1. Two graph structures on a pair of handcuffs, for which the identity is a smoothing from left to right

We close with several important families of graphs, examples of which are depicted in Figure 2.

Example 2.1. The *star graph* S_n is the cone on the discrete space $\{1, \dots, n\}$.

Example 2.2. The *theta graph* Θ_n is the unique graph with two vertices, n edges, and no self-loops.

Example 2.3. The *complete bipartite graph* $K_{m,n}$ is the join of the discrete spaces $\{1, \dots, m\}$ and $\{1, \dots, n\}$.




FIGURE 2. Examples of graphs

2.2. The Świątkowski complex. Our object of study in this paper is the homology of the configuration spaces of a graph Γ . Our primary weapon is a certain chain complex, which we now define—see [ADCK20, §2.2] for further discussion.

Definition 2.4. Let Γ be a graph. For $v \in V$, write $S(v) = \mathbb{Z}\langle \emptyset, v, h \in H(v) \rangle$. The *Świątkowski complex* is the $\mathbb{Z}[E]$ -module

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

equipped with the bigrading $|\emptyset| = (0, 0)$, $|v| = |e| = (0, 1)$, and $|h| = (1, 1)$, together with the differential determined by the equation $\partial(h) = e(h) - v(h)$.

Note that the differential ∂ preserves the second grading, which corresponds to the number of particles in a configuration. We refer to this auxiliary grading as *weight*.

We systematically omit all factors of $\emptyset$ and all tensor symbols when dealing with elements of $S(\Gamma)$, and we regard half-edge generators at different vertices as permutable up to sign.

Theorem 2.5 ([ADCK20, Thm. 2.10]). *There is a natural isomorphism of bigraded $\mathbb{Z}[E]$ -modules*

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

Several comments are in order. First, precursors to this result can be found in [Świ01] and [CL18]. Second, the action of $\mathbb{Z}[E]$ on the lefthand side arises from an E -indexed family of *edge stabilization maps*. Stabilization at e replaces the subconfiguration of particles lying in the closure of e with the collection of averages of consecutive particles and endpoints—see Figure 3 and [ADCK20, §2.2]. Third, regarding the implied functoriality, we direct the reader to [ADCK20, §2.3].




FIGURE 3. Edge stabilization

2.3. Reduction and explosion. The *reduced* Świątkowski complex $\tilde{S}(\Gamma)$ is obtained by replacing $S(v)$ in the definition of $S(\Gamma)$ with the submodule $\tilde{S}(v) \subseteq S(v)$ spanned by $\emptyset$ and all differences of half-edges. The inclusion $\tilde{S}(\Gamma) \subseteq S(\Gamma)$ is a quasi-isomorphism as long as Γ has no isolated vertices [ADCK19, Prop. 4.9]. Note that, for any $h_0 \in H(v)$, a basis for $\tilde{S}(v)$ is given by $\{\emptyset\} \cup \{h - h_0\}_{h_0 \neq h \in H(v)}$. In this way, a (non-canonical) basis for $\tilde{S}(\Gamma)$ may be obtained.

Given a graph Γ and $v \in V$, we write Γ_v for the graph obtained by exploding the vertex v —see Figure 4 and [ADCK20, Def. 2.12]—which we regard as a subgraph of a subdivision of Γ , uniquely up to isotopy. More generally, given a subset $W \subseteq V$, we write Γ_W for the graph obtained by exploding each of the vertices in W .

Proposition 2.6 ([ADCK20, Prop. 2.3]). *Fix $v \in V$ and $h_0 \in H(v)$. The sequence*

$$\cdots \rightarrow H_i(B_k(\Gamma_v)) \xrightarrow{\iota_*} H_i(B_k(\Gamma)) \xrightarrow{\psi} \bigoplus_{h_0 \neq h \in H(v)} H_{i-1}(B_{k-1}(\Gamma_v)) \xrightarrow{\delta} H_{i-1}(B_k(\Gamma_v)) \rightarrow \cdots$$

is exact. Here,

- (1) *the map ι_* is induced by the inclusion of Γ_v ,*




FIGURE 4. A local picture of vertex explosion

- (2) the map ψ is induced by the chain map on reduced Świątkowski complexes sending $b + \sum (h - h_0)a_h$ to (a_h) , where b involves no half-edge generators at v , and
- (3) on the summand indexed by h , the map δ is multiplication by the ring element $e(h) - e(h_0)$.

All maps shown are natural and compatible with edge stabilization.

The same exact sequence obtains with homology replaced by homology with coefficients in an arbitrary commutative ring.

We close this section by drawing the following simple, but useful, consequence.

Lemma 2.7. *Let v be a vertex of Γ and e_1 and e_2 edges lying in distinct components of Γ_v . The $(e_1 - e_2)$ -torsion submodule of $H_1(B(\Gamma))$ is contained in the image of $H_1(B(\Gamma_v))$.*

Proof. Our assumption implies that multiplication by $e_1 - e_2$ is injective on the third term in the exact sequence

$$\cdots \rightarrow H_1(B_k(\Gamma_v)) \rightarrow H_1(B_k(\Gamma)) \rightarrow \bigoplus_{d(v)-1} H_0(B_{k-1}(\Gamma_v)) \rightarrow \cdots,$$

and the claim follows. $\square$

3. GENERATORS AND RELATIONS

This section introduces loop, star, and theta classes, the atomic homology classes involved in Theorem 1.1, and explores some relations among them.

3.1. Loop classes and star classes. We begin with two basic types of class in $H_1(B(\Gamma))$. The reader is directed to [ADCK19, §5.1] for further details.

Example 3.1. Since $\Gamma = B_1(\Gamma)$ is a subspace of $B(\Gamma)$, an oriented cycle in Γ determines an element of $H_*(B(\Gamma))$, called a *loop class*. We denote loop classes generically by the letter β .

A standard chain level representative of a loop class is obtained by summing the differences of half-edges involved in the cycle in question. For example, the standard representative of the unique loop class in the graph L depicted in Figure 5, oriented clockwise, is $b = h - h' \in \tilde{S}(L)$.

Example 3.2. In view of the homotopy equivalence $S^1 \simeq B_2(S_3)$, the choice of half-edges h_1 , h_2 , and h_3 sharing a common vertex determines a *star class* in $H_1(B_2(\Gamma))$, which depends on the ordering only up to sign. We denote star classes by α or, e.g., α_{123} if we wish to emphasize the particular choice of half-edges.




FIGURE 5. The lollipop graph L

Writing e_j for the edge associated to h_j , a standard chain level representative for a star class is given by the sum

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

The alternative expression

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

in the basis for $\tilde{S}(\Gamma)$ privileging h_1 is also useful.

In what follows, we refer to the standard representatives introduced above as *loop cycles* and *star cycles* respectively.

Definition 3.3. The *support* of a star cycle a is the vertex $v(h)$, where h is any half-edge involved in a . The *support* of a loop cycle b is the union of the edges $e(h)$ and vertices $v(h)$, where h ranges over all half-edges involved in b .

In other words, the support of a star cycle is the essential vertex used in its definition, and the support of a loop cycle is the loop used in its definition. The reader is warned that support is not well-defined at the level of homology.

Proposition 3.4 ([ADCK19, Prop. 5.6]). *If Γ is connected, then $H_1(B(\Gamma))$ is generated over $\mathbb{Z}[E]$ by star classes and loop classes.*

Star classes and loop classes interact according to a relation called the *Q-relation*. Our notation will refer to the graph L of Figure 5, but functoriality propagates the relation to any graph with a subgraph isomorphic to a subdivision of L . Writing β for the clockwise oriented loop class and α for the counterclockwise oriented star class in L , we have the following.

Lemma 3.5 (Q-relation). *In the homology of the configuration spaces of the graph L , there is the relation $(e - e')\beta = \alpha$.*

Star classes at distinct vertices can also be related. Referring to the graph Θ_3 as shown in Figure 2, and writing α and α' for the clockwise oriented star classes at the top vertex and bottom vertices, respectively, we have the following relation.

Lemma 3.6 (θ -relation). *In the homology of the configuration spaces of the graph Θ_3 , there is the relation of star classes $\alpha - \alpha' = 0$.*

This last relation has the following amusing consequence, which is left as an exercise (or see [ADCK19, p. 60]).

Example 3.7. Any two star classes in $H_1(B_2(K_{3,3}))$ are equal, regardless of orientation. In particular, any such star class in is 2-torsion.

It is useful to distinguish those star classes involved in no instances of the θ -relation.

Definition 3.8. A star cycle with support v is *rigid* if it involves half-edges lying in multiple components of Γ_v . A star class is *rigid* if it has a rigid representative.

According to [ADCK20, Lem. 3.15], any star cycle representing a rigid star class is rigid.

In view of the homeomorphism $B(\Gamma_1 \sqcup \Gamma_2) \cong B(\Gamma_1) \times B(\Gamma_2)$, classes in the homology of $B(\Gamma)$ represented by cycles in the configuration spaces of disjoint subgraphs give rise to an *external product* class in $H_*(B(\Gamma))$ [ADCK19, Def. 5.10]. At the level of Świątkowski complexes, the external product is represented by the tensor product of representing cycles. The reader is cautioned that the external product may depend on the choice of representing cycles.

Note that stabilizations of external products of loop classes and star classes are represented by maps from tori.

3.2. Theta classes. In this section, we give an elementary description of the non-toric class in $H_2(B_3(\Theta_4))$ discovered in [CL18] and [WG17] terms of the combinatorics of the Świątkowski complex. We begin with a simple lemma.

Lemma 3.9. *In the following, $2 \leq i \leq 4$ and $2 \leq j < k \leq 4$.*

- (1) *The Abelian group $H_1(B_2(S_4))$ is freely generated by the classes α_{1jk} .*
- (2) *The Abelian group $H_1(B_3(S_4))$ is generated by the classes $(e_i - e_1)\alpha_{1jk}$ and $e_1\alpha_{1jk}$, subject only to the relation*

$$(e_4 - e_1)\alpha_{123} - (e_3 - e_1)\alpha_{124} + (e_2 - e_1)\alpha_{134} = 0.$$

Proof. Since both groups are spanned by stabilized star classes by Proposition 3.4, the relations of [ADCK21, Lem. 2.9, 2.10] imply generation and the validity of the relation. Since $H_1(B_2(S_4))$ and $H_1(B_3(S_4))$ are free Abelian with respective ranks 3 and 11 by [FS05, Cor. 4.2], the claim follows. $\square$

Proposition 3.10. *The group $H_2(B_3(\Theta_4))$ is free Abelian of rank 1.*

Proof. Consider the exact sequence

$$0 \rightarrow H_2(B_3(\Theta_4)) \xrightarrow{\psi} H_1(B_2(S_4))^{\oplus 3} \xrightarrow{\delta} H_1(B_3(S_4))$$

arising from Proposition 2.6 after exploding one of the vertices of Θ_4 , privileging the half-edge incident on e_1 . It follows from Lemma 3.9 and the explicit formula for δ given in Proposition 2.6 that the image of δ is free Abelian of rank 8, implying the claim. $\square$

Definition 3.11. Let Γ be a graph. A *theta class* is a class in $H_2(B_3(\Gamma))$ that is the image of a generator of $H_2(B_3(\Theta_4))$ under a topological embedding $\Theta_4 \rightarrow \Gamma$.

In what follows, it will be useful to have a second method of accessing theta classes. We subdivide one of the edges of Θ_4 once by adding a bivalent vertex v and write (abusively) Θ_v for the graph obtained by exploding this bivalent vertex.

Lemma 3.12. *In the long exact sequence for the vertex explosion Θ_v , the map ψ sends a generator of $H_2(B_3(\Theta_4))$ to the unique (up to sign) non-rigid star class in $H_1(B_2(\Theta_v))$.*

Proof. Denoting the star and theta classes in question by α and τ , respectively, the θ -relation implies that $\alpha \in \ker(\delta)$. Thus, by exactness and Proposition 3.10, we conclude that $\alpha = n\psi(\tau)$ for some $0 \neq n \in \mathbb{Z}$. The exact sequence in question is valid (and α nonzero) over any field, and carrying out the same argument over $\mathbb{F}_p$ shows that $p \nmid n$ for every prime p , whence $n = \pm 1$. $\square$

For the sake of completeness, we now give an explicit chain level representative for the generator of $H_2(B_3(\Theta_4))$, although we will not have cause to use it.

Construction 3.13. Denote the vertices of Θ_4 as v_1 and v_2 , and write $h_{i,j}$ for the half-edge incident on v_i and e_j . Define an element $A_2 \in \tilde{S}_2(\Theta_4)_3$ by the formula

$$A_2 = \sum_{\sigma \in \Sigma_4} \text{sgn}(\sigma)(e_1 - e_{\sigma(3)})(h_{1,\sigma(1)} - h_{1,1})(h_{2,\sigma(2)} - h_{2,1}).$$

Lemma 3.14. *The chain A_2 is a cycle, and $[A_2]$ generates $H_2(B_3(\Theta_4))$.*

Proof. For the first claim, consideration of the transpositions τ_{13} and τ_{23} shows that all terms cancel in the sum

$$\partial A_2 = \sum_{\sigma \in \Sigma_4} \text{sgn}(\sigma)(e_1 - e_{\sigma(3)}) [(e_{\sigma(1)} - e_1)(h_{2,\sigma(2)} - h_{2,1}) - (e_{\sigma(2)} - e_1)(h_{1,\sigma(1)} - h_{1,1})].$$

For the second claim, we privilege the half-edge $h_{2,1}$ and consider the long exact sequence arising from explosion of the vertex v_2 , calculating that

$$\begin{aligned} \psi([A_2])_{h_{2,4}} &= \left[\sum_{\sigma(2)=4} \text{sgn}(\sigma)(e_1 - e_{\sigma(3)})(h_{1,\sigma(1)} - h_{1,1}) \right] \\ &= \left[\sum_{\sigma(2)=4, \sigma(4)=1} \text{sgn}(\sigma)(e_1 - e_{\sigma(3)})(h_{1,\sigma(1)} - h_{1,1}) \right] \\ &= \alpha_{123}. \end{aligned}$$

Similarly, we have $\psi([A_2])_{h_{2,2}} = \alpha_{134}$ and $\psi([A_2])_{h_{2,3}} = -\alpha_{124}$. It follows from the calculation made in the proof of Proposition 3.10 that $\psi([A_2])$ generates $\ker(\delta)$, implying the claim. $\square$

Remark 3.15. One shows easily that $(e_i - e_j)[A_2]$ is a difference of two external products of star classes. This analogue of the Q -relation implies that, modulo tori, the action of $\mathbb{Z}[E]$ on a theta class factors through the quotient identifying e_i for $1 \leq i \leq 4$, implying that the graded dimension of the $\mathbb{Z}[E]$ -submodule generated by a theta class grows at a strictly slower rate than $\dim H_2(B_k(\Gamma))$ (see [ADCK20, Thm. 1.2]). Combined with Theorem 1.1, this observation provides a hands-on verification of the degree 2 planar case of [ADCK21, Thm. 1.1], which asserts that toric classes always dominate in the limit $k \rightarrow \infty$.

Remark 3.16. The star and theta classes are the first two examples of a uniform construction outputting an element of $H_n(B_{n+1}(\mathbb{K}_{n,n+2}))$ for every $n > 0$. We defer a systematic study of these classes to future work.

4. DECOMPOSING GRAPHS

In this section, we discuss two methods of reducing the complexity of a graph, which form the basis for various inductive arguments to come. The first technique involves the removal of a subset of vertices, while the second involves replacing a subgraph by an edge. The first is drawn from classical graph theory, and the second is discussed at greater leisure in [ADCK20].

4.1. Connectivity and cuts. Classical connectivity theory for graphs only behaves well after restricting to combinatorially well-behaved classes of graphs—simplicial graphs, for example. According to our conventions, not all graphs are simplicial complexes, since we allow self-loops and multiple edges, but a simple device will allow us to circumvent this difficulty.

Definition 4.1. A *minimal simplicial model* for Γ is a simplicial graph Γ_Δ homeomorphic to Γ such that any smoothing with source Γ_Δ and simplicial target is an isomorphism.

In view of the following standard result, we typically (and abusively) refer to Γ_Δ as *the* minimal simplicial model.

Proposition 4.2. *Every graph admits a minimal simplicial model, which is unique up to isomorphism.*

Concretely, as long as Γ has no component homeomorphic to S^1 , the minimal simplicial model Γ_Δ may be obtained by the following three-step process: first, smooth all bivalent vertices of Γ ; second, add a bivalent vertex to each self-loop of the resulting graph; third, add a bivalent vertex to all but one of each set of multiple edges of the resulting graph.

Definition 4.3. A non-singleton graph Γ is (topologically) *k-connected* if any two distinct vertices of Γ_Δ may be joined by k paths pairwise disjoint away from the endpoints. By convention, the singleton graph is 1-connected but not k -connected for any $k > 1$.

In the special cases $k = 1, 2, 3$, we at times use the respective terms *connected*, *biconnected*, and *triconnected*.

As illustrated by the classical theorem due to Menger [Men27], connectivity is intimately related to the concept of a cut in a graph.

Definition 4.4. A *k-cut* in Γ is a subset $S \subseteq V$ of cardinality k such that the complement of the open star of S has at least two connected components, each containing a vertex. If S is a k -cut, an *S-component* of Γ is the closure in Γ of a connected component of $\Gamma \setminus S$.

The set of k -cuts of a graph depends crucially on the combinatorial structure; for example, any tail may be subdivided to contain a bivalent 1-cut. This pathology cannot occur in Γ_Δ , so we define $N_1(\Gamma)$ to be one plus the number of 1-cuts in Γ_Δ , where Γ is a connected graph.

Observation 4.5. The parameter $N_1(\Gamma)$ has the following properties.

- (1) If $N_1(\Gamma) > 1$, then Γ_Δ admits a 1-cut $\{x\}$ such that $N_1(\Delta) < N_1(\Gamma)$ for every $\{x\}$ -component Δ .

- (2) The equality $N_1(\Gamma) = 1$ holds if and only if Γ is biconnected, a singleton, or homeomorphic to $K_{1,1}$.

As we shall see, the parameter $N_1(\Gamma)$ is useful in inductive arguments. We shall also make use of an analogue in the case $k = 2$, for which we require an auxiliary definition.

Definition 4.6. Let $\{x, y\}$ be a 2-cut in Γ and Δ an $\{x, y\}$ -component. The *completion* of Δ is the graph $\overline{\Delta}$ obtained from Δ by adding an edge e_{xy} joining x and y .

The following result is a consequence of the combinatorial decomposition theory of [CE80], the details of which we elide.

Theorem 4.7 (Cunningham–Edmonds). *There is a parameter $N_2(\Gamma)$ associated to a biconnected graph Γ with the following properties.*

- (1) If $N_2(\Gamma) > 1$, then Γ_Δ admits a 2-cut $\{x, y\}$ such that $N_2(\overline{\Delta}) < N_2(\Gamma)$ for every $\{x, y\}$ -component Δ .
- (2) The equality $N_2(\Gamma) = 1$ holds if and only if Γ is triconnected, a cycle, or homeomorphic to Θ_n for some $n \geq 3$.

4.2. Surgery. In most circumstances, configuration spaces only enjoy functoriality for continuous injections. As we now explain, extra functoriality is available at the level of homology for configuration spaces of graphs.

Definition 4.8. Let $\Delta \subseteq \Gamma$ be a connected subgraph with two distinct distinguished vertices $\{x, y\}$. The result of *surgery on Γ along Δ* (using the vertices $\{x, y\}$) is the graph obtained by replacing Δ with a single edge e_{xy} connecting x and y .




FIGURE 6. Depiction of a surgery

Although a surgery is not a graph morphism in the sense given above, we nevertheless think of it as a type of morphism, writing $\sigma : \Gamma \dashrightarrow \Gamma'$ when Γ' is the result of surgery on Γ .

Example 4.9. Given a 2-cut $\{x, y\}$ in a biconnected graph Γ , there is a canonical surgery onto the completion of any $\{x, y\}$ -component of Γ .

Our notation is justified by a certain non-obvious functoriality. Note that, given a path from x to y in Δ , there results a non-canonical topological embedding $\iota : \Gamma' \rightarrow \Gamma$, called a *section* of the surgery σ .

Proposition 4.10. *Let $\sigma : \Gamma \dashrightarrow \Gamma'$ be a surgery. There is a canonical map of bigraded Abelian groups $\sigma_* : H_*(B(\Gamma)) \rightarrow H_*(B(\Gamma'))$ that is a retraction of ι_* for*

any section $\iota : \Gamma' \rightarrow \Gamma$. Moreover, σ_* is compatible with edge stabilization via the ring homomorphism

$$\sigma_*(e) = \begin{cases} e & e \notin E(\Delta) \\ e_{xy} & e \in E(\Delta). \end{cases}$$

This result is a consequence of the proof of [ADCK20, Lem. 4.17], where it is shown that σ_* is induced by a map $S(\sigma)$ at the level of Świątkowski complexes. Concretely, the map $S(\sigma)$ sends edges of Δ and vertices of Δ different from x and y to e_{xy} ; sends half-edges of Δ incident on x to the unique half-edge of e_{xy} incident on x (resp. y); annihilates all other half-edges of Δ ; and acts as the identity on vertices, edges, and half-edges not lying in Δ .

We close with two results concerning the difference between a homology class α and its modification $\iota_*\sigma_*(\alpha)$ after surgery. In both statements, we begin with a decomposition $\Gamma = \Delta \cup \Delta'$, where Δ and Δ' are subgraphs with Δ connected.

Proposition 4.11. *Suppose that Δ and Δ' intersect in the single vertex x , choose a vertex $x \neq y \in \Delta$ arbitrarily, and let $\sigma : \Gamma \dashrightarrow \Gamma'$ denote the resulting surgery along Δ . Fix a star cycle a , a loop cycle b , and $p, q \in \mathbb{Z}[E]$. If the support of a lies $\Delta' \setminus \{x\}$, then $(1 - \iota_*\sigma_*)([pa]) = 0$ for any section $\iota : \Gamma' \rightarrow \Gamma$ of σ (resp. b , $[qb]$).*

Proof. Since x does not lie in the support of a , we may assume by connectivity of Δ that p involves only edges lying in Δ' , in which case the claim holds by inspection (resp. b , q). $\square$

Proposition 4.12. *Suppose that Δ and Δ' intersect in the pair of distinct vertices $\{x, y\}$, and let $\sigma : \Gamma \dashrightarrow \Gamma'$ denote the resulting surgery along Δ . Fix a loop cycle b in Γ and $q \in \mathbb{Z}[E]$, and suppose that the support of b does not lie entirely in Δ . There is a section $\iota : \Gamma' \rightarrow \Gamma$ of σ such that $(1 - \iota_*\sigma_*)([qb])$ is a sum of stabilized star classes represented by cycles with support lying in Δ .*

Proof. By assumption, the support of b intersects Δ either in a subset of $\{x, y\}$ or in a path from x to y . Choose ι such that this path coincides with $\iota(e_{xy})$ in the latter case and arbitrarily otherwise. By inspection, $[qb] = \iota_*\sigma_*([qb])$ in the special case where q involves only edges lying either in Δ' or in the support of b . By connectivity and the Q-relation, q may be assumed of this form at the cost of introducing star classes with the desired support property. $\square$

5. PESKY CYCLES

In this section, we reformulate Theorem 1.1 in terms of the surjectivity of a certain map. We then prove Proposition 5.10, which describes generators for the cokernel of the map in question. These “pesky” cycles are the main players in the remainder of the paper, which is devoted to the proof of their vanishing.

5.1. Reformulation. Given a graph Γ , we define $M(\Gamma) \subseteq H_2(B(\Gamma))$ to be the $\mathbb{Z}[E]$ -submodule generated by theta classes and external products of loop classes and star classes; thus, the conclusion of Theorem 1.1 is that $M(\Gamma) = H_2(B(\Gamma))$. Given a bivalent vertex w , we consider the long exact sequence

$$\cdots \rightarrow H_2(B_k(\Gamma_w)) \xrightarrow{\iota_*} H_2(B_k(\Gamma)) \xrightarrow{\psi} H_1(B_{k-1}(\Gamma_w)) \xrightarrow{\delta} H_1(B_k(\Gamma_w)) \rightarrow \cdots$$

of Proposition 2.6. We will prove the following result concerning this sequence.

Theorem 5.1. *For any connected planar graph Γ and bivalent vertex w , the restriction $\psi : M(\Gamma) \rightarrow \ker(\delta)$ is surjective.*

We now clarify the relationship between this result and the main theorem, beginning with the following simple observation.

Lemma 5.2. *Let Γ_1 and Γ_2 be planar graphs. The conclusion of Theorem 1.1 holds for $\Gamma_1 \sqcup \Gamma_2$ if and only if it holds for Γ_1 and Γ_2 .*

Proof. The claim follows from the homeomorphism $B(\Gamma_1 \sqcup \Gamma_2) \cong B(\Gamma_1) \times B(\Gamma_2)$ and the Künneth isomorphism. The latter holds integrally since $H_1(B(\Gamma_i))$ is torsion-free for $i \in \{1, 2\}$ by planarity [KP12, Cor. 3.6]. $\square$

Proposition 5.3. *If Γ is a graph with a bivalent vertex w such that the conclusion of Theorem 1.1 holds for Γ_w , then the conclusion of Theorem 1.1 also holds for Γ provided either*

- (1) Γ_w is connected and the conclusion of Theorem 5.1 holds for Γ and w , or
- (2) Γ_w is disconnected.

Proof. In the first case, we have $\iota_*(H_2(B(\Gamma_w))) = \iota_*(M(\Gamma_w)) \subseteq M(\Gamma)$. From the long exact sequence above, the map ψ induces an isomorphism

$$\frac{H_2(B(\Gamma))}{\iota_*(H_2(B(\Gamma_w)))} \cong \ker(\delta) \implies \frac{H_2(B(\Gamma))}{M(\Gamma)} \cong \frac{\ker(\delta)}{\psi(M(\Gamma))} = 0,$$

whence $H_2(B(\Gamma)) = M(\Gamma)$, as desired.

In the second case, Lemma 5.2 implies that the conclusion of Theorem 1.1 holds for each component of Γ_w , and the claim follows from [ADCK19, Prop. 5.22]. $\square$

Corollary 5.4. *Theorems 1.1 and 5.1 are equivalent.*

Proof. The forward implication is immediate from exactness. For the reverse implication, we proceed by induction on the first Betti number β_1 of Γ . The base case of $\beta_1 = 0$ is that of a tree, which is well known (see also Example 6.1). For the induction step, choose a bivalent vertex w such that Γ is connected, subdividing if necessary. Since the conclusion of Theorem 1.1 holds for Γ_w by induction, Proposition 5.3 yields the conclusion. $\square$

In pursuing Theorem 5.1, we will use the following criterion repeatedly. Here we abusively write w and w' for the vertices of Γ_w corresponding to the two half-edges of Γ incident on w .

Lemma 5.5 (Path argument). *Let γ be a path from w to w' in Γ_w . A loop or star class lies in $\psi(M(\Gamma))$ if it admits a representing cycle with support disjoint from γ .*

Proof. Let a denote the representing cycle. From the path γ , we obtain a loop in Γ and thereby a loop cycle. By the assumption on the support of a , the external product of a with this loop cycle is defined, and the homology class of the external product lies in $M(\Gamma) \cap \psi^{-1}([a])$ by inspection of the definition of ψ . $\square$

As an immediate application, we obtain the following special case of Theorem 5.1.

Corollary 5.6. *Let Γ be a connected graph with a bivalent vertex w . If the edges incident on w form a loop, then the conclusion of Theorem 5.1 holds for Γ and w .*




FIGURE 7. Schematic depiction of Notation 5.8

Proof. Let e and e' denote the edges incident on w . By assumption, there is a vertex $u \neq w$ such that e and e' are also incident on u . Then e and e' lie in distinct components of $(\Gamma_w)_u$, so Lemma 2.7 implies that $\ker(\delta)$ lies in the image of $H_1(B((\Gamma_w)_u))$, and the path argument applies. $\square$

We close with an analogue of the path argument for star classes.

Lemma 5.7. *If a is a non-rigid star cycle, then $[a] \in \psi(M(\Gamma))$.*

Proof. Let u denote the support of a . By non-rigidity, there is an embedding of a subdivision of Θ_3 into Γ_w under which u is the image of an essential vertex. Denote the image of the other essential vertex by u' , and choose a path in Γ_w from w to w' . If the path avoids u' , then the θ -relation and the path argument yield the conclusion. If the path contains u' , then our embedding extends to an embedding of a subdivision of Θ_4 into Γ , and Lemma 3.12 implies the claim. $\square$

5.2. Standard and pesky cycles. The goal of this section is to describe generators for the cokernel of the map of Theorem 5.1. Before giving a precise formulation in Proposition 5.10 below, we require a few preliminary definitions and results. We begin by establishing notation that we maintain for the remainder of the paper.

Notation 5.8 (see Figure 7). Let Γ be a connected graph. Fix an edge e with vertices v and v' , and subdivide e by adding a bivalent vertex w . Abusively, we write w and w' for the resulting vertices of Γ_w and e and e' for the resulting edges, so that e has vertices v and w (resp. e' , v' and w'). We write $\Gamma_e := \Gamma \setminus e$ for the graph given by the complement of the (open) edge e .

As long as Γ has an essential vertex, we may assume after smoothing that v and v' are essential in Γ and in Γ_w ; however, either vertex may be bivalent in Γ_e . In light of Corollary 5.6, we further assume that $v \neq v'$, i.e., that e is not a self-loop in Γ .

Definition 5.9. Let Γ_w be as in Notation 5.8.

- (1) A *standard cycle* is a weight-homogeneous 1-cycle $c \in S_1(\Gamma_w)$ of the form

$$c = \sum_{i=1}^r p_i a_i + \sum_{j=1}^s q_j b_j,$$

where the a_i are star cycles, the b_j loop cycles, and the p_i and q_j polynomials in the edges of Γ_w . The *support* of c is the union of the supports of the a_i and b_j .

- (2) We define the following potential properties of a standard cycle c .
(P) Every **path** from w to w' in Γ_w intersects the support of every a_i and every b_j .

- (E) If an **edge** e is involved in p_i (resp. q_j), then e is adjacent to (resp. contained in) the support of a_i (resp. b_j).
 - (S) No 1-cut of Γ_w separating e and e' is contained in the **support** of c .
 - (K) The class $[c]$ lies in the **kernel** of δ .
- (3) If c has all four of these properties, then we say that c is *pesky*.

The interest of pesky cycles lies in the following result.

Proposition 5.10. *The quotient $\ker(\delta)/\psi(M(\Gamma))$ is generated by the images of pesky cycles with no star summands.*

The proof requires a few simple lemmas.

Lemma 5.11. *If c is a standard cycle satisfying **K**, then c is homologous to a standard cycle satisfying **S** and **E**.*

Proof. Let W be the set of 1-cuts of Γ_w separating e and e' . Repeated use of Lemma 2.7 shows that $[c]$ lies in the image of $H_1(B((\Gamma_w)_W))$. By Proposition 3.4, we may represent $[c]$ by a standard cycle with support avoiding W , and the Q -relation and connectedness of Γ_w imply that we may take this standard cycle to satisfy **E**. $\square$

Lemma 5.12. *Let c be a standard cycle. Then $c = c_1 + c_2$, where $[c_1] \in \psi(M(\Gamma))$ and c_2 is a standard cycle satisfying **P**. Moreover, if c satisfies **E**, **S**, or **K**, then so does c_2 .*

Proof. Let I denote the set of indices i for which there is a path from w to w' in Γ_w avoiding the support of a_i (resp. J, j, b_j), and define

$$c_1 = \sum_{i \in I} p_i a_i + \sum_{j \in J} q_j b_j$$

$$c_2 = c - c_1.$$

The path argument implies that $[c_1] \in \psi(M(\Gamma))$, and, by construction, c_2 is a standard cycle satisfying **P** and further satisfying **E** or **S** if c did. For the last claim, we note that $\delta \circ \psi = 0$, so $\delta([c]) = \delta([c_2])$. $\square$

Lemma 5.13. *A pesky cycle contains no star summands.*

Proof. Supposing otherwise, it follows from **P** that the support of the star cycle in question is a 1-cut separating e and e' , violating **S**. $\square$

Proof of Proposition 5.10. The claim is trivial unless Γ has an essential vertex, in which case Lemmas 5.11, 5.12, and 5.13 imply the claim. $\square$

6. EXAMPLES AND A FIRST REDUCTION

The proof of Theorem 5.1 proceeds through two inductions. In this section, we carry out the first (and simplest) of these inductions after considering various examples related to base cases.

6.1. Rogues' gallery. We pause to consider a few (partially redundant) examples that will be of use in what follows.

Example 6.1. If Γ is a tree, then the conclusion of Theorem 1.1 holds for Γ . This well known claim follows (for example) by repeated application of [ADCK19, Prop. 5.22].

Example 6.2. If Γ is obtained from a tree by attaching self-loops, then the conclusion of Theorem 1.1 holds for Γ . This claim follows from Proposition 5.3 by induction on the number of self-loops using Example 6.1 as the base case and Corollary 5.6 for the induction step.

Example 6.3. If the first Betti number b_1 of Γ is equal to 1, then the conclusion of Theorem 1.1 holds for Γ . This claim follows from Example 6.1 and Proposition 5.10 after exploding a bivalent vertex w such that Γ_w is a tree. Indeed, a tree admits no loop cycles and no star cycles satisfying **S**, hence no nonzero pesky cycles.

Example 6.4. If Γ has exactly two essential vertices, each trivalent, then $H_2(B(\Gamma))$ is generated by external products of star classes and loop classes by [ADCK19, Prop. 5.25], so the conclusion of Theorem 1.1 holds in this case.

Example 6.5. More generally, the conclusion of Theorem 1.1 holds if Γ has exactly two essential vertices. Assume first that Γ has no self-loops (see Figure 8). Adding a bivalent vertex w (shown in green) to one of the non-tail edges, we observe that the support of every loop and star cycle in Γ_w contains one of the essential vertices of Γ_w , thereby violating **S**. Thus, Γ_w has no nonzero pesky cycles, and it follows from Proposition 5.10 that the conclusion of Theorem 5.1 holds. The claim in the case without self-loops now follows from Proposition 5.3 by induction on the first Betti number using Example 6.1 as a base case. The claim in general follows by repeated application of Corollary 5.6 as in Example 6.2.




FIGURE 8. The graph of Example 6.5 with auxiliary markings

Example 6.6. The conclusion of Theorem 1.1 holds for the graph obtained by adding an edge e_0 joining the blue vertices in the graph of Figure 8. The argument here is essentially the same; we induct on the first Betti number, reducing to the case $b_1 = 2$ by observing that any loop or star cycle in Γ_w violates either **P** or **S**. In the case $b_1 = 2$, adding the vertex w to e_0 again yields an explosion Γ_w with no nonzero pesky cycles, and, since the conclusion holds for Γ_w by Example 6.3, the claim follows.

Example 6.7. The conclusion of Theorem 1.1 holds for the graph obtained by adding an edge e_0 joining a red vertex in the graph of Figure 8 either to an essential vertex or to the other red vertex. As before, the claim follows by induction on b_1 via a sequence of explosions such that Γ_w has no nonzero pesky cycles.

We close with an example showing that the assumption of planarity in Theorem 1.1 is necessary.

Example 6.8. Let Δ be the union of a complete bipartite graph $K_{3,3}$ and a star graph S_3 along a set of three pairwise non-adjacent vertices—see Figure 9. Subdividing and exploding one of the edges S_3 , we make two claims: first, the loop cycle b shown in red is pesky; second, $M(\Delta)$ vanishes in weight 2. In light of Proposition 5.10, these claims together imply that the conclusion of Theorem 5.1, hence that of Theorem 1.1, fails for Δ .




FIGURE 9. The graph Δ and a loop cycle in red

For the first claim, only $\mathbf{K}$ is not immediate. Writing $\beta = [b]$, the Q-relation implies that

$$(e - e')\beta = (e - e_0)\beta - (e' - e_0)\beta = \alpha - \alpha',$$

where e_0 is a fixed edge in the support of b , and α and α' are star classes. Since $\alpha = \alpha'$ by Example 3.7, the claim follows.

For the second claim, $M(\Delta)$ is spanned in weight 2 by external products of loop classes, since theta classes or external products involving a star class require at least 3 particles (this much is true for any graph). On the other hand, there is no pair of disjoint loop cycles in Δ , since every loop in Δ involves at least 4 of the 7 vertices.

Remark 6.9. The calculation of Example 6.8 exhibits a previously unknown type of class.

6.2. Reduction to the biconnected case. We close this section with a reduction permitting the full use of Proposition 5.10. The reader is reminded that we maintain Notation 5.8.

Lemma 6.10. *If Theorem 5.1 holds under the further assumption that Γ_e is biconnected, then it holds in general.*

Proof. We proceed by induction on $N_1(\Gamma_e)$. In the base case of $N_1(\Gamma_e) = 1$, either Γ_e is biconnected, in which case our assumption applies; or Γ_e is an isolated vertex, in which case Γ is homeomorphic to a cycle graph; or Γ_e is homeomorphic to an interval. In this latter case, either Γ is homeomorphic to a cycle graph, or Γ has two essential vertices, each trivalent. The examples of Section 6.1 encompass these cases.

Given a 1-cut u of Γ_e (which is necessarily a 1-cut of Γ_w , and which may coincide with v or v'), [KP12, Lem. 3.11] supplies the decomposition

$$H_1(B_k(\Gamma_w)) \cong \left(\bigoplus_{\ell=1}^m H_1(B_k(\Delta_\ell)) \oplus \mathbb{Z}^{N(k, \Gamma_w, u)} \right)$$

of Abelian groups, where the Δ_ℓ are the u -components of Γ_w , and the last term is spanned by star classes at u (see also [ADCK19, Lem. 3.11]). We consider the relationship between the terms of this direct sum decomposition and those of a pesky cycle $c = \sum_{i=1}^r p_i a_i + \sum_{j=1}^s q_j b_j$.

There are two cases. If u separates e and e' , then **S** implies that u does not lie in the support of c . Thus, each term lies in one of the first m summands. Since $H_1(B_k(\Delta_\ell)) \subseteq \psi(M(\Gamma))$ by the path argument and Proposition 3.4 if Δ_ℓ does not contain e or e' , we may assume by naturality of ψ that $\ell = 2$. By Proposition 4.11, we may write $[c] = \iota_*^1 \sigma_*^1([c]) + \iota_*^2 \sigma_*^2([c])$, where σ^ℓ is surgery along Δ_ℓ using the vertices u and either w or w' , as appropriate. The claim now follows from $\mathbb{Z}[E]$ -linearity of σ_*^ℓ , the induction hypothesis, and naturality of ψ .

If u does not separate e and e' , we may assume by the path argument and naturality of ψ that $\ell = 1$, and the claim follows by induction. $\square$

The remainder of the paper is devoted to establishing Theorem 5.1 under the assumption that Γ_e is biconnected. Although necessarily more elaborate in its details, the strategy is essentially the same, namely to reduce to the triconnected case by induction on the parameter $N_2(\Gamma_e)$. The primary advantage afforded by biconnectivity is that, through **E** and Proposition 5.10, each summand of a pesky cycle may be taken to be a loop cycle stabilized at edges internal to its own support.

7. BASE CASES

The goal of this section is to prove Theorem 5.1 under the assumption that Γ_e is triconnected, a cycle graph, or a theta graph (we maintain Notation 5.8). These cases will form the basis for the induction carried out in Section 8.

7.1. The triconnected case. In this section, we prove the following result.

Lemma 7.1. *Suppose that Γ is planar and Γ_e is triconnected. If b is a loop cycle in Γ_w with support disjoint from $\{v, v'\}$, then $[b] \in \psi(M(\Gamma))$.*

Before turning to the proof, we establish notation. Using the connectivity assumption, we may find paths γ_1 and γ_2 from v to v' in Γ_e disjoint away from $\{v, v'\}$. This claim follows after applying the definition of triconnectivity to the nearest essential vertices to v and v' (which may or may not be v and v' themselves).

By the path argument, we may assume that b satisfies **P**, so the support of b intersects both paths, and we write x_i and y_i for the first and last vertices of the intersection with γ_i . Although no two points with different index coincide, it may be that $x_i = y_i$. Thus, Γ contains a topologically embedded *angel graph*—see Figure 10.




FIGURE 10. Angels

Note that the four points carry a natural cyclic ordering up to reversal of orientation, which must be (x_1, y_1, y_2, x_2) by planarity; in particular, there is no case in which the pairs of endpoints are linked.

The core of the argument is an analysis of such embedded graphs in the presence of sufficient connectivity. For the sake of narrative flow, this analysis is deferred to Section 7.2 below. There we prove Proposition 7.6, which implies that our embedded angel graph extends to an embedding of one of the graphs depicted in Figure 11 (we have shown only the “generic” case—in the case where $x_1 = y_1$ and $x_2 = y_2$, only the first possibility occurs, and so forth).




FIGURE 11. Some split angels

Proof of Lemma 7.1. We refer to Figure 11. The first case expresses b as the sum of loop cycles b_1 and b_2 to which the path argument applies—see Figure 12. Note that this case includes the degenerate case where $x_1 = y_1$ and $x_2 = y_2$.

The second case expresses b as the sum of loop cycles b_1 and b_2 such that the path argument applies to b_2 and the support of b_1 lies in an embedded angel graph with x_1 strictly closer to y_1 . Similar remarks apply to the third case mutatis mutandis, while the fourth case expresses b as the sum of two loop cycles, each of whose supports lies in an embedded angel graph with x_i strictly closer to y_i for both i . We now proceed by induction on the sum of the edge lengths of the arcs (x_1, y_1) and (x_2, y_2) in b , the base case being the degenerate case already considered. $\square$

We are now in a position to complete the proof of our base cases.

Proposition 7.2. *The conclusion of Theorem 5.1 holds under the further assumption that Γ_e is biconnected and $N_2(\Gamma_e) = 1$.*

Proof. If Γ_e is a cycle graph or homeomorphic to Θ_n for some $n \geq 3$, then Γ is among the examples considered in Section 6.1, so we may assume by Theorem 4.7 that Γ_e is triconnected.

Let $c = \sum_{j=1}^s q_j b_j$ be a pesky cycle. Each b_j satisfies **P** by definition, and **S** implies that the support of b_j does not intersect $\{v, v'\}$. Thus, Lemma 7.1 implies that $[b_j] \in \psi(M(\Gamma))$ for each j , whence $[c] \in \psi(M(\Gamma))$. Proposition 5.10 now implies that $\ker(\delta)/\psi(M(\Gamma)) = 0$, as desired. $\square$




FIGURE 12. Splitting a loop cycle

7.2. Angelic graphs. In this section, we supply the missing ingredient in the proof of Lemma 7.1.

We employ uniform notation when dealing with angel graphs, some of which is indicated in Figure 10. The corner vertices are denoted v and v' , and the edge connecting them is denoted e . The central loop is denoted S , and the two unique paths from v to v' avoiding e and involving at most two vertices of S are denoted γ_1 and γ_2 , respectively. Considering γ_i as oriented from v to v' , the first point of intersection of γ_i with S is denoted x_i and the last is denoted y_i (it may be that $x_i = y_i$).

Definition 7.3. A *split angel* is an angel together an additional edge connecting two distinct components of $S \setminus \{x_1, y_1, x_2, y_2\}$.

Some split angels are depicted in Figure 11.

Definition 7.4. An *angelic graph* is a topological embedding $\Omega \rightarrow \Gamma$, where Ω is an angel. If the embedding extends to an embedding of a split angel, then the angelic graph is said to be *split*.

We abuse notation in referring to an angelic graph by the letter Γ , as well as by identifying parts of Ω with their images in Γ .

Definition 7.5. The angelic graphs $\Omega_1 \rightarrow \Gamma$ and $\Omega_2 \rightarrow \Gamma$ are called *equivalent* if there are commuting diagrams of the following form:

$$\begin{array}{ccc} [0, 1] & \xrightarrow{e} & \Omega_1 \\ e \downarrow & & \downarrow \\ \Omega_2 & \longrightarrow & \Gamma \end{array} \qquad \begin{array}{ccc} S^1 & \xrightarrow{S} & \Omega_1 \\ s \downarrow & & \downarrow \\ \Omega_2 & \longrightarrow & \Gamma. \end{array}$$

Proposition 7.6. Let Γ be an angelic graph satisfying the following three conditions:

- (1) Γ is planar
- (2) $\Gamma \setminus e$ is triconnected, and
- (3) v and v' lie in distinct components of $\Gamma \setminus (e \cup S)$.

Then Γ is split up to equivalence.

The key input to the proof will be the following result, which should be thought of as an infinitesimal version of Proposition 7.6.

Lemma 7.7. Let Γ be an angelic graph satisfying the hypotheses of Proposition 7.6. Up to equivalence, each component of $S \setminus \{y_1, x_2\}$ contains an essential vertex.

Proof. If $x_1 \neq y_1$ and $x_2 \neq y_2$, then all four are essential, and there is nothing to show, so we may assume that $x_2 = y_2$. Denote the components in question by S' and S'' . We may assume that S' contains an essential vertex; indeed, in a minimal simplicial representative of Γ , S contains some vertex other than y_1 and x_2 , and assuming all such vertices to be bivalent implies that $\{y_1, x_2\}$ is a 2-cut, contradicting (2).

We may further assume that S'' contains no essential vertex, in which case (2) guarantees paths in Γ connecting S' to each component of the complement in Ω of the star of $\{y_1, x_2\}$, which furthermore avoid e , y_1 , x_2 , and S'' . For the last, we have used that S'' is contained in the open star of $\{y_1, x_2\}$ in a minimal simplicial representative.




FIGURE 13. The connected components S' and S'' of $S \setminus \{y_1, x_2\}$

If $x_1 = y_1$ and $x_2 = y_2$, then there are two components and hence two paths. By (3) may assume that these paths are disjoint away from S and that neither intersects both γ_1 and γ_2 away from S . Using these paths, we obtain a topologically embedded $K_{3,3}$, contradicting (1)—see Figure 14. We take this case as the base case in an induction on the edge length of $\gamma_1 \cap S$ as calculated in Γ (in Ω , this edge length is always either 0 or 1).

For the induction step, assume that $x_1 \neq y_1$, so there is only one path γ . A similar contradiction of (1) as above is achieved unless both endpoints of γ lie on γ_1 . Replacing a segment of γ_1 with γ now produces an equivalent angelic graph in which $\gamma_1 \cap S$ has strictly shorter edge length—see Figure 15. $\square$




FIGURE 14. Embedded copies of $K_{3,3}$



FIGURE 15. A path in red connecting the two components of the complement of the star of $\{y_1, x_2\}$ in Ω , together with the resulting shortening of $\gamma_1 \cap S$

Proof of Proposition 7.6. By Lemma 7.7, we may assume that each component of $S \setminus \{y_1, x_2\}$ contains an essential vertex. Thus, by (2), there is a path connecting these components and avoiding e , y_1 , and x_2 . By (3), we may assume that this path does not intersect both γ_1 and γ_2 away from S . There are now six possibilities, as depicted in Figure 16. We conclude that, if Γ is not split, then Γ is equivalent to an angelic graph in which the edge length either of $\gamma_1 \cap S$ or of $\gamma_2 \cap S$ is strictly smaller, as in the proof of Lemma 7.7. An induction on the sum of these edge lengths completes the proof. $\square$




FIGURE 16. Splitting or reducing arc length

8. INDUCTION STEP

In this section, we complete the proof of Theorem 5.1. Maintaining Notation 5.8, we assume that Γ_e is biconnected and proceed by induction on $N_2(\Gamma_e)$, the base case being Proposition 7.2.

8.1. Setup and reductions. We assume throughout this section and the next that Γ_e has a 2-cut $\{x, y\}$ with $\{x, y\}$ -components $\{\Delta_i\}_{i=1}^m$ such that $N_2(\overline{\Delta}_i) < N_2(\Gamma_e)$ for each i . Given a loop cycle b in Γ_w (whose support necessarily lies in Γ_e), we say that b is *local* if its support lies entirely in a single $\{x, y\}$ -component; otherwise, we say that b is *global* (see Figure 17). Note that the support of a global loop cycle is necessarily contained in the union of exactly two $\{x, y\}$ -components.




FIGURE 17. A local loop cycle b and a global loop cycle b'

We begin by dealing with the local case.

Lemma 8.1. *Suppose that Γ is planar. If b is a local loop cycle in Γ_w with support disjoint from $\{v, v'\}$, then $[b] \in \psi(M(\Gamma))$.*

Proof. Without loss of generality, the support of b lies in Δ_1 . By inspection, we have $[b] = \iota_* \sigma_*([b])$, where σ is the surgery along $\Gamma' = \bigcup_{i=2}^m \Delta_i$ using the vertices $\{x, y\}$, and ι is any section.

If one of v or v' , say v , is contained in Δ_1 , then we may regard the completion $\overline{\Delta}_1$ as Δ_e , where Δ is the graph obtained by adding an edge between v and a bivalent vertex v' added to the edge e_{xy} . The claim now follows from Example 4.9 and Proposition 7.2 by induction on $N_2(\Gamma_e)$. Note that we use the assumption that b is local to ensure that the loop cycle resulting from surgery again has support disjoint from $\{v, v'\}$, permitting the induction.

If neither v nor v' is contained in Δ_1 , then we may assume by the path argument that $v \in \Delta_2$ and $v' \in \Delta_3$. For the same reason, writing S for the support of b , we may further assume that $\{x, y\} \subseteq S$. Since $\Delta_1 \setminus \{x, y\}$ is connected, the two components of $S \setminus \{x, y\}$ are joined by a path in Δ_1 . Using this path, we may write $b = b_1 + b_2$, where each b_i is a loop cycle with support avoiding $\{v, v'\}$ and not containing $\{x, y\}$, hence yielding to the path argument. $\square$

A similar tactic dispenses with global loop cycles in certain situations.

Lemma 8.2. *Suppose that Γ is planar. If b is a global loop cycle in Γ_w with support disjoint from $\{v, v'\}$, and if v and v' lie in a common $\{x, y\}$ -component, then $[b] \in \psi(M(\Gamma))$.*

Proof. We begin by observing that $\{x, y\} \cap \{v, v'\} = \emptyset$ by our assumptions on b , so v and v' in fact lie in the same component of $\Gamma \setminus \{x, y\}$. If b does not intersect this component, then the claim follows from the path argument; otherwise, the inductive surgery argument of Lemma 8.1 applies. Here, we instead use the fact that v and v' lie in the same component of $\Gamma \setminus \{x, y\}$ to ensure that the loop cycle resulting from surgery again has support disjoint from $\{v, v'\}$. $\square$

8.2. Conclusion. We begin with a simple observation concerning star classes, which is a consequence of [ADCK19, Lemma C.14].

Lemma 8.3. *Let $\{x, y\}$ be a 2-cut in Δ with $\{x, y\}$ -components $\{\Delta_i\}_{i=1}^m$ and a and a' star cycles in Δ . Write i_j for the index such that the j th half-edge involved in a lies in Δ_{i_j} (resp. i'_j , a'). If $\{i_j\}_{j=1}^3 \neq \{i'_j\}_{j=1}^3$, then $[a_1]$ and $[a_2]$ are independent in the sense that*

$$m[a_1] + n[a_2] = 0 \implies m[a_1] = 0.$$

We now consider the following simple case.

Lemma 8.4. *Suppose that Γ is planar and the 2-cut $\{x, y\}$ is disjoint from and separates v and v' . Let b be a global loop cycle, e_0 an edge lying in the support of b , and m and n non-negative integers. If the cycle $ne_0^m b$ satisfies **K**, then $n = 0$.*

Proof. Suppose that $ne_0^m b$ satisfies **K**. Without loss of generality, the support of b lies in $\Delta_1 \cup \Delta_2$. By **K** and the Q-relation, we have

$$0 = ne_0^m(e - e')[b] = \pm ne_0^m([a] - [a']),$$

where a is a star cycle satisfying the following conditions with respect to v (resp. a' , v'):

- (1) if v lies in Δ_i for some $i \in \{1, 2\}$, then a involves at least two half-edges lying in Δ_i (resp. v' , a');
- (2) if v does not lie in Δ_i for any $i \in \{1, 2\}$, then a involves a single half-edge in each of Δ_1 , Δ_2 , and the $\{x, y\}$ -component of v .

These conditions guarantee that the hypothesis of Lemma 8.3 obtains. Since multiplication by e_0 is injective by [ADCK19, Prop. 5.21], it follows that n annihilates $[a] - [a']$ and hence each class separately. We conclude that $n = 0$ by planarity and [KP12, Cor. 3.6]. $\square$

Note that Lemma 8.4 does not require the support of b to be disjoint from $\{v, v'\}$.

Proof of Theorem 5.1. By Lemma 6.10 and Proposition 7.2, it suffices to establish the conclusion under the further assumption that Γ_e is biconnected with a 2-cut $\{x, y\}$ such that $N_2(\overline{\Delta_i}) < N_2(\Gamma_e)$ for each i . Consider the pesky cycle $c = \sum_{j=1}^s q_j b_j$, where each b_j is a loop cycle. By biconnectivity and Proposition 5.10, it suffices to show that $[c] \in \psi(M(\Gamma))$. By Lemmas 8.1 and 8.2, we may further assume that each b_j is global and that $\{x, y\}$ separates v and v' .

Unless $c = 0$, since b_1 is global satisfying **S**, its support intersects Δ_1 (up to symmetry) in a path γ from x to y avoiding v and v' . We call a loop cycle *special* if the intersection of its support with Δ_1 is also γ . The support of an arbitrary global loop cycle b intersects Δ_1 either in a path from x to y or in the set $\{x, y\}$. A cycle of the former type differs from a special loop cycle by a sum of local loop cycles, while a cycle of the latter type is the difference of two special loop cycles. Therefore, at the cost of introducing stabilized local loop cycles, which lie in $\psi(M(\Gamma))$ by Lemma 8.1, each b_j may be assumed special. Moreover, fixing an edge e_0 lying in the path γ , we may take $q_j = e_0^{k-1}$ up to homology for each j . (and up to non-rigid star cycles which can be ignored by Lemma 5.7).

Let σ denote surgery along $\Gamma' = \cup_{i=2}^m \Delta_i$ using the vertices $\{x, y\}$. By Proposition 4.12, there is a section ι such that $(1 - \iota_* \sigma_*)([c])$ is a sum of a stabilized star classes, which lie in $\psi(M(\Gamma))$ by Lemma 5.7 (these star classes are necessarily non-rigid by biconnectivity). By inspection, since each b_j is special, we have

$$\iota_* \sigma_*([c]) = \sum_{j=1}^s e_0^{k-1} \iota_* \sigma_*([b_j]) = n e_0^{k-1} [b],$$

where b is the loop cycle with support given by the union of γ and $\iota(e_{xy})$ (note that we may have $s \neq n$, since each b_j carries an orientation). By $\mathbb{Z}[E]$ -linearity and the fact that $\ker(\delta)$ is $(e - e')$ -torsion, we conclude that the cycle $n e_0^{k-1} b$ satisfies **K**, whence $n = 0$ by Lemma 8.4. It follows that $[c] \in \psi(M(\Gamma))$, as claimed. $\square$

REFERENCES

- [Abr00] A. Abrams, *Configuration spaces of braid groups of graphs*, Ph.D. thesis, UC Berkeley, 2000.
- [ADCK19] B. H. An, G. C. Drummond-Cole, and B. Knudsen, *Subdivisional spaces and graph braid groups*, Doc. Math. **24** (2019), 1513–1583.
- [ADCK20] ———, *Edge stabilization in the homology of graph braid groups*, Geom. Topol. **24** (2020), no. 1, 421–469.
- [ADCK21] ———, *Asymptotic homology of graph braid groups*, Geom. Topol., to appear (2021).
- [Arn69] V. Arnold, *The cohomology ring of the group of dyed braids*, Mat. Zametki **5** (1969), 227–231.
- [CE80] W. H. Cunningham and J. Edmonds, *A combinatorial decomposition theory*, Can. J. Math. **32** (1980), no. 3.
- [CEF15] T. Church, J. S. Ellenberg, and B. Farb, *FI-modules and stability for representations of symmetric groups*, Duke Math. J. **164** (2015), no. 9, 1833–1910.
- [Chu12] T. Church, *Homological stability for configuration spaces of manifolds*, Invent. Math. **188** (2012), no. 2, 465–504.
- [CL18] S. Chettih and D. Lütgehetmann, *The homology of configuration spaces of trees with loops*, Alg. Geom. Topol. **18** (2018), no. 4, 2443–2469.
- [FS05] D. Farley and L. Sabalka, *Discrete Morse theory and graph braid groups*, Alg. Geom. Topol. **5** (2005), 1075–1109.
- [KP12] K. H. Ko and H. W. Park, *Characteristics of graph braid groups*, Discrete Comput. Geom. **48** (2012), no. 4, 915–963.

- [McD75] D. McDuff, *Configuration spaces of positive and negative particles*, Topology **14** (1975), 91–107.
- [Men27] K. Menger, *Zur allgemeinen Kurventheorie*, Fund. Math. **10** (1927), no. 1, 96–115.
- [Świ01] J. Świątkowski, *Estimates for homological dimension of configuration spaces of graphs*, Colloq. Math. **89** (2001), no. 1, 69–79.
- [WG17] John Wiltshire-Gordon, *Models for configuration space in a simplicial complex*, Colloquium Mathematicum **155** (2017).

E-mail address: `anbyhee@knu.ac.kr`

DEPARTMENT OF MATHEMATICS EDUCATION, TEACHERS COLLEGE, KYUNGPOOK NATIONAL UNIVERSITY, DAEGU, SOUTH KOREA

E-mail address: `b.knudsen@northeastern.edu`

DEPARTMENT OF MATHEMATICS, NORTHEASTERN UNIVERSITY, BOSTON, MA 02115, USA