Groups Geom. Dyn. (Online first)
DOI 10.4171/GGD/823

A graphical description of the BNS-invariants of Bestvina–Brady groups and the RAAG recognition problem

Yu-Chan Chang and Lorenzo Ruffoni

Abstract. A finitely presented Bestvina–Brady group (BBG) admits a presentation involving only commutators. We show that if a graph admits a certain type of spanning tree, then the associated BBG is a right-angled Artin group (RAAG). As an application, we obtain that the class of BBGs contains the class of RAAGs. On the other hand, we provide a criterion to certify that certain finitely presented BBGs are not isomorphic to RAAGs (or more general Artin groups). This is based on a description of the Bieri–Neumann–Strebel invariants of finitely presented BBGs in terms of separating subgraphs, analogous to the case of RAAGs. As an application, we characterize when the BBG associated with a 2-dimensional flag complex is a RAAG in terms of certain subgraphs.

Contents

1. Introduction

Let Γ be a finite simplicial graph and denote its vertex set and edge set by $V(\Gamma)$ and $E(\Gamma)$, respectively. The associated *right-angled Artin group* (RAAG) A_Γ is the group defined by the following finite presentation:

$$A_\Gamma = \langle V(\Gamma) \mid [v, w] \text{ whenever } (v, w) \in E(\Gamma) \rangle.$$

RAAGs have been a central object of study in geometric group theory because of the beautiful interplay between algebraic properties of the groups and combinatorial properties of

Mathematics Subject Classification 2020: 20F36 (primary); 20J05, 20F65, 20F05 (secondary).
Keywords: Bestvina–Brady group, BNS-invariant, flag complex, right-angled Artin group, subspaces arrangement, tree 2-spanner.

the defining graphs, and also because they contain many interesting subgroups, such as the fundamental group of many surfaces and 3-manifolds, and more generally, especially cubulated groups (see [23]).

The *RAAG recognition problem* consists in deciding whether a given group is a RAAG. Several authors have worked on this problem for various classes of groups, for instance, the pure symmetric automorphism groups of RAAGs in [13, 24], the pure symmetric outer automorphism groups of RAAGs in [16], and a certain class of subgroups of RAAGs and RACGs in [15] and of mapping class groups in [25]. An analogous recognition problem for right-angled Coxeter groups has been considered in [14]

However, the RAAG recognition problem is not easy to answer in general, even when the given group shares some essential properties with RAAGs. For example, the group G with the following presentation

$$G = \langle a, b, c, d, e \mid [a, b], [b, c], [c, d], [b^{-1}c, e] \rangle$$

is finitely presented with only commutator relators; it is CAT(0) and splits as a graph of free abelian groups. However, it is not a RAAG (see [31, Example 2.8]). Even more is true: Bridson [6] showed that there is no algorithm to determine whether or not a group presented by commutators is a RAAG, answering a question by Day and Wade [16, Question 1.2].

In this article, we study the RAAG recognition problem for a class of normal subgroups of RAAGs, namely, the *Bestvina–Brady groups* (BBGs). Let $\chi\colon A_\Gamma \to \mathbb{Z}$ be the homomorphism sending all the generators to 1. The BBG defined on Γ is the kernel of χ and is denoted by BB_Γ. For example, the group G from above is the BBG defined on the *trefoil graph* (see Figure 1). BBGs were introduced and studied in [3], and they have become popular as a source of pathological examples in the study of finiteness properties and cohomology of groups. For instance, some BBGs are finitely generated but not finitely presented, and there are some BBGs that are potential counterexamples to either the Eilenberg–Ganea conjecture or the Whitehead conjecture.

Inspired by the example of the group G from above, we are interested in understanding how much a BBG can be similar to a RAAG without being a RAAG. In particular, we are interested in a criterion that can be checked directly on the defining graph. It is well known that two RAAGs are isomorphic if and only if their defining graphs are isomorphic (see [21]). However, this is not the case for BBGs. For instance, the BBG defined on a tree with n vertices is always the free group of rank $n-1$. Nevertheless, some features

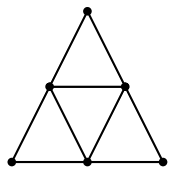

Figure 1. The trefoil graph.

of BBGs can still be seen directly from the defining graphs. For example, it was proved in [3] that BB_Γ is finitely generated if and only if Γ is connected, and BB_Γ is finitely presented if and only if the flag complex Δ_Γ associated with Γ is simply connected. When a BBG is finitely generated, an explicit presentation was found by Dicks and Leary [18]. More properties that have been discussed from a graphical perspective include various cohomological invariants in [19, 28, 31, 32], Dehn functions in [11], and graph of groups decompositions in [1, 12, 17].

In this paper, we add to this list a solution to the RAAG recognition problem for BBGs whose associated flag complexes are 2-dimensional (equivalently, the largest complete subgraphs of the defining graphs are triangles). Unless otherwise stated, we will always assume Γ is connected. We note that it is natural to make two additional assumptions. The first one is that Γ is biconnected, that is, it has no cut vertices (otherwise, one can split BB_Γ as the free product of the BBGs on the biconnected components of Γ; see Corollary 4.22). The second assumption is that the associated flag complex Δ_Γ is simply connected (otherwise, the group BB_Γ is not even finitely presented). Our solution to the RAAG recognition problem in dimension 2 is in terms of the presence or absence of two particular types of subgraphs. A *tree 2-spanner* T of the graph Γ is a spanning tree such that for any two vertices x and y, we have $d_T(x, y) \le 2d_\Gamma(x, y)$. A *crowned triangle* of the associated flag complex Δ_Γ is a triangle whose edges are not on the boundary of Δ_Γ (see Section 5 for the precise definition). For instance, the central triangle of the trefoil graph in Figure 1 is a crowned triangle.

Theorem A. *Let Γ be a biconnected graph such that Δ_Γ is 2-dimensional and simply connected. Then the following statements are equivalent.*

(1) Γ *admits a tree 2-spanner.*

(2) Δ_Γ *does not contain crowned triangles.*

(3) BB_Γ *is a RAAG.*

(4) BB_Γ *is an Artin group.*

Our proof of Theorem A relies on two conditions that are independent, in the sense that they work separately and regardless of the dimension of Δ_Γ. The first one is a sufficient condition for a BBG to be a RAAG that is based on the existence of a tree 2-spanner (see Section 1.1). The second one is a sufficient condition for any finitely generated group not to be a RAAG that is based on certain properties of the *Bieri–Neumann–Strebel invariant* (BNS-invariant) and may be of independent interest (see Section 1.3). We prove that these two conditions are equivalent when the flag complex Δ_Γ is 2-dimensional (see Section 1.4).

This allows one to recover the fact that the group G from above (i.e., the BBG defined on the trefoil graph from Figure 1) is not a RAAG. This was already known by the results of [16] or [31]. While the results in these two papers apply to groups that are more general

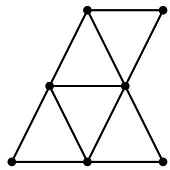

Figure 2. The extended trefoil graph. The BBG defined by it has this presentation: $\langle a, b, c, d, e, f \mid [a, b], [b, c], [c, d], [b^{-1}c, e], [e, f]\rangle$.

than the group G, they do not address the case of a very minor modification of that example, such as the BBG defined on the graph in Figure 2. This BBG shares all the properties with the group G described above. But again, it is not a RAAG by Theorem A since the defining graph contains a crowned triangle (see Example 5.14).

The features of the BNS-invariant that we use to show that a BBG is not a RAAG turn out to imply that the BBG cannot even be a more general Artin group. This relies on the theory of *resonance varieties* developed by Papadima and Suciu in [30, 31]. Roughly speaking, we show that for the BBGs under consideration in this paper, the resonance varieties are the same as the complements of the BNS-invariants (see Section 4.6).

1.1. The condition to be a RAAG: Tree 2-spanners

As observed in [31, Corollary 2.3], when BB_Γ is finitely presented, any spanning tree T of Γ provides a finite presentation whose relators are commutators. If T is a tree 2-spanner, then this presentation can actually be simplified to a standard RAAG presentation, and in particular, the group BB_Γ is a RAAG. We can even identify the defining graph for this RAAG in terms of the *dual graph* T^* of T, that is, the graph whose vertices are edges of T, and two vertices are adjacent if and only if the corresponding edges of T are contained in the same triangle of Γ. Note that the following result does not have any assumption on the dimension of Δ_Γ.

Theorem B. *If Γ admits a tree 2-spanner T, then BB_Γ is a RAAG. More precisely, the Dicks–Leary presentation can be simplified to the standard RAAG presentation with generating set $E(T)$. Moreover, we have $\mathrm{BB}_\Gamma \cong A_{T^*}$.*

Here are two applications. The first one is that if Γ admits a tree 2-spanner, then Δ_Γ is contractible (see Corollary 3.9). The second application is that for any graph Λ, the BBG defined on the cone over Λ is isomorphic to A_Λ, regardless of the structure of Λ (see Corollary 3.10). This means that the class of BBGs contains the class of RAAGs. That is, every RAAG arises as the BBG defined on some graph.

As we have mentioned before, two RAAGs are isomorphic if and only if their defining graphs are isomorphic, but this is not true for BBGs. However, when two graphs admit tree 2-spanners, the associated BBGs and the dual graphs completely determine each other.

Corollary 1. *Let Γ and Λ be two graphs admitting tree 2-spanners T_Γ and T_Λ, respectively. Then $\mathrm{BB}_\Gamma \cong \mathrm{BB}_\Lambda$ if and only if $T_\Gamma^* \cong T_\Lambda^*$.*

This provides new examples of non-isomorphic graphs defining isomorphic BBGs (see Example 3.7). On the other hand, when Γ does not admit a tree 2-spanner, the presentation for BB_Γ associated with any spanning tree is never a RAAG presentation. However, there might be a RAAG presentation not induced by a spanning tree. In order to obstruct this possibility, we need to look for invariants that do not depend on the choice of a generating set. We will consider the BNS-invariant $\Sigma^1(\mathrm{BB}_\Gamma)$ of BB_Γ.

1.2. The BNS-invariants of BBGs from the defining graphs

The BNS-invariant $\Sigma^1(G)$ of a finitely generated group G is a certain open subset of the character sphere $S(G)$, that is, the unit sphere in the space of group homomorphisms $\mathrm{Hom}(G, \mathbb{R})$. This invariant was introduced in [5] as a tool to study finiteness properties of normal subgroups of G with abelian quotients, such as kernels of characters. In general, the BNS-invariants are hard to compute.

The BNS-invariants of RAAGs have been characterized in terms of the defining graphs by Meier and VanWyk in [29]. The BNS-invariants of BBGs are less understood. In [32, Theorem 15.8], Papadima and Suciu gave a cohomological upper bound for the BNS-invariants of BBGs. Recently, Kochloukova and Mendonça have shown in [27, Corollary 1.3] how to reconstruct the BNS-invariant of a BBG from that of the ambient RAAG. However, an explicit description of the BNS-invariant of a BBG from its defining graph is still in need (recall that the correspondence between BBGs and graphs is not as explicit as in the case of RAAGs).

Since the vertices of Γ are generators for A_Γ, a convenient way to describe characters of A_Γ is via vertex-labelings. Inspired by this, in the present paper, we encode characters of BB_Γ as edge-labelings. This relies on the fact that the edges of Γ form a generating set for BB_Γ, under our standing assumption that Γ is connected (see [18] and Section 4.1.3). We obtain the following graphical criterion for a character of a BBG to belong to the BNS-invariant. The condition appearing in the following statement involves the *dead edge subgraph* $\mathcal{DE}(\chi)$ of a character χ of BB_Γ, which is the graph consisting of edges on which χ vanishes. This is reminiscent of the living subgraph criterion for RAAGs in [29]. However, it turns out that the case of BBGs is better understood in terms of the dead edge subgraph (see Example 4.16). An analogous dead subgraph criterion for RAAGs was considered in [1].

Theorem C (Graphical criterion for the BNS-invariant of a BBG). *Let Γ be a biconnected graph with Δ_Γ simply connected. Let $\chi \in \mathrm{Hom}(\mathrm{BB}_\Gamma, \mathbb{R})$ be a non-zero character. Then $[\chi] \in \Sigma^1(\mathrm{BB}_\Gamma)$ if and only if $\mathcal{DE}(\chi)$ does not contain a full subgraph that separates Γ.*

Theorem C allows one to work explicitly in terms of graphs with labeled edges. In particular, we show in Corollary 4.21 that the following are equivalent: The graph Γ is biconnected, the BNS-invariant $\Sigma^1(\mathrm{BB}_\Gamma)$ is non-empty, and BB_Γ *algebraically fibers* (i.e., it admits a homomorphism to $\mathbb{Z}$ with finitely generated kernel).

In the same spirit, we obtain the following graphical description for (the complement of) the BNS-invariants of BBGs. Here, a *missing subsphere* is a subsphere of the character sphere that is in the complement of the BNS-invariant (see Section 4.1 for details).

Theorem D (Graphical description of the BNS-invariant of a BBG). *Let Γ be a biconnected graph with Δ_Γ simply connected. Then $\Sigma^1(\mathrm{BB}_\Gamma)^c$ is a union of missing subspheres corresponding to full separating subgraphs. More precisely,*

(1) $\Sigma^1(\mathrm{BB}_\Gamma)^c = \bigcup_\Lambda S_\Lambda$, *where Λ ranges over the minimal full separating subgraphs of Γ.*

(2) *There is a bijection between maximal missing subspheres in $\Sigma^1(\mathrm{BB}_\Gamma)^c$ and minimal full separating subgraphs of Γ.*

In particular, as observed in [27, Corollary 1.4], the set $\Sigma^1(\mathrm{BB}_\Gamma)^c$ carries a natural structure of a rationally defined spherical polyhedron. A set of defining equations can be computed directly by looking at the minimal full separating subgraphs of Γ. This is analogous to the case of RAAGs (see [29]).

As a corollary of our description, we can identify the complement of the BNS-invariant with the first resonance variety (see Proposition 4.43). This improves the inclusion from [32, Theorem 15.8] to an equality. Once again, this is analogous to the case of RAAGs (see [30, Theorem 5.5]). It should be noted that there are groups for which the inclusion is strict (see [36]).

1.3. The condition not to be a RAAG: Redundant triangles

The BNS-invariant of a RAAG or BBG is the complement of a certain arrangement of subspheres of the character sphere. (Equivalently, one could consider the arrangement of linear subspaces given by the linear span of these subspheres.) The structural properties of this arrangement do not depend on any particular presentation of the group, so this arrangement turns out to be a useful invariant. In Section 4.4, inspired by the work of [16, 24], we consider the question of whether the maximal members in this arrangement are “in general position,” that is, whether they satisfy the inclusion-exclusion principle.

In [24], Koban and Piggott proved that the maximal members in the arrangement for a RAAG satisfy the inclusion-exclusion principle. Day and Wade in [16] developed a homology theory to detect when an arrangement does not satisfy the inclusion-exclusion principle. These results can be used together with our description of the BNS-invariants of BBGs to see that many BBGs are not RAAGs. However, some BBGs elude Day–Wade’s homology theory, such as the BBG defined on the graph in Figure 2. This motivated us to find an additional criterion to certify that a group G is not a RAAG. A more general result in Proposition 4.36 roughly says that if there are three maximal subspheres of $\Sigma^1(G)^c$ that are not “in general position,” then G is not a RAAG.

We are able to apply Proposition 4.36 to a wide class of BBGs. This is based on the notion of a *redundant triangle*. Loosely speaking, a redundant triangle is a triangle in Γ such that the links of its vertices are separating subgraphs of Γ that do not overlap too much (see Section 4.5 for the precise definition). The presence of such a triangle provides a triple of missing subspheres (in the sense of our graphical description; see Theorem D) that does not satisfy the inclusion-exclusion principle.

Theorem E. *Let Γ be a biconnected graph such that Δ_Γ is simply connected. If Γ has a redundant triangle, then* BB_Γ *is not a RAAG.*

We emphasize that Theorem E works without any assumptions on the dimension of Δ_Γ. On the other hand, the obstruction is 2-dimensional, in the sense that it involves a triangle, regardless of the dimension of Δ_Γ (see Example 5.17).

1.4. The 2-dimensional case: Proof of Theorem A

The two conditions described in Sections 1.1 and 1.3 are complementary when Γ is biconnected and Δ_Γ is 2-dimensional and simply connected. This follows from some structural properties enjoyed by 2-dimensional flag complexes.

In Proposition 5.11, we establish that Γ admits a tree 2-spanner if and only if Δ_Γ does not contain crowned triangles. The "if" direction relies on a decomposition of Δ_Γ into certain elementary pieces, namely, the cones over certain 1-dimensional flag complexes. It then follows from Theorem B that BB_Γ is a RAAG.

On the other hand, we show in Lemma 5.13 that every crowned triangle is redundant in dimension 2. It then follows from Theorem E that BB_Γ is not a RAAG. The theory of resonance varieties (see Section 4.6) allows us to conclude that BB_Γ cannot be a more general Artin group either.

Figure 3 illustrates the various implications. The only implication we do not prove directly is that if BB_Γ is a RAAG, then Γ has a tree 2-spanner. This implication follows from the other ones, and in particular, it means that one can write down the RAAG presentation for BB_Γ associated with the tree 2-spanner. This fact is a priori not obvious but quite satisfying.

For the sake of completeness, we note that Theorem A fails for higher-dimensional flag complexes (see Remark 5.15).

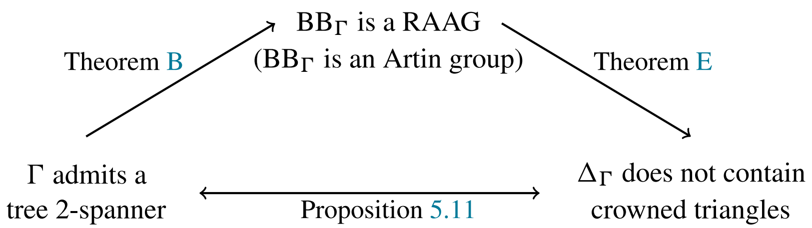

Figure 3. The implications in the proof of Theorem A.

1.5. Structure of the paper

The rest of the paper is organized as follows. In Section 2, we fix some terminology and give some background on BBGs. In Section 3, we study tree 2-spanners and use them to provide a sufficient condition for a BBG to be a RAAG (Theorem B). We also give many examples. In Section 4, we present a graphical criterion (Theorem C) and a graphical description (Theorem D) for the BNS-invariants of BBGs. We use this to provide a sufficient condition for a BBG not to be a RAAG (Theorem E). This is based on a study of the inclusion-exclusion principle for the arrangements that define the complement of the BNS-invariants. We discuss the relation with resonance varieties in Section 4.6. In Section 5, we provide a solution to the RAAG recognition problem for BBGs defined on 2-dimensional flag complexes (Theorem A). In the end, we include some observations about the higher-dimensional case.

2. Preliminaries

2.1. Notation and terminology

In this paper, unless otherwise stated, a *graph* Γ is a finite 1-dimensional simplicial complex, not necessarily connected. We denote by $V(\Gamma)$ the set of its *vertices* and by $E(\Gamma)$ the set of its *edges*. We do not fix any orientation on Γ, but we often need to work with *oriented edges*. If e is an oriented edge, then we denote its initial vertex and terminal vertex by ιe and τe, respectively; we denote by $\overline{e}$ the same edge with opposite orientation. We always identify edges of Γ with the unit interval and equip Γ with the induced length metric. A *subgraph* of Γ is a simplicial subcomplex, possibly not connected, possibly not full.

A *path*, a *cycle*, and a *complete graph* on n vertices are denoted by P_n, C_n, and K_n, respectively. (Note that by definition, there is no repetition of edges in a path or cycle.) A *clique* of Γ is a complete subgraph. A *tree* is a simply connected graph. A *spanning tree* of a graph Γ is a subgraph $T \subseteq \Gamma$ such that T is a tree and $V(T) = V(\Gamma)$.

The *link* of a vertex $v \in V(\Gamma)$, denoted by $\mathrm{lk}(v, \Gamma)$, is the full subgraph induced by the vertices that are adjacent to v. The *star* of v in Γ, denoted by $\mathrm{st}(v, \Gamma)$, is the full subgraph on $\mathrm{lk}(v, \Gamma) \cup \{v\}$. More generally, let Λ be a subgraph of Γ. The *link* of Λ is the full subgraph $\mathrm{lk}(\Lambda, \Gamma)$ induced by vertices at distance 1 from Λ. The *star* of Λ in Γ, denoted by $\mathrm{st}(\Lambda, \Gamma)$, is the full subgraph on $\mathrm{lk}(\Lambda, \Gamma) \cup V(\Lambda)$.

The *join* of two graphs Γ_1 and Γ_2, denoted by $\Gamma_1 * \Gamma_2$, is the full graph on $V(\Gamma_1) \cup V(\Gamma_2)$ together with an edge joining each vertex in $V(\Gamma_1)$ to each vertex in $V(\Gamma_2)$. A vertex in a graph that is adjacent to every other vertex is called a *cone vertex*. A graph that has cone vertex is called a *cone graph*. In other words, a cone graph Γ can be written as a join $\{v\} * \Gamma'$. In this case, we also say that Γ is a *cone over* Γ'. The *complement of* Λ *in* Γ is the full subgraph $\Gamma \setminus \Lambda$ spanned by $V(\Gamma) \setminus V(\Lambda)$. We say that Λ is *separating* if $\Gamma \setminus \Lambda$ is disconnected. A *cut vertex* of Γ is a vertex that is separating as a subgraph. A *cut edge* of Γ is an edge that is separating as a subgraph. A graph is *biconnected* if it has no cut vertices. If a graph is not biconnected, its *biconnected components* are the maximal biconnected subgraphs.

Given a graph Γ, the *flag complex* Δ_Γ on Γ is the simplicial complex obtained by gluing a k-simplex to Γ for every collection of $k+1$ pairwise adjacent vertices of Γ (for $k \geq 2$). The *dimension* of Δ_Γ is denoted by $\dim \Delta_\Gamma$ and defined to be the maximal dimension of a simplex in Δ_Γ. (If Δ_Γ 1-dimensional, then it coincides with Γ, and the following terminology agrees with the one introduced before.) If Z is a subcomplex of Δ_Γ, the *link* of Z in Δ_Γ, denoted by $\mathrm{lk}(Z, \Delta_\Gamma)$, is defined as the full subcomplex of Δ_Γ induced by the vertices at a distance one from Z. Similarly, the *star* of Z in Δ_Γ, denoted by $\mathrm{st}(Z, \Delta_\Gamma)$, is defined as the full subcomplex induced by $\mathrm{lk}(Z, \Delta_\Gamma) \cup Z$.

2.2. The Dicks–Leary presentation

Let Γ be a graph, and let A_Γ be the associated RAAG. Let $\chi\colon A_\Gamma \to \mathbb{Z}$ be the homomorphism sending all the generators to 1. The *Bestvina–Brady group* (BBG) on Γ, denoted by BB_Γ, is defined to be the kernel of χ. When Γ is connected, the group BB_Γ is finitely generated (see [3]) and has the following (infinite) presentation, called the *Dicks–Leary presentation*.

Theorem 2.1 ([18, Theorem 1]). *Let Γ be a graph. If Γ is connected, then BB_Γ is generated by the set of oriented edges of Γ, and the relators are words of the form $e_1^n \cdots e_l^n$ for each oriented cycle $(e_1, \dots, e_l)$, where $n, l \in \mathbb{Z}$, $n \neq 0$, and $l \geq 2$. Moreover, the group BB_Γ embeds in A_Γ via $e \mapsto \tau e(\iota e)^{-1}$ for each oriented edge e.*

For some interesting classes of graphs, the Dicks–Leary presentation can be considerably simplified. For instance, when the flag complex Δ_Γ on Γ is simply connected, the group BB_Γ admits the following finite presentation.

Corollary 2.2 ([18, Corollary 3]). *When the flag complex Δ_Γ on Γ is simply connected, the group BB_Γ admits the following finite presentation: The generating set is the set of the oriented edges of Γ, and the relators are $e\bar{e} = 1$ for every oriented edge e, and $e_i e_j e_k = 1$ and $e_k e_j e_i = 1$ whenever (e_i, e_j, e_k) form an oriented triangle (see Figure 4).*

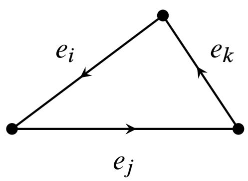

Figure 4. Oriented triangle.

Remark 2.3. In the notations of Corollary 2.2, it follows that e_i, e_j, and e_k generate a $\mathbb{Z}^2$ subgroup.

Example 2.4. If Γ is a tree on n vertices, then BB_Γ is a free group of rank $n-1$. If $\Gamma = K_n$ is a complete graph on n vertices, then $\mathrm{BB}_\Gamma = \mathbb{Z}^{n-1}$.

Moreover, as observed by Papadima and Suciu, the edge set of a spanning tree is already enough to generate the whole group.

Corollary 2.5 ([31, Corollary 2.3]). *Let T be a spanning tree of Γ. When the flag complex Δ_Γ on Γ is simply connected, the group BB_Γ admits a finite presentation in which the generators are the edges of T, and every defining relator is a commutator between words in the generators.*

Remark 2.6 (Oriented vs. unoriented edges). The presentation from Corollary 2.2 is very symmetric but clearly redundant because each (unoriented) edge appears twice. The orientation is just an accessory tool, and one can obtain a shorter presentation by choosing an arbitrary orientation for each edge e, dropping the relator $e\bar{e}$, and allowing inverses in the relators whenever needed. For instance, this is what happens in Corollary 2.5. Strictly speaking, each choice of orientation for the edges results in a slightly different presentation. However, switching the orientation of an edge simply amounts to replacing a generator with its inverse. Therefore, in the following sections, we will naively regard the generators in Corollary 2.5 as being given by unoriented edges of T, and we will impose a specific orientation only when needed in a technical argument.

3. BBGs that are RAAGs

When Γ is a tree or complete graph, the group BB_Γ is a free group or abelian group, respectively. Hence, it is a RAAG (see Example 2.4). In this section, we identify a wider class of graphs whose associated BBGs are RAAGs.

3.1. Tree 2-spanners

Let Γ be a connected graph. Recall from the introduction that a tree 2-spanner of Γ is a spanning tree T of Γ such that for all $x, y \in V(T)$, we have $d_T(x, y) \leq 2d_\Gamma(x, y)$. If Γ is a tree, then Γ is a tree 2-spanner of itself. Here, we are interested in more general graphs which admit tree 2-spanners. We start by proving some useful properties of tree 2-spanners.

Lemma 3.1. *Let T be a tree 2-spanner of Γ, and let $e \in E(\Gamma)$. Then either $e \in E(T)$ or there is a unique triangle (e, f, g) such that $f, g \in E(T)$.*

Proof. Write $e = (x, y)$, so $d_T(x, y) \leq 2d_\Gamma(x, y) = 2$. If e is not an edge of T, then $d_T(x, y) = 2$. So, there must be some $z \in V(T)$ adjacent to both x and y in Γ such that the edges $f = (x, z)$ and $g = (y, z)$ are in T. Obviously, the edges e, f, and g form a triangle.

To see that such a triangle is unique, let (e, f', g') be another triangle such that $f', g' \in E(T)$. Then (f, g, f', g') is a cycle in the spanning tree T, which leads to a contradiction. ∎

Lemma 3.2. *Let T be a tree 2-spanner of Γ. Then in every triangle of Γ, either no edge is in T or two edges are in T.*

Proof. Let (e, f, g) be a triangle in Γ, and assume by contradiction that $e \in E(T)$ but $f, g \notin E(T)$. Then by Lemma 3.1, the edges f and g are contained in uniquely determined triangles (f, f_1, f_2) and (g, g_1, g_2), respectively, with $f_1, f_2, g_1, g_2 \in E(T)$. Then (e, f_1, f_2, g_1, g_2) is a loop in T, which is absurd since T is a tree. ∎

Lemma 3.3. *Let T be a tree 2-spanner of Γ, and let (e, f, g) be a triangle in Γ with no edges from T. Then there are edges $e', f', g' \in E(T)$ that together with e, f, g form a K_4 in Γ.*

Proof. By Lemma 3.1, there are uniquely determined triangles (e, e_1, e_2), (f, f_1, f_2), and (g, g_1, g_2) such that $e_1, e_2, f_1, f_2, g_1, g_2 \in E(T)$. Let v_e be a common vertex shared by e_1 and e_2 and similarly define v_f and v_g. If at least two vertices among v_e, v_f, and v_g are distinct, then concatenating the edges $e_1, e_2, f_1, f_2, g_1, g_2$ gives a non-trivial loop in T, which is absurd. Thus, we have $v_e = v_f = v_g$. Therefore, there is a K_4 induced by the vertex v_e and the triangle (e, f, g). ∎

We establish the following result about the global structure of Δ_Γ. (We will prove in Corollary 3.9 that if Γ is a tree 2-spanner, then Δ_Γ is even contractible.)

Lemma 3.4. *If Γ has a tree 2-spanner, then Δ_Γ is simply connected.*

Proof. It is enough to check that every cycle of Γ bounds a disk in Δ_Γ. Let T be a tree 2-spanner of Γ, and let $C = (e_1, e_2, \ldots, e_n)$ be a cycle of Γ. If $n = 3$, then by construction C bounds a triangle in Δ_Γ. So, we may assume $n \geq 4$. If C contains a pair of vertices connected by an edge not in C, then C can be split into the concatenation of two shorter cycles. So, we assume that C contains no such a pair of vertices, that is, a chordless cycle. In particular, all edges of C are distinct.

For each $e_i \in E(C)$, either $e_i \in E(T)$ or $e_i \notin E(T)$. In the second case, by Lemma 3.1, there are two edges e_i^- and e_i^+ in $E(T)$ such that (e_i, e_i^-, e_i^+) form a triangle. We denote by w_i the common vertex of e_i^- and e_i^+. Note that $w_i \notin V(C)$ and $e_i^-, e_i^+ \notin E(C)$ because C is assumed to be chordless and of length $n \geq 4$. Let L be the loop obtained by the following surgery on C (see Figure 5, left): For each edge e_i, if $e_i \in E(T)$, then keep it;

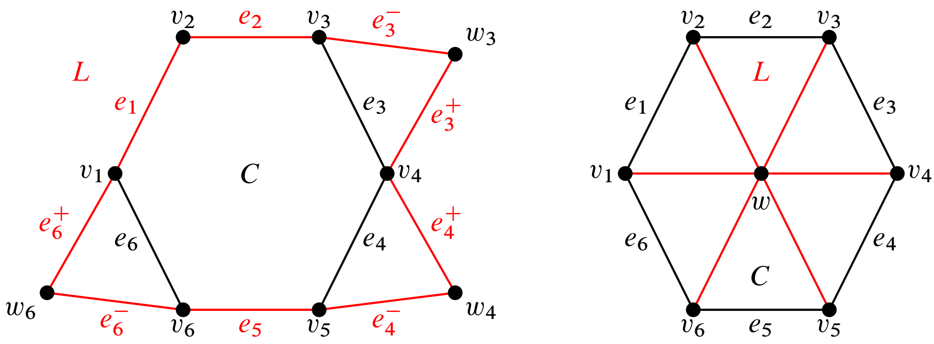

Figure 5. The construction of the loop L from the cycle C (left), and its contraction to a cone vertex (right).

otherwise, replace it with the concatenation of the two edges e_i^- and e_i^+. Then L is a loop made of edges of T. Since T is a tree, the loop L is contractible. Thus, if we start from a vertex of L and travel along the edges of L back to the starting vertex, then we must travel along each edge an even number of times in opposite directions.

Since $e_i^-, e_i^+ \notin E(C)$, each edge of C appears at most once in L. So, if some edge of C appears in L, then L is not contractible. This proves that $E(C) \cap E(T) = \varnothing$, and therefore, we have $L = (e_1^-, e_1^+, e_2^-, e_2^+, \dots, e_n^-, e_n^+)$. Once again, since the edges of L must appear an even number of times, the loop L contains a repeated edge. That is, we have $e_i^+ = e_{i+1}^-$ and $w_i = w_{i+1}$ for some i. Deleting the vertex v_{i+1} (and the repeated edge), we obtain a shorter cycle in T, made of edges from L. Iterating the process, we see that $w_1, \dots, w_n$ are all actually the same vertex, say w (see Figure 5, right). Notice that every vertex of C is adjacent to w, so C is entirely contained in $\mathrm{st}(w, \Gamma)$. Therefore, the cycle C bounds a triangulated disk in Δ_Γ as desired. ∎

In the next statement, we show that if Γ has a tree 2-spanner, then BB_Γ is a RAAG. Even more, the tree 2-spanner itself provides a RAAG presentation for BB_Γ. Let T be a tree 2-spanner for Γ. Recall from the introduction that the dual graph T^* of T is the graph whose vertices are edges of T, and two vertices are adjacent if and only if the corresponding edges of T are contained in the same triangle of Γ. Roughly speaking, the dual graph encodes the way in which T sits inside Γ.

Theorem B. *If Γ admits a tree 2-spanner T, then BB_Γ is a RAAG. More precisely, the Dicks–Leary presentation can be simplified to the standard RAAG presentation with generating set $E(T)$. Moreover, we have $\mathrm{BB}_\Gamma \cong A_{T^*}$.*

Proof. Let T be a tree 2-spanner of Γ. By Lemma 3.4, the flag complex Δ_Γ is simply connected. By Corollary 2.2, the Dicks–Leary presentation for BB_Γ is finite. The generators are the oriented edges of Γ, and the relators correspond to the oriented triangles in Γ. By Corollary 2.5, the presentation can be further simplified by discarding all edges not in T to obtain a presentation that only involves commutators between words in the generators. We explicitly note that to achieve this, one also needs to choose an arbitrary orientation for each edge of T (compare Remark 2.6). To ensure that the resulting presentation is a standard RAAG presentation, we need to check that it is enough to use relators that are commutators of edges of T (as opposed to commutators of more general words). In order to do this, we check what happens to the Dicks–Leary presentation from Corollary 2.2 when we remove a generator corresponding to an edge that is not in T. The relators involving such an edge correspond to the triangles of Γ that contain it. One of them is the special triangle from Lemma 3.1, and there might be other ones corresponding to other triangles.

Let $e \in E(\Gamma) \setminus E(T)$. By Lemma 3.1, we know that there is a unique triangle (e, f, g) with $f, g \in E(T)$. Then $(e^{\varepsilon_1}, f^{\varepsilon_2}, g^{\varepsilon_3})$ is an oriented triangle (in the sense of Figure 4) for some suitable $\varepsilon_j = \pm 1$, where the negative exponent stands for a reversal in the orientation. When we drop e from the generating set, the relations $e^{\varepsilon_1} f^{\varepsilon_2} g^{\varepsilon_3} = 1 = g^{\varepsilon_3} f^{\varepsilon_2} e^{\varepsilon_1}$ can be replaced by $f^{\varepsilon_2} g^{\varepsilon_3} = e^{-\varepsilon_1} = g^{\varepsilon_3} f^{\varepsilon_2}$, hence by the commutator $[f^{\varepsilon_2}, g^{\varepsilon_3}]$ (compare with Remark 2.3). But such a commutator can always be replaced by $[f, g]$. This is completely insensitive to the chosen orientation. This shows that the relators of the presentation from Corollary 2.2, which arise from the triangles provided by Lemma 3.1, are turned into commutators between generators in the presentation from Corollary 2.5.

We need to check what happens to the other type of relators. We now show that they follow from the former type of relators and hence can be dropped. As before, let $e \in E(\Gamma) \setminus E(T)$, and let (e, f, g) be the triangle from Lemma 3.1 having $f, g \in E(T)$. Let (e, f', g') be another triangle containing e. Since $e \notin E(T)$, by the uniqueness in Lemma 3.2 we have $e, f', g' \notin E(T)$. Therefore, by Lemma 3.3, there are $e'', f'', g'' \in E(T)$ that form a K_4 together with e, f', and g' (see the left picture of Figure 6). Up to relabeling, say that e'' is the edge of this K_4 that is disjoint from e. Then (e, f'', g'') is a triangle containing e with $f'', g'' \in E(T)$. Again, since the triangle (e, f, g) is unique, we have $\{f'', g''\} = \{f, g\}$. In particular, the triangles (e, f, g) and (e, f', g') are part of a common K_4 (see the right picture of Figure 6). The edges of this K_4 that are in T are precisely e'', f, and g, and any two of them commute by Remark 2.3. So, the relator

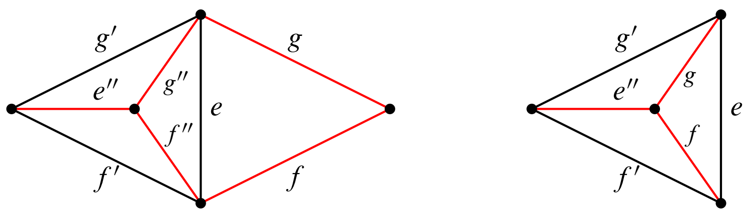

Figure 6. The graph on the left shows the triangle (e, f, g) and a K_4 consisting of the edges e, f', g', e'', f'', and g''. The red edges are in $E(T)$. The graph on the right illustrates the uniqueness of the triangle (e, f, g).

$ef'g'$ follows from the fact that e, f', and g' can be rewritten in terms of f, g, and e''. In particular, this relator can be dropped.

Therefore, the Dicks–Leary presentation for BB_Γ can be simplified to a presentation in which the generating set is $E(T)$, and the relators are commutators $[e_i, e_j]$, where e_i and e_j are in $E(T)$ and are contained in the same triangle of Δ_Γ. In particular, we have $\mathrm{BB}_\Gamma \cong A_{T^*}$. ∎

Remark 3.5. It is natural to ask which graphs admit a tree 2-spanner. The problem of determining whether a graph admits a tree 2-spanner is NP-complete (see [2]). However, if a graph contains a tree 2-spanner, then it can be found in linear time (see [9, Theorem 4.5]).

As a consequence, we have the following criterion to detect whether two BBGs are isomorphic in terms of the defining graphs in the special case where they admit tree 2-spanners.

Corollary 1. *Let Γ and Λ be two graphs admitting tree 2-spanners T_Γ and T_Λ, respectively. Then $\mathrm{BB}_\Gamma \cong \mathrm{BB}_\Lambda$ if and only if $T^*_\Gamma \cong T^*_\Lambda$.*

Proof. The result follows from Theorem B and the fact that two RAAGs are isomorphic if and only if their defining graphs are isomorphic (see Droms [21]). ∎

Remark 3.6. In general, non-isomorphic graphs can define isomorphic BBGs. For example, any two trees with n vertices define the same BBG (the free group of rank $n-1$). Notice that every tree is a tree 2-spanner of itself with a totally disconnected dual graph. Even when Γ admits a tree 2-spanner with a connected dual graph, the group BB_Γ does not determine Γ (see Example 3.7).

Example 3.7. Denote the graphs in Figure 7 by Γ and Λ. Let T_Γ and T_Λ be the tree 2-spanners of Γ and Λ, respectively, given by the red edges in the pictures. One can see that $\Gamma \ncong \Lambda$ as well as $T_\Gamma \ncong T_\Lambda$. However, the dual graphs T^*_Γ and T^*_Λ are isomorphic to the path on five vertices P_5. Thus, by Theorem B and Corollary 1, we have $\mathrm{BB}_\Gamma \cong A_{P_5} \cong \mathrm{BB}_\Lambda$.

The following graph-theoretic statements might be of independent interest. The first one says that any two tree 2-spanners for a graph Γ sit in the same way inside Γ (even though they do not have to be isomorphic as trees; see Example 3.7). The second one strengthens the conclusion of Lemma 3.4.

Figure 7. Non-isomorphic biconnected graphs that give isomorphic BBGs.

Corollary 3.8. *If T_1 and T_2 are tree* 2*-spanners for* Γ*, then* $T_1^* \cong T_2^*$*.*

Proof. Take $\Gamma = \Lambda$ in Corollary 1. ■

Corollary 3.9. *If* Γ *admits a tree* 2*-spanner, then* Δ_Γ *is contractible.*

Proof. By Theorem B, the group BB_Γ is isomorphic to the RAAG A_Λ on some graph Λ. The Salvetti complex associated with Λ is a finite classifying space for A_Λ, so the group $\mathrm{BB}_\Gamma \cong A_\Lambda$ is of type F. It follows from [3] that Δ_Γ is simply connected and acyclic. By the Hurewicz theorem, the homotopy group $\pi_k(\Delta_\Gamma)$ is trivial for $k \geq 1$. By the Whitehead theorem, we can conclude that Δ_Γ is contractible. ■

3.2. Joins and 2-trees

In this section, we describe some ways of constructing new graphs out of old ones in such a way that the BBG defined on the resulting graph is a RAAG.

3.2.1. Joins. Recall from Section 2.1 the definition of the join of two graphs. It corresponds to a direct product operation on the associated RAAGs. The following corollary can also be found in [31, Example 2.5] and [11, Proposition 3.4].

Corollary 3.10. *Let* Λ *be a graph. If* $\Gamma = \{v\} * \Lambda$*, then* $\mathrm{BB}_\Gamma \cong A_\Lambda$*.*

Proof. Since v is a cone vertex, the edges that are incident to v form a tree 2-spanner T of Γ. By Theorem B, we know that BB_Γ is a RAAG, namely, $\mathrm{BB}_\Gamma \cong A_{T^*}$. The result follows from the observation that $T^* \cong \Lambda$. ■

For instance, if Γ does not contain a full subgraph isomorphic to C_4 or P_3, then Γ is a cone (see the first lemma in [22]), and the previous corollary applies. Actually, in this case every subgroup of A_Γ is known to be a RAAG by the main theorem in [22].

Remark 3.11. Corollary 3.10 implies that the class of BBGs contains the class of RAAGs, that is, every RAAG arises as the BBG defined on some graph.

Remark 3.12. Corollary 3.10 indicates that the fact that BB_Γ is not a RAAG is not obviously detected by subgraphs in general. Indeed, if Γ is a cone over Λ, then BB_Γ is always a RAAG, regardless of the fact that BB_Λ is a RAAG or not.

Corollary 3.13. *Let* Λ *be a graph and* Γ' *a cone graph. If* $\Gamma = \Gamma' * \Lambda$*, then* BB_Γ *is a RAAG.*

Proof. Since Γ' is a cone graph, so is Γ. Therefore, the group BB_Γ is a RAAG by Corollary 3.10. ■

Corollary 3.14. *If* A_Γ *has non-trivial center, then* BB_Γ *is a RAAG.*

Proof. By [33, the centralizer theorem], when A_Γ has non-trivial center, there is a complete subgraph $\Gamma' \subseteq \Gamma$ such that each vertex of Γ' is adjacent to every other vertex of Γ. That is, the graph Γ decomposes as $\Gamma = \Gamma' * \Lambda$, where $V(\Lambda) = V(\Gamma) \setminus V(\Gamma)'$. Since a complete graph is a cone graph, the result follows from Corollary 3.13. ∎

Remark 3.15. BBGs defined on arbitrary graph joins are not necessarily isomorphic to RAAGs. For example, the cycle of length four C_4 is the join of two pairs of non-adjacent vertices. The associated RAAG is $\mathbb{F}_2 \times \mathbb{F}_2$, and the associated BBG is not a RAAG because it is not even finitely presented (see [3]).

3.2.2. Two-trees. Roughly speaking, a 2-*tree* is a graph obtained by gluing triangles along edges in a tree-like fashion. More formally, the class of 2-trees is defined recursively as follows: The graph consisting of a single edge is a 2-tree, and then a graph Γ is a 2-tree if it contains a vertex v such that the neighborhood of v in Γ is an edge and the graph obtained by removing v from Γ is still a 2-tree. The trefoil graph from Figure 1 is an explicit example of a 2-tree. A general 2-tree may not be a triangulation of a 2-dimensional disk as it can have branchings; see Figure 8 for an example. It is not hard to see that the flag complex on a 2-tree is simply connected. So, the associated BBG is finitely presented and has only commutator relators.

Cai showed that a 2-tree contains no trefoil subgraphs if and only if it admits a tree 2-spanner (see [8, Proposition 3.2]). The next corollary follows from Cai's result and Theorem B. In Section 5, we will prove a more general result that especially implies the converse of the following statement.

Corollary 3.16. *Let Γ be a 2-tree. If Γ is trefoil-free, then BB_Γ is a RAAG.*

Example 3.17 (A bouquet of triangles). Let Γ be the 2-tree shown in Figure 8. Since Γ does not contain trefoil subgraphs, the group BB_Γ is a RAAG by Corollary 3.16. The reader can check that the red edges form a tree 2-spanner of Γ.

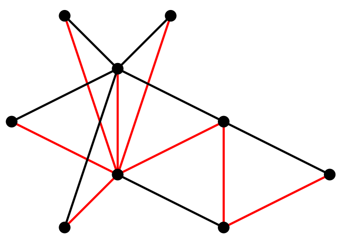

Figure 8. A 2-tree whose flag complex is not a triangulated disk.

4. BBGs that are not RAAGs

While in the previous section we have provided a condition on a graph Γ to ensure that BB_Γ is a RAAG, in this section, we want to obtain a condition on Γ to guarantee that BB_Γ is not a RAAG. The main technical tool consists of a description of the BNS-invariants of BBGs in terms of the defining graphs. Recall that we are always assuming that Γ is connected, and that this ensures that BB_Γ is finitely generated (see [3, 18]).

4.1. BNS-invariants of finitely generated groups

Let G be a finitely generated group. A *character* of G is a homomorphism $\chi\colon G \to \mathbb{R}$. Two characters χ_1 and χ_2 are equivalent, denoted by $\chi_1 \sim \chi_2$, whenever $\chi_1 = \lambda\chi_2$ for some positive real number λ. Denote by $[\chi]$ the equivalence class of χ. The set of equivalence classes of non-zero characters of G is called the *character sphere* of G:

$$S(G) = \{[\chi] \mid \chi \in \mathrm{Hom}(G, \mathbb{R}) \setminus \{0\}\}.$$

The character sphere naturally identifies with the unit sphere in $\mathrm{Hom}(G, \mathbb{R})$ (with respect to some background inner product), so by abuse of notation, we will often write $S(G) \subseteq \mathrm{Hom}(G, \mathbb{R})$. A character $\chi\colon G \to \mathbb{R}$ is called *integral*, *rational*, or *discrete* if its image is an infinite cyclic subgroup of $\mathbb{Z}$, $\mathbb{Q}$, or $\mathbb{R}$, respectively. In particular, an integral character is rational, a rational character is discrete, and the equivalence class of a discrete character always contains an integral representative.

Let $\mathcal{S}$ be a finite generating set for G, and let $\mathrm{Cay}(G, \mathcal{S})$ be the Cayley graph for G with respect to $\mathcal{S}$. Note that the elements of G are identified with the vertex set of the Cayley graph. For any character $\chi\colon G \to \mathbb{R}$, let $\mathrm{Cay}(G, \mathcal{S})_{\chi\geq 0}$ be the full subgraph of $\mathrm{Cay}(G, \mathcal{S})$ spanned by $\{g \in G \mid \chi(g) \geq 0\}$. Bieri, Neumann, and Strebel [5] introduced a geometric invariant of G, known as the *BNS-invariant* $\Sigma^1(G)$ of G, which is defined as the following subset of $S(G)$:

$$\Sigma^1(G) = \{[\chi] \in S(G) \mid \mathrm{Cay}(G, \mathcal{S})_{\chi\geq 0} \text{ is connected}\}.$$

They also proved that the BNS-invariant of G does not depend on the generating set $\mathcal{S}$.

The interest in $\Sigma^1(G)$ is due to the fact that it can detect finiteness properties of normal subgroups with abelian quotients, such as kernels of characters. For instance, the following statement can be taken as an alternative definition of what it means for a discrete character to belong to $\Sigma^1(G)$ (see [5, §4] or [35, Corollary A4.3]).

Theorem 4.1. *Let $\chi\colon G \to \mathbb{R}$ be a discrete character. Then* $\ker(\chi)$ *is finitely generated if and only if both* $[\chi]$ *and* $[-\chi]$ *are in* $\Sigma^1(G)$.

As a major motivating example, when G is the fundamental group of a compact 3-manifold M, the BNS-invariant $\Sigma^1(G)$ describes all the possible ways in which M fibers over the circle with fiber a compact surface (see [5, 34, 37]).

Remark 4.2. For each group G of interest in this paper, it admits an automorphism that acts as the antipodal map $\chi \mapsto -\chi$ on $\mathrm{Hom}(G, \mathbb{R})$. In this case, the BNS-invariant $\Sigma^1(G)$ is invariant under the antipodal map. Therefore, its rational points correspond exactly to discrete characters with finitely generated kernels.

Remark 4.3 (The complement and the missing subspheres). It is often the case that the BNS-invariant $\Sigma^1(G)$ is better described in terms of its complement in the character sphere $S(G)$. Moreover, for many groups of interest, the complement of the BNS-invariant is often a union of subspheres (see [4, 10, 26, 27, 29] for examples). In this paper, the *complement* of $\Sigma^1(G)$ is by definition $\Sigma^1(G)^c = S(G) \setminus \Sigma^1(G)$. A *great subsphere* is defined as a subsphere of $S(G)$ of the form $S_W = S(G) \cap W$, where W is a linear subspace of $\mathrm{Hom}(G, \mathbb{R})$ going through the origin. We say that a great subsphere S_W is a *missing subsphere* if $S_W \subseteq \Sigma^1(G)^c$. The subspace W is the linear span of S_W and is called a *missing subspace*.

4.1.1. The BNS-invariants of RAAGs. The BNS-invariants of RAAGs have a nice description given by Meier and VanWyk [29]. Let Γ be a graph and $\chi\colon A_\Gamma \to \mathbb{R}$ a character of A_Γ. Define the *living subgraph* $\mathcal{L}(\chi)$ of χ to be the full subgraph of Γ on the vertices v with $\chi(v) \neq 0$ and the *dead subgraph* $\mathcal{D}(\chi)$ of χ to be the full subgraph of Γ on the vertices v with $\chi(v) = 0$. Note that $\mathcal{L}(\chi)$ and $\mathcal{D}(\chi)$ are disjoint, and they do not necessarily cover Γ. A subgraph Γ' of Γ is *dominating* if every vertex of $\Gamma \setminus \Gamma'$ is adjacent to some vertex of Γ'.

Theorem 4.4 (A graphical criterion for $\Sigma^1(A_\Gamma)$ [29, Theorem 4.1]). *Let $\chi\colon A_\Gamma \to \mathbb{R}$ be a character. Then $[\chi] \in \Sigma^1(A_\Gamma)$ if and only if $\mathcal{L}(\chi)$ is connected and dominating.*

By Theorem 4.1, if χ is a discrete character, then $\mathcal{L}(\chi)$ detects whether $\ker(\chi)$ is finitely generated. Indeed, in a RAAG, the map sending a generator to its inverse is a group automorphism. Hence, the set $\Sigma^1(A_\Gamma)$ is invariant under the antipodal map $\chi \mapsto -\chi$ (see Remark 4.2).

We find it convenient to work with the following reformulation of the condition in Theorem 4.4. It previously appeared inside the proof of Corollary 3.4 of [1]. We include a proof for completeness. For the sake of clarity, in the following lemma, the graph Λ is a subgraph of $\mathcal{D}(\chi)$ that is separating as a subgraph of Γ, but it may not separate $\mathcal{D}(\chi)$ (see Figure 9 for an example).

Lemma 4.5. *Let $\chi\colon A_\Gamma \to \mathbb{R}$ be a non-zero character. Then the following statements are equivalent.*

1. *The living graph $\mathcal{L}(\chi)$ is either not connected or not dominating.*
2. *There exists a full subgraph $\Lambda \subseteq \Gamma$ such that Λ separates Γ and $\Lambda \subseteq \mathcal{D}(\chi)$.*

Proof. We begin by proving that (1) implies (2). If $\mathcal{L}(\chi)$ is not connected, then $\Lambda = \mathcal{D}(\chi)$ separates Γ. If $\mathcal{L}(\chi)$ is not dominating, then there is a vertex $v \in V(\Gamma)$ such that χ vanishes on v and on all the vertices adjacent to v. Since χ is non-zero, the vertex v is not a

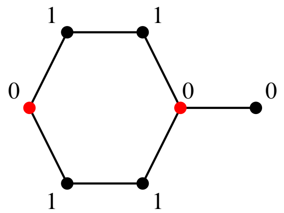

Figure 9. The subgraph Λ in Lemma 4.5 may not be a union of connected components of $\mathcal{D}(\chi)$. Here, the graph Λ is given by the two red vertices.

cone vertex. In particular, the graph $\Lambda = \mathrm{lk}(v, \Gamma)$ is a subgraph of $\mathcal{D}(\chi)$. Moreover, the subgraph Λ is a full separating subgraph of Γ, as desired.

To prove that (2) implies (1), assume that $\mathcal{L}(\chi)$ is connected and dominating. Let $\Lambda \subseteq \mathcal{D}(\chi)$ be a full subgraph of Γ, and let $u_1, u_2 \in V(\Gamma) \setminus V(\Lambda)$. We want to show that u_1 and u_2 can be connected in the complement of Λ. There are three cases. Firstly, if u_1 and u_2 are vertices of $\mathcal{L}(\chi)$, then they are connected by a path entirely in $\mathcal{L}(\chi)$. Secondly, if both u_1 and u_2 are vertices of $\mathcal{D}(\chi)$, then they are adjacent to some vertices in $\mathcal{L}(\chi)$, say v_1 and v_2, respectively. Then we can extend a path in $\mathcal{L}(\chi)$ between v_1 and v_2 to a path between u_1 and u_2 avoiding $V(\Lambda)$. Finally, suppose that u_1 is a vertex of $\mathcal{L}(\chi)$ and u_2 is a vertex of $\mathcal{D}(\chi)$, then u_2 is adjacent to a vertex v_2 of $\mathcal{L}(\chi)$. Again, we can extend a path in $\mathcal{L}(\chi)$ between u_1 and v_2 to a path between u_1 and u_2 avoiding $V(\Lambda)$. As a result, we have connected u_1 to u_2 with a path disjoint from Λ. This shows that Λ is not separating, which contradicts (2). ■

Notice that a subgraph Λ arising from Lemma 4.5 may not be connected and may not be equal to $\mathcal{D}(\chi)$. Also, it may not even be a union of connected components of $\mathcal{D}(\chi)$. This is in particular true when looking for a minimal such Λ (see Figure 9).

Remark 4.6. It follows from Theorem 4.4 that $\Sigma^1(A_\Gamma)^c$ is a rationally defined spherical polyhedron, given by a union of missing subspheres (see [29, Theorem 5.1]). Moreover, each (maximal) missing subsphere consists of characters that vanish on a (minimal) separating subgraph of Γ, thanks to Lemma 4.5 (see also [35, Proposition A4.14]). For example, the missing hyperspheres in $\Sigma^1(A_\Gamma)^c$ are in bijective correspondence with the cut vertices of Γ. We will further discuss the correspondence in Example 4.29.

4.1.2. The BNS-invariants of BBGs. As in the case of RAAGs, some elementary properties of the BNS-invariants of BBGs can be seen directly from the defining graph.

Example 4.7. The graph Γ is complete if and only if $\Sigma^1(\mathrm{BB}_\Gamma) = S(\mathrm{BB}_\Gamma)$. Indeed, in this case, the group BB_Γ is free abelian.

Example 4.8. At the opposite extreme, if Γ is connected and has a cut vertex, then $\Sigma^1(\mathrm{BB}_\Gamma) = \varnothing$ (see [32, Corollary 15.10]). Vice versa, if Γ has no cut vertices and BB_Γ is finitely presented, then we will prove in Corollary 4.20 that $\Sigma^1(\mathrm{BB}_\Gamma) \neq \varnothing$.

The following lemma shows that the BNS-invariant of a BBG is invariant under the antipodal map, as in the case of a RAAG.

Lemma 4.9. *For all* $\chi\colon \mathrm{BB}_\Gamma \to \mathbb{R}$*, if* $[\chi] \in \Sigma^1(\mathrm{BB}_\Gamma)$*, then* $[-\chi] \in \Sigma^1(\mathrm{BB}_\Gamma)$*.*

Proof. Choose an orientation for the edges of Γ, and let $f\colon \Gamma \to \Gamma$ be the map that reverses the orientation on each edge. Then f induces an automorphism $f_*\colon \mathrm{BB}_\Gamma \to \mathrm{BB}_\Gamma$ which sends every generator e to its inverse e^{-1}. Then the lemma follows from the fact that $-\chi = \chi \circ f_*$ (see Remark 4.2). ∎

Beyond these observations, not many explicit properties are known, and more refined tools are needed. We will use a recent result of Kochloukova and Mendonça that relates the BNS-invariant of a BBG to that of the ambient RAAG. The following statement is a particular case of [27, Corollary 1.3].

Proposition 4.10. *Let* Γ *be a connected graph, and let* $\chi : \mathrm{BB}_\Gamma \to \mathbb{R}$ *be a character. Then* $[\chi] \in \Sigma^1(\mathrm{BB}_\Gamma)$ *if and only if* $[\widehat{\chi}] \in \Sigma^1(A_\Gamma)$ *for every character* $\widehat{\chi} : A_\Gamma \to \mathbb{R}$ *that extends* χ*.*

Proposition 4.10 allows one to recover the previous observations as well as more properties of the BNS-invariants of BBGs, which are reminiscent of those of RAAGs. For instance, the complement of the BNS-invariant of a BBG is a rationally defined spherical polyhedron (see [27, Corollary 1.4] and compare with Remark 4.6).

4.1.3. Coordinates and labelings. Here, we want to describe a useful parametrization for $\mathrm{Hom}(A_\Gamma, \mathbb{R})$ and $\mathrm{Hom}(\mathrm{BB}_\Gamma, \mathbb{R})$ in terms of labeled graphs. This is based on the following elementary observation about a class of groups with a particular type of presentation that includes RAAGs and BBGs.

Lemma 4.11. *Let* G *be a group with a presentation* $G = \langle \mathcal{S} \mid \mathcal{R} \rangle$ *in which for each generator* s *and each relator* r*, the exponent sum of* s *in* r *is zero. Let* A *be an abelian group. Then there is a bijection between* $\mathrm{Hom}(G, A)$ *and* $\{f\colon \mathcal{S} \to A\}$*.*

Proof. Given a homomorphism $G \to A$, one obtains a function $\mathcal{S} \to A$ just by restriction. Conversely, let $f\colon \mathcal{S} \to A$ be any function and let $\widetilde{f}\colon \mathbb{F}(\mathcal{S}) \to A$ be the induced homomorphism on the free group on $\mathcal{S}$. Let $r \in \mathbb{F}(\mathcal{S})$ be a relator for G. Since the exponent sum of each generator in r is zero and A is abelian, we have that $\widetilde{f}(r)$ is trivial in A. Therefore, the homomorphism $\widetilde{f}\colon \mathbb{F}(\mathcal{S}) \to A$ descends to a well-defined homomorphism $G \to A$. ∎

A typical example of a relator in which the exponent sum of every generator is zero is a commutator. In particular, the standard presentation for a RAAG and the simplified Dicks–Leary presentation for a BBG in Corollary 2.5 are presentations of this type. We now show how Lemma 4.11 can be used to introduce nice coordinates on $\mathrm{Hom}(G, \mathbb{R})$ for these two classes of groups.

Let Γ be a connected graph, and let $V(\Gamma) = \{v_1, \dots, v_n\}$. By Lemma 4.11, a homomorphism $\chi \colon A_\Gamma \to \mathbb{R}$ is uniquely determined by its values on $V(\Gamma)$. Therefore, we get a natural identification

$$\mathrm{Hom}(A_\Gamma, \mathbb{R}) \to \mathbb{R}^{|V(\Gamma)|}, \quad \chi \mapsto (\chi(v_1), \dots, \chi(v_n)).$$

In other words, a character $\chi \colon A_\Gamma \to \mathbb{R}$ is the same as a labeling of $V(\Gamma)$ by real numbers. A natural basis for $\mathrm{Hom}(A_\Gamma, \mathbb{R})$ is given by the characters $\chi_1, \dots, \chi_n$, where $\chi_i(v_j) = \delta_{ij}$.

For BBGs, a similar description is available in terms of edge-labelings. Different from RAAGs, not every assignment of real numbers to the edges of Γ corresponds to a character. Indeed, the labels along an oriented cycle must sum to zero. So, assigning the labels on sufficiently many edges already determines the labels on the other ones. To find a clean description, we assume that the flag complex Δ_Γ is simply connected, and we fix a spanning tree T of Γ with $E(T) = \{e_1, \dots, e_m\}$. By Corollary 2.5, we know that the Dicks–Leary presentation can be simplified to have only $E(T)$ as a generating set and all relators are commutators. By Lemma 4.11, we get an identification

$$\mathrm{Hom}(\mathrm{BB}_\Gamma, \mathbb{R}) \to \mathbb{R}^{|E(T)|}, \quad \chi \mapsto (\chi(e_1), \dots, \chi(e_m)).$$

In other words, a character $\chi \colon \mathrm{BB}_\Gamma \to \mathbb{R}$ is encoded by a labeling of $E(T)$ by real numbers. To obtain a basis for $\mathrm{Hom}(\mathrm{BB}_\Gamma, \mathbb{R})$, one can take the characters $\chi_1, \dots, \chi_m$, where $\chi_i(e_j) = \delta_{ij}$, with respect to some arbitrary orientation of the edges of T (compare Remark 4.12).

Remark 4.12 (Computing a character on an edge). Note that there is a slight abuse of notation: Strictly speaking, in order to see an edge e as an element of BB_Γ, one needs to orient it. So, it only makes sense to evaluate a character χ on an oriented edge (see Remark 2.6). However, the value of $\chi(e)$ with respect to the two possible orientations just differs by a sign. Indeed, if we change the orientation of an edge and the sign of the corresponding label, then we obtain a different description of the same character of BB_Γ. In particular, it makes sense to say that a character vanishes or not on an edge, regardless of orientation. Moreover, given the symmetry of $\Sigma^1(\mathrm{BB}_\Gamma)$ under sign changes (see Remark 4.2 and Lemma 4.9), it is still quite meaningful and useful to think of a character as an edge-labeling for a spanning tree T.

As a result of the previous discussions, we obtain the following lemma, which we record for future reference.

Lemma 4.13. *Let Γ be a graph with Δ_Γ simply connected. Let T be a spanning tree of Γ. Fix an orientation for the edges of T. Then the following statements hold.*

(1) *A character $\chi \colon \mathrm{BB}_\Gamma \to \mathbb{R}$ is uniquely determined by its values on $E(T)$.*

(2) *Any assignment $E(T) \to \mathbb{R}$ uniquely extends to a character $\chi \colon \mathrm{BB}_\Gamma \to \mathbb{R}$.*

We conclude this section with a description of a natural restriction map. Recall from Theorem 2.1 that BB_Γ embeds in A_Γ via $e \mapsto \tau e(\iota e)^{-1}$ for each oriented edge e, where ιe and τe, respectively, denote the initial vertex and terminal vertex of e. We have an induced restriction map

$$r\colon \mathrm{Hom}(A_\Gamma, \mathbb{R}) \to \mathrm{Hom}(BB_\Gamma, \mathbb{R}), \quad \widehat{\chi} \mapsto r\widehat{\chi},$$

where $(r\widehat{\chi})(e) = \widehat{\chi}(\tau e) - \widehat{\chi}(\iota e)$. See Figure 12 for examples. The next result follows from the Dicks–Leary presentation in Theorem 2.1, so it holds without additional assumptions on Γ.

Lemma 4.14. *Let Γ be a connected graph. The restriction map $r\colon \mathrm{Hom}(A_\Gamma, \mathbb{R}) \to \mathrm{Hom}(BB_\Gamma, \mathbb{R})$ is a linear surjection, and its kernel consists of the characters defined by the constant functions $V(\Gamma) \to \mathbb{R}$.*

Proof. The map r is clearly linear. Let us prove that it is surjective. Let $\chi\colon BB_\Gamma \to \mathbb{R}$ be a character. Define a character $\widehat{\chi} \in \mathrm{Hom}(A_\Gamma, \mathbb{R})$ by prescribing its values on vertices as follows. Fix some $v_0 \in V(\Gamma)$ and choose $\widehat{\chi}(v_0) \in \mathbb{R}$ arbitrarily. Pick a vertex $v \in V(\Gamma)$ and choose an oriented path $p = e_1 \cdots e_k$ connecting v_0 to v. Here, we mean that p is oriented from v_0 to v and that all the edges along p are given the induced orientation. In particular, the edges of p can be seen as elements of BB_Γ. Define the value at v to be

$$\widehat{\chi}(v) = \widehat{\chi}(v_0) + \sum_{i=1}^{k} \chi(e_i). \tag{4.1}$$

We claim that $\widehat{\chi}\colon V(\Gamma) \to \mathbb{R}$ is well defined. Suppose that there is another oriented path $p' = e'_1 \cdots e'_h$ from v_0 to v. Then the loop $p(p')^{-1}$ gives a relator in the Dicks–Leary presentation for BB_Γ. Thus, we have

$$\chi(e_1 \cdots e_k (e'_1 \cdots e'_h)^{-1}) = 0.$$

In other words, we have

$$\sum_{i=1}^{k} \chi(e_i) = \sum_{j=1}^{h} \chi(e'_j).$$

Therefore, the value $\widehat{\chi}(v)$ does not depend on the choice of a path from v_0 to v, as desired. This provides that $\widehat{\chi} : A_\Gamma \to \mathbb{R}$ is a character. A direct computation shows that for each oriented edge e, we have $\chi(e) = \widehat{\chi}(\tau v) - \widehat{\chi}(\iota v)$. That is, the character $\widehat{\chi}$ is an extension of χ to A_Γ.

To describe the kernel of r, note that if $\widehat{\chi}$ is constant on $V(\Gamma)$, then for each oriented edge e we have $(r\widehat{\chi})(e) = \widehat{\chi}(\tau e) - \widehat{\chi}(\iota e) = 0$. Conversely, let $\widehat{\chi} \in \ker(r)$. It follows from (4.1) that $\widehat{\chi}(v) = \widehat{\chi}(w)$ for any choice of $v, w \in V(\Gamma)$. ∎

Note that the (non-zero) characters defined by the constant functions $V(\Gamma) \to \mathbb{R}$ all differ by a (multiplicative) constant. In particular, they all have the same kernel, which is precisely BB_Γ. The restriction map r has a natural linear section

$$s\colon \mathrm{Hom}(BB_\Gamma, \mathbb{R}) \to \mathrm{Hom}(A_\Gamma, \mathbb{R}), \quad \chi \mapsto s\chi,$$

defined as follows. Let $\widehat{\chi}$ be any extension of χ to A_Γ. Then define

$$s\chi = \widehat{\chi} - \frac{1}{|V(\Gamma)|} \sum_{v \in V(\Gamma)} \widehat{\chi}(v).$$

The image of s is a hyperplane W going through the origin that can be regarded as a copy of $\operatorname{Hom}(\mathrm{BB}_\Gamma, \mathbb{R})$ inside $\operatorname{Hom}(A_\Gamma, \mathbb{R})$.

Recall that if $V(\Gamma) = \{v_1, \dots, v_n\}$, then $\operatorname{Hom}(A_\Gamma, \mathbb{R})$ carries a canonical basis given by the characters $\chi_1, \dots, \chi_n$ such that $\chi_i(v_j) = \delta_{ij}$. Fix an inner product that makes this basis orthonormal. Then $\ker(r) = \operatorname{span}(1, \dots, 1)$, the hyperplane W is the orthogonal complement of $\ker(r)$, and the restriction map (or rather the composition $s \circ r$) is given by the orthogonal projection onto W.

It is natural to ask how this behaves with respect to the BNS-invariants, that is, whether r restricts to a map $\Sigma^1(A_\Gamma) \to \Sigma^1(\mathrm{BB}_\Gamma)$. In general, this is not the case. For instance, the set $\Sigma^1(\mathrm{BB}_\Gamma)$ could be empty even if $\Sigma^1(A_\Gamma)$ is not empty. On the other hand, the restriction map r maps each missing subspace of $\operatorname{Hom}(A_\Gamma, \mathbb{R})$ into one of the missing subspaces of $\operatorname{Hom}(\mathrm{BB}_\Gamma, \mathbb{R})$ (compare Remark 4.3 and Remark 4.28). Indeed, one way to reinterpret the content of Proposition 4.10 (a particular case of [27, Corollary 1.3]) is to say that if $\chi \in \operatorname{Hom}(\mathrm{BB}_\Gamma, \mathbb{R}) \cong W$, then $[\chi] \in \Sigma^1(\mathrm{BB}_\Gamma)$ if and only if the line parallel to $\ker(r)$ passing through χ avoids all the missing subspaces of $\operatorname{Hom}(A_\Gamma, \mathbb{R})$.

4.2. A graphical criterion for $\Sigma^1(\mathrm{BB}_\Gamma)$

In this subsection, we give a graphical criterion for the BNS-invariants of BBGs that is analogous to Theorem 4.4.

Let $\chi \in \operatorname{Hom}(\mathrm{BB}_\Gamma, \mathbb{R})$ be a non-zero character. An edge $e \in E(\Gamma)$ is called a *living edge* of χ if $\chi(e) \neq 0$; it is called a *dead edge* of χ if $\chi(e) = 0$. This is well defined, regardless of orientation, as explained in Remark 4.12. We define the *living edge subgraph*, denoted by $\mathcal{LE}(\chi)$, to be the subgraph of Γ consisting of the living edges of χ. The *dead edge subgraph* $\mathcal{DE}(\chi)$ is the subgraph of Γ consisting of the dead edges of χ. We will also say that χ *vanishes on* any subgraph of $\mathcal{DE}(\chi)$ because the associated labeling (in the sense of Section 4.1.3) is zero on each edge of $\mathcal{DE}(\chi)$. Notice that $\mathcal{LE}(\chi)$ and $\mathcal{DE}(\chi)$ cover Γ, but they are not disjoint; they intersect at vertices. Also, note that $\mathcal{LE}(\chi)$ and $\mathcal{DE}(\chi)$ are not always full subgraphs and do not have isolated vertices. Moreover, in general, the dead subgraph of an extension of χ is only a proper subgraph of $\mathcal{DE}(\chi)$. See Figure 10 for an example displaying all these behaviors.

The next lemma establishes a relation between the dead edge subgraph of a character of a BBG and the dead subgraphs of its extensions to the ambient RAAG. Note that the statement fails without the assumption that Λ is connected (see the example in Figure 10 once again).

Lemma 4.15. *Let Γ be a connected graph and let $\Lambda \subseteq \Gamma$ be a connected subgraph with at least one edge. Let $\chi \in \operatorname{Hom}(\mathrm{BB}_\Gamma, \mathbb{R})$ be a non-zero character. Then $\Lambda \subseteq \mathcal{DE}(\chi)$ if and only if there is an extension $\widehat{\chi} \in \operatorname{Hom}(A_\Gamma, \mathbb{R})$ of χ such that $\Lambda \subseteq \mathcal{D}(\widehat{\chi})$.*

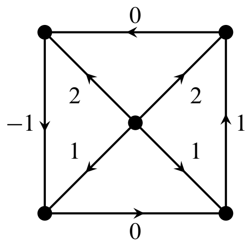

Figure 10. The dead edge subgraph $\mathcal{DE}(\chi)$ consists of a pair of opposite edges, and the living edge subgraph $\mathcal{LE}(\chi)$ consists of the remaining edges. Neither is a full subgraph. Moreover, the dead subgraph of any extension of χ is a proper subgraph of $\mathcal{DE}(\chi)$.

Proof. Suppose $\Lambda \subseteq \mathcal{DE}(\chi)$. By Lemma 4.14, there exists an extension of χ to A_Γ, unique up to additive constants. In particular, if we fix a vertex $v_0 \in V(\Lambda)$, then we can find an extension $\widehat{\chi} \in \mathrm{Hom}(A_\Gamma, \mathbb{R})$ such that $\widehat{\chi}(v_0) = 0$. Let $v \in V(\Lambda)$. Since Λ is connected, there is a path p connecting v_0 to v entirely contained in Λ. Since $\widehat{\chi}$ extends χ and χ vanishes on edges of p, a direct computation shows that $\widehat{\chi}(v) = 0$. Therefore, we have $\Lambda \subseteq \mathcal{D}(\widehat{\chi})$.

For the other direction, let $\widehat{\chi} \in \mathrm{Hom}(A_\Gamma, \mathbb{R})$ be an extension of χ such that $\Lambda \subseteq \mathcal{D}(\widehat{\chi})$. For every oriented edge $e = (v, w) \in E(\Lambda)$, we have $\chi(e) = \widehat{\chi}(vw^{-1}) = \widehat{\chi}(v) - \widehat{\chi}(w) = 0$. Thus, the edge e is in $\mathcal{DE}(\chi)$. Hence, we have $\Lambda \subseteq \mathcal{DE}(\chi)$. ∎

The main reason to focus on the dead edge subgraph instead of the living edge subgraph is that it is not clear how to transfer connectivity conditions from $\mathcal{L}(\widehat{\chi})$ to $\mathcal{LE}(\chi)$. On the other hand, the disconnecting features of $\mathcal{D}(\widehat{\chi})$ do transfer to $\mathcal{DE}(\chi)$. This is showcased by the following example.

Example 4.16. Let Γ be a cone over the path P_5 and consider a character $\widehat{\chi}\colon A_\Gamma \to \mathbb{Z}$ as shown in Figure 11. The living subgraph $\mathcal{L}(\widehat{\chi})$ is neither connected nor dominating. It follows from Theorem 4.4 that $[\widehat{\chi}] \notin \Sigma^1(A_\Gamma)$, and therefore, the restriction $\chi = \widehat{\chi}|_{\mathrm{BB}_\Gamma}\colon \mathrm{BB}_\Gamma \to \mathbb{Z}$ is not in $\Sigma^1(\mathrm{BB}_\Gamma)$ by Proposition 4.10. However, the living edge subgraph $\mathcal{LE}(\chi)$ is connected and dominating. On the other hand, note that $\mathcal{D}(\widehat{\chi})$ contains a full subgraph that separates Γ (compare with Lemma 4.5), and so does $\mathcal{DE}(\chi)$.

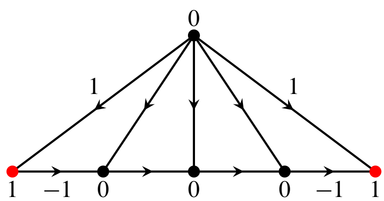

Figure 11. The living subgraph $\mathcal{L}(\widehat{\chi})$ consists of two red vertices. It is neither connected nor dominating. The living edge subgraph $\mathcal{LE}(\chi)$ (labeled by ± 1) is connected and dominating.

Our goal now is to show that the observations made in Example 4.16 about the dead edge subgraph hold in general. We will need the following general topological facts that we collect for the convenience of the reader. Here and in the following, "minimality" is always with respect to the inclusion of subgraphs. More precisely, a "minimal full separating subgraph" is a "full separating subgraph whose full subgraphs are not separating."

Lemma 4.17. *Let Γ be a biconnected graph with Δ_Γ simply connected.*

(1) *If $\Lambda \subseteq \Gamma$ is a connected full subgraph, then there is a bijection between the components of its complement and the components of its link.*

(2) *The link of every vertex is connected.*

(3) *If $\Lambda \subseteq \Gamma$ is a minimal full separating subgraph, then Λ is connected and not a single vertex.*

(4) *If $|V(\Gamma)| \geq 3$, then every edge is contained in at least one triangle. In particular, we have* $\dim \Delta_\Gamma \geq 2$.

Proof. (1). Let $A = \Delta_\Gamma \setminus \Delta_\Lambda$ (set-theoretic difference) and $B = \operatorname{st}(\Delta_\Lambda, \Delta_\Gamma)$. Then $\Delta_\Gamma = A \cup B$, and $\operatorname{lk}(\Delta_\Lambda, \Delta_\Gamma)$ deformation retracts to $A \cap B$. The Mayer–Vietoris sequence for reduced homology associated with this decomposition of Δ_Γ provides the following exact sequence:

$$\cdots \to H_1(\Delta_\Gamma) \to \widetilde{H}_0(\operatorname{lk}(\Delta_\Lambda, \Delta_\Gamma)) \to \widetilde{H}_0(A) \oplus \widetilde{H}_0(B) \to \widetilde{H}_0(\Delta_\Gamma) \to 0.$$

We have $H_1(\Delta_\Gamma) = 0 = \widetilde{H}_0(\Delta_\Gamma)$ since Δ_Γ is simply connected. Moreover, since Λ is connected, the subcomplex B is connected, and therefore, we obtain $\widetilde{H}_0(B) = 0$. This gives a bijection between $\widetilde{H}_0(\operatorname{lk}(\Delta_\Lambda, \Delta_\Gamma))$ and $\widetilde{H}_0(A)$, as desired.

(2). Take $\Lambda = v$ to be a single vertex. Since Γ is biconnected, the vertex v is not a cut vertex, so its complement is connected. Then the conclusion follows from (1).

(3). Let Λ be a minimal full separating subgraph of Γ. Then we can find two subcomplexes A and B of Δ_Γ such that $A \cup B = \Delta_\Gamma$ and $A \cap B = \Delta_\Lambda$. The Mayer–Vietoris sequence for reduced homology gives

$$\cdots \to H_1(\Delta_\Gamma) \to \widetilde{H}_0(\Delta_\Lambda) \to \widetilde{H}_0(A) \oplus \widetilde{H}_0(B) \to \widetilde{H}_0(\Delta_\Gamma) \to 0.$$

Arguing as in (1), we have $H_1(\Delta_\Gamma) = 0 = \widetilde{H}_0(\Delta_\Gamma)$ since Δ_Γ is simply connected. Therefore, we obtain $\widetilde{H}_0(\Delta_\Lambda) = \widetilde{H}_0(A) \oplus \widetilde{H}_0(B)$. Suppose by contradiction that Λ is disconnected. Then at least one of A or B is disconnected. Without loss of generality, say $A = A_1 \cup A'$, with A_1 a connected component of A and $A_1 \cap A' = \varnothing$. Let $B' = B \cup A'$ and let Λ' be the subgraph such that $\Delta_{\Lambda'} = A_1 \cap B'$. Then Λ' is a proper full subgraph of Λ which separates Γ, contradicting the minimality of Λ.

Finally, if by contradiction Λ were a single vertex, then it would be a cut vertex. But this is impossible because Γ is biconnected.

(4). Suppose by contradiction that there is an edge $e = (u, v)$ in Δ_Γ that is not contained in a triangle. Since Γ has at least three vertices, at least one endpoint of e, say v, is adjacent to at least another vertex different from u. Since e is not contained in a triangle, the vertex u is an isolated component of $\mathrm{lk}(v, \Gamma)$. Therefore, the subgraph $\mathrm{lk}(v, \Gamma)$ has at least two components, and hence, the vertex v is a cut vertex of Γ by (1). This contradicts the fact that Γ is biconnected. ∎

We now give a graphical criterion for a character to belong to the BNS-invariant of a BBG that is analogous to the living subgraph criterion in [29], or rather to the dead subgraph criterion Lemma 4.5 (see also [1, Corollary 3.4]).

Theorem C (Graphical criterion for the BNS-invariant of a BBG). *Let Γ be a biconnected graph with Δ_Γ simply connected. Let $\chi \in \mathrm{Hom}(\mathrm{BB}_\Gamma, \mathbb{R})$ be a non-zero character. Then $[\chi] \in \Sigma^1(\mathrm{BB}_\Gamma)$ if and only if $\mathcal{DE}(\chi)$ does not contain a full subgraph that separates Γ.*

Proof. Let $[\chi] \in \Sigma^1(\mathrm{BB}_\Gamma)$. Suppose by contradiction that there is a full subgraph $\Lambda \subseteq \mathcal{DE}(\chi)$ that separates Γ. Up to passing to a subgraph, we can assume that Λ is a minimal full separating subgraph. So, by (3) in Lemma 4.17, we can assume that Λ is connected. By Lemma 4.15, there is an extension $\hat{\chi} \in \mathrm{Hom}(A_\Gamma, \mathbb{R})$ of χ such that $\Lambda \subseteq \mathcal{D}(\hat{\chi})$. Since Λ separates Γ, by Lemma 4.5, we have $[\hat{\chi}] \notin \Sigma^1(A_\Gamma)$, and therefore $[\chi] \notin \Sigma^1(\mathrm{BB}_\Gamma)$ by Proposition 4.10. Hence, we reach a contradiction.

Conversely, assume $[\chi] \notin \Sigma^1(\mathrm{BB}_\Gamma)$. Then by Proposition 4.10, there is an extension $\hat{\chi} \in \mathrm{Hom}(A_\Gamma, \mathbb{R})$ of χ such that $[\hat{\chi}] \notin \Sigma^1(A_\Gamma)$. So, the living subgraph $\mathcal{L}(\hat{\chi})$ is either disconnected or not dominating. Equivalently, by Lemma 4.5, the dead subgraph $\mathcal{D}(\hat{\chi})$ contains a full subgraph Λ which separates Γ. Note that every edge of Λ is contained in $\mathcal{DE}(\chi)$ because $\Lambda \subseteq \mathcal{D}(\hat{\chi})$. A priori, the subgraph Λ could have some components consisting of isolated points. Once again, passing to a subgraph, we can assume that Λ is a minimal full separating subgraph. By (3) in Lemma 4.17, we can also assume that Λ is connected and not reduced to a single vertex. Therefore, we have $\Lambda \subseteq \mathcal{DE}(\chi)$. This completes the proof. ∎

We give two examples to illustrate that the hypotheses of Theorem C are optimal. Here, characters are represented by labels in the sense of Section 4.1.3.

Example 4.18 (Simple connectedness is needed). Consider the cycle of length four $\Gamma = C_4$ (see the left-hand side of Figure 12). Then Δ_Γ is not simply connected. Note that in this case, the group BB_Γ is finitely generated but not finitely presented (see [3]). Let $\hat{\chi} \in \mathrm{Hom}(A_\Gamma, \mathbb{R})$ be the character of A_Γ that sends two non-adjacent vertices to 0 and the other two vertices to 1. Let $\chi = \hat{\chi}|_{\mathrm{BB}_\Gamma} \in \mathrm{Hom}(\mathrm{BB}_\Gamma, \mathbb{R})$ be the restriction of $\hat{\chi}$ to BB_Γ. Then the dead edge subgraph $\mathcal{DE}(\chi)$ is empty. In particular, it does not contain any subgraph that separates Γ. However, the living subgraph $\mathcal{L}(\hat{\chi})$ consists of two opposite vertices, which is not connected. Thus, we have $[\hat{\chi}] \notin \Sigma^1(A_\Gamma)$. Hence, by Proposition 4.10, we obtain $[\chi] \notin \Sigma^1(\mathrm{BB}_\Gamma)$.

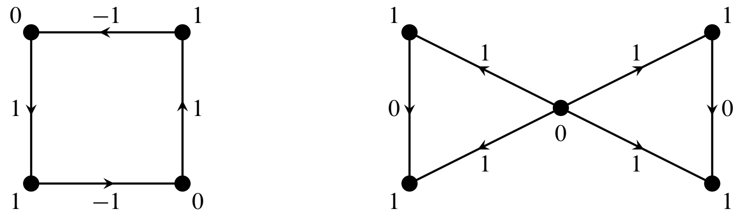

Figure 12. Theorem C does not hold on a graph with Δ_Γ not simply connected (left), nor on a graph with a cut vertex (right).

Example 4.19 (Biconnectedness is needed). Let Γ be the graph obtained by gluing two triangles at a vertex (see the right-hand side of Figure 12). Let $\hat{\chi} \in \mathrm{Hom}(A_\Gamma, \mathbb{R})$ be the character that sends the cut vertex to 0 and all the other vertices to 1. Let $\chi = \hat{\chi}|_{\mathrm{BB}_\Gamma} \in \mathrm{Hom}(\mathrm{BB}_\Gamma, \mathbb{R})$ be the restriction of $\hat{\chi}$ to BB_Γ. Then the dead edge subgraph $\mathcal{DE}(\chi)$ consists of the two edges that are not incident to the cut vertex. In particular, it does not contain any subgraph that separates Γ. However, the living subgraph $\mathcal{L}(\hat{\chi})$ is not connected (also notice that $\mathcal{L}(\hat{\chi}) = \mathcal{DE}(\chi)$). Thus, we have $[\hat{\chi}] \notin \Sigma^1(A_\Gamma)$. Hence, Proposition 4.10 implies $[\chi] \notin \Sigma^1(\mathrm{BB}_\Gamma)$.

As mentioned in Example 4.8, the graph Γ in Example 4.19 has a cut vertex, and hence, the BNS-invariant $\Sigma^1(\mathrm{BB}_\Gamma)$ is empty (see [32, Corollary 15.10]). As promised, we now show the following result.

Corollary 4.20. *Let Γ be a biconnected graph with Δ_Γ simply connected. Then, we have $\Sigma^1(\mathrm{BB}_\Gamma) \neq \varnothing$.*

Proof. Let T be a spanning tree of Γ. Assign an orientation to each edge of T and write $E(T) = \{e_1, \dots, e_m\}$. Let

$$\chi \colon E(T) \to \mathbb{R}, \quad \chi(e_k) = 10^k, \quad k = 1, \dots, m.$$

Then this defines a character thanks to Lemma 4.13. We claim that χ does not vanish on any edge of Γ. Indeed, let $e \in E(\Gamma)$. The claim is clear if $e \in E(T)$. Suppose $e \notin E(T)$, and let $(e_{i_1}, \dots, e_{i_p})$ be the path in T between the endpoints of e. Then $(e, e_{i_1}, \dots, e_{i_p})$ is a cycle in Γ, and hence, the element $e e_{i_1} \cdots e_{i_p}$ is a relator in BB_Γ by Theorem 2.1. Therefore, we have

$$0 = \chi(e) + \chi(e_{i_1}) + \dots + \chi(e_{i_p}) = \chi(e) \pm 10^{k_{i_1}} \pm \dots \pm 10^{k_{i_p}},$$

where the signs are determined by the orientations of the corresponding edges. The sum $\pm 10^{k_{i_1}} \pm \dots \pm 10^{k_{i_p}}$ is never zero since all the exponents are different. Thus, we have $\chi(e) \neq 0$. This proves the claim. It immediately follows that $\mathcal{DE}(\chi) = \varnothing$, and therefore, we have $[\chi] \in \Sigma^1(\mathrm{BB}_\Gamma)$ by Theorem C. ∎

As a summary, we have the following corollary. Most implications are well known. Our contribution is that (1) implies (2). Recall that a finitely generated group G *algebraically fibers* if there is a surjective homomorphism $G \to \mathbb{Z}$ whose kernel is finitely generated.

Corollary 4.21. *Let Γ be a connected graph such that Δ_Γ is simply connected. Then the following statements are equivalent.*

(1) *Γ is biconnected.*

(2) $\Sigma^1(\mathrm{BB}_\Gamma) \neq \varnothing$.

(3) *BB_Γ does not split as a free product.*

(4) *BB_Γ is 1-ended.*

(5) *BB_Γ algebraically fibers.*

Proof. The equivalence of (1) and (2) is given by Corollary 4.20 and [32, Corollary 15.10]. Given that BB_Γ is torsion-free, the equivalence of (3) and (4) is just Stallings theorem about the ends of groups (see [7, Theorem I.8.32]). The fact that (2) implies (3) is discussed in [35, Example 3 in A2.1a]. The fact that (3) implies (1) can be seen directly from the Dicks–Leary presentation from Theorem 2.1.

Finally, we show that (5) is equivalent to (2). It follows from Theorem 4.1 that BB_Γ algebraically fibers if and only if there exists a discrete character $\chi : \mathrm{BB}_\Gamma \to \mathbb{R}$ such that both $[\chi]$ and $[-\chi]$ are in $\Sigma^1(\mathrm{BB}_\Gamma)$. This is actually equivalent to just requiring that $[\chi] \in \Sigma^1(\mathrm{BB}_\Gamma)$, because $\Sigma^1(\mathrm{BB}_\Gamma)$ is symmetric (see Lemma 4.9). Note that the points of the character sphere $S(\mathrm{BB}_\Gamma)$ given by the equivalence classes of discrete characters are exactly the rational points. In particular, since $\Sigma^1(\mathrm{BB}_\Gamma)$ is an open subset of the character sphere $S(\mathrm{BB}_\Gamma)$ (see [5, Theorem A]), it is non-empty if and only if it contains the equivalence class of a discrete character. ∎

We record the following consequence for future reference. It will reduce our discussion about the RAAG recognition problem to the case of biconnected graphs.

Corollary 4.22. *Let Γ be a connected graph with Δ_Γ simply connected, and let $\Gamma_1, \ldots, \Gamma_n$ be its biconnected components. Then BB_Γ is a RAAG if and only if BB_{Γ_i} is a RAAG for all $i = 1, \ldots, n$.*

Proof. It is clear from the Dicks–Leary presentation that BB_Γ is the free product of the BB_{Γ_i}. Moreover, since Γ_i is biconnected, each BB_{Γ_i} is freely indecomposable (see Corollary 4.21). If all the BB_{Γ_i} are RAAGs, then BB_Γ is a RAAG because the free product of RAAGs is a RAAG. This proves one implication. For the converse implication, suppose that BB_Γ is a RAAG, say $\mathrm{BB}_\Gamma = A_\Lambda$ for some graph Λ. Let $\Lambda_1, \ldots, \Lambda_m$ be the connected components of Γ. Then $\mathrm{BB}_\Gamma = A_\Lambda$ can also be written as the free product of the RAAGs A_{Λ_j}, each of which is freely indecomposable. It follows that $m = n$, and for each i, there is some j such that $\mathrm{BB}_{\Gamma_i} \cong A_{\Lambda_j}$. ∎

Remark 4.23. We conclude this subsection by observing that when Γ is a chordal graph, the statement in Theorem C can also be obtained as follows. By [1, §3.2], the group BB_Γ splits as a finite graph of groups. More precisely, the vertex groups correspond to the BBGs on the maximal cliques of Γ, and the edge groups correspond to BBGs on the minimal separating subgraphs of Γ (that are also cliques because Γ is chordal). In particular, all these groups are finitely generated free abelian groups. Hence, one can apply the results from [10, §2].

4.3. A graphical description of $\Sigma^1(\mathrm{BB}_\Gamma)$

We now provide a graphical description of $\Sigma^1(\mathrm{BB}_\Gamma)$, that is, a way to compute the BNS-invariant of BB_Γ in terms of subgraphs of Γ.

Recall from Remark 4.6 that $\Sigma^1(A_\Gamma)^c$ is given by an arrangement of missing subspheres parametrized by the separating subgraphs of Γ. Thanks to [27, Corollary 1.4], we know that $\Sigma^1(\mathrm{BB}_\Gamma)^c$ is also an arrangement of missing subspheres. Moreover, the restriction map $r\colon \mathrm{Hom}(A_\Gamma, \mathbb{R}) \to \mathrm{Hom}(\mathrm{BB}_\Gamma, \mathbb{R})$ sends the missing subspheres of $\Sigma^1(A_\Gamma)^c$ to those of $\Sigma^1(\mathrm{BB}_\Gamma)^c$ (see the discussion after Lemma 4.14). So, it makes sense to look for a description of the missing subspheres of $\Sigma^1(\mathrm{BB}_\Gamma)^c$ in terms of subgraphs of Γ, analogous to the one available for $\Sigma^1(A_\Gamma)^c$.

However, recall from Example 3.7 that BB_Γ does not completely determine Γ, so it is a priori not clear that $\Sigma^1(\mathrm{BB}_\Gamma)^c$ should admit such a description. Moreover, the restriction map is not always well behaved with respect to the vanishing behavior of characters, in the sense that the dead edge subgraph of a character can be strictly larger than the dead subgraph of any of its extensions (see Figure 10). To address this, we need a way to construct characters with prescribed vanishing behavior.

For any subgraph Λ of Γ, we define the following linear subspace of $\mathrm{Hom}(\mathrm{BB}_\Gamma, \mathbb{R})$

$$W_\Lambda = \{\chi\colon \mathrm{BB}_\Gamma \to \mathbb{R} \mid \chi(e) = 0,\ \forall\, e \in E(\Lambda)\} = \{\chi\colon \mathrm{BB}_\Gamma \to \mathbb{R} \mid \Lambda \subseteq \mathcal{DE}(\chi)\}$$

and the great subsphere S_Λ given by the following intersection

$$S_\Lambda = W_\Lambda \cap S(\mathrm{BB}_\Gamma).$$

Note that if a character χ of BB_Γ vanishes on a spanning tree of Γ, then χ is trivial (see Lemma 4.13). In other words, if Λ is a spanning tree, then $W_\Lambda = 0$ and $S_\Lambda = \varnothing$. We look for a condition on Λ such that $W_\Lambda \neq 0$ and $S_\Lambda \neq \varnothing$. Notice that the following lemma applies as soon as $V(\Lambda) \neq V(\Gamma)$, and that if it applies to Λ, then it also applies to all of its subgraphs.

Lemma 4.24. *Let Γ be a graph with Δ_Γ simply connected, and let $\Lambda \subseteq \Gamma$ be a subgraph. Assume that there is an edge $e_0 \in E(\Gamma)$ with at least one endpoint not in $V(\Lambda)$. Then there exists a character $\chi\colon \mathrm{BB}_\Gamma \to \mathbb{R}$ such that $\chi(e_0) = 1$ and $\chi(e) = 0$ for all $e \in E(\Lambda)$. In particular, we have $[\chi] \in S_\Lambda$.*

Proof. Let T_Λ be a spanning forest of Λ, that is, a subgraph of Λ that is a disjoint union of trees and contains all vertices of Λ (note that we are not assuming that Λ is connected). Observe $e_0 \notin E(\Lambda)$ by assumption. Therefore, we can extend $T_\Lambda \cup \{e_0\}$ to a spanning tree T of Γ. Orient the edges of T arbitrarily and label the edges of T_Λ by 0 and all the remaining edges of T by 1. By Lemma 4.13, this defines a character $\chi\colon \mathrm{BB}_\Gamma \to \mathbb{R}$. By construction, we have $\chi(e_0) = 1$ and $\chi(e) = 0$ for $e \in E(T_\Lambda)$. Let $e \in E(\Lambda) \setminus E(T_\Lambda)$. Since T_Λ is a spanning forest of Λ, there is a unique path p in T_Λ from τe to ιe. Then ep is a cycle in Γ, and therefore, it is a relator in the Dicks–Leary presentation for BB_Γ. Since χ vanishes on p, it must also vanish on e, as desired. ∎

Remark 4.25. Notice that if two subgraphs Λ and Λ' have the same edge sets, then $W_\Lambda = W_{\Lambda'}$ because these subspaces only depend on the edge sets. In particular, we have $S_\Lambda = S_{\Lambda'}$. This is the reason why we use the strict inclusion $\subsetneq$ instead of the weak inclusion $\subseteq$ in the statement (2) of the following lemma.

Lemma 4.26. *Let Γ be a biconnected graph with Δ_Γ simply connected, and let Λ and Λ' be full separating subgraphs. Then we have the following statements.*

(1) *S_Λ is a missing subsphere, that is, we have $S_\Lambda \subseteq \Sigma^1(\mathrm{BB}_\Gamma)^c$.*

(2) *$\Lambda' \subsetneq \Lambda$ if and only if $S_\Lambda \subsetneq S_{\Lambda'}$.*

Proof. (1). If $[\chi] \in S_\Lambda$, then $\mathcal{DE}(\chi)$ contains Λ, which is a separating subgraph. Then the statement follows from Theorem C.

(2). The implication $\Lambda' \subsetneq \Lambda \Rightarrow S_\Lambda \subsetneq S_{\Lambda'}$ follows from the definitions. For the reverse implication $S_\Lambda \subsetneq S_{\Lambda'} \Rightarrow \Lambda' \subsetneq \Lambda$ we argue as follows. The inclusion $S_\Lambda \subsetneq S_{\Lambda'}$ implies that a character vanishing on Λ must also vanish on Λ'. We need to show that Λ' is a proper subgraph of Λ.

By contradiction, suppose that Λ' is not a subgraph of Λ. Notice that if $\Lambda' \setminus \Lambda$ consists of isolated vertices, then $S_\Lambda = S_{\Lambda'}$ (see Remark 4.25). So, we can assume that there is an edge $e_0 \in E(\Lambda') \setminus E(\Lambda)$. Since Λ is full, the edge e_0 cannot have both endpoints in Λ. By Lemma 4.24, there is a character $\chi\colon \mathrm{BB}_\Gamma \to \mathbb{R}$ with $\chi(e_0) = 1$ and $\chi(e) = 0$ for all $e \in E(\Lambda)$. This is a character that vanishes identically on Λ but not on Λ', which is absurd. ∎

Recall that if Λ is a separating subgraph, then $S_\Lambda \neq \varnothing$.

Theorem D (Graphical description of the BNS-invariant of a BBG). *Let Γ be a biconnected graph with Δ_Γ simply connected. Then $\Sigma^1(\mathrm{BB}_\Gamma)^c$ is a union of missing subspheres corresponding to full separating subgraphs. More precisely,*

(1) *$\Sigma^1(\mathrm{BB}_\Gamma)^c = \bigcup_\Lambda S_\Lambda$, where Λ ranges over the minimal full separating subgraphs of Γ.*

(2) *There is a bijection between maximal missing subspheres of $\Sigma^1(\mathrm{BB}_\Gamma)^c$ and minimal full separating subgraphs of Γ.*

Proof. (1). We start by proving that $\Sigma^1(\mathrm{BB}_\Gamma)^c = \bigcup_\Lambda S_\Lambda$, where Λ ranges over the full separating subgraphs of Γ. If Λ is a full separating subgraph, then we know that $S_\Lambda \subseteq \Sigma^1(\mathrm{BB}_\Gamma)^c$ by Lemma 4.26 (1). So one inclusion is clear. Vice versa, let $[\chi] \in \Sigma^1(\mathrm{BB}_\Gamma)^c$. Then by Theorem C we have that $\mathcal{DE}(\chi)$ contains a full separating subgraph Λ. In particular, the character χ vanishes on Λ, hence $[\chi] \in S_\Lambda$. This proves the other inclusion. To see that one can restrict to Λ ranging over minimal full separating subgraphs, just observe that the latter correspond to maximal missing subspheres by Lemma 4.26 (2). This completes the proof of (1).

(2). By (1), we know that $\Sigma^1(\mathrm{BB}_\Gamma)^c$ is a union of maximal missing subspheres. Notice that this is a finite union because Γ has only finitely many subgraphs. So, each maximal missing subsphere S is of the form $S = S_\Lambda$ for Λ a minimal full separating subgraph.

Vice versa, let Λ be a minimal full separating subgraph of Γ. We know from Lemma 4.26 (1) that S_Λ is a missing subsphere. We claim that S_Λ is a maximal missing subsphere in $\Sigma^1(\mathrm{BB}_\Gamma)^c$. Let S be a maximal missing subsphere in $\Sigma^1(\mathrm{BB}_\Gamma)^c$ such that $S_\Lambda \subseteq S$. By the previous paragraph, we know that $S = S_{\Lambda'}$ for some minimal full separating subgraph Λ'. If we had $S_\Lambda \subsetneq S = S_{\Lambda'}$, then by Lemma 4.26 (2) it would follow that $\Lambda' \subsetneq \Lambda$. But this would contradict the minimality of Λ. Thus, we have $S_\Lambda = S_{\Lambda'} = S$. Hence, the missing subsphere S_Λ is maximal. ■

The following example establishes a correspondence between the cut edges in Γ and the missing hyperspheres (the missing subspheres of codimension one) in $\Sigma^1(\mathrm{BB}_\Gamma)^c$. It should be compared with the case of RAAGs, where the correspondence is between the cut vertices of Γ and the missing hyperspheres in $\Sigma^1(A_\Gamma)^c$ (compare Remark 4.6 and Example 4.29).

Example 4.27 (Hyperspheres). Let Γ be a biconnected graph with Δ_Γ simply connected. Let e be a cut edge of Γ. Notice that e is a minimal separating subgraph since Γ is biconnected, and it is also clearly full. So by Theorem D we know that S_e is a maximal missing subsphere in $\Sigma^1(\mathrm{BB}_\Gamma)^c$. We want to show that the subspace $W_e = \mathrm{span}(S_e)$ is a hyperplane. To see this, let T be a spanning tree of Γ with $E(T) = \{e_1, \dots, e_m\}$, and let y_i be the coordinate dual to e_i in the sense of Section 4.1.3. This means that $y_i(\chi) = \chi(e_i)$ for all $\chi \in \mathrm{Hom}(\mathrm{BB}_\Gamma, \mathbb{R})$. Note that W_{e_i} is the hyperplane given by the equation $y_i = 0$. If $e \in E(T)$, then $e = e_i$ for some $i = 1, \dots, m$ and $W_e = W_{e_i}$ is a hyperplane. If $e \notin E(T)$, then there is a unique path $(e_{j_1}, \dots, e_{j_p})$ in T connecting the endpoints of e. Since $(e_{j_1}, \dots, e_{j_p}, e)$ is a cycle in Γ, the word $e_{j_1} \cdots e_{j_p} e$ is a relator in the Dicks–Leary presentation. So, we have $\chi(e_{j_1}) + \cdots + \chi(e_{j_p}) + \chi(e) = 0$. Therefore, we obtain $\chi(e) = 0$ if and only if $y_{j_1}(\chi) + \cdots + y_{j_p}(\chi) = 0$. This means that W_e is the hyperplane defined by the equation $y_{j_1} + \cdots + y_{j_p} = 0$.

Vice versa, let $S \subseteq \Sigma^1(\mathrm{BB}_\Gamma)^c$ be a hypersphere. We claim that $S = S_e$ for some cut edge e. To see this, let $[\chi] \in S$. By Theorem C, we know that $\mathcal{DE}(\chi)$ contains a full subgraph Λ that separates Γ. Since Γ is biconnected, the subgraph Λ must contain at least one edge. In particular, the character χ vanishes on $E(\Lambda)$, and therefore, we have

$[\chi] \in \bigcap_{e\in E(\Lambda)} S_e$. This proves $S \subseteq \bigcap_{e\in E(\Lambda)} S_e$. However, by the discussion above, we know that S_e is a hypersphere. Since S is also a hypersphere, the subgraph Λ must consist of a single edge e only. In particular, it is a cut edge.

Remark 4.28. The linear span of the arrangement of the missing subspheres of $\Sigma^1(G)^c$ gives rise to a subspace arrangement in $\mathrm{Hom}(G, \mathbb{R})$. The main difference between RAAGs and BBGs is that the arrangement for a RAAG is always "in general position," while the arrangement for a BBG is not. We will discuss the details in the next section.

4.4. The inclusion-exclusion principle

Given a group G, one can consider the collection of maximal missing subspheres. That is, the maximal great subspheres of the character sphere $S(G)$ that are in the complement of the BNS-invariant $\Sigma^1(G)$ (see Remark 4.3). Additionally, one can also consider the collection of maximal missing subspaces in $\mathrm{Hom}(G, \mathbb{R})$, that is, the linear spans of the maximal missing subspheres. This provides an arrangement of (great) subspheres in $S(G)$ and an arrangement of (linear) subspaces in $\mathrm{Hom}(G, \mathbb{R})$ that can be used as invariants for G. For instance, these arrangements completely determine the BNS-invariant when G is a RAAG or a BBG (see Remark 4.6 or Theorem D, respectively). Moreover, in the case of RAAGs, these arrangements satisfy a certain form of the inclusion-exclusion principle (see Section 4.4.1). This fact can be used to detect when a group G is not a RAAG. We take this point of view from the work of Koban and Piggott in [24] and Day and Wade in [16]. The former focuses on the subsphere arrangement, while the latter focuses on the subspace arrangement. In this section, we find it convenient to focus on the subspace arrangement.

Let V be a real vector space. (The reader should think $V = \mathrm{Hom}(G, \mathbb{R})$ for a group G.) For convenience, we fix some background inner product on V. All arguments in the following are combinatorial and do not depend on the choice of inner product. We say that a finite collection of linear subspaces $\{W_j\}_{j\in J}$ of V satisfies the *inclusion-exclusion principle* if the following equality holds:

$$\dim\left(\sum_{j=1}^{|J|} W_j\right) = \sum_{k=1}^{|J|}(-1)^{k+1}\left(\sum_{I\subset J, |I|=k} \dim\left(\bigcap_{j\in I} W_j\right)\right). \tag{4.2}$$

Notice that if an arrangement satisfies (4.2), then any linearly equivalent arrangement also satisfies (4.2). Here are two examples. The first is a RAAG, and the collection of maximal subspaces in the complement of its BNS-invariant satisfies the inclusion-exclusion principle. The second is a BBG, and the collection of maximal subspaces in the complement of its BNS-invariant does not satisfy the inclusion-exclusion principle. Note that this BBG is known to be not isomorphic to any RAAG by [31].

Example 4.29 (Trees). Let Γ be a tree on n vertices, and let $\{v_1, \ldots, v_m\}$ be the set of cut vertices of Γ. Then it follows that $\Sigma^1(A_\Gamma)$ is obtained from $S(A_\Gamma) = S^n$ by removing the

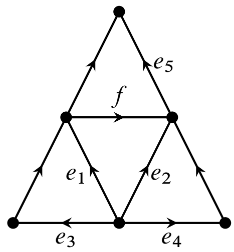

Figure 13. An oriented trefoil graph with a spanning tree.

hyperspheres S_i defined by $x_i = 0$ for $i = 1, \ldots, m$ (see Section 4.1.3). The associated missing subspaces satisfy the inclusion-exclusion principle (4.2).

Example 4.30 (The trefoil). Let Γ be the (oriented) trefoil graph with a choice of a spanning tree T whose edge set is $E(T) = \{e_1, e_2, e_3, e_4, e_5\}$ (see Figure 13). We consider the three cut edges e_1, e_2, and f. By Example 4.27, we have that S_{e_1}, S_{e_2}, and S_f are missing hyperspheres in $\Sigma^1(\mathrm{BB}_\Gamma)^c$. By Theorem D, we have $\Sigma^1(\mathrm{BB}_\Gamma)^c = S_{e_1} \cup S_{e_2} \cup S_f$. If $y_1, \ldots, y_5$ are the dual coordinates on $\mathrm{Hom}(\mathrm{BB}_\Gamma, \mathbb{R}) \cong \mathbb{R}^5$ with respect to T (in the sense of Section 4.1.3), then S_{e_1}, S_{e_2}, and S_f are given by $y_1 = 0$, $y_2 = 0$, and $y_1 - y_2 = 0$, respectively. To see the latter, first note that we have a relator $e_1 f = e_2 = f e_1$ in the Dicks–Leary presentation. Then for any character $\chi \in \mathrm{Hom}(\mathrm{BB}_\Gamma, \mathbb{R})$, we have $\chi(e_1) + \chi(f) = \chi(e_2)$. Thus, we obtain $\chi(f) = 0$ if and only if $\chi(e_1) = \chi(e_2)$. Therefore, the hypersphere S_f is defined by $y_1 = y_2$, that is, the equation $y_1 - y_2 = 0$. A direct computation shows that the associated missing subspaces do not satisfy the inclusion-exclusion principle (4.2).

It is natural to ask whether the phenomenon from Example 4.30 is actually a general obstruction for a BBG to be a RAAG. In [16], Day and Wade developed a homology theory $H_*(\mathcal{V})$ for a subspace arrangement $\mathcal{V}$ in a vector space that is designed to measure the failure of the inclusion-exclusion principle for $\mathcal{V}$. They proved that if G is a RAAG, then $H_k(\mathcal{V}_G) = 0$ for all $k > 0$, where $\mathcal{V}_G$ denotes the arrangement of maximal subspaces corresponding to the maximal missing spheres in $\Sigma^1(G)^c$ (see [16, Theorem B]).

Given our description of the BNS-invariant for BBGs from Section 4.1.3 and Theorem D, we can determine that certain BBGs are not RAAGs. For example, a direct computation shows that the group $G = \mathrm{BB}_\Gamma$ from Example 4.30 has $H_1(\mathcal{V}_G) \neq 0$. On the other hand, there are BBGs that cannot be distinguished from RAAGs in this way, as in the next example.

Example 4.31 (The extended trefoil). Let Γ be the trefoil graph with one extra triangle attached (see Figure 14). Imitating Example 4.30, we choose a spanning tree T whose edge set is $E(T) = \{e_1, e_2, e_3, e_4, e_5, e_6\}$. By Theorem D, we have $\Sigma^1(\mathrm{BB}_\Gamma)^c = S_{e_1} \cup S_{e_2} \cup S_f \cup S_{e_5}$. If $y_1, \ldots, y_6$ are the dual coordinates on $\mathrm{Hom}(\mathrm{BB}_\Gamma, \mathbb{R}) \cong \mathbb{R}^6$ with

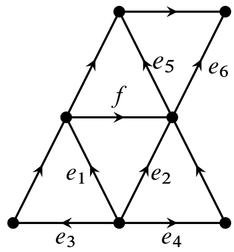

Figure 14. The extended trefoil: a new example of a BBG that is not a RAAG.

respect to T (in the sense of Section 4.1.3), then these missing hyperspheres are defined by the hyperplanes given by $y_1 = 0$, $y_2 = 0$, $y_1 - y_2 = 0$, and $y_5 = 0$, respectively. A direct computation shows that $H_k(\mathcal{V}_{\mathrm{BB}_\Gamma}) = 0$ for all $k \geq 0$, that is, these homology groups look like the homology groups for the arrangement associated with a RAAG. However, we will show that this BBG is not a RAAG in Example 5.14.

Our goal now is to obtain a criterion to detect when a BBG is not a RAAG that is still based on a certain failure of the inclusion-exclusion principle in the complement of the BNS-invariant. The obstruction always involves only a collection of three subspaces, regardless of the complexity of the graph. So, we find it convenient to introduce the following notation:

$$\begin{aligned}\mathrm{IEP}(W_1, W_2, W_3) = \dim W_1 + \dim W_2 + \dim W_3 \\ - \dim(W_1 \cap W_2) - \dim(W_1 \cap W_3) - \dim(W_2 \cap W_3) \\ + \dim(W_1 \cap W_2 \cap W_3). \qquad (4.3)\end{aligned}$$

4.4.1. RAAG behavior. The following lemma states that the arrangement defining the BNS-invariant of any RAAG satisfies the inclusion-exclusion principle. This is due to the fact that in this case, the missing subspaces are effectively described in terms of sets of vertices of Γ and the inclusion-exclusion principle holds for subsets of a given set. The argument follows the proof of Lemma 5.3 of [24]. We include a proof for completeness.

Lemma 4.32. *Let Γ be a connected graph. Let $\{W_j\}_{j \in J}$ be a collection of maximal missing subspaces in* $\mathrm{Hom}(A_\Gamma, \mathbb{R})$. *Then $\{W_j\}_{j \in J}$ satisfies* (4.2). *In particular, when $J = \{1, 2, 3\}$ we have* $\dim(W_1 + W_2 + W_3) = \mathrm{IEP}(W_1, W_2, W_3)$.

Proof. Recall that the subspace W_j corresponds to a minimal full separating subgraph Λ_j of Γ (see Remark 4.6). Moreover, the dimension of W_j is equal to the number of vertices in the complement $A_j = \Gamma \setminus \Lambda_j$ of Λ_j (those vertices provide a basis for W_j, in the sense of Section 4.1.3.) It follows that one can compute the dimensions of the various intersections, and obtain the following equalities

$$\dim\left(\sum_{j=1}^{|J|} W_j\right) = \left|\bigcup_{j=1}^{|J|} V(A_j)\right| = \sum_{k=1}^{|J|}(-1)^{k+1}\left(\sum_{\substack{I\subset J\\ |I|=k}}\left|\bigcap_{j\in I} V(A_j)\right|\right)$$
$$= \sum_{k=1}^{|J|}(-1)^{k+1}\left(\sum_{\substack{I\subset J\\ |I|=k}}\dim\left(\bigcap_{j\in I} W_j\right)\right).$$

This means precisely that $\{W_j\}_{j\in J}$ satisfies (4.2), as desired. ∎

4.4.2. Non-RAAG behavior. We now want to identify a condition that is not compatible with the property established in Section 4.4.1 for the arrangement associated with a RAAG. More precisely, we look for a sharp lower bound for the term $\mathrm{IEP}(W_1, W_2, W_3)$. The key condition is the one in Lemma 4.34. It is inspired by [16], and it could be interpreted in the setting of the homology theory introduced in that paper (see Remark 4.33). For the reader's convenience, we provide a self-contained exposition.

Let V be a real vector space of dimension n. Once again, the reader should think of the case $V = \mathrm{Hom}(G,\mathbb{R}) \cong \mathbb{R}^n$ for some group G with n generators. We fix some inner product, an orthonormal basis $\{e_1,\dots,e_n\}$, and the corresponding coordinates $\{y_1,\dots,y_n\}$, that is, $y_i(e_j) = \delta_{ij}$. Consider three subspaces of V given by the following systems of equations:

$$W_1 = \left\{y_1 = 0, \sum_{i=1}^{n}\lambda_{ij}^1 y_i = 0 \text{ for } j = 1,\dots,m_1\right\},$$
$$W_2 = \left\{y_2 = 0, \sum_{i=1}^{n}\lambda_{ij}^2 y_i = 0 \text{ for } j = 1,\dots,m_2\right\},$$
$$W_3 = \left\{y_1 - y_2 = 0, \sum_{i=1}^{n}\lambda_{ij}^3 y_i = 0 \text{ for } j = 1,\dots,m_3\right\}, \tag{4.4}$$

where for $k \in \{1,2,3\}$, we have $\lambda_{ij}^k \in \mathbb{R}$, and m_k is a non-negative integer (possibly zero, in which case it is understood that the subspace is just given by the first equation, as in Example 4.30). Without loss of generality, we assume that each set of equations is minimal. That is, we have $\dim W_k = n - (m_k + 1)$.

We now proceed to compute the term $\mathrm{IEP}(W_1, W_2, W_3)$ defined in (4.3). In the naive system of equations that defines the intersection $W_1 \cap W_2 \cap W_3$ (i.e., the one obtained by putting all the equations together), there is an obvious linear relation among the equations $y_1 = 0$, $y_2 = 0$, and $y_1 - y_2 = 0$. This can cause the dimension of $W_1 \cap W_2 \cap W_3$ to be higher than expected. From this perspective, one of the three equations is redundant. We find it convenient to work with the orthogonal complements. For $i, j \in \{1,2,3\}$, $i \neq j$, consider the following natural maps:

$$\begin{aligned} I_{ij} &: W_i^{\perp} \cap W_j^{\perp} \longrightarrow W_i^{\perp} \oplus W_j^{\perp}, \quad I_{ij}(u) = (u, -u), \\ F_{ij} &: W_i^{\perp} \oplus W_j^{\perp} \longrightarrow W_i^{\perp} + W_j^{\perp}, \quad F_{ij}(\zeta_i, \zeta_j) = \zeta_i + \zeta_j, \\ J_{ij} &: W_i^{\perp} \oplus W_j^{\perp} \longrightarrow W_1^{\perp} \oplus W_2^{\perp} \oplus W_3^{\perp}, \end{aligned} \tag{4.5}$$

where the last one is the natural inclusion (e.g., $J_{12}(\zeta_1, \zeta_2) = (\zeta_1, \zeta_2, 0)$). These maps fit in the diagram in Figure 15, where the first row is exact.

Let $K_{ij} \subseteq W_1^{\perp} \oplus W_2^{\perp} \oplus W_3^{\perp}$ be the image of $J_{ij} \circ I_{ij}$. By construction, we have $K_{ij} \cong (W_i + W_j)^{\perp} = W_i^{\perp} \cap W_j^{\perp}$. Finally, consider the vector $\xi = (-e_1, e_2, e_1 - e_2) \in W_1^{\perp} \oplus W_2^{\perp} \oplus W_3^{\perp}$. We say that a triple of subspaces $\{W_1, W_2, W_3\}$ as above is a *redundant triple of subspaces* if $\xi \notin K_{12} + K_{23} + K_{13}$.

Remark 4.33. Although we will not need it, we observe that the condition $\xi \notin K_{12} + K_{23} + K_{13}$ described above can be interpreted in the sense of the subspace arrangement homology introduced in [16] as follows. Consider the arrangement $\mathcal{W}^{\perp}$ given by the orthogonal complements $\{W_1^{\perp}, W_2^{\perp}, W_3^{\perp}\}$. Then $\{W_1, W_2, W_3\}$ is a redundant triple of subspaces precisely when ξ defines a non-trivial class in $H_1(\mathcal{W}^{\perp})$.

Lemma 4.34. *In the above notation, if $\{W_1, W_2, W_3\}$ is a redundant triple of subspaces, then it does not satisfy the inclusion-exclusion principle. More precisely,*

$$\dim(W_1 + W_2 + W_3) + 1 \le \mathrm{IEP}(W_1, W_2, W_3).$$

Proof. We will compute all the terms that appear in $\mathrm{IEP}(W_1, W_2, W_3)$ (see (4.3)). The exactness of the first row of the diagram in Figure 15 yields that

$$\dim(W_i^{\perp} + W_j^{\perp}) = \dim(W_i^{\perp} \oplus W_j^{\perp}) - \dim(W_i^{\perp} \cap W_j^{\perp}) = 2 + m_i + m_j - \dim K_{ij}.$$

It follows that

$$\begin{aligned} \dim(W_i \cap W_j) &= n - \dim((W_i \cap W_j)^{\perp}) \\ &= n - \dim(W_i^{\perp} + W_j^{\perp}) \\ &= n - (2 + m_i + m_j) + \dim K_{ij}. \end{aligned} \tag{4.6}$$

$$\begin{array}{ccccccccc} 0 & \longrightarrow & (W_i + W_j)^{\perp} = W_i^{\perp} \cap W_j^{\perp} & \xrightarrow{I_{ij}} & W_i^{\perp} \oplus W_j^{\perp} & \xrightarrow{F_{ij}} & W_i^{\perp} + W_j^{\perp} & \longrightarrow & 0 \\ & & & & \Big\downarrow J_{ij} & & & & \\ & & & & W_1^{\perp} \oplus W_2^{\perp} \oplus W_3^{\perp} & & & & \end{array}$$

Figure 15. The diagram for Lemma 4.34.

We deal with the triple intersection similarly. Consider the map

$$F : W_1^{\perp} \oplus W_2^{\perp} \oplus W_3^{\perp} \longrightarrow W_1^{\perp} + W_2^{\perp} + W_3^{\perp}, \quad F(\zeta_1, \zeta_2, \zeta_3) = \zeta_1 + \zeta_2 + \zeta_3.$$

We have $\dim(W_1^{\perp} \oplus W_2^{\perp} \oplus W_3^{\perp}) = 3 + m_1 + m_2 + m_3$. Since F is surjective, its codomain has dimension $3 + m_1 + m_2 + m_3 - \dim(\ker F)$. It follows that

$$\begin{aligned} \dim(W_1 \cap W_2 \cap W_3) &= n - \dim((W_1 \cap W_2 \cap W_3)^{\perp}) \\ &= n - \dim(W_1^{\perp} + W_2^{\perp} + W_3^{\perp}) \\ &= n - (3 + m_1 + m_2 + m_3) + \dim(\ker F). \end{aligned} \tag{4.7}$$

Using $\dim W_k = n - (m_k + 1)$, (4.6) and (4.7), we obtain

$$\mathrm{IEP}(W_1, W_2, W_3) = n + \dim(\ker F) - \dim K_{12} - \dim K_{13} - \dim K_{23}. \tag{4.8}$$

We now claim that $\dim(\ker F) \geq 1 + \dim K_{12} + \dim K_{13} + \dim K_{23}$. The vector $\xi = (-e_1, e_2, e_1 - e_2)$ is in $\ker F$, and K_{ij} is a subspace of $\ker F$ by definition. A direct computation shows that $K_{ij} \cap K_{ik} = 0$.

By assumption, we also have $\xi \notin K_{12} + K_{13} + K_{23}$. Therefore, the direct sum $\mathrm{span}(\xi) \oplus K_{12} \oplus K_{13} \oplus K_{23}$ is a subspace of $\ker F$. This proves the claim. Then it follows from (4.8) that

$$\begin{aligned} \mathrm{IEP}(W_1, W_2, W_3) &= n + \dim(\ker F) - \dim K_{12} - \dim K_{13} - \dim K_{23} \\ &\geq n + 1 \geq \dim(W_1 + W_2 + W_3) + 1. \end{aligned}$$

This completes the proof. ■

On the other hand, if $\{W_1, W_2, W_3\}$ is not a redundant triple of subspaces, then we have the dichotomy in the following statement. This criterion will be useful in the proof of Theorem E.

Lemma 4.35. *In the above notation, if $\{W_1, W_2, W_3\}$ is not a redundant triple of subspaces, then one of the following situations occurs:*

(1) *either $e_1, e_2 \in W_j^{\perp}$ for all $j = 1, 2, 3$, or*

(2) *there exists some $i \geq 3$ such that $e_i \notin W_j$ for all $j = 1, 2, 3$.*

Proof. Recall that K_{ij} is the image of the natural map $J_{ij} \circ I_{ij} : W_i^{\perp} \cap W_j^{\perp} \to W_1^{\perp} \oplus W_2^{\perp} \oplus W_3^{\perp}$ (see Figure 15 at the beginning of Section 4.4.2). We have an induced map

$$\begin{aligned} &K : (W_1^{\perp} \cap W_2^{\perp}) \oplus (W_2^{\perp} \cap W_3^{\perp}) \oplus (W_1^{\perp} \cap W_3^{\perp}) \to W_1^{\perp} \oplus W_2^{\perp} \oplus W_3^{\perp}, \\ &K(a, b, c) = (a + c, -a + b, -b - c), \end{aligned}$$

whose image is precisely $K_{12} + K_{23} + K_{13}$. Since $\{W_1, W_2, W_3\}$ is not a redundant triple of subspaces, we have $\xi = (-e_1, e_2, e_1 - e_2) \in \mathrm{Im}(K)$. This means that there exist $a =$

$\sum_{i=1}^{n} a_i e_i \in W_1^{\perp} \cap W_2^{\perp}$, $b = \sum_{i=1}^{n} b_i e_i \in W_2^{\perp} \cap W_3^{\perp}$, and $c = \sum_{i=1}^{n} c_i e_i \in W_1^{\perp} \cap W_3^{\perp}$, such that $a + c = -e_1$, $-a + b = e_2$, and $-b - c = e_1 - e_2$, where $a_i, b_i, c_i \in \mathbb{R}$. A direct computation shows that a, b, and c must satisfy the following relations:

$$a_1 = b_1 = -1 - c_1, \quad a_2 = -c_2 = b_2 - 1, \quad \text{and} \quad a_i = b_i = -c_i \quad \text{for } i \geq 3. \tag{4.9}$$

Note that if a_i, b_i, and c_i are equal to zero for all $i \geq 3$, then $a = a_1 e_1 + a_2 e_2 \in W_1^{\perp} \cap W_2^{\perp}$. Since $e_1 \in W_1^{\perp}$, we have $e_2 \in W_1^{\perp}$. Similar arguments show that e_1 and e_2 also belong to $W_2^{\perp}$ and $W_3^{\perp}$. Therefore, we are in case (1).

If (1) does not occur, then we can reduce to the case that one of a, b, and c has at least one non-zero coordinate along e_i for some $i \geq 3$. But $a_i \neq 0$ implies that $W_1^{\perp}$ and e_i are not orthogonal, so we have $e_i \notin W_1$. Thanks to (4.9), we also know that b and c have non-zero coordinates along e_i. Then a similar argument shows that $e_i \notin W_2, W_3$. Therefore, we are in case (2). ∎

Finally, we obtain a criterion to certify that a group is not a RAAG.

Proposition 4.36. *Let G be a finitely generated group. Suppose that there exist three maximal missing subspaces W_1, W_2, and W_3 in* $\mathrm{Hom}(G, \mathbb{R})$. *If they form a redundant triple of subspaces, then G is not a RAAG.*

Proof. Since $\{W_1, W_2, W_3\}$ is a redundant triple of subspaces, by Lemma 4.34 we have that $\dim(W_1 + W_2 + W_3) + 1 \leq \mathrm{IEP}(W_1, W_2, W_3)$. Assume by contradiction that G is a RAAG. Then by Lemma 4.32 we have $\dim(W_1 + W_2 + W_3) = \mathrm{IEP}(W_1, W_2, W_3)$. This leads to a contradiction. ∎

The fact that certain BBGs are not isomorphic to RAAGs can be obtained via the methods in [16] or [31], such as the BBG defined on the trefoil graph in Example 4.30. Proposition 4.36 allows us to obtain new examples that were not covered by previous criteria, such as the BBG defined on the extended trefoil (see Examples 4.31 and 5.14).

4.5. Redundant triples for BBGs

The purpose of this section is to find a general graphical criterion to certify that a BBG is not a RAAG. The idea is to start from a triangle in the flag complex Δ_Γ and find suitable subspaces of the links of its vertices that induce a redundant triple of subspaces in the complement of $\Sigma^1(\mathrm{BB}_\Gamma)$. Let τ be a triangle in Δ_Γ with vertices (v_1, v_2, v_3). Let e_j be the edge opposite to v_j. We say that τ is a *redundant triangle* if for each $j = 1, 2, 3$, there exists a subgraph $\Lambda_j \subseteq \mathrm{lk}(v_j, \Gamma)$ such that:

(1) $e_j \in E(\Lambda_j)$;

(2) Λ_j is a minimal separating subgraph of Γ;

(3) $\Lambda_1 \cap \Lambda_2 \cap \Lambda_3$ is the empty subgraph.

Example 4.37. The central triangle in the trefoil graph in Figure 1 is redundant. However, if we consider the cone over the trefoil graph, then the central triangle in the base trefoil graph is not redundant. Redundant triangles can appear in higher-dimensional complexes (see Example 5.17).

The purpose of this section is to prove the following theorem.

Theorem E. *Let Γ be a biconnected graph such that Δ_Γ is simply connected. If Γ has a redundant triangle, then* BB_Γ *is not a RAAG.*

We start by considering a redundant triangle τ with a choice of subgraph Λ_j of the link $\mathrm{lk}(v_j, \Gamma)$ as in the above definition of redundant triangle. We denote by $W_j = W_{\Lambda_j}$ the induced subspace of $V = \mathrm{Hom}(\mathrm{BB}_\Gamma, \mathbb{R})$. By Theorem D, we know that $W_j = W_{\Lambda_j}$ is a maximal subspace in the complement of $\Sigma^1(\mathrm{BB}_\Gamma)$. We want to show that $\{W_1, W_2, W_3\}$ is a redundant triple of subspaces. To do this, we will choose some suitable coordinates on V, that is, a suitable spanning tree for Γ. Notice that different spanning trees correspond to different bases on $\mathrm{Hom}(\mathrm{BB}_\Gamma, \mathbb{R})$. In particular, the linear isomorphism class of the arrangement of missing subspaces does not depend on these choices, and we can work with a convenient spanning tree.

To construct a desired spanning tree, we will need the following terminology. Let $v \in V(\Gamma)$. The *spoke* of v in Γ is the subgraph $\mathrm{spoke}(v)$ consisting of the edges that contain v. Note that $\mathrm{spoke}(v)$ is a spanning tree of $\mathrm{st}(v)$. Let Λ be a subgraph of $\mathrm{lk}(v)$. We define the *relative star* of v with respect to Λ to be the full subgraph $\mathrm{st}(v, \Lambda)$ of $\mathrm{st}(v)$ generated by $\{v\} \cup V(\Lambda)$. We define the *relative spoke* of v with respect to Λ to be the subgraph $\mathrm{spoke}(v, \Lambda)$ of $\mathrm{spoke}(v)$ consisting of the edges that connect v to a vertex of Λ. Note that $\mathrm{spoke}(v, \Lambda)$ is a spanning tree of $\mathrm{st}(v, \Lambda)$. We now construct a spanning tree T for Γ as follows.

- Let $T_3 = \mathrm{spoke}(v_3, \Lambda_3)$. Since we chose Λ_3 to contain e_3, we have $v_1, v_2 \in V(\Lambda_3)$ and $e_1, e_2 \in E(T_3)$.
- Let $Z_2 = \mathrm{spoke}(v_2, \Lambda_2 \setminus \mathrm{st}(v_3, \Lambda_3))$ and let $T_2 = T_3 \cup Z_2$. Notice that T_2 is a spanning tree of $\mathrm{st}(v_2, \Lambda_2) \cup \mathrm{st}(v_3, \Lambda_3)$.
- Let $Z_1 = \mathrm{spoke}(v_1, \Lambda_1 \setminus (\mathrm{st}(v_2, \Lambda_2) \cup \mathrm{st}(v_3, \Lambda_3)))$ and let $T_1 = T_2 \cup Z_1$. Notice that T_1 is a spanning tree of $\mathrm{st}(v_1, \Lambda_1) \cup \mathrm{st}(v_2, \Lambda_2) \cup \mathrm{st}(v_3, \Lambda_3)$.
- Finally, extend T_1 to a spanning tree T for Γ.

To fix notation, say $E(T) = \{f_1, f_2, \dots, f_n\}$. Without loss of generality, say $f_1 = e_1$ and $f_2 = e_2$. Fix an arbitrary orientation for the edges of T. With respect to the associated system of coordinates the subspaces W_1, W_2, and W_3 are given by equations of the form (4.4):

$$W_1 = \{y_1 = 0, \dots\}, \quad W_2 = \{y_2 = 0, \dots\}, \quad \text{and} \quad W_3 = \{y_1 - y_2 = 0, \dots\}.$$

Recall that $\{\chi_f \mid f \in E(T)\}$ is a basis for $\mathrm{Hom}(\mathrm{BB}_\Gamma, \mathbb{R})$, where $\chi_f : \mathrm{BB}_\Gamma \to \mathbb{R}$ is the character defined by $\chi_f(e) = 1$ if $f = e$ and $\chi_f(e) = 0$ if $f \neq e$. We also fix a background inner product with respect to which $\{\chi_f \mid f \in E(T)\}$ is an orthonormal basis. We now proceed to prove some technical lemmas that will be used to recognize the edges $f \in E(T)$ for which the associated character χ_f is in one of the subspaces W_1, W_2, and W_3. This is needed to use Lemma 4.35. We start with the following general fact.

Lemma 4.38. *Let $v \in \Gamma$, and let Λ be a subgraph of $\mathrm{lk}(v)$. Let T_Λ be a spanning tree of $\mathrm{st}(v, \Lambda)$. If $f \notin E(T_\Lambda)$, then $\chi_f \in W_\Lambda$.*

Proof. Suppose $f \notin E(T_\Lambda)$. Then $\chi_f = 0$ on T_Λ. Since T_Λ is a spanning tree of $\mathrm{st}(v, \Lambda)$, the character χ_f vanishes on $\mathrm{st}(v, \Lambda)$ by Lemma 4.13. In particular, it vanishes on Λ, hence $\chi_f \in W_\Lambda$. ∎

We now proceed to use Lemma 4.38 for each Λ_j, with respect to a suitable choice of spanning tree for $\mathrm{st}(v_j, \Lambda_j)$.

Lemma 4.39. *Let $f \in E(T)$. If $f \notin E(T_3)$, then $\chi_f \in W_3$.*

Proof. Since T_3 is a spanning tree of $\mathrm{st}(v_3, \Lambda_3)$, the statement follows directly from Lemma 4.38. ∎

Lemma 4.40. *Let $f \in E(T)$. If $f \notin E(Z_2)$, $f \neq e_1$, and f does not join v_3 to a vertex in $\Lambda_2 \cap \Lambda_3$, then $\chi_f \in W_2$.*

Proof. We construct a spanning tree for $\mathrm{st}(v_2, \Lambda_2)$ as follows. First, note that Z_2 is a spanning tree of $\mathrm{st}(v_2, \Lambda_2 \setminus \mathrm{st}(v_3, \Lambda_3))$ by construction. If u is a vertex in $\mathrm{st}(v_2, \Lambda_2)$ but not in Z_2, then $u = v_3$ or $u \in V(\Lambda_3)$. Let T_2' be the result of extending Z_2 with the edge $e_1 = (v_2, v_3)$ and all the edges that join v_3 to the vertices in $\Lambda_2 \cap \Lambda_3$. This gives a spanning subgraph T_2' of $\mathrm{st}(v_2, \Lambda_2)$. Note that T_2' is a tree because it is a subgraph of T. By the choice of f, we have $f \notin E(T_2')$. Then it follows from Lemma 4.38 that $\chi_f \in W_2$. ∎

Lemma 4.41. *Let $f \in E(T)$. If $f \notin E(Z_1)$, $f \neq e_1, e_2$, and f does not join v_3 to a vertex in $\Lambda_1 \cap \Lambda_3$ nor v_2 to a vertex in $\Lambda_1 \cap \Lambda_2$, then $\chi_f \in W_1$.*

Proof. We construct a spanning tree for $\mathrm{st}(v_1, \Lambda_1)$ as follows. First, note that Z_1 is a spanning tree for $\mathrm{st}(v_1, \Lambda_1 \setminus (\mathrm{st}(v_2, \Lambda_2) \cup \mathrm{st}(v_3, \Lambda_3)))$ by construction. If u is a vertex in $\mathrm{st}(v_1, \Lambda_1)$ but not in Z_1, then either $u = v_2$, $u = v_3$, $u \in V(\Lambda_2)$, or $u \in V(\Lambda_3)$. Let T_1' be the result of extending Z_1 with the edges $e_1 = (v_2, v_3)$, $e_2 = (v_1, v_3)$, all the edges that join v_3 to the vertices in $\Lambda_1 \cap \Lambda_3$, and all the edges that join v_2 to the vertices in $\Lambda_1 \cap \Lambda_2$. This gives a spanning subgraph T_1' of $\mathrm{st}(v_1, \Lambda_1)$. Note that T_1' is a tree because it is a subgraph of T. By the choice of f, we have $f \notin E(T_1')$. Then it follows from Lemma 4.38 that $\chi_f \in W_1$. ∎

Lemma 4.42. *Let $f \in E(T)$. If $\chi_f \notin W_j$ for all $j = 1, 2, 3$, then $f = e_1$ or $f = e_2$.*

Proof. By contradiction, suppose that there is an edge $f \neq e_1, e_2$ such that $\chi_f \notin W_j$ for all $j = 1, 2, 3$. Since $\chi_f \notin W_3$, we know that $f \in E(T_3)$ by Lemma 4.39. In particular, we have $v_3 \in V(f)$. Since $f \neq e_1, e_2$, this implies that $v_1, v_2 \notin V(f)$, and in particular this means that $f \notin E(Z_1), E(Z_2)$. The assumption $\chi_f \notin W_2$ implies that f joins v_3 to a vertex in $\Lambda_2 \cap \Lambda_3$, thanks to Lemma 4.40. Similarly, the assumption $\chi_f \notin W_1$ implies that f joins v_3 to a vertex in $\Lambda_1 \cap \Lambda_3$, thanks to Lemma 4.41. Therefore, we have obtained that f connects v_3 to a vertex in $\Lambda_1 \cap \Lambda_2 \cap \Lambda_3$. But this is absurd because this intersection is empty, by condition (3) in the definition of redundant triangle. ■

We are now ready for the proof of Theorem E.

Proof of Theorem E. Recall by construction that the subspaces W_1, W_2, and W_3 are given by equations of the form (4.4) with respect to the coordinates defined by the spanning tree T constructed above. Suppose by contradiction that $\{W_1, W_2, W_3\}$ is not a redundant triple By Lemma 4.35, one of the following cases occurs:

(1) either $\chi_{e_1}, \chi_{e_2} \in W_j^{\perp}$ for all $j = 1, 2, 3$, or

(2) there exists some $i \geq 3$ such that $\chi_{f_i} \notin W_j$ for all $j = 1, 2, 3$.

We claim that neither of these two situations can occur in our setting. To see that (1) does not occur, observe that $\chi_{e_1} + \chi_{e_2} \in W_3$ and $\chi_{e_1} + \chi_{e_2}$ are not orthogonal to χ_{e_1}, so $\chi_{e_1} \notin W_3^{\perp}$. The same is true for χ_{e_2}. On the other hand, (2) does not occur by Lemma 4.42. We have reached a contradiction, so $\{W_1, W_2, W_3\}$ is a redundant triple of subspaces. Then it follows from Proposition 4.36 that BB_Γ is not a RAAG. ■

We will use Theorem E in Section 5 to prove that certain BBGs are not isomorphic to RAAGs (see Theorem A for the case in which Δ_Γ is 2-dimensional and Example 5.17 for a higher-dimensional example).

4.6. Resonance varieties for BBGs

The goal of this section is to show that for a finitely presented BBG, the complement of its BNS-invariant coincides with the restriction of its first real resonance variety to the character sphere.

Let $A = H^*(G, \mathbb{R})$ be the cohomology algebra of G over $\mathbb{R}$. For each $a \in A^1 = H^1(G, \mathbb{R}) = \mathrm{Hom}(G, \mathbb{R})$, we have $a^2 = 0$. So, we can define a cochain complex (A, a)

$$(A, a) : A^0 \to A^1 \to A^2 \to \cdots,$$

where the coboundary is given by the right-multiplication by a. The *(first) resonance variety* is defined to be the set of points in A^1 so that the above chain complex fails to be exact, that is,

$$\mathcal{R}_1(G) = \{a \in A^1 \mid H^1(A, a) \neq 0\}.$$

In many cases of interest, the resonance variety $\mathcal{R}_1(G)$ is an affine algebraic subvariety of the vector space $A^1 = H^1(G, \mathbb{R}) = \mathrm{Hom}(G, \mathbb{R})$. For G a RAAG or a BBG, these varieties have been computed in [30] and [31], respectively. These varieties turn out to be defined by linear equations; that is, they are subspace arrangements. Following the notation in [31], let Γ be a finite graph. For any $U \subseteq V(\Gamma)$, let H_U be the set of characters $\chi : A_\Gamma \to \mathbb{R}$ vanishing (at least) on all the vertices in the complement of U. In our notation, this means $U \subseteq \mathcal{D}(\widehat{\chi})$. Moreover, let H'_U be the image of H_U under the restriction map $r\colon \mathrm{Hom}(A_\Gamma, \mathbb{R}) \to \mathrm{Hom}(\mathrm{BB}_\Gamma, \mathbb{R})$ (see Section 4.1.3).

Proposition 4.43. *Let Γ be a biconnected graph with Δ_Γ simply connected. Then* $\Sigma^1(\mathrm{BB}_\Gamma)^c = \mathcal{R}_1(\mathrm{BB}_\Gamma) \cap S(\mathrm{BB}_\Gamma)$.

Proof. By [31, Theorem 1.4], we have that $\mathcal{R}_1(\mathrm{BB}_\Gamma)$ is the union of the subspaces H'_U, where U runs through the maximal collections of vertices that induce disconnected subgraphs. Similarly, it follows from Theorem D that $\Sigma^1(\mathrm{BB}_\Gamma)^c$ is the union of the subspheres S_Λ, where Λ runs through the minimal separating subgraphs. Note that U is a maximal collection of vertices inducing a disconnected subgraph precisely when the subgraph Λ induced by $V(\Gamma) \setminus U$ is a minimal separating subgraph. So, it is enough to show that for each such U, we have $H'_U = W_\Lambda$, where W_Λ is the linear span of the sphere S_Λ, as defined above in Section 4.3.

To show this equality, let $\chi : \mathrm{BB}_\Gamma \to \mathbb{R}$. Then $\chi \in H'_U$ if and only if there is an extension $\widehat{\chi}$ of χ to A_Γ such that $\widehat{\chi} \in H_U$. This means that $\Lambda \subseteq \mathcal{D}(\widehat{\chi})$. Note that Λ is connected by Lemma 4.17 (3). So, by Lemma 4.15, we have $\Lambda \subseteq \mathcal{D}(\widehat{\chi})$ if and only if $\Lambda \subseteq \mathcal{DE}(\chi)$, which is equivalent to $\chi \in W_\Lambda$. ∎

Remark 4.44. We note that in general one has the inclusion $\Sigma^1(G)^c \subseteq \mathcal{R}_1(G) \cap S(G)$ thanks to [32, Theorem 15.8]. However, the equality does not always hold; see [36, §8] for some examples, such as the Baumslag–Solitar group $\mathrm{BS}(1,2)$.

We now recall a construction that reduces an Artin group to a RAAG (see [20, §11.9] or [31, §9]). Let (Γ, m) be a weighted graph, where $m\colon E(\Gamma) \to \mathbb{N}$ is an assignment of positive integers on the edge set. We denote by $A_{\Gamma,m}$ the associated *Artin group*. When $m = 2$ on every edge, it reduces to $A_{\Gamma,m} = A_\Gamma$. The *odd contraction* of (Γ, m) is an unweighted graph $\widetilde{\Gamma}$ defined as follows. Let Γ_{odd} be the graph whose vertex set is $V(\Gamma)$ and edge set is $\{e \in E(\Gamma) \mid m(e) \text{ is an odd number}\}$. The vertex set $V(\widetilde{\Gamma})$ of $\widetilde{\Gamma}$ is the set of connected components of Γ_{odd}, and two vertices C and C' are connected by an edge if there exist adjacent vertices $v \in V(C)$ and $v' \in V(C)'$ in the original graph Γ.

Corollary 4.45. *Let Γ be a biconnected graph such that Δ_Γ is simply connected. If Γ has a redundant triangle, then* BB_Γ *is not an Artin group.*

Proof. Let τ be a redundant triangle, with chosen minimal full separating subgraphs $\{\Lambda_1, \Lambda_2, \Lambda_3\}$. Let $W_j = W_{\Lambda_j}$ be the subspace of $V = \mathrm{Hom}(\mathrm{BB}_\Gamma, \mathbb{R})$ defined by Λ_j.

Arguing as in the proof of Theorem E, we have that W_j is a maximal missing subspace in the complement of $\Sigma^1(\mathrm{BB}_\Gamma)$, and that $\{W_1, W_2, W_3\}$ is a redundant triple in subspaces of V. By Lemma 4.34, we have

$$\dim(W_1 + W_2 + W_3) + 1 \leq \mathrm{IEP}(W_1, W_2, W_3).$$

Now, assume by contradiction that BB_Γ is isomorphic to an Artin group $A_{\Gamma',m}$. Let $\widetilde{\Gamma}'$ be the odd contraction of (Γ', m). Then $A_{\widetilde{\Gamma}'}$ is a RAAG. Notice that $A_{\Gamma',m}$ and $A_{\widetilde{\Gamma}'}$ have the same abelianization. Hence, the three spaces $\mathrm{Hom}(A_{\Gamma',m}, \mathbb{R})$, $\mathrm{Hom}(A_{\widetilde{\Gamma}'}, \mathbb{R})$, and V can be identified together. In particular, the three character spheres $S(A_{\Gamma,m})$, $S(A_{\widetilde{\Gamma}'})$, and $S(\mathrm{BB}_\Gamma)$ can be identified as well.

Arguing as in [31, Proposition 9.4], there is an ambient isomorphism of the resonance varieties $\mathcal{R}_1(\mathrm{BB}_\Gamma) \cong \mathcal{R}_1(A_{\Gamma',m}) \cong \mathcal{R}_1(A_{\widetilde{\Gamma}'})$, seen as subvarieties of V. Since $A_{\widetilde{\Gamma}'}$ is a RAAG, by [30, Theorem 5.5] we have $\Sigma^1(A_{\widetilde{\Gamma}'})^c = \mathcal{R}_1(A_{\widetilde{\Gamma}'}) \cap S(A_{\widetilde{\Gamma}'})$. Similarly, since BB_Γ is a BBG, by Proposition 4.43 we have $\Sigma^1(\mathrm{BB}_\Gamma)^c = \mathcal{R}_1(\mathrm{BB}_\Gamma) \cap S(\mathrm{BB}_\Gamma)$. It follows that we have an ambient isomorphism of the complements of the BNS-invariant $\Sigma^1(\mathrm{BB}_\Gamma)^c \cong \Sigma^1(A_{\widetilde{\Gamma}'})^c$ (seen as arrangements of subspheres in $S(\mathrm{BB}_\Gamma)$), as well as an ambient isomorphism of the associated arrangements of (linear) subspaces of V. In particular, the arrangement of maximal missing subspaces of BB_Γ inside V is ambient isomorphic to the arrangement of maximal missing subspaces of a RAAG. Applying Lemma 4.32 to the triple $\{W_1, W_2, W_3\}$ gives

$$\mathrm{IEP}(W_1, W_2, W_3) = \dim(W_1 + W_2 + W_3).$$

This leads to a contradiction. ■

5. BBGs on 2-dimensional flag complexes

If Δ_Γ is a simply connected flag complex of dimension 1, then Γ is a tree. In this case, the group BB_Γ is a free group generated by all the edges of Γ, and in particular, it is a RAAG. The goal of this section is to determine what happens in dimension 2. Namely, we will show that the BBG defined on a 2-dimensional complex is a RAAG if and only if a certain poison subgraph is avoided. We will discuss some higher-dimensional examples at the end (see Examples 5.16 and 5.17).

Throughout this section, we assume that Γ is a biconnected graph such that Δ_Γ is 2-dimensional and simply connected unless otherwise stated. Note that by Lemma 4.17 this implies that Δ_Γ is homogeneous of dimension 2. We say that

- An edge e is a *boundary edge* if it is contained in exactly one triangle. Denote by $\partial\Delta_\Gamma$ the *boundary* of Δ_Γ. This is a 1-dimensional subcomplex consisting of boundary edges. An edge e is an *interior edge* if $e \cap \partial\Delta_\Gamma = \varnothing$. Equivalently, none of its vertices is on the boundary.

- A *boundary vertex* is a vertex contained in $\partial\Delta_\Gamma$. Equivalently, it is contained in at least one boundary edge. A vertex v is an *interior vertex* if it is contained only in edges that are not boundary edges.
- A triangle τ is an *interior triangle* if $\tau \cap \partial\Delta_\Gamma = \varnothing$. A triangle τ is called a *crowned triangle* if none of its edges is in $\partial\Delta_\Gamma$. This is weaker than being an interior triangle because a crowned triangle can have vertices in $\partial\Delta_\Gamma$. If τ is a crowned triangle, each of its edges is contained in at least one triangle different from τ.

Remark 5.1. We will prove in Lemma 5.13 that in dimension 2, a crowned triangle is redundant in the sense of Section 4.5. If $\partial\Delta_\Gamma$ is empty, then every triangle is crowned, simply because no edge can be a boundary edge. Note that a vertex is either a boundary vertex or an interior vertex, but we might have edges which are neither boundary edges nor interior edges. For example, the trefoil graph (see Figure 1) has no interior edges, but only six of its nine edges are boundary edges. Moreover, it has no interior triangles, but it has one crowned triangle. Notice that a crowned triangle is contained in a trefoil subgraph of Γ, but the trefoil subgraph is not necessarily a full subgraph of Γ (see Figure 16).

5.1. Complexes without crowned triangles

The goal of this section is to provide a characterization of complexes without crowned triangles.

Lemma 5.2. *A vertex $v \in V(\Gamma)$ is an interior vertex if and only if for each vertex w in $\mathrm{lk}(v,\Gamma)$, its degree in $\mathrm{lk}(v,\Gamma)$ is at least 2.*

Proof. First of all, notice that for a vertex w in $\mathrm{lk}(v,\Gamma)$, its degree in $\mathrm{lk}(v,\Gamma)$ is equal to the number of triangles of Δ_Γ that contain the edge (v,w).

Suppose that $v \in V(\Gamma)$ is an interior vertex, and let w be a vertex of $\mathrm{lk}(v,\Gamma)$. Since v is interior, the edge (v,w) is not a boundary edge, hence it is contained in at least two triangles. Therefore, the vertex w has degree at least 2 in $\mathrm{lk}(v,\Gamma)$. Conversely, let $v \in V(\Gamma)$ and let $e = (v,w)$ be an edge containing v, where w is some vertex in $\mathrm{lk}(v,\Gamma)$. Since the

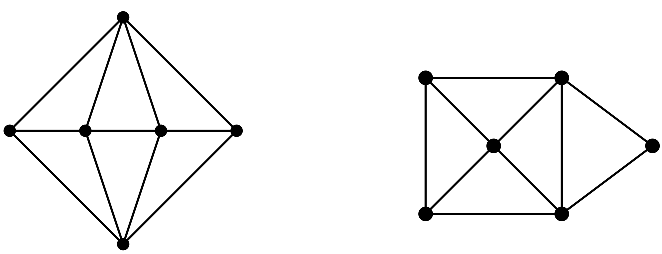

Figure 16. Graphs that contain crowned triangles, but the resulting trefoil subgraphs are not full subgraphs.

degree of w in $\mathrm{lk}(v,\Gamma)$ is at least 2, the edge e must be contained in at least two triangles. Thus, the edge e is not in $\partial\Delta_\Gamma$. Hence, the vertex v is an interior vertex. ∎

Lemma 5.3. *Let Γ be a biconnected graph such that Δ_Γ is 2-dimensional and simply connected. If Δ_Γ has no crowned triangles, then Δ_Γ has no interior triangles, no interior edges, and has at most one interior vertex.*

Proof. Since an interior triangle is automatically a crowned triangle, it is clear that Δ_Γ has no interior triangles.

For the second statement, assume that there is an interior edge $e = (u, v)$ of Δ_Γ. Since e is an interior edge, it is contained in at least two triangles. Let τ be a triangle containing e. Let w be the third vertex of τ, and let e_1 and e_2 be the other two edges of τ. Since u and v are interior vertices, we have that e_1 and e_2 are not in $\partial\Delta_\Gamma$. So, no edge of τ is a boundary edge. That is, the triangle τ is a crowned triangle, a contradiction.

Finally, let v be an interior vertex. By definition, none of the edges containing v is in $\partial\Delta_\Gamma$. We claim that $\Gamma = \mathrm{st}(v,\Gamma)$, and in particular, there are no other interior vertices. First, take a triangle τ containing v, then the two edges of τ that meet at v are not in $\partial\Delta_\Gamma$. The third edge of τ must be a boundary edge; otherwise, the triangle τ would be a crowned triangle. This shows that all vertices in $\mathrm{lk}(v,\Gamma)$ are in $\partial\Delta_\Gamma$. Now, assume by contradiction that there is a vertex u at distance two from v. Let w be a vertex in $\mathrm{lk}(v,\Gamma)$ that is adjacent to u. Note that w is a boundary vertex. Since $\mathrm{lk}(w,\Gamma)$ is connected by Lemma 4.17 (2), there is a path p in $\mathrm{lk}(w,\Gamma)$ from u to a vertex u' in $\mathrm{lk}(v,\Gamma) \cap \mathrm{lk}(w,\Gamma)$ (see Figure 17). Then the path p, together with the edges (w,u) and (w,u'), bounds a triangulated disk in Δ_Γ. Then the edge (w,u') is contained in more than one triangle, and therefore, it is not a boundary edge, and the triangle formed by the vertices v, w, and u' is a crowned triangle, a contradiction. ∎

Before we prove the next result, we give some terminologies on graphs. A graph Γ is called an *edge-bonding* of two graphs Γ_1 and Γ_2 if it is obtained by identifying two

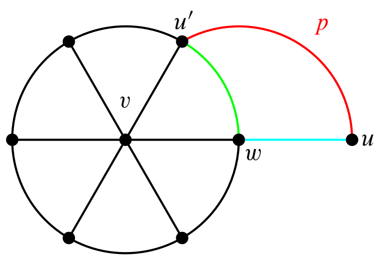

Figure 17. The path p and the edges (w,u) and (w,u') bound a triangulated disk in Δ_Γ. This implies that the edge (w,u') is not a boundary edge, and the vertices v, w, and u' form a crowned triangle.

edges $e_1 \in E(\Gamma_1)$ and $e_2 \in E(\Gamma_2)$. If e denotes the image of e_1 and e_2 in Γ, we also write $\Gamma = \Gamma_1 \cup_e \Gamma_2$ and say that e is the *bonding edge*.

Remark 5.4. Since an edge-bonding involves identifying two edges from two different graphs, one can perform several edge-bondings of a collection of graphs simultaneously. In particular, if one performs a sequence of edge-bondings, then the result can actually be obtained by a simultaneous edge-bonding.

We also note that there are two ways of identifying e_1 and e_2 that can result in two different graphs. However, this will not be relevant in the following.

Our goal is to decompose a given graph as an edge-bonding of certain elementary pieces that we now define. A *fan* is a cone over a path. Let Γ_0 be a connected graph having no vertices of degree 1 and whose associated flag complex Δ_{Γ_0} is 1-dimensional. Note that Γ_0 contains no triangles. The cone over such a Γ_0 is called a *simple cone* (see Figure 18 for an example).

Remark 5.5. Fans and simple cones could be further decomposed via edge-bonding by disconnecting them along a cut edge. For example, a fan can be decomposed into triangles. However, we will not take this point of view. Instead, it will be convenient to decompose a graph into fans and simple cones and regard them as elementary pieces.

It follows from Corollary 3.10 that the BBG defined on a fan or simple cone is a RAAG. Here are some further properties of fans and simple cones that follow directly from the definitions.

Lemma 5.6. *Let Γ be a fan or a simple cone. The following statements hold.*

(1) *The flag complex Δ_Γ is 2-dimensional, simply connected, and contractible.*

(2) *The flag complex Δ_Γ has no interior edges, no interior triangles, and no crowned triangles.*

(3) *If $\Gamma = \{v\} * P$ is a fan over a path P with endpoints u and w, then $\partial\Delta_\Gamma = P \cup \{(v,u),(v,w)\}$, and there are no interior vertices.*

(4) *If $\Gamma = \{v\} * \Gamma_0$ is a simple cone over Γ_0, then $\partial\Delta_\Gamma = \Gamma_0$, and the cone vertex v is the only interior vertex.*

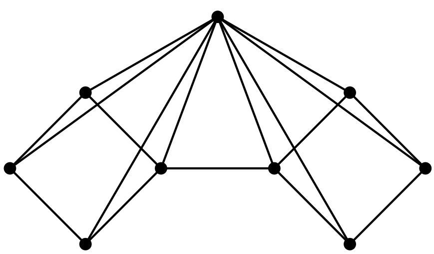

Figure 18. A simple cone.

Lemma 5.7. *Let Γ be a biconnected graph such that Δ_Γ is 2-dimensional and simply connected. Suppose that Δ_Γ has no crowned triangles. Then Γ decomposes as edge-bondings of fans and simple cones.*

Proof. We argue by induction on the number of cut edges of Γ. Suppose that Γ has no cut edges. By Lemma 5.3, the complex Δ_Γ contains at most one interior vertex. We claim that if Δ_Γ contains no interior vertices, then Γ is a fan. Let $v \in V(\Gamma)$. Since v is a boundary vertex, its link has degree 1 vertices by Lemma 5.2. Moreover, since Γ has no cut edges, the link of v has no cut vertices. Then $\mathrm{lk}(v, \Gamma)$ must be a single edge, and therefore, the graph Γ is a triangle, which is a fan. Thus, the claim is proved. If Δ_Γ contains one interior vertex u, then $\Gamma = \mathrm{st}(u, \Gamma)$ as in the proof of Lemma 5.3. So, the graph Γ is the cone over $\mathrm{lk}(u, \Gamma)$. Since u is an interior vertex, its link has no degree 1 vertices. Note that the flag complex on $\mathrm{lk}(u, \Gamma)$ is 1-dimensional; otherwise, the dimension of Δ_Γ would be greater than 2. Thus, the graph $\Gamma = \mathrm{st}(u, \Gamma)$ is a simple cone. This proves the base case of induction.

Suppose that the conclusion holds for graphs having n cut edges, $n \geq 1$. Assume that Γ has $n + 1$ cut edges. Let e be a cut edge of Γ. Cutting along e gives some connected components $\Gamma_1, \ldots, \Gamma_k$. Each of these components, as a full subgraph of Γ, satisfies all the assumptions of the lemma and has at most n cut edges. By induction, the subgraphs $\Gamma_1, \ldots, \Gamma_k$ are edge-bondings of fans and simple cones. Therefore, the graph Γ, as an edge-bonding of $\Gamma_1, \ldots, \Gamma_k$, is also an edge-bonding of fans and simple cones. ∎

Remark 5.8. The decomposition in Lemma 5.7 is not unique (for instance, it is not maximal; see Remark 5.5). We do not need this fact in this paper.

We now proceed to study the ways in which one can perform edge-bondings of fans and simple cones. Recall from Section 4.5 that the spoke of a vertex v is the collection of edges containing v. When Γ is a fan, write $\Gamma = \{v\} * P_n$, where P_n is the path on n labeled vertices (see Figure 19). We call the edges (v, w_1) and (v, w_n) *peripheral edges*, and the edges (w_1, w_2) and (w_{n-1}, w_n) are called *modified-peripheral edges*. A *peripheral triangle* is a triangle containing a peripheral edge and a modified-peripheral edge.

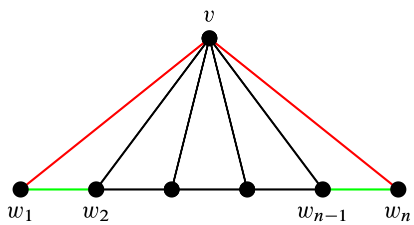

Figure 19. The red edges are peripheral edges, and the green edges are modified-peripheral edges. The left-most and right-most triangles are peripheral triangles.

We say that an edge of a fan is *good* if either it belongs to the spoke or it is a modified-peripheral edge. Similarly, we say that an edge of a simple cone is *good* if it belongs to the spoke. We say an edge is *bad* if it is not good. Note that a bad edge is necessarily a boundary edge (see Lemma 5.6). We extend this definition to more general graphs as follows: Let Γ be a graph obtained via an edge-bonding on a collection of fans and simple cones, and let $e \in E(\Gamma)$ be a bonding edge of Γ. We say that e is *good* if it is good in each fan component or simple cone component of Γ that contains e. We say that e is *bad* otherwise. These concepts are motivated by the fact that forming edge-bonding along good edges does not create crowned triangles; see the following example.

Example 5.9. Let Γ_1 and Γ_2 be a fan and a simple cone, respectively. If we form the edge-bonding of Γ_1 and Γ_2 by identifying a good edge in each of them, the resulting graph has no crowned triangles. The situation is analogous if Γ_1 and Γ_2 are both fans or both simple cones.

Lemma 5.10. *Let $\Gamma = \Gamma_1 \cup_e \Gamma_2$, where Γ_1 is a fan or a simple cone, and Γ_2 is any graph obtained via edge-bonding of fans and simple cones. If e is a bad edge of Γ_1, then Γ contains a crowned triangle.*

Proof. If $e \in E(\Gamma_1)$ is bad, then it is in $\partial\Delta_{\Gamma_1}$. In particular, there is a unique triangle τ of Γ_1 containing e (namely, the cone over e), and the other two edges of τ are not boundary edges (in the case of a fan, recall that a modified-peripheral edge is good). When we form an edge-bonding along e, the edge e is no longer a boundary edge in Γ, so τ becomes a crowned triangle in Γ. ■

Proposition 5.11. *Let Γ be a biconnected graph such that Δ_Γ is 2-dimensional and simply connected. Then Γ admits a tree 2-spanner if and only if Δ_Γ does not contain crowned triangles.*

Proof. Let T be a tree 2-spanner of Γ. Suppose by contradiction that Γ contains a crowned triangle τ whose edges are e, f, and g. By Lemma 3.2, either two of e, f, and g are in $E(T)$ or none of them is in $E(T)$. If e, f, and g are not in $E(T)$, then by Lemma 3.3, the graph Γ contains a K_4. This contradicts the fact that Δ_Γ is 2-dimensional. Now consider the case that $e \notin E(T)$ and f and g are in $E(T)$. Since τ is a crowned triangle, the edge e is not on the boundary of Δ_Γ, and there is another triangle τ' based on e that is different from τ. Denote the other edges of τ' by f' and g'. Note that f' and g' cannot be in $E(T)$ by the uniqueness part of Lemma 3.1. This means that none of the edges of τ' is in $E(T)$. Again, by Lemma 3.3 we obtain a K_4, hence a contradiction. Therefore, the graph Γ has no crowned triangles.

Conversely, suppose that Δ_Γ has no crowned triangles. By Lemma 5.7 the graph Γ decomposes as edge-bondings of some fans and simple cones $\Gamma_1, \dots, \Gamma_m$. Let $\{e_1, \dots, e_n\}$ be the set of bonding edges. Note that by Remark 5.4 these edge-bonding operations can be performed simultaneously. Since Γ has no crowned triangles, by Lemma 5.10, each of the edges in $\{e_1, \dots, e_n\}$ is good. We now construct a tree 2-spanner for Γ. We do this by constructing a tree 2-spanner T_i for each Γ_i and then gluing them together. For a simple cone component Γ_i, choose T_i to be the spoke. For a fan component Γ_i, write $\Gamma_i = \{v_i\} * P_{n_i}$ and order the vertices of P_{n_i} as $w_1, \dots, w_{n_i}$. Define T_i to consist of the edges $(v_i, w_{n_2}), \dots, (v_i, w_{n_{i-1}})$, together with two more edges, one from each peripheral triangle, chosen as follows. If the peripheral edge or the modified-peripheral edge in a peripheral triangle is involved in some edge-bondings, then choose that edge to be in T_i. If none of them is involved in any edge-bonding, then choose either one of them. Note that it is not possible that both the peripheral edge and the modified-peripheral edge of the same peripheral triangle are involved in edge-bondings; otherwise, the graph Γ would contain a crowned triangle. In all the cases, this provides a tree 2-spanner T_i in Γ_i. Moreover, if e is a bonding edge for Γ that appears in a component Γ_i, then e is in T_i. It follows from [9, Theorem 4.4] that $T = \bigcup_{i=1}^{m} T_i$ is a tree 2-spanner of Γ. ∎

Remark 5.12. When Γ is a 2-tree (recall from Section 3.2.2), the flag complex Δ_Γ is a biconnected contractible 2-dimensional flag complex. In [8], Cai showed that a 2-tree admits a tree 2-spanner if and only if it does not contain a trefoil subgraph (see Figure 1). Proposition 5.11 generalizes Cai's result to any biconnected and simply connected 2-dimensional flag complex. Note that a trefoil subgraph in a 2-tree is necessarily full, but this is not the case in general (see Figure 16).

5.2. The RAAG recognition problem in dimension 2

In this section, we provide a complete answer to the RAAG recognition problem on 2-dimensional complexes. In other words, we completely characterize the graphs Γ such that BB_Γ is a RAAG, under the assumption $\dim \Delta_\Gamma = 2$.

Observe that a RAAG is always finitely presented (recall that all graphs are finite in our setting). On the other hand, by [3, the Main Theorem 3], a BBG is finitely presented precisely when the defining flag complex is simply connected. Therefore, we can assume that Δ_Γ is simply connected. Moreover, by Corollary 4.22 we can assume that Γ is also biconnected. Note that RAAGs are actually groups of type F, so one could even restrict to the case that Δ_Γ is contractible, thanks to [3, the main theorem]; compare this with Corollary 3.9. However, we do not need this fact. We start by showing that in dimension 2, any crowned triangle is redundant.

Lemma 5.13. *If* $\dim(\Delta_\Gamma) = 2$*, then every crowned triangle is a redundant triangle.*

Proof. Let τ be a crowned triangle with edges e_1, e_2, e_3 and vertices v_1, v_2, v_3, where v_j is opposite to e_j. Since τ is a crowned triangle, no edge e_j is a boundary edge. Hence, there is another triangle τ_j adjacent to τ along e_j. Let u_j be the vertex of τ_j not in τ. If u_j were adjacent to v_j, then we would have a K_4, which is impossible since $\dim \Delta_\Gamma = 2$. Thus, the vertices v_j and u_j are not adjacent. As a consequence, we can choose a full subgraph $\Lambda_j \subseteq \mathrm{lk}(v_j, \Gamma)$ that contains e_j and is a minimal full separating subgraph of Γ.

Finally, note that the intersection $\Lambda_1 \cap \Lambda_2 \cap \Lambda_3$ cannot contain any vertex. Otherwise, we would see a K_4, which is against the assumption that Δ_Γ is 2-dimensional. ■

Theorem A. *Let Γ be a biconnected graph such that Δ_Γ is 2-dimensional and simply connected. Then the following statements are equivalent.*

(1) *Γ admits a tree 2-spanner.*

(2) *Δ_Γ does not contain crowned triangles.*

(3) *BB_Γ is a RAAG.*

(4) *BB_Γ is an Artin group.*

Proof. The implications (1) $\Leftrightarrow$ (2) follow from Proposition 5.11. Moreover, the implication (1) $\Rightarrow$ (3) is Theorem B. The implication (3) $\Rightarrow$ (4) is obvious.

We prove the implication (3) $\Rightarrow$ (2) as follows. Assume that Δ_Γ contains a crowned triangle τ. Then by Lemma 5.13 we know that τ is also a redundant triangle. Then it follows from Theorem E that BB_Γ is not a RAAG. The implication (4) $\Rightarrow$ (2) is obtained in a similar way, using Corollary 4.45 instead of Theorem E. ■

Papadima and Suciu in [31, Proposition 9.4] showed that if Δ_Γ is a certain type of triangulation of the 2-disk (which they call *extra-special triangulation*), then BB_Γ is not a RAAG. Those triangulations always contain a crowned triangle, so Theorem A recovers Papadima–Suciu's result and extends it to a wider class of graphs, such as arbitrary triangulations of disks (see Example 5.14), or even flag complexes that are not triangulations of disks (see Example 3.17).

Example 5.14 (The extended trefoil continued). Let Γ be the graph in Figure 14. Since Γ contains a crowned triangle, the group BB_Γ is not a RAAG by Theorem A. Note that this fact does not follow from [31]: The flag complex Δ_Γ is a triangulation of the disk but not an extra-special triangulation. This fact also does not follow from [16], because all the subspace arrangement homology groups vanish for this group BB_Γ, that is, they look like those of a RAAG (as observed in Example 4.31).

Remark 5.15. The criterion for a BBG to be a RAAG from Theorem B works in any dimension. On the other hand, Theorem A fails for higher-dimensional complexes. Indeed, the mere existence of a crowned triangle is not very informative in higher-dimensional cases (see Example 5.16). However, the existence of a redundant triangle is an obstruction for a BBG to be a RAAG even in higher-dimensional complexes (see Example 5.17).

Example 5.16 (A crowned triangle in dimension 3 does not imply that the BBG is not a RAAG). Let Γ be the cone over the trefoil graph in Figure 1. Then Δ_Γ is 3-dimensional and Γ contains a crowned triangle (the one sitting in the trefoil graph). However, this crowned triangle is not a redundant triangle, and the group BB_Γ is actually a RAAG by Corollary 3.10.

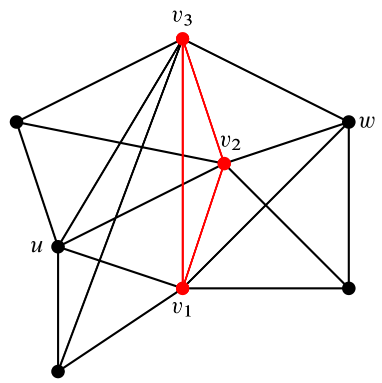

Figure 20. A 3-dimensional complex that contains a redundant triangle.

Example 5.17 (A redundant triangle in dimension 3 implies that the BBG is not a RAAG). Consider the graph Γ in Figure 20. Then Δ_Γ is 3-dimensional and every 3-simplex has a 2-face in $\partial\Delta_\Gamma$. However, we can show that this BB_Γ is not a RAAG. The triangle induced by the vertices v_1, v_2, and v_3 is a redundant triangle. Indeed, the full subgraphs Λ_1, Λ_2, and Λ_3 induced by the sets of vertices $\{u, v_2, v_3\}$, $\{u, v_1, v_3\}$, and $\{v_1, v_2, w\}$, respectively, satisfy condition (3) in the definition of redundant triangle. Then it follows from Theorem E that this BB_Γ is not a RAAG.

Acknowledgments. We thank Tullia Dymarz for bringing [27] to our attention and Daniel C. Cohen, Pallavi Dani, Max Forester, Wolfgang Heil, Alexandru Suciu, and Matthew Zaremsky for helpful conversations. We thank the School of Mathematics at the Georgia Institute of Technology for their hospitality during a visit in which part of this work was done. We thank the referee for the very careful reading and helpful comments. The second author acknowledges support from the AMS and the Simons Foundation.

References

[1] E. M. Barquinero, L. Ruffoni, and K. Ye, Graphical splittings of Artin kernels. *J. Group Theory* **24** (2021), no. 4, 711–735 Zbl 1485.20095 MR 4279130

[2] M. W. Bern, *Network design problems: Steiner trees and spanning k-trees*. PhD thesis, University of California, Berkeley, CA, 1987

[3] M. Bestvina and N. Brady, Morse theory and finiteness properties of groups. *Invent. Math.* **129** (1997), no. 3, 445–470 Zbl 0888.20021 MR 1465330

[4] R. Bieri and J. R. J. Groves, The geometry of the set of characters induced by valuations. *J. Reine Angew. Math.* **1984** (1984), no. 347, 168–195 Zbl 0526.13003 MR 0733052

[5] R. Bieri, W. D. Neumann, and R. Strebel, A geometric invariant of discrete groups. *Invent. Math.* **90** (1987), no. 3, 451–477 Zbl 0642.57002 MR 0914846

[6] M. R. Bridson, On the recognition of right-angled Artin groups. *Glasg. Math. J.* **62** (2020), no. 2, 473–475 Zbl 1481.20116 MR 4085052

[7] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*. Grundlehren Math. Wiss. 319, pp. 643, Springer, Berlin, 1999 Zbl 0988.53001 MR 1744486

[8] L. Cai, On spanning 2-trees in a graph. *Discrete Appl. Math.* **74** (1997), no. 3, 203–216 Zbl 0883.05040 MR 1444941

[9] L. Cai and D. G. Corneil, Tree spanners. *SIAM J. Discrete Math.* **8** (1995), no. 3, 359–387 Zbl 0832.05037 MR 1341545

[10] C. H. Cashen and G. Levitt, Mapping tori of free group automorphisms, and the Bieri-Neumann-Strebel invariant of graphs of groups. *J. Group Theory* **19** (2016), no. 2, 191–216 Zbl 1368.20048 MR 3466593

[11] Y.-C. Chang, Identifying Dehn functions of Bestvina-Brady groups from their defining graphs. *Geom. Dedicata* **214** (2021), 211–239 Zbl 07402681 MR 4308276

[12] Y.-C. Chang, Abelian splittings and JSJ-decompositions of finitely presented Bestvina-Brady groups. *J. Group Theory* **26** (2023), no. 4, 677–692 Zbl 1523.20062 MR 4609850

[13] R. Charney, K. Ruane, N. Stambaugh, and A. Vijayan, The automorphism group of a graph product with no SIL. *Illinois J. Math.* **54** (2010), no. 1, 249–262 Zbl 1243.20047 MR 2776995

[14] C. Cunningham, A. Eisenberg, A. Piggott, and K. Ruane, Recognizing right-angled Coxeter groups using involutions. *Pacific J. Math.* **284** (2016), no. 1, 41–77 Zbl 1397.20046 MR 3530862

[15] P. Dani and I. Levcovitz, Right-angled Artin subgroups of right-angled Coxeter and Artin groups. *Algebr. Geom. Topol.* **24** (2024), no. 2, 755–802 Zbl 07846434 MR 4735045

[16] M. B. Day and R. D. Wade, Subspace arrangements, BNS invariants, and pure symmetric outer automorphisms of right-angled Artin groups. *Groups Geom. Dyn.* **12** (2018), no. 1, 173–206 Zbl 1456.20043 MR 3781420

[17] P. Deshpande and M. Roy, On the structure of finitely presented Bestvina-Brady groups. *Internat. J. Algebra Comput.* **34** (2024), no. 1, 69–85 Zbl 07817036 MR 4716440

[18] W. Dicks and I. J. Leary, Presentations for subgroups of Artin groups. *Proc. Amer. Math. Soc.* **127** (1999), no. 2, 343–348 Zbl 0923.20032 MR 1605948

[19] A. Dimca, S. Papadima, and A. I. Suciu, Quasi-Kähler Bestvina-Brady groups. *J. Algebraic Geom.* **17** (2008), no. 1, 185–197 Zbl 1176.20037 MR 2357684

[20] A. Dimca, Ş. Papadima, and A. I. Suciu, Topology and geometry of cohomology jump loci. *Duke Math. J.* **148** (2009), no. 3, 405–457 Zbl 1222.14035 MR 2527322

[21] C. Droms, Isomorphisms of graph groups. *Proc. Amer. Math. Soc.* **100** (1987), no. 3, 407–408 Zbl 0619.20015 MR 0891135

[22] C. Droms, Subgroups of graph groups. *J. Algebra* **110** (1987), no. 2, 519–522 Zbl 0625.20026 MR 0910401

[23] F. Haglund and D. T. Wise, Special cube complexes. *Geom. Funct. Anal.* **17** (2008), no. 5, 1551–1620 Zbl 1155.53025 MR 2377497

[24] N. Koban and A. Piggott, The Bieri-Neumann-Strebel invariant of the pure symmetric automorphisms of a right-angled Artin group. *Illinois J. Math.* **58** (2014), no. 1, 27–41 Zbl 1332.20044 MR 3331840

[25] T. Koberda, Right-angled Artin groups and a generalized isomorphism problem for finitely generated subgroups of mapping class groups. *Geom. Funct. Anal.* **22** (2012), no. 6, 1541–1590 Zbl 1282.37024 MR 3000498

[26] D. H. Kochloukova, On the Bieri-Neumann-Strebel-Renz Σ^1-invariant of even Artin groups. *Pacific J. Math.* **312** (2021), no. 1, 149–169 Zbl 1535.20266 MR 4298796

[27] D. H. Kochloukova and L. Mendonça, On the Bieri-Neumann-Strebel-Renz Σ-invariants of the Bestvina-Brady groups. *Forum Math.* **34** (2022), no. 3, 605–626 Zbl 1497.20056 MR 4415959

[28] I. J. Leary and M. Saadetoğlu, The cohomology of Bestvina-Brady groups. *Groups Geom. Dyn.* **5** (2011), no. 1, 121–138 Zbl 1246.20046 MR 2763781

[29] J. Meier and L. VanWyk, The Bieri-Neumann-Strebel invariants for graph groups. *Proc. Lond. Math. Soc. (3)* **71** (1995), no. 2, 263–280 Zbl 0835.20037 MR 1337468

[30] S. Papadima and A. I. Suciu, Algebraic invariants for right-angled Artin groups. *Math. Ann.* **334** (2006), no. 3, 533–555 Zbl 1165.20032 MR 2207874

[31] S. Papadima and A. I. Suciu, Algebraic invariants for Bestvina-Brady groups. *J. Lond. Math. Soc. (2)* **76** (2007), no. 2, 273–292 Zbl 1176.20038 MR 2363416

[32] S. Papadima and A. I. Suciu, Bieri-Neumann-Strebel-Renz invariants and homology jumping loci. *Proc. Lond. Math. Soc. (3)* **100** (2010), no. 3, 795–834 Zbl 1273.55003 MR 2640291

[33] H. Servatius, Automorphisms of graph groups. *J. Algebra* **126** (1989), no. 1, 34–60 Zbl 0682.20022 MR 1023285

[34] J. Stallings, On fibering certain 3-manifolds. In *Topology of 3-manifolds and related topics*. Proc. The Univ. of Georgia Institute, 1961, pp. 95–100, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1961 Zbl 1246.57049 MR 0158375

[35] R. Strebel, Notes on the Sigma invariants. [v1] 2012, [v2] 2013, arXiv:1204.0214v2

[36] A. I. Suciu, Sigma-invariants and tropical varieties. *Math. Ann.* **380** (2021), no. 3-4, 1427–1463 Zbl 1482.20027 MR 4297190

[37] W. P. Thurston, A norm for the homology of 3-manifolds. *Mem. Amer. Math. Soc.* **59** (1986), no. 339, i–vi and 99–130 Zbl 0585.57006 MR 0823443

Received 8 May 2023.

Yu-Chan Chang
Department of Mathematics and Computer Science, Wesleyan University, 265 Church Street, Middletown, CT 06459, USA; ychang01@wesleyan.edu

Lorenzo Ruffoni
Department of Mathematics, Tufts University, 177 College Avenue, Medford, MA 02155; current address: Binghamton University, State University of New York, 4400 Vestal Parkway East, Binghamton, NY 13902, USA; lruffoni@binghamton.edu