

RIGHT-ANGLED ARTIN GROUPS AND THE COHOMOLOGY BASIS GRAPH

RAMÓN FLORES¹, DELARAM KAHROBAEI^{2,3,4,5}, THOMAS KOBERDA⁶  AND
CORENTIN LE COZ⁷

¹*Department of Geometry and Topology, University of Seville, Seville, Spain*

²*Departments of Mathematics and Computer Science, Queens College, CUNY, New York, USA*

³*Department of Computer Science, University of York, York, UK*

⁴*Department of Computer Science and Engineering, New York University, Tandon School of Engineering, New York, USA*

⁵*The Initiative for the Theoretical Sciences, CUNY Graduate Center, New York, USA*

⁶*Department of Mathematics, University of Virginia, Charlottesville, VA, USA*

⁷*Department of Mathematics: Algebra and Geometry (WE01), Universiteit Gent, Ghent, Belgium*

Corresponding author: Ramon Flores, email: ramonjflores@us.es

(Received 16 November 2023)

Abstract Let Γ be a finite graph and let $A(\Gamma)$ be the corresponding right-angled Artin group. From an arbitrary basis \mathcal{B} of $H^1(A(\Gamma), \mathbb{F})$ over an arbitrary field, we construct a natural graph $\Gamma_{\mathcal{B}}$ from the cup product, called the *cohomology basis graph*. We show that $\Gamma_{\mathcal{B}}$ always contains Γ as a subgraph. This provides an effective way to reconstruct the defining graph Γ from the cohomology of $A(\Gamma)$, to characterize the planarity of the defining graph from the algebra of $A(\Gamma)$ and to recover many other natural graph-theoretic invariants. We also investigate the behaviour of the cohomology basis graph under passage to elementary subminors and show that it is not well-behaved under edge contraction.

Keywords: determinant; cohomology; minor; right-angled Artin group; planar graph; outerplanar graph

2020 Mathematics subject classification: 15A15; 20F36; 05C10; 05C83

1. Introduction

This paper forms part of a program to study the relationship between the algebraic structure of right-angled Artin groups and the combinatorial structure of graphs, and specifically how one can extract combinatorial properties of a graph Γ from an abstract group G which is isomorphic to $A(\Gamma)$. The methods of this paper investigate the interplay between group theory, linear algebra, algebraic topology, combinatorics and commutative algebra which arise in the study of graphs and right-angled Artin groups. The reader is directed to [30] for background and commentary.



Translational tools between combinatorial properties and algebraic properties are interesting from a purely theoretical point of view, and they also arise in applied contexts such as group based cryptography (cf. [16], for instance). Computational tractability motivates the present work to a high degree.

Good characterizations of many combinatorial properties of graphs via the algebraic structure of right-angled Artin groups have been obtained by various authors, for instance, being a non-trivial join [35], being disconnected [3], containing a square [25, 27], being a co-graph [27, 28], being a finite tree or complete bipartite graph [24], admitting a non-trivial automorphism [16], being k -colourable [17], fitting in a sequence of expanders [19], admitting a Hamiltonian path or cycle [18] and being (outer)planar [20]. In this paper, we present a new perspective on characterizing (outer)planarity of the underlying graph, through the cohomology of the right-angled Artin group, which makes use of the properties of the Colin de Verdière invariant [11] and permits to describe other graph properties in terms of groups, as for example being a linear forest or being linklessly embeddable in \mathbb{R}^3 . Moreover, our method also characterizes when the complement of the graph has some of these properties. The key and more difficult point of our argument is Theorem 1.1, which establishes a certain embedding of graphs.

Let us now give a more precise setup. Let Γ be a finite *simplicial* graph. That is, Γ is undirected, and its geometric realization is a one-dimensional simplicial complex. Such graphs are also sometimes called *simple*. This paper focusses on the problem of extracting combinatorial information about Γ from the associated right-angled Artin group

$$A(\Gamma) = \langle V(\Gamma) \mid [v, w] = 1, \{v, w\} \in E(\Gamma) \rangle.$$

Let \mathbb{F} be an arbitrary field. We consider the cohomology ring $H^*(A(\Gamma), \mathbb{F})$. It is well-known that $H^*(A(\Gamma), \mathbb{F})$ can be recovered from an arbitrary presentation of $A(\Gamma)$, see § 2.5. Moreover, $A(\Gamma)$ is *1-formal*, meaning that $H^*(A(\Gamma), \mathbb{F})$, and in particular the restriction of the cup product to the first degree, completely determines Γ . The fact that the right-angled Artin group determines the underlying graph up to isomorphism is obtained as the main result of [14], and a related rigidity result was established by [34]. An alternative proof that the cohomology ring of a right-angled Artin group determines the underlying graph is given as Theorem 6.4 in [29], cf. Theorem 15.2.6 of [30].

1.1. The cohomology basis graph and the defining graph

We will be interested in effective ways of reconstructing Γ from $A(\Gamma)$, especially through $H^*(A(\Gamma), \mathbb{F})$. For this, let \mathcal{B} be an arbitrary basis of $H^1(A(\Gamma), \mathbb{F})$. The *cohomology basis graph* associated with $\Gamma_{\mathcal{B}}$ is the graph whose vertices are elements of \mathcal{B} , and whose edge relation is given by having a non-trivial cup product.

The main result of this paper is the following. We give an algebraic topology–combinatorics version, though there are many other equivalent formulations, see § 3.2. Here and in what follows, we do not require *subgraphs* to be full.

Theorem 1.1. *Let Γ be a finite simplicial graph and let \mathcal{B} be an arbitrary basis for $H^1(A(\Gamma), \mathbb{F})$. Then Γ is a subgraph of $\Gamma_{\mathcal{B}}$.*

The importance of this result for us is two-fold: first, it is the cornerstone on which almost all results in the present paper are based and permits the use of minor-monotonicity of the Colin de Verdière invariant to provide group-theoretic characterizations of many graph properties, cf. Theorem 1.5. Second, it has equivalent formulations in other contexts (such as in commutative algebra), wherein the corresponding results were previously unknown.

In the course of the proof, we develop a novel perspective on the computation of the determinant of an invertible matrix. This method greatly generalizes the approach used in [18] and highlights the role of certain graphs that arise naturally from the structure of the minors of the matrix, cf. §3. This is where the main difficulty in establishing Theorem 1.1 lies: the main result implicitly finds a bijection between the vertices of Γ and the basis \mathcal{B} which has good algebraic properties, though in general no canonical bijection exists.

Since it is easy to see that Γ and $\Gamma_{\mathcal{B}}$ have the same number of vertices, Theorem 1.1 says that Γ can be obtained from $\Gamma_{\mathcal{B}}$ by deleting edges:

Corollary 1.2. *For a finite simplicial graph Γ , we have that $\Gamma \cong \Gamma_{\mathcal{B}}$ for any basis \mathcal{B} of $H^1(A(\Gamma), \mathbb{F})$ which minimizes the number of edges in $\Gamma_{\mathcal{B}}$.*

From Corollary 1.2, one can give an *a priori* bound on the complexity of reconstructing Γ from $H^*(A(\Gamma), \mathbb{F})$, since one can apply the corollary to a field with two elements, over which there are only finitely many bases.

Recall that to every finite graph Γ , we may associate the *Colin de Verdière invariant* $\mu(\Gamma)$, which is a natural number that characterizes disconnected graphs, forests, outerplanar graphs, planar graphs and many other classes of graphs. The reader may find the definition and basic properties of $\mu(\Gamma)$ in §2.3.

Theorem 1.3. *For a finite simplicial graph Γ and natural number k , we have that $\mu(\Gamma) \leq k$ if and only if there exists a basis \mathcal{B} of $H^1(A(\Gamma), \mathbb{F})$ such that $\mu(\Gamma_{\mathcal{B}}) \leq k$.*

From the computability of the cohomology ring (cf. §2.5), we have the following consequence:

Corollary 1.4. *From an arbitrary finitely presented group $G = \langle S \mid R \rangle$ such that $G \cong A(\Gamma)$ for some finite simplicial graph Γ , the value of $\mu(\Gamma)$ is computable from $\langle S \mid R \rangle$.*

Recall that a graph is *planar* if its geometric realization can be embedded in the plane, and *outerplanar* if it can be embedded in the plane in such a way that every vertex is adjacent to the unbounded component of the complement. Moreover, a graph is *linklessly embeddable* in \mathbb{R}^3 if there is an embedding of the graph in \mathbb{R}^3 such that no pair of cycles are linked after being embedded; observe that this property can be thought of as a three-dimensional analogue of planarity.

Obtaining the following consequence was another motivation for carrying out the present work.

Theorem 1.5. *Let P be a property of graphs that is characterized by excluding a class of forbidden subgraphs. A finite simplicial graph Γ has property P if and only if there exists a basis \mathcal{B} of $H^1(A(\Gamma), \mathbb{F})$ such that $\Gamma_{\mathcal{B}}$ has property P .*

Moreover, let P one of the following graph properties:

- *Emptiness (having no edges).*
- *Being a linear forest (union of disjoint paths).*
- *Planarity.*
- *Outerplanarity.*
- *Linkless embeddability.*

Then a finite simplicial graph Γ has property P if and only if there exists a basis \mathcal{B} of $H^1(A(\Gamma), \mathbb{F})$ such that $\Gamma_{\mathcal{B}}$ has property P .

One must be careful in generalizing Theorem 1.5 to forbidden minors, since the cohomology basis graph does not behave well under taking minors of Γ , see §5. Theorem 1.5 behaves well for properties which are monotone with respect to taking minors, such as having Colin de Verdière invariant bounded by a fixed integer, see [10] for a discussion of graphs characterized by forbidden subgraphs.

Analogous characterizations of some of the properties enumerated in Theorem 1.5 for the complement of a given graph are also possible, see Proposition 3.6.

As was mentioned already, the formulation of the previous theorem makes no reference to any distinguished set of generators of the group $A(\Gamma)$. Moreover, information about graph properties can be effectively obtained out of any presentation of the associated right-angled Artin group, via the basis cohomology graph associated to that presentation and using \mathbb{F}_2 -coefficients (see Example 4.3 and the discussion at the end of §2.5). This last observation contrasts with the recent results of Gheorghiu in [20], at least from the computational point of view. Indeed, Gheorghiu finds an intrinsic characterization of right-angled Artin groups on planar graphs for instance, but it is not clear whether his methods are effective. In this vein, we note that Theorem 1.5 also furnishes a linear algebraic characterization of planarity of finite simple graphs, in the spirit of Maclane (cf. [33]).

The full extent of information which can be gleaned from the cohomology basis graph is not yet clear.

Question 1.6. *Let Γ and $A(\Gamma)$ be as above.*

- (1) *What is the effect of different fields on the structure of the cohomology basis graph? Note that depending on the characteristic, different edges may appear or be deleted.*
- (2) *What further data about the defining graph can be extracted from the cohomology basis graph? For instance, how can the cycles of Γ be read off?*
- (3) *If one considers all possible bases for the cohomology of $A(\Gamma)$, one obtains a partially ordered set under inclusion of subgraphs. Which subgraphs between Γ and the complete graph on the vertices of Γ occur? To what extent does the answer depend on Γ ?*

The paper is structured as follows: in §2, we provide background about right-angled Artin groups and the Colin de Verdière invariant, as well as different aspects of the

cohomology of right-angled Artin groups that are relevant to our analysis. Section 3 contains the proofs of the main results. In § 4, we describe some specific examples which illustrate the difficulties in establishing the main result. We conclude with § 5, where we investigate the relationship between the minors of the defining graph and of the cohomology basis graphs in more detail.

2. Background

2.1. Notation and terminology

We follow generally accepted conventions and notation for graphs, see [13], for instance. Of particular interest will be *graph minors*. A graph Λ is an *elementary minor* of Γ if Λ is obtained from Γ by deleting a vertex, deleting an edge or contracting an edge. We say that Λ is a *minor* of Γ if there is a sequence $\{\Gamma_0, \dots, \Gamma_n\}$ of graphs such that

$$\Gamma = \Gamma_0, \quad \Gamma_n = \Lambda, \quad \Gamma_{i+1} \text{ is an elementary minor of } \Gamma_i \text{ for all } i.$$

We will adopt some slightly non-standard linear algebra terminology. For a matrix A , we will write entries a_i^j , where i indicates the row and j indicates the column. We will write $\{a_1, \dots, a_n\}$ for the rows of a matrix. A *minor* of A is simply a (possibly empty) square submatrix of A obtained by deleting some (possibly empty) collection of rows and columns of A . The *dimension* of such a minor is just the number of rows or columns in the minor.

We will write S_n for the symmetric group on n letters and σ for an arbitrary element of S_n .

2.2. Cohomology of right-angled Artin groups

We recall some basic facts about the structure of the cohomology algebra of a right-angled Artin group $A(\Gamma)$. The result recorded here is easy and well-known and follows from standard methods in geometric topology together with the fact that the Salvetti complex associated with Γ is a classifying space for $A(\Gamma)$. More details can be found in [19, 30], for instance.

Let $V(\Gamma) = \{v_1, \dots, v_n\}$ and $E(\Gamma) = \{e_1, \dots, e_m\}$ be the vertices and edges of Γ and write \smile for the cup product pairing on $H^*(A(\Gamma), \mathbb{F})$.

Lemma 2.1. *Let Γ be a finite simplicial graph. Then there are bases $\{v_1^*, \dots, v_n^*\}$ for $H^1(A(\Gamma), \mathbb{F})$ and $\{e_1^*, \dots, e_m^*\}$ for $H^2(A(\Gamma), \mathbb{F})$ such that:*

- (1) *We have $v_i^* \smile v_j^* = 0$ if and only if $\{v_i, v_j\} \notin E(\Gamma)$;*
- (2) *We have $v_i^* \smile v_j^* = \pm e_k^*$ whenever $\{v_i, v_j\} = e_k \in E(\Gamma)$.*

Let

$$w_i = \sum_{j=1}^n a_i^j v_j^*$$

for $i = 1, 2$, and for coefficients a_i^j in a field. From Lemma 2.1, we observe that $w_1 \smile w_2 \neq 0$ if and only if there is a pair of indices j and k such that $v_j^* \smile v_k^* \neq 0$ and the matrix

$$\begin{pmatrix} a_1^j & a_1^k \\ a_2^j & a_2^k \end{pmatrix}$$

is non-singular. This fact will be used implicitly throughout the rest of this paper.

2.3. The Colin de Verdière invariant

The Colin de Verdière invariant of a graph is an invariant arising from spectral graph theory, which gives a vast generalization of classical planarity criteria for graphs. General references on the Colin de Verdière invariant are [11] and [36]. We give a brief summary of the definition and main properties of this invariant for the convenience of the reader, which can be found in the aforementioned references.

We represent a finite simple connected graph Γ by its vertex set $V = \{1, \dots, n\}$ and its edge set E . We consider symmetric real $n \times n$ matrices M such that the following three conditions hold:

- For all distinct indices $1 \leq i, j \leq n$, we have $M_i^j < 0$ if $\{i, j\} \in E$ and $M_i^j = 0$ otherwise;
- M has exactly one negative eigenvalue of multiplicity one;
- There is no non-zero symmetric real $n \times n$ matrix X such that $MX = 0$ and such that $X_i^j = 0$ whenever $i = j$ or $M_i^j \neq 0$.

The Colin de Verdière invariant $\mu(\Gamma)$ is the largest corank of any M satisfying these conditions. Here, for a symmetric matrix $m \times m$ of rank r , the *corank* is equal to $m - r$.

Although the original definition of the invariant assumes the graph to be connected, it is easy to extend the definition to disconnected non-empty graphs by taking the maximum of the value of the invariant on the components; for empty graphs, the invariant is defined to be zero, see [31].

The following result illustrates the power of this invariant:

Theorem 2.2. *Let Γ be a finite simple graph such that $\mu(\Gamma) \leq k$. Then:*

- $k = 0$ if and only if Γ has no edges.
- $k = 1$ if and only if Γ is a union of disjoint paths.
- $k = 2$ if and only if Γ is outerplanar.
- $k = 3$ if and only if Γ is planar.
- $k = 4$ if and only if Γ is linklessly embeddable in \mathbb{R}^3 .

It is also possible to describe the planarity properties of the complement of the graph in terms of this invariant. Recall that two vertices in a graph are *twins* if they have the same set of neighbours.

Theorem 2.3. *Let Γ be a finite simple graph of n vertices with no twin vertices, and such that $\mu(\Gamma) \geq n - k$. Then:*

- $k=5$ if and only if the complement of Γ is planar.
- $k=4$ if and only if the complement of Γ is outerplanar.

The previous result can be refined further. For example, the ‘only if’ implications are true without assuming any conditions on the vertices. Moreover, there are also (weaker) characterizations of the complement being a linear forest or linklessly embeddable in \mathbb{R}^3 , also using the Colin de Verdière invariant, see [31].

2.4. Right-angled Artin groups and formality

As was noted already, the isomorphism type of $A(\Gamma)$ determines the isomorphism type of the defining graph Γ . In [29], a proof was given that passed through the cohomology rings of right-angled Artin groups, cf. [30]. That is, the cohomology of the right-angled Artin group up to dimension 2, together with the cup product pairing, determines the isomorphism type of the underlying graph Γ . This fact fits into a much broader theory of formality, in particular, 1-formality. This is a phenomenon closely related to the de Rham fundamental group, quadratic presentability of the Mal’cev algebra and the minimal 1-model, see [2]. For general Artin groups, 1-formality is a consequence of the work of Kapovich and Millson [26]. Categorical perspectives on right-angled Artin groups and their defining graphs are investigated by Grossack [22].

2.5. Effectiveness, automaticity and computation

For a general finitely presented group, one can algorithmically recover the cohomology ring structure in degree 1 (i.e. products of elements in degree 1), using a standard five-term exact sequence from group cohomology (arising from the Lyndon–Hochschild–Serre spectral sequence, cf. [6]). Explicitly, if G is a finitely generated group and H is its abelianization, then one can compute the kernel of the cup product map

$$\bigwedge^2 H \longrightarrow H^2(G)$$

via the lower central series. In particular, it can be decided which products in \bigwedge^2 are trivial or not. Thus, from any generating set for the first cohomology of a right-angled Artin group, one can recover the structure of the associated cohomology basis graph without appealing to the duals of Artin generators. If we consider coefficients in \mathbb{F}_2 and generating sets whose cardinality is the rank of the group, this method and Corollary 1.2 provide an effective way of reconstructing Γ from the data of $H^*(A(\Gamma), \mathbb{F})$, and assuming all queries take a unit amount of time, one can construct a naive algorithm to compute Γ whose complexity is $O(e^{|\mathcal{V}(\Gamma)|^2})$. The complexity of recovering Γ from an arbitrary finite presentation of $A(\Gamma)$ is a somewhat different matter.

We remark that, in general, computing the second cohomology of a finitely presented group is difficult. Indeed, one can bootstrap the unsolvability of the word problem in

general finitely presented groups to prove that it is generally undecidable whether or not $H^2(G, \mathbb{Z}) = 0$, cf. [21].

For a right-angled Artin group, these pathologies do not occur. Indeed, right-angled Artin groups are *biautomatic*, cf. [5, 7–9, 12]. The exact definition of this property is irrelevant for our purposes, and we direct the reader to the seminal text [15].

The property of automatic or biautomatic does not depend on the underlying presentation, though passing between different automatic structures can be mysterious. For a finitely presented group which is known to be biautomatic, practically finding a biautomatic structure is not always clear. From a given biautomatic structure on a group G , it is a theorem of Bridson–Reeves [4] that there is an algorithm to construct a finite dimensional approximation to a classifying space for G (i.e. a finite dimensional skeleton of BG). Thus, for an arbitrary finitely presented group which is abstractly isomorphic to a right-angled Artin group, there is an algorithm which computes all of $H^*(G, \mathbb{Z})$.

Observe that any presentation of a right-angled Artin group defines a basis for the cohomology, and then the arguments above allow us to directly compute the cohomology basis graph. Any bound on the Colin de Verdière invariant for this graph immediately gives the same bound for the defining graph. In particular, we obtain a planarity test for the defining graph whose input is any presentation of the right-angled Artin group, see Example 4.3.

3. Γ -null-connectedness and the proof of the main results

The ideas we use to establish the main result expand and generalize the constructions developed by the first three authors in [18], which in turn are partially inspired by classical expansions of the determinant relying on the computation of 2×2 minors, such as Laplace expansion, the Dodgson condensation formula and the Sylvester formula [1].

3.1. Null-connectedness

For the remainder of this section, we fix a finite simple graph Γ , with vertices $\{1, \dots, n\}$, as well as $A \in M_n(\mathbb{F})$. The indices in the labelling of columns of A will be identified with the vertices of Γ .

Recall that we write row vectors of A as $a_r = (a_r^1, \dots, a_r^n)$. We say two rows a_r and a_s of A are Γ -null-connected if the minor

$$\begin{pmatrix} a_r^i & a_r^j \\ a_s^i & a_s^j \end{pmatrix}$$

is singular whenever $\{i, j\}$ is an edge of Γ .

A submatrix M of A will be called a Γ -1-block if the following conditions are satisfied.

- (1) M has at least two rows and two columns.
- (2) All entries of M are non-zero.
- (3) The indices of the columns occurring in M span a connected subgraph of Γ .

- (4) The rows of A which meet M span a connected graph with the Γ -null-connectedness adjacency relation.
- (5) M is maximal with respect to these conditions, in the sense that there is no submatrix N of A which properly contains M and which satisfies the previous conditions.

A Γ -1-minor is a minor of A with at least two rows, which is contained in a Γ -1-block.

Lemma 3.1. *Let M be a Γ -1-block in A . Then the row space of M is one-dimensional.*

Proof. Clearly the row space of M is at least one-dimensional. Let $\{i, j\}$ be indices of columns in M that span an edge of Γ , and let a_1 and a_2 be null-connected rows of A . We have that the matrix

$$\begin{pmatrix} a_1^i & a_1^j \\ a_2^i & a_2^j \end{pmatrix}$$

is singular and has only non-zero entries, whereby it follows that the vector (a_2^i, a_2^j) is a non-zero scalar multiple $\lambda_{i,j}$ of (a_1^i, a_1^j) . By definition, the indices of the columns meeting M span a connected subgraph $\Lambda \subset \Gamma$. For every edge of Λ spanned by vertices $\{s, t\}$, we obtain a non-zero scalar $\lambda_{s,t}$ relating (a_2^s, a_2^t) and (a_1^s, a_1^t) . Moreover, if two edges of Λ share a vertex then the two corresponding scalars must coincide. It follows by induction on the diameter of the subgraph Λ and from the fact that all entries in M are non-zero that the scalars $\lambda_{s,t}$ depend only on a_1 and a_2 and not on the indices s and t of the columns. It follows that the two rows a_1 and a_2 of M are scalar multiples of each other. Since the rows of M span a connected graph under the null-connectivity relation, we see that any two rows of M are scalar multiples of each other, the desired conclusion. \square

The following property of Γ -1-blocks is crucial for canonical sorting of summands making up the determinant. Again, a similar statement and proof are found as Lemma 2.10 in [18].

Lemma 3.2. *Let M_1 and M_2 be two distinct Γ -1-blocks of A . Then M_1 and M_2 are disjoint as submatrices of A .*

Proof. Let $I \subset \{1, \dots, n\}$ be the set of columns meeting M_1 and $J \subset \{1, \dots, n\}$ be the set of columns meeting M_2 , and let $\{\ell_1, \dots, \ell_s\}$ be the indices of the rows in M_2 . Write Λ_I and Λ_J for the corresponding connected subgraphs of Γ . Suppose $a_{\ell_1}^k$ is an entry appearing in both M_1 and M_2 . We will show that in this case $M_1 = M_2$.

By definition, for all $i \in I$, we have that $a_{\ell_1}^i$ is a non-zero entry of M_1 , and that $a_{\ell_m}^k$ is a non-zero entry of M_2 . By Lemma 3.1, the row spaces of both M_1 and M_2 are one-dimensional. We have $k \in I \cap J$ and a_{ℓ_1} is a row meeting both M_1 and M_2 .

If a_{ℓ_m} is another row meeting M_2 then a_{ℓ_1} and a_{ℓ_m} lie in the same Γ -null-connected component of the rows of A . Suppose first that they are actually Γ -null-connected. Then we have that the matrix

$$\begin{pmatrix} a_{\ell_1}^k & a_{\ell_1}^i \\ a_{\ell_m}^k & a_{\ell_m}^i \end{pmatrix}$$

is necessarily singular and consists of all non-zero entries since $a_{\ell_m}^k \neq 0$. By the same argument as in Lemma 3.1 (using the connectivity of Λ_I), the rows $a_{\ell_1}^I$ and $a_{\ell_m}^I$, consisting of entries in the columns indexed by I , are non-zero scalar multiples of each other. Since the rows of M_2 span a connected graph under the Γ -null-connectivity relation, we obtain that $a_{\ell_1}^I$ and $a_{\ell_m}^I$ are non-zero scalar multiples of each other for all $1 \leq m \leq s$. By the maximality condition on M_2 and the connectivity of Λ_I and Λ_J , we have that $I \subseteq J$. By symmetry, $I = J$. By the maximality conditions on M_1 and M_2 , we obtain that $M_1 = M_2$, the desired conclusion. \square

Lemma 3.2 allows us to canonically partition the entries of a matrix into three different types:

- (1) Non-zero entries that lie in a Γ -1-block.
- (2) Non-zero entries that do not lie in a Γ -1-block.
- (3) Zero entries.

Next, we need to generalize to this context the notion of Γ -1-track defined in [18]. By definition, this is a sequence $\{A_1, \dots, A_k\}$ of minors of A with the following properties.

- (1) For each i , the minor A_i is either a 1×1 submatrix or a Γ -1-minor.
- (2) Each column of A meets exactly one A_i .
- (3) For $i \neq j$, it is not the case that A_i and A_j belong to a common Γ -1-minor.

Two Γ -1-tracks are said to be *different* if they consist of different sets of submatrices of A .

The preceding definitions serve to sort the summands that make up the determinant of the matrix A , the latter of which is simply viewed as a signed combination of products of matrix entries. We write $\mathbf{a} = (a_{\sigma(1)}^1, \dots, a_{\sigma(n)}^n)$ for an arbitrary string of non-zero entries of A . Such a string *belongs* to a Γ -1-track $\{A_1, \dots, A_k\}$ if for all i there exists a j such that $a_{\sigma(i)}^i$ is an entry of A_j .

Lemma 3.3. *Let $\mathbf{a} = (a_{\sigma(1)}^1, \dots, a_{\sigma(n)}^n)$ be a string of non-zero entries of A . Then \mathbf{a} belongs to a unique Γ -1-track in A .*

Proof. For B a Γ -1-block, we let $\{b_1, \dots, b_s\}$ be the (possibly empty) set of entries of \mathbf{a} which lie in B . Each row and each column of A contain exactly one entry of \mathbf{a} , and so $\{b_1, \dots, b_s\}$ defines a Γ -1-minor A_B in B of dimension exactly s . By construction, for distinct Γ -1-blocks B_1 and B_2 , the Γ -1-minors A_{B_1} and A_{B_2} are disjoint. The remaining entries of \mathbf{a} , say $\{c_1, \dots, c_t\}$, are 1×1 non-zero matrices that belong to no Γ -1-block. Thus, the Γ -1-track associated with \mathbf{a} is

$$\{A_B\}_{B \in \mathcal{B}} \cup \{c_1, \dots, c_t\},$$

where \mathcal{B} ranges over all Γ -1-blocks of A . The disjointness of distinct Γ -1-blocks guarantees that this is actually a Γ -1-track.

It is clear that this Γ -1-track in A is unique. Indeed, the constituents $\{c_1, \dots, c_t\}$ and $\{A_B\}_{B \in \mathcal{B}}$ are canonically defined and hence unique. \square

Recall the Leibniz formula for the determinant of an $n \times n$ matrix, namely

$$\det A = \sum_{\sigma \in S_n} \text{sgn}(\sigma) \prod_{i=1}^n a_{\sigma(i)}^i.$$

For a given Γ -1-track \mathcal{T} of A , we write $(\det A)_{\mathcal{T}}$ for the restriction of the sum defining the determinant to permutations σ such that $\mathbf{a} = (a_{\sigma(1)}^1, \dots, a_{\sigma(n)}^n)$ belongs to \mathcal{T} .

Lemma 3.4. *Let \mathcal{T} be a Γ -1-track of A . Suppose \mathcal{T} contains a minor of dimension at least two. Then*

$$(\det A)_{\mathcal{T}} = 0.$$

Proof. Let M be such a minor. Without loss of generality, M consists of the top left $k \times k$ minor of A , and we may identify S_k with the group of permutations of the rows of M , and we extend these permutations by the identity. If \mathbf{a} belongs to \mathcal{T} and $\tau \in S_k$, then so does the string \mathbf{a}^{τ} , which is obtained by applying τ to the row indices of \mathbf{a} . Since the signature of a permutation is a homomorphism, we have that the contribution of \mathbf{a}^{τ} to $(\det A)_{\mathcal{T}}$ differs from that of \mathbf{a} by $\text{sgn}(\tau)$. Since the row space of a Γ -1-block is one-dimensional, we have that the product of entries of \mathbf{a} and \mathbf{a}^{τ} are equal. It is now immediate that $(\det A)_{\mathcal{T}} = 0$, since exactly half the permutations in S_k have signature 1 and half have signature -1 . \square

In the case of \mathbb{F}_2 , Lemma 3.4 could be proved by simply noting that if \mathcal{T} contains a minor of dimension at least two then an even number of distinct strings \mathbf{a} belongs to \mathcal{T} .

Corollary 3.5. *Suppose A is invertible. Then there exists a reordering of the rows of A such that for all edges $\{i, j\}$ of Γ , we have that a_i and a_j are not Γ -null-connected.*

Proof. We suppose the contrary and argue that $\det A = 0$. Let $\mathbf{a} = (a_{\sigma(1)}^1, \dots, a_{\sigma(n)}^n)$ be a string of non-zero entries of A . By assumption, there is an edge $\{i, j\}$ of Γ which witnesses the fact that σ fails to be a suitable reordering. By reordering the rows by σ^{-1} , we may assume $\mathbf{a} = (a_1^1, \dots, a_n^n)$. We have that the matrix

$$M = \begin{pmatrix} a_i^i & a_j^i \\ a_i^j & a_j^j \end{pmatrix}$$

is singular, because these rows must be Γ -null-connected. Since the matrix M has non-zero diagonal entries and is singular, it must lie inside of a Γ -1-block of A . It follows that

the unique Γ -1-track \mathcal{T} to which \mathbf{a} belongs contains a Γ -1-minor of dimension at least two, and so $(\det A)_{\mathcal{T}} = 0$. The choice of \mathbf{a} was arbitrary, and so each such string belongs to a Γ -1-track \mathcal{T} such that $(\det A)_{\mathcal{T}} = 0$. Now, the uniqueness of the track to which \mathbf{a} belongs implies that $\det A$ is the sum of $(\det A)_{\mathcal{T}}$, where \mathcal{T} varies over all possible Γ -1-tracks. The desired result now follows. \square

Proof of Theorem 1.1. We fix notation and write $\{v_1^*, \dots, v_n^*\}$ for the standard basis for $H^1(A(\Gamma), \mathbb{F})$, and we let $A \in \text{GL}_n(\mathbb{F})$ be arbitrary. We let $\mathcal{B} = \{w_1, \dots, w_n\}$ be the result of applying A , viewed as a change of basis, so that

$$w_i = \sum_{j=1}^n a_i^j v_j^*.$$

We let σ be a reordering of the rows of A as guaranteed by Corollary 3.5, and we relabel the vectors $\{w_1, \dots, w_n\}$ according to σ . Computing $w_i \smile w_j$, we see that this cup product is zero if and only if the rows a_i and a_j of A are Γ -null-connected. It follows that the cohomology basis graph $\Gamma_{\mathcal{B}}$ contains Γ as a subgraph, as desired. \square

Observe that Corollary 1.2 is immediate from Theorem 1.1, as for an inclusion of graphs $\Gamma' \subseteq \Gamma$ with the same number of vertices, equality in the number of edges implies isomorphism. Theorem 1.3 is a consequence of Theorem 1.1 and the minor-monotonicity of the Colin de Verdière invariant. As right-angled groups are biautomatic, Corollary 1.4 follows. Finally, Theorem 1.5 is implied in turn by Theorem 1.3, together with the properties of the invariant discussed in § 2.3.

We conclude by stating a property concerning complements, which is analogous to Theorem 1.5:

Proposition 3.6. *Let Γ be a finite simple graph with no twin vertices. Then the complement of Γ is planar (respectively outerplanar) if and only if for every basis \mathcal{B} of $H^1(A(\Gamma), \mathbb{F})$, the complement of the cohomology basis graph $\Gamma_{\mathcal{B}}$ is planar (respectively outerplanar).*

Proof. We prove the case of planarity; the argument for outerplanarity is analogous. By Theorem 2.3, the complement of Γ is planar if and only if $\mu(\Gamma) \geq n - 5$, where $n = |V(\Gamma)|$. By the minor monotonicity of μ , we see that $\mu(\Gamma_{\mathcal{B}}) \geq n - 5$ for every cohomology basis graph $\Gamma_{\mathcal{B}}$. By Theorem 2.3, the complement of every such $\Gamma_{\mathcal{B}}$ is planar. The other implication is immediate. \square

3.2. A reformulation in terms of edge ideals

There are equivalent reformulations of the main Theorem 1.1 in other contexts, which to our knowledge were also open questions; this was communicated to the authors by Van Tuyl [38]. Here we discuss a perspective from commutative algebra and in the next section from graph theory.

One fruitful context for investigating the interplay between combinatorics of graphs and commutative algebra is through *chutters* and *monomial ideals*, and especially *edge*

ideals, which are in turn related to the theory of polyhedral products and Stanley–Reisner rings; the reader is directed to [23, 32] for background.

We consider $k = \mathbb{F}_2$, the field with two elements, and a polynomial ring $R = k[x_1, \dots, x_n]$. We let $I \subset R$ be an ideal generated by square-free monomials of degree 2, and we let Γ_I be the graph having I as its edge ideal. A matrix $A \in \text{GL}_n(k)$ determines a linear change of variables $x_i \mapsto w_i$ that preserves degrees of polynomials. We let I' be the complement of I , which is to say that I' is generated by all degree 2 monomials that do not lie in I , and we let $\bar{R} = R/I'$.

One can now consider the ideal $J \subset k[w_1, \dots, w_n]$ generated by products $w_i w_j$, with $i \neq j$, and its image $\bar{J} \subset \bar{R}$. We write Γ_J for the graph with vertices $\{w_1, \dots, w_n\}$, and with an edge whenever the monomial $w_i w_j$ survives in \bar{J} . These graphs have been extensively studied in the literature, see, for example, [37] and references therein.

The following result can easily be seen to follow from Theorem 1.1:

Corollary 3.7. *The graph Γ_I is a subgraph of Γ_J .*

Conversely, observe that every finite simple graph can be represented by Γ_I for some edge ideal. Moreover, every change of basis can be effected by an invertible matrix. Thus, if Γ_I is always a subgraph of Γ_J as above, then Theorem 1.1 holds over a field of characteristic 2.

3.3. A reformulation in classical graph theory

Let $A \in \text{GL}_n(\mathbb{F}_2)$, and let Γ be a fixed graph on n vertices $\{v_1, \dots, v_n\}$. We now define a new finite simple graph with vertex set $\{e_1, \dots, e_n, w_1, \dots, w_n\}$. The edge relation is given as follows:

- (1) We place an edge between e_i and e_j precisely when there is an edge between v_i and v_j .
- (2) We place an edge between w_i and e_j precisely when the entry a_i^j of A is non-zero.
- (3) For $i \neq j$, we place a new edge between w_i and w_j if and only if there exist indices k and ℓ such that the following conditions hold:
 - (a) There is an edge between v_k and v_ℓ .
 - (b) There is a path of length 3 between w_i and w_j that contains the edge between e_k and e_ℓ induced by the previous condition and the two edges arising from (2) above, and there is no path of length two in the subgraph induced by $\{w_i, w_j, e_k, e_\ell\}$.

We write Γ' for the graph spanned by $\{w_1, \dots, w_n\}$, and let \mathcal{B} be the basis for $H^1(A(\Gamma), \mathbb{F}_2)$ induced by the rows of A ; by construction, we may naturally identify elements of \mathcal{B} with $\{w_1, \dots, w_n\}$. It is straightforward to see that w_i and w_j are adjacent in Γ' if and only if w_i and w_j are adjacent in $\Gamma_{\mathcal{B}}$: indeed, this can be seen from explicitly writing out the cup product in terms of the basis of $H^1(A(\Gamma), \mathbb{F}_2)$ coming from the vertices of Γ , cf. § 2.2. Thus, Theorem 1.1 is equivalent to:

Corollary 3.8. *The graph Γ is a subgraph of Γ' .*

4. Some examples

In this section, we describe explicit examples that we find illustrative for understanding the difficulties that arise in attempts to prove the main result directly.

One of the basic issues that makes finding a proof of Theorem 1.1 non-trivial is the ‘global’ nature of the assertion it makes. The result says that from an arbitrary basis for $H^1(A(\Gamma), \mathbb{F})$, one can find an assignment between these arbitrary basis vectors and vectors in the standard basis which respects the cup product structure. Experiments suggest and careful consideration shows that there is no canonical way to realize such a bijection, and this is the reason that inductive strategies do not seem to work; even if one assumes the existence of such a bijection for a proper subgraph, extending by even one vertex seems technically impossible. These issues already appear in the following simple example:

Example 4.1. Consider the defining graph Γ with vertices $\{v_1, v_2, v_3, v_4\}$ and edges

$$\{v_1, v_2\}, \{v_1, v_3\}, \{v_1, v_4\},$$

and the basis $\mathcal{B} = \{w_1, w_2, w_3, w_4\}$ for the first cohomology of $A(\Gamma)$ given by the change of basis matrix

$$A = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

Now, the associated cohomology basis graph $\Gamma_{\mathcal{B}}$ has edges

$$\{w_1, w_2\}, \{w_1, w_3\}, \{w_1, w_4\},$$

which is isomorphic to Γ . Hence, the natural assignment $v_i \mapsto w_i$ induces an isomorphism of graphs, and in particular an inclusion $\Gamma \subseteq \Gamma_{\mathcal{B}}$.

Now, we add the edge $\{v_2, v_4\}$ to the defining graph, obtaining a new graph Λ , and we retain the basis \mathcal{B} . The cohomology basis graph $\Lambda_{\mathcal{B}}$ has edges

$$\{w_1, w_2\}, \{w_1, w_3\}, \{w_1, w_4\}, \{w_2, w_3\},$$

and the previous assignment $v_i \mapsto w_i$ do not extend to the new graphs. Of course, Theorem 1.1 ensures that there exists another suitable assignment for the new graphs, for instance one sending v_3 to w_4 and v_4 to w_3 , but this is an *ad hoc* modification that is difficult to make canonical.

In searching for a formula defining a canonical bijection between the vertices of the graph and the arbitrary basis vectors, one encounters many reasonable-sounding linear algebraic claims which end up being false. One might hope, for example, that if $A \in GL_n(\mathbb{F}_2)$ then there is a reordering of the rows of A so that all the principal minors

are non-singular; one might hope for this to be the case for just the two-dimensional principal minors, and it is not difficult to see how this claim would imply Theorem 1.1. A counterexample to the claim is given by the following:

Example 4.2. Consider the invertible matrix

$$A = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 1 \end{pmatrix}.$$

Then, a straightforward analysis shows that in every reordering of the rows there is at least one 2-dimensional principal minor which is singular.

One further example that we give in this section concerns the practicality of the planarity test described at the end of §2.5.

Example 4.3. Consider the group

$$G = \langle x_1, x_2, x_3, x_4, x_5 \mid x_1x_2x_5x_4 = x_2x_5x_4x_1, x_3x_2x_5x_4 = x_2x_5x_4x_3, x_4x_5 = x_5x_4 \rangle.$$

This groups is abstractly isomorphic to a right-angled Artin group on a graph with five vertices. If we consider the dual generators

$$\{x_1^*, x_2^*, x_3^*, x_4^*, x_5^*\}$$

in $H^1(G; \mathbb{F}_2)$, it is not hard to check using the arguments of §2.5 that the only non-trivial products of two of these elements are $x_2^*x_i^*$ for every $i \neq 2$ and $x_4^*x_5^*$. Hence, the cohomology graph is a star with one additional edge, which is planar. Thus, the defining graph is also planar, by Theorem 1.3.

5. Graph minors and the cohomology basis graph

In this last section, we investigate the behaviour of the cohomology graph $\Gamma_{\mathcal{B}}$ under taking minors of Γ . This is particularly relevant in light of Wagner’s theorem (i.e. a graph is planar if and only if it does not admit K_5 or $K_{3,3}$ as a minor), though we will show that the cohomology basis graph and graph minors do not interact in a sufficiently nice way to make this a viable approach to characterizing planarity. Thus, we have another justification for Theorem 1.1 being the ‘correct’ approach to understanding planarity through cohomological means.

Notice that if Λ is an elementary minor of Γ obtained by deleting an edge or vertex then there is a natural inclusion of Λ into Γ . If Λ is an elementary minor under contracting an edge with vertices v_1 and v_2 , then Λ is equipped with a special vertex which we can formally and suggestively label $v_1 + v_2$.

Let \mathcal{B} be an arbitrary basis for $H^1(A(\Gamma), \mathbb{F})$. We begin by writing the elements of \mathcal{B} in terms of the cohomology classes that are dual to vertices of Γ , i.e. $\{w_1^*, \dots, w_n^*\}$. If Λ is an

elementary minor of Γ , then there is a natural way to obtain a basis \mathcal{B}' for $H^1(A(\Lambda), \mathbb{F})$. Then, we have the following moves:

- (1) If Λ is obtained from Γ by deleting an edge of Γ then $\mathcal{B}' = \mathcal{B}$;
- (2) If Λ is obtained from Γ by deleting the vertex corresponding to the cohomology class v_i^* , then the vectors of \mathcal{B}' are simply the vectors of \mathcal{B} with every occurrence of v_i^* deleted and any repeats or trivial vectors discarded and, if necessary, one further vector discarded to ensure that \mathcal{B}' is linearly independent;
- (3) If Λ is obtained from Γ by contracting the edge connecting vertices corresponding to v_i^* and v_j^* , then the vectors of \mathcal{B}' are vectors of \mathcal{B} with v_i^* and v_j^* replaced by $v_i^* + v_j^*$, and with any repeats or trivial vectors discarded and, if necessary, one further vector discarded to ensure that \mathcal{B}' is linearly independent.

To fix terminology, if \mathcal{B} is an arbitrary basis for $H^1(A(\Gamma), \mathbb{F})$, we will call a basis \mathcal{B}' for $H^1(A(\Lambda), \mathbb{F})$ obtained by one of the previous three moves an *elementary minor basis*.

We first make some remarks about the naturality of the transformation from \mathcal{B} to \mathcal{B}' . For edge deletions, there is little to say. If Λ is obtained from Γ by deleting an edge between vertices v_1 and v_2 , then there is a natural homomorphism $\phi: A(\Lambda) \rightarrow A(\Gamma)$ which is defined by simply declaring that v_1 and v_2 commute with each other. In this case, we get a natural map

$$\phi^*: H^1(A(\Gamma), \mathbb{F}) \rightarrow H^1(A(\Lambda), \mathbb{F}),$$

and $\phi^*(\mathcal{B}) = \mathcal{B}'$.

The case of vertex deletion is similar, using the fact that the inclusion $\Lambda \rightarrow \Gamma$ induces a map $\psi: A(\Lambda) \rightarrow A(\Gamma)$, and consequently a map

$$\psi^*: H^1(A(\Gamma), \mathbb{F}) \rightarrow H^1(A(\Lambda), \mathbb{F}).$$

We then consider $\psi^*(\mathcal{B})$, which contains a basis of $H^1(A(\Lambda), \mathbb{F})$, and so we discard a vector if necessary to obtain \mathcal{B}' .

Though \mathcal{B}' is obtained from \mathcal{B} in a reasonably natural way, there is still a matter of the choice of which vector to discard. Proposition 5.1 shows that the choice made is absorbed by a choice of minor on the cohomology basis graph side.

The case of edge contraction is slightly different, since there is a natural map

$$\chi: A(\Gamma) \rightarrow A(\Lambda),$$

though not the other way. Let v_1 and v_2 be two adjacent vertices of Γ that are identified in Λ , wherein we call the resulting vertex v_0 . Write $\{v_1^*, v_2^*, v_0^*\}$ for the corresponding dual cohomology classes. Under the natural induced map $H^1(A(\Lambda), \mathbb{F}) \rightarrow H^1(A(\Gamma), \mathbb{F})$, the cohomology class v_0^* is sent to the cohomology class $v_1^* + v_2^*$. Thus, $H^1(A(\Lambda), \mathbb{F})$ is canonically isomorphic to a subspace of $H^1(A(\Gamma), \mathbb{F})$ with the identification $v_0^* \mapsto v_1^* + v_2^*$. We let

$$\varphi: H^1(A(\Gamma), \mathbb{F}) \rightarrow H^1(A(\Gamma), \mathbb{F})$$

be the endomorphism which is the identity on all generators dual to vertices other than v_1 and v_2 , and otherwise $v_1^*, v_2^* \mapsto v_1^* + v_2^*$. Finally, we let

$$\chi_*: H^1(\varphi(A(\Gamma), \mathbb{F})) \rightarrow H^1(A(\Lambda), \mathbb{F})$$

be the map that is the identity on all generators dual to vertices other than v_1 and v_2 , which sends $v_1^* + v_2^* \mapsto v_0^*$. Then, $\chi_* \circ \varphi(\mathcal{B})$ contains a basis for $H^1(A(\Lambda), \mathbb{F})$, and so we discard a vector if necessary to obtain \mathcal{B}' .

Proposition 5.1. *Let Γ be a graph, let Λ an elementary minor of Γ obtained by either edge deletion or vertex deletion, let \mathcal{B} be an arbitrary basis for $H^1(A(\Gamma), \mathbb{F})$ and let \mathcal{B}' be an elementary minor basis for $H^1(A(\Lambda), \mathbb{F})$. Then the graph $\Gamma_{\mathcal{B}'}$ is a minor of the graph $\Gamma_{\mathcal{B}}$.*

Proof. We verify the claim for the two moves that we allow.

Edge deletion. In this case, $\mathcal{B} = \mathcal{B}'$. Suppose $b_1, b_2 \in \mathcal{B}$ are adjacent to each other in $\Gamma_{\mathcal{B}}$, and write these basis vectors as a sum of cohomology classes dual to the vertices of Γ . Abusing notation slightly, there is an edge $\{v_1, v_2\}$ of Γ such that v_1 but not v_2 occurs in the expression of b_1 , and v_2 but not v_1 occurs in the expression of b_2 . If the edge $\{v_1, v_2\}$ persists in Λ then b_1 and b_2 remain adjacent in $\Gamma_{\mathcal{B}'}$. If the edge $\{v_1, v_2\}$ does not persist in Λ , then b_1 and b_2 may or may not remain adjacent in $\Gamma_{\mathcal{B}'}$, contingent on the existence of another edge of Γ that witnesses the continued adjacency of b_1 and b_2 .

If b_1 and b_2 are non-adjacent in $\Gamma_{\mathcal{B}}$, then we wish to argue that these vertices remain non-adjacent in $\Gamma_{\mathcal{B}'}$. Let $\{v_1, v_2\}$ denote an arbitrary edge of Γ . Necessarily, one of the following three possibilities holds, up to relabelling vertices or basis elements:

- (1) Neither v_1 nor v_2 occurs in the expressions for b_1 and b_2 ;
- (2) Both v_1 and v_2 occur in both expressions for b_1 and b_2 ;
- (3) The class v_1 occurs in the expression for b_1 but v_2 does not occur in the expression for b_2 .

Now, there is a pair of vertices $\{v_i, v_j\}$ which span an edge of Γ but such that $v_i^* \smile v_j^* = 0$ in $H^1(A(\Lambda))$, with no other cup products between dual vertex basis vectors being changed. It follows immediately then that pairs of non-adjacent vertices in $\Gamma_{\mathcal{B}}$ remain non-adjacent in $\Gamma_{\mathcal{B}'}$.

Vertex deletion. Retaining notation from above, let $\psi^*(\mathcal{B}) \subset H^1(A(\Lambda))$ be the image of \mathcal{B} under the map induced by the inclusion $\Lambda \rightarrow \Gamma$, and let $b_1, b_2 \in \mathcal{B}$. For $i \in \{1, 2\}$, writing $\psi^*(b_i)$ and b_i in terms of the vertex duals, we simply have that a summand v is deleted from $\psi^*(b_i)$ if it occurs in b_i . If $\psi^*(b_1) = \psi^*(b_2)$ then we will begin by deleting one of them (which is vertex deletion in $\Gamma_{\mathcal{B}}$) and then proceed by applying ψ^* , which will then yield \mathcal{B}' without any further deletions.

Suppose that b_1 and b_2 are adjacent in $\Gamma_{\mathcal{B}}$, and that this adjacency is witnessed only by edges in Γ that are incident to v . Then after deleting v from Γ , all these edges are severed, in which case $\psi^*(b_1) \smile \psi^*(b_2) = 0$. If the adjacency is witnessed by an edge that is not incident to v , then $\psi^*(b_1) \smile \psi^*(b_2) \neq 0$.




Figure 1. The dumbbell graph Γ .

Suppose that b_1 and b_2 are non-adjacent in $\Gamma_{\mathcal{B}}$. Then for an arbitrary edge $\{v_1, v_2\}$ of Γ , we have the three possibilities as in the case of edge deletion. It is straightforward to check that the three possibilities persist after applying ψ^* , in which case $\psi^*(b_1) \smile \psi^*(b_2) = 0$.

Let Γ_{ψ^*} be the graph obtained by taking vertices to be elements $\psi^*(\mathcal{B})$ and adjacency to be given by non-vanishing cup product. Then the preceding argument shows that Γ_{ψ^*} is a minor of $\Gamma_{\mathcal{B}}$. The basis \mathcal{B}' is obtained by (possibly) deleting an element of $\psi^*(\mathcal{B})$, in which case we see that $\Gamma_{\mathcal{B}'}$ is a minor of Γ_{ψ^*} , as desired. \square

It is not generally true that if Λ is obtained from Γ by edge contraction then $\Gamma_{\mathcal{B}'}$ is a minor of $\Gamma_{\mathcal{B}}$. Consider for instance the dumbbell graph Γ from Figure 1. It has $n + m + 1$ edges.

For compactness of notation, we will conflate names of vertices and corresponding dual cohomology classes and cease writing asterisk superscripts. To simplify notation further, we will write sums of cohomology classes multiplicatively. Let

$$\mathcal{B} = \{b_1, \dots, b_m, a_1 b_1, v_1 a_1, \dots, v_1 a_n, v_2 a_1\}.$$

It is straightforward to verify that this is indeed a basis for $H^1(A(\Gamma), \mathbb{F})$. The edges of $\Gamma_{\mathcal{B}}$ are of the following form:

- (1) $\{b_i, v_2 a_1\}$ for $1 \leq i \leq m$;
- (2) $\{v_2 a_1, v_1 a_j\}$ for $1 \leq j \leq n$;
- (3) $\{v_2 a_1, a_1 b_1\}$;
- (4) $\{v_1 a_i, v_1 a_j\}$ for $i \neq j$;
- (5) $\{a_1 b_1, v_1 a_j\}$ for $1 \leq j \leq n$.

We check easily that $\Gamma_{\mathcal{B}}$ has

$$\frac{n^2}{2} + \frac{3n}{2} + m + 1$$

total edges. Now to compute \mathcal{B}' and $\Gamma_{\mathcal{B}'}$, we introduce the symbol z for $v_1 + v_2$. Each occurrence of v_1 and v_2 is replaced by z . The vertices $v_2 a_1$ and $v_1 a_1$ become identical, and so we delete one of them. Then we have

$$\mathcal{B}' = \{b_1, \dots, b_m, a_1 b_1, z a_1, \dots, z a_n\}.$$

The edges of $\Gamma_{\mathcal{B}'}$ are of the following form:

- (1) $\{b_i, za_j\}$ for $1 \leq i \leq m$ and $1 \leq j \leq n$;
- (2) $\{za_i, za_j\}$ for $i \neq j$;
- (3) $\{a_1 b_1, za_j\}$ for $1 \leq j \leq n$.

We check easily that $\Gamma_{\mathcal{B}'}$ has

$$\frac{n^2}{2} + mn + \frac{n}{2}$$

total edges. Thus, the difference between the total number of edges of $\Gamma_{\mathcal{B}'}$ and $\Gamma_{\mathcal{B}}$ is $(m-1)(n-1) - 2$. Evidently, this difference can be made positive by choosing the parameters n and m suitably. Now, if $\Gamma_{\mathcal{B}'}$ were a minor of $\Gamma_{\mathcal{B}}$ then $\Gamma_{\mathcal{B}'}$ would have fewer edges than $\Gamma_{\mathcal{B}}$, which is a contradiction.

Acknowledgements. We thank Y. Antolín, O. Cárdenas, M. Casals-Ruiz, M. Chudnovsky, J. Gálvez, M. Gheorghiu, P. Giménez, J. González-Meneses, C. Johnson, A. Márquez, J. Martínez, S. Morey, F. Muro, C. Valencia and A. Viruel for helpful comments. We are in particular very grateful to K. Li, M. Tsatsomeros and A. van Tuyl for their interest in our work. The authors are grateful to an anonymous referee who provided numerous helpful comments which greatly improved the paper.

Funding Statement. Ramón Flores is supported by FEDER-MEC grant PID2020-117971GB-C21 of the Spanish Ministry of Science. Thomas Koberda is partially supported by NSF Grant DMS-2002596. Delaram Kahrobaei is supported in part by an, ONR grant, 62909-24-1-2002. Corentin Le Coz is supported by the FWO and the F.R.S.–FNRS under the Excellence of Science (EOS) program (project ID 40007542).

References

- (1) F. F. Abeles, Chiò's and Dodgson's determinantal identities, *Linear Algebra Appl.* **454** (2014), 130–137. MR3208413
- (2) J. Amorós, M. Burger, K. Corlette, D. Kotschick and D. Toledo, *Fundamental groups of compact Kähler manifolds*, Mathematical Surveys and Monographs, Volume 44 (American Mathematical Society, Providence, RI, 1996). MR1379330
- (3) N. Brady and J. Meier, Connectivity at infinity for right angled Artin groups, *Trans. Amer. Math. Soc.* **353**(1) (2001), 117–132. MR1675166
- (4) M. R. Bridson and L. Reeves, On the algorithmic construction of classifying spaces and the isomorphism problem for biautomatic groups, *Sci. China Math.* **54**(8) (2011), 1533–1545. MR2824957
- (5) E. Brieskorn and K. Saito, Artin-Gruppen und Coxeter-Gruppen, *Invent. Math.* **17** (1972), 245–271. MR323910
- (6) K. S. Brown, *Cohomology of groups*, Graduate Texts in Mathematics, Volume 87 (Springer-Verlag, New York, 1994). Corrected reprint of the 1982 original. MR1324339
- (7) R. Charney, Artin groups of finite type are biautomatic, *Math. Ann.* **292**(4) (1992), 671–683. MR1157320

- (8) R. Charney, Geodesic automation and growth functions for Artin groups of finite type, *Math. Ann.* **301**(2) (1995), 307–324. MR1314589
- (9) R. Charney, An introduction to right-angled Artin groups, *Geom. Dedicata* **125** (2007), 141–158. MR2322545
- (10) G. Chartrand, D. Geller and S. Hedetniemi, Graphs with forbidden subgraphs, *J. Combin. Theory Ser. B* **10** (1971), 12–41, doi:10.1016/0095-8956(71)90065-7
- (11) Y. C. de Verdière, Sur un nouvel invariant des graphes et un critère de planarité, *J. Combin. Theory Ser. B* **50**(1) (1990), 11–21. MR1070462
- (12) P. Deligne, Les immeubles des groupes de tresses généralisés, *Invent. Math.* **17** (1972), 273–302. MR422673
- (13) R. Diestel, *Graph theory*, 5th edition, Graduate Texts in Mathematics, Volume 173 (Springer, Berlin, 2017). MR3644391
- (14) C. Droms, Isomorphisms of graph groups, *Proc. Amer. Math. Soc.* **100**(3) (1987), 407–408. MR891135
- (15) D. B. A. Epstein, J. W. Cannon, D. F. Holt, S. V. F. Levy, M. S. Paterson and W. P. Thurston, *Word processing in groups* (Jones and Bartlett Publishers, Boston, MA, 1992). MR1161694
- (16) R. Flores, D. Kahrobaei and T. Koberda, Algorithmic problems in right-angled Artin groups: complexity and applications, *J. Algebra* **519** (2019), 111–129. MR3874519
- (17) R. Flores, D. Kahrobaei and T. Koberda, An algebraic characterization of k -colorability, *Proc. Amer. Math. Soc.* **149**(5) (2021), 2249–2255. MR4232214
- (18) R. Flores, D. Kahrobaei and T. Koberda, Hamiltonicity via cohomology of right-angled Artin groups, *Linear Algebra Appl.* **631** (2021), 94–110. MR4308807
- (19) R. Flores, D. Kahrobaei, and T. Koberda, Expanders and right-angled Artin groups, *J. Topol. Anal.* **16**(2) (2024), 155–179, doi:10.1142/S179352532150059X. To appear.
- (20) M. Gheorghiu, Loose ear decompositions and their applications to right-angled artin groups, arXiv:2307.08459, 2023.
- (21) C. Gordon, Some embedding theorems and undecidability questions for groups, *Combinatorial and geometric group theory (Edinburgh, 1993)*, London Mathematical Society Lecture Note Series, Volume 204, pp. 105–110 (Cambridge University Press, Cambridge, 1995). MR1320278
- (22) C. Grossack, The right angled artin group functor as a categorical embedding, arXiv:2309.06614, 2024.
- (23) H. Hà and A. Van Tuyl, Monomial ideals, edge ideals of hypergraphs, and their graded Betti numbers, *J. Algebraic Combin.* **27** (2008), 215–245. MR2375493
- (24) S. Hermiller and Z. Šunić, Poly-free constructions for right-angled artin groups, *J. Group Theory* **10** (2007), 117–138. MR2288463
- (25) M. Kambites, On commuting elements and embeddings of graph groups and monoids, *Proc. Edinb. Math. Soc. (2)* **52**(1) (2009), 155–170. MR2475886
- (26) M. Kapovich and J. J. Millson, On representation varieties of Artin groups, projective arrangements and the fundamental groups of smooth complex algebraic varieties, *Inst. Hautes Études Sci. Publ. Math.* **88** (1998), 5–95. MR1733326
- (27) S.-H. Kim and T. Koberda, Embedability between right-angled Artin groups, *Geom. Topol.* **17**(1) (2013), 493–530. MR3039768
- (28) S.-H. Kim and T. Koberda, Free products and the algebraic structure of diffeomorphism groups, *J. Topol.* **11**(4) (2018), 1054–1076. MR3989437
- (29) T. Koberda, Right-angled Artin groups and a generalized isomorphism problem for finitely generated subgroups of mapping class groups, *Geom. Funct. Anal.* **22**(6) (2012), 1541–1590. MR3000498

- (30) T. Koberda, Geometry and combinatorics via right-angled Artin groups, *In the tradition of Thurston II, Geometry and groups*, pp. 475–518 (Springer, Cham, 2022). MR4472060
- (31) A. Kotlov, L. Lovász and S. Vempala, The Colin de Verdière number and sphere representations of a graph, *Combinatorica* **17**(4) (1997), 483–521.
- (32) S. Morey and R. H. Villarreal, *Edge ideals: algebraic and combinatorial properties*, Progress in commutative algebra, Volume 1, pp. 85–126 (de Gruyter, Berlin, 2012). MR2932582
- (33) P. O’Neil, A short proof of Mac Lane’s planarity theorem, *Proc. Amer. Math. Soc.* **37** (1973), 617–618. MR313098
- (34) L. Sabalka, On rigidity and the isomorphism problem for tree braid groups, *Groups Geom. Dyn.* **3**(3) (2009), 469–523. MR2516176 (2010g:20074)
- (35) H. Servatius, Automorphisms of graph groups, *J. Algebra* **126**(1) (1989), 34–60. MR1023285 (90m:20043)
- (36) H. van der Holst, L. Lovász and A. Schrijver, *The Colin de Verdière graph parameter*, Graph theory and combinatorial biology (Balatonlelle, 1996), Bolyai Society Mathematical Studies, Volume 7, pp. 29–85 (János Bolyai Mathematical Society, Budapest, 1999). MR1673503
- (37) A. Van Tuyl, *A beginner’s guide to edge and cover ideals*, Monomial ideals, computations and applications, Lecture Notes in Mathematics, Volume 2083, pp. 63–94 (Springer, Heidelberg, 2013). MR3184120
- (38) A. Van Tuyl, Personal communication, 2023.