Chapter 15
Geometry and Combinatorics via Right-Angled Artin Groups

Thomas Koberda

Abstract We survey the relationship between the combinatorics and geometry of graphs and the algebraic structure of right-angled Artin groups. We concentrate on the defining graph of the right-angled Artin group and on the extension graph associated to the right-angled Artin group. Additionally, we discuss connections with geometric group theory and complexity theory.

Keywords Right-angled Artin group · Extension graph · Graph expanders · Hamiltonian graph · k-Colorability · Graph automorphism · Acylindrical group action · Quasi-isometry · Commensurability · Mapping class group · Curve graph

2020 Mathematics Subject Classification Primary: 20F36, 20F65, 05C50; Secondary: 05C45, 05C48, 05C60, 68Q15, 03D15

15.1 Introduction

In this paper, we survey the interplay between the algebraic structure of right-angled Artin groups, the combinatorics of graphs, and geometry. Throughout the paper, let Γ be a finite simplicial graph, and we write $V(\Gamma)$ and $E(\Gamma)$ for the set of vertices and edges of Γ, respectively. The *right-angled Artin group* on Γ, denoted by $A(\Gamma)$, is the group defined by

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

T. Koberda (✉)
Department of Mathematics, University of Virginia, Charlottesville, VA, USA
https://sites.google.com/view/koberdat

K. Ohshika, A. Papadopoulos (eds.), *In the Tradition of Thurston II*,
https://doi.org/10.1007/978-3-030-97560-9_15

15.1.1 Scope of This Survey

Right-angled Artin groups interpolate between free groups and abelian groups, and they exhibit a wide range of complex phenomena. Moreover, they are simple enough that their structure is relatively tractable, and hence one can come to understand these groups fairly well. They are prototypical examples of CAT(0) groups, and they serve as toy examples that mirror many important properties of and inform conjectures about more complicated groups, such as mapping class groups. Some well-known and difficult conjectures about mapping class groups, such as the characterization of convex cocompact subgroups, admit complete, tractable analogues in the case of right-angled Artin groups; see [88] for a detailed discussion.

In this article, we will concentrate on some specific aspects of right-angled Artin groups, which we will outline in the remainder of this section. For a survey of the general properties of right-angled Artin groups, the reader is directed to [36].

Some of the basic questions we will discuss are as follows.

Question 15.1.1 *What is the exact relationship between the group theoretic structure of the group $A(\Gamma)$ and the combinatorial structure of Γ?*

The reader will find that there are two answers to Question 15.1.1, the trivial one and the nontrivial one. The trivial one will be a consequence of Theorem 15.2.6 below, which shows that Γ is completely determined by the cohomology algebra of $A(\Gamma)$, and in fact by the degree one and two parts together with the cup product pairing. Thus, one can in principle recover Γ from $A(\Gamma)$, so that any combinatorial properties of Γ is automatically determined by the algebraic structure of $A(\Gamma)$. Conversely, the algebraic structure of $A(\Gamma)$ is, in a sense that is so general as to render it almost meaningless, "known" by the graph Γ.

There is a more interesting approach to Question 15.1.1 that seeks to find a dictionary between the combinatorics of Γ and the algebra of $A(\Gamma)$, by passing between specific graph-theoretic and group-theoretic properties that are analogous. This line of inquiry yields some otherwise nonobvious insights that have applications outside of geometric group theory, such as in cryptography and complexity theory. Some sample results we will discuss in the sequel are the following:

Theorem 15.1.2 ([51]) *Let Γ be a finite simplicial graph. Then Γ admits a nontrivial automorphism if and only if the outer automorphism group $\mathrm{Out}(A(\Gamma))$ contains a finite nonabelian group.*

Theorem 15.1.3 ([54]) *Let Γ be a finite simplicial graph with n vertices. Then Γ admits a k-coloring if and only if $A(\Gamma)$ surjects to a product*

$$F_{n_1} \times \cdots \times F_{n_k},$$

where F_{n_i} is a free group of rank n_i, and

$$\sum_{i=1}^{k} n_i = n.$$

Theorem 15.1.4 ([55]) *Let Γ be a finite simplicial graph. Then Γ admits a Hamiltonian cycle if and only if the cohomology algebra of $A(\Gamma)$ is a Hamiltonian vector space.*

In Theorem 15.1.4, *Hamiltonicity* of a vector space means that there is a bilinear form satisfying certain "connectivity" conditions. We direct the reader to Sect. 15.3.3 for precise definitions.

Whereas the graph Γ is evidently intimately related to the structure of $A(\Gamma)$, the graph Γ is not always ideally suited for the study of the internal structure of $A(\Gamma)$, since there is no natural interesting action of $A(\Gamma)$ on Γ. However, one can augment Γ in the "smallest way possible" in order to get a graph on which $A(\Gamma)$ acts. The key idea is to conflate a vertex v of Γ with an element of $A(\Gamma)$. One can then consider the set

$$V(\Gamma^e) = \{v^g \mid v \in V(\Gamma), g \in A(\Gamma)\} \subset A(\Gamma)$$

of all conjugates of vertices of Γ, where here we write $v^g = g^{-1}vg$. It is true though largely irrelevant that $V(\Gamma^e)$ is not canonically defined as a subset of $A(\Gamma)$, since automorphisms of $A(\Gamma)$ need not preserve the set of conjugates of given vertex generators of $A(\Gamma)$.

We build a graph Γ^e, called the *extension graph* of Γ (cf. [78]), by putting an edge $\{v^g, w^h\}$ between vertices in $V(\Gamma^e)$ whenever $[v^g, w^h] = 1$ in $A(\Gamma)$. The group $A(\Gamma)$ now acts in a canonical way on Γ^e, i.e. by conjugation.

Question 15.1.5 *What is the relationship between the structure of $A(\Gamma)$ and the structure of Γ^e? What is the geometry of the action of $A(\Gamma)$ on Γ^e?*

The graph Γ^e, though algebraically defined, is very closely related to Hagen's *contact graph* [63], which encodes the intersection pattern between hyperplanes in a natural CAT(0) cube complex on which $A(\Gamma)$ acts. This, together with an analogy between the extension graph and the curve graph associated to a hyperbolic surface of finite type, is an entry point into the theory of hierarchically hyperbolic spaces (HHSs) and hierarchically hyperbolic groups (HHGs) (see [15, 16], for instance). We will largely avoid discussing that aspect of the theory in this paper.

The extension graph carries a large amount of data about the subgroup structure of $A(\Gamma)$. A sample result we will discuss is the following:

Theorem 15.1.6 ([78]) *Suppose Γ has no triangles, and let Λ be an arbitrary finite simplicial graph. Then $A(\Lambda)$ occurs as a subgroup of $A(\Gamma)$ if and only if Λ occurs as a subgraph of Γ^e.*

The action of $A(\Gamma)$ on Γ^e by conjugation, though perhaps simple at first glance, serves to unify the group theory of $A(\Gamma)$, the geometry of Γ^e, and the intrinsic CAT(0) geometry of $A(\Gamma)$. We state the following result that we will discuss in some detail, and we will defer definitions of the terminology until then.

Theorem 15.1.7 ([79]) *The action of $A(\Gamma)$ on Γ^e is acylindrical. Moreover, the following are equivalent.*

(1) The element $g \in A(\Gamma)$ acts loxodromically on Γ^e.
(2) The element $g \in A(\Gamma)$ acts as a rank one isometry of the universal cover of the Salvetti complex of Γ.
(3) The element $g \in A(\Gamma)$ is not conjugate into a join subgroup of $A(\Gamma)$.

We will not give detailed proofs of most of the results in this survey. We will give proof sketches where it is feasible, and we will strive to give complete references. As already suggested above, we will omit large parts of the theory and neglect various viewpoints. The specific topics discussed herein undoubtedly reflect the idiosyncratic tastes of the author.

15.1.2 Notation and Terminology

Most of the notation and terminology used in this survey is standard or nearly standard. All graphs will be undirected and simplicial unless otherwise noted, so that in particular there are no double edges nor edges that start and end at a single vertex. The *complement* of a graph Γ is the complement of Γ in the complete graph on the vertices of Γ; that is, complete all the missing edges of Γ and then delete the edges that were present in Γ. Two vertices are therefore connected by an edge in the complement of Γ if and only if they are not connected by an edge in Γ.

A graph Γ is a *join* if its complement graph is disconnected. The join of graphs Γ_1 and Γ_2 is written $\Gamma_1 * \Gamma_2$, and every vertex of Γ_1 is adjacent to every vertex of Γ_2. The join of two graphs mimics the geometric join in topology: if A and B are topological spaces, then the join $A * B$ is the quotient of $A \times B \times I$ that collapses $A \times B \times \{0\}$ to A and $A \times B \times \{1\}$ to B. For us, a subgraph Λ of a graph Γ is always *full*, which is to say Λ contains all edges that are present in Γ. A *clique* is a complete graph, and a *k-clique* is a complete graph on k vertices. The set $V(\Gamma)$, viewed as a subset of $A(\Gamma)$, is called the set of *vertex generators* of $A(\Gamma)$. The *link* of a vertex $v \in V(\Gamma)$ is written $\mathrm{Lk}(v)$ and consists of the vertices that are adjacent to v. If $\emptyset \neq S \subset V(\Gamma)$ then

$$\mathrm{Lk}(S) = \bigcap_{s \in S} \mathrm{Lk}(s).$$

The *star* of v is given by $\mathrm{St}(v) = \mathrm{Lk}(v) \cup \{v\}$. The *degree* of a vertex v is given by $|\mathrm{Lk}(v)|$. A vertex v is *isolated* if $\mathrm{Lk}(v)$ is empty. A graph is *totally disconnected* if every vertex is isolated. A *path* in Γ is a tuple of vertices $\overline{p} = (v_1, \ldots, v_k)$ in $V(\Gamma)$ such that $\{v_i, v_{i+1}\} \in E(\Gamma)$ for all suitable indices. The parameter k is arbitrary, and the *length* of the path $\overline{p}$ is $k - 1$. A *cycle* or *circuit* is a path for which $v_1 = v_k$ and for which $v_i \neq v_{i+2}$ for all suitable indices. A graph is *connected* if for all pairs of vertices $v, w \in V(\Gamma)$, there is a path in Γ such that $v = v_1$ and $w = v_k$.

The *rank of a linear map* is the dimension of its image, and the *rank of a group* is the minimal number of generators of the group. The identity element of a group is denoted 1 with an exception in the case of additive abelian groups when it is written 0.

Let Γ be a (possibly infinite) graph. We build a graph Γ_k, called the *clique graph* of Γ as follows. We start with the vertices and edges of Γ. For every complete subgraph $K \subset V(\Gamma)$ with at least two vertices, we add a new vertex v_K to $V(\Gamma_k)$. If K_1 and K_2 are cliques such that $V(K_1) \cup V(K_2)$ also spans a complete subgraph of Γ, then we add an edge $\{v_{K_1}, v_{K_2}\}$ to $E(\Gamma_k)$. Finally, we add an edge between each vertex of the form v_K and the vertices making up K. The resulting graph is the clique graph. It is helpful to illustrate the clique graph with an example: if Γ is a graph without triangles then the only cliques with two or more vertices are the edges of Γ. In this case, the clique graph of Γ is just Γ with an extra vertex v_e for each edge $e \in E(\Gamma)$, and two edges connecting v_e to the two vertices of Γ spanning e. Thus, Γ_k is just a copy of Γ with a "fin" hanging off each edge.

15.1.3 A Remark About Generators

When we specify a right-angled Artin group, we will write $A(\Gamma)$. Since $A(\Gamma)$ as an abstract group determines Γ up to isomorphism, the specification of Γ (viewed as an abstract graph) does not constitute a choice of generators for $A(\Gamma)$. However, once we speak of particular generators of $A(\Gamma)$, we have implicitly chosen an identification of $V(\Gamma)$ with a set of generators for $A(\Gamma)$. The author has taken pains to avoid ambiguities that could cause confusion for the reader.

15.2 The Cohomology Ring of a Right-Angled Artin Group

A central role in the dictionary between algebra and combinatorics is played by the cohomology of a right-angled Artin group. Recall that the cohomology of a group G is defined to be the cohomology of a $K(G, 1)$, which is unique up to homotopy equivalence (see [64], for instance). A right-angled Artin group has a very easy to describe $K(G, 1)$, and a large number of natural retractions allows for an efficient calculation of the cohomology algebra. For the entirety of this section, R will denote a commutative ring with a unit, unless otherwise noted.

15.2.1 The Topology of the Salvetti Complex

We will write $\mathscr{S}(\Gamma)$ for the Salvetti complex of Γ, and we construct it as follows (cf. [36]). Let Γ be a graph with $n = |V(\Gamma)|$. We fix a bijection between $V(\Gamma)$

and $\{1, \ldots, n\}$. Consider now the unit cube $[0, 1]^n \subset \mathbb{R}^n$. We build a certain subset $S \subset [0, 1]^n$, cube by cube. For $1 \leq i \leq n$, we write J_i for the unit segment in the i^{th} coordinate direction, emanating from the origin. We include J_i in S for all i. Now, if $K \subset \{1, \ldots, n\}$ consists of a collection of vertices which span a complete subgraph of Γ, then we include the subcube of $[0, 1]^n$ spanned by $\{J_i\}_{i \in K}$ in S.

Once S has been constructed in this way, we set $\mathscr{S}(\Gamma)$ to be the image of S in $\mathbb{R}^n/\mathbb{Z}^n$, where $\mathbb{Z}^n$ acts on $\mathbb{R}^n$ by usual integer translations. Thus, the complex $\mathscr{S}(\Gamma)$ is realized as a subcomplex of an n-dimensional torus.

Proposition 15.2.1 *The following are properties of* $\mathscr{S}(\Gamma)$.

1. *The fundamental group of* $\mathscr{S}(\Gamma)$ *is isomorphic to* $A(\Gamma)$.
2. *The universal cover of* $\mathscr{S}(\Gamma)$ *is contractible.*

That the fundamental group of $\mathscr{S}(\Gamma)$ is isomorphic to $A(\Gamma)$ is a straightforward calculation using Van Kampen's Theorem. That the universal cover of $\mathscr{S}(\Gamma)$ is contractible is much less obvious, and follows from the fact that $\mathscr{S}(\Gamma)$ admits the structure of a locally CAT(0) cube complex. To delve into the details would take us far afield, and we shall content ourselves to direct the reader to some references, such as [29, 62, 119]. The crucial point here is that the homology and cohomology of $\mathscr{S}(\Gamma)$ are in fact a invariants of $A(\Gamma)$, since $\mathscr{S}(\Gamma)$ is a $K(G, 1)$ for $G = A(\Gamma)$.

The homology of $\mathscr{S}(\Gamma)$ is easily calculated by a standard Mayer–Vietoris argument. In our construction of $\mathscr{S}(\Gamma)$ above, we obtain a distinguished k-subtorus of $\mathscr{S}(\Gamma)$ for every k-subclique of Γ. When two such distinguished subtori (corresponding to subcliques K_1 and K_2 of Γ) meet, they meet along the distinguished subtorus corresponding to the intersection $K_1 \cap K_2$ (which is just the basepoint in case this intersection is empty). Thus, we see that:

Proposition 15.2.2 *Let* R *be a ring. Then* $H_k(A(\Gamma), R) \cong R^{N_k}$, *where* N_k *denotes the number of* k*-cliques in* Γ.

Here and throughout, we always assume that the $A(\Gamma)$ action on the ring of coefficients is trivial, so that our homology and cohomology groups are always untwisted. Computation of the twisted groups is much more complicated; cf. [40, 73]. In particular, the rank of the abelianization of $A(\Gamma)$ is the number of vertices of Γ, and the dimension of the second homology coincides with the number of edges.

The cohomology groups of $A(\Gamma)$ have the same ranks as the homology groups, and the formal structure of $A(\Gamma)$ (or of $\mathscr{S}(\Gamma)$) allows one to give a satisfactory description of the cohomology algebra of $A(\Gamma)$. For this, we let T_k denote the k-dimensional torus. As is standard, the cohomology algebra of T_k with coefficients in R is $\bigwedge(R^k)$, the exterior algebra of R^k.

Proposition 15.2.3 *Let* $\Lambda \subset \Gamma$ *be a subgraph. Then the map* $A(\Gamma) \longrightarrow A(\Lambda)$ *defined by the identity for vertices* $\lambda \in V(\Lambda)$ *and by* $v \mapsto 1$ *otherwise is a retraction of groups.*

Of course, the fact that Salvetti complexes are classifying spaces for right-angled Artin groups means that Proposition 15.2.3 admits a dual statement for spaces. That

is, there are natural retractions $\mathscr{S}(\Gamma) \longrightarrow \mathscr{S}(\Lambda)$ which induce the corresponding maps on fundamental groups whenever $\Lambda \subset \Gamma$ is a subgraph.

Specializing to the case where $K \subset \Gamma$ is a k-clique, we get a natural surjective map

$$A(\Gamma) \longrightarrow A(K) \cong \mathbb{Z}^k,$$

and thus an induced injective map on cohomology $\bigwedge(R^k) \longrightarrow H^*(A(\Gamma), R)$. Suppose we have a decomposition of graphs $\Gamma = \Lambda_1 \cup \Lambda_2$ and $\Theta = \Lambda_1 \cap \Lambda_2$. For technical reasons, we suppose that every edge between Λ_1 and Λ_2 is realized by an edge between Λ_i and Θ for $i \in \{1, 2\}$. We obtain a natural commutative diagram of retractions.

$$\begin{array}{ccc} A(\Gamma) & \longrightarrow & A(\Lambda_1) \\ \downarrow & & \downarrow \\ A(\Lambda_2) & \longrightarrow & A(\Theta) \end{array}$$

Replacing the retractions by inclusions of groups, $A(\Gamma)$ acquires the structure of a graph of groups with vertex groups $A(\Lambda_1)$ and $A(\Lambda_2)$ and edge group $A(\Theta)$ (cf. [105]). Without the assumption that every edge between Λ_1 and Λ_2 be realized by Θ, this previous assertion would no longer be true.

Dualizing, we get a commutative diagram on cohomology.

$$\begin{array}{ccc} H^*(A(\Gamma), R) & \longleftarrow & H^*(A(\Lambda_1), R) \\ \uparrow & & \uparrow \\ H^*(A(\Lambda_2), R) & \longleftarrow & H^*(A(\Theta), R) \end{array}$$

In category theory language, $H^*(A(\Gamma), R)$ is the pushout of the corresponding diagram. Again, the technical hypothesis on the decomposition of Γ is hidden in this last assertion, since the assertion follows from the Mayer–Vietoris sequence and would be false without this hypothesis (cf. for example when Γ is a complete graph and Λ_1 and Λ_2 are both proper subgraphs).

These considerations show that one can describe the cohomology algebra of $A(\Gamma)$ entirely in terms of exterior algebras by inductively building up Γ from its cliques. In particular, one can take an appropriate exterior algebra for each maximal clique in Γ, and identify exterior subalgebras corresponding to intersections of maximal cliques. The simplest cliques are the 1-cliques, and a retraction

$$A(\Gamma) \longrightarrow \langle v \rangle \cong \mathbb{Z}$$

for $v \in V(\Gamma)$ allows us to identify preferred generators $\{v^* \mid v \in V(\Gamma)\}$ for $H^1(A(\Gamma), R)$, which we will refer to as the *dual 1-classes* to the vertex generators.

These dual 1-classes can be interpreted as dual to certain natural subspaces of $\mathscr{S}(\Gamma)$, though we will not require this point of view here.

We clearly have that $H^*(A(\Gamma), R)$ is generated by its degree one part. Now let $v, w \in V(\Gamma)$ and let v^* and w^* be the corresponding dual 1-classes. There is a retraction $A(\Gamma) \longrightarrow \langle v, w\rangle$, and the target group is either $\mathbb{Z}^2$ or F_2, corresponding to the cases where $\{v, w\} \in E(\Gamma)$ and where $\{v, w\} \notin E(\Gamma)$ respectively. In the first case, the cup product $v^* \smile w^*$ is nontrivial and in the second case, the cup product vanishes.

The most important consequence of the previous discussion for us in the sequel is the following, which characterizes the degree one and degree two parts of the cohomology of $A(\Gamma)$ together with the cup product pairing:

Proposition 15.2.4 *Let Γ be a finite simplicial graph with $V(\Gamma) = \{v_1, \ldots, v_n\}$ and let $E(\Gamma) = \{e_1, \ldots, e_m\}$. Then there are bases $\{v_1^*, \ldots, v_n^*\}$ and $\{e_1^*, \ldots, e_m^*\}$ for $H^1(A(\Gamma), R)$ and $H^2(A(\Gamma), R)$ respectively, such that:*

(1) $v_i^* \smile v_j^* = 0$ *if* $\{v_i, v_j\} \notin E(\Gamma)$.
(2) $v_i^* \smile v_j^* = \pm e_\ell^*$ *if* $\{v_i, v_j\} = e_\ell$.

The description of $H^1(A(\Gamma), R)$ and $H^2(A(\Gamma), R)$ furnished by Proposition 15.2.4 will be essential in describing many of the correspondences between the group theoretic structure of $A(\Gamma)$ and the combinatorics of Γ.

15.2.2 Vector Spaces with a Vector-Space Valued Pairing

In the sequel, it is sometimes convenient to consider vector spaces equipped with a bilinear vector-space valued pairing. We will write $q : V \times V \longrightarrow W$ for such a pairing, where V and W are both finite dimensional vector spaces over the same field F. The pairing q is intended to generalize the cup product pairing

$$\smile : H^1(A(\Gamma), F) \times H^1(A(\Gamma), F) \longrightarrow H^2(A(\Gamma), F),$$

and so we will always adopt the assumption that q is either symmetric or anti-symmetric unless otherwise noted. This assumption on q is mostly for convenience, since relaxing some sort of symmetry assumption only adds unnecessary layers of complication that do not enrich the underlying theory in a meaningful way.

We will say that the triple (V, W, q) is *pairing-connected*, if for all nontrivial direct sum decompositions $V \cong V_0 \oplus V_1$, there are vectors $v_0 \in V_0$ and $v_1 \in V_1$ such that $q(v_0, v_1) \neq 0$. With this terminology, we can formulate and prove an entry in the algebra-combinatorics dictionary.

Proposition 15.2.5 (See [53]) *Let Γ be a finite simplicial graph, let*

$$V = H^1(A(\Gamma), F), \quad W = H^2(A(\Gamma), F),$$

and let q be the cup product pairing. Then Γ is connected if and only if (V, W, q) is pairing-connected.

Connectedness of Γ has another, simpler characterization in terms of $A(\Gamma)$, as we shall indicated below; namely, Γ is connected if and only if $A(\Gamma)$ is freely indecomposable; see Theorem 15.3.2. The (mostly complete) proof of Proposition 15.2.5 will illustrate the principle that many results that related the algebra of $A(\Gamma)$ with the combinatorics of Γ have an easy direction and a less easy direction.

Proof of Proposition 15.2.5 Suppose first that (V, W, q) is pairing-connected, and let $\Gamma = \Gamma_0 \cup \Gamma_1$ be a purported separation of Γ. Let V_i denote the span of the vertices $\{v_j^* \mid v_j \in V(\Gamma_i)\}$ for $i \in \{0, 1\}$. Pairing connectedness implies that there are vectors $w_i \in V_i$ such that $q(w_0, w_1) \neq 0$. Writing w_0 and w_1 in terms of the preferred basis vectors, we see that there are vertices $x_i \in V(\Gamma_i)$ such that $q(x_0^*, x_1^*) \neq 0$, which implies that $\{x_0, x_1\} \in E(\Gamma)$ by Proposition 15.2.4, a contradiction.

Suppose conversely that Γ is connected, and let $V \cong V_0 \oplus V_1$ be a nontrivial direct sum decomposition that witnesses the failure of (V, W, q) to be pairing-connected. Let $\{x_1, \ldots, x_m\}$ be a sequence of vertices of Γ such that every vertex of Γ appears on this list, and such that for all suitable i we have $\{x_i, x_{i+1}\} \in E(\Gamma)$. We allow this list to have repeats.

Let

$$w_0 = \sum_{i=1}^{n} \alpha_i v_i^* \in V_0, \; w_1 = \sum_{i=1}^{n} \beta_i v_i^* \in V_1$$

be expressions for nonzero vectors with respect to the standard dual basis for V. If $\{v_i, v_j\} \in E(\Gamma)$ then the expression $q(w_0, w_1) = 0$ implies that $\alpha_i \beta_j = \alpha_j \beta_i$. The two sides of this last equation are either both zero or both nonzero, and in the latter case we have that the pairs (α_i, α_j) and (β_i, β_j) are proportional. In this case, since $\{v_1^*, \ldots, v_n^*\}$ is a basis for V, we may perturb w_0 or w_1 within the respective vector spaces V_0 and V_1 in order to obtain vectors for which the coefficients corresponding to v_i^* and v_j^* are not proportional. Thus, the condition $q(w_0, w_1)$ implies that

$$\alpha_i \beta_j = \alpha_j \beta_i = 0.$$

With these observations, we can complete the proof. Let w_0 be as above. Relabeling if necessary, we have $v_1 = x_1$ and $v_2 = x_2$. Without loss of generality, we may assume that $\alpha_1 \neq 0$. Now let $w_1 \in V_1$ be expressed as above. If $\beta_2 \neq 0$ then $\alpha_1 \beta_2 \neq 0$, a conclusion that was ruled out by the considerations in the previous paragraph. Thus, $\beta_2 = 0$, and since w_1 was arbitrary, the coefficient of x_2^* vanishes for all vectors in V_1. Then, we may find a vector in V_0 whose coefficient α_2 is nonzero, and arguing symmetrically, we see that the coefficient β_1 is zero for all vectors in V_1. By induction on m and using the fact every vertex of Γ occurs on the list $\{x_1, \ldots, x_m\}$, we see that V_1 must be the zero vector space. This is a contradiction. □

15.2.3 The Cohomology Ring of $A(\Gamma)$ Determines Γ

We are now ready to state and prove a central fact about the cohomology of $A(\Gamma)$, namely that it determines the isomorphism type of Γ.

Theorem 15.2.6 *Let Γ be a finite simplicial graph, let $V = H^1(A(\Gamma), F)$, let $W = H^2(A(\Gamma), F)$, and let q be the cup product pairing. Then the triple (V, W, q) determines Γ up to isomorphism.*

One essential point in Theorem 15.2.6 is that the triple (V, W, q) is considered abstractly, without any further data such as bases. Before giving a proof of Theorem 15.2.6, we can make several observations about special instances of the result. First, the dimension of $V = H^1(A(\Gamma), F)$ coincides with $|V(\Gamma)|$, and the dimension of $W = H^2(A(\Gamma), F)$ coincides with $|E(\Gamma)|$, as is immediate from Proposition 15.2.4. Moreover, the first and second cohomology of $A(\Gamma)$ together with the cup product pairing identify complete graphs. To see this, it is convenient to introduce a map $V \longrightarrow \mathrm{Hom}(V, W)$, defined by $v \mapsto f_v$, and where $f_v(v') = v \smile v'$. The graph Γ is complete if and only if for all $v \in H^1(A(\Gamma), F)$, the rank of the image of f_v is $\dim H^1(A(\Gamma), F) - 1$. We leave the verification of this last claim as a straightforward exercise for the reader.

The fact that $A(\Gamma)$ determines the graph Γ uniquely is well-known. See [44, 85, 102] for several perspectives. The proof offered here that gives uniqueness of Γ via the cohomology algebra of $A(\Gamma)$ fits into the theory of *cohomological uniqueness*. In the context of cohomological uniqueness, one is often concerned with the question of whether or not a particular space (often decorated with adjectives such as p–completeness, where p is a prime) is determined up to homotopy equivalence by its cohomology (with various groups of coefficients). In our setting, Theorem 15.2.6 implies that among Salvetti complexes associated to finite simplicial graphs, the integral (or rational) cohomology of the space determines the space up to homotopy equivalence, and its defining graph up to isomorphism. Moreover, only the ring structure on the cohomology algebra is required, and only in degrees one and two. The reader is directed to [42, 101, 115, 116] for a more detailed discussion of cohomological uniqueness.

Another perspective on the cup product pairing determining $A(\Gamma)$ and Γ uniquely is given by the theory of 1-formality in the sense of D. Sullivan, a property which is shared notably with Kähler groups, and into which we will not delve in deeply. The reader is directed to Chapter 3 of [4], for instance.

Proof of Theorem 15.2.6 We will actually prove a stronger statement. Suppose $\Gamma \longrightarrow \Lambda$ is obtained by deleting vertices (so that Λ is a subgraph of Γ), with an induced retraction $A(\Gamma) \longrightarrow A(\Lambda)$ defined by sending the vertices $V(\Gamma) \setminus V(\Lambda)$ to the identity. Thus, we obtain triples $(V_\Gamma, W_\Gamma, q_\Gamma)$ and $(V_\Lambda, W_\Lambda, q_\Lambda)$ corresponding to the cohomologies of these groups, and a map of triples

$$i_{\Lambda,\Gamma} \colon (V_\Lambda, W_\Lambda, q_\Lambda) \longrightarrow (V_\Gamma, W_\Gamma, q_\Gamma),$$

which is injective on the level of vector spaces, and q_Λ is extended by q_Γ.

Claim The triple

$$\{(V_\Lambda, W_\Lambda, q_\Lambda), (V_\Gamma, W_\Gamma, q_\Gamma), i_{\Lambda,\Gamma}\}$$

uniquely determines graphs Λ and Γ, together with an injection of graphs $\Lambda \longrightarrow \Gamma$. The theorem will then follow from the special case where $V(\Lambda) = \varnothing$.

We proceed by induction on $(|V(\Gamma)|, |V(\Lambda)|)$, ordered lexicographically, the cases where $|V(\Gamma)| \in \{1, 2\}$ being easy consequences of the remarks preceding the proof. We now suppose the claim has been established for all graphs with at most n vertices, and we suppose that Γ has $n+1$ vertices. We consider the (possibly trivial) subspace

$$V_0 \subset V = H^1(A(\Gamma), F)$$

spanned by vectors for which f_v has rank zero. It is immediate from Proposition 15.2.4 that a vector $w \in V_0$ is in the span on vectors dual to vertices of degree zero in Γ. The quotient V/V_0 is isomorphic to $H^1(A(\Gamma'), F)$, where Γ' is the result of deleting all the vertices of Γ that have degree zero.

The natural map $A(\Gamma) \longrightarrow A(\Gamma')$ given by sending isolated vertices to the identity induces a map $H^1(A(\Gamma', F) \longrightarrow H^1(A(\Gamma), F)$, which identifies V/V_0 with a subspace of V. The cup product on $H^1(A(\Gamma), F)$ restricts to

$$q\colon H^1(A(\Gamma'), F) \times H^1(A(\Gamma'), F) \longrightarrow W.$$

Thus, if $V_0 \neq 0$ then Γ' satisfies the conclusion of the claim by induction, and Γ is obtained from Γ' by adding $\dim V_0$ many isolated vertices. We may therefore assume that Γ has no isolated vertices.

We now consider a vector $v \in V$ such that the rank of f_v is minimized.

Case 1 Suppose first that the linear span U of v coincides with the span of a vector dual to a vertex x of Γ, as furnished by Proposition 15.2.4. Then V/U coincides with the first cohomology of $A(\Gamma')$, where Γ' is obtained from Γ by deleting x. The map $V \longrightarrow V/U$ is induced by the inclusion $\Gamma' \longrightarrow \Gamma$ and the corresponding injection $A(\Gamma') \longrightarrow A(\Gamma)$.

Writing Z for the image of f_v, we have that W/Z coincides with the second cohomology of $A(\Gamma')$, and the cup product pairing descends to a bilinear map

$$\overline{q}\colon V/U \times V/U \longrightarrow W/Z,$$

which coincides with the cup product pairing on the cohomology of $A(\Gamma')$. By induction, the triple $(V/U, W/Z, \overline{q})$ determines Γ' uniquely.

Let $N \subset V$ be the kernel of f_v. Then N is spanned by the dual vector v associated to the vertex x and the duals of the vertices which are not adjacent of

x. If $N = U$ then Γ is the join of x and Γ'. If not, then we pass to the quotient V/N, which coincides with $H^1(A(\mathrm{Lk}(x)), F)$. Again, the map $V \longrightarrow V/N$ is induced by the inclusion of $A(\mathrm{Lk}(x)) \longrightarrow A(\Gamma)$. Passing to a suitable quotient W/Y of W as above, we can recover the cup product pairing on the cohomology of $A(\mathrm{Lk}(x))$, and thus recover $\mathrm{Lk}(x)$, by induction. Finally, we use the full strength of the induction hypothesis to obtain an injection $i_x\colon \mathrm{Lk}(x) \longrightarrow \Gamma'$. The graph Γ is now reconstructed by attaching x to each vertex in the image of i_x.

To complete the induction, let

$$\{(V_\Lambda, W_\Lambda, q_\Lambda), (V_\Gamma, W_\Gamma, q_\Gamma), i_{\Lambda,\Gamma}\}$$

be a triple satisfying the hypotheses of the claim. We quotient out the degree one part of the cohomology V_Λ and V_Γ by U, and the map $i_{\Lambda,\Gamma}$ descends to the quotients by hypothesis. By induction, we obtain an injection of graphs $\Lambda' \longrightarrow \Gamma'$, where the primed graphs are obtained by deleting the vertex x. The links of x in Γ' and Λ' can be determined as above, whence we can reconstruct Γ.

Case 2 Suppose that $v \in V$ is arbitrary such that the rank k of f_v is minimized, and suppose that v is supported on the duals of two or more vertices, so that

$$v = \sum_{i=1}^{m} \alpha_i x_i^*,$$

where all indices have nonzero coefficients and $m \geq 2$. It is clear that for all i, the degree of x_i must coincide with k, by an easy application of Proposition 15.2.4. Consider the vertices x_1 and x_2. Observe that Proposition 15.2.4 again implies that there cannot be a vertex that is distinct from both x_1 and x_2 and that is adjacent to x_1 but not to x_2. Thus, every vertex that is adjacent to x_1 and distinct from x_2 is also adjacent to x_2. By symmetry, the same statement holds after switching the roles of x_1 and x_2. The argument now bifurcates into two subcases, according to whether x_1 and x_2 are adjacent or not.

Subcase 1 Suppose first that x_1 and x_2 are adjacent, and suppose that $m \geq 3$. Suppose that x_3 is not adjacent to x_1. Then since the degrees of x_2 and x_3 are the same and coincide with the rank k of f_v, we have that

$$|\mathrm{Lk}(x_2) \cup \mathrm{Lk}(x_3)| \geq k + 1.$$

This violates the minimality of the choice of v, since then Proposition 15.2.4 implies that the rank of f_v is at least $k + 1$. It follows that x_1 and x_3 are adjacent, and by symmetry we have that x_2 and x_3 are adjacent. By a straightforward induction, we have that $\{x_1, \ldots, x_m\}$ form a clique, and for all i, a vertex $y \notin \{x_1, \ldots, x_m\}$ adjacent to x_i is adjacent to all the vertices $\{x_1, \ldots, x_m\}$. Observe that if

$$v' = \sum_{i=1}^{m} \beta_i x_i^*$$

is another linear combination of dual vectors, then nonzero linear combinations w of v and v' also satisfy that the rank of f_w is equal to k. We set $V_{\min}$ to be a maximal vector subspace of V that contains v and such that for all $0 \neq w \in V_{\min}$, the rank of f_w is equal to k.

It is straightforward now to show that $V_{\min}$ is generated by $\{x_1^*, \ldots, x_\ell^*\}$, where $\{x_1, \ldots, x_\ell\}$ form an ℓ-clique such that

$$\mathrm{Lk}(x_i) \setminus \{x_j\} = \mathrm{Lk}(x_j) \setminus \{x_i\}$$

for all i and j.

We may now proceed as in Case 1 above, treating $\{x_1, \ldots, x_\ell\}$ as a single vertex, and replacing the subspace U by the subspace $V_{\min}$.

Subcase 2 We now have that x_1 and x_2 are not adjacent. If $m \geq 3$, then the argument in Subcase 1 above implies that x_3 is adjacent to neither x_1 nor x_2. We thus conclude that $\{x_1, \ldots, x_m\}$ form a totally disconnected subgraph of Γ, and $\mathrm{Lk}(x_i) = \mathrm{Lk}(x_j)$ for all i and j. We construct a vector space $V_{\min}$ as in Subcase 1 and conclude that it is generated by $\{x_1^*, \ldots, x_\ell^*\}$, where $\{x_1, \ldots, x_\ell\}$ form a totally disconnected graph and such that the links of any two vertices on this list coincide. We again reduce to Case 1.

□

Some remarks about Theorem 15.2.6 are in order. For one, one need only consider the degree one and degree two parts of the cohomology and not the full cohomology algebra, and this is not surprising since a graph is determined by its vertices and its edges, and a graph determines the corresponding right-angled Artin group. Second, in Case 2 of the proof, the vertices $\{x_1, \ldots, x_\ell\}$ are indistinguishable from each other, in the sense of graph automorphisms. That is, every permutation of $\{x_1, \ldots, x_\ell\}$ is realized by a graph automorphism of Γ, and therefore it is reasonable that one can treat this collection of vertices as a single vertex. Moreover, in the two subcases, $\{x_1, \ldots, x_\ell\}$ generates either an abelian or a free subgroup of $A(\Gamma)$. The full group of automorphisms of $\mathbb{Z}^\ell$ or of F_ℓ embeds in the group $\mathrm{Aut}(A(\Gamma))$ (cf. Sect. 15.3.5 below). Finally, in the proofs of Subcases 1 and 2, we obtain a vector space $V_{\min}$, which either comes from a clique or a totally disconnected subgraph. These two cases can be checked linear algebraically by whether the cup product pairing is trivial or not on $V_{\min}$.

15.3 Translating Between Group Theory and Combinatorics

In this section, we will describe some of the results and ideas that go into translation between the algebraic structure of $A(\Gamma)$ and the combinatorics of Γ. As we have remarked already, the abstract structure of $A(\Gamma)$ determines completely the nature of Γ, passing perhaps through cohomology (Theorem 15.2.6). We will seek clean, definitive results characterizing aspects of the combinatorial structure of Γ in terms of the algebra of $A(\Gamma)$. In the process, we will gain insight into both structures.

15.3.1 Elementary Properties

We begin with some of the first properties of graphs, and how these properties are reflected in $A(\Gamma)$. In Proposition 15.2.5, we have that pairing-connectedness of the triple (V, W, q) characterizes the connectedness of Γ. One can characterize the connectedness of Γ and its complement directly from the group theory of $A(\Gamma)$, without reference to the cohomology algebra, as follows.

Theorem 15.3.1 ([106]) *The group $A(\Gamma)$ splits as a nontrivial direct product if and only if Γ splits as a nontrivial join.*

Recall that a graph Γ splits as a nontrivial join if and only if the complement of Γ is disconnected. Dually, we have the following fact:

Theorem 15.3.2 ([26]) *The group $A(\Gamma)$ splits as a nontrivial free product if and only if Γ is disconnected.*

Both Theorems 15.3.1 and 15.3.2 are easy in one direction. If Γ is disconnected, then $A(\Gamma)$ admits a presentation of the form

$$A(\Gamma) = \langle V(\Gamma_1) \cup V(\Gamma_2) \mid \mathrm{R}_1 \cup \mathrm{R}_2 \rangle,$$

where Γ_1 and Γ_2 are nonempty and disjoint subgraphs of Γ, and where R_i only contains generators from Γ_i for $i \in \{1, 2\}$. It follows then immediately that $A(\Gamma) \cong A(\Gamma_1) * A(\Gamma_2)$.

If Γ splits as a join $\Gamma_1 * \Gamma_2$, then every vertex of Γ_1 is adjacent to every vertex of Γ_2. We have that $A(\Gamma_1)$ and $A(\Gamma_2)$ are subgroups of $A(\Gamma)$, and together generate the whole group. Moreover, they normalize each other and have trivial intersection (this last point is not completely trivial and requires some argument if one wishes to be pedantic, but we shall sweep it under the rug). It follows that $A(\Gamma_1)$ and $A(\Gamma_2)$ generate a direct product.

The converse directions are more complicated, and we outline the main ideas for the convenience of the reader.

Sketch of Proof of Theorem 15.3.2 We use the characterization of free products that follows from the work of Stallings [109, 110]. Let G be a finitely generated group with Cayley graph X. Recall that the set of *ends* of G is the inverse limit of $\pi_0(X \setminus K)$, where K ranges over all compact subgraphs of X. A group has zero, one, two, or infinitely many ends. As right-angled Artin groups are torsion-free (as follows from Proposition 15.2.1 for instance), we have that a right-angled Artin group $A(\Gamma)$ splits as a nontrivial free product if and only if it has infinitely many ends. It thus suffices to argue that a connected graph Γ yields a group with finitely many ends. For a graph with a single vertex, we have $A(\Gamma)$ is $\mathbb{Z}$ and hence has two ends. A straightforward argument shows that if G and H are both infinite groups then $G \times H$ has one end. Thus, we have that all nontrivial joins of graphs yield right-angled Artin groups with one end, and by induction we suppose that all connected graphs with at most n vertices yield groups with at most two ends. Let

$v \in V(\Gamma)$. Then there is a proper subgraph Λ of Γ such that

$$\Gamma = \Lambda \cup_{\mathrm{Lk}(v)} \mathrm{St}(v).$$

We have that $\mathrm{Lk}(v)$ is not empty since Γ is connected. Thus, we have that

$$A(\Gamma) = A(\Lambda) *_{A(\mathrm{Lk}(v))} A(\mathrm{St}(v)).$$

If Λ is connected then $A(\Gamma)$ is an amalgamated product of two finite-ended groups over an infinite subgroup (cf. [105]), whence one can prove directly that $A(\Gamma)$ is one-ended. If Λ is disconnected, then one can argue component-by-component of Λ to obtain the same conclusion. □

For Theorem 15.3.1, we require a basic result about the structure of centralizers of elements in $A(\Gamma)$. Let w be a word in the vertices of Γ and their inverses. We say that w is *reduced* if w cannot be shortened by applications of free reductions and moves of the form $[v_1^{\pm 1}, v_2^{\pm 1}]$ for $\{v_1, v_2\} \in E(\Gamma)$. We say that w is *cyclically reduced* if it remains reduced after allowing cyclic permutations of the letters occurring in w. It is true but not trivial that the moves of free reduction and commutation solve the word problem in right-angled Artin groups, and that cyclic reduction solves the conjugacy problem (see especially [39], cf. [31, 65, 114, 120]).

The *support* of w is written $\mathrm{supp}(w)$ and is defined to be the set of vertices which are required (possibly inverted) to express w. It is not completely trivial but true that the support of w is well-defined in the sense that for reduced words, $w_1 = w_2$ in $A(\Gamma)$ implies that $\mathrm{supp}(w_1) = \mathrm{supp}(w_2)$.

Theorem 15.3.3 ([106]) *Let $1 \neq w \in A(\Gamma)$ be cyclically reduced. Then the centralizer of w lies in $\langle \mathrm{supp}(w) \cup \mathrm{Lk}(\mathrm{supp}(w)) \rangle$. If the centralizer of w is not cyclic then either $\mathrm{Lk}(\mathrm{supp}(w))$ is nonempty, or $\mathrm{supp}(w)$ decomposes as a nontrivial join.*

Armed with Theorem 15.3.3, we can illustrate the other direction of Theorem 15.3.1.

Proof of Theorem 15.3.1 Suppose that $A(\Gamma) \cong G \times H$ for nontrivial groups G and H. Then since $A(\Gamma)$ is torsion-free, we have that every nontrivial element of $A(\Gamma)$ contains a copy of $\mathbb{Z}^2$ in its centralizer. Writing $V(\Gamma) = \{v_1, \ldots, v_n\}$, we have that $w = v_1 \cdots v_n$ is cyclically reduced and has noncyclic centralizer. Moreover, $\mathrm{Lk}(\mathrm{supp}(w)) = \emptyset$, so that Theorem 15.3.3 implies that $\mathrm{supp}(w) = \Gamma$ splits as a nontrivial join. □

Theorem 15.3.3 has several other important consequences that relate the algebraic structure of $A(\Gamma)$ to the combinatorics of Γ. First, we have the following.

Theorem 15.3.4 *The cohomological dimension of $A(\Gamma)$ coincides with the size of the maximal clique in Γ.*

Theorem 15.3.4 follows from standard ideas about cohomological dimension (cf. [30]), using the description of the Salvetti complex as a union of tori together with the fact that it is aspherical by Proposition 15.2.1. We have that the maximal dimensional cells in $\mathscr{S}(\Gamma)$ have the same dimension as the maximal size of a clique in Γ, say k. Moreover, this k-cell is the top dimensional cell in a subtorus of dimension k, which has nontrivial cohomology in degree k. Finally, the retraction $\mathscr{S}(\Gamma) \longrightarrow (S^1)^k$ implies that the degree k cohomology of $\mathscr{S}(\Gamma)$ is also nontrivial. It follows that k is also the cohomological dimension of $A(\Gamma)$.

The cohomological dimension and maximal clique size also describe the rank of a maximal abelian subgroup.

Theorem 15.3.5 *The maximal clique size of* Γ *coincides with the rank of a maximal abelian subgroup of* $A(\Gamma)$.

Theorem 15.3.5 is also a consequence of general facts about cohomological dimension. Clearly, if the maximal clique size of Γ is k then $A(\Gamma)$ contains a copy of $\mathbb{Z}^k$. Since $\mathscr{S}(\Gamma)$ is k-dimensional and aspherical, it follows that no cover of $\mathscr{S}(\Gamma)$ can have fundamental group $\mathbb{Z}^{k+1}$, so there are no abelian subgroups of rank exceeding k.

For another perspective, suppose Γ is connected and $G < A(\Gamma)$ is an abelian subgroup of rank $k \geq 2$. Conjugating if necessary, at least one nontrivial element of G is cyclically reduced, so that Theorem 15.3.3 implies that all nontrivial elements of G are supported on a subgraph J of Γ that splits as a nontrivial join. Writing $A(J) \cong A(J_1) \times A(J_2)$, we may restrict the projections $A(J) \longrightarrow A(J_i)$ for each i to G.

Now, suppose first that Γ has no triangles (i.e. 3-cliques). Then J_1 and J_2 cannot have any edges, since otherwise Γ would have a triangle. It follows then that $A(J_i)$ is free for $i \in \{1, 2\}$, and so the image of G in $A(J_i)$ is cyclic for each i. It follows that G has rank at most two. Thus, we may assume by induction that if the maximal clique size of Γ is at most $k \geq 2$ then the maximal abelian subgroup has rank at most k. Supposing Γ has maximal clique size $k+1$, then J_1 and J_2 have maximal clique sizes $k_1 > 0$ and $k_2 > 0$, which satisfy $k_1 + k_2 \leq k + 1$. It follows by induction that the ranks of the images of G in $A(J_1)$ and $A(J_2)$ are at most k_1 and k_2, so that G has rank at most $k+1$. This gives an alternate proof of Theorem 15.3.5.

The final elementary combinatorial property of graphs we will discuss is the maximal degree of a vertex. This property is essential in the theory of expander graphs, which will be discussed below. In the sequel we will use a different characterization of the maximal degree that is understood through cohomology, though the following is a significantly cleaner statement.

Proposition 15.3.6 *Let* Γ *be a graph and let* d *denote the maximum valence of a vertex of* Γ. *Then the rank of the centralizer of a nontrivial element of* $A(\Gamma)$ *is at most* $d+1$. *Conversely, if for all elements* $1 \neq g \in A(\Gamma)$ *the centralizer of* g *has rank at most* $d+1$, *then the maximum degree of a vertex of* Γ *is at most* d.

The proof of Proposition 15.3.6 is a fairly straightforward application of Theorem 15.3.3, and we leave it as an exercise for the reader.

15.3.2 k-Colorability

From the point of computational complexity, one of the most basic and difficult questions one can pose about a graph Γ is about its colorability. A *(vertex) coloring* of a graph Γ is a function $\kappa : V(\Gamma) \longrightarrow X$, where X is a finite set of *colors*, and where $\{v, w\} \in E(\Gamma)$ implies that $\kappa(v) \neq \kappa(w)$. A classical result of Brooks [43] says that the minimal size of X is at most the maximal degree of a vertex of Γ plus one. If Γ is not an odd length cycle or a clique then the bound can be improved to the maximal degree of a vertex. The minimal size of X is called the *chromatic number* of Γ, and we say that Γ is $|X|$-colorable.

A graph that is 2-colorable is called *bipartite*. Determining if a graph is bipartite is easy from a computational point of view, and can be accomplished by a sorting algorithm that runs in a period of time that is bounded by a polynomial in the size of the set of vertices. However, the problem of determining if a graph is 3-colorable is extremely difficult from a computational standpoint, and is *NP-complete* (see Sect. 15.3.7 below).

We remark that there is a related notion of *edge coloring*, which is a function $\epsilon : E(\Gamma) \longrightarrow X$ such that if $v \in V(\Gamma)$ is incident to both e_1 and e_2, then $\epsilon(e_1) \neq \epsilon(e_2)$. It is clear that the minimal size of X for a valid edge coloring is bounded below by the maximal degree of a vertex of Γ. A result of Vizing [43] shows that Γ admits an edge coloring with $|X|$ the maximal degree of a vertex of Γ plus one. Thus, giving sharp or almost sharp estimates on edge colorability of a graph is an essentially local problem, whereas determining vertex colorability is an essentially global problem.

Let Γ be a k-colorable graph. Choose a k-coloring of Γ, and add an edge to Γ for every pair of vertices with different colors, naming the result Λ. Observe that the vertices of Λ are partitioned as

$$V(\Lambda) = V_1 \cup \cdots \cup V_k,$$

where there are no edges between vertices in V_i for each i, and where for $i \neq j$, each vertex of V_i is adjacent to each vertex of V_j. It follows that $A(\Lambda)$ is a product of free groups, and that $A(\Lambda)$ is a quotient of $A(\Gamma)$. It turns out that these elementary considerations characterize k-colorable graphs.

Theorem 15.3.7 *Let Γ be a finite graph with N vertices. Then Γ is k-colorable if and only if there is a surjective map*

$$A(\Gamma) \longrightarrow \prod_{i=1}^{k} F_{n_i},$$

where for each i the group F_{n_i} is free of rank n_i, and where

$$\sum_{i=1}^{k} n_i = N.$$

We have already established the "only if" direction, which is easy. The reverse direction is somewhat more substantial, owing to the fact that the surjective homomorphism need not send vertex generators of Γ to a free factor of one of the free groups occurring on the right hand side.

Sketch of Proof of Theorem 15.3.7 We identify the product of free groups with $A(\Lambda)$, where Λ is a k-fold join of completely disconnected graphs, say $\{\Lambda_1, \ldots, \Lambda_k\}$. If $g \in A(\Lambda)$, then g can be written uniquely as a product of $g_1 \cdots g_k$, where $g_i \in A(\Lambda_i)$. One then shows that if $g = g_1 \cdots g_k$ and $h = h_1 \cdots h_k$ are elements of $A(\Lambda)$ that commute, then for each i, the images of g_i and h_i in the abelianization of $A(\Lambda_i)$ are rational multiples of each other.

Now, the surjective map $A(\Gamma) \longrightarrow A(\Lambda)$ induces an isomorphism

$$\phi\colon H_1(A(\Gamma), \mathbb{Q}) \longrightarrow H_1(A(\Lambda), \mathbb{Q}),$$

which can be expressed as a matrix A with respect to the vertex generators of both graphs. We will view the rows of A as expressions for $\phi(v)$ for $v \in V(\Gamma)$, in terms of the vertex generators of Λ. We arrange the columns so that the first $|V(\Lambda_1)|$ columns correspond to vertices of Λ_1, followed by the vertices of Λ_2, and so on.

Write $A = (A_1 \mid \cdots \mid A_k)$, where the columns of A_i correspond to the vertices of Λ_i, and therefore the column space of A_i has dimension $|V(\Lambda_i)|$. Note that since A is invertible, the row space of A_1 has dimension $|V(\Lambda_1)| = n_1$.

It is an exercise in linear algebra to show that there is an $n_1 \times n_1$ minor B_1 of A_1 and a $(N - n_1) \times (N - n_1)$ minor C of $(A_2 \mid \cdots \mid A_k)$ such that both B_1 and C are invertible.

By induction, one permutes the rows of A to obtain a block matrix $B = (B_1 \mid \cdots \mid B_k)$ such that the diagonal $n_i \times n_i$ blocks $\{C_1, \ldots, C_k\}$ of B are invertible. This row permutation is simply a permutation of the vertices of Γ. One defines a coloring of the vertices by setting $\kappa(v_i) = j$ if in the matrix expression B of ϕ, we have that the row $\phi(v_i)$ meets the block C_j. That is, the vertices corresponding to the first n_1 rows are assigned color 1, the next n_2 are assigned color 2, and so on.

To check that this is a valid coloring, suppose v and w are adjacent in Γ. Then $[v, w] = 1$ in $A(\Gamma)$. For each block B_i, we may consider the restriction of the rows $\phi(v)$ and $\phi(w)$ to the columns in B_i. In B_i, these two rows are rational multiples of each other. If v and w were assigned the same color then in some block B_i, the rows both meet B_i in the diagonal sub-block C_i. Since C_i is invertible, this is a contradiction. Thus, we see that adjacent vertices of Γ are assigned different colors, and so the coloring of Γ is valid. □

Unpacking the final check that κ is a valid coloring in the proof of Theorem 15.3.7, it is not difficult to see that in fact one can relax the condition that the homomorphism $A(\Gamma) \longrightarrow A(\Lambda)$ be surjective, and replace it with the condition that it be surjective on the level of rational homology. From a practical point of view, this is a useful observation. Indeed, checking that a linear map is surjective is relatively easy, but maps to direct products of free groups are much less well-behaved, since the subgroup structure of the latter is very complicated [97].

15.3.3 Hamiltonicity

In addition to computing the chromatic number of a finite graph, a classical NP-complete problem in graph theory is deciding whether a given connected graph admits a Hamiltonian cycle. Here, a *Hamiltonian cycle* is a circuit in Γ that visits every vertex of Γ exactly once. A graph that admits a Hamiltonian cycle is simply called Hamiltonian. Much like vertex colorability versus edge colorability, there is a notion of a circuit in Γ that traverses every edge exactly once, called an *Eulerian cycle*. It is a standard fact that a connected graph admits an Eulerian cycle if and only if each vertex has even degree. Thus, determining whether a graph admits an Eulerian cycle is a purely local question, and the existence of a Hamiltonian cycle is a global question, impervious to local methods. We direct the reader to [43] for background on Eulerian and Hamiltonian paths and cycles in graphs.

Let $(x_0, \ldots, x_n)$ denote a Hamiltonian cycle in Γ, and let

$$\{x_0^*, \ldots, x_n^*\} \subset V = H^1(A(\Gamma), F)$$

denote the corresponding dual classes. Proposition 15.2.4 implies that

$$x_i^* \smile x_{i+1}^* \neq 0$$

for all i, where the indices are considered cyclically modulo n. This is the fundamental observation when it comes to characterizing Hamiltonicity of Γ in terms of the intrinsic algebra of $A(\Gamma)$.

Let (V, W, q) be a triple consisting of a vector space V equipped with a vector-space-valued (i.e. W-valued) (anti)-symmetric bilinear pairing. We will assume that V is finite dimensional. We say that (V, W, q) is *Hamiltonian* if for all bases $\{v_0, \ldots, v_n\}$ of V, there is a permutation $\sigma \in S_{n+1}$ such that for all i, we have $q(v_{\sigma(i)}, v_{\sigma(i+1)}) \neq 0$.

Setting

$$V = H^1(A(\Gamma), F), \quad W = H^2(A(\Gamma), F), \quad q = \smile,$$

suppose that (V, W, q) is Hamiltonian. Then there is a basis $\{x_0^*, \ldots, x_n^*\}$ consisting of classes dual to the vertices of Γ. The Hamiltonicity of the triple immediately implies the existence of a permutation σ such that

$$x_{\sigma(i)}^* \smile x_{\sigma(i+1)}^* \neq 0$$

for all relevant indices, which immediately implies that Γ admits a Hamiltonian cycle.

Theorem 15.3.8 (See [55]) *Let Γ and (V, W, q) be as above. Then Γ admits a Hamiltonian cycle if and only if (V, W, q) is Hamiltonian.*

The reader may check as an easy exercise that the Hamiltonicity of a triple (V, W, q) implies that the triple is in fact pairing-connected, so that if (V, W, q) is Hamiltonian then Γ is automatically connected by Proposition 15.2.5.

Establishing Theorem 15.3.8 is tricky, and requires significantly more insight than Theorem 15.3.7, for instance. We will attempt to briefly convey the main ideas to the reader in the remainder of this subsection. The reader is directed to [55] for a full account.

In order to establish Theorem 15.3.8, it is clearly sufficient to show that if Γ is Hamiltonian then the triple (V, W, q) is also Hamiltonian. One may begin with the standard dual basis $\{x_0^*, \ldots, x_n^*\}$ for V and hope to bootstrap it to show that (V, W, q) is Hamiltonian. One can begin with a change of basis matrix A, which transforms $\{x_0^*, \ldots, x_n^*\}$ to a given basis $\{v_0, \ldots, v_n\}$ for V. We write $A = (a_i^j)$, where the subscript refers to the row and the superscript refers to the column of a given entry.

We leave it as an easy exercise for the reader to show the following:

Lemma 15.3.9 *The triple (V, W, q) is Hamiltonian if for all $A \in GL_{n+1}(A)$, there is a permutation $\sigma \in S_{n+1}$ such that for all $0 \leq i \leq n$, there exists a $0 \leq j \leq n$ such that*

$$A_i^j = \begin{pmatrix} a_{\sigma(i)}^j & a_{\sigma(i)}^{j+1} \\ a_{\sigma(i+1)}^j & a_{\sigma(i+1)}^{j+1} \end{pmatrix}$$

is invertible, where all indices are considered cyclically.

Lemma 15.3.9 gives rise to a natural definition that one can associate to matrices (which need not be invertible, or even square). The *two-row graph* $\mathscr{G}(A)$ of a matrix A is a graph whose vertices are the rows $\{r_0, \ldots, r_n\}$ of A, and whose columns are given by the relation $\{r_i, r_j\} \in E(\mathscr{G}(A))$ if the matrix

$$A_{i,j}^k = \begin{pmatrix} a_i^k & a_i^{k+1} \\ a_j^k & a_j^{k+1} \end{pmatrix}$$

is invertible for some k.

It is clear from Lemma 15.3.9 that (V, W, q) is Hamiltonian provided that $\mathscr{G}(A)$ is itself Hamiltonian for all suitable matrices A. To get a feel for $\mathscr{G}(A)$, the reader is encouraged to prove directly that $\mathscr{G}(A)$ is connected whenever A is invertible. The heart of the proof of Theorem 15.3.8 is the following:

Lemma 15.3.10 *Let A be an invertible matrix. Then $\mathscr{G}(A)$ is Hamiltonian.*

Lemma 15.3.10 is a curious fact in its own right, and its proof is fairly involved. Producing a Hamiltonian cycle directly in $\mathscr{G}(A)$ appears to be a difficult problem itself, and which has the feel of an NP-complete problem (though this is by no means a theorem). Thus, one needs to use more indirect methods to find a Hamiltonian cycle in $\mathscr{G}(A)$.

The key idea is to analyze block submatrices of a matrix A which consist of nonzero entries with one-dimensional row spaces. One can consider maximal such blocks, which give rise to a partition of the products of entries of A which contribute to the determinant of A, according to the standard Leibniz formula. Using certain symmetries, one can then argue that if no Hamiltonian cycle exists in $\mathscr{G}(A)$ then all summands in the determinant of A cancel in pairs, and hence the determinant of A is zero.

15.3.4 Graph Expanders

In this subsection, we leave behind individual graphs, and concentrate on families of graphs known as graph expanders. Graph expanders are sequences of connected graphs that are uniformly sparse and uniformly difficult to separate. Expander families find applications in a myriad of different fields, such as knot theory, spectral graph theory and spectral geometry, probabilistic computation, and network theory. We direct the reader to [3, 23, 24, 67, 89, 92–94] for references relevant to this section.

A sequence of finite graphs $\{\Gamma_i\}_{i\in\mathbb{N}}$ is called a *graph expander family* if the following conditions are satisfied:

(1) There is a d such that for all i, the maximum degree of a vertex of Γ_i is at most d.
(2) We have $|V(\Gamma_i)| \longrightarrow \infty$.
(3) The Cheeger constant of Γ_i is uniformly bounded away from zero, independently of i.

Here, the *Cheeger constant* of a graph Γ is defined by considering subsets $A \subset V(\Gamma)$ such that $|A| \leq |V(\Gamma)|/2$, and by looking at ∂A, which is defined to be the set of vertices of $V(\Gamma) \setminus A$ that are adjacent to a vertex of A. The *isoperimetric constant* of A is defined to be

$$c_A = \frac{|\partial A|}{|A|},$$

and the Cheeger constant c is the minimum of c_A as A ranges over all admissible subsets of $V(\Gamma)$. From this point of view, it is clear why the Cheeger constant measures the difficulty in separating Γ: in order to completely cut a set $A \subset V(\Gamma)$ out of Γ, one has to sever at least $c \cdot |A|$ edges. The reader may check that if $\{\Gamma_i\}_{i\in\mathbb{N}}$ forms a graph expander family then the Cheeger constant inequality implies that each graph in the family is connected.

By associating the standard cohomology triple

$$V_i = H^1(A(\Gamma_i), F),\ W_i = H^2(A(\Gamma_i), F),\ q_i = \smile,$$

some of the defining properties of graph expanders translate almost immediately. Namely, we have $|V(\Gamma_i)| \longrightarrow \infty$ if and only if $\dim V_i \longrightarrow \infty$, and Γ_i is connected if and only if (V_i, W_i, q_i) is q_i-pairing-connected.

The remaining conditions for defining graph expanders require some new ideas. The degree of a vertex is already characterized in terms of centralizers via Proposition 15.3.6 above. Since centralizers of elements are less transparently cohomological objects, we first translate this notion of degree into linear algebra. Let (V, W, q) be a vector space with a vector-space-valued bilinear pairing. If $\emptyset \neq S \subset V$ and B is a basis for V, we write

$$d_B(S) = \max_{s \in S} |\{b \in B \mid q(s, b) \neq 0\}|.$$

To get rid of the dependence on B, we set $d(S)$ to be the minimum of $d_B(S)$, taken over all possible bases. To get rid of the dependence on S, we set $d(V)$ to be the minimum of $d(S)$, taken over all S which span V. The quantity $d(V)$ is called the *q-valence* of V.

A reader who has understood the ideas in the proof of Theorem 15.2.6 will have no trouble proving the following fact:

Proposition 15.3.11 *Let (V, W, q) be the usual cohomological triple associated to $A(\Gamma)$, and let d be the maximum degree of a vertex of Γ. Then $d(V) = d$.*

It remains to properly define the Cheeger constant of the triple (V, W, q). Again, a reader who has absorbed the ideas in the proof of Theorem 15.2.6 could probably guess the definition. Let $Z \subset V$ be a vector space with $0 \neq \dim Z \leq (\dim V)/2$. We will write C for the orthogonal complement of Z, which is to say the set of vectors $v \in V$ such that $q(v, z) = 0$ for all $z \in Z$. The isoperimetric constant of Z is defined to be

$$c_Z = \frac{\dim V - \dim Z - \dim C + \dim(C \cap Z)}{\dim Z}.$$

The Cheeger constant c_V of the triple (V, W, q) is taken to be the infimum of c_Z as Z varies over all nonzero subspaces of V of dimension at most half of that of V.

Let $\{x_1, \ldots, x_n\}$ denote the vertices of Γ and $\{x_1^*, \ldots, x_n^*\}$ be the dual generators of $H^1(A(\Gamma), F)$. If $B \subset \{x_1, \ldots, x_n\}$, write B^* for the corresponding subset of $\{x_1^*, \ldots, x_n^*\}$. The following is an exercise for the reader:

Proposition 15.3.12 *Let $\emptyset \neq B \subset V$, and let $Z \subset V$ be generated by B^*. Then*

$$c_Z = \frac{|\partial B|}{|B|}.$$

Thus, the Cheeger constant of Γ is bounded below by c_V. *A priori*, there are many more subspaces of V than there are subgraphs of Γ, so that in principle c_V could be strictly smaller than the Cheeger constant of Γ.

Theorem 15.3.13 (See [53]) *Let* $\{\Gamma_i\}_{i\in\mathbb{N}}$ *be a sequence of graphs, and let*

$$\{(V_i, W_i, q_i)\}_{i\in\mathbb{N}}$$

be the corresponding cohomological triples. We have that $\{\Gamma_i\}_{i\in\mathbb{N}}$ *forms a family of expanders if and only if:*

(1) There is a d *such that the* q_i*-valence of* V_i *is bounded above by* d.
(2) We have $\dim V_i \longrightarrow \infty$.
(3) There is an $\epsilon > 0$ *such that for all* i, *we have* $c_{V_i} \geq \epsilon$.

As in the case of graph expander families, pairing-connectedness in the sequence of cohomological triples may be assumed but is actually redundant since it is a consequence of the Cheeger constant bound. An abstract sequence of triples $\{(V_i, W_i, q_i)\}_{i\in\mathbb{N}}$ is called a family of *vector space expanders* (not to be confused with *dimensional expanders*, cf. [23, 24, 93]). In light of the preceding discussion, in order to establish Theorem 15.3.13, it suffices to show that for each i, the Cheeger constant c_{V_i} coincides with the Cheeger constant of Γ_i. Unfortunately, the author does not know a conceptually simple proof of this fact. The proof given in [53] involves a rather technical sorting argument, and so we will not comment on it any further.

15.3.5 Graph Automorphisms

One of the most basic questions one can ask about a graph (and indeed about a relation) is how symmetric it is. Symmetry is measured by the richness of the automorphism group, and the smaller the size of the automorphism group, however it is measured, the less symmetric the object.

The automorphisms of graphs are of great interest in graph theory [17, 43, 58], and in complexity theory as well [7]. Many finite graphs are highly symmetric. For instance, the automorphism group of a k-clique is the full symmetric group on k letters. Many other graphs have no nontrivial automorphisms. For instance, take a path of length five, with vertices labeled linearly as $\{a, b, c, d, e, f\}$, and add another vertex g which is adjacent only to d. The resulting graph Γ has no nontrivial automorphisms, as is readily verified by an exhaustive check. See Fig. 15.1.

Observe that a nontrivial automorphism of a graph Γ gives rise to a non-inner automorphism of $A(\Gamma)$. Moreover, if $v \in V(\Gamma)$, then the function $v \mapsto v^{-1}$ extends to a non-inner automorphism of $A(\Gamma)$ via the identity on the remaining vertices. It

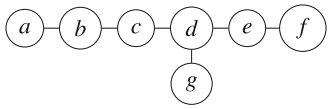

Fig. 15.1 A graph that has no nontrivial automorphisms

is easy to see that the group $\mathrm{Aut}(A(\Gamma))$, and in fact $\mathrm{Out}(A(\Gamma))$, contains a subgroup isomorphic to

$$\mathrm{Aut}(\Gamma) \ltimes (\mathbb{Z}/2\mathbb{Z})^{|V(\Gamma)|}.$$

Thus, if Γ admits a nontrivial automorphism, then $\mathrm{Out}(A(\Gamma))$ contains a nonabelian finite subgroup.

Theorem 15.3.14 (See [51]) *Let Γ be a finite simplicial graph. We have that Γ admits a nontrivial automorphism if and only if $\mathrm{Out}(A(\Gamma))$ contains a finite nonabelian subgroup.*

The "only if" direction follows from the discussion preceding Theorem 15.3.14. The converse is significantly harder and requires a more careful analysis of $\mathrm{Out}(A(\Gamma))$.

A result of M. Laurence ([90], cf. [106]) says that $\mathrm{Aut}(A(\Gamma))$ is generated by automorphisms of the following type.

(1) Vertex inversions.
(2) Graph automorphisms.
(3) Partial conjugations.
(4) Dominated transvections.

Graph automorphisms have already been discussed, and *vertex inversions* have been mentioned above as arising from the map $v \mapsto v^{-1}$ for some $v \in V(\Gamma)$. A *partial conjugation* is given by considering a vertex $v \in V(\Gamma)$ whose star $\mathrm{St}(v)$ separates Γ. The automorphism acts by conjugation by v on one component of $\Gamma \setminus \mathrm{St}(v)$ and by the identity on the remaining components of Γ.

To define *dominated transvections*, we say that a vertex $v \in V(\Gamma)$ *dominates* a vertex $w \in V(\Gamma)$ if $\mathrm{Lk}(w) \subset \mathrm{St}(v)$. Then, the map $w \mapsto wv$ extends to an automorphism of $A(\Gamma)$ via the identity on the remaining vertices. Domination is clearly a relation on vertices of Γ that can be determined from the combinatorics of Γ.

Sketch of Proof of Theorem 15.3.14 We suppose that Γ admits no nontrivial automorphisms. A theorem of Toinet [113] implies that if $\phi \in \mathrm{Aut}(A(\Gamma))$ has finite order then ϕ acts nontrivially on $H_1(A(\Gamma), \mathbb{Z}) \cong \mathbb{Z}^{|V(\Gamma)|}$. Thus, it suffices to consider the action of automorphisms on $H_1(A(\Gamma), \mathbb{Z})$, and the effect of partial conjugations is then trivial.

Next, one shows that if there is a cycle $\{v_1, v_2, \ldots, v_k, v_1\}_{k \geq 2}$ where v_i dominates v_{i+1} (with the indices considered cyclically), then Γ admits a nontrivial automorphism, specifically an automorphism that exchanges two vertices of Γ. It follows that no such cycles exist. We may therefore order the vertices of Γ in such a way that if $v_i < v_j$ then v_j cannot dominate v_i. If we then write the image of $\mathrm{Out}(A(\Gamma))$ in $\mathrm{GL}_n(\mathbb{Z})$ with respect to the corresponding ordered basis for $H_1(A(\Gamma), \mathbb{Z})$, the result is a group of upper triangular integer matrices. Such a group has only abelian finite subgroups (coming from diagonal matrices with entries ± 1).

Thus, if Γ has no automorphisms then $\mathrm{Out}(A(\Gamma))$ has only abelian finite subgroups. □

We remark that in the proof of Theorem 15.3.14, one of the key observations is that a graph with a domination cycle admits a nontrivial automorphism. The converse of this statement is false. The 5-cycle C_5 admits many automorphisms, but no two vertices dominate each other.

15.3.6 Some Further Entries in the Combinatorics–Algebra Dictionary

There are a number of other results relating the combinatorics of graphs to the algebraic structure of groups which we will not discuss in detail for the sake of space. We briefly mention two results appearing in [66]. Recall that a group G is *poly-free* if there is a finite length subnormal filtration of G by subgroups such that successive quotients are free. Hermiller–Šunić proved that a right-angled Artin group is always poly-free, and that the length of the poly-free filtration is bounded above by the chromatic number of the defining graph. In the same paper, they established that $A(\Gamma)$ is a semidirect product of two finitely generated free groups if and only if Γ is a tree or a *complete bipartite graph*, which is to say a join of two completely disconnected graphs. Moreover, for a connected graph Γ with at least two vertices, the poly-free length of $A(\Gamma)$ is exactly two if and only if there is a subset $D \subset V(\Gamma)$ such that no pair of elements of D spans an edge, and every circuit in Γ meets D in at least two vertices. It is an interesting direction for future research to investigate the relationship between the normal structure of $A(\Gamma)$ and the combinatorics of Γ, and it appears that this subject is largely unexplored.

15.3.7 Usefulness Beyond Group Theory and Combinatorics

The various correspondences between combinatorics of graphs and algebraic structures of groups have theoretical and practical applications beyond the structural framework of Question 15.1.1 and its refinements. Here, we record some specific examples.

15.3.7.1 Complexity of Problems in Combinatorial Group Theory

One of the main applications of the foregoing discussion is in the domain of complexity theory, which is hardly surprising in light of the fact that many computationally difficult problems (i.e. NP-complete problems, cf. [6, 56, 98]) are formulated in a finitistic way, with reference to only combinatorial structures.

Consider a right-angled Artin group $A(\Gamma)$, and a homomorphism

$$\phi\colon A(\Gamma) \longrightarrow F_{n_1} \times F_{n_2} \times F_{n_3},$$

where F_{n_i} denotes a free group of rank n_i, and where

$$n_1 + n_2 + n_3 = |V(\Gamma)|.$$

If Γ is specified (e.g. by a list of vertices and pairs of adjacent vertices) and ϕ is specified in terms of the image of each vertex of Γ with respect to a fixed free basis of each of the free group factors in the target of ϕ, then it is easy to check if ϕ is a homomorphism that is surjective on the level of first rational homology. Indeed, it suffices to check first that ϕ is well-defined, meaning that adjacent vertices in Γ are sent to commuting elements of $F_{n_1} \times F_{n_2} \times F_{n_3}$, which can be performed efficiently. The latter claim results from the fact that centralizers of elements in $F_{n_1} \times F_{n_2} \times F_{n_3}$ are straightforward to describe, and because the word problem is efficiently solvable. Then, one must check that ϕ is surjective on the level of first rational homology, which is an easy linear algebra problem. In light of Theorem 15.3.7, the data specifying the homomorphism ϕ forms a (short) certificate of the fact that Γ is 3-colorable. Since the 3-colorability of Γ and the existence of this homomorphism are equivalent, the problem of deciding whether such a homomorphism exists is NP-complete. To state this conclusion formally:

Proposition 15.3.15 *Let Γ be a finite graph with $|V(\Gamma)| = N$, and let*

$$\{F_{n_1}, \ldots, F_{n_k}\}$$

be free groups such that

$$\sum_{i=1}^{k} n_i = N.$$

Write $G = \prod_i F_{n_i}$.

(1) If $k = 2$ then the problem of deciding whether or not there exists a homomorphism $A(\Gamma) \longrightarrow G$ that is surjective on first rational homology is in P.
(2) If $k = 3$ then the problem of deciding whether or not there exists a homomorphism $A(\Gamma) \longrightarrow G$ that is surjective on first rational homology is NP-complete.
(3) The problem of finding the minimal k for which there exist free groups $\{F_{n_1}, \ldots, F_{n_k}\}$ as above and a homomorphism $A(\Gamma) \longrightarrow G$ that is surjective on first rational homology is NP-complete.

Finding explicit examples of NP-complete problems is always of interest in complexity theory, and given the profusion of them in graph theory, Proposition 15.3.15 is just a taste of the richness of the available theory arising in the context of groups.

15.3.7.2 Hamiltonicity Testing

Continuing in the theme of NP-complete problems, it is well-known that deciding if a finite graph admits a Hamiltonian path or Hamiltonian cycle is NP-complete, as we have mentioned above. The ideas surrounding Theorem 15.3.8 can be used to certify that certain graphs are not Hamiltonian, in a purely finitistic linear algebraic way.

To expand on this a bit, first note that the field over which cohomology is considered is arbitrary. In particular, we may assume that the underlying field is just the field with two elements. Under this assumption, all the relevant vector spaces become finite sets, and are hence amenable to combinatorial techniques.

Consider then the standard cohomological triple (V, W, q) for a right-angled Artin group $A(\Gamma)$. In order to show that Γ is not Hamiltonian, it suffices to find a single basis for V which witnesses the claim that (V, W, q) is not Hamiltonian. Thus, such a basis can be used as a short certificate that a graph contains no Hamiltonian circuit.

15.3.7.3 Linear Algebraic Detection of Graph Expanders

Considering cohomology with coefficients in a field with two elements allows a finitary and algebraic way to check if a sequence of graphs is a family of expanders. Moreover, it is shown in [53] that there are families of vector space expanders that do not arise from the cohomology of families of graph expanders. Thus, the theory of vector space expanders is *a priori* richer than the theory of graph expanders. Some practical applications of expanders can be found in [35, 59], for instance.

15.3.7.4 Interactive Proof Systems

Many interactive proof systems function as a way for a prover to demonstrate a proposition to a skeptical verifier. Using an unbiased random bit sent by the verifier, the prover sends a response that is conditioned on the value of the random bit. In this way, the verifier's ignorance of the prover's private information is balanced by the prover's ignorance of the value of the bit that will be sent by verifier, and this balance can be used to communicate the existence of knowledge without revealing its content. This is, for instance, the idea behind *zero-knowledge proof protocols*, in which the prover holds a certificate for an instance of an NP-complete problem, and convinces the verifier of the fact that she is in possession of a valid certificate without revealing the certificate itself. Any NP-complete problem can be used as a platform. Thus, linear algebraic versions of Hamiltonicity as in Theorem 15.3.8 and Proposition 15.3.15 are suitable for formulating a zero-knowledge proof protocol. A detailed explanation of a platform using Theorem 15.3.8 is given in [55]. For general background on interactive proofs and zero-knowledge proof protocols, we refer the reader to [6, 8, 22, 60, 100].

15.3.7.5 Group-Based Cryptosystems

Many cryptosystems rely on computational problems that are difficult to solve directly, which is why many modern cryptographic protocols assume $P \neq NP$. The theme of this section has been the translation of combinatorial properties of graphs, and especially computationally interesting ones, into algebraic language. This immediately suggests numerous potential group-based cryptosystems, a topic which has been developing rapidly in recent decades. Explicit cryptosystems using right-angled Artin groups as a platform have been proposed in [50], for example. Translating the graph homomorphism problem (which is NP-complete) into an instance of the subgroup homomorphism problem for right-angled Artin groups, one can formulate a secure authentication scheme, for instance. For further discussion of specific cryptosystems and for a biased sample of the literature, we direct the reader to [49, 51, 52, 75, 84, 99].

15.4 The Extension Graph and Its Properties

We now leave the world of the finite graph Γ and its relationship with $A(\Gamma)$, and turn to the (usually) infinite extension graph Γ^e. We recall that Γ^e is a development of Γ into a graph on which $A(\Gamma)$ acts by conjugation. So, we fix an identification of the vertices of Γ with generators for $A(\Gamma)$, set the vertices of Γ^e to be the collection of all conjugates of $V(\Gamma)$ by elements of $A(\Gamma)$, and set the edge relation to be commutation inside of $A(\Gamma)$. The reader will find that the ideas here, though still fundamentally relating combinatorics to algebra, are quite different from those in Sect. 15.3.

15.4.1 Basic Properties of the Extension Graph

Some properties of the extension graph are easy to prove. For instance:

Proposition 15.4.1 *The extension graph Γ^e is finite if and only if Γ is complete.*

Others are somewhat less obvious. We note some which will be useful in the sequel, and which otherwise will give the reader a better idea of how the extension graph functions.

Proposition 15.4.2 (See [78]) *The extension graph Γ^e enjoys the following properties:*

(1) The graph Γ^e is connected if and only if Γ is connected.
(2) The graph Γ^e is connected and of infinite diameter if and only if Γ is connected, has at least two vertices, and is not a join.

(3) The size of a maximal clique in Γ *and* Γ^e *coincide.*
(4) If Λ *is a subgraph of* Γ *then* Λ^e *is a subgraph of* Γ^e.
(5) The graph Γ^e *is k-colorable if and only if* Γ *is k-colorable.*

The proof of item (2) of Proposition 15.4.2 we will provide probably illustrates the diversity of methods that can be used in investigating right-angled Artin groups.

Sketch of Proof of Proposition 15.4.2, (2) Consider a collection of disjoint compact annuli $\{A_v \mid v \in V(\Gamma)\}$, one for each vertex of Γ. Glue two such annuli A_v and A_w together along a disk if the vertices v and w are *not* adjacent in Γ. We do this in such a way that the result is an orientable surface Σ with boundary. A key observation is that since Γ is not a join, its complement graph X is connected. Therefore, Σ is a connected surface. Since Σ was built out of at least two annuli, an easy Euler characteristic computation shows that Σ is of hyperbolic type (i.e. admits a complete hyperbolic metric of finite volume). We will name the core curves of the annuli in the construction $\{\gamma_1, \ldots, \gamma_n\}$.

The (isotopy class of the) homeomorphism of Σ given by cutting Σ open along γ_i and re-gluing with a full right-handed twist is called a (right-handed) *Dehn twist* about γ_i, and is denoted by T_i. Recall that the group of isotopy classes of (orientation preserving) homeomorphisms of Σ is called the *mapping class group* of Σ, and is written $\mathrm{Mod}(\Sigma)$ [47]. A result of the author [85] shows that there is an $N > 0$ such that for all $k \geq N$, the subgroup of $\mathrm{Mod}(\Sigma)$ generated by $\{T_1^k, \ldots, T_n^k\}$ is isomorphic to $A(\Gamma)$.

The surface Σ has an associated *curve graph* $\mathscr{C}(\Sigma)$, which is of infinite diameter. This curve graph consists of isotopy classes of embedded, essential, nonperipheral loops on Σ, with the edge relation being disjoint realization. There are certain mapping classes ψ which have the property that for any vertex c of $\mathscr{C}(\Sigma)$, the graph distance between c and $\psi^k(c)$ tends to infinity as k tends to infinity [112]. These mapping classes are called *pseudo-Anosov*, and are typical inside of $\mathrm{Mod}(\Sigma)$.

In particular, realizing $A(\Gamma) < \mathrm{Mod}(\Sigma)$ as above, there is an element $g \in A(\Gamma)$ whose realization as a mapping class is pseudo-Anosov. Moreover, the realization $A(\Gamma) < \mathrm{Mod}(\Sigma)$ is compatible with a realization of $\Gamma^e \subset \mathscr{C}(\Sigma)$. Specifically, if $v \in V(\Gamma)$ is associated to a Dehn twist about γ_i and if $h \in A(\Gamma)$ corresponds to the mapping class group ψ_h, then the vertex v^h is sent to $\psi_h(\gamma_i)$.

Now, since we have a map $\Gamma^e \longrightarrow \mathscr{C}(\Sigma)$ which respects the edge relation, general facts about graph homomorphisms imply that it cannot be distance increasing. Thus, if $d_{\mathscr{C}(\Sigma)}(\gamma_i, \psi_h^k(\gamma_i))$ tends to infinity then $d_{\Gamma^e}(v, v^{h^k})$ also tends to infinity. The conclusion now follows. □

It turns out that mapping class groups of surfaces are extremely useful tools for probing right-angled Artin groups, and that many of their properties can be paired analogously. This is a theme that will recur in this section, and we will comment more on it below.

15.4.2 The Extension Graph and Subgroups

One useful property of the extension graph, and for which it was developed in the first place, is that the extension graph classifies right-angled Artin subgroups of a right-angled Artin group. Classically, we know that subgroups of finitely generated free abelian groups are again free abelian (by the classification of finitely generated modules over a principal ideal domain) and subgroups of free groups are always free (by the Nielsen–Schreier Theorem). Since right-angled Artin groups interpolate between these two extremes, it is therefore a natural question whether (finitely generated) subgroups of right-angled Artin groups are again right-angled Artin groups, and if so what sorts of right-angled Artin groups they are.

It is not true that subgroups of right-angled Artin groups are again right-angled Artin groups. There are many different subgroups of right-angled Artin groups, ranging from surface groups [38, 107] to hyperbolic 3-manifold groups [1, 2, 117–119] to many arithmetic lattices in rank one Lie groups [18], all the way to groups with various exotic finiteness properties [19]. It is in fact known that every finitely generated subgroup of $A(\Gamma)$ is again a right-angled Artin group if and only if Γ has no subgraph isomorphic to a square or to a path of length three, by a result of Droms [45].

It is difficult to characterize all subgroups of right-angled Artin groups, even finitely presented ones (see [28]). Some general known facts are that every nonabelian subgroup of a right-angled Artin group contains a nonabelian free group by a result of Baudisch [10], and in fact any such subgroup surjects to a nonabelian free group by a result of Antolín–Minasyan [5]. A nonabelian subgroup of a right-angled Artin group must surject to $\mathbb{Z}^2$ [46, 87]. Solvable subgroups of right-angled Artin groups are automatically finitely generated and free abelian, by the Flat Torus Theorem [29].

Given the difficulty of understanding general subgroups of right-angled Artin groups, it is therefore interesting and natural to wonder which subgroups of $A(\Gamma)$ are of the form $A(\Lambda)$, and what sorts of graphs Λ can occur. To the author's knowledge, there is no clean, complete answer available, though the partial answers are satisfying and useful for many applications.

Theorem 15.4.3 *Let $\Lambda < \Gamma^e$ be a finite subgraph. Then there is an injective homomorphism $A(\Lambda) \longrightarrow A(\Gamma)$.*

The injection in Theorem 15.4.3 is quite explicit; one simply views vertices of Λ as elements in $A(\Gamma)$ and passes to a sufficiently high power. Theorem 15.4.3 first appeared in a paper of Kim and the author [78], though apparently this fact was already known to experts in combinatorial group theory. One approach to proving Theorem 15.4.3 does not require ideas beyond those that go into item (2) of Proposition 15.4.2. Once the extension graph has been embedded in the curve graph in a way that preserves both adjacency and non-adjacency, the author's result from [85] about powers of mapping classes applies and gives the desired result.

Fig. 15.2 The square

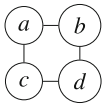

Unfortunately, Theorem 15.4.3 does not admit an easy converse. The first examples disproving the obvious naïve converse appeared in the work of Casals-Ruiz–Duncan–Kazachkov [33], and a large class of examples was produced by Kim and the author [81]. With some further assumptions on Γ, one can formulate a converse to Theorem 15.4.3.

Theorem 15.4.4 (See [78]) *Suppose* Γ *has no* 3*-cliques, and suppose that* $A(\Lambda) < A(\Gamma)$. *Then* Λ *is a subgraph of* Γ^e.

Theorem 15.4.4 is a corollary of a more general result, which is the most general converse to Theorem 15.4.3 that is known to the author.

Theorem 15.4.5 (See [78]) *Suppose that* $A(\Lambda) < A(\Gamma)$. *Then* Λ *is a subgraph of the clique graph* $(\Gamma^e)_k$.

The basic idea behind Theorem 15.4.5 is again to use mapping class groups, though it is significantly more complicated than Theorem 15.4.3 and Proposition 15.4.2. One builds certain "partial" pseudo-Anosov mapping classes in the image of $A(\Lambda)$ and builds an embedding of a larger graph X into Γ^e, which contains Λ in its clique graph. Incidentally, Theorem 15.4.5 has a natural analogue for mapping class groups: if a right-angled Artin group $A(\Gamma)$ embeds in a mapping class group $\mathrm{Mod}(\Sigma)$, then Γ embeds as a subgraph of $\mathscr{C}(\Sigma)_k$, the clique graph of the curve graph ([80], cf. [82]). We will avoid giving further details here.

Theorem 15.4.5 admits several other corollaries that can serve as converses to Theorem 15.4.3, and also allows one to prove many results that relate the combinatorics of Γ to the structure of $A(\Gamma)$. Given the conclusion of Theorem 15.4.5, we leave the following result (originally due to Kambites [76], who offered a combinatorial argument that is very different in flavor from the ideas expounded here) as an exercise for the reader:

Proposition 15.4.6 *Let* Γ *be a finite graph. Then* Γ *contains a square if and only if* $F_2 \times F_2 < A(\Gamma)$.

Here, by *square* we mean a graph with four vertices and a cyclic adjacency relation (Fig. 15.2).

15.4.3 A Characterization of Cographs via Right-Angled Artin Groups and the Geometry of the Extension Graph

An important class of graphs that occurs naturally in graph theory is the class of *cographs*, or P_4-free graphs (see [74, 103, 111] for some early references

Fig. 15.3 The graph P_4

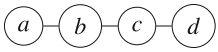

Fig. 15.4 The graph X_6

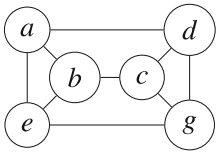

introducing cographs). These are simply the graphs that do not have the path P_4 of length three as a subgraph (Fig. 15.3).

Right-angled Artin groups on cographs can be characterized algebraically, and right-angled Artin groups provide a perspective on cographs that insight into one of their most fundamental properties, i.e. recursive definition.

Theorem 15.4.7 ([78]) *Let Γ be a finite connected graph. The following are equivalent:*

(1) The graph Γ has no (full) subgraph isomorphic to P_4.
(2) The graph Γ^e has no (full) subgraph isomorphic to P_4.
(3) The graph Γ is either a single vertex or splits as a nontrivial join.

Corollary 15.4.8 *The graph Γ is a cograph if and only if $A(\Gamma)$ does not contain a copy of $A(P_4)$.*

In particular, Theorem 15.4.7 shows that a right-angled Artin group cannot contain "hidden" copies of $A(P_4)$. If $A(\Gamma)$ contains $A(P_4)$ then one can decide simply from looking at the graph Γ. This is in contrast to other classes of graphs. For instance, $A(P_4)$ contains a copy of $A(P_5)$, where P_5 denotes the path of length four. Thus, there can be hidden copies of $A(P_5)$. For a more striking example, one may consider X_6, the complement graph of the hexagon, also known as the triangular prism. This graph contains no cycle C_5 of length 5, though Kim proved that $A(C_5) < A(X_6)$ [77]; also, C_5 is a subgraph of the extension graph X_6^e, and so Kim's result follows from Theorem 15.4.3 (Fig. 15.4).

We leave the proofs of Theorem 15.4.7 and Corollary 15.4.8 as an exercise for the reader, as they follow from Proposition 15.4.2 and some elementary combinatorial group theory considerations.

Let Γ be a connected graph such that Γ^e has finite diameter. By Theorem 15.4.7 (or even just by Proposition 15.4.2), the graph Γ splits as a nontrivial join. If Γ is a cograph and Λ is a join factor of Γ, then Λ must also be a cograph and hence Λ^e also has finite diameter, whence it follows that Λ must also split as a nontrivial join.

Let $\mathscr{K}_0$ denote a singleton vertex. For $i > 0$, we set $\mathscr{K}_{2i-1}$ to be the collection of all finite graphs obtained as (possibly trivial) joins of elements of $\mathscr{K}_{2i-2}$. We set $\mathscr{K}_{2i}$ to be the collection of all finite graphs obtained as disjoint unions of elements

of $\mathscr{K}_{2i-1}$. Clearly for $i \leq j$ we have $\mathscr{K}_i \subset \mathscr{K}_j$, and we set

$$\mathscr{K} = \bigcup_{i \geq 0} \mathscr{K}_i.$$

Clearly, if $\Gamma \in \mathscr{K}$ then Γ is a cograph. Conversely, the preceding remarks and an easy induction on $|V(\Gamma)|$ show that if Γ is a cograph then $\Gamma \in \mathscr{K}$. This coincides with the recursive description of cographs.

Since $\mathscr{K}$ is built up recursively, we can give the following characterization of $A(\Gamma)$ for $\Gamma \in \mathscr{K}$, which results immediately from the preceding discussion:

Corollary 15.4.9 *We have $\Gamma \in \mathscr{K}$ if and only if $A(\Gamma)$ is an element of the smallest class of groups that:*

(1) Contains $\mathbb{Z}$;
(2) Is closed under finite direct products;
(3) Is closed under finite free products.

For example, note that $\Gamma \in \mathscr{K}_0$ if and only if $A(\Gamma) \cong \mathbb{Z}$. We have $\Gamma \in \mathscr{K}_1$ if and only if $A(\Gamma) \cong \mathbb{Z}^n$ for some n. We have $\Gamma \in \mathscr{K}_2$ if and only if $A(\Gamma)$ is a free product of free abelian groups. A graph Γ lies in $\mathscr{K}_3$ if and only if $A(\Gamma)$ is a direct product of free products of free abelian groups. The following characterizes $\mathscr{K}_i$ for $i \leq 3$:

Proposition 15.4.10 (See [83]) *A graph Γ lies in $\mathscr{K}_i$ for $i \leq 3$ if and only if $A(\Gamma)$ has no subgroup isomorphic to $(F_2 \times \mathbb{Z}) * \mathbb{Z}$.*

As an aside, we note that the hierarchy $\mathscr{K}$ and the associated right-angled Artin groups is closely related to the theory of right-angled Artin group actions on the interval and on the circle. It turns out that $A(P_4)$ does not act faithfully by C^2 diffeomorphisms on I or S^1 [9], so any right-angled Artin group admitting such an action must have its underlying graph in $\mathscr{K}$. By a result of Kim and the author [83], a right-angled Artin group $A(\Gamma)$ admits a faithful C^2 action on I or S^1 if and only if it admits a faithful C^∞ such action, if and only if $\Gamma \in \mathscr{K}_3$.

15.4.4 More on the Geometry of the Extension Graph

As we have suggested in this section, and in particular in the discussion about Proposition 15.4.2, the extension graph of Γ plays a role analogous to that of the curve graph $\mathscr{C}(\Sigma)$ of a surface, with the role of the mapping class group in the latter context played by the group $A(\Gamma)$ in the former context.

The graph $\mathscr{C}(\Sigma)$ is very complicated in both its local and its global structure. One of the most important foundational results about the global structure of $\mathscr{C}(\Sigma)$ is a result of Masur and Minsky which asserts that $\mathscr{C}(\Sigma)$ is *δ-hyperbolic* [96], see also [57, 61]. That is, there is a $\delta \geq 0$ so that in any geodesic triangle in $\mathscr{C}(\Sigma)$, a δ-neighborhood of two of the sides of the triangle contains the third.

Perhaps the easiest example of an infinite diameter δ-hyperbolic metric space is an infinite diameter tree, which is 0-hyperbolic. There are many other examples of δ-hyperbolic spaces that are not trees, such the usual hyperbolic spaces. For most surfaces, the curve graph $\mathscr{C}(\Sigma)$ is far from being a tree; it has one end, whereas for example a locally finite tree that admits a proper and cocompact action by an infinite group will have at least two ends, as follows from Bass–Serre Theory [105].

The geometry of the extension graph is something in between the curve graph and a tree. To state a precise result, we need the notion of a quasi-isometry. Let $f: X \longrightarrow Y$ be a function between metric spaces. Then we say that f is a *quasi-isometry* if there are constants $\lambda \geq 1$ and $C \geq 0$ such that for all $x, z \in X$, we have

$$\frac{1}{\lambda} \cdot d_X(x, z) - C \leq d_Y(f(x), f(z)) \leq \lambda \cdot d_X(x, z) + C,$$

and where for all $y \in Y$ there exists an $x \in X$ such that

$$d_Y(f(x), y) \leq C.$$

Here, the distance functions are all interpreted in the relevant spaces. A quasi-isometry can be thought of a function that is bi-Lipschitz on a large scale. For instance, the integers equipped with the metric induced from the real line are quasi-isometric to the real line, and any two finite-diameter metric spaces are quasi-isometric to each other, but an infinite-diameter metric space is not quasi-isometric to a finite-diameter metric space.

The relation induced by quasi-isometry is an equivalence relation on metric spaces, and so one often speaks of the quasi-isometry class of a metric space. The quasi-isometry class of a finitely generated group is the quasi-isometry class of its Cayley graph, equipped with the graph metric; see [41] for more details, for example.

In coarse geometry, one often searches for properties of metric spaces that are invariant under quasi-isometry. Examples of such properties include δ-hyperbolicity and the number of ends.

A metric space is called a *quasi-tree* if it contains a 0-hyperbolic metric space in its quasi-isometry class. Whereas simplicial trees are 0-hyperbolic, the converse is not quite true: a geodesic metric space is 0-hyperbolic if and only if it is an $\mathbb{R}$-tree. We will not discuss $\mathbb{R}$-trees any further, since they are not necessary for our discussion. We specialize the definition of a quasi-tree slightly: if Γ is a graph equipped with the graph metric, we call it a quasi-tree if it contains a simplicial tree in its quasi-isometry class.

Theorem 15.4.11 (See [78]) *Let Γ be a connected graph. Then Γ^e is a quasi-tree, and is in particular δ-hyperbolic. More precisely:*

(1) If Γ splits as a nontrivial join, then Γ^e has finite diameter and is hence quasi-isometric to a point.

(2) If Γ does not split as a nontrivial join then Γ^e is quasi-isometric to a regular simplicial tree of countable degree.

More interesting than the mere description of the quasi-isometry type of the extension graph is the interaction between group elements in $A(\Gamma)$ and Γ^e. Here, the analogy between the mapping class group and $A(\Gamma)$ develops further, with the natural isometric action of $A(\Gamma)$ on Γ^e mirroring many of the properties of the natural isometric action of $\mathrm{Mod}(\Sigma)$ on $\mathscr{C}(\Sigma)$.

The classical *Nielsen–Thurston classification* [47, 112] says that a mapping class is either finite order, reducible (i.e. some power fixes the homotopy class of an essential nonperipheral loop on the surface Σ), or pseudo-Anosov. As discussed around Proposition 15.4.2, this lattermost type of mapping class is characterized by the fact that every orbit of its action on $\mathscr{C}(\Sigma)$ is unbounded. Finite order and reducible mapping classes are characterized by every orbit in $\mathscr{C}(\Sigma)$ being bounded (and in fact having a periodic point in $\mathscr{C}(\Sigma)$). Algebraically, a reducible mapping class has a copy of $\mathbb{Z}^2$ in its centralizer [21], whereas a pseudo-Anosov mapping classes have virtually cyclic centralizers [48].

Further insight into the action of $\mathrm{Mod}(\Sigma)$ is provided by a result of Bowditch [25], which says that the action of $\mathrm{Mod}(\Sigma)$ on $\mathscr{C}(\Sigma)$ is *acylindrical*. Acylindricity is a notion of proper discontinuity for group actions on non-proper metric spaces which are not properly discontinuous. Following Bowditch (cf. [86, 104]) we say that an action of a group G on a metric space X is acylindrical if for all $r > 0$ there exist constants R and N such that for all pairs of points $x, y \in X$ with $d(x, y) \geq R$, we have

$$|\{g \in G \mid d(gx, x), d(gy, y) \leq r\}| \leq N.$$

In other words, the r-quasi-stabilizer of R-separated points is uniformly finite. Bowditch showed that if X is a δ-hyperbolic graph and G acts isometrically and acylindrically on X then each $g \in G$ is either *elliptic* or *loxodromic*. The former of these means that some (equivalently every) orbit of G on X is bounded. A loxodromic element is characterized by having a positive asymptotic translation distance in X. Moreover, the asymptotic translation length is bounded away from zero by a constant that depends only on the hyperbolicity and acylindricity constants. The Nielsen–Thurston classification can be thus recast in terms of acylindricity: a mapping class is pseudo-Anosov if and only if it is loxodromic as an isometry of $\mathscr{C}(\Sigma)$.

For extension graphs, one has a picture that is analogous to curve graphs.

Theorem 15.4.12 (See [79]) *Let Γ be a connected graph with at least two vertices. The action of $A(\Gamma)$ on Γ^e is acylindrical. An element $1 \neq g \in A(\Gamma)$ is elliptic if and only if g is conjugate into a subgroup $A(J)$, where J is a subgraph of Γ that is a nontrivial join. Equivalently, g is elliptic if and only if its centralizer in $A(\Gamma)$ is noncyclic.*

An element $1 \neq g \in A(\Gamma)$ *is loxodromic if and only if its centralizer is cyclic. An element* g *is cyclically reduced and loxodromic if and only if* $\mathrm{supp}(g)$ *is not contained in a subgraph of* Γ *that splits as a nontrivial join.*

The join/non-join dichotomy for graphs and their associated right-angled Artin groups runs deep, and analogies between $A(\Gamma)$ and Γ^e with $\mathrm{Mod}(\Sigma)$ and $\mathscr{C}(\Sigma)$ are extensive. Many (but not all; see [88]) of the instances of these analogies can be and have been incorporated into the theory of hierarchically hyperbolic groups.

A further equivalence in Theorem 15.4.12 is given by a result of Behrstock–Charney [11], which asserts that a nontrivial element of $A(\Gamma)$ is loxodromic if and only if, when viewed as a deck transformation of the universal cover of the Salvetti complex $\mathscr{S}(\Gamma)$, it acts as a *rank one isometry*. That is, the corresponding deck group element has an axis that does not bound a half-plane (cf. [29]).

15.4.5 The Extension Graph as a Quasi-Isometry and Commensurability Invariant

A basic problem in geometric group theory is to sort groups into quasi-isometry classes. For right-angled Artin groups, the natural question is to decide when two right-angled Artin groups $A(\Gamma)$ and $A(\Lambda)$ are quasi-isometric. Much progress on understanding the quasi-isometric classification of right-angled Artin groups has been made, for instance by Behrstock–Neumann [12], Behrstock–Januszkiewicz–Neumann [13], Bestvina–Kleiner–Sageev [20], Huang [70], and Margolis [95] (see also [32, 71]). Thus, we can consider the following equivalence relation on finite graphs: Γ is equivalent to Λ if $A(\Gamma)$ and $A(\Lambda)$ are quasi-isometric to each other. Other than the cases we have cited, understanding this equivalence relation in full is still unresolved.

Certainly two right-angled Artin groups that are isomorphic to each other will be quasi-isometric to each other, and from Theorem 15.2.6, we know that if $A(\Gamma)$ and $A(\Lambda)$ are isomorphic to each other then Γ and Λ are isomorphic as graphs. There is yet another equivalence relation on finite graphs that is coarser than isomorphism and yet finer than quasi-isometry.

If $H < G$ are groups with G finitely generated and $[G : H] < \infty$, then with respect to any finite generating sets for G and H, the inclusion of H into G is a quasi-isometry on the level of Cayley graphs, as is readily verified. It follows that if G and H are finitely generated groups, and both G and H contain a finite index subgroup isomorphic to K, then G and H are quasi-isometric. In this case, we say that G and H are *commensurable*. Like quasi-isometry, commensurability is an equivalence relation on groups. It is well known that commensurability of groups is a strictly finer equivalence relation than quasi-isometry. For instance, one can take closed hyperbolic 3-manifolds whose volumes are not rational multiples of each other. Then, the corresponding fundamental groups are both quasi-isometric to hyperbolic space, but are not commensurable [57]. Even among right-angled Artin groups, commensurability is a strictly finer equivalence relation (see [34, 70]).

It is easy to produce pairs of non-isomorphic graphs which give rise to commensurable right-angled Artin groups. For instance, consider a graph Γ and $v \in V(\Gamma)$. There is a surjective homomorphism $A(\Gamma) \longrightarrow \mathbb{Z}/2\mathbb{Z}$ that sends v to the nontrivial element in $\mathbb{Z}/2\mathbb{Z}$ and sends the remaining vertices to the identity. It is an exercise in combinatorial group theory for the reader to prove that the kernel of this homomorphism is isomorphic to $A(\Lambda)$, where Λ is obtained by taking two copies of Γ and identifying them along $\mathrm{St}(v)$. If v is not central in $A(\Gamma)$ then it is easy to see that Λ and Γ fail to be isomorphic graphs, but $A(\Gamma)$ and $A(\Lambda)$ are clearly commensurable. This construction can be repeated *ad infinitum*, generally producing infinite families of non-isomorphic graphs whose associated right-angled Artin groups are all commensurable.

There are pairs of graphs which give rise to commensurable right-angled Artin groups, but for which a commensuration between them is less obvious. The reader is challenged to prove for themself that the groups $A(P_4)$ and $A(P_5)$ are commensurable, where as before P_4 and P_5 denote the paths of length three and length four respectively (cf. [34]). The fact that $A(P_4)$ and $A(P_5)$ are commensurable also shows that the extension graph is hopeless as a complete commensurability invariant. Again, the reader is encouraged to convince themself that the extension graphs of P_4 and P_5 are not isomorphic to each other. It turns out that in both cases, the corresponding extension graphs are trees, and what distinguishes them in their isomorphism type is the location of degree one vertices.

So, let us consider a connected graph Γ with no degree one vertices. In order to identify the extension graph algebraically and in an unambiguous way, it would help to be able to identify vertices and their conjugates, up to powers. For this, it helps to assume that Γ is connected, has no triangles, and has no squares. Under these assumptions, if v is a vertex of Γ then v contains a nonabelian free group in its centralizer. Conversely, suppose that $g \in A(\Gamma)$ has a nonabelian free group in its centralizer. Then, since Γ has no triangles and no squares, every nontrivial join in Γ is merely the star of a vertex of Γ, and the structure of such a star is the join of a single vertex and a completely disconnected graph. It follows that if g has a nonabelian free group in its centralizer, then g is conjugate to a nonzero power of a vertex generator of Γ. It follows that maximal cyclic subgroups of $A(\Gamma)$ whose centralizers contain nonabelian free groups are in bijection with conjugates of vertex generators of $A(\Gamma)$. Since the adjacency relation in Γ^e is just commutation in $A(\Gamma)$, we immediately obtain:

Theorem 15.4.13 (See [79]) *Let Γ be a finite connected graph with no degree one vertices, no triangles, and no squares. Then the extension graph Γ^e is a commensurability invariant for $A(\Gamma)$. That is, if $A(\Gamma)$ is commensurable with $A(\Lambda)$ then $\Gamma^e \cong \Lambda^e$.*

Incidentally, the analogy between right-angled Artin groups and mapping class groups persists here as well, since the curve graph can be obtained from the mapping class group in the same way that the extension graph is obtained from $A(\Gamma)$. Specifically, let T be a Dehn twist about a simple closed curve on Σ. Then T is centralized by two maximal rank torsion-free abelian subgroups of $\mathrm{Mod}(\Sigma)$ which

intersect in a copy of $\mathbb{Z}$. This can be used to algebraically characterize a (nonzero power of a) Dehn twist as an element of $\mathrm{Mod}(\Sigma)$. A Dehn twist unambiguously identifies the homotopy class of a simple closed curve on Σ, and the adjacency relation in $\mathscr{C}(\Sigma)$ coincides with commutation of Dehn twists in $\mathrm{Mod}(\Sigma)$. Thus, the curve graph can be recovered algebraically from $\mathrm{Mod}(\Sigma)$. It follows in particular that automorphisms of $\mathrm{Mod}(\Sigma)$ induce automorphisms of $\mathscr{C}(\Sigma)$, a fact which can be used to prove various rigidity results (see [27, 72, 91], for instance).

As we have seen, we can have commensurable right-angled Artin groups with non-isomorphic extension graphs. It is also possible to have two right-angled Artin groups whose extension graphs are isomorphic and yet the groups are not quasi-isometric to each other (see Example 5.22 in [68]). So, there is similarly no hope that extension graphs form a complete quasi-isometry invariant for right-angled Artin groups.

Recall from Sect. 15.3.5 that a full set of generators for $\mathrm{Aut}(A(\Gamma))$ is known, and from the description of these generators, it is immediate that $\mathrm{Out}(A(\Gamma))$ is finite if and only if $\mathrm{Aut}(A(\Gamma))$ admits no nontrivial partial conjugations and no dominated transvections. Graphs for which $\mathrm{Out}(A(\Gamma))$ is finite can thus be identified through a finitary combinatorial analysis, since it suffices to check that there are no separating stars of vertices and no pairs of vertices where one dominates the other (see [37] for a discussion of the genericity of this phenomenon).

The following result was established by Huang [69]:

Theorem 15.4.14 *Suppose* Γ *is a graph for which* $\mathrm{Out}(A(\Gamma))$ *is finite. The following are equivalent:*

(1) The group $A(\Lambda)$ *is quasi-isometric to* $A(\Gamma)$.
(2) The group $A(\Lambda)$ *is isomorphic to a finite index subgroup of* $A(\Gamma)$.
(3) The graphs Λ^e *and* Γ^e *are isomorphic.*

Thus, in the case of finite groups of outer automorphisms, quasi-isometry, commensurability, and isomorphism of extension graphs are equivalent conditions to place on a right-angled Artin group. Here again, the analogy with mapping class groups persists. If two mapping class groups of surfaces are quasi-isometric, then except for some sporadic cases, the resulting mapping class groups are in fact isomorphic to each other [14]. Thus again excluding some sporadic cases, quasi-isometry, commensurability, and isomorphism of mapping class groups are equivalent. Finally, aside from some sporadic cases, isomorphism of curve graphs is equivalent to isomorphism of mapping class groups [108].

15.5 Further Directions

Much remains to be understood in the relationship between combinatorics and algebra via the lens of right-angled Artin groups. As the reader has certainly come to understand, it is not just some property of groups that one seeks to analogize a

property of graphs; one wants it to be a clean and natural statement about groups that reflects the particular flavor of the property in question. Therefore, it is not likely one could produce a satisfactory omnibus result, since some subjective notions of beauty and philosophical considerations enter into the picture.

With these musings, we close by giving some particular open questions of interest. Some are well-known open problems, and we make no claim to having been the first to pose them.

Question 15.5.1 *What is the full quasi-isometric classification of right-angled Artin groups? What about the commensurability classification of right-angled Artin groups? What sorts of combinatorial objects serve as complete invariants for these equivalence relations?*

Some specific natural combinatorial properties we have not discussed are of interest in graph theory.

Question 15.5.2 *What algebraic property of $A(\Gamma)$ is equivalent to the planarity of Γ?*

Closely related to Question 15.5.2 is the problem of determining whether a graph Λ is a subdivision of a graph Γ by examining the relationship between the groups $A(\Lambda)$ and $A(\Gamma)$, which to the knowledge of the author is also open.

A graph is *self-complementary* if it is isomorphic to its complement graph. A singleton vertex is self-complementary, as are the path P_4 of length three and the cycle C_5 of length five. A question that is a particular favorite of the author is the following:

Question 15.5.3 *What algebraic property of $A(\Gamma)$ is equivalent to the statement that Γ is self-complementary?*

Following the remarks in Sect. 15.3.6 above and the results of [66], we have the following.

Question 15.5.4 *What is the relationship between the normal subgroup structure of $A(\Gamma)$ and the combinatorics of Γ?*

Finally, we have the following more open-ended question.

Question 15.5.5 *Is there a synthesis between the ideas in Sect. 15.3 and algebraic graph theory? How can one formulate spectral graph theory in terms of right-angled Artin groups?*

Some of the discussion in this survey is a step towards an answer to Question 15.5.5. For one, the Cheeger constant c of a graph can be viewed as a spectral invariant of a graph, as it controls the spectral gap of the graph via the Cheeger inequality due to Dodziuk and Alon–Milman (see [89] for a detailed discussion): if λ_2 is the second largest eigenvalue of a d–regular connected graph Γ then

$$\frac{1}{2}(d-\lambda_2) \leq c \leq \sqrt{2d(d-\lambda_2)}.$$

The content of Theorem 15.3.13 is that the Cheeger constant of a graph Γ can be read off from the cohomology algebra of $A(\Gamma)$. It is natural to ask how one might recover more information about the eigenvalues of the adjacency matrix of Γ from the group theory of $A(\Gamma)$.

We hope that this survey will encourage further investigations in these directions.

Acknowledgements The author is partially supported by an Alfred P. Sloan Foundation Research Fellowship, by NSF Grant DMS-1711488, and by NSF Grant DMS-2002596. The author thanks K. Ohshika and A. Papadopoulos for inviting him to write this survey, and to R. Flores for providing numerous invaluable comments on an earlier draft.

References

1. I. Agol, Criteria for virtual fibering. J. Topol. **1**(2), 269–284 (2008)
2. I. Agol, The virtual Haken conjecture. Doc. Math. **18**, 1045–1087 (2013). With an appendix by I. Agol, D. Groves, J. Manning
3. N. Alon, Eigenvalues and expanders. Combinatorica **6**, 83–96 (1986). *Theory of Computing* (Singer Island, FL, 1984)
4. J. Amorós, M. Burger, K. Corlette, D. Kotschick, D. Toledo, *Fundamental Groups of Compact Kähler Manifolds*. Mathematical Surveys and Monographs, vol. 44 (American Mathematical Society, Providence, RI, 1996)
5. Y. Antolín, A. Minasyan, Tits alternatives for graph products. J. Reine Angew. Math. **704**, 55–83 (2015)
6. S. Arora, B. Barak, *Computational Complexity: A Modern Approach* (Cambridge University Press, Cambridge, 2009)
7. L. Babai, Graph isomorphism in quasipolynomial time [extended abstract], in *STOC'16—Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing* (ACM, New York, 2016), pp. 684–697
8. L. Babai, S. Moran, Arthur-Merlin games: a randomized proof system, and a hierarchy of complexity classes. J. Comput. Syst. Sci. **36**, 254–276 (1988). *17th Annual ACM Symposium on the Theory of Computing* (Providence, RI, 1985), pp. 254–276
9. H. Baik, S. Kim, T. Koberda, Unsmoothable group actions on compact one-manifolds. J. Eur. Math. Soc. **21**(8), 2333–2353 (2019)
10. A. Baudisch, Subgroups of semifree groups. Acta Math. Acad. Sci. Hungar. **38**(1–4), 19–28 (1981)
11. J. Behrstock, R. Charney, Divergence and quasimorphisms of right-angled Artin groups. Math. Ann. **352**(2), 339–356 (2012)
12. J.A. Behrstock, W.D. Neumann, *Quasi-isometric classification of graph manifold groups*, Duke Math. J. **141**(2), 217–240 (2008)
13. J.A. Behrstock, T. Januszkiewicz, W.D. Neumann, Quasi-isometric classification of some high dimensional right-angled Artin groups. Groups Geom. Dyn. **4**(4), 681–692 (2010)
14. J. Behrstock, B. Kleiner, Y. Minsky, L. Mosher, Geometry and rigidity of mapping class groups. Geom. Topol. **16**(2), 781–888 (2012)
15. J. Behrstock, M.F. Hagen, A. Sisto, Hierarchically hyperbolic spaces, I: curve complexes for cubical groups. Geom. Topol. **21**(3), 1731–1804 (2017)
16. J. Behrstock, M.F. Hagen, A. Sisto, Hierarchically hyperbolic spaces II: combination theorems and the distance formula. Pac. J. Math. **299**(2), 257–338 (2019)
17. L.W. Beineke, R.J. Wilson (eds.), *Topics in Algebraic Graph Theory, Encyclopedia of Mathematics and Its Applications*, vol. 102 (Cambridge University Press, Cambridge, 2004)

18. N. Bergeron, F. Haglund, D.T. Wise, Hyperplane sections in arithmetic hyperbolic manifolds. J. Lond. Math. Soc. (2) **83**(2), 431–448 (2011)
19. M. Bestvina, N. Brady, Morse theory and finiteness properties of groups. Invent. Math. **129**(3), 445–470 (1997)
20. M. Bestvina, B. Kleiner, M. Sageev, The asymptotic geometry of right-angled Artin groups. I. Geom. Topol. **12**(3), 1653–1699 (2008)
21. J.S. Birman, A. Lubotzky, J. McCarthy, Abelian and solvable subgroups of the mapping class groups. Duke Math. J. **50**(4), 1107–1120 (1983)
22. M. Blum, How to prove a theorem so no one else can claim it, in *Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Berkeley, Calif., 1986)* (Amer. Math. Soc., Providence, RI, 1987), pp. 1444–1451
23. J. Bourgain, Expanders and dimensional expansion. C. R. Math. Acad. Sci. Paris **347**(7–8), 357–362 (2009)
24. J. Bourgain, A. Yehudayoff, Expansion in $\mathrm{SL}_2(\mathbb{R})$ and monotone expanders. Geom. Funct. Anal. **23**(1), 1–41 (2013)
25. B.H. Bowditch, Tight geodesics in the curve complex. Invent. Math. **171**(2), 281–300 (2008)
26. N. Brady, J. Meier, Connectivity at infinity for right angled Artin groups. Trans. Am. Math. Soc. **353**(1), 117–132 (2001)
27. T.E. Brendle, D. Margalit, Commensurations of the Johnson kernel. Geom. Topol. **8**, 1361–1384 (2004)
28. M. Bridson, On the subgroups of right-angled Artin groups and mapping class groups. Math. Res. Lett. **20**(2), 203–212 (2013)
29. M.R. Bridson, A. Haefliger, *Metric Spaces of Non-positive Curvature, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*, vol. 319 (Springer, Berlin, 1999)
30. K.S. Brown, *Cohomology of Groups, Graduate Texts in Mathematics*, vol. 87 (Springer, New York, 1982)
31. P. Cartier, D. Foata, *Problèmes combinatoires de commutation et réarrangements*. Lecture Notes in Mathematics, No. 85 (Springer, Berlin, 1969)
32. M. Casals-Ruiz, Embeddability and quasi-isometric classification of partially commutative groups. Algebr. Geom. Topol. **16**(1), 597–620 (2016)
33. M. Casals-Ruiz, A. Duncan, I. Kazachkov, Embedddings between partially commutative groups: two counterexamples. J. Algebra **390**, 87–99 (2013)
34. M. Casals-Ruiz, I. Kazachkov, A. Zakharov, On commensurability of right-angled Artin groups I: RAAGs defined by trees of diameter 4. Rev. Mat. Iberoam. **35**(2), 521–560 (2019)
35. D.X. Charles, K.E. Lauter, E.Z. Goren, Cryptographic hash functions from expander graphs. J. Cryptol. **22**(1), 93–113 (2009)
36. R. Charney, An introduction to right-angled Artin groups. Geom. Dedicata **125**, 141–158 (2007)
37. R. Charney, M. Farber, Random groups arising as graph products. Algebr. Geom. Topol. **12**(2), 979–995 (2012)
38. J. Crisp, B. Wiest, Embeddings of graph braid and surface groups in right-angled Artin groups and braid groups. Algebr. Geom. Topol. **4**, 439–472 (2004)
39. J. Crisp, E. Godelle, B. Wiest, The conjugacy problem in subgroups of right-angled Artin groups. J. Topol. **2**(3), 442–460 (2009)
40. M.W. Davis, The cohomology of a Coxeter group with group ring coefficients. Duke Math. J. **91**(2), 297–314 (1998)
41. P. de la Harpe, *Topics in Geometric Group Theory*. Chicago Lectures in Mathematics (University of Chicago Press, Chicago, IL, 2000)
42. A. Díaz, A. Ruiz, A. Viruel, Cohomological uniqueness of some p-groups. Proc. Edinb. Math. Soc. **56**(2), 449–468 (2013)
43. R. Diestel, *Graph Theory*, 5th edn. Graduate Texts in Mathematics, vol. 173 (Springer, Berlin, 2017)
44. C. Droms, Isomorphisms of graph groups. Proc. Am. Math. Soc. **100**(3), 407–408 (1987)

45. C. Droms, Subgroups of graph groups. J. Algebra **110**(2), 519–522 (1987)
46. G. Duchamp, D. Krob, The lower central series of the free partially commutative group. Semigroup Forum **45**(3), 385–394 (1992)
47. B. Farb, D. Margalit, *A Primer on Mapping Class Groups*. Princeton Mathematical Series, vol. 49 (Princeton University Press, Princeton, NJ, 2012)
48. A. Fathi, F. Laudenbach, V. Poénaru (eds.), *Travaux de Thurston sur les surfaces* (Société Mathématique de France, Paris, 1991). Reprint of *Travaux de Thurston sur les surfaces* (Soc. Math. France, Paris, 1979) [MR0568308 (82m:57003)]. Astérisque No. 66–67 (1991)
49. B. Fine, D. Kahrobaei, G. Rosenberger (eds.), *Computational and Combinatorial Group Theory and Cryptography*. Contemporary Mathematics, vol. 582 (American Mathematical Society, Providence, RI, 2012)
50. R. Flores, D. Kahrobaei, Cryptography with right-angled Artin groups (2016). Preprint, arXiv:1610.06495
51. R. Flores, D. Kahrobaei, T. Koberda, Algorithmic problems in right-angled Artin groups: complexity and applications. J. Algebra **519**, 111–129 (2019)
52. R. Flores, D. Kahrobaei, T. Koberda, A cryptographic application of the Thurston norm. Int. J. Comput. Math. Comput. Syst. Theory **5**(1), 15–24 (2020)
53. R. Flores, D. Kahrobaei, T. Koberda, Expanders and right-angled Artin groups. J. Topol. Anal. (2020). to appear
54. R. Flores, D. Kahrobaei, T. Koberda, An algebraic characterization of k–colorability. Proc. Am. Math. Soc. **149**(5), 2249–2255 (2021)
55. R. Flores, D. Kahrobaei, T. Koberda, Hamiltonicity via cohomology of right-angled Artin groups. Linear Algebra Appl. **631**, 94–110 (2021).
56. M.R. Garey, D.S. Johnson, *Computers and Intractability* (W. H. Freeman, San Francisco, CA, 1979). A Guide to the Theory of NP-Completeness, A Series of Books in the Mathematical Sciences
57. É. Ghys, P. de la Harpe (eds.), *Sur les groupes hyperboliques d'après Mikhael Gromov*. Progress in Mathematics, vol. 83 (Birkhäuser, Boston, MA, 1990). Papers from the Swiss Seminar on Hyperbolic Groups held in Bern, 1988
58. C. Godsil, G. Royle, *Algebraic Graph Theory*. Graduate Texts in Mathematics, vol. 207 (Springer, New York, 2001)
59. O. Goldreich, R. Impagliazzo, L. Levin, R. Venkatesan, D. Zuckerman, Security preserving amplification of hardness, in *31st Annual Symposium on Foundations of Computer Science, Vol. I, II (St. Louis, MO, 1990)* (IEEE Comput. Soc. Press, Los Alamitos, CA, 1990), pp. 318–326
60. O. Goldreich, S. Micali, A. Wigderson, Proofs that yield nothing but their validity, or All languages in NP have zero-knowledge proof systems. J. Assoc. Comput. Mach. **38**(3), 691–729 (1991)
61. M. Gromov, Hyperbolic groups, in *Essays in Group Theory*. Math. Sci. Res. Inst. Publ., vol. 8 (Springer, New York, 1987), pp. 75–263
62. M. Gromov, *Metric Structures for Riemannian and Non-Riemannian Spaces*. Progress in Mathematics, vol. 152 (Birkhäuser, Boston, MA, 1999). Based on the 1981 French original, With appendices by M. Katz, P. Pansu, S. Semmes, Translated from the French by Sean Michael Bates
63. M.F. Hagen, Weak hyperbolicity of cube complexes and quasi-arboreal groups. J. Topol. **7**(2), 385–418 (2014)
64. A. Hatcher, *Algebraic Topology* (Cambridge University Press, Cambridge, 2002)
65. S. Hermiller, J. Meier, Algorithms and geometry for graph products of groups. J. Algebra **171**(1), 230–257 (1995)
66. S. Hermiller, Z. Šunić, Poly-free constructions for right-angled Artin groups. J. Group Theory **10**(1), 117–138 (2007)
67. S. Hoory, N. Linial, A. Wigderson, Expander graphs and their applications. Bull. Am. Math. Soc. (N.S.) **43**(4), 439–561 (2006)

68. J. Huang, Quasi-isometry classification of right-angled Artin groups II: several infinite out cases (2016). Preprint, arXiv:1603.02372
69. J. Huang, Quasi-isometric classification of right-angled Artin groups I: the finite out case. Geom. Topol. **21**(6), 3467–3537 (2017)
70. J. Huang, Commensurability of groups quasi-isometric to RAAGs. Invent. Math. **213**(3), 1179–1247 (2018)
71. J. Huang, B. Kleiner, Groups quasi-isometric to right-angled Artin groups. Duke Math. J. **167**(3), 537–602 (2018)
72. N.V. Ivanov, Automorphism of complexes of curves and of Teichmüller spaces. Int. Math. Res. Notices **14**, 651–666 (1997)
73. C. Jensen, J. Meier, The cohomology of right-angled Artin groups with group ring coefficients. Bull. Lond. Math. Soc. **37**(5), 711–718 (2005)
74. H.A. Jung, On a class of posets and the corresponding comparability graphs. J. Combinatorial Theory Ser. B **24**(2), 125–133 (1978)
75. D. Kahrobaei, V. Shpilrain (eds.), *Algorithmic Problems of Group Theory, Their Complexity, and Applications to Cryptography*. Contemporary Mathematics, vol. 633 (American Mathematical Society, Providence, RI, 2015)
76. M. Kambites, On commuting elements and embeddings of graph groups and monoids. Proc. Edinb. Math. Soc. (2) **52**(1), 155–170 (2009)
77. S. Kim, Co-contractions of graphs and right-angled Artin groups. Algebr. Geom. Topol. **8**(2), 849–868 (2008)
78. S. Kim, T. Koberda, Embedability between right-angled Artin groups. Geom. Topol. **17**(1), 493–530 (2013)
79. S. Kim, T. Koberda, The geometry of the curve graph of a right-angled Artin group. Int. J. Algebra Comput. **24**(2), 121–169 (2014)
80. S. Kim, T. Koberda, An obstruction to embedding right-angled Artin groups in mapping class groups. Int. Math. Res. Not. **14**, 3912–3918 (2014)
81. S. Kim, T. Koberda, Anti-trees and right-angled Artin subgroups of braid groups. Geom. Topol. **19**(6), 3289–3306 (2015)
82. S. Kim, T. Koberda, Right-angled Artin groups and finite subgraphs of curve graphs. Osaka J. Math. **53**(3), 705–716 (2016)
83. S. Kim, T. Koberda, Free products and the algebraic structure of diffeomorphism groups. J. Topol. **11**(4), 1054–1076 (2018)
84. K.H. Ko, S.J. Lee, J.H. Cheon, J.W. Han, J. Kang, C. Park, *New Public-Key Cryptosystem Using Braid Groups*. Advances in cryptology—CRYPTO 2000 (Santa Barbara, CA), Lecture Notes in Comput. Sci., vol. 1880 (Springer, Berlin, 2000), pp. 166–183
85. T. Koberda, Right-angled Artin groups and a generalized isomorphism problem for finitely generated subgroups of mapping class groups. Geom. Funct. Anal. **22**(6), 1541–1590 (2012)
86. T. Koberda, What is...an acylindrical group action?. Notices Am. Math. Soc. **65**(1), 31–34 (2018)
87. T. Koberda, Actions of right-angled Artin groups in low dimensions, in *Handbook of Group Actions*, ed. by L. Ji, A. Papadopoulos, S.-T. Yau, Chap 6, vol. V. ALM 48 (International Press, Boston; Higher Education Press, Beijing, 2020), pp. 227–260
88. T. Koberda, J. Mangahas, S.J. Taylor, The geometry of purely loxodromic subgroups of right-angled Artin groups. Trans. Am. Math. Soc. **369**(11), 8179–8208 (2017)
89. E. Kowalski, *An Introduction to Expander Graphs*. Cours Spécialisés [Specialized Courses], vol. 26 (Société Mathématique de France, Paris, 2019)
90. M.R. Laurence, A generating set for the automorphism group of a graph group. J. Lond. Math. Soc. (2) **52**(2), 318–334 (1995)
91. C.J. Leininger, D. Margalit, Abstract commensurators of braid groups. J. Algebra **299**(2), 447–455 (2006)
92. A. Lubotzky, *Discrete Groups, Expanding Graphs and Invariant Measures*. Modern Birkhäuser Classics (Birkhäuser, Basel, 2010). With an appendix by J. D. Rogawski, Reprint of the 1994 edition

93. A. Lubotzky, E. Zelmanov, Dimension expanders. J. Algebra **319**(2), 730–738 (2008)
94. A. Lubotzky, R. Phillips, P. Sarnak, Ramanujan graphs. Combinatorica **8**(3), 261–277 (1988)
95. A. Margolis, Quasi-isometry classification of right-angled Artin groups that split over cyclic subgroups. Groups Geom. Dyn. **14**(4), 1351–1417 (2020)
96. H.A. Masur, Y.N. Minsky, Geometry of the complex of curves. I. Hyperbolicity. Invent. Math. **138**(1), 103–149 (1999)
97. K.A. Mihaĭlova, The occurrence problem for direct products of groups. Dokl. Akad. Nauk SSSR **119**, 1103–1105 (1958)
98. M.L. Minsky, *Computation: Finite and Infinite Machines* (Prentice-Hall, Englewood Cliffs, NJ, 1967). Prentice-Hall Series in Automatic Computation.
99. A. Myasnikov, V. Shpilrain, A. Ushakov, *Group-Based Cryptography*. Advanced Courses in Mathematics (CRM Barcelona, Birkhäuser, Basel, 2008)
100. A. Rosen, *Concurrent Zero-Knowledge: With Additional Background and a Foreword by Oded Goldreich (Information Security and Cryptography)* (Springer, Berlin, 2006)
101. A. Ruiz, A. Viruel, Cohomological uniqueness, Massey products and the modular isomorphism problem for 2 -groups of maximal nilpotency class. Trans. Am. Math. Soc. **365**(7), 3729–3751 (2013)
102. L. Sabalka, On rigidity and the isomorphism problem for tree braid groups. Groups Geom. Dyn. **3**(3), 469–523 (2009)
103. D. Seinsche, On a property of the class of n-colorable graphs. J. Comb. Theory Ser. B **16**, 191–193 (1974)
104. Z. Sela, Acylindrical accessibility for groups. Invent. Math. **129**(3), 527–565 (1997)
105. J.-P. Serre, *Arbres, amalgames*, SL_2 (Société Mathématique de France, Paris, 1977). Avec un sommaire anglais, Rédigé avec la collaboration de Hyman Bass, Astérisque, No. 46
106. H. Servatius, Automorphisms of graph groups. J. Algebra **126**(1), 34–60 (1989)
107. H. Servatius, C. Droms, B. Servatius, Surface subgroups of graph groups. Proc. Am. Math. Soc. **106**(3), 573–578 (1989)
108. K.J. Shackleton, Combinatorial rigidity in curve complexes and mapping class groups. Pac. J. Math. **230**(1), 217–232 (2007)
109. J. Stallings, On torsion-free groups with infinitely many ends. Ann. Math. (2) **88**, 312–334 (1968)
110. J. Stallings, *Group Theory and Three-Dimensional Manifolds*. A James K. Whittemore Lecture in Mathematics given at Yale University, 1969, Yale Mathematical Monographs, 4 (Yale University Press, New Haven, 1971)
111. D.P. Sumner, Dacey graphs. J. Aust. Math. Soc. **18**, 492–502 (1974)
112. W.P. Thurston, On the geometry and dynamics of diffeomorphisms of surfaces. Bull. Am. Math. Soc. (N.S.) **19**(2), 417–431 (1988)
113. E. Toinet, Conjugacy p-separability of right-angled Artin groups and applications. Groups Geom. Dyn. **7**(3), 751–790 (2013)
114. L. VanWyk, Graph groups are biautomatic. J. Pure Appl. Algebra **94**(3), 341–352 (1994)
115. A. Viruel, Homotopy uniqueness of $B\mathrm{G}_2$. Manuscripta Math. **95**(4), 471–497 (1998)
116. A. Viruel, Mod 3 homotopy uniqueness of BF_4. J. Math. Kyoto Univ. **41**(4), 769–793 (2001)
117. D.T. Wise, Research announcement: the structure of groups with a quasiconvex hierarchy. Electron. Res. Announc. Math. Sci. **16**, 44–55 (2009)
118. D.T. Wise, The structure of groups with a quasiconvex hierarchy (2011)
119. D.T. Wise, From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry, in *CBMS Regional Conference Series in Mathematics*, vol. 117 (Published for the Conference Board of the Mathematical Sciences, Washington, DC, 2012)
120. C. Wrathall, The word problem for free partially commutative groups. J. Symb. Comput. **6**(1), 99–104 (1988)