

Biautomatic structures in systolic Artin groups

Newcastle 2nd July 2018

1 Introduction

The work reported in this article was motivated by a recent article of Huang and Osajda [13] proving the biautomaticity of a class of Artin groups, namely those of almost large type, which contains all Artin groups of large type. Previously biautomaticity had been proved by Brady and McCammond for many, but not all, Artin groups of large type, specifically for those of extra-large type [18], all 3-generated examples, and all others for which the associated Coxeter diagram can be oriented in such a way as to exclude certain oriented subdiagrams [1]. Biautomaticity had also been proved for all Artin groups of finite type, by Charney [2, 3]; for these groups (which are Garside groups), symmetric, geodesic biautomatic structures were found, over the Garside generators.

We were particularly interested in the results of [13] since we already knew that all Artin groups of almost large type (in fact all in the slightly larger class of sufficiently large type) had shortlex automatic structures over their standard generating sets [10, 11]; but these structures are not, in general, biautomatic. Our methods were combinatorial, based on rewriting.

The methods used in [1] and [2, 3] to prove biautomaticity for those two types of Artin groups are quite distinct from each other. Brady and McCammond's approach is geometric, relying on the construction of a piecewise Euclidean non-positively curved 2-complex on which the Artin group acts discretely and fixed point freely, followed by the application of results of Gersten and Short [7, 8]. But Charney's approach uses the Garside structure of Artin groups of finite type.

Huang and Osajda's article [13] proves biautomaticity of an Artin group of almost large type via the construction of a systolic complex on which the group acts simplicially, properly discontinuously and cocompactly; hence by definition, the group is systolic. The biautomaticity of these Artin groups

then follows immediately from [14, Theorem E], that all systolic groups are biautomatic. We observe that, in fact, the universal covers of the 2-complexes of [1] are systolic, although the construction of those 2-complexes does not seem to follow the same pattern as that of the systolic complexes in [13]; hence in some sense the results of [13] generalise those of [1]. Our aim in this work has been to examine the construction of [13] and see what we can learn from it.

To a large extent our approach has been experimental, through computation and examination of examples. In particular, we have written programs in GAP that construct the biautomatic structures arising from the systolic complexes that are acted on by Artin groups of almost large type in which all edge labels in the associated Coxeter diagram are either at most 4 or equal to infinity.

In order to describe and analyse the construction of [13] we need a variety of background material on automatic and biautomatic groups, on Artin groups, and on systolic complexes, and we have tried to give sufficient detail to make this article accessible to a range of readers.

SR: added
this short
paragraph
of explana-
tion

After this introductory section, Section 2 introduces the concepts and notations we shall use from group theory. In Section 2.1 we define automatic and biautomatic structure for groups, and introduce Artin groups. while Section 2.2 explains briefly how group actions may be used to deduce such structures.

Section 3 introduces the concepts and notation of systolic complexes and systolic groups, and is based on [14]. The basic concepts are introduced in Section 3.1, then Section 3.2 explains the concepts of directed and allowable geodesics that are vital to the construction of biautomatic structures in systolic groups. Section 4 is devoted to a description of the construction of biautomatic structures for systolic groups, again with reference to [14]. Section 4.1 describes generating sets \mathcal{B} and \mathcal{A} over which biautomatic structures will be defined, and descriptions of those biautomatic structures, and of the automata that define them, are given in the subsequent sections. In fact we introduce two slightly different biautomatic structures, over the two alphabets \mathcal{B} and \mathcal{A} . The language \mathcal{L}' over \mathcal{B} is the structure of [14], but we prefer in this article to work the languages \mathcal{L} and \mathcal{L}^0 , over \mathcal{A} , which are closely related to \mathcal{L}' .

Huang and Osajda's construction of a systolic complex \mathcal{X} for an Artin group G of almost large type is described briefly at the beginning of Section 5. The complex \mathcal{X} is formed out of the Cayley graph \mathcal{G} for G by replacing subgraphs \mathcal{G}_{ij} , corresponding to 2-generator subgroups G_{ij} of G , by systolic complexes \mathcal{X}_{ij} , and the construction of those subcomplexes is described in Section 5.2.

We illustrate the construction of \mathcal{X} by computing some examples in Section 6. We describe the complexes for the right-angled Artin groups \mathbb{Z}^2 , $\mathbb{Z}^2 * \mathbb{Z}$ and $\mathbb{F}_2 \times \mathbb{Z}$ in Sections 6.1, 6.2, 6.3, and the 3-string braid group $G(A_2)$ in Section 6.4.

Section 7 is devoted to the identification of some properties found in the biautomatic structure of any Artin group of almost large type. Here, we may assume that the Artin group is non-free; for a free group in its natural presentation, the associated systolic complex is simply the Cayley graph over the standard generating set. Corollary 7.3 shows that, given a sensible selection of the orbit representatives upon which the language \mathcal{L}^0 depends, the generating set \mathcal{A} (as well as the related set \mathcal{B}) consists precisely of the union of inverse closed Garside generating sets for the 2-generated Artin subgroups G_{ij} of G , together with a symbol representing the identity element. Subsequent results examine properties of the language \mathcal{L}^0 over \mathcal{A} . It is clear from our examples that \mathcal{L}^0 is not, in general, geodesic over \mathcal{A} . But we prove in Proposition 7.4 that the \mathcal{A} -length of a word in \mathcal{L}^0 is no greater than its geodesic length over the standard Artin generating set X , and no less than $1/\max\{2, M - 2\}$ of that geodesic length, where M is the maximum of all finite $m_{i,j}$. We observe too, from our examples, that the language \mathcal{L}^0 is not in general prefix closed, but from a result about extensions of directed geodesics that we prove in Proposition 7.5 we can deduce Corollary 7.6, which associates to each prefix of a word in \mathcal{L}^0 a word in \mathcal{L}^0 differing from it by a bounded amount at its end.

Finally, Section 8 describes our development of algorithms that construct systolic complexes and then biautomatic structures for those systolic Artin groups that we were able to handle.

2 Concepts and notation from group theory

2.1 Introducing biautomatic groups and Artin groups

All the groups considered in this article will be finitely presented. Let $G = \langle X \mid R \rangle$ be such a group. We define its *Cayley graph* $\mathcal{G} = \mathcal{G}(G, X)$ to be the graph with vertex set G and, for each $x \in X$, directed edges labelled x and x^{-1} from g to gx and from gx to g , respectively. These two directed edges have the same underlying undirected edge.

A finitely presented group $G = \langle X \mid R \rangle$ is *automatic* if there exists

- (1) a regular set L of words over X providing a complete set of representatives of the elements of G , and

(2) an associated integer k ,

satisfying the following condition:

whenever words u, v in L represent elements $g, h \in G$ satisfying $ga =_G h$ with $a \in X \cup X^{-1} \cup \{1\}$, and $\gamma_1(u), \gamma_1(v)$ are the paths in $\mathcal{G}(G, X)$ that are traced out by u, v from the identity vertex, then $\gamma_1(u)$ and $\gamma_1(v)$ fellow travel at distance k .

By definition, a set of words is *regular* if it is the language of a finite state automaton. We say that two paths in a graph *fellow travel* at distance k if, for any i , the distance in the graph between the vertices at distance i from the initial vertex of each path is at most k .

An automatic group G is *biautomatic* if there exist L, k as above satisfying the following additional condition:

whenever words $u, v \in L$ represent elements g and $h \in G$ satisfying $ah =_G g$ with $a \in X \cup X^{-1}$, and $\gamma_1(u)$ and $\gamma_a(v)$ are the paths in $\mathcal{G}(G, X)$ that are traced out by u from the identity vertex and by v from a , then $\gamma_1(u)$ and $\gamma_a(v)$ fellow travel at distance k .

We refer to [4, 12] for more detail. But we note in particular that the properties of automaticity and biautomaticity are independent of the choice of finite generating set; that is, if a group G has an automatic or biautomatic structure over some finite generating set X then it has a corresponding structure over any other finite generating set. We note also that any hyperbolic group is biautomatic; over any generating set the set of all geodesic words is the language of a biautomatic structure [19].

SR: moved and abbreviated reference to Paspoglu

The language L of an automatic structure for a group G is called *symmetric* if whenever $w \in L$ then $w^{-1} \in L$; when L is symmetric then the structure satisfies both fellow traveller conditions and so G is biautomatic. L is called *geodesic* if every word in L is a minimal length representative of the group element it represents. L is called *prefix closed* if whenever $w \in L$ then every prefix of w is also in L .

In this article we consider biautomatic structures for certain types of Artin groups. The standard presentation for an Artin group over its standard generating set $X = \{a_1, \dots, a_n\}$ is as

$$\langle a_1, \dots, a_n \mid m_{ij}(a_i, a_j) = m_{ji}(a_j, a_i) \text{ for each } i \neq j \rangle,$$

where the integers m_{ij} are the entries in a Coxeter matrix (a symmetric $n \times n$ matrix (m_{ij}) with entries in $\mathbb{N} \cup \{\infty\}$, where $m_{ii} = 1$, and $m_{ij} \geq 2$, $\forall i \neq j$), and where for generators a, a' and $m \in \mathbb{N}$ we define ${}_m(a, a')$ to be the word that is a product of m alternating a 's and a' 's and starts with a . Adding the relations $a_i^2 = 1$ to those for the Artin group defines the associated Coxeter group, which is more commonly presented as

$$\langle a_1, \dots, a_n \mid (a_i a_j)^{m_{ij}} = 1 \text{ for each } i, j \rangle.$$

The standard presentation of an Artin group can be described using an associated *Coxeter diagram* Γ , with n nodes, where nodes corresponding to a_i and a_j are joined by an edge labelled by m_{ij} .

An Artin group $G(\Gamma)$ is said to be of *spherical* or *finite* type if the associated Coxeter group is finite, of *dihedral type* if the associated Coxeter group is dihedral (or, equivalently, the standard generator set has two elements), and of *large type* if $m_{ij} \geq 3$ for all $i \neq j$. The group $G(\Gamma)$ is defined in [13] to have *almost large type* if (1) in any triangle of edges of Γ either at least one edge is labelled ∞ or no edge has label 2, and (2) in any square of edges of Γ either at least one edge is labelled ∞ or at most one edge has label 2. So, in particular, all Artin groups of large type are of almost large type.

Artin groups of various types are proved automatic [2, 18, 1, 21, 9, 10, 11], but before the results of [13] biautomatic structures were only known to exist for all those of finite type, as well as for many (but not all) those of large type. Automatic structures for which the associated language L is a set of shortlex geodesics over the standard generating set were already known for all almost large Artin groups and, slightly more generally, for all *sufficiently large* Artin groups [11]).

2.2 Using group actions to derive automatic structures

We recall a basic tenet of geometric group theory commonly referred to as the Milnor-Švarc lemma: if a group $G = G(X)$ has a ‘nice’ (properly discontinuous and compact) discrete, isometric action on a metric space \mathcal{X} , then its Cayley graph $\mathcal{G} = \mathcal{G}(G, X)$ is quasi-isometric to \mathcal{X} . Specifically, for any vertex $x_0 \in \mathcal{X}$, there is a quasi-isometry from \mathcal{G} to \mathcal{X} that maps the identity vertex of \mathcal{G} to x_0 , and each vertex g of \mathcal{G} to the image gx_0 in \mathcal{X} of x_0 under the action of g .

This is at the basis of results such as those that follow, of Gersten and Short, and of Niblo and Reeves.

Theorem 2.1 (Gersten&Short, 1990,1991 [7, 8]). *The fundamental groups of piecewise Euclidean 2-complexes of types $A_1 \times A_1$, A_2 , B_2 and G_2 (cor-*

responding to tessellations of the Euclidean plane by squares, equilateral triangles, and triangles with angles $(\pi/2, \pi/4, \pi/4)$, $(\pi/2, \pi/3, \pi/6)$ are biautomatic.

This result is stated in terms of automatic rather than biautomatic structures in [8, Theorem1], but it is made clear in the introduction of that paper that the construction gives biautomatic structures.

It follows from Gersten and Short's results that a group acting discretely and fixed point freely on a complex of any of the types listed must be biautomatic.

Theorem 2.2 (Niblo&Reeves 1998, Theorem 5.3 of [16]). *If \mathcal{X} is a simply connected and non-positively curved cube complex, and G acts effectively, cellularly, properly discontinuously and cocompactly on \mathcal{X} , then G is biautomatic.*

A cube complex is a metric polyhedral complex in which each cell is isometric to a Euclidean cube and the glueing maps are isometries. It is proved in [17] that any finite rank Coxeter group G acts properly discontinuously by isometries on an appropriate cube complex; in the many cases where the action is cocompact, the Coxeter group is proved biautomatic.

The results of [7, 8] are used in the proof of the following result.

Theorem 2.3 (Brady&McCammond, 2000 [1]). *Various Artin groups of large type, including all that are 3-generated, act appropriately on piecewise Euclidean non-positively curved (i.e. locally $CAT(0)$) 2-complexes of types A_2 or B_2 , and hence are biautomatic.*

The 2-complexes constructed in the proof of this theorem are simplicial, made by attaching angles and lengths to presentation complexes for non-standard presentations for the Artin groups, in which all relators have length 3. It can be seen that their universal covers are systolic [14], as defined below.

3 Systolic complexes

3.1 Basic concepts

We need the definition of a systolic complex, and related notation, which we shall use throughout this article. In this we generally follow [14]. Suppose that \mathcal{X} is a simplicial complex, and $\mathcal{S} = \mathcal{S}(\mathcal{X})$ its set of simplices (the

vertices of its ‘barycentric subdivision’). We denote by \mathcal{X}_1 the 1-skeleton of \mathcal{X} , that is, the graph on its vertices whose edges are the 1-cells. The set \mathcal{S} itself has the structure of a simplicial complex, whose simplices are sets of simplices of \mathcal{X} that can be ordered into an ascending chain of simplices that are related by inclusion; in particular its set $E(\mathcal{S})$ of 1-cells (edges) is the set of pairs $\{\rho, \sigma\}$ of distinct simplices of \mathcal{X} for which either $\sigma \subset \rho$ or $\rho \subset \sigma$. We denote the 1-skeleton of \mathcal{S} by \mathcal{S}_1 .

We define a *cycle* in \mathcal{X} to be a 1-dimensional subcomplex of \mathcal{X} consisting of the vertices and edges of a closed path in \mathcal{X}_1 in which all edges are distinct and only the first and last vertices coincide; the length of the cycle is defined to be the number of 1-cells in it. We call a cycle a *full cycle* if it is full as a subcomplex, that is, if any simplex spanned by a subset of its vertices is within the cycle. This means that no two non-adjacent vertices within the cycle are joined by an edge (or rather, there are no chords between non-adjacent vertices of the cycle). Note that, since a cycle is 1-dimensional, a full cycle in a simply connected complex cannot have length 3. We define the *systole* of \mathcal{X} , $\text{sys}(\mathcal{X})$, to be the minimum length of a full cycle.

We write \mathcal{X}_σ for the *link* in \mathcal{X} of a simplex σ (the set of simplices τ such that τ is disjoint from σ , and such that τ and σ span a simplex $\sigma * \tau$ of \mathcal{X}), and for any subcomplex \mathcal{Y} of \mathcal{X} containing σ , we write $\text{Res}(\sigma, \mathcal{Y})$ for the *residue* of σ in \mathcal{Y} (the union of all closed simplices that contain σ ; this is sometimes called the *closed star* of σ in \mathcal{Y}). We define the 1-ball $B_1(\sigma, \mathcal{Y})$ of σ in \mathcal{Y} to be the union of all closed simplices within \mathcal{Y} that intersect σ non-trivially.

We say that \mathcal{X} is κ -large if $\text{sys}(\mathcal{X}) \geq \kappa$ and also $\text{sys}(\mathcal{X}_\sigma) \geq \kappa$ for every simplex σ , *locally κ -large* if the residue of every simplex of \mathcal{X} is κ -large, κ -systolic if it is connected, simply connected and locally κ -large. We call a group κ -systolic if it acts properly discontinuously and cocompactly by automorphisms on a κ -systolic complex. (Proper discontinuity means that stabilisers are finite.) We abbreviate 6-systolic as *systolic*. By [14] κ -systolic complexes with $\kappa \geq 6$ are κ -large (but this isn’t true for $\kappa = 4, 5$). And (locally) κ -large implies (locally) m -large for $\kappa \geq m$.

We note the following results from [14]

Theorem 3.1 (Januszkiewicz&Swiatkowski, Theorem A [14]). *The 1-skeleton of a 7-systolic complex is hyperbolic; any 7-systolic group is hyperbolic.*

Theorem 3.2 (Januszkiewicz&Swiatkowski, Theorem E [14]). *Any systolic group is biautomatic.*

In order to prove a group G biautomatic, we need to define a regular language L of words (over some selected generating set) that maps onto the group and has appropriate fellow travelling properties. When G acts appropriately

on a systolic complex \mathcal{X} , we fix a vertex (0-cell) x_0 of \mathcal{X} . Then for each element $g \in G$ we need to select a path (or paths) within the complex \mathcal{X} from x_0 to gx_0 , and then associate with each such path a word over the selected generated set that represents g . It turns out that directed geodesics between pairs of vertices x_0, gx_0 have some very good properties; we use the Milnor-Švarc lemma to translate those properties into properties for the words in our selected language.

3.2 Directed and allowable geodesics

Suppose that \mathcal{X} is a locally 6-large (but not necessarily systolic) simplicial complex. We call a sequence of simplices $\sigma_0, \dots, \sigma_n$ from $\mathcal{S}(\mathcal{X})$ a *directed geodesic* from σ_0 to σ_n if

- (1) for all i with $0 \leq i \leq n-1$, σ_i and σ_{i+1} are disjoint and span a simplex $\sigma_i * \sigma_{i+1}$ of \mathcal{X} , and
- (2) for all i with $0 \leq i \leq n-2$, $\text{Res}(\sigma_i, \mathcal{X}_{\sigma_{i+1}}) \cap B_1(\sigma_{i+2}, \mathcal{X}_{\sigma_{i+1}}) = \emptyset$.

We note that

Lemma 3.3. (1) *Every subpath of a directed geodesic is a directed geodesic,*
(2) *If $\sigma_0, \dots, \sigma_n$ is a directed geodesic, and v is a vertex of the simplex σ_n , then $\sigma_0, \dots, \sigma_{n-1}, v$ is a directed geodesic.*

Proof. It follows straight from the definition that a sequence of simplices is a directed geodesic if and only if the same is true of every subsequence of length at most 3, and so (1) is immediate.

(2) is immediate once we observe that if $\sigma_n \in \mathcal{X}_{\sigma_{n-1}}$ and v is a vertex of σ_n , then $B_1(v, \mathcal{X}_{\sigma_{n-1}}) \subseteq B_1(\sigma_n, \mathcal{X}_{\sigma_{n-1}})$. \square

Note that the reverse of a directed geodesic need not be a directed geodesic.

If \mathcal{X} is systolic, then for any pair of vertices v, w of \mathcal{X} , there is a unique directed geodesic $\gamma(v, w) = \sigma_0, \sigma_1, \dots, \sigma_n$ from $v = \sigma_0$ to $w = \sigma_n$ [14, Lemma 9.7], and it is geodesic, that is $d_{\mathcal{X}_1}(v, w) = n$ [14, Corollary 9.8]. That directed geodesic is the *projection ray* from v to w (as defined in [14]); it is proved in [14, Lemma 9.3] that every projection ray in a systolic complex is a directed geodesic, and in [14, Proposition 9.6] that every directed geodesic $\sigma_0, \dots, \sigma_n$ is a projection ray on its final simplex σ_n .

Given vertices v, w , we call an *allowable geodesic* [14, Section 11] an infinite sequence of vertices u_0, u_1, \dots such that if $\sigma_0, \sigma_1, \dots, \sigma_n$ is the directed geodesic from $v = \sigma_0$ to $w = \sigma_n$, then $u_0 = v, u_i = w$ for $i \geq n$,

and $u_i \in \sigma_i$ for $0 < i < n$. Such a sequence u_0, u_1, \dots, u_n is a geodesic path in the 1-skeleton of \mathcal{X} [14, Fact 11.1]. It is proved in [14, Prop 11.2] that a pair of allowable geodesics joining (respectively) vertex v to vertex w and vertex p to vertex q must follow travel at distance at most $3 \max\{\text{dist}_{\mathcal{X}_1}(v, p), \text{dist}_{\mathcal{X}_1}(w, q)\} + 1$.

From a directed geodesic $\gamma = \gamma(v, w)$ as above, we can define an associated *polygonal path*

$$\hat{\gamma} = \hat{\gamma}(v, w) = \sigma_0, \sigma_0 * \sigma_1, \sigma_1, \sigma_1 * \sigma_2, \sigma_2, \dots, \sigma_{n-1}, \sigma_{n-1} * \sigma_n, \sigma_n,$$

and note that any two adjacent simplices in a polygonal path are related by (alternately) inclusion or reverse inclusion. So a polygonal path is a path within the graph \mathcal{S}_1 , whereas consecutive simplices on a directed geodesic are at distance 2 as vertices of \mathcal{S}_1 .

4 Biautomatic structures for G

4.1 Alternate generating sets \mathcal{B} and \mathcal{A} for G

Now suppose that G is a systolic group associated with the systolic complex \mathcal{X} . Let K be the set of orbits of the action of G on $\mathcal{S}(\mathcal{X})$, and let V_0 be a set of representatives of the orbits. Choose $v_0 \in V_0$ to be a 0-cell. The 0-cell v_0 is a natural choice for x_0 , which we use as we apply the Milnor-Švarc lemma to embed G in \mathcal{X} as described in Section 2.2.

For $\sigma \in \mathcal{S}(\mathcal{X})$, let $\bar{\sigma}$ denote its representative in V_0 . For $\sigma \in \mathcal{S}(\mathcal{X})$, define $\Lambda_\sigma := \{g \in G : \sigma = g\bar{\sigma}\}$, the ‘set of labels of σ ’; note that, since we have $v_0 \in V_0$, we have $\Lambda_{v_0} = G_{v_0}$. Then for any pair (σ, τ) of simplices from $\mathcal{S}(\mathcal{X})$, define $\Lambda_{\sigma, \tau} := \Lambda_\sigma^{-1} \Lambda_\tau$, the ‘set of labels of (σ, τ) ’; Note that if stabilisers are all trivial, then the sets Λ_σ and $\Lambda_{\sigma, \tau}$ are all singletons. Note also that, for any $g \in G$, $\Lambda_{g\sigma, g\tau} = \Lambda_{\sigma, \tau}$, that $G_{v_0} \Lambda_{v_0, \sigma} = \Lambda_{v_0, \sigma}$, and that for all ρ, σ, τ , $\Lambda_{\rho, \sigma} \Lambda_{\sigma, \tau} = \Lambda_{\rho, \tau}$. We define an alphabet \mathcal{B} , as in [14], to be a (finite) set of symbols representing the union of G_{v_0} together with all the sets $\Lambda_{\sigma, \tau}$ for which σ and τ are simplices related by inclusion or reverse inclusion. Note that \mathcal{B} includes a symbol representing the identity element 1.

We also define an alphabet \mathcal{A} , a set of symbols representing the union of G_{v_0} with all the sets $\Lambda_{\rho, \tau}$ for which ρ, τ are disjoint simplices that together span a simplex.

We will define biautomatic structures for G over each of the two alphabets. The first, which we call \mathcal{L}' , over the alphabet \mathcal{B} , and associated with polygonal paths, is the structure described in [14]. We prefer the second, \mathcal{L} , over

the alphabet \mathcal{A} , and associated with directed geodesics. It is closely associated with \mathcal{L}' , but the words in it (over \mathcal{A}) are generally shorter than those in \mathcal{L}' (over \mathcal{B}). When the subgroup G_{v_0} is trivial, we can simplify \mathcal{L} further by deleting the first symbol in every word; we shall call that simplified language \mathcal{L}^0 .

4.2 A structure over \mathcal{B}

The following biautomatic structure over \mathcal{B} is (essentially) described both in [14] and in [20]. We have made some small adjustments after noticing some differences between the two descriptions; we believe that there is an error in the suggestion in [14] that any occurrences of the symbol 1 in a word should be deleted from it before it is included in the language; for the language of words shortened in this way would not satisfy the fellow traveller conditions.

Given $g \in G$, let $\gamma = \gamma(v_0, gv_0)$ be the unique directed geodesic from v_0 to gv_0 , and $\hat{\gamma} = \hat{\gamma}(v_0, gv_0)$ the corresponding polygonal path. Write $\gamma = \sigma_0 = v_0, \dots, \sigma_i, \dots, \sigma_n = gv_0$, $\hat{\gamma} = \sigma'_0 = v_0, \dots, \sigma'_j, \dots, \sigma'_{2n} = gv_0$ (so that $\sigma'_{2i} = \sigma_i$). Following [14], we associate to $\hat{\gamma}$ all sequences of elements $g_j \in G$ ($j = 0, \dots, 2n$) with $g_j \in \Lambda_{\sigma'_j}$ for each j (so $g_0 \in G_{v_0}$) and $g_{2n} = g$, and then all words $g_0\nu_1 \cdots \nu_{2n}$ over \mathcal{B} , where $\nu_j \in \Lambda_{\sigma'_{j-1}, \sigma'_j}$ ($j = 1, \dots, 2n$) is defined by $\nu_j := g_{j-1}^{-1}g_j$.

We see that the word $g_0\nu_1 \cdots \nu_{2n}$ represents the element g , and that its proper prefixes represent $g_0, g_1, \dots, g_{2n-1}$. We define \mathcal{L}' to be the set of all such words $g_0\nu_1 \cdots \nu_{2n}$. If all stabilisers are trivial then \mathcal{L}' contains a unique representative of each element of G , but otherwise some elements will admit more than one representative.

That \mathcal{L}' defines a biautomatic structure for G is proved in [14]. Regularity of \mathcal{L}' is verified by explicit construction of an automaton that recognises \mathcal{L}' . Fellow travelling of appropriate pairs of words in \mathcal{L}' is inherited via the Milnor-Švarc lemma from the fellow travelling of related pairs of allowable geodesics.

4.3 A structure over \mathcal{A}

We can form a biautomatic structure for G over the alphabet \mathcal{A} , directly from the directed geodesics, as follows.

Given $g \in G$, again let $\gamma = \gamma(v_0, gv_0) = \sigma_0 = v_0, \sigma_1, \dots, \sigma_n = gv_0$ be the unique directed geodesic from v_0 to gv_0 . Now, associate to γ all sequences of

elements h_i ($i = 0, \dots, n$) with $h_i \in \Lambda_{\sigma_i}$ for each i (so $h_0 \in G_{v_0}$), and $h_n = g$, and then all words $h_0\mu_1 \cdots \mu_n$ over \mathcal{A} , where $\mu_i \in \Lambda_{\sigma_{i-1}, \sigma_i}$ ($i = 1, \dots, n$) is defined by $\mu_i := h_{i-1}^{-1}h_i$; we define the language \mathcal{L} over \mathcal{A} to be the set of all such words.

Again, we see that the word $h_0\mu_1 \cdots \mu_n$ represents the element g , and that its proper prefixes represent h_0, h_1, \dots, h_{n-1} . Clearly, each of the generators μ_j is equal in G to a word of length 2 over \mathcal{B} , and so there is a natural map from \mathcal{L} to \mathcal{L}' , in which generators of \mathcal{A} are replaced by words of length 2 over \mathcal{B} .

Because of the relationship between \mathcal{L}' and \mathcal{L} , it is clear that \mathcal{L} is also a biautomatic structure.

In addition, we note that the word $w = h_0\mu_1 \cdots \mu_n$ representing g has length $n + 1$, while the directed geodesic $\gamma(v_0, gv_0)$ has length n , the same as the length of an allowable geodesic from v_0 to gv_0 .

4.4 Constructing the automaton that recognises \mathcal{L}'

The description of the automaton \mathcal{M}' that recognises \mathcal{L}' is taken from [14, Page 50], with a small amount of adjustment.

We need the notation introduced in Section 4.1. Recall that K denotes the set of orbits $G\sigma$ of the simplices σ in the set $\mathcal{S} = \mathcal{S}(\mathcal{X})$, under the action of G . Recall also that \mathcal{S} itself has the structure of a simplicial complex. We denote by $E(K)$ the set of orbits of G on the set $E(\mathcal{S})$ of edges of \mathcal{S} ; its elements are sets $G\{\rho, \sigma\} := \{g\rho, g\sigma\} : g \in G$, for which $\{\rho, \sigma\} \in E(\mathcal{S})$. Now, as above let $\bar{\rho}, \bar{\sigma}$ be the representatives in the set V_0 of the orbits $G\rho, G\sigma$, and choose $\lambda_\rho \in \Lambda_\rho, \lambda_\sigma \in \Lambda_\sigma$; so $\rho = \lambda_\rho\bar{\rho}, \sigma = \lambda_\sigma\bar{\sigma}$. Then, although it is not necessarily true that $\{\bar{\rho}, \bar{\sigma}\} \in E(\mathcal{S})$, we have $\{\lambda_\rho\bar{\rho}, \lambda_\sigma\bar{\sigma}\} \in E(\mathcal{S})$, and it follows that $\{\lambda^{-1}\bar{\rho}, \bar{\sigma}\}, \{\bar{\rho}, \lambda\bar{\sigma}\} \in E(\mathcal{S})$ for all $\lambda \in \Lambda_{\rho, \sigma}$.

In the automaton \mathcal{M}' , we have a unique start state, which we label v_0 , and for each $h \in G_{v_0}$ we have a state (v_0, h) . In addition, for each element $G\{\rho, \sigma\}$ of $E(K)$, and each $\lambda \in \Lambda_{\rho, \sigma}$, we define states $(G(\rho, \sigma), \lambda)$. For ease of notation, we refer to the state $(G(\rho, \sigma), \lambda)$ as (ρ, σ, λ) , but of course, for any $g \in G$, the triple $(g\rho, g\sigma, \lambda)$ represents that same state. The accept states are $(v_0, 1)$ and all states of the form $(\sigma, v_0, \lambda) = (g\sigma, gv_0, \lambda)$. The arrows are as follows:

- (1) an arrow labelled h from v_0 to (v_0, h) , for each $h \in G_{v_0} = \Lambda_{v_0}$,
- (2) for each $\sigma \in \mathcal{S}$ with $v_0 \subset \sigma$, and each $\lambda \in \Lambda_{v_0, \sigma}$, an arrow labelled λ from (v_0, h) to the state (v_0, σ, λ) ,

(3) an arrow labelled μ from state (ρ, σ, λ) to state (σ, τ, μ) , provided that:

- either (i) $\lambda^{-1}\bar{\rho}$ and $\mu\bar{\tau}$ are disjoint and span $\bar{\sigma}$,
- or (ii) $\bar{\sigma}$ is a proper face of both $\lambda^{-1}\bar{\rho}$ and $\mu\bar{\tau}$ and
 $\text{Res}((\lambda^{-1}\bar{\rho})_{\bar{\sigma}}, \mathcal{X}_{\bar{\sigma}}) \cap B_1((\mu\bar{\tau})_{\bar{\sigma}}, \mathcal{X}_{\bar{\sigma}}) = \emptyset$.

We note that all paths from the start state to accept states have odd length, alternating after the start between states (ρ, σ, λ) with ρ contained in σ and those with σ contained in ρ . We note that all paths of length 1 from the start state lead to states of the form (v_0, h) , all paths of odd length greater than 1 from the start state lead to states of the form (ρ, σ, λ) with $\rho \supset \sigma$, and all paths of even length greater than 0 from the start state lead to states of the form (ρ, σ, λ) with $\rho \subset \sigma$.

4.5 Constructing the automaton that recognises \mathcal{L}

We modify the construction of \mathcal{M}' to construct an automaton \mathcal{M} that recognises \mathcal{L} ; we use just the start state of \mathcal{M}' and those states at odd distance from it in \mathcal{M}' . We label the unique start state v_0 , and have accepting states (v_0, h) , for each $h \in G_{v_0}$, as before. Then for each element $G\{\rho, \sigma\}$ of $E(K)$ with $\rho \supset \sigma$, and each $\lambda \in \Lambda_{\rho, \sigma}$, we define a state (ρ, σ, λ) . As above, if $g \in G$, the triples (ρ, σ, λ) and $(g\rho, g\sigma, \lambda)$ represent the same state. As above, the accepting states are $(v_0, 1)$ and all states of the form $(\sigma, v_0, \lambda) = (g\sigma, gv_0, \lambda)$. The arrows are as follows:

- (1) an arrow labelled h from v_0 to (v_0, h) , for each $h \in G_{v_0} = \Lambda_{v_0}$,
- (2) for each $\rho \in \mathcal{S}$ with $v_0 \subset \rho$, and $\lambda \in \Lambda_{v_0, \rho}$, where $\sigma \subset \rho$ and $\mu \in \Lambda_{\rho, \sigma}$, an edge labelled $\lambda\mu$ from (v_0, h) to the state (ρ, σ, μ) , provided that $\lambda^{-1}v_0$ and $\mu\bar{\sigma}$ are disjoint and span $\bar{\rho}$,
- (3) an arrow labelled $\mu\nu$ from state (ρ, σ, λ) to state (τ, ν, ν) , where $\sigma \subset \tau$ and $\mu \in \Lambda_{\sigma, \tau}$, provided that:
 - (i) the simplices $\mu^{-1}\bar{\sigma}$ and $\nu\bar{\nu}$ are disjoint and span $\bar{\tau}$;
 - (ii) the simplex $\bar{\sigma}$ is a proper face of both $\lambda^{-1}\bar{\rho}$ and $\mu\bar{\tau}$; and
 - (iii) $\text{Res}((\lambda^{-1}\bar{\rho})_{\bar{\sigma}}, \mathcal{X}_{\bar{\sigma}}) \cap B_1((\mu\bar{\tau})_{\bar{\sigma}}, \mathcal{X}_{\bar{\sigma}}) = \emptyset$.

We observe that, if the vertex stabilisers are trivial then the first letter in every element of \mathcal{L} is simply the letter 1, representing the identity element. In this case we might prefer to delete it; we shall denote by \mathcal{L}^0 the language we get from \mathcal{L} by this simple adjustment. Clearly the new language also defines a biautomatic structure. It is elementary to modify the description of \mathcal{L} above, to find an automaton that recognises \mathcal{L}^0 , as follows: we no longer need the state labelled $(v_0, 1)$, or the type 1 transitions from v_0 to

$(v_0, 1)$, the state v_0 becomes an accepting state, and we replace the type 2 transitions by transitions on the same letter, and with the same target, but with source v_0 rather than $(v_0, 1)$.

5 Building systolic complexes and biautomatic structures for Artin groups

5.1 The basic construction

Let G be an Artin group of almost large type, defined in its natural presentation over the standard generating set X . The systolic complex \mathcal{X} of [13] is built out of the Cayley graph \mathcal{G} for G over X ; the vertex v_0 is set to be the vertex of \mathcal{G} that corresponds to the identity element. For each (parabolic) subgroup $G_{ij} = \langle x_i, x_j \rangle$ of G with $i \neq j$ (which is itself a 2-generated Artin group) and $m_{ij} \neq \infty$, we build a systolic complex \mathcal{X}_{ij} out of the Cayley graph \mathcal{G}_{ij} for G_{ij} . This involves adjoining new vertices and edges and also new higher dimensional simplices to \mathcal{G}_{ij} . So \mathcal{G}_{ij} embeds in \mathcal{X}_{ij} , and the action of G_{ij} on \mathcal{X}_{ij} is inherited from its action on \mathcal{G}_{ij} . We describe this construction in more detail for the cases $m_{ij} = 2, 3$ in Sections 6.1 and 6.4 below.

For each vertex v of \mathcal{G} , there is a subgraph of \mathcal{G} isomorphic to \mathcal{G}_{ij} with base point v . We construct the systolic complex \mathcal{X} for G by replacing each of these subgraphs for which $m_{ij} < \infty$ by the corresponding systolic complex \mathcal{X}_{ij} . When we make this replacement we leave unchanged those vertices and edges that already belong to \mathcal{G} . So, for each v and each such pair i, j , we are adjoining a collection of new vertices, edges and higher dimensional simplices to \mathcal{X} . The action of G on \mathcal{G} extends naturally to an action on \mathcal{X} .

Since G acts fixed point freely on \mathcal{G} , in this action of G on the systolic complex \mathcal{X} all simplex stabilisers are trivial. So the subgroup G_{v_0} is the identity subgroup, and the sets Λ_σ and $\Lambda_{\sigma,\rho}$ are all singletons; we'll denote the unique element of Λ_σ by λ_σ , the unique element of $\Lambda_{\sigma,\tau}$ by $\lambda_{\sigma,\tau}$.

5.2 The modification of \mathcal{G}_{ij} that produces \mathcal{X}_{ij}

The process that transforms the Cayley graph \mathcal{G}_{ij} for the 2-generator Artin group G_{ij} into a systolic complex \mathcal{X}_{ij} , when $m_{ij} < \infty$, is described in some detail in [13, Section 3]. We note only a few points here that we need in order to identify features of the biautomatic structures.

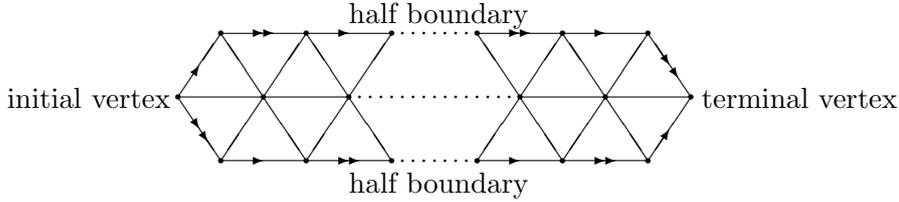


Figure 1: Triangulation of a precell

Under the action of G_{ij} , the Cayley graph \mathcal{G}_{ij} contains vertices in a single orbit, and there are two orbits on undirected edges. These correspond to the directed edges labelled $a_i^{\pm 1}$ and those labelled $a_j^{\pm 1}$.

A circuit of length $2m_{ij}$ labelled by the word that is the concatenation of $a_i a_j \cdots$ and the inverse of $a_j a_i \cdots$ (both of length m_{ij}) starts at each vertex; such a circuit is called a *precell* in [13]. We call the two vertices on the boundary of a precell Π that are (respectively) the sources and targets of two directed edges (labelled x_i and x_j) its *initial* and *terminal vertices*, and we call the two portions of the boundary that consist of paths running from the initial vertex to the terminal vertex its two *half boundaries*. We note that if two precells have intersecting boundaries, then that intersection consists either of a single vertex, which is the initial vertex of one and the terminal vertex of the other, or of a path that is a proper prefix of a half boundary of one and a proper suffix of a half boundary of the other.

The first step of the process that transforms \mathcal{G}_{ij} to \mathcal{X}_{ij} is the triangulation of each such precell by the addition of $m_{ij} - 2$ new vertices, $5m_{ij} - 9$ new edges, and $4m_{ij} - 6$ new 2-cells, as shown in Figure 1. In [13] the new vertices are called *interior vertices*, and the vertices inherited from \mathcal{G}_{ij} *real vertices*; we shall use the same convention here. In fact we shall call both vertices and edges created during the triangulation of a precell Π *interior vertices and edges of Π* .

But this process produces some *bad links* when $m_{ij} > 2$. At this stage, the link of each of the interior vertices is a hexagon, but some of the real vertices have links that are 4-cycles or 5-cycles, and so *bad*. Figure 2 shows what happens at the junction of two precells when $m_{ij} = 3$.

Those bad links are corrected by the addition of new edges, and then the attachment of higher dimensional simplices as necessary, in order to ensure that \mathcal{X}_{ij} continues to be *flag* (that is, any set of pairwise incident vertices forms a simplex).

This whole process of addition of edges and higher dimensional simplices that terminates in the systolic complex \mathcal{X}_{ij} is described in detail in [13, Section 3.2]. The following is clear from that description; part (i) is stated at the end of the second paragraph, and part (ii) is clear from the description

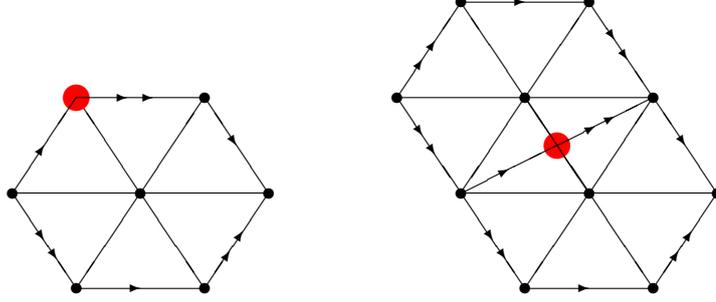


Figure 2: Triangulating the Cayley graph produces bad links at some vertices

of the addition of edges in zigzag patterns.

Lemma 5.1. (i) Any edge of $\mathcal{X}_{ij} \setminus \mathcal{G}_{ij}$ is an interior edge of a single precell or joins two interior vertices within two distinct overlapping precells Π, Π' whose boundaries intersect in a path π of at least two edges, where π contains the initial vertex of one of Π, Π' , and the terminal vertex of the other.

(ii) If Π, Π' are overlapping precells, and u_1, u_2, u_3 are interior vertices with $u_1 \in \Pi$, $u_2, u_3 \in \Pi'$, and such that $\{u_1, u_2\}$ and $\{u_1, u_3\}$ are edges, then $\{u_2, u_3\}$ is an edge within Π' .

Lemma 5.2. Suppose that G is non-free, and that v, v' are real vertices of \mathcal{X} at distance n within the graph \mathcal{X}_1 . Then v, v' are at distance at most $n \max(2, M - 2)$ within \mathcal{G} , where $M := \max_{i,j}(\{m_{ij} : m_{ij} \neq \infty\})$.

Proof. Let $v, v_1, \dots, v_n = v'$ be the sequence of vertices of \mathcal{X} on a path of length n in \mathcal{X}_1 . It follows from the triangle inequality that it is sufficient to prove the result when all vertices v_i with $1 \leq i < n$ (i.e those that are strictly between v and v') are interior.

If $n = 1$, then either v and v' are joined by an edge in \mathcal{X}_1 , or they are joined by an internal edge in a precell of a subcomplex \mathcal{X}_{ij} with $m_{ij} = 2$. In either case $d_{\mathcal{G}}(v, v') \leq 2$, and the result holds.

If $n > 1$, then there are interior vertices between v and v' on the path, and the path must lie within an \mathcal{X}_{ij} subcomplex, for some i, j with finite $m_{ij} > 2$. Let $m := m_{ij}$.

If $n = 2$, then the edges $\{v, v_1\}$ and $\{v_1, v'\}$ are interior edges of a precell Π of which v, v' are boundary vertices. In that case $d_{\mathcal{G}}(v, v') \leq m$ and so, since $m \leq \max(4, 2(m - 2))$, the result holds.

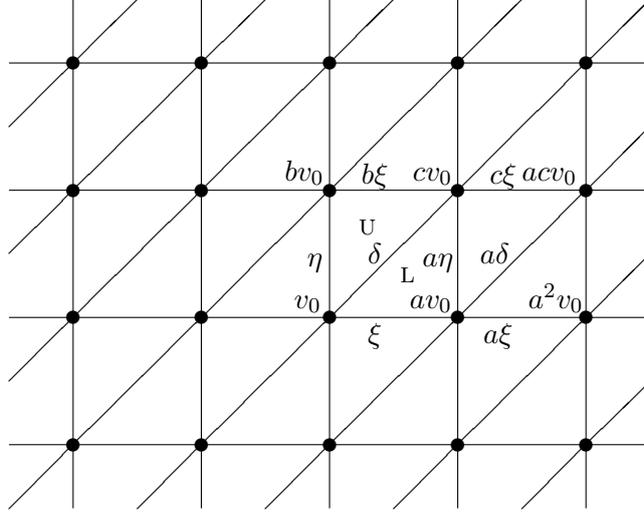


Figure 3: Systolic complex for \mathbb{Z}^2

So now suppose that $n > 2$. Then each of the interior vertices v_i is in a unique precell Π_i and, if $\Pi_i \neq \Pi_{i+1}$, then the boundaries of Π_i and Π_{i+1} intersect in a path containing at least two edges.

Let u_i, w_i be the initial and terminal vertices, respectively, of the precell Π_i . Then v is on the boundary of Π_1 , so $d(v, u_1) + d(v, w_1) = m$ and similarly $d(v', u_{n-1}) + d(v', w_{n-1}) = m$. Hence either $d(v, u_1) + d(v', u_{n-1}) \leq m$ or $d(v, w_1) + d(v', w_{n-1}) \leq m$ (or both). In the first case, we define $v'_i := u_i$ for each $i = 1, \dots, n-1$, and in the second case $v'_i := w_i$ for each $i = 1, \dots, n-1$. So in either case we have $d_{\mathcal{G}}(v'_i, v) + d_{\mathcal{G}}(v'_{n-1}, v') \leq m$. If $v'_i \neq v'_{i+1}$ then, since Π_i and Π_{i+1} overlap, sharing at least two edges of their boundaries, we must have $d_{\mathcal{G}}(v'_i, v'_{i+1}) \leq m - 2$. We deduce that

$$\begin{aligned} d_{\mathcal{G}}(v, v') &\leq d_{\mathcal{G}}(v, v'_1) + \sum_{i=1}^{n-2} d_{\mathcal{G}}(v'_i, v'_{i+1}) + d_{\mathcal{G}}(v'_{n-1}, v') \\ &\leq m + (n-2)(m-2) \leq n \max(2, m-2). \end{aligned}$$

as required. \square

6 Examples of Artin group complexes

6.1 The dihedral Artin group $\mathbb{Z}^2 = \langle a, b \mid ab = ba \rangle$

Here the systolic complex \mathcal{X} is formed from the tessellation of the plane by the integer lattice, with each of the squares subdivided into two right-

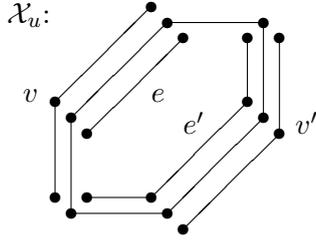


Figure 4: Link of vertex u , residues of v, e , and a 1-ball disjoint from each

angled triangles (2-cells) by a diagonal (an interior edge) running bottom left to top right. The generator a translates one unit to the right, and b one unit upwards. So the vertices are all (x, y) with $x, y \in \mathbb{Z}$, the edges are all pairs $\{(x, y), (x+1, y)\}$, $\{(x, y), (x, y+1)\}$ and $\{(x, y), (x+1, y+1)\}$ and the 2-cells are all triples $\{(x, y), (x, y+1), (x+1, y+1)\}$ and $\{(x, y), (x+1, y), (x+1, y+1)\}$. There are six orbits of the action of G on \mathcal{X} , and $V_0 := \{v_0, \xi, \eta, \delta, \mathbb{U}, \mathbb{L}\}$ is a set of orbit representatives, where $v_0 = (0, 0)$ is the single vertex; $\xi = \{(0, 0), (1, 0)\}$, $\eta = \{(0, 0), (0, 1)\}$, and $\delta = \{(0, 0), (1, 1)\}$ are edges parallel to the x -axis, the y -axis and the diagonal; and $\mathbb{U} = \{(0, 0), (0, 1), (1, 1)\}$ and $\mathbb{L} = \{(0, 0), (1, 0), (1, 1)\}$ are upper and lower triangular 2-cells, as indicated in Figure 3.

We have $\mathcal{B} = \mathcal{A} = \{1, a, b, d, A, B, D\}$, where we use the symbol 1 to represent the identity element, d to represent the product ab , and A, B, D to denote the inverses of a, b, d respectively.

In order to construct \mathcal{L}^0 we need to identify the directed geodesics. We note that the simplices in a directed geodesic $\sigma_0, \dots, \sigma_n$ can only be of dimensions 0 or 1.

The link of an edge (1-dimensional simplex) is a pair of vertices. So if σ_{i+1} is an edge, then σ_i and σ_{i+2} must both be vertices.

On the other hand, the link \mathcal{X}_u of a vertex (0-dimensional simplex) u is the set of vertices and edges of a hexagon.

If v, v' are vertices of that hexagon, then each of $\text{Res}(v, \mathcal{X}_u)$ and $B_1(v', \mathcal{X}_u)$ consists of the three vertices and two edges on a path on the perimeter of the hexagon. For edges e, e' , $\text{Res}(e, \mathcal{X}_u)$ consists of e and the two vertices on e , whereas $B_1(e', \mathcal{X}_u)$ consists of the four vertices and three edges on a path on the perimeter of the hexagon, and contains the 1-balls of both of the vertices on e' . It follows from the conditions on a directed geodesic that, if σ_{i+1} is a vertex, then either

- (i) σ_i and σ_{i+2} are both vertices, opposite each other within the hexagon,

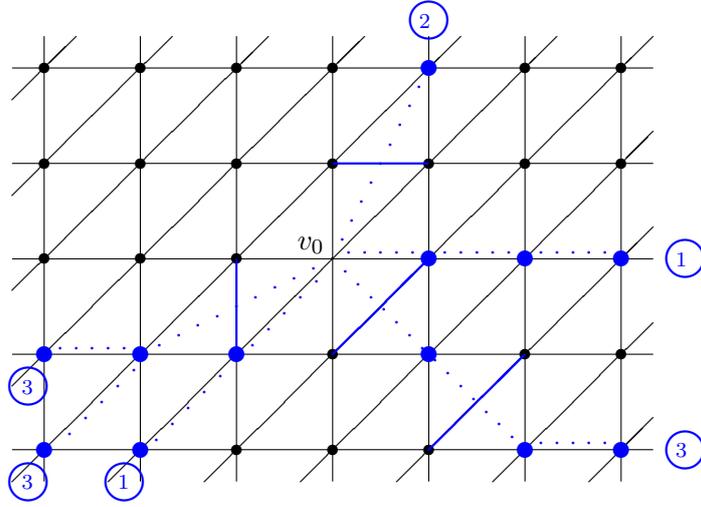


Figure 5: Directed geodesics in complex of \mathbb{Z}^2

or

- (ii) σ_i and σ_{i+2} are both edges, opposite each other within the hexagon,
or
- (iii) σ_i is an edge and σ_{i+2} is one of the two vertices on the edge opposite σ_i within the hexagon.

Hence we see that the directed geodesics of the form $\gamma(v_0, gv_0)$ are of three different types, which can be described in terms of the diagram:

- (1) sequences of consecutive vertices $v_0, v_1, v_2, \dots, v_n$ on one of the six rays through v_0 ,
- (2) sequences $v_0, \sigma_1, v_1, \sigma_2, v_2, \dots, \sigma_k, v_k$ of alternately vertices and edges, where the vertices are all on a line out of v_0 , the edges are bisected by that line, and the vertices and edges are listed in the order in which they meet the line,
- (3) sequences $v_0, \sigma_1, v_1, \dots, \sigma_k, v_k, v_{k+1}, \dots, v_{k+m}$, for which the sequence $v_0, \sigma_1, v_1, \dots, \sigma_k, v_k$ is a sequence of type 2, and $v_k, v_{k+1}, \dots, v_{k+m}$ is the sequence of consecutive vertices on a ray that starts at v_k , one of the two rays through v_k that are adjacent to the continuation of the line from v_0 to v_k .

Considering these three types, we can describe all the words in the language \mathcal{L}^0 . Corresponding to directed geodesics of type 1, we have all words α^n for $\alpha \in \mathcal{A}$. Corresponding to directed geodesics of type 2, we have all words of one of the forms

$$(ad)^n, (bd)^n, (Ab)^n, (DA)^n, (DB)^n, (Ba)^n.$$

Corresponding to directed geodesics of type 3, we have all words of one of the forms

$$(ad)^k a^m, (ad)^k d^m, (bd)^k b^m, (bd)^k d^m, (Ab)^k A^m, (Ab)^k b^m, \\ (DA)^k A^m, (DA)^k D^m, (DB)^k B^m, (DB)^k D^m, (Ba)^k a^m, (Ba)^k B^m.$$

6.2 The right-angled Artin group $\mathbb{Z}^2 * \mathbb{Z}$

Let G be the Artin group $\langle a, b, c \mid ab = ba \rangle$, isomorphic to $\mathbb{Z}^2 * \mathbb{Z}$, and let $\mathcal{A} = \{1, a, b, c, d, A, B, C, D\}$, where d represents the product ab , and A, B, C, D represent the inverses of a, b, c, d . Let G_{ab} be the subgroup $\langle a, b \rangle$, let \mathcal{X}_{ab} be the complex for G_{ab} described in Section 6.1, and let \mathcal{L}_{ab}^0 be the associated language over $\{1, a, b, d, A, B, D\}$. The systolic complex \mathcal{X} for G that is built according to the construction of [13] can be considered as a ‘tree’ of copies of \mathcal{X}_{ab} ; at each vertex of any given copy of \mathcal{X}_{ab} are attached edges labelled by c and its inverse C , whose targets are the basepoints (identity vertices) of further copies of \mathcal{X}_{ab} , but at most one edge joins distinct copies of \mathcal{X}_{ab} . Directed geodesics between two vertices of \mathcal{X} are formed as concatenations of directed geodesics between vertices in copies of \mathcal{X}_{ab} and paths along edges labelled c or C . The biautomatic language \mathcal{L}^0 consists of all words of the form $w_1 \cdots w_k$, with $k \geq 1$, where for each j , either $w_j \in \mathcal{L}_{ab}^0$ or $w_j = c^{i_j}$, for $i_j \neq 0$, alternately, and for $j > 1$, w_j is non-empty.

6.3 The right-angled Artin group $\mathbb{F}_2 \times \mathbb{Z}$

Let G be the Artin group $\langle a, b, c \mid ab = ba, ac = ca \rangle$, isomorphic to $\mathbb{F}_2 \times \mathbb{Z}$, and let $\mathcal{A} = \{1, a, b, c, d, e, A, B, C, D, E\}$, where d, e represent the products ab and ac , and A, B, C, D, E represent the inverses of a, b, c, d, e . Let G_{ab}, G_{ac} be the subgroups $\langle a, b \rangle$ and $\langle a, c \rangle$, let $\mathcal{X}_{ab}, \mathcal{X}_{ac}$ be the complexes for G_{ab}, G_{ac} described in Section 6.1, and let $\mathcal{L}_{ab}^0, \mathcal{L}_{ac}^0$ be the associated languages over $\{1, a, b, d, A, B, D\}$ and $\{1, a, c, e, A, C, E\}$. The systolic complex \mathcal{X} for G that is built according to the construction of [13] can be considered as a ‘tree’ of copies of \mathcal{X}_{ab} and \mathcal{X}_{ac} ; a copy of \mathcal{X}_{ac} is attached along each ‘ray’ of a -edges within any given copy of \mathcal{X}_{ab} , and similarly a copy of \mathcal{X}_{ab} is attached along each ‘ray’ of a -edges within any given copy of \mathcal{X}_{ac} , but copies of \mathcal{X}_{ab} and \mathcal{X}_{ac} can only intersect in a single a -ray, and distinct copies of either \mathcal{X}_{ab} or \mathcal{X}_{ac} must have trivial intersection.

Since every subsequence of a directed geodesic must be a directed geodesic, we can see that a directed geodesic between two vertices of \mathcal{X} must be a sequence $\sigma_1 = \sigma_{i_0}, \dots, \sigma_{i_1}, \dots, \sigma_{i_j}, \dots, \sigma_{i_k} = \sigma_n$, for which the subsequences $\gamma_j = \sigma_{i_{j-1}}, \dots, \sigma_{i_j}$ are directed geodesics within subcomplexes \mathcal{X}_{ab} and \mathcal{X}_{ac}

alternately, and so in particular $\sigma_{i_1}, \dots, \sigma_{i_{k-1}}$ are all on intersection rays. The final subsequence might be along an intersection ray, but none of the others. And, apart from $\sigma_0 = \sigma_{i_0}$ and $\sigma_n = \sigma_{i_k}$, the simplices σ_{i_j} need not necessarily be vertices.

To understand the precise form of such a sequence we need to consider what must happen when a simplex σ_i is on a ray of a -edges, while $\sigma_{i-1}, \sigma_{i+1}$ are not on the ray, but each is within (a distinct) one of the two subcomplexes containing that ray.

When σ_i is a vertex $v = gv_0$, then its link is the union of two intersecting hexagons, as shown in Figure 6. In this example there are five orbits of G on edges, and we denote by $\xi = \{v_0, av_0\}$, $\eta = \{v_0, bv_0\}$, $\theta = \{v_0, cv_0\}$, $\delta = \{v_0, dv_0\}$, $\zeta = \{v_0, ev_0\}$, the representatives of those orbits that pass through the vertex v_0 . The vertices gav_0 and gAv_0 are on the intersection ray, and the outer and inner hexagons are within the two subcomplexes that intersect on that ray.

Suppose first that σ_{i-1} is a vertex (not on the intersection ray). Within the link, the residue of any one of those eight vertices has the same shape, is a path of length 3 containing three vertices and two edges. This residue is disjoint from the 1-ball of precisely two vertices but of no edges from the opposite subcomplex.

So, for example, if $\sigma_{i-1} = gBv_0$ (and $\sigma_i = gv_0$), then $\sigma_{i+1} = gcv_0$ or gEv_0 ; in the corresponding word, this situation gives us $\lambda_i = b$ and $\lambda_{i+1} = c$ or E , but not $\lambda_{i+1} = C$ or e , so we see that words of the form $b^r c^s$ label directed geodesics that are sequences of vertices, but those of the form $b^r c^s$ do not.

Now suppose that σ_{i-1} is an edge. If σ_{i-1} is one of the four edges in the orbit of ξ , then its residue within the link consists of just two vertices, and is disjoint from the 1-ball of two edges and of four vertices within the opposite subcomplex; for example, if $\sigma_{i-1} = gD\xi$, then σ_{i+1} could be any one of $gC\xi$, $gE\xi$, gcv_0 , gev_0 , gEv_0 , or gCv_0 .

But if σ_{i-1} is one of the remaining eight edges, then its residue within the link is disjoint from the 1-ball of two edges and two vertices in the opposite subcomplex; for example, if $\sigma_{i-1} = gB\delta$, then σ_{i+1} could be any one of $gA\zeta$, $gE\theta$, gcv_0 or gEv_0 .

When σ_i is an edge of the intersection ray, then its link is a set of four vertices, and σ_{i-1} and σ_{i+1} can be any vertices within that link in distinct subcomplexes.

From the above analysis, we see from the description in Section 6.1 that each of the subsequences γ_j can be either (1) a sequence of vertices, or (2)

description of \mathcal{L}^0 . Instead, we shall just give some examples of words in the language. In the following description, when we refer to the ‘precell based at g ’ for an element $g \in G$, we mean the precell of which the initial vertex is the real vertex labelled g .

Replaced “
by ‘ and ’
by ’ for consistency

We have $a^n, A^n, b^n, B^n \in \mathcal{L}^0$ for all $n \geq 0$, and these label corresponding directed geodesics through real vertices of the complex. But the accepted word for the group element ab is $1c$. A directed geodesic with that label goes from v_0 to a 2-cell that contains the interior vertices of the precells based at $1, a$ and B , and from there to the real vertex labelled ab . Similarly, the accepted word for ba is $1d$. The word for $a^{-1}b^{-1}$ is $D1$, and labels a corresponding directed geodesic in the reverse direction. Similarly, $C1 \in \mathcal{L}^0$.

The accepted word for aba is $1e$, with corresponding directed geodesic starting from v_0 and passing through the interior vertex of the precell based at 1 . In fact $(1e)^n \in \mathcal{L}^0$ for all $n \geq 0$. The word for $(aba)^{-n}$ is $(E1)^n$, and labels the corresponding directed geodesic in the opposite direction.

The word in \mathcal{L}^0 for ab^{-1} is Cd . It seems easiest to describe a corresponding directed geodesic with this label as one that starts at the real vertex labelled c , proceeds to the interior edge (1-cell) that joins the interior vertex of the precell based at 1 to the real vertex labelled e , and from there to the real vertex labelled d . Similarly, the word in \mathcal{L}^0 for $a^{-1}b$ labels a directed geodesic from the real vertex labelled a to the real vertex labelled b , passing through an interior edge within the precell based at 1 .

The word in \mathcal{L}^0 for a^2b is $1ac$, with directed geodesic going from 1 to the interior edge joining the real vertex labelled a and the interior vertex for the precell based at 1 , and from there to the interior vertex of the precell based at a , and finally to the real vertex labelled a^2b . A directed geodesic labelled $1ae$ has the same first two edges but the final edge leads to the real vertex labelled $(ab)^2 =_G a^2ba =_G ae$.

These examples suggest that it might be easier to understand the nature of the directed geodesics rather than the words that label them, which is not surprising given that the labelling words depend on a choice of orbit representatives of the simplices, whereas the directed geodesics themselves have no such dependency.

7 Properties of the Artin group structures

From now on we suppose that G is a systolic Artin group, and \mathcal{X} the associated systolic complex. We will always assume that v_0 is the vertex of \mathcal{X}

that is equal to the basepoint of the embedded Cayley graph \mathcal{G} , and labelled by the identity element of G . Since all simplex stabilisers are trivial, each of the sets Λ_σ is a singleton set; we denote by λ_σ its single element.

It is clear that the structure of \mathcal{L}^0 must depend on our choice of the set V_0 of orbit representatives, since this determines the labels. But we can make a choice of V_0 that imposes a sensible structure, as we see below.

Suppose first that we have a 2-generator Artin group. Let Π_0 be the unique precell that has v_0 as its initial vertex. We say that V_0 is *based on* Π_0 if

- (1) V_0 contains v_0 and the two real edges through v_0 on the boundary of Π_0 ;
- (2) V_0 contains all the interior vertices and edges of Π_0 ;
- (3) any simplex in V_0 contains at least one vertex within Π_0 .

The following lemma is straightforward to prove, and is used in the proof of the proposition that follows.

Lemma 7.1. *Suppose that v_0 is the identity vertex, and V_0 is based on Π_0 . Then*

- (i) *the label λ_v of a real vertex of \mathcal{X} is the same as its label as a vertex of \mathcal{G} ;*
- (ii) *the label λ_v of an interior vertex of \mathcal{X} is the same as the label of the initial vertex of the unique precell containing v ; that is, $\lambda_v = g$ where v is an interior vertex of $g\Pi_0$.*

Proposition 7.2. *Let G be the 2-generator Artin group*

$$\langle a, b \mid \overbrace{ab \cdots}^m = \overbrace{ba \cdots}^m \rangle.$$

If V_0 is based on Π_0 , then each of the sets \mathcal{A} and \mathcal{B} consists of (a set representing) the complete set of Garside generators of G , together with their inverses and the identity element.

Proof. We represent the set of Garside generators of G as the set of all alternating products of a, b of length less than m , together with the unique alternating product of length m that begins with a , which we denote by Δ (see, e.g. [3]). If g is a Garside generator, then so are $g^{-1}\Delta$ and Δg^{-1} .

We start by observing a correspondence between the vertices other than v_0 on the boundary of Π_0 and the Garside generators; for each real vertex $v \neq v_0$ on the boundary of Π_0 the label λ_v of v (which maps v_0 to v) is a Garside generator, and for each Garside generator g , gv_0 is on the boundary of Π_0 . The generator Δ is the label of the terminal vertex of Π_0 .

We see that a, b, A, B must all be in \mathcal{A} , by considering the directed geodesics of length 1 from v_0 to each of av_0, bv_0, Av_0, Bv_0 , and by considering the related polygonal path, we see that each of a, b, A, B must also be in \mathcal{B} .

We observe that a directed geodesic of length $m - 1$ joins v_0 to Δv_0 , through the $m - 2$ interior vertices of Π_0 . Since all interior vertices and edges of Π_0 are in V_0 , we see that the directed geodesic corresponds to the word $1 \cdots 1\Delta$ of length $m - 1$ over \mathcal{A} in \mathcal{L}^0 , and then that the polygonal path defined by that directed geodesic corresponds to the word $11 \cdots 11\Delta$ of length $2m - 1$ over \mathcal{B} in \mathcal{L}' . Similarly the word $\Delta^{-1}1 \cdots 1$ of length $m - 1$ over \mathcal{A} in \mathcal{L}^0 and the word $\Delta^{-1}11 \cdots 11$ of length $2m - 1$ over \mathcal{B} in \mathcal{L}' label the directed geodesic and polygonal path from v_0 to $\Delta^{-1}v_0$. So we see that $1, \Delta, \Delta^{-1}$ are all in both \mathcal{A} and \mathcal{B} .

Now let v, w be vertices, one on each of the two half boundaries of Π_0 , at distances $i, i - 1$ from v_0 , for some $1 < i < m$. Then λ_v, λ_w are Garside generators, and in fact λ_w is the maximal proper suffix of λ_v . A path of length 2 joins v to w in \mathcal{X}_1 , through an interior vertex u of Π_0 ; both w, u, v and v, u, w are directed geodesics, and their images under λ_w^{-1} and λ_v^{-1} are directed geodesics from v_0 to $\lambda_w^{-1}v$ and $\lambda_v^{-1}w$ respectively, labelled by words $\lambda_w^{-1}\lambda_v$ and $\lambda_v^{-1}\lambda_w$ in \mathcal{L}^0 . We deduce that $\lambda_v, \lambda_v^{-1}$ (and also $\lambda_w, \lambda_w^{-1}$) are in \mathcal{A} , and similarly in \mathcal{B} .

Conversely we need to verify that whenever σ, ρ are consecutive simplices in a directed geodesic or polygonal path then $\lambda_\sigma^{-1}\lambda_\rho$ is equal to either 1, a Garside generator, or the inverse of such. We have $\sigma = \lambda_\sigma\bar{\sigma}$ and $\rho = \lambda_\rho\bar{\rho}$. Then it follows from condition (2) above on V_0 that each of $\bar{\sigma}, \bar{\rho}$ must contain a vertex of Π_0 . It follows that σ contains a vertex v with $\lambda_v = \lambda_\sigma$, and ρ a vertex w with $\lambda_w = \lambda_\rho$. Now, since σ, ρ are consecutive vertices in either a directed geodesic or a polygonal path, v, w are vertices within a simplex of \mathcal{X} (within $\sigma * \rho$ in the first case, and one of σ, ρ in the second). If $v = w$, then $\lambda_\sigma = \lambda_v = \lambda_\rho$. Otherwise, $\{v, w\}$ is an edge of \mathcal{X}_1 . If both v and w are real, then they are adjacent vertices of \mathcal{G} , and so $\lambda_\sigma^{-1}\lambda_\rho = \lambda_v^{-1}\lambda_w \in \{a, b, A, B\}$. If v is real and w is interior, then v must be a vertex on the boundary of the unique precell $\lambda_w\Pi_0$ that contains w , and then $\lambda_v = \lambda_w g$, so $\lambda_\sigma^{-1}\lambda_\rho = g^{-1}$ where either $g = 1$ or g is a Garside generator. If v is interior and w real, the same argument gives $\lambda_\sigma^{-1}\lambda_\rho = g$, with $g = 1$ or g a Garside generator.

Finally, if v and w are both interior, then the unique precells $\lambda_v\Pi_0$ and $\lambda_w\Pi_0$ that contain them are either equal, in which case $\lambda_v = \lambda_w$, or they have boundaries that intersect in a path of length at least 2 that connects the initial vertex of one of the two precells (labelled by λ_v or λ_w) to the terminal vertex of the other (labelled by $\lambda_w\Delta$ or $\lambda_v\Delta$). The element g labelling such a path must be a Garside generator, and we have either $\lambda_w\Delta = \lambda_v g$ or $\lambda_v\Delta = \lambda_w g$, from which we see that $\lambda_\sigma^{-1}\lambda_\rho = \lambda_v^{-1}\lambda_w$ is equal either to

Δg^{-1} , which is also a Garside generator, or to its inverse. \square

We easily deduce the following corollary from Proposition 7.2.

Corollary 7.3. *Let G be an Artin group of almost large type, and \mathcal{X} the systolic complex for G constructed as in [13]. Select a vertex v_0 of \mathcal{X} , and suppose that, for each i, j , Π_0^{ij} is a precell within \mathcal{X}_{ij} with initial vertex v_0 . Suppose that V_0 is chosen such that orbit representatives for each parabolic subgroup G_{ij} of G are based on Π_0^{ij} . Then each of the sets \mathcal{A} and \mathcal{B} consists of (a set representing) the union of complete sets of Garside generators for the subgroups G_{ij} , together with their inverses and the identity element.*

We now consider properties of the language \mathcal{L}^0 over \mathcal{A} .

We can find upper and lower bounds on the lengths of words in \mathcal{L}^0 . The upper bounds are general for a systolic group with trivial vertex stabilisers, but the lower bound is specific to systolic Artin groups.

Proposition 7.4. *Let $G = \langle X \rangle$ be systolic, and $g \in G$, and let \mathcal{L} be the language of a biautomatic structure over \mathcal{A} .*

- (i) *If $w' \in \mathcal{L}$ represents g , then $|w'|_{\mathcal{A}} \leq |g|_X + 1$.*
- (ii) *If G has trivial vertex stabilisers and $w \in \mathcal{L}^0$ represents g , then $|w|_{\mathcal{A}} \leq |g|_X$.*
- (iii) *If G is a non-free Artin group of almost large type and $w \in \mathcal{L}^0$ represents g , then $|g|_X \leq \max\{2, M - 2\}|w|_{\mathcal{A}}$, where M is the maximum of those m_{ij} that are finite.*

Proof. For (i) and (ii), we note that if $\sigma_0, \dots, \sigma_n$ is the directed geodesic that corresponds to $w' \in \mathcal{L}$, then $|w'| = n + 1$. But now if $\gamma = v_0, \dots, v_n$ is a corresponding allowable geodesic (so that $v_i \in \sigma_i$), then it follows from [14, Fact 11.1] that γ is a geodesic in the 1-skeleton of \mathcal{X} . So, since the Cayley graph \mathcal{G} embeds naturally within \mathcal{X}_1 , we have $n \leq |g|_X$, and hence (i) follows. Now if $w \in \mathcal{L}^0$ representing g , then the word $1w$ represents g in \mathcal{L} , and so (ii) is an immediate consequence of (i). Finally, (iii) follows by Lemma 5.2 \square

Proposition 7.5. *If $\sigma_0, \dots, \sigma_n$ is a directed geodesic, then there is a directed geodesic $\sigma_0, \dots, \sigma_{n-1}, v_n$, for which v_n is a vertex. Furthermore, if v_n is an interior vertex, then there is a directed geodesic*

$$\sigma_0, \dots, \sigma_{n-1}, v_n, v_{n+1}, \dots, v_{n+k},$$

with v_{n+k} real and $k \leq M - 2$, where M is the maximum of those m_{ij} that are finite.

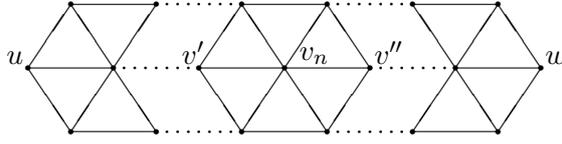


Figure 7: Extending a directed geodesic past v_n

Proof. The existence of a vertex v_n such that $\sigma_0, \dots, \sigma_{n-1}, v_n$ is a directed geodesic follows directly from Lemma 3.3 (2); then σ_{n-1} is within the link \mathcal{X}_{v_n} of v_n . If v_n is not real, then it is an interior vertex within a precell Π of some \mathcal{X}_{ij} subcomplex. By Lemma 5.1 the link of v_n is also within that subcomplex: it intersects the precell in a hexagon, as shown in Figure 7. Let u, w, v_n, v', v'' be the vertices marked in that figure.

Let π be the intersection of $\text{Res}(\sigma_{n-1}, \mathcal{X}_{v_n})$ with that hexagon. We claim that π is either a closed path of two consecutive edges and their vertices, or is a single closed edge. That is clear if σ_{n-1} is within Π , in which case σ_{n-1} is either a vertex of edge of the hexagon. Otherwise, σ_{n-1} contains a vertex u' outside Π that is joined by an edge to v_n . Then, by Lemma 5.1, u' must be an interior vertex in a precell Π' that intersects Π along a path on one of its half boundaries. Then every vertex of π apart from u' is joined to u' by an edge, and so in this case our claim follows by Lemma 5.1.

Now, except when π consists of the two closed edges through either v' or v'' , at least one of the four real vertices of the hexagon is distance at least two from π , and so its 1-ball is disjoint from π . In that case, we can choose v_{n+1} to be such a vertex, and $\sigma_0, \dots, \sigma_{n-1}, v_n, v_{n+1}$ is a directed geodesic with v_{n+1} real.

Otherwise, without loss of generality, π consists of the two closed edges through v' . Then the 1-ball about v'' is disjoint from π , and we can choose v_{n+1} to be v'' . Now $\sigma_0, \dots, \sigma_{n-1}, v_n, v_{n+1}$ is a directed geodesic, but v_{n+1} is interior. But now, by the same argument as above, we can extend the directed geodesic with vertices $v_{n+2}, \dots, v_{n+k} = w$ as successive vertices on the path within the triangulated precell from v'' to w . \square

We note that, since the reverse of a directed geodesic is not in general a directed geodesic, we should not expect \mathcal{L}^0 to be symmetric. And although every prefix of a directed geodesic is also a directed geodesic, its final simplex is not in general a vertex, and so we do not expect (or observe, in our examples) prefix closure. However, we have the following result.

Corollary 7.6. *Suppose that $w \in \mathcal{L}^0$, and that w_1 is an prefix of w , and let w_2 be the maximal prefix of w_1 (formed by deleting its last letter). Then there exists a word $w' \in \mathcal{L}^0$, that has w_2 as a prefix, with $|w'| - |w_2| \leq$*

$\max_{i,j}(\{m_{ij} - 1 : m_{ij} \neq \infty\})$.

Proof. We let $\sigma_0, \dots, \sigma_n$ be the directed geodesic corresponding to w' , and choose $w' \in \mathcal{L}^0$ to correspond to the directed geodesic $\sigma_0, \dots, \sigma_{n-1}, v_n, \dots, v_{n+k}$ that is constructed in Proposition 7.5. \square

8 Computer code

We have written computer code that constructs the biautomatic structures described in [13] for certain Artin groups G of almost large type. Currently our programs can handle such groups provided that all of the entries m_{ij} , $i \neq j$ in the associated Coxeter matrix lie in $\{2, 3, 4, \infty\}$. This code is written in GAP [6] and makes use of the KBMAG [15] package for computing shortlex automatic structures and performing operations on finite state automata.

The biautomatic structures that we construct are over the generating set \mathcal{A} that was described in Section 4.1 above for general systolic groups. With this code we can explicitly construct the word acceptor and right and left multiplier automata and verify computationally that they do indeed define a biautomatic structure. We can also reduce words over the standard generating set $X = \{a_1, \dots, a_n\}$ of G to their equivalent normal form words over \mathcal{A} .

As we saw in Corollary 7.3, with appropriate choice of the set V_0 of orbit representatives, the set \mathcal{A} is equal to (symbols representing) the set $\{1\} \cup Y \cup Y^{-1}$, where Y is the union of X and, for each i, j with $1 \leq i < j \leq n$, the sets corresponding to $\{a_i a_j\}$, $\{a_i a_j, a_j a_i, a_i a_j a_i\}$, and $\{a_i a_j, a_j a_i, a_i a_j a_i, a_j a_i a_j, a_i a_j a_i a_j\}$, when $m_{ij} = 2, 3$ and 4 , respectively. In fact we only construct the left and right multiplier automata for the generators in $\{1\} \cup X \cup X^{-1}$, but this is sufficient to verify the correctness of the structures, and the multipliers for the additional generators could be constructed as composite automata (defined in [5, Section 5]) if required.

The code works as follows. The authors proved in [10] and [11] that G is shortlex automatic over X , and we start by using KBMAG to construct this shortlex automatic structure. Then we enumerate words in the accepted language up to some chosen length l (it turns out that $l = 7$ is sufficient in general, or $l = 4$ if there are no i, j with $m_{ij} = 4$). So we have effectively constructed the ball of radius l about the base point in the Cayley graph \mathcal{G} of G over X . Using this, we construct the simplices in part of the systolic complex \mathcal{X} , where v_0 is the base point of \mathcal{G} .

Now, provided that l is large enough, this part of the complex contains

enough information to enable us to construct the automaton, as described in 4.5; since vertex stabilisers are trivial, we construct the automaton \mathcal{M}^0 that accepts \mathcal{L}^0 . To do this, we first have to choose orbit representatives V_0 for the action of G on \mathcal{X} , which must satisfy the hypotheses of Corollary 7.3. In general, the representatives are chosen to be close in \mathcal{X} to v_0 .

The automaton \mathcal{M}^0 is in general non-deterministic, but we can use the automata manipulation functions in KBMAG to construct a deterministic automaton with a minimal number of states that accepts the same language.

We next describe how we construct the multiplier automata. In general, let W be a finite subset of a group G , and let u and v be words over some generating set X of G . For $0 \leq i \leq \ell(u)$, let $u(i)$ denote the prefix of u of length i , and put $u(i) = u$ for $i \geq \ell(u)$. Then, for $a \in G$, we write $u \sim_{W,a} v$ if $u(i)^{-1}av(i) \in W$ for all $i \geq 0$.

For the language L of an automatic structure of G over X , there is a finite set $W_R \subset G$ such that, for all $a \in X \cup X^{-1}$, the language $\{(u, v) : u, v \in L, ua =_G v, u \sim_{W_R,1} v\}$ is regular and has both of its components projecting onto L . These are the languages of the right multipliers in the automatic structure. An automaton accepting this language is described in [4, Section 6.3].

If in addition L is the language of a biautomatic structure for G , then there is a finite set $W_L \subset G$ such that, for all $a \in X \cup X^{-1}$, the language $\{(u, v) : u, v \in L, u =_G av, u \sim_{W_L,a} v\}$ is regular and has both of its components projecting onto L . These are the languages of the left multipliers in the biautomatic structure.

If we are given the word acceptor of an automatic structure and a candidate for the set W_R , then we can construct the associated right multiplier automata and then use the *axiom checking process* described in [4, Section 5.1] to verify that this really is an automatic structure and hence that W_R has the required properties. If the axiom checking process fails, then we can attempt to find specific instances of failure and thereby add further words to W_R , after which we try again. This is essentially the same method that is described in more detail in [5] for constructing shortlex automatic structures. As mentioned earlier, in our code we just construct the automata for generators in $\{1\} \cup X \cup X^{-1}$, which is sufficient for the axiom checking.

Furthermore, if we are given a candidate for W_L , then we can construct the associated left multipliers. We can use these to check the condition

$$\forall u \in L \exists v \in L (u =_G av, u \sim_{W_L,a} v)$$

for each $a \in X$. If the condition holds, then we have verified that the structure is indeed biautomatic. If not, then we can find specific words

$u \in L$ for which the condition fails, compute words $v \in L$ with $u =_G av$, and then add new elements to W_L to make the condition $u \sim_{W_L, a} v$ hold for these words.

Note that to do all of this in our specific situation for Artin groups (in which the accepted language of the word acceptor is \mathcal{L}^0), we need to be able to test words over \mathcal{A} for equality in G (i.e. to solve the word problem in G), but we can use the shortlex automatic structure to do that. We also need to be able to reduce words to normal form; that is, given a word u find a word v with $v \in \mathcal{L}^0$ such that $u =_G v$. A method for doing this that uses the set W_R is described in the proof of Theorem 4.1 of [5].

In practice it turned out that $W_R = \mathcal{A}$ was adequate for the construction of the right multipliers for the generators in $\{1\} \cup X \cup X^{-1}$, and also for word reduction. But the set W_L needs to be significantly larger, and we were able to construct that using the error correcting process described above.

The principal reason that we have not attempted to handle Artin groups with $m_{ij} > 4$ is that the complex \mathcal{X} becomes increasingly more difficult to describe explicitly with increasing m_{ij} . We found that, for $m_{ij} = 2, 3$ and 4, the highest dimension of a simplex in the complex is m_{ij} and it seems plausible that the same is true for higher m_{ij} .

References

- [1] T. Brady and J.P. McCammond, Three-generator Artin groups of large type are biautomatic, *J. Pure Appl. Alg.* 151 (2000) 1–9.
- [2] R. Charney, Artin groups of finite type are biautomatic, *Math. Annalen* 292 (1992) 671–683.
- [3] R. Charney, Geodesic automation and growth functions for Artin groups of finite type. *Math. Ann.* 301 (1995), no. 2, 307–324
- [4] David B. A. Epstein J.W. Cannon, D.F. Holt, S. Levy, M.S. Patterson and W. Thurston, *Word processing in groups*, Jones and Bartlett, 1992.
- [5] D.B.A. Epstein, D.F. Holt and S.E. Rees, The use of Knuth-Bendix methods to solve the word problem in automatic groups, *J. of Symbolic Computation*, 12 (1991),397–414.
- [6] The GAP Group. *GAP — Groups, Algorithms, and Programming, Version 4.8.10*, 2018. (<http://www.gap-system.org>).
- [7] S.M. Gersten and H.B. Short, Small cancellation theory and automatic groups. *Invent. Math.* 102 (1990), no. 2, 305–334.

- [8] S.M. Gersten and H. Short, Small cancellation theory and automatic groups. II. *Invent. Math.* 105 (1991), no. 3, 641–662.
- [9] S. Hermiller and J. Meier, Algorithms and geometry for graph products of groups, *J. Alg.* 171 (1995).
- [10] D.F. Holt and S. Rees, Artin groups of large type are automatic with regular geodesics, *Proc. London Math. Soc.* 104 (2012) 486–512.
- [11] D.F. Holt and S. Rees, Shortlex automaticity and geodesic regularity in Artin groups, *Groups, Complex. Cryptol.* 5 (2013) 1–23.
- [12] D.F. Holt, S. Rees and C.E. Röver, *Groups, languages and automata*, LMS student texts 88, CUP, 2017.
- [13] J. Huang and D. Osajda, Large-type Artin groups are systolic, [arXiv:1706.05473](https://arxiv.org/abs/1706.05473)
- [14] T. Januszkiewicz and J. Swiatkowski, Simplicial nonpositive curvature, *Publ. Math. Inst. Hautes Etudes Sci.* No. 104 (2006), 1–85.
- [15] D.F. Holt, KBMAG—*Knuth-Bendix in Monoids and Automatic Groups*, software package (1995), available from (<http://homepages.warwick.ac.uk/~mareg/download/kbmag2//>)
- [16] G.A. Niblo and L.D. Reeves, The geometry of cube complexes and the complexity of their fundamental groups, *Topology* 37 (1998) 621–633.
- [17] G.A. Niblo and L.D. Reeves, Coxeter groups act on CAT(0) cube complexes, *J. Group Theory* 6 (2003), no. 3, 399–413.
- [18] D. Peifer, Artin groups of extra-large type are biautomatic, *J. Pure Appl. Alg.* 110 (1996) 15–56.
- [19] P. Papasoglu, Strongly geodesically automatic groups are hyperbolic, *Invent. Math.* 121 (1995), 323–334.
- [20] J. Swiatkowski, Regular path systems and (bi)automatic groups, *Geometriae Dedicata* 118 (2006) 23–48.
- [21] L. VanWyk, Graph groups are biautomatic, *J. Pure App. Alg.* 94 (1994).