TRANSACTIONS OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 367, Number 5, May 2015, Pages 3549–3577
S 0002-9947(2014)06218-1
Article electronically published on October 1, 2014

DIVERGENCE IN RIGHT-ANGLED COXETER GROUPS

PALLAVI DANI AND ANNE THOMAS

Abstract. Let W be a 2-dimensional right-angled Coxeter group. We characterise such W with linear and quadratic divergence, and construct right-angled Coxeter groups with divergence polynomial of arbitrary degree. Our proofs use the structure of walls in the Davis complex.

1. Introduction

The divergence of a pair of geodesics is a classical notion related to curvature. Roughly speaking, given a pair of geodesic rays emanating from a basepoint, their divergence measures, as a function of r, the length of a shortest "avoidant" path connecting their time-r points. A path is *avoidant* if it stays at least distance r away from the basepoint. In [15], Gersten used this idea to define a quasi-isometry invariant of spaces, also called divergence. We recall the definitions of both notions of divergence in Section 2.

The divergence of every pair of geodesics in Euclidean space is a linear function, and it follows from Gersten's definition that any group quasi-isometric to Euclidean space has linear divergence. In a δ-hyperbolic space, any pair of non-asymptotic rays diverges exponentially; thus the divergence of any hyperbolic group is exponential. In symmetric spaces of non-compact type, the divergence is either linear or exponential, and Gromov suggested in [16] the same should be true in CAT(0) spaces.

Divergence has been investigated for many important groups and spaces, and contrary to Gromov's expectation, quadratic divergence is common. Gersten first exhibited quadratic divergence for certain CAT(0) spaces in [15]. He then proved in [14] that the divergence of the fundamental group of a closed geometric 3-manifold is either linear, quadratic or exponential, and characterised the (geometric) ones with quadratic divergence as the fundamental groups of graph manifolds. Kapovich–Leeb [17] showed that all graph manifold groups have quadratic divergence. More recently, Duchin–Rafi [13] established that the divergence of Teichmüller space and the mapping class group is quadratic (for mapping class groups this was also obtained by Behrstock in [5]). Druţu–Mozes–Sapir [12] have conjectured that the divergence of lattices in higher rank semisimple Lie groups is always linear, and proved this conjecture in some cases. Abrams et al. [1] and independently Behrstock–Charney [2] have shown that if A_Γ is the right-angled

Received by the editors February 13, 2013 and, in revised form, June 12, 2013.
2010 *Mathematics Subject Classification.* Primary 20F65; Secondary 20F55.
The first author was supported in part by Louisiana Board of Regents Support Fund Contract LEQSF(2011-14)-RD-A-06 and NSF Grant No. DMS-1207868.
This research of the second author was supported in part by ARC Grant No. DP110100440, and the second author was also supported in part by an Australian Postdoctoral Fellowship.

Artin group associated to a graph Γ, the group A_Γ has either linear or quadratic divergence, and its divergence is linear if and only if Γ is (the 1-skeleton of) a join.

In this work we study the divergence of 2-dimensional right-angled Coxeter groups. Our first main result is Theorem 1.1 below, which characterises such groups with linear and quadratic divergence in terms of their defining graphs. This result can be seen as a step in the quasi-isometry classification of (right-angled) Coxeter groups, about which very little is known.

We note that by [10], every right-angled Artin group is a finite index subgroup of, and therefore quasi-isometric to, a right-angled Coxeter group. However, in contrast to the setting of right-angled Artin groups, where one sees only linear and quadratic divergence, even the class of 2-dimensional right-angled Coxeter groups exhibits a greater variety of divergence functions. For example, there exist 2-dimensional right-angled Coxeter groups that are hyperbolic, and therefore have exponential divergence. Our second main result provides further evidence of this phenomenon: in Theorem 1.2 below, we construct right-angled Coxeter groups with divergence polynomial of any degree.

Given a finite simplicial graph Γ, the associated *right-angled Coxeter group* W_Γ has generating set S the vertices of Γ, and relations $s^2 = 1$ for all $s \in S$ and $st = ts$ whenever s and t are adjacent vertices. We restrict our attention to W_Γ one-ended and of dimension 2; equivalently, Γ is connected, triangle-free and has no separating vertices or edges. The group W_Γ acts geometrically on its Davis complex Σ_Γ. As Σ_Γ is a CAT(0) square complex, W_Γ is a CAT(0) group. We investigate divergence by considering geodesics and paths in the Cayley graph of W_Γ with respect to the generating set S. This Cayley graph may be identified with the 1-skeleton of the Davis complex Σ_Γ, and we use many properties of walls in the Davis complex to determine upper and lower bounds on lengths of avoidant paths. See Section 3 for details and further background on W_Γ and Σ_Γ, including references.

By Moussong's Theorem [9, Corollary 12.6.3], W_Γ is hyperbolic if and only if Γ has no embedded cycles of length four. In order to investigate divergence for W_Γ not hyperbolic, we consider the set of embedded four-cycles in Γ. Each such four-cycle induces a family of isometrically embedded flats in Σ_Γ. In Section 4 we define an explicit, easy-to-check condition, which we call $\mathcal{CFS}$, on the graph Γ. If Γ is $\mathcal{CFS}$, then Σ_Γ has a distinguished collection of flats coming from a specific class of four-cycles in Γ, with these flats intersecting along infinite bands, such that each generator of W_Γ is in the four-cycle for at least one such flat.

Theorem 1.1. *Let Γ be a finite, simplicial, connected, triangle-free graph which has no separating vertices or edges. Let W_Γ be the associated right-angled Coxeter group.*

1. *The group W_Γ has linear divergence if and only if Γ is a join.*
2. *The group W_Γ has quadratic divergence if and only if Γ is $\mathcal{CFS}$ and is not a join.*

Note that part (1) is equivalent to saying that W_Γ has linear divergence if and only if it is reducible, since for Γ triangle-free, W_Γ is reducible if and only if Γ is a join. Our proof of part (1) is similar to that of the corresponding result for A_Γ in [1].

To establish a quadratic upper bound on divergence when the graph Γ is $\mathcal{CFS}$, we construct, given a pair of geodesic segments based at a common point, an avoidant path between their endpoints which travels only in flats from the distinguished collection of flats guaranteed by the $\mathcal{CFS}$ condition. Since the divergence within a flat is linear, the quadratic upper bound comes from showing that this path only needs to pass through linearly many flats. As pointed out by the referee, this quadratic upper bound could also be obtained using the thickness machinery developed by Behrstock–Druţu [3]. (See Remark 4.8.)

The more delicate direction of part (2) of Theorem 1.1 is proving that $\mathcal{CFS}$ graphs are exactly the class of graphs for which there is a quadratic upper bound on divergence. We in fact establish a cubic lower bound on divergence when Γ is not $\mathcal{CFS}$. To obtain lower bounds on the lengths of avoidant paths, we consider van Kampen diagrams whose boundaries consist of a pair of geodesic segments with common basepoint and an avoidant path between their endpoints. The fact that the defining graph is not $\mathcal{CFS}$ has certain implications on the cell structure of the van Kampen diagram, which force a lower bound on the length of its boundary (and therefore of the avoidant path).

In contrast with the classes of groups discussed above, right-angled Coxeter groups may have divergence other than linear, quadratic or exponential. We prove:

Theorem 1.2. *For all $d \geq 1$, there is a right-angled Coxeter group W_d with divergence polynomial of degree d.*

In [14], Gersten asked whether polynomial divergence of degree ≥ 3 is possible for CAT(0) groups. Macura [19] constructed a family of CAT(0) groups G_d with divergence polynomial of degree $d \geq 2$. These groups G_d are the same as the "hydra groups" investigated by Dison–Riley [11]. Behrstock–Druţu [3] subsequently obtained examples of CAT(0) groups H_d with divergence polynomial of any degree $d \geq 2$, with H_d the amalgamated free product of two copies of H_{d-1} along an infinite cyclic subgroup. The groups W_d that we construct are not of this form. Most recently, Behrstock–Hagen [4] used a similar construction to that of [3] to obtain fundamental groups of CAT(0) cube complexes with divergence polynomial of any degree. Theorem 1.2 provides an answer to Gersten's question within a well-known class of CAT(0) groups.

We prove Theorem 1.2 in Section 5, where we inductively construct a family of graphs Γ_d such that $W_d = W_{\Gamma_d}$ has divergence polynomial of degree d. We prove upper and lower bounds on the divergence of W_d in Propositions 5.1 and 5.3 respectively. As discussed in Remark 5.2, the upper bound for the divergence of W_d could also be derived from thickness considerations. Our arguments to obtain the lower bounds on divergence are considerably shorter than Macura's.

After proving Theorem 1.2, we noticed that Macura's group G_d and our group W_d both act geometrically on a CAT(0) square complex with all vertex links equal to the graph Γ_d (namely the Cayley 2-complex for G_d, and the Davis complex for W_d, respectively). A natural question is thus whether G_d and W_d are commensurable. Since our techniques for addressing this question are quite different to those used to prove Theorems 1.1 and 1.2, we discuss this question in Appendix A. We first show in Proposition A.8, using covering theory and complexes of groups, that G_2 and W_2 are commensurable. While attempting to prove commensurability of G_d and W_d for $d > 2$, we were surprised to discover that their corresponding square complexes are not in fact isometric (see Corollary A.10). Hence the strategy of finding a common finite cover to establish commensurability fails. We do not know whether G_d and W_d are commensurable or even quasi-isometric for $d > 2$.

2. Divergence

In this section we recall Gersten's definition of divergence as a quasi-isometry invariant from [15]. We restrict ourselves to spaces which are one-ended.

Let (X,d) be a one-ended geodesic metric space. For $p \in X$, let $S(p,r)$ and $B(p,r)$ denote the sphere and open ball of radius r about p. A path in X is said to be (p,r)-*avoidant* if it lies in $X - B(p,r)$. Then, given a pair of points $x, y \in X - B(p,r)$, the (p,r)-*avoidant distance* $d^{\text{av}}_{p,r}(x,y)$ between them is the infimum of the lengths of all (p,r)-avoidant paths connecting x and y.

Now fix a basepoint $e \in X$. In the rest of the paper we will write r-*avoidant* or simply *avoidant* for (e,r)-avoidant, and $d^{\text{av}}(x,y)$ for $d^{\text{av}}_{e,r}(x,y)$, indicating the basepoint and radius only if they differ from e and r.

For each $0 < \rho \leq 1$, let

$$\delta_\rho(r) = \sup_{x,y \in S(e,r)} d^{\text{av}}_{\rho r}(x,y).$$

Then the *divergence* of X is defined to be the resulting collection of functions

$$\text{div}_X = \{\delta_\rho \,|\, 0 < \rho \leq 1\}.$$

The spaces X that we will consider (Cayley graphs of right-angled Coxeter groups) have the geodesic extension property (i.e. any finite geodesic segment can be extended to an infinite geodesic ray). It is not hard to show that in a metric space X with this property, $\delta_\rho \simeq \delta_1$ for all $0 < \rho \leq 1$, where $\simeq$ is the equivalence on functions generated by

$$f \preceq g \iff \exists\, C > 0 \text{ such that } f(r) \leq Cg(Cr + C) + Cr + C.$$

Thus in this paper, we think of div_X as a function of r, defining it to be equal to δ_1. We say that the divergence of X is *linear* if $\text{div}_X(r) \simeq r$, *quadratic* if $\text{div}_X(r) \simeq r^2$, and so on.

The divergence of X is then, up to the relation $\simeq$, a quasi-isometry invariant which is independent of the chosen basepoint (see [15]). Thus it makes sense to define the divergence of a finitely generated group to be the divergence of one of its Cayley graphs.

The divergence of a pair of geodesic rays α and β with the same initial point p, or of a bi-infinite geodesic γ, are defined as, respectively,

$$\text{div}_{\alpha,\beta}(r) = d^{\text{av}}_{p,r}(\alpha(r), \beta(r)) \text{ and } \text{div}_\gamma(r) = d^{\text{av}}_{\gamma(0),r}(\gamma(-r), \gamma(r)).$$

Note that in a geodesic metric space X, if $\text{div}_{\alpha,\beta}(r) \leq f(r)$ for all pairs of geodesic rays in X with initial point e, then $\text{div}_X(r) \leq f(r)$. On the other hand, if there exists a pair of geodesic rays (or a bi-infinite geodesic) such that $\text{div}_{\alpha,\beta}(r) \succeq f(r)$, then $\text{div}_X(r) \succeq f(r)$. Finally, if X is CAT(0) and $\text{div}_{\alpha,\beta}(r) \geq f(r)$, then, using the fact that projections do not increase distances, one can show that $d^{\text{av}}_{p,r}(\alpha(s), \beta(t)) \geq f(r)$ for any $s, t \geq r$. These observations will be used repeatedly in the proofs.

3. Coxeter groups and the Davis complex

In this section, we recall definitions and results concerning right-angled Coxeter groups (Section 3.1) and their associated Davis complexes (Section 3.2). Section 3.3 then gives a careful discussion of walls in the Davis complex. Section 3.4 discusses paths in the Cayley graph of W_Γ and their relationship to walls in the Davis complex. We mostly follow Davis' book [9].

3.1. **Right-angled Coxeter groups.** Let Γ be a finite simplicial graph with vertex set S and let W_Γ be the associated right-angled Coxeter group, as defined in the introduction. The group W_Γ is *reducible* if S can be written as a disjoint union $S_1 \sqcup S_2$ of non-empty subsets such that $W_1 := \langle S_1 \rangle$ commutes with $W_2 := \langle S_2 \rangle$, in which case $W = W_1 \times W_2$.

In this paper we restrict ourselves to Γ triangle-free. Then it is easy to see that W_Γ is reducible if and only if Γ is a join (i.e. a complete bipartite graph). Also, with this assumption, W_Γ is one-ended if and only if Γ is connected and has no separating vertices or edges (see Theorem 8.7.2 of [9]).

Given $T \subseteq S$, the subgroup $W_T := \langle T \rangle$ of W_Γ is called a *special subgroup*. By convention, $W_\emptyset$ is the trivial group. If Λ is an induced subgraph of Γ with vertex set T, we may write W_Λ for the special subgroup W_T. Denote by C_2 the cyclic group of order 2 and by D_∞ the infinite dihedral group. Then for each $s \in S$, the special subgroup $W_{\{s\}}$ is isomorphic to C_2. If s and t are adjacent vertices, then $W_{\{s,t\}} \cong C_2 \times C_2$, while if s and t are non-adjacent vertices, we have $W_{\{s,t\}} \cong D_\infty$.

Example 3.1. Suppose $T = \{s, t, u, v\} \subset S$ is such that s, t, u and v are, in cyclic order, the vertices of an embedded four-cycle in Γ. Then W_T is reducible with

$$W_T = W_{\{s,u\}} \times W_{\{t,v\}} \cong D_\infty \times D_\infty.$$

Now suppose T_1 and T_2 are distinct subsets of S such that $T_1 \cap T_2 = \{s, t, u\}$, with s and u both adjacent to t. Since Γ is triangle-free, this implies that s and u are not connected by an edge. Then

$$W_{T_1 \cap T_2} = W_{\{s,u\}} \times W_{\{t\}} \cong D_\infty \times C_2$$

and $W_{T_1 \cup T_2}$ splits as the amalgamated free product

$$W_{T_1 \cup T_2} = W_{T_1} *_{W_{T_1 \cap T_2}} W_{T_2} \cong W_{T_1} *_{D_\infty \times C_2} W_{T_2}.$$

A special subgroup W_T is said to be a *spherical* special subgroup if W_T is finite. The set of *spherical subsets* of S, denoted $\mathcal{S}$, is the set of subsets $T \subseteq S$ such that W_T is spherical. (The reason for the terminology "spherical" is that if W_T is finite, then W_T acts as a geometric reflection group on the unit sphere in $\mathbb{R}^{|T|}$; see Theorem 6.12.9 of [9].) It follows from the paragraph before Example 3.1 that for Γ triangle-free, the only spherical subsets of S are the empty set, the sets $\{s\}$ for $s \in S$, and the sets $\{s, t\}$ where s and t are adjacent vertices. The corresponding spherical special subgroups of W are isomorphic to the trivial group, C_2, and $C_2 \times C_2$ respectively.

A *word* in the generating set S is a finite sequence $\mathbf{s} = (s_1, \ldots, s_k)$ where each $s_i \in S$. We denote by $w(\mathbf{s}) = s_1 \cdots s_k$ the corresponding element of W. The *support* of a word $\mathbf{s}$ is the set of generators which appear in $\mathbf{s}$. A word $\mathbf{s}$ is said to be *reduced* if the element $w(\mathbf{s})$ cannot be represented by any shorter word, and a word $\mathbf{s}$ is *trivial* if $w(\mathbf{s})$ is the trivial element. We will later by abuse of notation write $s_1 \cdots s_k$ for both words and group elements. A word $\mathbf{s}$ in the generating set S of a right-angled Coxeter group is reduced if and only if it cannot be shortened by a sequence of operations of either deleting a subword of the form (s, s), with $s \in S$, or replacing a subword (s, t) such that $st = ts$ by the subword (t, s). (This is a special case of Tits' solution to the word problem for Coxeter groups; see Theorem 3.4.2 of [9].)

3.2. **The Davis complex.** From now on, Γ is a finite, simplicial, connected, triangle-free graph with no separating vertices or edges, and $W = W_\Gamma$ is the associated right-angled Coxeter group. In this section, we discuss the Davis complex for W.

By our assumptions on Γ, we may define the *Davis complex* $\Sigma = \Sigma_\Gamma$ to be the Cayley 2-complex for the presentation of W_Γ given in the introduction, in which all disks bounded by a loop with label s^2 for $s \in S$ have been shrunk to an unoriented edge with label s. Then the vertex set of Σ is W_Γ and the 1-skeleton of Σ is the Cayley graph $\mathcal{C}_\Gamma$ of W with respect to the generating set S. Since all relators in this presentation other than $s^2 = 1$ are of the form $stst = 1$, Σ is a square complex. We call this cellulation of Σ the *cellulation by big squares*, with the *big squares* being the 2-cells. Note that the link of each vertex in this cellulation is the graph Γ.

We next define the *cellulation by small squares* of Σ to be the first square subdivision of the cellulation by big squares, with the *small squares* being the squares obtained by subdividing each big square into four. We will use both of these cellulations in our proofs.

We now assign *types* $T \in \mathcal{S}$ to the vertices of the cellulation by small squares. If σ is also a vertex of the cellulation by big squares, then σ has type $\emptyset$. If σ is the midpoint of an edge in the cellulation by big squares, then since $\mathcal{C}_\Gamma$ is the 1-skeleton of the cellulation by big squares, σ is the midpoint of an edge connecting g and gs for some $g \in W$ and $s \in S$, and we assign type $\{s\} \in \mathcal{S}$ to σ. Finally, if σ is the centre of a big square, then σ is assigned type $\{s,t\} \in \mathcal{S}$, where two of the vertices adjacent to σ have type $\{s\}$, and two of the vertices adjacent to σ have type $\{t\}$.

Consider Σ with the cellulation by small squares. The group W naturally acts on the left on Σ, preserving types, so that the stabiliser of each vertex of type $T \in \mathcal{S}$ is a conjugate of the finite group W_T. Let σ be the vertex of type $\emptyset$ corresponding to the identity element of W. The *base chamber* K is the union of the set of small squares which contain σ. Any translate of K by an element of W is called a *chamber*. For each $T \in \mathcal{S}$, we denote by σ_T the unique vertex of type $T \in \mathcal{S}$ in the base chamber. The quotient of Σ by the action of W is the base chamber K, and the W-stabiliser of σ_T is precisely the spherical special subgroup W_T.

For $s \in S$, the *mirror* K_s is the union of the set of edges in the base chamber which contain $\sigma_{\{s\}}$ but not $\sigma_\emptyset$. The mirror K_s is thus the star graph of valence n, where n is the cardinality of the set $\{t \in S \mid st = ts, t \neq s\}$. Note that $n \geq 2$, since Γ has no isolated vertices or vertices of valence one. The *centre* of the mirror K_s is the vertex $\sigma_{\{s\}}$. Any translate of K_s by an element of W is called a *panel (of type s)*.

Let Σ be the Davis complex cellulated by either big or small squares. We now metrise Σ so that each big square is a unit Euclidean square, hence each small square is a Euclidean square of side length $\frac{1}{2}$. By [9, Theorem 12.2.1], this piecewise Euclidean cubical structure on Σ is CAT(0). Since the group W acts on Σ with compact quotient K and finite stabilisers, W is a CAT(0) group.

Let W_T be a special subgroup of W. Then the Cayley graph of W_T (with respect to the generating set T) embeds isometrically in $\mathcal{C}_\Gamma \subset \Sigma$. Hence for each $g \in W$ and each special subgroup W_T of W, left-multiplication of the Cayley graph of W_T by g results in an isometrically embedded copy of the Cayley graph of W_T in $\mathcal{C}_\Gamma \subset \Sigma$, which contains the vertex g. We will refer to this copy as the *Cayley graph of W_T based at g*. For each special subgroup W_T of W, and each coset gW_T, there is also an isometrically embedded copy of Σ_T in Σ. If Θ is an induced subgraph of Γ, we may denote by Σ_Θ the Davis complex for the special subgroup W_Θ, and by $\mathcal{C}_\Theta$ the Cayley graph for W_Θ with generating set the vertices of Θ.

Remark 3.2. Suppose that T is the set of vertices of an embedded four-cycle in Γ, so that $W_T \cong D_\infty \times D_\infty$. Then each copy of Σ_T in Σ is an isometrically embedded copy of the Euclidean plane (tessellated by either big or small squares). Now consider Σ with the cellulation by big squares and let T_1 and T_2 be sets of vertices of embedded four-cycles in Γ such that $W_{T_1 \cup T_2}$ splits over $W_{T_1 \cap T_2} \cong D_\infty \times C_2$. Then each intersection of a copy of the flat Σ_{T_1} with a copy of the flat Σ_{T_2} in Σ is an infinite band of big squares corresponding to a copy of $\Sigma_{T_1 \cap T_2}$. To be precise, this infinite band of big squares is the direct product $\mathbb{R} \times [0,1]$ tessellated by squares of side length 1.

3.3. **Walls.** Consider the Davis complex $\Sigma = \Sigma_\Gamma$ with the cellulation by small squares. Recall that an element $r \in W = W_\Gamma$ is a *reflection* if $r = gsg^{-1}$ for some $g \in W$ and $s \in S$. A *wall* in Σ is defined to be the fixed set of a reflection $r \in W$. For each reflection r, the wall associated to r separates Σ, and r interchanges the two components of the complement. Each wall is a totally geodesic subcomplex of the CAT(0) space Σ, hence each wall is contractible. By the construction of Σ, each wall in Σ is a union of panels, and so is contained in the 1-skeleton of Σ. Hence each wall of Σ is a tree.

We now assign *types* $s \in S$ to the walls. To show that this can be done in a well-defined fashion, suppose first that $gsg^{-1} = s'$, where $g \in W$ and $s, s' \in S$. Fix a reduced word $(s_1, \ldots, s_k)$ for g, and consider the trivial word $\mathbf{s} = (s_1, \ldots, s_k, s, s_k, \ldots, s_1, s')$, which corresponds to the equation $gsg^{-1}s' = 1$. Since $\mathbf{s}$ is non-reduced, by Tits' solution to the word problem for W (see the final paragraph of Section 3.1 above), we must be able to reduce $\mathbf{s}$ to the empty word by a sequence of operations of deleting repeated letters, and swapping ut for tu, where $u, t \in S$ are adjacent vertices. It follows that the number of instances of each letter in $\mathbf{s}$ must be even. Thus $s = s'$; in other words, no two distinct elements of S are conjugate in W. Hence for any reflection $r \in W$, there is a unique $s \in S$ so that $r = gsg^{-1}$ for some $g \in W$. It is thus well-defined to declare the type of the wall which is the fixed set of the reflection $r = gsg^{-1}$ to be s. A wall of type s is a union of panels of type s, and in fact is a maximal connected union of panels of type s. So if each panel of type s is a star-graph of valence $n \geq 2$, each wall of type s will be a $(2, n)$-biregular tree.

For each generator $s \in S$, we denote by H_s the unique wall of type s which contains a panel of the base chamber, and by gH_s, for $g \in W$, the unique translate of the wall H_s which contains a panel of the chamber gK. If H is a wall of type s, then all walls that intersect H are of types which commute with s (and are not equal to s). Since Γ is triangle-free, there are no triples of pairwise intersecting walls. All intersections of walls consist of two walls intersecting at right angles at the centre of some big square, thus subdividing it into four small squares.

3.4. **Paths.** A *path* in $\mathcal{C}_\Gamma$ is a map from an interval (finite or infinite) to $\mathcal{C}_\Gamma$, such that each integer is mapped to a vertex of $\mathcal{C}_\Gamma$ and consecutive integers are mapped to adjacent vertices. Given a path α, we may use $\alpha(i)$ to denote either the image vertex in $\mathcal{C}_\Gamma$ or the group element in W_Γ associated with that vertex.

As noted in Section 3.2, the Cayley graph $\mathcal{C}_\Gamma$ is the 1-skeleton of the cellulation of Σ by big squares. In this cellulation, each edge of $\mathcal{C}_\Gamma$ crosses a unique wall in Σ. Thus the length of a path in $\mathcal{C}_\Gamma$ is equal to its number of wall-crossings (note that a path may cross a given wall more than once). We will sometimes describe paths using the labels of the walls they cross. For example, by the statement "α is the geodesic ray emanating from (or based at) g labelled $a_1a_2a_3\dots$" we will mean that α is a geodesic path such that $\alpha(0) = g$ and $\alpha(i) = ga_1a_2\dots a_i$ for $i > 0$. The path will be a geodesic if each subsegment $a_i \dots a_j$ is reduced. We will often use the fact that a path is a geodesic if and only if it does not cross any wall twice (compare Lemma 3.2.14 and Theorem 3.2.16 of [9]). If α is a geodesic, we will use the notation $\alpha_{[i_1,i_2]}$ to denote the part of α that lies between $\alpha(i_1)$ and $\alpha(i_2)$, including these endpoints. The *support* of a path is the set of labels of the walls that it crosses.

Since Γ is triangle-free, the set of all generators that commute with a given one, say a, generate a special subgroup W_T of W_Γ which is a free product of finitely many copies of C_2. Thus the Cayley graph of W_T (with generating set T) is a tree. Now consider a wall gH_a of type a. There is a copy of the Cayley graph of W_T based at g which runs parallel to the wall gH_a, at constant distance $\frac{1}{2}$ from this wall. We say that a path emanating from g *runs along* the wall gH_a if it is a path in this copy of the Cayley graph of W_T. Equivalently, the path emanates from g and has support contained in the set of generators labelling the link of a in Γ.

Another fact that will be used repeatedly is the following: Suppose γ is a geodesic segment, and η is any path between its endpoints. Let H be a wall that is crossed by γ. Then η crosses H at least once. This is because H (like any wall) separates the Davis complex, and γ, being a geodesic, crosses H exactly once. Thus the endpoints of γ are in different components of the complement of H. Since η is a (continuous) path connecting them, η must cross H.

4. Linear and quadratic divergence in right-angled Coxeter groups

In this section we prove Theorem 1.1 of the introduction. We characterise the defining graphs of 2-dimensional right-angled Coxeter groups with linear and quadratic divergence in Sections 4.1 and 4.2 respectively.

All the graphs Γ considered in this section satisfy our standing assumptions: they are connected, simplicial, triangle-free and have no separating vertices or edges. Recall from Section 2 that the divergence of W_Γ is by definition the divergence of one of its Cayley graphs. We denote by div_Γ the divergence of the Cayley graph $\mathcal{C}_\Gamma \subset \Sigma_\Gamma$. All distances below will be measured in the Cayley graph $\mathcal{C}_\Gamma$, that is, using the word metric on W_Γ with respect to the generating set S, and all paths considered will be in $\mathcal{C}_\Gamma$.

4.1. **Linear divergence.** In this section we prove the following result.

Theorem 4.1. *The divergence* div_Γ *is linear (i.e.* $\mathrm{div}_\Gamma(r) \simeq r$*) if and only if* Γ *is a join.*

As noted in Section 3.1, the graph Γ is a join if and only if W_Γ is reducible (that is, W splits as a direct product of special subgroups). It is proved in [1, Lemma 7.2] that a direct product $H \times K$ has linear divergence if both H and K have the geodesic extension property. This property certainly holds for right-angled Coxeter groups. Thus if Γ is a join, W_Γ has linear divergence.

In Proposition 4.3 below, we prove that when Γ is not a join, the Cayley graph of W_Γ contains a bi-infinite geodesic γ such that $\mathrm{div}_\gamma(r) \succeq r^2$. This completes the proof of Theorem 4.1, as $\mathrm{div}_\Gamma(r) \succeq r^2$ in this case.

Definition 4.2 (The word w and bi-infinite geodesic γ)**.** Recall that the *complementary graph* of Γ, denoted by Γ^c, is the graph with the same vertex set as Γ, in which two vertices are connected by an edge if and only if they are not connected by an edge in Γ. Since Γ is not a join, Γ^c is connected. Choose a loop in Γ^c which visits each vertex (possibly with repetitions). Choose a vertex a_1 on this loop, and let $w = a_1 \cdots a_k$ be the word formed by the vertices of this loop in the order encountered along the loop, where a_k is the last vertex encountered before the loop closes up at its starting point a_1. We assume that the loop is never stationary at a vertex, so that $a_i \neq a_{i+1}$ for any i. Then w is a word in the generators of Γ such that no two consecutive generators commute, and a_k does not commute with a_1. It follows that w^n is reduced for all $n \in \mathbb{Z}$. Let γ be the bi-infinite geodesic in $\mathcal{C}_\Gamma$ which passes through e and is labelled by $\dots wwww \dots$, so that $\gamma(0) = e$, $\gamma(i) = a_1 \cdots a_i$ for $1 \le i \le k$, $\gamma(-1) = a_k$, and so on.

Proposition 4.3. *If Γ is not a join, and γ is the bi-infinite geodesic in $\mathcal{C}_\Gamma$ from Definition 4.2, then $\mathrm{div}_\gamma(r) \succeq r^2$.*

The idea of the proof is similar to that of the corresponding result for right-angled Artin groups in Lemma 7.3 in [1], although we write it in terms of crossings of walls rather than van Kampen diagrams. We include the proof here because it sets the stage for the proof of Proposition 4.9.

Proof. It is enough to obtain a lower bound on $d^{\mathrm{av}}(\gamma(-nk), \gamma(nk))$ as a quadratic function of n (where k is the length of the word w from Definition 4.2). Let η be an arbitrary avoidant path from $\gamma(-nk)$ to $\gamma(nk)$. Since $\gamma_{[-nk,nk]}$ is a geodesic and η is a path with the same endpoints, η must cross each wall crossed by γ at least once. For notational convenience, we will focus on the walls $w^i H_{a_1}$ for $0 \le i \le n-1$ which are crossed by $\gamma_{[0,nk]}$. Now let $(g_i, g_i a_1)$ be the edge of $\mathcal{C}_\Gamma$ at which η first crosses $w^i H_{a_1}$, where g_i is the vertex in the component of the complement of $w^i H_{a_1}$ containing e. Let η_i be the part of η between g_i and g_{i+1} (so that the first edge of η_i is $(g_i, g_i a_1)$).

For $0 \le i \le n-1$, let ν_i denote the geodesic connecting w^i and g_i which runs along $w^i H_{a_i}$, and let H_i be the first wall crossed by ν_i, with type a_j for some j. We claim that H_i does not intersect ν_{i+1}. Since a_j belongs to the support of w, the segment of γ between w^i and w^{i+1} crosses a wall of type a_j. By the construction of w, this wall cannot intersect $w^i H_{a_1}$. It is therefore distinct from H_i and consequently separates H_i from ν_{i+1}. It follows that no subsequent wall crossed by ν_i intersects ν_{i+1} either. Thus each wall crossed by ν_i separates g_i and g_{i+1} into distinct components. Since η_i is a path from g_i to g_{i+1}, it must cross all of these walls. Thus $\ell(\eta_i) \ge \ell(\nu_i) \ge k(n-i)$, and

$$\ell(\eta) \ge \sum_{i=0}^{n-1} \ell(\eta_i) \ge \sum_{i=0}^{n-1} k(n-i) \ge \frac{k}{2} n^2,$$

which completes the proof. □

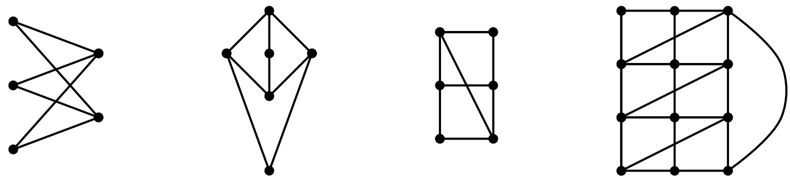

FIGURE 4.1. Some $\mathcal{CFS}$ graphs. (The middle two are actually the same graph.)

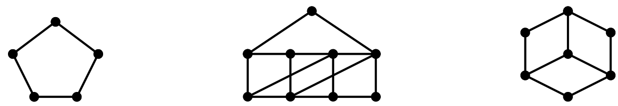

FIGURE 4.2. Some non-$\mathcal{CFS}$ graphs. (The four-cycle graph of the first one is empty. In the second one the four-cycle graph is connected but does not have full support, while in the third, the four-cycle graph has full support, but is not connected and does not have a component with full support.)

4.2. **Quadratic divergence.** We first introduce the $\mathcal{CFS}$ terminology for the graphs which give rise to right-angled Coxeter groups with quadratic divergence. The main result of this section is Theorem 4.6 below.

Given a graph Γ, define the associated *four-cycle graph* Γ^4 as follows. The vertices of Γ^4 are the embedded loops of length four (i.e. *four-cycles*) in Γ. Two vertices of Γ^4 are connected by an edge if the corresponding four-cycles in Γ share a pair of adjacent edges. For example, if Γ is the join $K_{2,3}$, then Γ^4 is a triangle. Given a subgraph Θ of Γ^4, we define the *support* of Θ to be the collection of vertices of Γ (i.e. generators of W_Γ) that appear in the four-cycles in Γ corresponding to the vertices of Θ.

Definition 4.4 ($\mathcal{CFS}$)**.** A graph Γ is said to be $\mathcal{CFS}$ if there exists a component of Γ^4 whose support is the entire vertex set of Γ, i.e., there is a "$\mathcal{C}$omponent with $\mathcal{F}$ull $\mathcal{S}$upport".

Figures 4.1 and 4.2 show some examples of $\mathcal{CFS}$ graphs and non-$\mathcal{CFS}$ graphs respectively. Note that any join is $\mathcal{CFS}$. The last example in Figure 4.1 shows that the four-cycle graph of a $\mathcal{CFS}$ graph need not be connected. However, the following observation will be useful in what follows:

Observation 4.5. The graph Γ is $\mathcal{CFS}$ if and only it if has a subgraph Λ such that Λ^4 is connected, and the support of Λ^4 is the vertex set of Γ. The graph Λ is obtained from Γ by (possibly) deleting some edges, while keeping all the vertices.

We now characterise the graphs which give rise to right-angled Coxeter groups with quadratic divergence.

Theorem 4.6. *The divergence* $\operatorname{div}_\Gamma$ *is quadratic (i.e.* $\operatorname{div}_\Gamma(r) \simeq r^2$*) if and only if* Γ *is* $\mathcal{CFS}$ *and not a join.*

In Proposition 4.7 below we obtain a quadratic upper bound on $\operatorname{div}_\Gamma$ when Γ is a $\mathcal{CFS}$ graph. On the other hand, Proposition 4.3 above shows that if Γ is not a join, then there is a quadratic lower bound on div_Γ. This proves one direction of Theorem 4.6. The other direction follows from Proposition 4.9 below, in which we show that if Γ is not $\mathcal{CFS}$, then $\mathcal{C}_\Gamma$ contains a bi-infinite geodesic whose divergence is at least cubic.

Proposition 4.7. *If* Γ *is* $\mathcal{CFS}$*, then* $\mathrm{div}_\Gamma(r) \preceq r^2$.

Proof. By Example 3.1, a four-cycle in Γ corresponds to a subgroup W' isomorphic to $D_\infty \times D_\infty$. Recall from Section 3.2 that for every $g \in W$, there is an isometrically embedded copy of the Cayley graph of W' based at $g \in \mathcal{C}_\Gamma$. By Theorem 4.1, $\mathrm{div}_{D_\infty \times D_\infty}(r) \simeq r$. In fact it is not hard to see directly that given a pair of geodesic rays α and β emanating from e in $\mathcal{C}_{D_\infty \times D_\infty}$, there is an r-avoidant path connecting $\alpha(r)$ and $\beta(r)$ of length at most $2r$.

Step 1. We first address the case that Γ^4 has a single component. Fix a 4-cycle Θ in Γ and a geodesic ray α emanating from $e \in \mathcal{C}_\Gamma$ whose support is contained in the set of vertex labels of Θ. Thus α lies in the copy of $\mathcal{C}_\Theta$ based at e. We show below that if β is an arbitrary geodesic ray in $\mathcal{C}_\Gamma$ emanating from e, then $\mathrm{div}_{\alpha,\beta}(r) \leq Mr^2$ for every r, where $M = 2\,\mathrm{diam}(\Gamma^4)$. This proves the quadratic upper bound on div_Γ, since it implies that if β_1 and β_2 are arbitrary geodesic rays based at e, then $\mathrm{div}_{\beta_1,\beta_2}(r) \leq 2Mr^2$.

Now let β be an arbitrary geodesic ray labelled $b_1b_2b_3\ldots$ and emanating from e. We first divide $\beta_{[0,r]}$ into *pieces* as follows, and then carry out induction on the number of pieces. Starting at b_1, choose the first piece to be the maximal word $b_1\cdots b_i$ such that $\{b_1, b_2, \ldots, b_i\}$ is contained in the set of vertex labels of a single 4-cycle of Γ. Now repeat this procedure starting at b_{i+1}, and continue until $\beta_{[0,r]}$ is exhausted.

If $\beta_{[0,r]}$ consists of a single piece, then $b_1, \ldots, b_r$ are among the vertices of a single 4-cycle Θ' of Γ. Since Γ^4 is connected, it contains a path connecting the fixed vertex Θ to Θ'. Let $\Theta = \Theta_1, \Theta_2, \ldots, \Theta_l = \Theta'$ be the vertices of Γ^4 along this path. For each $1 \leq i \leq l-1$, since Θ_i and Θ_{i+1} are joined by an edge in Γ^4, the intersection $W_{\Theta_i} \cap W_{\Theta_{i+1}}$ is isomorphic to $W_{\Theta_i \cap \Theta_{i+1}} \cong C_2 \times D_\infty$.

Recall from Remark 3.2 that each Σ_{Θ_i} is an isometrically embedded Euclidean plane tesselated by big squares, and Σ_{Θ_i} and $\Sigma_{\Theta_{i+1}}$ intersect in an infinite band of big squares corresponding to a copy of $\Sigma_{\Theta_i \cap \Theta_{i+1}}$. We proceed below by introducing geodesic rays ν_i based at e, where ν_i lies in $\mathcal{C}_{\Theta_i \cap \Theta_{i+1}} \subset \Sigma_{\Theta_i \cap \Theta_{i+1}}$ for $1 \leq i \leq l-1$. Since successive geodesics in the sequence $\alpha = \nu_0, \nu_1, \ldots, \nu_{l-1}, \nu_l = \beta$ lie in a Euclidean plane, there are linear length avoidant paths between them, and concatenating these gives an avoidant path between α and β.

Let ν denote the geodesic in $\mathcal{C}_{C_2 \times D_\infty}$ that is based at the identity and labelled $g_1g_2g_1g_2\ldots$, where g_1 and g_2 are the generators of the D_∞ factor. For $1 \leq i \leq l-1$, let ν_i denote the image of this geodesic in the copy of $\mathcal{C}_{\Theta_i \cap \Theta_{i+1}}$ based at e in $\mathcal{C}_\Gamma$ (for some identification of g_1 and g_2 with the Coxeter generators of the D_∞ factor of $W_{\Theta_i} \cap W_{\Theta_{i+1}}$). Define $\nu_0 = \alpha$ and $\nu_l = \beta$, and observe that for $1 \leq i \leq l$, the geodesics ν_{i-1} and ν_i are supported on a single 4-cycle of Γ, namely Θ_i. Thus $\nu_{i-1}(r)$ and $\nu_i(r)$ can be connected by an avoidant path of length at most $2r$ in the copy of $\mathcal{C}_{\Theta_i}$ based at e. Concatenating all of these paths, one obtains an r-avoidant path connecting $\alpha(r)$ and $\beta(r)$, with length at most $2rl \leq Mr$, since $l \leq \mathrm{diam}(\Gamma^4)$.

We now induct on the number of pieces of $\beta_{[0,r]}$ to show that $d^{\mathrm{av}}(\alpha(r), \beta(r))$ is at most Mr times the number of pieces. Suppose $\beta_{[0,r]}$ has $k+1$ pieces and is labelled by $w_1 w_2 \dots w_k w_{k+1}$, where each w_i is a piece. Then it is not hard to construct a word w such that:

(1) the support of w is contained in the support of the 4-cycle corresponding to the piece w_k;
(2) the word $w_k w$ is reduced; and
(3) $|w| = |w_{k+1}|$ (so that $|w_1 w_2 \cdots w_k w| = r$), where $|w|$ is the length of w.

It follows that the path μ emanating from e labelled $w_1 w_2 \dots w_k w$ is a geodesic of length r with k pieces. By the inductive hypothesis, there is an r-avoidant path connecting $\alpha(r)$ to $\mu(r)$ of length at most Mkr.

Further, if $s = r - |w_{k+1}|$, then $\beta_{[s,r]}$ and $\mu_{[s,r]}$ are supported on 4-cycles Ψ and Ψ' respectively, and $\beta(s) = \mu(s)$. A more careful version of the construction for the base case yields an r-avoidant path from $\mu(r)$ to $\beta(r)$, as follows. As before, choose a path in Γ^4 which visits the vertices $\Psi = \Psi_1, \Psi_2, \dots, \Psi_m = \Psi'$, and for each i, choose a geodesic ray ν_i emanating from $\beta(s)$ in the copy of $\mathcal{C}_{W_{\Psi_i} \cap W_{\Psi_{i+1}}}$ based at $\beta(s)$, but this time require ν_i to have the additional property that $\beta_{[0,s]}$ concatenated with ν_i is a geodesic. (This will be true for at least one of the two possibilities for ν_i.) Now the construction from the base case (applied with basepoint $\beta(s)$ instead of e) yields a path that avoids not only the ball of radius $|w_{k+1}|$ based at $\beta(s)$, but also the ball of radius r based at e. The length of this path is at most $M|w_{k+1}| \leq Mr$. Concatenating the paths from $\alpha(r)$ to $\mu(r)$ and from $\mu(r)$ to $\beta(r)$, one has the desired r-avoidant path, with length clearly bounded above by $M(k+1)r$.

Finally, since the total number of pieces is bounded above by r, the length of this avoidant path is bounded above by Mr^2.

Step 2. Now suppose that Γ is $\mathcal{CFS}$ but Γ^4 is not connected. Then by Observation 4.5, there exists a subgraph Λ of Γ, such that Λ^4 is connected, and Γ is obtained from Λ by adding edges (between vertices that are at least distance 3 apart in Λ). Since the effect of adding edges is to add more commuting relations in the presentation, there is a natural quotient map $q : W_\Lambda \to W_\Gamma$. Hence if β_1 and β_2 are arbitrary geodesic rays emanating from e in $\mathcal{C}_\Gamma$, they have pullbacks β_1' and β_2' which are geodesic rays emanating from e in $\mathcal{C}_\Lambda$.

We claim that the pushforward of the r-avoidant path constructed in Step 1 between $\beta_1'(r)$ and $\beta_2'(r)$ is r-avoidant in $\mathcal{C}_\Gamma$. The path was constructed by concatenating several subpaths, each of which was r-avoidant in a subgraph $\mathcal{C}_\Psi$, where Ψ is a single four-cycle. The claim follows from the observation that if Ψ is an embedded four-cycle in Λ, then it is an embedded four-cycle in Γ, and the composition of the induced map $q : \mathcal{C}_\Lambda \to \mathcal{C}_\Gamma$ with the inclusion $\mathcal{C}_\Psi \hookrightarrow \mathcal{C}_\Lambda$ is actually an isometric embedding of $\mathcal{C}_\Psi$ into $\mathcal{C}_\Gamma$. □

Remark 4.8. Proposition 4.7 is a special case of the upper bound on divergence given by Theorem 4.9 of [3]. To see this, suppose Γ is $\mathcal{CFS}$ and let $\mathcal{H}$ be the collection of special subgroups of W_Γ generated by the embedded four-cycles in Γ which are the vertices of a component of Γ^4 with full support. Then it is easy to see that W_Γ is strongly algebraically thick of order at most 1 with respect to $\mathcal{H}$. Hence by results in [3], the divergence of W_Γ is at most quadratic. In fact, together with Proposition 4.3 above, one sees that W_Γ is strongly algebraically thick of order exactly equal to 1 if and only if Γ is $\mathcal{CFS}$ but not a join.

We now show that graphs which are not $\mathcal{CFS}$ give rise to right-angled Coxeter groups with super-quadratic divergence. If Γ is not $\mathcal{CFS}$, then, in particular, it is not a join, and there is a word w (of length k) and bi-infinite geodesic γ in $\mathcal{C}_\Gamma$ as described in Definition 4.2. We show that in this setting, the divergence of γ is at least cubic.

Proposition 4.9. *If Γ is not $\mathcal{CFS}$, and γ is the bi-infinite geodesic in $\mathcal{C}_\Gamma$ from Definition 4.2, then* $\operatorname{div}_\gamma(r) \succeq r^3$.

Proof. Let η be an arbitrary avoidant path from $\gamma(-nk)$ to $\gamma(nk)$. We begin exactly as in the first paragraph of the proof of Proposition 4.3 and define the subpaths η_i of η as we did there. However, this time we use the fact that Γ is not $\mathcal{CFS}$ to obtain a quadratic lower bound on $\ell(\eta_i)$. This is a consequence of the following lemma, which is proved separately below.

Lemma 4.10. *Suppose Γ is a graph that is not $\mathcal{CFS}$ and w is the word from Definition 4.2. Let α be an arbitrary geodesic ray emanating from e that travels along H_{a_1} and let β be a path emanating from e consisting of a geodesic segment labelled w followed by an arbitrary geodesic ray emanating from w that travels along wH_{a_1}. Then β is a geodesic, and for any $r > 2k$,*

$$\operatorname{div}_{\alpha,\beta}(r) \geq \frac{1}{16} r^2.$$

Note that γ crosses the wall $w^i H_{a_1}$ at the edge $(w^i, w^i a_1)$. Let ν_i denote the geodesic segment that connects w^i to g_i and runs along $w^i H_{a_1}$. Let μ_i be the path emanating from w^i consisting of the part of γ between w^i and w^{i+1} concatenated with ν_{i+1}. Lemma 4.10, applied with basepoint w^i instead of e, implies that μ_i is a geodesic, and that for $0 \leq i \leq n-2$, and $n > 2$,

$$\ell(\eta_i) \geq d^{\text{av}}_{w^i}(g_i, g_{i+1}) \geq d^{\text{av}}_{w^i}(\nu_i(kn - ki), \mu_i(kn - ki)) \geq \frac{k^2}{16}(n-i)^2.$$

For the middle inequality above, we use the observation in the last paragraph of Section 2. In conclusion,

$$\ell(\eta) \geq \sum_{i=0}^{n-2} \ell(\eta_i) \geq \sum_{i=0}^{n-2} \frac{k^2}{16}(n-i)^2.$$

This is a cubic function of n. □

Proof of Lemma 4.10. We first show that β is a geodesic ray. Since $\beta_{[0,k]}$ (which is labelled by w) and $\beta_{[k,\infty]}$ are geodesics, the only way β can fail to be a geodesic is if there is a wall which intersects both of these. Recall that $w = a_1 \cdots a_k$, so that the walls crossed by $\beta_{[0,k]}$ are $\beta(i-1)H_{a_i}$ for $1 \leq i \leq k$, where $\beta(0) = e$ and $\beta(i) = a_1 \cdots a_i$. By construction, a_i and a_{i+1} don't commute for any i (mod k), so it follows that these walls are pairwise disjoint, and are all disjoint from wH_{a_1}. On the other hand every wall that intersects $\beta_{[k,\infty]}$ necessarily crosses wH_{a_1}, since $\beta_{[k,\infty]}$ runs along wH_{a_1}. It follows that no wall can cross both $\beta_{[0,k]}$ and $\beta_{[k,\infty]}$. Similarly, since α is a geodesic emanating from e along the wall H_{a_1}, the same argument shows that no wall can cross both $\beta_{[0,k]}$ and α, a fact that will be useful later in this proof.

To obtain a lower bound on $\operatorname{div}_{\alpha,\beta}$, choose an arbitrary r-avoidant path η between $\alpha(r)$ and $\beta(r)$. Then one obtains a loop in $\mathcal{C}_\Gamma$ by concatenating $\alpha_{[0,r]}$, followed by η, and followed by $\beta_{[0,r]}$ traversed in the negative direction. There is a van Kampen diagram D with boundary label equal to the word encountered along this loop.

Note that by construction, $\alpha_{[0,r]}$, $\beta_{[0,r]}$ and η do not have any common edges in $\mathcal{C}_\Gamma$. It follows that every edge of ∂D is part of a 2-cell of D, and that D is homeomorphic to a disk. We will abuse notation and use α, β and η to denote the parts of ∂D that are labelled by these paths.

There is a label-preserving combinatorial map from D to Σ_Γ with the cellulation by big squares. Under this map, edges and vertices of D go to edges and vertices of $\mathcal{C}_\Gamma$, which is the 1-skeleton of the cellulation by big squares. We may assume that each 2-cell in D is a square, since any 2-cell with boundary label of the form s^2 maps to an edge of Σ_Γ and can therefore be collapsed to an edge in D. Thus the map takes each 2-cell of D homeomorphically to a big square of Σ_Γ. Further, if we metrise each square of D as $[0,1] \times [0,1]$, then we can arrange that the restriction of this map to a square of D is an isometry onto its image of a big square.

We will work primarily with a cell structure on D that is dual to the one just described. We first define *walls of* D, record some of their properties, and then use them to define the dual structure on D. The dual structure is then used to divide D into *strips*, and we will show that the length of a strip is a lower bound on the length of η. We then finish the proof by inductively estimating the lengths of the strips. The fact that Γ is not $\mathcal{CFS}$ is used to show that the lengths of strips grow quadratically.

Walls of D**.** Recall that each big square in Σ_Γ is subdivided into four small squares by a pair of (segments of) walls which intersect at the centre of the big square. For each square in D, we pull back this pair of segments to D, and label them with the type of walls they came from. The types of the two wall-segments in a square of D are necessarily distinct. Now suppose there are two squares in D which share an edge ϵ. By construction, both squares contain a wall-segment that intersects ϵ at its midpoint, and these wall-segments must have the same label. To see this, recall that the image of ϵ in Σ_Γ is the side of a big square, and the midpoint of such a side cannot be the point of intersection of a pair of walls. Thus, starting at any wall-segment in a square of D, one can continue it through adjacent squares until it eventually meets ∂D. We call a path constructed in this way a *wall of* D, and the type of wall is the type of any of its wall-segments. Walls are similar to corridors: if one "fattens up" a wall of type a by taking the union of the squares containing its individual wall-segments, then one has an a-corridor of D.

Two walls of D intersect each other at most once; they intersect only if their types commute and are distinct. A wall of D cannot intersect itself, as this would require there to be a square in D in which both wall-segments have the same type. Thus each wall of D is an embedded interval connecting a pair of points on ∂D. We record the following observation for future use.

Observation 4.11. Every wall of D has at least one endpoint on η.

To see this, recall from the first paragraph of this proof that in Σ_Γ, and therefore in D, any wall intersecting $\beta_{[0,k]}$ is disjoint from both $\beta_{[k,r]}$ and $\alpha_{[0,r]}$. Thus any wall in D with an endpoint on $\beta_{[0,k]}$ has its other endpoint on η. Now suppose there is a wall P in D with one endpoint on α and the other on $\beta_{[k,r]}$. Then P separates D, putting $\beta_{[0,k]}$ and η in different components. This implies that every wall with an endpoint on $\beta_{[0,k]}$ intersects P. However, one of these walls has the same type as P, since $\beta_{[0,k]}$, which is labelled by w, has as its support the full vertex set of Γ. This is a contradiction.

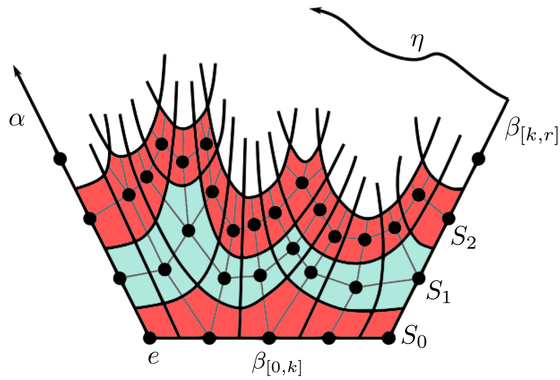

FIGURE 4.3. The van Kampen diagram D. The light edges and bold dots are the 1-cells and 0-cells, respectively, of the original cell structure on D. The walls of D, which bound the dual 2-cells, are shown in bold. The strips are shaded.

The dual cell structure on D**.** We now define the dual structure on D. Its 1-skeleton is the union of the walls of D, together with ∂D; see Figure 4.3. Thus the vertices are points of intersection of a pair of walls (i.e. centres of squares in the original structure) or points of intersection of a wall with ∂D. Removing the vertices from the 1-skeleton yields several components; the edges are the closures of these components. The 2-cells are the closures of the components of the complement of the 1-skeleton in D. We use the terms *dual cells* and *original cells* to distinguish between cells from the two structures on D. A dual cell is called a *boundary cell* if it intersects ∂D. Otherwise it is called an *interior cell*. Since Γ is a triangle-free graph, it is easy to see that the boundary of any interior dual 2-cell is a polygon with at least four sides.

Strips in D**.** We now use the dual structure to define strips S_i in D, for $0 \leq i < (r-k)/2$.

Define the 0th strip S_0 to be the union of all the dual 2-cells intersecting $\beta_{[0,k]}$. Define the *top boundary* B_0 of S_0, by $B_0 = \partial S_0 \setminus \partial D$. Let $\epsilon_{\alpha(j)}$ (respectively $\epsilon_{\beta(j)}$) denote the dual edge of ∂D containing the original vertex $\alpha(j)$ (respectively $\beta(j)$). Observe that:

(1) S_0 is connected and consists of an ordered collection of dual 2-cells, each intersecting the previous one in a dual edge, and going from $\epsilon_{\alpha(0)}$ to $\epsilon_{\beta(k)}$.
(2) If Q is a wall that forms part of B_0, then S_0 is contained in a single component of $D \setminus Q$.
(3) B_0 is connected, and all but the first and last dual edges of B_0 are interior edges.

Note that (1) follows from the fact that every edge of $\beta_{[0,k]}$ is part of a 2-cell, and that D is homeomorphic to a disk. If (2) fails, then Q crosses S_0 and has an endpoint on $\beta_{[0,k]}$. On the other hand, since it is part of B_0, it contributes to the boundary of a boundary 2-cell, and two of the boundary edges of this 2-cell are parts of walls P_1 and P_2 which intersect $\beta_{[0,k]}$. In order to intersect S_0, the wall Q must cross either P_1 or P_2. This is a contradiction, since by construction, no two walls with endpoints on $\beta_{[0,k]}$ intersect each other. Finally, (3) follows from (1), together with the fact that the construction forces B_0 to consist solely of parts of walls.

Now suppose S_{i-1} and its top boundary B_{i-1} have been defined, with properties analogous to (1)–(3) above. In particular, the 2-cells of S_{i-1} go from $\epsilon_{\alpha(i-1)}$ to $\epsilon_{\beta(k+i-1)}$. Define S_i to be the union of all the dual 2-cells intersecting B_{i-1}. Then S_i contains the dual 2-cells whose boundaries contain the edges $\epsilon_{\alpha(i)}$ and $\epsilon_{\beta(k+i)}$. Define the top boundary B_i to be $\partial S_i \setminus \{B_{i-1}, \epsilon_{\alpha(i)}, \epsilon_{\beta(k+i)}\}$. We claim that if $i < (r-k)/2$, then S_i has properties analogous to (1)–(3) above.

To see (1), note that property (1) for S_{i-1} implies that S_{i-1}, and therefore B_{i-1}, separates D. Let D_i be the closure of the component of $D \setminus B_{i-1}$ not containing S_{i-1} (so that ∂D_i consists of B_{i-1} and a part of ∂D). By property (3) for B_{i-1}, all but the first and last dual edges of B_{i-1} are interior edges of D, so every edge of B_{i-1} is part of a 2-cell in D_i and D_i is homeomorphic to a disk. It follows that S_i is connected and consists of an ordered collection of dual 2-cells, each intersecting the previous one in a dual edge, going from $\epsilon_{\alpha(i)}$ to $\epsilon_{\beta(k+i)}$.

An argument involving intersections of walls similar to the S_0 case proves property (2) for S_i.

Property (3) would fail for B_i if one of the dual 2-cells of S_i other than the first and the last is a boundary cell, as this would mean that B_i contains part of $\alpha_{[i+1,r]}$, $\beta_{[k+i+1,r]}$, or η. (Note that $\alpha_{[0,i-1]}$ and $\beta_{[0,k+i-1]}$ cannot be part of B_i since B_{i-1} separates S_i from these parts of ∂D.)

We first rule out $\alpha_{[i+1,r]}$ and $\beta_{[k+i+1,r]}$. Let A_i denote the wall of D with an endpoint at the intersection of $\epsilon_{\alpha(i)}$ and $\epsilon_{\alpha(i+1)}$. Note that $\alpha_{[i+1,r]}$ cannot cross A_i by construction. Now A_i is a part of B_i, so by property (2) for S_i it separates $\alpha_{[i+1,r]}$ from S_i. This implies that S_i cannot have any boundary cells intersecting $\alpha_{[i+1,r]}$. By the same argument, S_i does not have any boundary cells intersecting $\beta_{[k+i+1,r]}$.

The map from D to $\mathcal{C}_\Gamma$ takes each original vertex contained in a dual cell of S_i into $B(e, k+2i) \subset \mathcal{C}_\Gamma$. To see this observe that each original vertex of S_0 is mapped into $B(e,k)$, and for $j > 0$, the image of an original vertex in S_j is at most distance two from the image of the vertices of S_{j-1}. So if $i < (r-k)/2$, then the original vertices of S_i are mapped into $B(e, r-1)$, and therefore cannot be vertices of η, which is r-avoidant. Thus S_{i-1} does not have any boundary cells intersecting η. This shows that all but the first and last 2-cells of S_i are interior cells, which implies (3).

Lengths of strips. Define the *length* of S_i, denoted $\ell(S_i)$, to be the number of interior dual 2-cells in it.

Claim 4.12. For $i < (r-k)/2$, we have $\ell(S_i) \leq \ell(\eta)$.

Proof. Let P be a wall of D which is *transverse* to S_i, meaning that it crosses S_i at least once, intersecting both B_{i-1} and B_i. We now show that P crosses S_i at most twice. Further, the number of times P crosses S_i is equal to the number of endpoints of P on η.

Suppose P crosses S_i at least twice. Starting at the endpoint of P on η (guaranteed by Observation 4.11), follow P until its second crossing of S_i, and let Q denote the top boundary wall at the second crossing. By property (2) for S_i, we know that S_i is contained in a single component of $D \setminus Q$. Thus, in order to cross S_i again, P

FIGURE 4.4. A strip which has no large 2-cells is either a single row of squares (left) or a sequence of such rows (right).

would have to cross Q a second time, which is impossible. So P crosses S_i at most twice.

Now suppose the second endpoint of P is on α. Since Q can cross neither P (a second time) nor S_i, it must also have an endpoint on α. This is a contradiction, since by construction, two walls with endpoints on α cannot intersect each other. By the same argument, P cannot have an endpoint on β. Thus, if P crosses S_i twice, it has two endpoints on η. If P crosses S_i exactly once, then Observation 4.11 and the fact that S_i separates D putting η in a single component imply that P has exactly one endpoint on η, completing the proof of the second statement above.

Thus there is an injective map from the set of transverse intersections of walls with S_i into the set of walls crossed by η in ∂D. This proves the claim, as the number of such transverse intersections is $\ell(S_i) + 1$, and the number of walls crossed by η in ∂D (and therefore in Σ_Γ) is $\ell(\eta)$. □

Lower bounds. We now inductively obtain lower bounds on the lengths of strips. Define an interior dual 2-cell to be *large* if its boundary has five or more sides.

Claim 4.13. Every strip has at least one large 2-cell.

Proof. If not, then there is a strip S_i built entirely out of squares. There are two possibilities: either this strip consists of a single row of squares, or it consists of a sequence of such rows of squares, with each such row connected to the next at right angles as in Figure 4.2.

Since two walls intersect only if the corresponding generators commute, it is possible to reconstruct a subgraph of Γ using S_i, as follows. The vertices of this subgraph are the labels of the walls which meet S_i (either transversely or as part of B_i or B_{i-1}). We add an edge between two such vertices of Γ whenever the corresponding walls intersect in S_i. It is easy to see that a single row of squares reconstructs a join subgraph, while a sequence of rows of squares meeting at right angles reconstructs a $\mathcal{CFS}$ subgraph. Every wall which has an endpoint on $\beta_{[0,k]}$ crosses S_i, since by Observation 4.11 its other endpoint is on η, and S_i separates η from $\beta_{[0,r]}$. Thus the generators corresponding to walls with endpoints on $\beta_{[0,k]}$ are vertices of the subgraph constructed above. But the support of $\beta_{[0,k]}$ is the entire vertex set of Γ, so we obtain a $\mathcal{CFS}$ subgraph of Γ which uses all the vertices of Γ. Then Γ itself is $\mathcal{CFS}$, by Observation 4.5. This is a contradiction. □

An interior dual 2-cell in S_i intersects S_{i-1} in either an edge or a vertex. Define the 2-cell to be *skew* if this intersection is a vertex. Let u_i denote the number of skew 2-cells in S_i.

Claim 4.14. For $1 \leq i < (r-k)/2$, we have $u_i \geq i$.

Proof. To see that $u_1 \geq 1$, note that B_0 cannot consist of a single wall, by Observation 4.11. So it contains at least one pair of walls that intersect at a point and

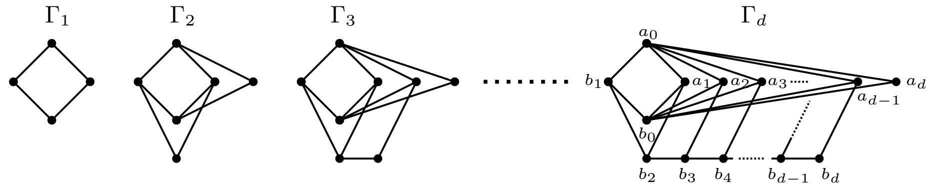

FIGURE 5.1

then pass through S_1, giving rise to a skew 2-cell whose closure intersects S_0 in the point of intersection of the walls.

For the inductive step, observe that a skew 2-cell in S_{i-1} whose boundary is a j-gon gives rise to $j-3$ skew 2-cells in S_i. Since each interior dual 2-cell has at least 4 sides, $j-3 \geq 1$. Similarly, a non-skew large 2-cell in S_{i-1} whose boundary is a j-gon gives rise to $j-4 \geq 1$ skew 2-cells in S_i. By Claim 4.13, every strip has at least one large 2-cell. Now if one of the skew cells in S_{i-1} is large, it gives rise to at least two skew cells in S_i, and we have $u_i \geq u_{i-1}+1$. Otherwise there is a non-skew large cell in S_{i-1}, which gives rise to a skew cell in S_i which does not come from a skew cell of S_{i-1}, and we have the same relation. It follows that $u_i \geq i$ for $1 \leq i < (r-k)/2$. □

There is a map from the 2-cells of S_i to the 2-cells of S_{i-1} defined as follows. The image of a skew 2-cell c is the unique 2-cell in S_{i-1} which shares a vertex with c. The image of a non-skew 2-cell c is the unique 2-cell of S_{i-1} which shares an edge with c. This is surjective by property (1) for S_i. The cardinality of the preimage is at least 1 for a non-skew 2-cell, and at least 3 for a skew 2-cell of S_{i-1}. Thus one has the relation $\ell(S_i) \geq \ell(S_{i-1}) + 2u_{i-1}$, since the length of a strip is the number of interior 2-cells in it. Then, using Claim 4.14, we have:

$$\ell(S_i) \geq \ell(S_{i-1}) + 2u_{i-1} \geq \cdots \geq \sum_{j=1}^{i} 2u_j \geq 2\sum_{j=1}^{i} j \geq (i)(i+1) \geq i^2.$$

Finally, if $r > 2k$, then $r/4 < (r-k)/2$, and by Claim 4.12, we have $\ell(\eta) \geq \ell(S_{r/4}) \geq \frac{1}{16}r^2$. □

5. HIGHER-DEGREE POLYNOMIAL DIVERGENCE IN RIGHT-ANGLED COXETER GROUPS

We now prove Theorem 1.2 of the introduction, by producing examples to show that the divergence of a 2-dimensional right-angled Coxeter group can be a polynomial of any degree. More precisely, if Γ_d is the sequence of graphs shown in Figure 5.1 ($d \geq 1$), then we show that $\mathrm{div}_{\Gamma_d}(r) \simeq r^d$. We prove the upper and lower bounds on $\mathrm{div}_{\Gamma_d}(r)$ in Propositions 5.1 and 5.3 respectively.

Proposition 5.1. $\mathrm{div}_{\Gamma_d}(r) \preceq r^d$.

Proof. Observe that the statement is true for $d=1$ and 2, as Γ_1 is a join and Γ_2 is a $\mathcal{CFS}$ graph. We proceed by induction on d. Assume that there is a constant C such that if μ and ν are arbitrary geodesic rays based at e in $\mathcal{C}_{\Gamma_{d-1}}$, then $d^{\mathrm{av}}(\mu(r), \nu(r)) \leq Cr^{d-1}$ for any r.

Now let α and β be an arbitrary pair of geodesic rays based at e in $\mathcal{C}_{\Gamma_d}$. If neither of them crosses any walls of type a_d or b_d, then they actually lie in the copy of $\mathcal{C}_{\Gamma_{d-1}}$ based at e, and the induction hypothesis yields the desired avoidant path.

Thus we may assume that at least one of them, say α, crosses a wall of type a_d or b_d. Let $H_1, \dots, H_k$ be the ordered set of walls of type a_d or b_d that α crosses between e and $\alpha(r)$, and let x_i denote the type of H_i. Then the label on $\alpha_{[0,r]}$ is $w_1 x_1 w_2 x_2 \cdots w_k x_k w_{k+1}$, where each w_i is a (possibly empty) word in the letters $a_0, a_1, \dots, a_{d-1}, b_0, b_1, \dots, b_{d-1}$, and each x_i is a_d or b_d. For $1 \leq i \leq k$, let g_i denote the word $w_1 x_1 w_2 x_2 \cdots w_i$. Then there exists a geodesic ray λ_i emanating from g_i with the following properties:

(1) The path emanating from e consisting of the segment labelled g_i followed by λ_i is a geodesic.
(2) The geodesic λ_i runs along H_i. (That is, the support of λ_i is either $\{a_0, b_0\}$ or $\{a_{d-1}, b_{d-1}\}$, depending on whether x_i is a_d or b_d, respectively.)

If $x_i = a_d$, the label of λ_i must be of the form $a_0 b_0 a_0 b_0 \dots$ or $b_0 a_0 b_0 a_0 \dots$. Choose the former if the projection of g_i to the group $\langle a_0, b_0 \rangle$ ends with b_0 and the latter otherwise. This guarantees that there is no cancellation when g_i is concatenated with the label of λ_i. The case $x_i = b_d$ is similar.

For $1 \leq i \leq k$, let ν_i be the geodesic ray emanating from $g_i x_i$ with the same label as λ_i. (See Figure 5.2.) For $0 \leq i \leq k-1$, let μ_i be the geodesic ray emanating from $g_i x_i$ (or e when $i = 0$) consisting of the segment with label w_{i+1} followed by λ_{i+1}. The choice of the λ_i guarantees that these are geodesics. Finally, define μ_k to be the infinite part of α emanating from $g_k x_k$.

If β does not cross any walls of type a_d or b_d, then define $\mu'_0 = \beta$. Otherwise define $H'_1, \dots, H'_l$, as well as x'_i, g'_i, u'_i, ν'_i, and μ'_i analogous to the corresponding objects for α.

By construction, the supports of $\nu_i, \mu_i, \nu'_i, \mu'_i$ are contained in $\{a_0, a_1, \dots, a_{d-1}, b_0, b_1, \dots, b_{d-1}\}$. Thus there exist paths η_i connecting $\nu_i(2r)$ and $\mu_i(2r)$ with length at most $C(2r)^{d-1}$, which avoid a ball of radius $2r$ based at $g_i x_i$, and therefore avoid a ball of radius r based at e. Similarly, there are r-avoidant paths η'_i and η_0 connecting $\nu'_i(2r)$ and $\mu'_i(2r)$ and $\mu_0(2r)$ and $\mu'_0(2r)$ respectively, each with length at most $C(2r)^{d-1}$.

For each i, the points $\mu_i(2r)$ and $\nu_{i+1}(2r)$ are connected by an edge, as are $\mu'_i(2r)$ and $\nu'_{i+1}(r)$. Using these $k+l$ edges to connect η_i, η'_i and η_0, one obtains an r-avoidant path between $\mu_k(2r)$ and $\mu'_k(2r)$. Finally, η is constructed by attaching the segment of α from $\alpha(r)$ to $\mu_k(2r)$ and the segment of β from $\beta(r)$ to $\mu'_k(2r)$, each with length at most $2r$. Since k and l are at most r, we have

$$\ell(\eta) \leq 4r + k + l + (k+l+1)C(2r)^{d-1} \leq 6r + (2r+1)C2^{d-1}r^{d-1} \leq C'r^d,$$

where $C' = 6 + 2^{d+1}C$. □

Remark 5.2. This upper bound could also be obtained by arguments in [3], as the group W_d is strongly algebraically thick of order at most $d-1$. To see this, for each $n \geq d \geq 1$ define a right-angled Coxeter group $W_{n,d}$ to be the special subgroup of W_n generated by the set $\{a_0, a_1, \dots, a_n, b_0, b_1, \dots, b_d\}$. Note that $W_{d,d} = W_d$. Now $W_{n,2}$ is strongly algebraically thick of order at most 1 since its defining graph is $\mathcal{CFS}$ (see Remark 4.8 above). By induction on d, the group $W_{n,d}$ is strongly algebraically thick of order at most $d-1$ with respect to $\mathcal{H} = \{W_{n,d-1},\, b_d W_{n,d-1} b_d\}$. Hence in particular, W_d is strongly algebraically thick of order at most $d-1$.

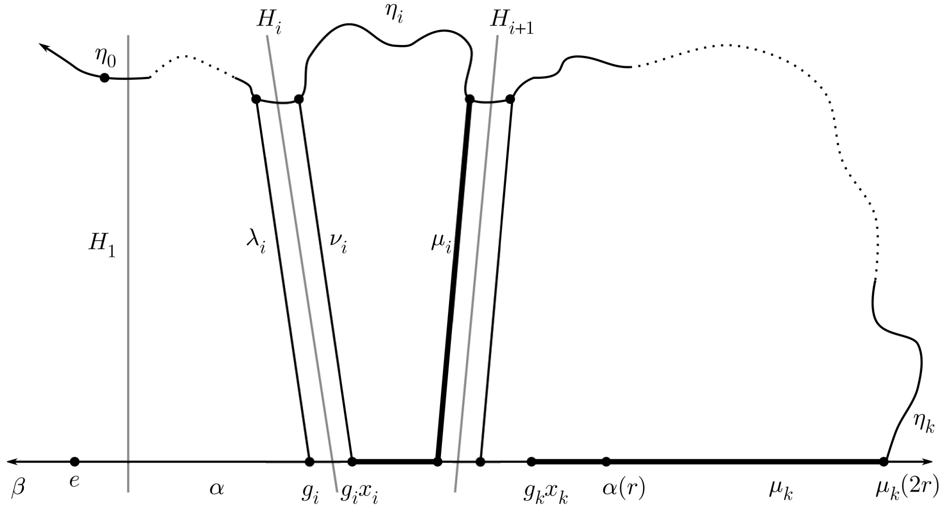

FIGURE 5.2. Construction of the avoidant path η. The geodesic rays μ_i and μ_k are shown in bold.

Proposition 5.3. $\mathrm{div}_{\Gamma_d}(r) \succeq r^d$.

Proof. We prove the lower bound by producing a pair of geodesic rays in $\mathcal{C}_{\Gamma_d}$ whose divergence is bounded below by a constant multiple of r^d. This will follow from a more general statement about the divergence of certain pairs of geodesics in $\mathcal{C}_{\Gamma_{d+2}}$.

For $1 \leq n \leq d$, let α_n and β_n be any geodesic rays in $\mathcal{C}_{\Gamma_{d+2}}$ satisfying the following conditions:

(1) α_n emanates from e and travels along $H_{b_{n+1}}$; and
(2) β_n emanates from e and travels along one of H_{a_n}, H_{b_n}, or $H_{b_{n+2}}$. (Note that $\{a_n, b_n, b_{n+2}\}$ is exactly the set of types of walls which can intersect $H_{b_{n+1}}$.)

Then we show below that

$$d^{\mathrm{av}}(\alpha_n(r), \beta_n(r)) \geq \frac{1}{2^{n(n+1)}} r^n. \tag{5.1}$$

When $n = d$, one can take α_d to be the geodesic ray based at e with label $b_d a_d b_d a_d \dots$, as this travels along $H_{b_{d+1}}$, and β_d to be the geodesic ray based at e with label $b_{d-1}a_{d-1}b_{d-1}a_{d-1}\dots$, as this travels along H_{b_d}. Observe that these geodesics are actually in the copy of $\mathcal{C}_{\Gamma_d}$ based at e. Any avoidant path between $\alpha_d(r)$ and $\beta_d(r)$ in $\mathcal{C}_{\Gamma_d}$ remains avoidant under the isometric inclusion $\mathcal{C}_{\Gamma_d} \hookrightarrow \mathcal{C}_{\Gamma_{d+2}}$, and therefore has length bounded below by $(1/2^{d(d+1)})r^d$, by (5.1). This completes the proof of the proposition.

We establish (5.1) by proving the following equivalent statement by induction on k: for all $1 \leq k \leq d$ and all $k \leq n \leq d$, if α_n and β_n satisfy the conditions (1) and (2) respectively, then $d^{\mathrm{av}}(\alpha_n(r), \beta_n(r)) \geq (1/2^{k(k+1)})r^k$.

Observe that for any n, if α_n and β_n are chosen as above, then α_n concatenated with β_n at e is a bi-infinite geodesic, since α_n and β_n have disjoint supports, regardless of the type of wall along which β_n travels. Thus any avoidant path between

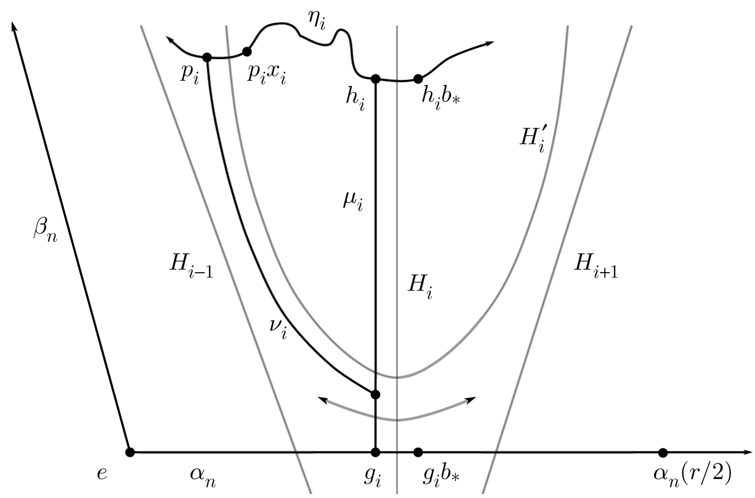

FIGURE 5.3. Construction of μ_i and ν_i. Here $\mu_i(0) = g_i$ and $\nu_i(0) = \mu_i(1)$.

$\alpha_n(r)$ and $\beta_n(r)$ must cross the $2r$ walls crossed by this bi-infinite geodesic. This proves the case $k = 1$, as $d^{\text{av}}(\alpha_n(r), \beta_n(r)) \geq 2r > (1/4)r$ for all $n \geq 1$.

Now suppose $n \geq k+1$, and let η be an avoidant path connecting $\beta_n(r)$ to $\alpha_n(r)$. Focus on the $r/2$ walls that α_n crosses between $\alpha_n(0)$ and $\alpha_n(r/2)$. Each of these is of type a_n, b_n, or b_{n+2}. Since two consecutive walls cannot be of the same type, at most half of these walls are of type a_n. Thus, in this range, α_n (and hence η) crosses at least $r/4$ walls of type b_*, where the subscript is either n or $n+2$. Call them $H_1, \ldots, H_l$, where $l \geq r/4$. Let (g_i, g_ib_*) be the edge where α_n crosses H_i and let (h_i, h_ib_*) be the first edge where η crosses H_i, going from $\beta_n(r)$ to $\alpha_n(r)$. Let μ_i denote the unique geodesic connecting g_i to h_i that travels along H_i. (See Figure 5.3.) Define $\mu_0 = \beta_n$ and $h_0 = \beta_n(r)$.

For $1 \leq i \leq l$, let H_i' denote the second wall crossed by μ_i starting at g_i. Note that H_i' intersects H_i, and therefore cannot intersect α_n, since no two walls crossed by α_n intersect. We claim that H_i' also does not intersect μ_{i-1}. The support of μ_i is contained in either $\{a_{n-1}, b_{n-1}, b_{n+1}\}$ or $\{a_{n+1}, b_{n+1}, b_{n+3}\}$, depending on the type of H_i. If the first wall crossed by μ_i doesn't intersect μ_{i-1}, then H_i' can't either, as H_i' is separated from μ_{i-1} by the first wall. Otherwise, the first wall crossed by μ_i has to be of type b_{n+1}, which means that H_i' is not of type b_{n+1}. If the types of H_i and H_{i-1} are different, then the type of H_i' is not in the support of μ_{i-1}, so H_i' cannot intersect μ_{i-1}. Finally, if the types of H_i and H_{i-1} are the same, then they must be separated by a wall of type a_n, since α_n can't cross two consecutive walls of the same type. Now H_i' can't intersect this wall, since it is not of type a_0, b_0 or b_{n+1}. So H_i' can't intersect μ_{i-1} either.

It follows that for $1 \leq i \leq l$, the wall H_i' separates the points h_{i-1} and h_i, since the path formed by concatenating μ_{i-1}, the part of α_n between g_{i-1} and g_i, and μ_i crosses H_i' exactly once. Now η contains a sub-path connecting h_{i-1} and h_i, so η must cross H_i'. Let (p_i, p_ix_i) be the first edge along which it crosses H_i', where x_i is the type of H_i'. Let η_i denote the part of η between p_i and h_i, and let ν_i denote the unique geodesic connecting $\mu_i(1)$ to p_i, which travels along H_i'.

Observe that μ_i is a geodesic that travels along a wall of type b_n or b_{n+2}, and ν_i is a geodesic that travels along a wall that intersects it. This means that the pair μ_i and ν_i is either of the form α_{n-1} and β_{n-1} or α_{n+1} and β_{n+1} (if we allow the geodesics to emanate from $\mu(1)$ instead of e). Since $n-1 \geq k$, the inductive hypothesis applies, and we have that $d^{\mathrm{av}}_{\mu(1)}(\mu_i(s+1), \nu_i(s)) \geq (1/2^{k(k+1)})s^k$ for all s. Since we are restricted to the walls $H_1, \ldots, H_l$ crossed by α_n between e and $\alpha_n(r/2)$, we know that $|g_i| \leq r/2$. On the other hand, since h_i and p_i are r-avoidant, the lengths of μ_i and ν_i are at least $r/4$. By the observation and the end of Section 2, $\ell(\eta_i) \geq (1/2^{k(k+1)})(r/4)^k$ for all i. So, since $l \geq 4$, we have

$$\ell(\eta) \geq \sum_{i=1}^{l} \ell(\eta_i) \geq l\left(\frac{1}{2^{k(k+1)}}\right)\left(\frac{r}{4}\right)^k \geq \left(\frac{r}{4}\right)\left(\frac{r^k}{2^{k(k+1)+k}}\right) = \frac{1}{2^{(k+1)(k+2)}} r^{k+1},$$

as required. □

Appendix A. Relationship with examples of Macura

In this appendix we discuss the relationship between our constructions of CAT(0) groups with divergence polynomial of any degree, and those of Macura [19].

For $d \geq 2$, we denote by G_d the group constructed in [19] with presentation

$$G_d = \langle a_0, a_1, \ldots, a_d \mid a_0 a_1 = a_1 a_0 \text{ and } a_i^{-1} a_0 a_i = a_{i-1} \text{ for } 2 \leq i \leq d \rangle.$$

Let X_d be the presentation 2-complex for this presentation of G_d. Then X_d has a single vertex v, $d+1$ oriented edges labeled by $a_0, a_1, \ldots, a_d$, and d squares with boundary labels $a_0 a_1 a_0^{-1} a_1^{-1}$ and $a_i^{-1} a_0 a_i a_{i-1}^{-1}$ for $2 \leq i \leq d$. Equip X_d with the metric such that each square is a unit Euclidean square. Then the universal cover $\widetilde{X_d}$ is a CAT(0) square complex, in which the link of every vertex is the graph Γ_d from Figure 5.1 above. The link of any vertex in the Davis complex for $W_d = W_{\Gamma_d}$ with the cellulation by big squares is also Γ_d. This observation is why we consider the relationship between G_d and W_d. To avoid confusion with Macura's notation, in this section we relabel the vertices of Γ_d by $s_{i^+} = a_i$ and $s_{i^-} = b_i$ for $0 \leq i \leq d$.

We would like to use covering theory to investigate common finite index subgroups of G_d and W_d. Any finite index subgroup of G_d is the fundamental group of a finite square complex Q such that there is a combinatorial covering map $\Psi : Q \to X_d$. However, since the group W_d has torsion a more sophisticated covering theory is needed; as we explain below, its finite index subgroups correspond to finite-sheeted covers of complexes of groups. We first recall some background on complexes of groups in Section A.1. We then use this theory to show in Section A.2 that W_2 and G_2 are commensurable, and to explain in Section A.3 why for $d > 2$ we do not know if W_d and G_d are commensurable.

A.1. **Complexes of groups.**

We adapt the theory of complexes of groups and their coverings to our situation. The general theory and details can be found in [7, Chapter III.$\mathcal{C}$]. Throughout this section, $W = W_\Gamma$ is a right-angled Coxeter group with Γ satisfying the hypotheses of Theorem 1.1, and Σ is the associated Davis complex.

Let Y be a square complex. Assume that the edges of Y may be oriented so that:

- (∗) for each square of Y, if the positively oriented edge labels of this square are a, b, a' and b', then $b' a' a^{-1} b^{-1}$ is the boundary label.

For an oriented edge e of Y, we denote by $i(e)$ its initial vertex and by $t(e)$ its terminal vertex.

Examples A.1. Two important examples of square complexes with edge orientations satisfying $(*)$ are the following:

(1) Let Y be the chamber K with the cellulation by small squares. For all pairs of spherical subsets $T' \subsetneq T$, we orient the edge of Y connecting the vertices $\sigma_{T'}$ and σ_T so that this edge has initial vertex $\sigma_{T'}$ and terminal vertex σ_T. Note that every edge incident to $\sigma_\emptyset$ has initial vertex $\sigma_\emptyset$.
(2) Similarly, if $Y = \Sigma$ with the cellulation by small squares, then the edges of Σ may be oriented by inclusion of type.

Now suppose that Y and Z are square complexes with edge orientations satisfying $(*)$.

Definition A.2. A *non-degenerate morphism* $f : Y \to Z$ is a map taking vertices to vertices and edges to edges, such that:

(1) for each square of Y, the restriction of f to this square is a bijection onto a square of Z; and
(2) for each vertex σ of Y, the restriction of f to the set of edges with initial vertex σ is a bijection onto the set of edges of Z with initial vertex $f(\sigma)$.

For example, if $Y = \Sigma$ and $Z = K$ with the orientations specified in Examples A.1 above, then the quotient map $f : Y \to Z$ induced by the action of W on Σ is a non-degenerate morphism.

Definition A.3. Let Y be a square complex with edge orientations satisfying $(*)$. A *complex of groups* $\mathcal{G}(Y) = (G_\sigma, \psi_e)$ over Y consists of:

(1) a group G_σ for each vertex σ of Y, called the *local group* at σ; and
(2) a monomorphism $\psi_e : G_{i(e)} \to G_{t(e)}$ along each edge e of Y.

A complex of groups is *trivial* if each local group is trivial.

Example A.4. We construct a canonical complex of groups $\mathcal{W}(K)$ over K as follows. For each spherical subset $T \in \mathcal{S}$, the local group at the vertex σ_T is the special subgroup W_T. All monomorphisms along edges are inclusions.

The complex of groups $\mathcal{W}(K)$ in Example A.4 is canonically induced by the action of W on Σ. More generally, if G is a subgroup of W, then the action of G on Σ induces a complex of groups $\mathcal{G}(Y)$ over $Y = G\backslash\Sigma$, such that for each vertex σ of Y, the G-stabiliser of each lift $\overline{\sigma}$ of σ in Σ is a conjugate of the local group G_σ of $\mathcal{G}(Y)$. A complex of groups is *developable* if it is isomorphic to a complex of groups induced by a group action. Complexes of groups, unlike graphs of groups, are not in general developable.

See [7] for the definition of the *fundamental group* $\pi_1(\mathcal{G}(Y))$ and *universal cover* of a (developable) complex of groups $\mathcal{G}(Y)$. The universal cover of $\mathcal{G}(Y)$ is a connected, simply-connected square complex X, equipped with an action of $G = \pi_1(\mathcal{G}(Y))$ so that $Y = G\backslash X$.

Examples A.5.

(1) The complex of groups $\mathcal{W}(K)$ has fundamental group W and universal cover Σ.

(2) Let $\mathcal{G}(Y)$ be the trivial complex of groups over a square complex Y. Then $\pi_1(\mathcal{G}(Y))$ is the (topological) fundamental group of Y, and $\pi_1(\mathcal{G}(Y))$ acts freely on the universal cover of $\mathcal{G}(Y)$.

If a complex of groups $\mathcal{G}(Y)$ is developable, then each local group G_σ naturally embeds in the fundamental group $\pi_1(\mathcal{G}(Y))$.

We now discuss coverings of complexes of groups. We will only need to construct coverings $\mathcal{G}(Y) \to \mathcal{W}(K)$ where $\mathcal{G}(Y)$ is a trivial complex of groups, and so do not give the general definition, which is considerably more complicated.

Definition A.6. Let Y be a square complex with edge orientations satisfying $(*)$. Let $\mathcal{G}(Y)$ be the trivial complex of groups over Y. A *covering* of complexes of groups $\Phi : \mathcal{G}(Y) \to \mathcal{W}(K)$ consists of:

(1) a non-degenerate morphism $f : Y \to K$; and
(2) for each edge e of Y, with $f(t(e)) = \sigma_T$, an element $\phi(e) \in W_T$;

such that for each vertex σ of Y and each edge e' of K, with $t(e') = f(\sigma) = \sigma_T$ and $i(e') = \sigma_{T'}$, the map

$$\Phi_{\sigma/e'} : \{e \in f^{-1}(e') \mid t(e) = \sigma\} \to W_T/W_{T'}$$

induced by $e \mapsto \phi(e)$ is a bijection.

Observe that if e' is an edge of K with $t(e') = \sigma_T$ and $i(e') = \sigma_{T'}$, then $|T| = |T'| + 1$; hence if $T = T' \cup \{t\}$ we have $W_T/W_{T'} \cong \langle t \rangle \cong C_2$. So the condition in Definition A.6 that $\Phi_{\sigma/e'}$ is a bijection is equivalent to the condition that the set $\{e \in f^{-1}(e') \mid t(e) = \sigma\}$ has two elements, say e_1 and e_2, such that without loss of generality $\phi(e_1) \in W_{T'}$ and $\phi(e_2) \in tW_{T'}$. In particular, it suffices to put $\phi(e_1) = 1$ and $\phi(e_2) = t$. A covering $\Phi : \mathcal{G}(Y) \to \mathcal{W}(K)$ as in Definition A.6 is *finite-sheeted* if Y is finite.

The following result is a special case of a general theorem on functoriality of coverings of complexes of groups. The general result is implicit in [7], and stated and proved explicitly in [18].

Theorem A.7. *Let K_d be the chamber for W_d, cellulated by small squares. Let $\mathcal{W}(K_d)$ be the complex of groups over K_d described in Example* A.4 *above, with fundamental group W_d. Then any subgroup of W_d is the fundamental group of a complex of groups $\mathcal{G}(Y')$ (not necessarily trivial) over a square complex Y', such that there is a covering of complexes of groups $\Phi : \mathcal{G}(Y') \to \mathcal{W}(K_d)$. Moreover, a subgroup of W_d has finite index if and only if it is the fundamental group of $\mathcal{G}(Y')$ such that there is a finite-sheeted covering $\Phi : \mathcal{G}(Y') \to \mathcal{W}(K_d)$.*

A.2. **Commensurability in the case $d = 2$.** We now use covering theory to prove the following.

Proposition A.8. *The groups G_2 and W_2 are commensurable.*

Proof. Denote by Z_2 the first square subdivision of the presentation 2-complex X_2. We will construct a finite square complex Y such that:

(1) there is a combinatorial covering map $\Psi : Y \to Z_2$; and
(2) there is a covering of complexes of groups $\Phi : \mathcal{G}(Y) \to \mathcal{W}(K_2)$, where $\mathcal{G}(Y)$ is the trivial complex of groups over Y.

Since $\mathcal{G}(Y)$ is the trivial complex of groups, the fundamental group of $\mathcal{G}(Y)$ is just the (topological) fundamental group of Y. It follows that G_2 and W_2 are commensurable.

The square complex Y will be the first square subdivision of the square complex Q constructed below. We will show that there is an 8-sheeted combinatorial covering map from Q to X_2, which implies (1). See Figure A.1; the complex Q is obtained by carrying out some further edge identifications on this square complex.

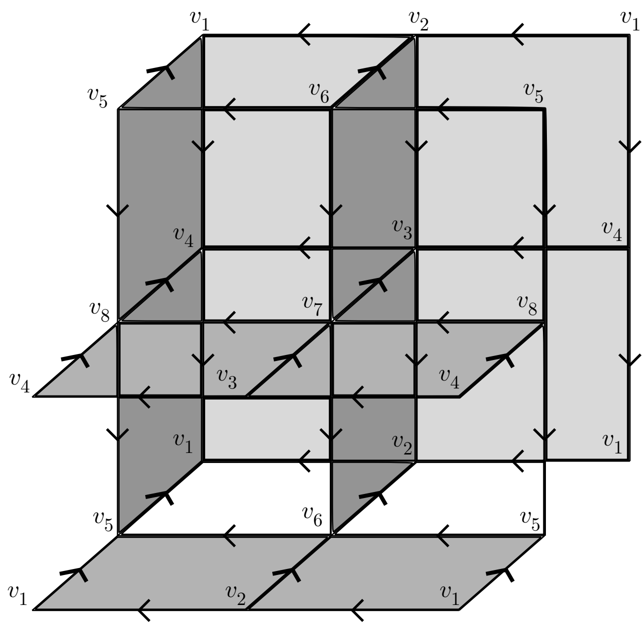

FIGURE A.1. The square complex Q, with vertices labelled and edges oriented, prior to some edge identifications. All squares except for the four squares with vertex set $\{v_5, v_6, v_7, v_8\}$ are shaded.

The complex Q has 8 vertices $v_1, \ldots, v_8$, each of which get mapped to the vertex v of X_2. There are 24 oriented edges of Q which form three families as follows. Here, $a_{i,j} = (v_k, v_l)$ means that the edge $a_{i,j}$ is the unique edge of Q with initial vertex v_k and terminal vertex v_l.

(1) The following 8 edges get mapped to the edge a_0 of X_2: $a_{0,1} = (v_1, v_2)$, $a_{0,2} = (v_2, v_1)$, $a_{0,3} = (v_4, v_3)$, $a_{0,4} = (v_3, v_4)$, $a_{0,5} = (v_6, v_7)$, $a_{0,6} = (v_7, v_6)$, $a_{0,7} = (v_5, v_8)$, $a_{0,8} = (v_8, v_5)$.

(2) The following 8 edges get mapped to the edge a_1 of X_2: $a_{1,1} = (v_1, v_4)$, $a_{1,2} = (v_4, v_1)$, $a_{1,3} = (v_2, v_3)$, $a_{1,4} = (v_3, v_2)$, $a_{1,5} = (v_5, v_6)$, $a_{1,6} = (v_6, v_5)$, $a_{1,7} = (v_8, v_7)$, $a_{1,8} = (v_7, v_8)$.

(3) The following 8 edges get mapped to the edge a_2 of X_2: $a_{2,1} = (v_1, v_5)$, $a_{2,2} = (v_5, v_1)$, $a_{2,3} = (v_4, v_8)$, $a_{2,4} = (v_8, v_4)$, $a_{2,5} = (v_3, v_7)$, $a_{2,6} = (v_7, v_3)$, $a_{2,7} = (v_2, v_6)$, $a_{2,8} = (v_6, v_2)$.

TABLE 1. Types of vertices in Y which are midpoints of edges of Q

Midpoint of edges	Type	Midpoint of edges	Type	Midpoint of edges	Type
$a_{1,1}, a_{1,3}, a_{0,5}, a_{0,7}$	0^+	$a_{0,1}, a_{0,3}, a_{1,5}, a_{1,7}$	1^+	$a_{2,2}, a_{2,4}, a_{2,6}, a_{2,8}$	2^+
$a_{1,2}, a_{1,4}, a_{0,6}, a_{0,8}$	0^-	$a_{0,2}, a_{0,4}, a_{1,6}, a_{1,8}$	1^-	$a_{2,1}, a_{2,3}, a_{2,5}, a_{2,7}$	2^-

We then attach 16 squares along the following edge labels, forming two families as follows:

(1) The following 8 squares get mapped to the square of X_2 attached along $a_0 a_1 a_0^{-1} a_1^{-1}$:

$$a_{0,1}a_{1,3}a_{0,3}^{-1}a_{1,1}^{-1}, \quad a_{0,2}a_{1,1}a_{0,4}^{-1}a_{1,3}^{-1}, \quad a_{0,3}a_{1,4}a_{0,1}^{-1}a_{1,2}^{-1}, \quad a_{0,4}a_{1,2}a_{0,2}^{-1}a_{1,4}^{-1},$$

$$a_{0,7}a_{1,7}a_{0,5}^{-1}a_{1,5}^{-1}, \quad a_{0,8}a_{1,5}a_{0,6}^{-1}a_{1,7}^{-1}, \quad a_{0,5}a_{1,8}a_{0,7}^{-1}a_{1,6}^{-1}, \quad a_{0,6}a_{1,6}a_{0,8}^{-1}a_{1,8}^{-1}.$$

(2) The following 8 squares get mapped to the square of X_2 attached along $a_2^{-1} a_0 a_2 a_1^{-1}$:

$$a_{2,1}^{-1}a_{0,1}a_{2,7}a_{1,5}^{-1}, \quad a_{2,7}^{-1}a_{0,2}a_{2,1}a_{1,6}^{-1}, \quad a_{2,3}^{-1}a_{0,3}a_{2,5}a_{1,7}^{-1}, \quad a_{2,5}^{-1}a_{0,4}a_{2,3}a_{1,8}^{-1},$$

$$a_{2,2}^{-1}a_{0,7}a_{2,4}a_{1,1}^{-1}, \quad a_{2,4}^{-1}a_{0,8}a_{2,2}a_{1,2}^{-1}, \quad a_{2,8}^{-1}a_{0,5}a_{2,6}a_{1,3}^{-1}, \quad a_{2,6}^{-1}a_{0,6}a_{2,8}a_{1,4}^{-1}.$$

This completes the construction of Q, together with a combinatorial covering $Q \to X_2$.

Now let Y be the first square subdivision of Q and let $\mathcal{G}(Y)$ be the trivial complex of groups over Y. We assign types $T \in \mathcal{S}$ to the vertices of Y, as follows. If a vertex of Y is one of the vertices of Q, it has type $\emptyset$. Next consider the vertices of Y which are midpoints of edges of Q. Table 1 shows the assigned types of these vertices. To simplify notation, we write $i^\pm$ for the type $\{s_{i^\pm}\} \in \mathcal{S}$, for $i = 0, 1, 2$.

Finally, consider the vertices of Y which are at the centres of squares of Q. Let σ be such a vertex. Then for some pair of types i^{ε_i} and j^{ε_j} with $i, j \in \{0, 1, 2\}$, $i \neq j$, and $\varepsilon_i, \varepsilon_j \in \{\pm\}$, two of the vertices of Y which are adjacent to σ are of type i^{ε_i}, and two of the vertices of Y which are adjacent to σ are of type j^{ε_j}. Moreover, $\{i^{\varepsilon_i}, j^{\varepsilon_j}\} \in \mathcal{S}$. We then assign type $\{i^{\varepsilon_i}, j^{\varepsilon_j}\}$ to the vertex σ.

After assigning these types, it may be verified that Y is obtained by taking 8 copies of the chamber K_2 and gluing together certain pairs of mirrors of the same type. We note also that the above assignment of types allows us to orient the edges of Y in the same way as in K_2, that is, an edge a has initial vertex of type T' and terminal vertex of type T if and only if $T' \subsetneq T$.

Next, define $f : Y \to K_2$ to be the only possible type-preserving morphism. It may be checked that f is a non-degenerate morphism. We construct a covering of complexes of groups $\Phi : \mathcal{G}(Y) \to \mathcal{W}(K_2)$ over f. In order to define the elements $\phi(a)$ for the edges a of Y, we put an equivalence relation, parallelism, on the set of edges of Y, so that if a and b are parallel, then we will have $\phi(a) = \phi(b)$. The relation is generated by saying that two edges are *parallel* if they are opposite edges of a (small) square of Y. The values of $\phi(a)$ for representatives a of certain parallelism classes of edges in Y are specified in Table 2. For all edges a of Y which are *not* parallel to an edge appearing in Table 2, we put $\phi(a) = 1$.

To verify that Φ is a covering of complexes of groups, we simplify notation and write s for the vertex $\sigma_{\{s\}}$ of the chamber K_2. For each vertex s_{i^ε} of K_2, where

TABLE 2. Non-trivial values of $\phi(a)$, for representatives a of certain parallelism classes of edges

Vertex $i(a)$	Type of $t(a)$	$\phi(a)$	Vertex $i(a)$	Type of $t(a)$	$\phi(a)$
v_1	0^+	s_{0^+}	v_1	2^-	s_{2^-}
v_4	0^-	s_{0^-}	v_4	2^-	s_{2^-}
v_1	1^+	s_{1^+}	v_5	2^+	s_{2^+}
v_2	1^-	s_{1^-}	v_6	2^+	s_{2^+}

$i \in \{0,1,2\}$ and $\varepsilon \in \{\pm\}$, there is a unique edge b of K_2 such that s_{i^ε} is the terminal vertex of b. Fix a vertex $\sigma \in f^{-1}(s_{i^\varepsilon})$. Then there are two edges a_1 and a_2 of Y with terminal vertex σ such that $f(a_1) = f(a_2) = b$. By construction, without loss of generality we have $\phi(a_1) = s_{i^\epsilon}$ and $\phi(a_2) = 1$. Therefore $\Phi_{\sigma/b}$ is a bijection to $\langle s_{i^\varepsilon}\rangle \cong C_2$, as required. Now consider a vertex σ_T of K_2 where $T \in \mathcal{S}$ with $|T| = 2$. Write $T = \{i^{\varepsilon_i}, j^{\varepsilon_j}\}$. If b is an edge of K_2 with terminal vertex σ_T, then without loss of generality b has initial vertex of type i^{ε_i}. Fix a vertex $\sigma \in f^{-1}(\sigma_T)$. Then there are two edges a_1 and a_2 of Y with terminal vertex σ such that $f(a_1) = f(a_2) = b$. By construction, without loss of generality we have $\phi(a_1) = s_{j^{\epsilon_j}}$ and $\phi(a_2) = 1$. Thus $\Phi_{\sigma/b}$ is a bijection to $W_T/\langle s_{i^{\varepsilon_i}}\rangle \cong \langle s_{j^{\varepsilon_j}}\rangle \cong C_2$, as required. Therefore Φ is a covering of complexes of groups. □

A.3. **Discussion of case $d > 2$.** We conclude with a brief discussion of whether G_d and W_d are commensurable when $d > 2$.

The following result shows that the strategy used to prove that G_2 and W_2 are commensurable cannot be implemented for $d > 2$. The proof of Proposition A.9 uses covering theory for complexes of groups, and may be found in the first version of this paper on the arXiv. We denote by Z_d the first square subdivision of X_d.

Proposition A.9. *If $d > 2$ there is no square complex Y such that both of the following conditions hold:*

(1) *there is a combinatorial covering map $\Psi : Y \to Z_d$; and*
(2) *there is a covering of complexes of groups $\Phi : \mathcal{G}(Y) \to \mathcal{W}(K_d)$.*

Note that Proposition A.9 does not require Y to be finite. In particular, it follows that:

Corollary A.10. *For $d > 2$, the universal cover $\widetilde{X_d}$ is not isometric to the Davis complex Σ_d.*

This is surprising, since both $\widetilde{X}_d$ and Σ_d are CAT(0) square complexes with all vertex links the graph Γ_d.

For $d > 2$ we do not know if G_d and W_d are commensurable, or even quasi-isometric. If they are commensurable, then there are finite square complexes Y and Y' such that there is a combinatorial covering map $Y \to Z_d$ and a covering of complexes of groups $\mathcal{G}(Y') \to \mathcal{W}(K_d)$, with $\pi_1(Y) \cong \pi_1(\mathcal{G}(Y'))$. Moreover, since G_d is torsion-free it is not hard to show that such a $\mathcal{G}(Y')$ must be the trivial complex of groups over Y', hence $\pi_1(\mathcal{G}(Y')) = \pi_1(Y')$. Now by Corollary A.10, the universal covers of Y and Y' are not isometric. Hence if there is some Mostow-type rigidity result which implies that the isomorphism $\pi_1(Y) \cong \pi_1(Y')$ is induced by an isometry of universal covers, we would obtain that W_d and G_d are not in fact commensurable. However, the only Mostow-type rigidity results for CAT(0) square complexes that we know of are Theorem 1.4.1 of [8], for certain uniform lattices on products of trees, and Corollary 1.8 of [6], concerning right-angled Artin groups, and neither of these results can be applied here.

Acknowledgements

We thank the University of Sydney for travel support. We also thank the organisers of the 2012 Park City Mathematics Institute Summer Program on Geometric Group Theory, during which part of this work was undertaken, and an anonymous referee for helpful comments.

References

[1] Aaron Abrams, Noel Brady, Pallavi Dani, Moon Duchin, and Robert Young, *Pushing fillings in right-angled Artin groups*, J. Lond. Math. Soc. (2) **87** (2013), no. 3, 663–688, DOI 10.1112/jlms/jds064. MR3073670
[2] Jason Behrstock and Ruth Charney, *Divergence and quasimorphisms of right-angled Artin groups*, Math. Ann. **352** (2012), no. 2, 339–356, DOI 10.1007/s00208-011-0641-8. MR2874959
[3] J. Behrstock and C. Drutu, *Divergence, thick groups, and short conjugators*, (2011).
[4] J. Behrstock and M. F. Hagen, *Cubulated groups: thickness, relative hyperbolicity, and simplicial boundaries*, (2012).
[5] Jason A. Behrstock, *Asymptotic geometry of the mapping class group and Teichmüller space*, Geom. Topol. **10** (2006), 1523–1578, DOI 10.2140/gt.2006.10.1523. MR2255505 (2008f:20108)
[6] Mladen Bestvina, Bruce Kleiner, and Michah Sageev, *The asymptotic geometry of right-angled Artin groups. I*, Geom. Topol. **12** (2008), no. 3, 1653–1699, DOI 10.2140/gt.2008.12.1653. MR2421136 (2009f:20061)
[7] Martin R. Bridson and André Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 319, Springer-Verlag, Berlin, 1999. MR1744486 (2000k:53038)
[8] Marc Burger, Shahar Mozes, and Robert J. Zimmer, *Linear representations and arithmeticity of lattices in products of trees*, Essays in geometric group theory, Ramanujan Math. Soc. Lect. Notes Ser., vol. 9, Ramanujan Math. Soc., Mysore, 2009, pp. 1–25. MR2605353 (2011k:20050)
[9] Michael W. Davis, *The geometry and topology of Coxeter groups*, London Mathematical Society Monographs Series, vol. 32, Princeton University Press, Princeton, NJ, 2008. MR2360474 (2008k:20091)
[10] Michael W. Davis and Tadeusz Januszkiewicz, *Right-angled Artin groups are commensurable with right-angled Coxeter groups*, J. Pure Appl. Algebra **153** (2000), no. 3, 229–235, DOI 10.1016/S0022-4049(99)00175-9. MR1783167 (2001m:20056)
[11] W. Dison and T. Riley, *Hydra groups*, Comment. Math. Helv. **88** (2013), no. 3, 507–540. MR3093501
[12] Cornelia Druţu, Shahar Mozes, and Mark Sapir, *Divergence in lattices in semisimple Lie groups and graphs of groups*, Trans. Amer. Math. Soc. **362** (2010), no. 5, 2451–2505, DOI 10.1090/S0002-9947-09-04882-X. MR2584607 (2011d:20084)
[13] Moon Duchin and Kasra Rafi, *Divergence of geodesics in Teichmüller space and the mapping class group*, Geom. Funct. Anal. **19** (2009), no. 3, 722–742, DOI 10.1007/s00039-009-0017-3. MR2563768 (2011c:30115)
[14] S. M. Gersten, *Divergence in* 3*-manifold groups*, Geom. Funct. Anal. **4** (1994), no. 6, 633–647, DOI 10.1007/BF01896656. MR1302334 (95h:57013)
[15] S. M. Gersten, *Quadratic divergence of geodesics in* CAT(0) *spaces*, Geom. Funct. Anal. **4** (1994), no. 1, 37–51, DOI 10.1007/BF01898360. MR1254309 (95h:53057)
[16] M. Gromov, *Asymptotic invariants of infinite groups*, Geometric group theory, Vol. 2 (Sussex, 1991), London Math. Soc. Lecture Note Ser., vol. 182, Cambridge Univ. Press, Cambridge, 1993, pp. 1–295. MR1253544 (95m:20041)
[17] M. Kapovich and B. Leeb, 3*-manifold groups and nonpositive curvature*, Geom. Funct. Anal. **8** (1998), no. 5, 841–852, DOI 10.1007/s000390050076. MR1650098 (2000a:57040)

[18] Seonhee Lim and Anne Thomas, *Covering theory for complexes of groups*, J. Pure Appl. Algebra **212** (2008), no. 7, 1632–1663, DOI 10.1016/j.jpaa.2007.10.012. MR2400734 (2009d:57002)

[19] Nataša Macura, *CAT(0) spaces with polynomial divergence of geodesics*, Geom. Dedicata **163** (2013), 361–378, DOI 10.1007/s10711-012-9754-9. MR3032700

Department of Mathematics, Louisiana State University, Baton Rouge, Louisiana 70803

School of Mathematics and Statistics, University of Glasgow, 15 University Gardens, Glasgow G12 8QW, United Kingdom