5-2018

Divergence of CAT(0) Cube Complexes and Coxeter Groups

Ivan Levcovitz

The Graduate Center, City University of New York

DIVERGENCE OF CAT(0) CUBE COMPLEXES AND COXETER GROUPS

Ivan Levcovitz

A dissertation submitted to the Graduate Faculty in Mathematics in partial fulfillment of the requirements for the degree of Doctor of Philosophy, The City University of New York, The Graduate Center.

2018

DIVERGENCE OF CAT(0) CUBE COMPLEXES AND COXETER GROUPS

Ivan Levcovitz

This manuscript has been read and accepted for the Graduate Faculty in Mathematics in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

Date	Jason Behrstock Chair of Examining Committee
Date	Ara Basmajian Executive Officer

Supervisory Committee:

Jason Behrstock

Abhijit Champanerkar

Olga Kharlampovich

THE CITY UNIVERSITY OF NEW YORK

Abstract

DIVERGENCE OF CAT(0) CUBE COMPLEXES AND COXETER GROUPS

Ivan Levcovitz

Advisor: Jason Behrstock

We provide geometric conditions on a pair of hyperplanes of a CAT(0) cube complex that imply divergence bounds for the cube complex. As an application, we characterize right-angled Coxeter groups with quadratic divergence and show right-angled Coxeter groups cannot exhibit a divergence function between quadratic and cubic. This generalizes a theorem of Dani-Thomas that addressed the class of 2–dimensional right-angled Coxeter groups. This characterization also has a direct application to the theory of random right-angled Coxeter groups. As another application of the divergence bounds obtained for cube complexes, we provide an inductive graph theoretic criterion on a right-angled Coxeter group's defining graph which allows us to recognize arbitrary integer degree polynomial divergence for many infinite classes of right-angled Coxeter groups. We also provide similar divergence results for some classes of Coxeter groups that are not right-angled. Finally, we discuss thick structures on right-angled Coxeter groups and show that for n larger than 1, there are right-angled Coxeter groups that are thick of order n but are algebraically thick of strictly larger order, answering a question of Behrstock-Druţu-Mosher.

Acknowledgments

I would first like to thank my advisor Jason Behrstock. He inspired me to pursue interesting problems and to keep at them, even when the path to their resolution was not in sight. He taught me to think and write about mathematics more clearly and how to discover new problems of my own (taught a man how to fish). It has been an honor to learn from a great mathematician. This thesis is a product of his patience and time commitment to me.

I thank Abhijit Champanerkar and Olga Kharlampovich for serving on my defense committee. Their questions and suggestions have pointed me towards new directions. I thank the many professors at the Graduate Center for passing their knowledge down to me including my committee members, Ara Basmajian, Khalid Bou-Rabee, Joseph Maher, Ilya Kofman and many others. They have shaped me into the mathematician I am today. I thank the Graduate Center for its support for my time there.

I would also like to thank Pallavi Dani, Mark Hagen, Tim Susse and Anne Thomas for very helpful conversations regarding this thesis topic and for their excellent suggestions.

I am lucky to have had the opportunity to learn from the many other students at my department. These include, and are not limited to, Chris Arretines, Daniel Berlyne, Jacob Russell-Madonia and Hai Yu.

Thank you to my family and many friends for always being supportive and for an often healthy distraction from mathematics. Lastly, I owe infinite thanks to my wonderful partner Jillian. She has been there throughout the highs and lows of my doctorate and has never ceased to be supportive. This thesis would also not have been possible without her.

The results in Sections 3 - 7 and Section 9 are to appear Algebraic & Geometric Topology [Leva]. Results in Section 8 appear in [Levb].

Contents

List of Figures

Chapter 1

Introduction

Geometric group theory explores the interplay between finitely generated, infinite groups and the geometry of spaces on which they act. In this theory, quasi-isometries provide a natural notion for the coarse equivalence of metric spaces. Two spaces are quasi-isometric if one can be mapped onto the other by a function differing from an isometry by an additive and multiplicative constant. Much of a group's algebra and geometry can be recovered from only the quasi-isometry class of its Cayley graph. With this in mind, Gromov proposed a program to study groups up to quasi-isometry [Gro93].

In this thesis, we study a natural quasi-isometry invariant, the divergence function. In particular, we explore the divergence of CAT(0) cube complexes and apply our results to the class of right-angled Coxeter groups. We also provide results regarding the divergence of the more general Coxeter groups.

Given a metric space X and a positive number r, the *divergence function* $\mathrm{Div}(X, r) = \mathrm{Div}(X)$ is the supremum over all lengths of minimal paths, which avoid a ball of radius r, connecting two points that are distance roughly r apart. One may roughly think of the divergence function as a measure of the best upper bound on the rate a pair of geodesic rays can stray apart from one another. For a finitely generated group G, $\mathrm{Div}(G)$ is the divergence function applied to the Cayley graph of G endowed with the word metric.

Groups which are δ–hyperbolic, an important class of groups possessing properties of negative curvature, all exhibit at least exponential divergence. On the other end of the spectrum, the divergence of $\mathbb{Z}^n$ is linear for $n \geq 2$. In some sense, the divergence function measures the presence of negative curvature in a given group.

Gromov conjectured that groups with non-positive curvature, such as those which act geometrically on a CAT(0) space, should exhibit either linear or exponential divergence [Gro93]. This turns out not to be the case. Many important classes of groups such as 3–manifold groups, the mapping class group of a closed surface of genus $g \geq 2$ and right-angled Artin groups have been shown to exhibit quadratic divergence [Ger94, KL98, Beh06, BC12]. More recently, CAT(0) groups exhibiting polynomial divergence of degree d, for any integer d have been found [BD14, BH16, Mac13] and, in particular, these include right-angled Coxeter groups [DT15]. Additionally, there are constructions of more exotic infinitely presented groups (although not CAT(0)) with divergence function not a polynomial

[GS, OOS09].

Right-angled Artin groups have played a central role in contemporary mathematics. Incredibly, Agol and Wise show that the fundamental group of every hyperbolic 3–manifold is virtually a subgroup of a right-angled Artin group [Wis11, Ago13]. In terms of their divergence, these groups satisfy a certain trichotomy: each right-angled Artin group either exhibits linear, quadratic or infinite divergence with these occurrences classified by simple properties of the group's defining graph [BC12]. The fundamental groups of 3–manifolds exhibit a similar trichotomy as well [Ger].

Right-angled Coxeter groups form another large class of CAT(0) groups. Associated to any simplicial graph Γ is a right-angled Coxeter group, W_Γ, whose presentation consists of an order 2 generator for each vertex of Γ with the relation that two generators commute if there is an edge between the corresponding vertices of Γ.

Despite their simple presentation, right-angled Coxeter groups form a wide class of groups. For one, both free groups and free abelian groups are finite index subgroups of a right-angled Coxeter group. Furthermore, every right-angled Artin group is finite index in some right-angled Coxeter group [DJ00], and the converse is not true. As divergence is a quasi-isometry invariant, the class of right-angled Coxeter groups contains groups of linear, quadratic and infinite divergence. However, even more is true for these groups.

In fact, for any positive integer degree, Dani-Thomas surprisingly provide an example of a 2–dimensional right-angled Coxeter group exhibiting polynomial divergence of the given degree [DT15]. This raises the question of which divergence functions are possible for right-angled Coxeter groups. Additionally, there is the question of which properties of right-angled Coxeter groups give rise to their broader spectrum of divergence functions and how can these properties be recognized through these groups' defining graphs.

The class of groups that act geometrically on a CAT(0) cube complex is vast and includes both right-angled Artin groups and right-angled Coxeter groups. More generally, we ask which properties of CAT(0) cube complexes give rise to different divergence functions.

The Rank Rigidity Theorem shows the existence of a rank one isometry in an irreducible, essential, locally compact CAT(0) cube complex with cocompact automorphism group [CS11] (the result actually holds under more general assumptions as well). As a consequence, the divergence of these spaces is either linear or at least quadratic [Hag].

We prove for the case of right-angled Coxeter groups, there is an additional gap between quadratic and cubic divergence, and we characterize exactly which right-angled Coxeter groups exhibit quadratic divergence.

Theorem 7.3.1 *Suppose the graph Γ is not a nontrivial join. The right-angled Coxeter group W_Γ exhibits quadratic divergence if and only if Γ is CFS. If Γ is not CFS, then the divergence of W_Γ is at least cubic.*

The CFS condition ("constructed from squares") is a purely graph-theoretic condition which can be computationally checked. We say a join graph is nontrivial if both graphs in the join decomposition are not cliques. We note that the divergence of W_Γ is linear if and only if Γ is a non-trivial join [BFRHS]. For the case when Γ does not contain triangles, the above theorem is a result of Dani-Thomas [DT15]. Such groups are precisely those whose Davis complex, a natural CAT(0) space a Coxeter group acts on, is 2–dimensional. Theorem 7.3.1 thus generalizes Dani-Thomas's result to right-angled Coxeter groups of arbitrary dimension.

An important application of Theorem 7.3.1 is to the theory of random right-angled Coxeter groups. Let $\Gamma(n, p(n))$ be a random n-vertex graph containing an edge between a given pair of vertices with probability $p(n)$. A random right-angled Coxeter group is simply the right-angled Coxeter group defined by a random graph. Behrstock–Falgas-Ravry–Hagen–Susse [BFRHS] give a threshold theorem for when a random graph is CFS with probability 1. Combining their result with Theorem 7.3.1, we obtain a threshold function for the transition between quadratic to at least cubic divergence in random right-angled Coxeter groups.

Theorem 1.0.1 (Behrstock–Falgas-Ravry–Hagen–Susse, Levcovitz)**.** *Suppose* $p(n)$ *is a probability density function bounded away from 1 and let* $\epsilon > 0$*. Let* $\Gamma = \Gamma(p(n), n)$ *be a random graph. If* $p(n) > n^{-\frac{1}{2}+\epsilon}$*, then the right-angled Coxeter group* W_Γ *asymptotically almost surely exhibits quadratic divergence. If* $p(n) < n^{-\frac{1}{2}-\epsilon}$*, then* W_Γ *asymptotically almost surely exhibits at least cubic divergence.*

We note that a random right-angled Coxeter group given as above asymptotically almost surely has dimension larger than two. Thus the generality of Theorem 7.3.1 is needed to obtain Theorem 1.0.1.

Strongly thick metric spaces of order d form an important class of spaces which can be constructed through a d–step inductive gluing procedure, with initial pieces of linear divergence. An important consequence is that these spaces must have divergence bounded above by a polynomial of degree $d + 1$ [BD14]. There are not many general results in the opposite direction giving lower bounds on divergence, and a goal of this thesis is to introduce criteria which imply such lower bounds. Having such criteria then allows us to give the exact divergence, up to an equivalence of functions, for many spaces.

We apply the following strategy to study the divergence in CAT(0) cube complexes. First, we define the *hyperplane divergence function*, HDiv, that, for each pair of non-intersecting hyperplanes, gives the length of a shortest path between these hyperplanes that avoids a ball of radius r about a basepoint. We then give conditions on a pair of non-intersecting hyperplanes that imply a lower bound on their corresponding hyperplane divergence function. The proof for these lower bounds involve the use of disk diagrams. Finally, we show how the hyperplane divergence function for a pair of such hyperplanes actually implies a lower bound on the divergence of the entire CAT(0) cube complex:

Theorem 5.2.6 *Let X be an essential, locally compact CAT(0) cube complex* *with cocompact automorphism group. Suppose $HDiv(\mathcal{Y}, \mathcal{Z}) \succeq F(r)$ for a pair of non-intersecting hyperplanes $\mathcal{Y}$ and $\mathcal{Z}$ in X. It then follows that $Div(X) \succeq rF(r)$.*

Consequently, this process reduces the problem of finding a lower bound on divergence to finding a pair of hyperplane with certain separation properties. Through this strategy, we prove the following theorem which gives lower bounds on divergence as a consequence of the existence of certain types of pairs of non-intersecting hyperplanes (these hypotheses on hyperplanes are defined in Section 5).

Theorem 5.1.2 and 6.0.2 *Suppose X is an essential, locally compact CAT(0) cube complex with cocompact automorphism group. Let $\mathcal{Y}$ and $\mathcal{Z}$ be non-intersecting hyperplanes in X.*

1. *If $\mathcal{Y}$ and $\mathcal{Z}$ are k–separated, then $Div(X)$ is bounded below by a quadratic function.*

2. *If $\mathcal{Y}$ and $\mathcal{Z}$ are k–chain separated, then $Div(X) \succeq \frac{1}{2}R^2 \log_2(\log_2(R))$.*

3. *If X contains a pair of degree d k–separated hyperplanes, then $Div(X)$ is bounded below by a polynomial of degree $d+1$.*

4. *Suppose X has k–alternating geodesics. If $\mathcal{Y}$ and $\mathcal{Z}$ are symbolically k–chain separated then $Div(X)$ is bounded below by a cubic function.*

The aforementioned classification of quadratic divergence in right-angled Coxeter groups is an application of 4 above. Furthermore, as an application of 3 we give graph-theoretic criteria which imply polynomial lower bounds on divergence of right angled Coxeter groups.

Theorem 7.4.3 *Suppose the graph Γ contains a rank n pair (s,t), then $Div(W_\Gamma)$ is bounded below by a polynomial of degree $n+1$.*

Here a rank n pair (s,t) is a pair of non-adjacent vertices $s,t \in \Gamma$ which satisfy a certain inductive graph-theoretic criteria. By the above theorem and the machinery of thickness, we provide exact bounds on the divergence of a wide range of right-angled Coxeter groups. The above theorem, in particular, applies to the examples given in [DT15].

In order to provide upper bounds on thickness, we introduce the *hypergraph index* of a right-angled Coxeter group. The hypergraph index of any right-angled Coxeter group can be directly computed from the group's defining graph and is either a non-negative integer or ∞. It is roughly a measure of the complexity of the geometry of group. The hypergraph index yields an upper bound for a right-angled Coxeter group's order of thickness, order of *algebraic* thickness, and divergence function.

Theorem 8.3.2 and 8.3.1 *Suppose the right-angled Coxeter group, W_Γ, has hypergraph index $h \neq \infty$, then W_Γ is thick of order at most h, algebraically thick of order at most $2h-1$ and the divergence of W_Γ is bounded above by a polynomial of degree $h+1$.*

For both $n=0$ and $n=1$, thickness of order n, algebraic thickness of order n, polynomial divergence of degree $n+1$ and hypergraph index n are all equivalent notions in the setting of right-angled Coxeter groups (see section 8 for an overview). Actually, the following conjecture seems to hold for all groups whose divergence and thickness we can compute:

Conjecture 8.5 *Let Γ be a simplicial graph and W_Γ the corresponding right-angled Coxeter group. The following are equivalent:*

1. *Γ has hypergraph index n.*
2. *W_Γ is thick of order n.*
3. *The divergence of W_Γ is a polynomial of degree $n+1$.*

One may then ask if algebraic thickness of order n and thickness of order n are equivalent notions in right-angled Coxeter groups, as this is true for $n=0$ and $n=1$. In fact, in the paper where thick groups are originally defined, the authors ask if the order of algebraic thickness of any finitely generated group is equivalent to the group's order of thickness [BDM09, Question 7.7]. Sisto provided a negative answer to this question by demonstrating an example of a group which is thick of order 1 but is not algebraically thick of order 1 [BD14]. We give a first negative answer to this question for the case of higher orders of thickness (see Theorem 8.4.1 for a more detailed statement):

Theorem 1.0.2. *Given any integer $n > 1$, there are right-angled Coxeter groups that are thick of order n, but are algebraically thick of order strictly larger than n.*

Finally, we explore the divergence in the setting of Coxeter groups (not necessarily right-angled). Given an edge-labeled simplicial graph Γ there is a corresponding Coxeter group W_Γ. An adaptation of our techniques allow us to prove results in this general case. For instance, we provide the following polynomial lower bound.

Theorem 9.1.4 *Let Γ be an even triangle-free Coxeter graph. Suppose (u, v) is a rank n pair, then the divergence of the Coxeter group W_Γ is bounded below by a polynomial of degree $n+1$ in r.*

By the above theorem and the results from [BHS17], for any positive integer degree, we can conclude there are infinite classes of Coxeter groups that are not right-angled and which have polynomial divergence of the given degree. This shows the existence of higher degree polynomial divergence in the general class of Coxeter groups is abundant.

Theorem 9.1.4 is actually proven in a more general setting as we only need Γ to be triangle-free and even for some neighborhood of the vertex u. For a precise statement see Section 9. For an edge-labeled graph Γ representing a Coxeter group, we let $\hat{\Gamma}$ denote the graph obtained by collapsing odd labeled edges of Γ to a point. We prove the following:

Theorem 9.2.2 *Let Γ be a Coxeter graph. If the diameter of $\hat{\Gamma}$ is larger than 2, then W_Γ has at least quadratic divergence.*

In particular, the above theorem shows that if W_Γ is an even Coxeter group where Γ has diameter larger than 2, then the divergence of W_Γ is at least quadratic.

This thesis is organized as follows. Section 2 provides general background material, including a background on general definitions, divergence functions, CAT(0) cube complexes, Coxeter groups and thick spaces. Section 3 provides necessary background on disk diagrams in CAT(0) cube complexes. In this section we set the notation and results regarding disk diagram structures that are used throughout the article.

In section 4, we introduce several notions of separation for a pair of non-intersecting hyperplanes in a CAT(0) cube complex. The consequences of these separation properties on the divergence of a CAT(0) cube complex are explored in section 5 and 6. The hyperplane divergence function is defined there as well.

In section 7 we apply the results obtained for CAT(0) cube complexes to the setting of right-angled Coxeter groups.

In section 8, we define the hypergraph index of a right-angled Coxeter group and explore its connection to thick and algebraically thick structures.

Finally, in section 9, we explore the divergence of Coxeter groups that are not necessarily right-angled.

Chapter 2

Background

2.1 General Definitions

2.1.1 Coarse Geometry

Given a metric space X, we will always use $B_p(r)$ to denote the ball of radius r about the point $p \in X$.

Definition 2.1.1 (quasi-isometry). Let X and Y be metric spaces. A *(k, c)–quasi-isometry* is a not necessarily continuous map $f : X \to Y$, such that for all $a, b \in X$ we have:

$$\frac{1}{k} d_X(a, b) - c \leq d_Y(f(a), f(b)) \leq k d_X(a, b) + c$$

Quasi-isometries provide a natural notion of equivalence in a coarse geometric setting. For a detailed background on quasi-isometries and geometric group theory in general see [BH99].

2.1.2 Graphs

Many graphs are referenced throughout this thesis, usually corresponding to different Coxeter groups. We present here some of the basic definitions regarding the graphs involved. All graphs considered are simplicial unless otherwise noted.

For Γ a graph, $V(\Gamma)$ and $E(\Gamma)$ are respectively the vertex set and edge set of Γ.

Definition 2.1.2 (Link and Star)**.** For $s \in V(\Gamma)$, the *link of s*, $Link(s) \subset V(\Gamma)$, is the set of vertices in Γ connected to s by an edge. The *star of s* is the set $Star(s) = Link(s) \cup s$.

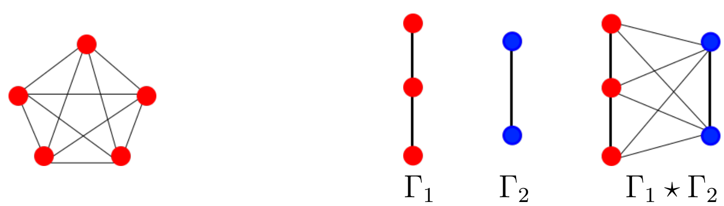

Figure 2.1: The graph on the left is a 5-clique. On the right a graph join is shown.

Definition 2.1.3 (Clique)**.** A *clique* in Γ is a subgraph whose vertices are all pairwise adjacent. A k–clique is a clique with k vertices. .

Definition 2.1.4 (Graph Join)**.** A graph Γ is a *join* if it has as subgraphs Γ_1, Γ_2 such that $V(\Gamma_1) \cup V(\Gamma_2) = V(\Gamma)$ and for every $v_1 \in V(\Gamma_1)$, $v_2 \in V(\Gamma_2)$, $(v_1, v_2) \in E(\Gamma)$. Graph joins are denoted as $\Gamma = \Gamma_1 \star \Gamma_2$.

2.2 Divergence

In this section we first present the intuitive notion of geodesic divergence. We then give the definition for the divergence of a metric space which is studied throughout this thesis.

2.2.1 Divergence of a Geodesic

The geodesic divergence function provides a measure of how quickly two ends of a geodesic stray apart from one another.

Definition 2.2.1 (Geodesic Divergence)**.** Let X be a metric space. Fix constants $0 < \delta \leq 1$, $\lambda \geq 0$ and consider the linear function $\rho(r) = \delta r - \lambda$. Let $\beta : \mathbb{R} \to X$ be a bi-infinite geodesic. The geodesic divergence of β is the function $\mathrm{GDiv}(\beta, \delta, \lambda)(r)$ whose value at r is the length of a shortest path connecting the points $\beta(r)$ and $\beta(-r)$ which does not intersect the ball $B_{\beta(0)}(\rho(r))$.

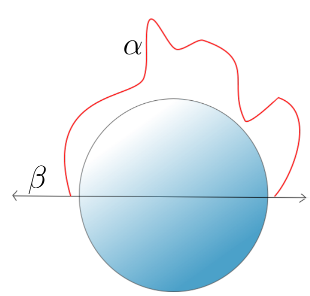

Figure 2.2: The figure above shows a shortest path α (in red) which avoids the ball $B_{\beta(0)}(\rho(r))$ and connects $\beta(r)$ to $\beta(-r)$ for some r. Thus, $\mathrm{GDiv}(\beta, \delta, \lambda)(r) = |\alpha|$.

If X is a δ–hyperbolic space, then for any bi-infinite geodesic β, $\mathrm{GDiv}(\beta, \delta, \lambda)(r)$

can always be bound below by an exponential function [BH99]. On the other hand, if we have a finitely generated group $G = G_1 \times G_2$ with G_1 and G_2 both infinite, then it is not hard to show that the divergence of any bi-infinite geodesic in the Cayley graph of G is bounded above by a linear function. In some sense the divergence of a geodesic measures how "hyperbolic" its behavior is.

2.2.2 Divergence of a Metric Space

For many non-hyperbolic metric spaces, the geodesic divergence function very much depends on the geodesic chosen. We wish now to provide a definition for the divergence of a metric space which provides a measure for the fastest rate a pair of geodesics can diverge from one another. The divergence of a metric space, defined this way, has the added benifit that it is a quasi-isometry invariant for finitely generated groups.

Definition 2.2.2 (Divergence of a Metric Space)**.** Let X be a metric space. Fix constants $0 < \delta \leq 1$, $\lambda \geq 0$ and consider the linear function $\rho(r) = \delta r - \lambda$. Let $a, b, c \in X$ and set $k = d(c, \{a, b\})$.

$$\mathrm{div}_\lambda(a, b, c, \delta)$$

is the length of the shortest path in X from a to b which avoids the ball $B_c(\rho(k))$.

$$\mathrm{Div}^X_\lambda(r, \delta)$$

is the supremum of $\mathrm{div}_\lambda(a, b, c, \delta)$ over all a, b, c with $d(a, b) \leq r$.

We set $f(r) \asymp g(r)$ if there exists a C such that:

$$\frac{1}{C}g(\frac{r}{C}) - Cr - C < f(r) < Cg(Cr) + Cr + C$$

Up to this equivalence relation on functions and under mild assumptions on the metric space, divergence is a quasi-isometry invariant. See [DMS10, Lemma 3.4] for the relevant hypotheses. For instance, the Cayley graph of a finitely generated group satisfies such hypotheses. In this thesis, we will only consider spaces which satisfy these hypotheses.

Furthermore, under the same hypotheses there is an appropriate choice of δ and λ so that for $\delta' < \delta$ and $\lambda' > \lambda$, we have $\mathrm{Div}^X_\lambda(r, \delta) \asymp \mathrm{Div}^X_{\lambda'}(r, \delta')$. For this reason, we will often suppress δ and λ from the notation and say the *divergence* of X is the function $\mathrm{Div}(X) = \mathrm{Div}^X_\lambda(r, \delta)$. For a finitely generated group G, $\mathrm{Div}(G)$ will mean the divergence of the Cayley graph of G.

2.3 CAT(0) Cube Complexes

A *CAT(0) cube complex*, X, is a simply connected cell complex whose cells consist of Euclidean unit cubes, $[-\frac{1}{2}, \frac{1}{2}]^d$, of varying dimension d. Additionally, the link of each vertex is a flag complex (i.e., any set of vertices which are pairwise connected by an edge, span a simplex). X with the induced metric is a CAT(0) space. X is *finite-dimensional* if there is an upper bound on the dimension of cubes in X. For a detailed account of CAT(0) cube complexes see [CS11] and [Wis11].

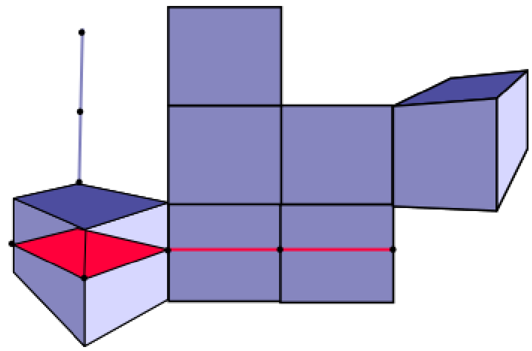

Figure 2.3: An example of a 3 dimensional cube complex. The red complex is a hyperplane.

A *midcube* $Y \subset [-\frac{1}{2}, \frac{1}{2}]^d$ is the restriction of a coordinate to 0. A *hyperplane* $\mathcal{H} \subset X$ is a connected subspace with the property that for each cube C in X, $\mathcal{H} \cap C$ is a midcube or $\mathcal{H} \cap C = \emptyset$. It follows $X - \mathcal{H}$ consists of exactly two distinct components. A *half-space* is the closure of such a component. We denote the two half-spaces associated to $\mathcal{H}$ by $\mathcal{H}^+$ and $\mathcal{H}^-$. The carrier of a hyperplane, $N(\mathcal{H})$, is the set of all cubes in X which have non-trivial intersection with $\mathcal{H}$.

A CAT(0) cube complex X is *essential* if all its half-spaces contain arbitrarily large balls of X. If X is one-ended and essential, then every hyperplane is unbounded.

2.3.1 Core CAT(0) Cube Complex Results

The following core lemmas, whose proofs are found in [CS11], are used throughout this paper.

We will work exclusively with the combinatorial metric on the 1–skeleton of X.

A *combinatorial geodesic* is a geodesic in the 1–skeleton of X under this metric and a *combinatorial path* is a path in the 1–skeleton of X. We often drop the word "combinatorial" from these definitions.

The following lemma allows us to work, up to the coarse equivalence of quasi-isometries, with the combinatorial metric.

Lemma 2.3.1. *Suppose X is a finite-dimensional CAT(0) cube complex. X is quasi-isometric to its 1–skeleton endowed with the combinatorial metric.*

This next lemma is very useful when we must find a large set of non-intersecting hyperplanes which cross some geodesic.

Lemma 2.3.2. *Suppose X is a finite-dimensional CAT(0) cube complex. For each $k > 0$, there exists a number $N(k)$ such that any combinatorial geodesics of length $N(k)$ in X must cross a set of pairwise non-intersecting hyperplanes $\{\mathcal{H}_1, ..., \mathcal{H}_k\}$.*

2.3.2 Double Skewering Lemma

We present a version here of the "Double Skewering Lemma" proven in [CS11]. This lemma is used later to prove Theorem 5.2.6. Given a pair of non-intersecting hyperplanes, the double skewering lemma will later allow us to construct an infinite chain of hyperplanes by translating the original pair by the automorphism group's action.

Lemma 2.3.3 (Double Skewering Lemma [CS11])**.** *Let X be an essential finite-dimensional, locally compact CAT(0) cube complex with cocompact automorphism* *group* $Aut(X)$*. Let* $\mathcal{Y}^+ \subset \mathcal{Z}^+$ *be two half-spaces in* X*. There exists a* $\gamma \in Aut(X)$ *such that* $\gamma\mathcal{Z}^+ \subset \mathcal{Y}^+ \subset \mathcal{Z}^+$*.*

Remark 2.3.3.1. The above Lemma is actually stated in [CS11] for the more general setting that X is a finite-dimensional CAT(0) cube complex with $\mathrm{Aut}(X)$ acting essentially without fixed point at infinity.

2.4 Coxeter Groups

We only give a brief background on Coxeter groups. For an extensive background we refer the reader to [BB05] and [Dav08].

A Coxeter group is defined by the presentation:

$$W = \langle s_1, s_2, ..., s_n | (s_i s_j)^{m(s_i, s_j)} = 1 \rangle$$

where $m(s_i, s_i) = 1$ and $m(s_i, s_j) = m(s_j, s_i) \in \{2, 3, ..., \infty\}$ when $i \neq j$. If $m(s_i, s_j) = \infty$ then no relation of the form $(s_i s_j)^m = 1$ is imposed.

Given a presentation for a Coxeter group, there is a corresponding labeled *Coxeter graph* Γ. The vertices of Γ are elements of S. There is an edge between s_i and s_j if and only if $m(s_i, s_j) \neq \infty$. This edge is labeled by $m(s_i, s_j)$ if $m(s_i, s_j) \geq 3$. If $m(s_i, s_j) = 2$ no label is placed on the corresponding edge. Conversely, given such an edge-labeled graph Γ, we have the Coxeter group W_Γ.

In the literature, there are many different conventions for associating a graph to a Coxeter group presentation. The given convention was chosen to make the theorems in this paper easier to state.

A *right-angled Coxeter group (RACG)* is a Coxeter group with generating set S where $m(s,t) \in \{\infty, 2\}$ for s,t distinct elements in S. An *even Coxeter group* is a Coxeter group given by a Coxeter graph where each edge either has an even label or no label.

We will often want to consider subgroups of a Coxeter group W_Γ corresponding to subgraphs of Γ. The full subgraph of $T \subset V(\Gamma)$ is the graph with vertex set T with a labeled edge (t_1, t_2) if and only if (t_1, t_2) is an edge of Γ with the same label.

2.4.1 Core Coxeter Groups Results

We provide some essential lemmas in the theory of Coxeter groups. We do not prove these here, and instead refer the reader to [BB05] and [Dav08].

Definition 2.4.1 (Induced Subgroup)**.** Let W_Γ be a Coxeter group with generating set $V(\Gamma) = S$. For $T \subset S$, let W_T be the subgroup of W generated by the induced subgraph of T.

The notation in this definition is justified by the following result.

Lemma 2.4.2. *Let W_Γ be a Coxeter group with generating set $V(\Gamma) = S$. Given $T \subset S$, let $\Delta \subset \Gamma$ be the induced subgraph of T. The subgroup W_T is indeed isomorphic to the Coxeter group W_Δ. Furthermore, W_T is convex in respect to the word metric of W_Γ.*

Coxeter groups provide very nice combinatorial objects. One example of this is that there is a straightforward combinatorial algorithm to solve the word problem in these groups. We describe this next.

A version of the following definition and theorem is found in Davis's book [Dav08].

Definition 2.4.3. Define the following two operations on a possibly not reduced word $w = s_1 s_2 ... s_n \in W$.

Elementary reduction: Delete a subword of the form ss where $s \in S$.

Swap: Replace the alternating word $stst...$ of length $m(s,t)$ with the alternating word $tsts...$ of length $m(s,t)$.

A word is *M-reduced* if it cannot be shorted further by a sequence of the above two operations.

The following lemma is due to Tits and can be found in [Dav08, Theorem 3.4.2].

Lemma 2.4.4 (Tit's Solution to the Word Problem)**.** *Let (W,S) be a Coxeter group.*

i) A word w is a reduced expression if and only if it is M-reduced

ii) Two reduced expressions u and v represent the same element of W if and only if one can be transformed into the other by a sequence of swaps.

The following lemma is very well known. It shows how one can recognize whether a right-angled Coxeter group is finite or is a product from properties of the group's defining graph. Nevertheless, we provide a proof here for completeness.

Lemma 2.4.5. *Let W_Γ be a right-angled Coxeter group.*

1. *W_Γ is finite if and only if Γ is a clique.*

2. *Given induced subgraphs Γ_1, Γ_2 of Γ, $W_\Gamma = W_{\Gamma_1} \times W_{\Gamma_2}$ if and only if $\Gamma = \Gamma_1 \star \Gamma_2$.*

Proof.

1. Suppose Γ is not a clique. Hence, there are two vertices $u, v \in \Gamma$ such that (u, v) is not an edge of $E(\Gamma)$. Let T be the induced subgraph of Γ which only contains the vertices u and v. It follows W_T is the infinite dihedral group, and therefore W_Γ is infinite. Hence, for W_Γ to be finite, Γ must be a clique.

   In the other direction, if Γ is a clique, then W_Γ is the product of $|\Gamma|$ copies of $\mathbb{Z}_2$, and so is finite.

2. Suppose first that $W_\Gamma = W_{\Gamma_1} \times W_{\Gamma_2}$. It follows that every generator in Γ_1 must commute with every generator of Γ_2. Since if two generators commute there is an edge between the corresponding vertices of Γ, we have that $\Gamma = \Gamma_1 \star \Gamma_2$.

   For the other direction, suppose $\Gamma = \Gamma_1 \star \Gamma_2$. It follows every word in W_{Γ_1} commutes with every word in W_{Γ_2}. Furthermore, the generators corresponding to vertices of Γ_1 and Γ_2 generate Γ. Hence, $W_\Gamma = W_{\Gamma_1} \times W_{\Gamma_2}$.

$\square$

2.4.2 The Davis Complex of a RACG

The *Davis complex*, Σ_Γ, is a natural CAT(0) cell complex which the Coxeter group W_Γ acts geometrically. In this thesis, we will only make use of the Davis complex for right-angled Coxeter groups.

Suppose W_Γ is a RACG. For every k–clique, $T \subset \Gamma$, the induced subgroup W_T is isomorphic to the direct product of k copies of $\mathbb{Z}_2$. It follows that the Cayley graph of W_T is isometric to a unit k–cube. The Davis complex Σ_Γ is constructed in the following way. The 1-skeleton of Σ_Γ is the Cayley graph of W_Γ where edges are given unit length. Additionally, for each k–clique, $T \subset \Gamma$, and coset, gW_T, we glue a unit k–cube to $gW_T \subset \Sigma_\Gamma$. The Davis Complex for a RACG is naturally a CAT(0) cube complex.

Throughout this thesis, much like the Cayley graph of W_Γ, we will assume that 1-cells of Σ_Γ are labeled by letters of Γ corresponding to the associated generator. Furthermore, vertices of Σ_Γ are labeled by group elements of W_Γ.

Now suppose $\mathcal{H}$ is a hyperplane in Σ_Γ, the Davis complex of a RACG. It is readily checked that 1-cells dual to $\mathcal{H}$ are labeled by the same letter $t \in \Gamma$. Accordingly, we say $\mathcal{H}$ is of type t. Furthermore, $N(\mathcal{H})$ is isometric to $\Sigma_t \times \Sigma_{Link(t)}$, where Σ_t is a 1-cell labelled by the generator t and $\Sigma_{Link(t)}$ is the Davis complex corresponding to $W_{Link(t)}$, the subgroup associated to the induced subgraph $Link(t)$.

Let $\mathcal{H}$ and $\mathcal{H}'$ be two crossing hyperplanes of corresponding types s and s'. An important consequence is that $m(s, s') = 2$, i.e. s and s' are connected by an edge in Γ. This fact will be used freely throughout this work.

2.4.3 When is the Davis Complex of a RACG Essential?

We provide a simple graph-theoretic criteria for when the Davis complex of a RACG is essential.

Lemma 2.4.6. *Let W_Γ be a RACG. The Davis complex Σ_Γ is essential if and only if $\Gamma \neq \Gamma' \star K$ where K is a clique and Γ' is any induced subgraph of Γ.*

Proof. First suppose Σ_Γ is essential, and for a contradiction assume $\Gamma = \Gamma' \star K$, with K a clique. Consequently, $W_\Gamma = W_{\Gamma'} \times W_K$. It follows Σ_Γ decomposes as the product $\Sigma_\Gamma = \Sigma_{\Gamma'} \times \Sigma_K$. Let $\mathcal{H}$ be a hyperplane dual to an edge of Σ_Γ corresponding to a generator $k \in K$. Note that one of the carriers of $\mathcal{H}$ is exactly the subcomplex corresponding to a coset $gW_{\Gamma'}$ with $g \in id \times W_K$. Let v be a vertex in W_Γ. It follows $v = g_1 \times g_2$ with $g_1 \in W_{\Gamma'}, g_2 \in W_K$, and so is at a distance at most $|W_K|$ from $\mathcal{H}$. This contradicts Σ_Γ being essential.

Now assume that $\Gamma \neq \Gamma' \star K$. Let $\mathcal{H}$ be a hyperplane through an edge of Σ_Γ corresponding to the vertex $v \in \Gamma$. By the group action, we can assume the chosen edge is adjacent to the identity vertex. It follows that there must be some generator u which does not commute with v, since otherwise $\Gamma = v \star (\Gamma - v)$. Now consider the bi-infinite geodesic $...uvuvu...$ in Σ_Γ. Since u and v do not commute, it follows that hyperplanes crossing edges labelled by u cannot cross hyperplanes that cross edges labelled by v. Hence, we have a bi-infinite chain of non-intersecting hyperplanes which contains $\mathcal{H}$. This shows that there are points arbitrarily far from $\mathcal{H}$ in both components of $\Sigma_\Gamma - \mathcal{H}$. Therefore, Σ_Γ is essential. $\square$

2.5 Thick Spaces

Thick metric spaces were first defined in [BDM09], and these authors prove the degree of thickness of a metric space is an important quasi-isometry invariant. Furthermore, they also show that thickness is an obstruction for a space to be relatively hyperbolic.

We work with the "strong" thickness definitions from [BD14]. As we will never make reference to the weaker notions of thickness, we will drop the word "strongly" from our definitions.

The authors of [BD14] showed there is a strong connection between the degree of thickness of a metric space and its divergence function This connection is especially relevant for this thesis.

The definitions and cited results of this section are only utilized later in Section 8. The reader may wish to skip this section until then.

2.5.1 Definitions

X will denote a metric space and $Y \subset X$ a subspace. Y is *C–path connected* if for any $y_1, y_2 \in Y$ there exists a path from y_1 to y_2 in $N_C(Y)$. Y is *(C, L)–quasi-convex* if for any $y_1, y_2 \in Y$, there exists an (L, L)–quasi-geodesic in $N_C(Y)$ connecting y_1

and y_2.

Roughly, X forms a tight network of spaces in respect to the subsets $\{Y_\alpha\}_{\alpha \in A}$ if these subsets coarsely cover X. Furthermore, any two subsets can be connected by a sequence of subsets such that consecutive subsets in this sequence coarsely intersect in an infinite diameter set. This is formally defined below.

Definition 2.5.1 (Tight network of subspaces)**.** [BD14, Definition 4.1]

Given $C > 0$ and $L > 0$, X is a (C, L)–*tight network with respect to a collection* $\{Y_\alpha\}_{\alpha \in \mathcal{A}}$ *of subsets* if the following hold:

a) Every $Y \in \{Y_\alpha\}_{\alpha \in \mathcal{A}}$ with the induced metric is (C, L)–quasi-convex

b) $X = \cup_{\alpha \in \mathcal{A}} N_C(Y_\alpha)$

c) For every $Y, Y' \in \{Y_\alpha\}$ and any $x \in X$ such that $N_{3C}(x)$ intersects both Y and Y', there exists a sequence of length $n \leq L$

$$Y = Y_1,\ Y_2, ...,\ Y_{n-1},\ Y_n = Y'$$

with $Y_i \in \{Y_\alpha\}$ such that for all $1 \leq i < n$, $N_C(Y_i) \cap N_C(Y_{i+1})$ is of infinite diameter, L–path connected and intersects $N_L(x)$.

A metric space is *wide* if every one of its asymptotic cones has cutpoints, and, additionally, every point in the space is uniformly near to a (L, L)–quasi-geodesic. The following definition provides a uniform version of this notion.

Definition 2.5.2 (Uniformly wide)**.** [BD14, Definition 4.11] A collection of metric spaces, $\{Y_\alpha\}_{\alpha \in A}$, is (C, L)–*uniformly wide* if:

1. There exists $C, L \geq 0$ such that for every $Y \in \{Y_\alpha\}_{\alpha \in A}$ and for every $y \in Y$, y is in the C neighborhood of some bi-infinite (L, L)–quasi-geodesic in Y.

2. Given any sequence of metric spaces (Y_i, d_i) in $\{Y_\alpha\}$, any ultrafilter ω, any sequence of scaling constants (s_i) and any sequence of basepoints (b_i) with $b_i \in Y_i$, it follows that the ultralimit $\lim_\omega (Y_i, b_i, \frac{1}{s_i} d_i)$ does not have cut-points.

Metric thickness of a space X, defined below, provides an inductive decomposition of X into tight network of spaces. The base case consists of a set of uniformly wide spaces.

Definition 2.5.3 (Metric thickness)**.** [BD14, Definition 4.13] A family of metric spaces is (C, L)*–thick of order zero* if it is (C, L)–uniformly wide.

Given $C \geq 0$ and $k \in \mathbb{N}$ we say that a metric space X is (C, L)*–thick of order at most k with respect to a collection of subsets* $\{Y_\alpha\}$ if

1. X is a (C, L)–tight network with respect to $\{Y_\alpha\}$.

2. The subsets in $\{Y_\alpha\}$ endowed with the restriction of the metric on X compose a family of spaces that are (C, L)–thick of order at most $k - 1$.

Furthermore, X is said to be *thick of order k* (with respect to $\{Y_\alpha\}$) if it is (C, L)–thick of order at most k (with respect to $\{Y_\alpha\}$) and for no choices of C, L and $\{Y_\alpha\}$ is X (C, L)–thick of order at most $k - 1$.

The following definitions give an algebraic version for thickness. The algebraic condition often implies stronger results (see [BD14]).

Definition 2.5.4 (Tight algebraic network of subgroups)**.** [BD14, Definition 4.1] Let $C > 0$, G a finitely generated group and $\mathcal{H}$ a set of subgroups of G. G is a *C–tight algebraic network with respect to $\mathcal{H}$* if the following hold:

a) Every $H \in \mathcal{H}$ is M–quasi-convex

b) The union of all subgroups in $\mathcal{H}$ generates a finite index subgroup of G.

c) For every $H, H' \in \{\mathcal{H}\}$, there exists a sequence

$$H_1 = H, H_2, ..., H_{n-1}, H_n = H'$$

with $H_i \in \{\mathcal{H}\}$ such that for all $1 \leq i < n$, $H_i \cap H_{i+1}$ is infinite and is M–path connected.

By [BD14, Proposition 4.3], if G admits a tight algebraic network of subgroups in respect to $\mathcal{H}$ then G is a tight network of subspaces with respect to the left cosets of groups in $\mathcal{H}$.

Definition 2.5.5 (Algebraic thickness)**.** [BD14, Definition 4.13] Let G be a finitely generated group. G is *algebraically thick of order zero* if it is wide. Given $C \geq 0$, G is *C–algebraically thick of order at most k with respect to a finite collection of subgroups $\mathcal{H}$* if

1. G is a C–tight algebraic network with respect to $\mathcal{H}$.

2. Every $H \in \mathcal{H}$ is algebraically thick of order at most $k - 1$.

G is algebraically thick of order k if it is algebraically thick of order at most k and is not algebraically thick of order $k-1$.

2.5.2 Relevant Results

We cite here some relevant results regarding thickness of spaces. An important result is that the degree of thickness of a space provides a polynomial upper bound on the space's divergence:

Theorem 2.5.6 ([BD14], Corollary 4.17). *Suppose X is thick of order at most d, then for any $0 < \delta < \frac{1}{54}$ and $\lambda \geq 0$, $Div_{\lambda}^{X}(r, \delta) \prec x^{d+1}$.*

Coxeter groups provide a nice example of thick spaces. In fact, the following is true:

Theorem 2.5.7 ([BHS17], Theorem A.1). *Let Γ be a Coxeter group. Then either W_Γ is relatively hyperbolic or W_Γ is algebraically thick (and therefore also metrically thick) of some degree.*

The above theorem tells us that one-ended Coxeter groups either have exponential divergence or have their divergence bounded above by a polynomial of some degree. Furthermore, these authors develop a graph-theoretic algorithm which provides an explicit algebraically thick structure for a given non-relatively hyperbolic Coxeter group. This algorithm is especially easily applied for the RACG case.

Later in Section 8 we provide a different algorithm which gives an explicit thick structure for a non-relatively hyperbolic RACG. Our algorithm is similar in nature to that of [BHS17]; however, the thick structures we provide are not always algebraically thick. Nevertheless, by not requiring algebraic thickness we can provide better bounds on the divergence of these spaces using Theorem 2.5.6.

Chapter 3

Disk Diagrams in CAT(0) Cube Complexes

In this chapter we review disk diagrams in CAT(0) cube complexes. Many of the ideas presented in this chapter originated in [Hag13] and [Wis11]. For our purposes, we require some modifications of definitions and lemmas from the mentioned works. We try to outline throughout the differences and similarities in the given definitions. For completeness we include proofs of these modified claims, even though many of the arguments are very similar.

3.1 Background

A *disk diagram* D is a contractible finite 2–dimensional cube complex with a fixed planar embedding $P\colon D \to \mathbb{R}^2$. The *area* of D is the number of 2–cells it contains.

By compactifying $\mathbb{R}^2$, $S^2 = \mathbb{R}^2 \cup \infty$, we can extend $P\colon D \to S^2$, giving a cellulation of S^2. The *boundary path of* D, ∂D, is the attaching map of the cell in this cellulation containing ∞. Note that this is not necessarily the topological boundary.

Let X be a CAT(0) cube complex. We say D is a *disk diagram in* X, if D is a disk diagram and there is a fixed continuous combinatorial map of cube complexes $F\colon D \to X$. By a lemma of Van Kampen, for every null-homotopic closed combinatorial path $p\colon S^1 \to X$, there exists a disk diagram D in X such that $\partial D = p$.

Suppose D is a disk diagram in a CAT(0) cube complex X and t is a 1–cell of D. A *dual curve* H *dual to* t, is a concatenation of midcubes in D which contains a midcube in D which intersects t. The image of H under the map $F\colon D \to X$ lies in some hyperplane $\mathcal{H} \subset X$. We also have that every edge in D is dual to exactly one maximal dual curve.

In our notation, we denote a dual curve by a capital letter and its corresponding hyperplane by the corresponding script letter.

An *end* of a dual curve H in D is a point of intersection of H with ∂D. Maximal dual curves either have no ends or two ends. The *carrier* $N(H)$ of H is the set of 2-cubes in D containing H.

3.2 Hyperplane-Path Sequences and Pathologies

Suppose we have oriented combinatorial paths p_1, p_2, ..., p_n in a disk diagram D and that $p_i \cap p_{i+1} \neq \emptyset$ for $1 \leq i < n$. We then define a new oriented path $p = p_1 * p_2 * ... * p_n$ by beginning at the first point of p_1, followed by p_1 until its first intersection with p_2; followed by p_i until its first intersection with p_{i+1}. In this definition, we further assume the orientations are chosen such that this construction is possible (i.e., we can always follow p_i along its orientation until its intersection with p_{i+1}). Furthermore, different choices of path orientations could produce different paths. We note that this construction is only used in Lemma 5.2.3 and Lemma 5.2.4, and the relevant path orientations there are given.

To every closed loop formed by a concatenation of combinatorial paths and hyperplanes, we wish to associate a disk diagram with boundary path this loop. This notion is formally defined below. If such a diagram is chosen appropriately, the dual curves associated to it behave nicely. We call such nicely behaved diagrams *combed* and define them later in this section.

Definition 3.2.1 (Hyperplane-Path Sequence)**.** In the following definition, we work modulo $n + 1$ (i.e, $n + 1 = 0$). Let $\bar{A} = \{A_0, A_1, ..., A_n\}$ be a sequence such that for each i, A_i is either a hyperplane or an oriented combinatorial path in a CAT(0) cube complex X. Furthermore, for $0 \leq i \leq n$, if A_i and A_{i+1} are both hyperplanes then they intersect. If A_i and A_{i+1} are both combinatorial paths, then the endpoint of A_i is the initial point of A_{i+1}. If A_i is a hyperplane and A_{i+1} is a path, then the beginning point of A_{i+1} lies on $N(A_i)$. Similarly, if A_i is a path and A_{i+1} is a hyperplane, then the endpoint of A_i lies on $N(A_{i+1})$. We call $\bar{A}$ a *hyperplane-path sequence.*

Given a hyperplane-path sequence $\bar{A} = \{A_0, ..., A_n\}$, let $\bar{P} = \{P_0, ..., P_n\}$ be a sequence of combinatorial paths where $P_i = A_i$ if A_i is a combinatorial path and P_i is a combinatorial geodesic in $N(A_i)$ if A_i is a hyperplane. Furthermore, assume $P = P_0 * P_1 * ... * P_n$ defines a loop. A disk diagram D is *supported by* $\bar{A}$ if $\partial D = P$ for some choice of $\bar{P}$. Often this choice is given and we say D *is supported by* $\bar{A}$ *with boundary path* $\bar{P}$. For an example of a disk diagram supported by a hyperpane path sequence see Figure 5.1.

Remark 3.2.1.1. Diagrams supported by a hyperplane-path sequence are a special case of diagrams with fixed carriers defined in [Hag13]. The difference is that consecutive hyperplanes in a hyperplane-path sequence must intersect, where in [Hag13] they either intersect or osculate. Most of this section can be modified to allow for osculating hyperplanes; however, there was no need for such sequences in this paper.

Definition 3.2.2 (Nongons, Bigons, Monogons and Oscugons)**.** A *nongon* is a dual curve of length greater than one which begins and ends on the same dual 1–cell. A *bigon* is a pair of dual curves which intersect at their first and last containing squares. A *monogon* is a dual curve which intersects itself in its first and last square. An *oscugon* is a dual curve which starts at the dual 1–cell e_1, ends at the dual 1–cell

e_2, such that $e_1 \neq e_2$, $e_1 \cap e_2 \neq \emptyset$ and e_1, e_2 are not contained in a common square. A disk diagram *without pathologies* is one that does not contain a nongon, bigon, monogon or oscugon.

The following is proved in [Wis11, Corollary 2.4]:

Lemma 3.2.3 ([Wis11])**.** *Suppose D is a disk diagram in a CAT(0) cube complex X, then D does not contain monogons.*

We now wish to discuss the idea of *boundary combinatorics* in a disk diagram D. Suppose D and D' are two disk diagrams in a CAT(0) cube complex X. Let $p \subset \partial D$ and $p' \subset \partial D'$ be subcomplexes. We say p and p' are *equal boundary complexes* if the canonical maps $p \to X$ and $p' \to X$ are the same combinatorial maps.

Suppose $p \subset \partial D$ and $p' \subset \partial D'$ are equal boundary complexes, and let $i\colon\ p \to p'$ be the canonical isomorphism between them. We say p and p' have *equal boundary combinatorics* if given any pair of edges $e_1, e_2 \in p$ dual to a common dual curve in D, it follows $i(e_1)$ and $i(e_2)$ are dual to a common dual curve in D'.

In particular, D and D' have *equal boundary* if ∂D and $\partial D'$ are equal combinatorial complexes. Additionally, D and D' have *equal boundary combinatorics* if ∂D and $\partial D'$ have equal boundary combinatorics.

The following is also proved in [Wis11, Lemma 2.3]:

Lemma 3.2.4 ([Wis11])**.** *Let D be a disk diagram in a CAT(0) cube complex X. There exists a disk diagram D' in X with no pathologies and equal boundary combinatorics to D.*

3.3 Combed Diagrams

Definition 3.3.1 (Combed Diagram)**.** Let D be a disk diagram supported by hyperplane-path sequence $\bar{A} = \{A_0, ..., A_n\}$ with boundary path $\bar{P} = \{P_0, ..., P_n\}$. We say D is *combed* if the following properties hold:

1. D has no pathologies.

2. If A_i is a hyperplane, no two dual curves dual to P_i intersect. In particular, no dual curve has both ends on P_i.

3. If A_i and A_{i+1} are both hyperplanes, no dual curve dual to P_i intersects P_{i+1}.

The arguments in the next two lemmas are essentially the same as those in the proof of [Hag13, Lemma 2.11]. However, for our purposes, we often require a statement regarding the boundary combinatorics of a given disk diagram. To be self-contained and to guarantee the arguments are valid in our context, we provide them here.

Lemma 3.3.2. *Suppose D is a disk diagram supported by hyperplane-path sequence $\bar{A} = \{A_0, ..., A_n\}$ with boundary path $\bar{P} = \{P_0, ..., P_n\}$, satisfying properties 1 and 2 of combed diagrams (Definition 3.3.1). There exists a combed diagram D' with boundary path $\bar{P}' = \{P'_0, ..., P'_n\}$, where P'_i is a connected subsegment of P_i. Furthermore, $\bar{P}'$ has equal boundary combinatorics as its image in $\bar{P}$.*

Proof. Suppose some dual curve C intersects both P_i and P_{i+1}. Let v be the vertex where P_i meets P_{i+1}. Let e_1 be the edge in P_i which intersects v and e_2 the edge in P_{i+1} which intersects v.

By property 2 of combed diagrams, every dual curve to P_i between v and C intersects P_{i+1} as well. It follows that the dual curve, K, to e_1 intersects e_2. Let Q be the combinatorial path in $N(K) - K$ that forms a loop based at v. Note that Q cannot contain any edge of ∂D. If $e_1 \neq e_2$, then any dual curve to Q must intersect Q twice, forming a bigon. However, as D does not contain bigons, it follows $e_1 = e_2$. Hence, we obtain a new diagram from D by simply deleting the edge e_1. By repeating this process we obtain the desired diagram D'. $\square$

Lemma 3.3.3. *Let D be a disk diagram in a CAT(0) cube complex X supported by hyperplane-path sequence $\bar{A} = \{A_0, ..., A_n\}$ with boundary path $\bar{P} = \{P_0, ..., P_n\}$. Suppose for some i, $\mathcal{A} = A_i$ is a hyperplane. There exists a disk diagram D' also supported by $\bar{A} = \{A_0, ..., A_n\}$ with boundary path $\bar{P}' = \{P_0, ..., P_{i-1}, P_i', P_{i+1}, ..., P_n\}$ such that $\partial D - P_i$ has the same boundary combinatorics as $\partial D' - P_i'$. Additionally, no two dual curves dual to P_i' in D' intersect.*

Proof. By Lemma 3.2.4 we may assume D has no pathologies. Set $P = P_i$ and suppose two dual curves to P, C_1 and C_2, intersect.

A dual curve cannot have two ends on P. For otherwise P would cross the same hyperplane twice, contradicting P being geodesic. In particular, $C_1 \neq C_2$. Let e_1 be the edge on P dual to C_1 and e_2 the edge on P dual to C_2. If $d_P(e_1, e_2) = d > 0$, it follows there is another dual curve to P, C_3, between e_1 and e_2. C_3 must then either intersect C_1 or intersect C_2 (C_3 cannot have both ends on P). Proceeding this way, we can then assume that e_1 and e_2 are distinct adjacent edges.

Let S be a square where C_1 and C_2 intersect. There are two cases. First suppose that S does not contain e_1 and e_2 as edges. As C_1 intersects C_2, it follows, e_1 and e_2 must both lie in another square of X, say S'. For if this were not the case, the hyperplanes associated to the dual curves C_1 and C_2 would both cross and osculate (i.e., have dual adjacent edges that are not in a common square). However, this is not possible in a CAT(0) cube complex (see [Wis11, Section 6b]).

As the link of vertices in X are flag complexes, it follows $S' \subset N(\mathcal{A})$. We can then form a new disk diagram, D', by attaching S' to D along the edges e_1 and e_2. This modifies the path P into a new path P', that is still geodesic and is still in $N(\mathcal{A})$. In D', the dual curves C_1 and C_2 now form a bigon. By [Wis11, Lemma 2.3], there is a another disk diagram D'', with the same boundary combinatorics as D' and no pathologies, such that $\text{Area}(D'') \leq \text{Area}(D') - 2 = \text{Area}(D) - 1$. We have thus produced a diagram, D'', with the desired boundary combinatorics that is of area strictly smaller than D.

For the second case, suppose e_1 and e_2 are edges of S. Label the other edges of S as e_3 and e_4. Let $\mathcal{H}_1$ and $\mathcal{H}_2$ respectively be hyperplanes which pass through e_1 and e_2. If $\mathcal{A} = \mathcal{H}_1$, then it follows that $e_1, e_2, e_3 \subset N(\mathcal{A})$. Hence, $S \subset N(\mathcal{A})$. Alternatively, suppose $\mathcal{H}_1$, $\mathcal{H}_2$, and $\mathcal{A}$ are distinct. Since the link of vertices in X are flag complexes, it follows S is in a 3-cube which is contained in $N(\mathcal{A})$. Either way, $S \subset N(\mathcal{A})$.

Let P' be the path in D which is the same as P with e_1, e_2 replaced with e_3, e_4. P' is still geodesic. Furthermore, $P' \subset N(\mathcal{A})$. Let $D' \subset D$ be the disk diagram obtained as a subdiagram of D by replacing P with P'. It follows D' is a diagram of smaller area than D. Furthermore, the boundary combinatorics of $\partial D - P$ are not affected.

In both cases we are able to produce a smaller area disk diagram with the desired boundary combinatorics. Therefore, by iterating this process we are guaranteed to eventually have a diagram with the conclusion of the lemma. □

The following lemma guarantees the existence of a combed diagram.

Lemma 3.3.4. *Let D be a disk diagram supported by hyperplane-path sequence $\bar{A}$. There exists a combed disk digram D' also supported by $\bar{A}$.*

Proof. The lemma follows by applying Lemma 3.2.4 to D to get a disk diagram with no pathologies. We then repeatedly apply Lemma 3.3.3 to the resulting diagram to obtain a diagram satisfying properties 1 and 2 of combed diagrams. Finally, we apply Lemma 3.3.2 to obtain a combed diagram. □

Given a maximal dual curve C in a disk diagram D, we want to construct a new combed diagram with the hyperplane $\mathcal{C}$ in its support. Furthermore, we want the appropriate boundary combinatorics of D preserved in this new diagram. The following technical lemma guarantees the existence of such a diagram.

Lemma 3.3.5. *Let D be a combed disk diagram in a CAT(0) cube complex X supported by hyperplane-path sequence $\bar{A} = \{A_0, ..., A_n\}$ and with boundary path $\bar{P} = \{P_0, ..., P_n\}$.*

Suppose C is a maximal dual curve from P_i to P_j with $i < j$ and let $\mathcal{C} \subset X$ be the hyperplane associated to C. Let P be a combinatorial path in $N(C)$ from P_i to P_j. Let P_i' be the subsegment of P_i between P_{i-1} and P, and let P_j' be the subsegment of P_j between P and P_{j+1}. Let A_i' and A_j' be the corresponding supports of P_i' and P_j' respectively ($A_i' = \mathcal{A}_i$ if A_i is a hyperplane and $A_i = P_i'$ otherwise).

Let P' be any combinatorial geodesic in $N(\mathcal{C})$ connecting the endpoints of P. There exists a combed disk diagram D' supported by

$$\bar{A}' = \{A_0, ..., A_{i-1}, A_i', \mathcal{C}, A_j', A_{j+1}, ..., A_n\}$$

with boundary path

$$\bar{P}' = \{P_0, ..., P_{i-1}, P_i', P', P_j', P_{j+1} ..., P_n\}$$

such that $\partial D' - P'$ has the same boundary combinatorics as the corresponding subset of ∂D.

Proof. Let D_1 be a combed diagram with boundary path $\{P', P\}$. Let D_2 be the subdiagram of D with boundary path $\{P_0, ..., P_i', P, P_j', ..., P_n\}$. Note that D_2 is still combed. We may form a new disk diagram D_3 by gluing D_1 to D_2 along P. D_3 is supported by $\{A_0, ..., A_i', \mathcal{C}, A_j', ..., A_n\}$ with boundary path $\{P_0, ..., P_i', P', P_j', ...P_n\}$.

By Lemma 3.2.4, we may assume D_3 has no pathologies and $\partial D_3 - P'$ has the same combinatorics as the corresponding subset of ∂D.

All that is left to prove is that properties 2 and 3 of combed diagrams (Definition 3.3.1) hold. Property 3 clearly still holds for dual curves which do not intersect P'. Assume A_i is a hyperplane. Since D is combed, no dual curve to P'_i in D_2 intersects P. Since the combinatorics of D_2 are preserved in D_3, this is still the case in D_3. In particular, no dual curve to P'_i intersects P' in D_3. Therefore, property 3 holds in D_3.

Assume property 2 is false in D_3. We then have two intersecting dual curves, C_1 and C_2, that are dual to the same boundary path in D_3 with hyperplane support. By the preservation of boundary combinatorics and the fact that D_3 is combed, it follows each of these curves must have an endpoint on P'. We can then modify D_3 using Lemma 3.3.3 to obtain a new diagram D' satisfying the claim. $\square$

Chapter 4

Hyperplane Separation Properties

For the remainder of this article, X will denote a CAT(0) cube complex. Furthermore, $\mathcal{Y}$ and $\mathcal{Z}$ will always denote a pair of non-intersecting unbounded hyperplanes in X.

We will discuss different definitions for separation properties of a given pair of non-intersecting hyperplanes. In the next section we will explore the relationship between these separation properties and the divergence of X.

4.1 Strongly Separated and k–Separated Hyperplanes

The following definition describes a first notion of separation of hyperplanes.

Definition 4.1.1 ([BC12]). $\mathcal{Y}$ and $\mathcal{Z}$ are *strongly separated* if no hyperplane intersects them both.

A *minimal geodesic* g between hyperplanes $\mathcal{Y}$ and $\mathcal{Z}$ is a combinatorial geodesic with endpoints on $N(\mathcal{Y})$ and $N(\mathcal{Z})$ such that $|g|$ is minimal over all such geodesics. The following lemma shows minimal geodesics between strongly separated hyperplanes have the same endpoints.

Lemma 4.1.2. *Suppose $\mathcal{Y}$ and $\mathcal{Z}$ are strongly separated hyperplanes. Let g_1 and g_2 be minimal geodesics between $\mathcal{Y}$ and $\mathcal{Z}$. It follows that $N(\mathcal{Y}) \cap g_1 = N(\mathcal{Y}) \cap g_2$. Consequently, g_1 and g_2 have the same endpoints.*

Proof. Let $N(\mathcal{Y}) \cap g_1 = v_1$ and $N(\mathcal{Y}) \cap g_2 = v_2$. Suppose, for a contradiction, that $v_1 \neq v_2$. Let h be a combinatorial geodesic from v_1 to v_2 in $N(\mathcal{Y})$ and let $\mathcal{H}$ be a hyperplane intersecting h. $\mathcal{H}$ cannot intersect $\mathcal{Z}$ since $\mathcal{Y}$ and $\mathcal{Z}$ are strongly separated. It follows $\mathcal{H}$ must either intersect g_1 or g_2. This is a contradiction (see [Wis11, Remark 3.12]). $\square$

The following definition gives a slight generalization of strongly separated hyperplanes.

Definition 4.1.3. $\mathcal{Y}$ and $\mathcal{Z}$ are k–separated if at most k hyperplanes intersect both $\mathcal{Y}$ and $\mathcal{Z}$. In particular, a pair of strongly separated hyperplanes are 0–separated.

The following two lemmas describe how minimal geodesics and hyperplanes intersecting a pair of k–separated hyperplanes behave nicely.

Lemma 4.1.4. *Suppose $\mathcal{Y}$ and $\mathcal{Z}$ are k–separated, then $d(\mathcal{H}_1 \cap \mathcal{Y}, \mathcal{H}_2 \cap \mathcal{Y}) \leq k$ for every pair of hyperplanes $\mathcal{H}_1$, $\mathcal{H}_2$ which intersect both $\mathcal{Y}$ and $\mathcal{Z}$. Furthermore, either $\mathcal{Y}$ and $\mathcal{Z}$ are strongly separated, or every minimal geodesic g connecting $\mathcal{Y}$ to $\mathcal{Z}$ lies in the carrier of a hyperplane which intersects both $\mathcal{Y}$ and $\mathcal{Z}$.*

Proof. Suppose hyperplanes $\mathcal{H}_1$ and $\mathcal{H}_2$ each intersect both $\mathcal{Y}$ and $\mathcal{Z}$. For a contradiction, suppose that $d(\mathcal{H}_1 \cap \mathcal{Y}, \mathcal{H}_2 \cap \mathcal{Y}) > k$. Let D be a combed disk diagram supported by $\{\mathcal{Y}, \mathcal{H}_1, \mathcal{Z}, \mathcal{H}_2\}$. Every dual curve to $\mathcal{Y}$ must intersect $\mathcal{Z}$. However, there are more than k such dual curves and hence more than k hyperplanes intersecting both $\mathcal{Y}$ and $\mathcal{Z}$. This contradicts $\mathcal{Y}$ and $\mathcal{Z}$ being k-separated.

To prove the lemma's second claim, suppose g is a minimal geodesic from $\mathcal{Y}$ to $\mathcal{Z}$, and assume that $\mathcal{Y}$ and $\mathcal{Z}$ are not strongly separated. Suppose $\mathcal{H}$ is a hyperplane intersecting both $\mathcal{Y}$ and $\mathcal{Z}$ and let D be a combed diagram supported by the hyperplane-path sequence $\{\mathcal{Y}, g, \mathcal{Z}, \mathcal{H}\}$ with boundary path $\{Y, g, Z, H\}$. Every dual curve to H must intersect g. Since g is geodesic, $|g| = |H|$ and it follows no dual curve to $\mathcal{Y}$ can intersect g. So every dual curve to Y intersects Z. It follows D is an Euclidean rectangle. Hence, there is a dual curve C from Y to Z which has g as part of its boundary path. So, g is contained in the carrier of the hyperplane $\mathcal{C}$ which intersects both $\mathcal{Y}$ and $\mathcal{Z}$. □

Definition 4.1.5. If infinitely many hyperplanes intersect both $\mathcal{Y}$ and $\mathcal{Z}$ then we say $\mathcal{Y}$ and $\mathcal{Z}$ are ∞-*connected.*

4.2 k–Chain Separated Hyperplanes

A pair of non-intersecting hyperplanes are k–chain connected (formally defined below) if there is an appropriate sequence of sets of hyperplanes connecting them. Hyperplanes that are not k–chain connected, k–chain separated hyperplanes, provide a generalization of the notion of k–separated hyperplanes.

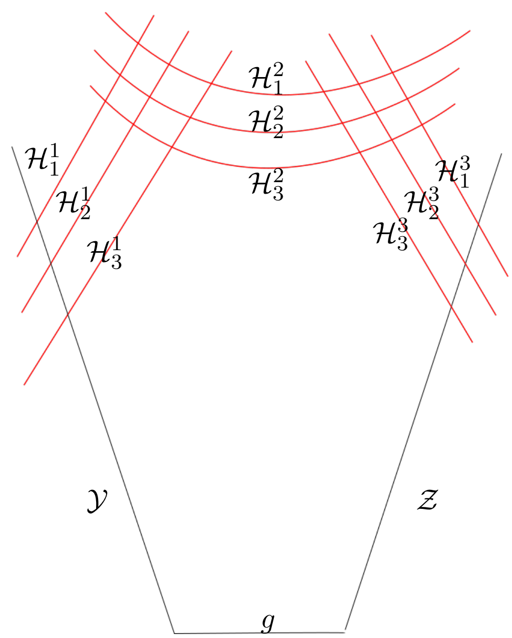

Figure 4.1: The hyperplanes $\mathcal{Y}$ and $\mathcal{Z}$ above are 3–chain connected.

Definition 4.2.1 (k–chain connected)**.** $\mathcal{Y}$ and $\mathcal{Z}$ are *k–chain connected* if there exists a sequence of length k sequences of hyperplanes:

$$S_1 = \{\mathcal{H}_1^1, \mathcal{H}_2^1, ..., \mathcal{H}_k^1\}$$

$$S_2 = \{\mathcal{H}_1^2, \mathcal{H}_2^2, ..., \mathcal{H}_k^2\}$$

$$...$$

$$S_m = \{\mathcal{H}_1^m, \mathcal{H}_2^m, ..., \mathcal{H}_k^m\}$$

satisfying the following properties:

I For each i, hyperplanes in S_i pairwise do not intersect.

II For each $i < m$, each hyperplane in S_i intersects each hyperplane in S_{i+1}.

III Every hyperplane in S_1 intersects $\mathcal{Y}$ and every hyperplane in S_m intersects $\mathcal{Z}$.

Definition 4.2.2. $\mathcal{Y}$ and $\mathcal{Z}$ are *k–chain separated* if they are not k–chain connected.

4.3 Symbolically k–Chain Separated Hyperplanes

Definition 4.3.1. The hyperplanes $\mathcal{H}$ and $\mathcal{H}'$ are of the same type, if they are in the same orbit of $Aut(X)$. The hyperplanes $\mathcal{H}$ and $\mathcal{H}'$ are of non-intersecting type if $g\mathcal{H} \cap \mathcal{H}' = \emptyset$ for all $g \in Aut(X)$.

Definition 4.3.2. Let $Q = \{\mathcal{H}_1, ..., \mathcal{H}_m\}$ be a sequence of hyperplanes. Define

$$Type(Q) = (T_1, ..., T_m)$$

where T_i is the hyperplane type (orbit class) of $\mathcal{H}_i$. Note that the tuple $Type(Q)$ is ordered.

The next set of definitions provide a further strengthening of the notion of k–chain separated hyperplanes which allows us to prove stronger divergence bounds in the next section. The following definitions were created with the key example of right-angled Coxeter groups in mind.

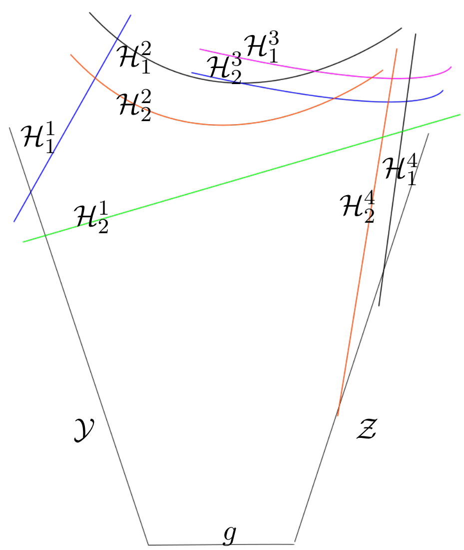

Figure 4.2: Hyperplanes $\mathcal{Y}$ and $\mathcal{Z}$ are symbolically 2–chain connected. The hyperplane colors signify their type.

Definition 4.3.3 (Symbolically k–chain Connected)**.** $\mathcal{Y}$ and $\mathcal{Z}$ are *symbolically* *k–chain connected* if there exists a sequence of length k sequences of hyperplanes:

$$S_1 = \{\mathcal{H}_1^1, \mathcal{H}_2^1, ..., \mathcal{H}_k^1\}$$

$$S_2 = \{\mathcal{H}_1^2, \mathcal{H}_2^2, ..., \mathcal{H}_k^2\}$$

$$...$$

$$S_m = \{\mathcal{H}_1^m, \mathcal{H}_2^m, ..., \mathcal{H}_k^m\}$$

satisfying the following five properties:

I For each $i \leq m$ and $j < k$, H_j^i and H_{j+1}^i are of non-intersecting type.

II For each $1 < i \leq m$, every hyperplane in S_i intersects $\mathcal{H}_1^{i-1}$.

III For each $i < m$ and $j \leq k$, there exists an integer $c(i,j)$ such that $i < c(i,j) \leq m$ and $\mathcal{H}_j^i$ intersects every hyperplane in $S_{c(i,j)}$.

Additionally, for all $j, j' \leq k$ and for all $i < m$, $Type(S_{c(i,j)}) = Type(S_{c(i,j')})$.

IV Every hyperplane in S_1 intersects $\mathcal{Y}$ and every hyperplane in S_m intersects $\mathcal{Z}$.

V Let g be a minimal geodesic from $\mathcal{Y}$ to $\mathcal{Z}$. For all $i \leq m$ and $j \leq k$, g and $\mathcal{H}_{j-1}^i$ lie in different half-spaces of $\mathcal{H}_j^i$.

Remark 4.3.3.1. By Lemma 4.1.4, property V in the definition above necessarily implies that $m > 1$. Furthermore, it follows that if property V is true for a minimal geodesic from $\mathcal{Y}$ to $\mathcal{Z}$, then it is true for all minimal geodesics from $\mathcal{Y}$ to $\mathcal{Z}$.

Definition 4.3.4 (symbolically k–chain separated)**.** Two non-intersecting hyperplanes are *symbolically k–chain separated* if they are not symbolically k–chain connected and are not k–chain connected.

Chapter 5

Divergence of CAT(0) Cube Complexes

We define the *hyperplane divergence function*, HDiv, which measures the length of a shortest path between two hyperplanes which avoids a ball centered on one of the hyperplanes. In this chapter we obtain lower bounds on the Div function from lower bounds on the HDiv function.

Definition 5.0.1. Let g be a minimal geodesic from $\mathcal{Y}$ to $\mathcal{Z}$ and set $p = \mathcal{Y} \cap g$. $\mathrm{HDiv}_g(\mathcal{Y}, \mathcal{Z})(r)$ is the length of a shortest path from $\mathcal{Y}$ to $\mathcal{Z}$ which avoids the ball $B_p(r)$.

Remark 5.0.1.1. If X is one-ended then $\mathrm{HDiv}_g(\mathcal{Y}, \mathcal{Z})(r)$ always takes finite values. Additionally, if $\mathcal{Y}$ and $\mathcal{Z}$ are k–separated and g_1, g_2 are different minimal geodesics from $\mathcal{Y}$ to $\mathcal{Z}$, then by Lemma 4.1.2 and Lemma 4.1.4 we have that $\mathrm{HDiv}_{g_1}(\mathcal{Y}, \mathcal{Z})(r-$

$k) \leq \mathrm{HDiv}_{g_2}(\mathcal{Y},\mathcal{Z})(r) \leq \mathrm{HDiv}_{g_1}(\mathcal{Y},\mathcal{Z})(r+k)$. Hence, up to the usual equivalence on divergence functions, for k–separated hyperplanes it is often not relevant which minimal geodesic is used.

5.1 Divergence Theorems

This section is devoted to proving the following two theorems which provide a connection between the hyperplane separation properties defined in the previous section and divergence in X. Theorem 5.1.1 gives bounds on $\mathrm{HDiv}(\mathcal{Y},\mathcal{Z})$, and Theorem 5.1.2 gives bounds on $\mathrm{Div}(X)$.

Theorem 5.1.1. *The following are true:*

1. *Suppose X is finite-dimensional and locally compact. $\mathcal{Y}$ and $\mathcal{Z}$ are ∞-connected if and only if $HDiv(\mathcal{Y},\mathcal{Z})$ is constant.*

2. *If $\mathcal{Y}$ and $\mathcal{Z}$ are k–separated, then $HDiv(\mathcal{Y},\mathcal{Z})$ is at least linear.*

3. *If $\mathcal{Y}$ and $\mathcal{Z}$ are k–chain separated and X is finite-dimensional, then $HDiv(\mathcal{Y},\mathcal{Z})(R) \succeq \frac{1}{2}R\log_2(\log_2 R)$.*

Theorem 5.1.2. *Suppose X is essential, locally compact and with cocompact automorphism group.*

1. *If $\mathcal{Y}$ and $\mathcal{Z}$ are k–separated, then $Div(X)$ is bounded below by a quadratic function.*

2. *If $\mathcal{Y}$ and $\mathcal{Z}$ are k–chain separated, then $Div(X) \succeq \frac{1}{2}R^2 \log_2(\log_2(R))$.*

3. *Suppose X has k–alternating geodesics (Definition 5.1.3). If $\mathcal{Y}$ and $\mathcal{Z}$ are symbolically k–chain separated then $HDiv(\mathcal{Y}, \mathcal{Z})$ is bounded below by a quadratic function and $Div(X)$ is bounded below by a cubic function.*

Definition 5.1.3. X has *k–alternating geodesics*, if there exists a constant M so that every geodesic of length M in X crosses a set of hyperplanes $\{\mathcal{H}_1, ..., \mathcal{H}_k\}$ such that $\mathcal{H}_i$ and $\mathcal{H}_{i+1}$ are of non–intersecting type for all $i < k$.

5.2 Divergence Theorems Proofs

In this section we will give the proofs of the stated results of the last section. This will be done through a series of lemmas.

5.2.1 k-Separated Hyperplane HDiv Bounds

Lemma 5.2.1. *Suppose X is finite-dimensional, locally compact and that $\mathcal{Y}$ and $\mathcal{Z}$ are ∞-connected. There exists a constant c such that for all r and choice of geodesic g, $HDiv_g(\mathcal{Y}, \mathcal{Z})(r) = c$.*

Proof. Fix a hyperplane H_1 intersecting both $\mathcal{Y}$ and $\mathcal{Z}$, and let $\mathcal{H}_2$ be another hyperplane intersecting both $\mathcal{Y}$ and $\mathcal{Z}$ a distance at least r from $\mathcal{H}_1$. Let D be a combed disk diagram supported by $\{\mathcal{Y}, \mathcal{H}_1, \mathcal{Z}, \mathcal{H}_2\}$. Every dual curve to $\mathcal{H}_1$ must intersect $\mathcal{H}_2$ and every dual curve to Y must intersect Z. Hence, D is an Euclidean strip of dimension $r \times d(\mathcal{Y}, \mathcal{Z})$ and the lemma follows. $\square$

Lemma 5.2.2. *Suppose $\mathcal{Y}$ and $\mathcal{Z}$ are k–separated. There exists a constant c such that $HDiv(\mathcal{Y}, \mathcal{Z})(r) \geq r + c$.*

Proof. Let g be any minimal geodesic from $\mathcal{Y}$ to $\mathcal{Z}$ and $p = g \cap \mathcal{Y}$. Let α be a path from $\mathcal{Y}$ to $\mathcal{Z}$ which avoids the ball $B_p(r)$. Let D be a combed disk diagram supported by $\{\mathcal{Y}, \alpha, \mathcal{Z}, g^{-1}\}$. At most $d = |g|$ dual curves to $\mathcal{Y}$ can intersect g and at most k dual curves to $\mathcal{Y}$ can intersect $\mathcal{Z}$. Hence, at least $r - d - k$ dual curves to $\mathcal{Y}$ intersect α. Therefore, $|\alpha| \geq r - d - k$. $\square$

5.2.2 Chain Separated Hyperplane HDiv Bounds

Most of the technical work in proving the results of this section is done in the next two lemmas.

Lemma 5.2.3. *Suppose X is finite-dimensional and $\mathcal{Y}$, $\mathcal{Z}$ are k–chain separated. Set $d = d_X(\mathcal{Y}, \mathcal{Z})$. There exists a constant $R_0(d, k)$ such that for $R > R_0$, $HDiv(\mathcal{Y}, \mathcal{Z}) \geq \frac{1}{2} R \log_2(\log_2 R)$.*

Proof. By Lemma 2.3.2, let $K > 0$ be the constant, only depending on k and X, such that a geodesic of length K in X must intersect at least k pairwise non-intersecting hyperplanes. Let g be a minimal geodesic from $\mathcal{Y}$ to $\mathcal{Z}$ and set $p = g \cap Y$. Fix

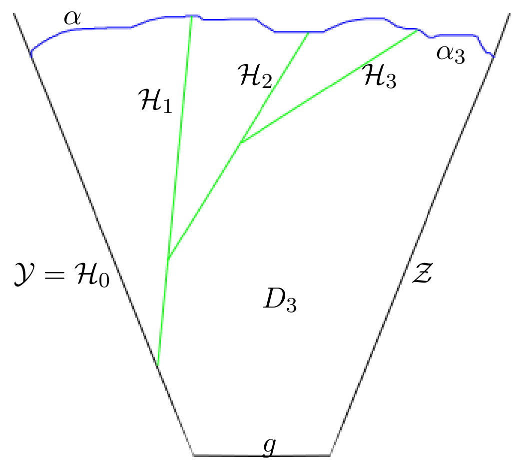

Figure 5.1: Graphic for the proof Lemma 5.2.3. The disk diagram D_3 is supported by the hyperplane-path sequence $\{\mathcal{H}_0, \mathcal{H}_1, \mathcal{H}_2, \mathcal{H}_3, \alpha_3, \mathcal{Z}, g^{-1}\}$.

$R > 0$ and let $r = \log_2(\log_2 R)$. Let α be a combinatorial path from $\mathcal{Y}$ to $\mathcal{Z}$ which avoids the ball $B_p(R)$.

Set $\mathcal{H}_0 = \mathcal{Y}$. Let D_0 be a combed disk diagram with boundary supported by the hyperplane-path sequence $\bar{A}_0 = \{\mathcal{H}_0, \alpha, \mathcal{Z}, g^{-1}\}$ and with boundary path $\bar{P}_0 = \{H_0, \alpha, Z, g^{-1}\}$. Orient H_0 from g to α. For $i \in \mathbb{Z}_{\geq 0}$ set $c_i = 2^{i+1}r(K + d)$ where $d = |g|$.

For $n \leq r$, define inductively H_n as the c_{n-1}'th dual curve to H_{n-1} in D_{n-1} and assume H_n intersects α. Orient H_n from H_{n-1} to α. Define α_n as the subpath of α from H_n to Z and β_n as the subpath of α from $\mathcal{H}_{n-1}$ to $\mathcal{H}_n$. Define D_n as the combed diagram with boundary supported by

$$\bar{A}_n = \{\mathcal{H}_0, \mathcal{H}_1, ..., \mathcal{H}_n, \alpha_n, \mathcal{Z}, g^{-1}\}$$

and with boundary path:

$$\bar{P}_n = \{H_0, H_1, ..., H_n, \alpha_n, Z, g^{-1}\}$$

obtained from D_{n-1} by Lemma 3.3.5. For $j \leq n$, define the paths in D_n:

$$T_j = H_{j+1} * H_{j+2} * ... * H_n * \alpha_n \text{ (set } T_n = \alpha_n)$$

$$B_j = Z^{-1} * g^{-1} * H_0 * H_1 * ... * H_{j-1}$$

$$B'_j = g^{-1} * H_0 * H_1 * ... * H_{j-1}$$

We will show through induction the following are true for all $n \leq r$:

A D_n is well defined. In particular, H_n intersects α.

B For $j < n$, rK dual curves to H_j in D_n, intersect T_j.

C For $n > 0$, $|\beta_n| \geq R - c_n$

Given the diagram D_{n-1}, in order to define D_n we must first show H_{n-1} has at least c_{n-1} dual curves emanating from it in D_{n-1}. Since D_n only needs to be defined for $n \leq r$, we can do this by showing $|H_r| > 0$ in D_r. Because α avoids the R-ball about p, we have:

$$\begin{aligned}
|H_r| &\geq R - \sum_{j=0}^{r-1} c_j \\
&= R - \sum_{j=0}^{r-1} 2^{j+1} r(K+d) \\
&= R - 2r(K+d) \sum_{j=0}^{r-1} 2^j
\end{aligned}$$

Using the formula for a geometric series:

$$\begin{aligned}|H_r| &\geq R - 2r(K+d)(2^r - 1)\\ &= R - \Big(2(K+d)\log_2(\log_2 R)\Big)\Big(2\log_2 R - 1\Big)\end{aligned}$$

There then exists a constant $m_1(k, d)$, such that for $R > m_1$, we have that $|H_r| \geq 0$. Thus when $R > m_1$, $H_n \geq c_n$ for $n < r$.

We now turn to the base case, $n = 0$. Hypothesis A follows immediately. Let Q denote the set of dual curves to H_0. At most d of these can intersect g. Additionally, we cannot have K dual curves in Q intersect Z. For if they did, this would imply $\mathcal{Y}$ and $\mathcal{Z}$ are k–chain connected, a contradiction. We then have that $R - K - d$ dual curves to H_0 intersect $\alpha = T_0 \subset D_0$.

The following inequalities imply that $R - K - d > rK$:

$$R > c_0 = 2r(K+d) > rK + K + d$$

Hypothesis B is then true, settling the base case.

For the general case, assume $n + 1 \leq r$ and that hypotheses A, B and C are satisfied for any $n' < n + 1$. Let $Q = \{Q_1, Q_2, ..., Q_m\}$ be the set of dual curves in D_n emanating from H_n ordered by the orientation on H_n. It follows at most $C = \sum_{j=0}^{n-2} c_j + d$ dual curves in Q can intersect B'_n (the sum does not go to $n - 1$ since the diagram is combed). Using the formula for a geometric series we have:

$$C = c_{n-1} - 2r(K+d) + d \leq c_{n-1}$$

Additionally, we cannot have K curves in Q intersect Z. For then, there is a subset of k of these dual curves, $S_1 = \{H_1^1, ..., H_k^1\} \subset Q$, corresponding to pairwise non-intersecting hyperplanes, which intersect Z. By induction hypothesis B, k dual curves to H_{n-1}, $S_2 = \{H_1^2, H_2^2, ..., H_k^2\}$, corresponding to pairwise non-intersecting hyperplanes, intersect every curve in S_1. Now, $H_k^1 * H_k^2$ is a path from H_{n-1} to Z. By the induction hypothesis B and the pigeonhole principle, k dual curves emanating from H_{n-2}, $S_3 = \{H_1^3, ..., H_k^3\}$, intersect either H_k^1 or H_k^2. Hence, every curve in S_3 intersects every curve in S_1 or S_2. Proceeding this way we can show $\mathcal{Y}$ is k–chain connected to $\mathcal{Z}$, a contradiction.

It follows for $j \geq c_{n-1} + K + d \geq C + K + d$, Q_j must intersect α. In particular, $H_{n+1} = Q_{c_n}$ must intersect α. Hence, using Lemma 3.3.5 we can define D_{n+1}, proving hypothesis A.

Note that for j such that, $c_{n-1} + K + d \leq j \leq c_n$, Q_j must intersect α. A direct calculation gives that there are at least rK of such curves. Because of this, and the boundary combinatorics preservation property of Lemma 3.3.5, hypothesis B is satisfied in D_{n+1}.

We are left to prove hypothesis C. Note that for $j > c_n$, Q_j intersects β_{n+1} in

D_n. By using the fact that α avoids $B_p(R)$ we have:

$$|\beta_{n+1}| \geq R - \sum_{j=0}^{n} c_j = R - 2r(K+d)\sum_{j=0}^{n} 2^j$$

$$= R - 2r(K+d)(2^{n+1}-1)$$

$$= R - c_{n+1} + 2r(K+d) > R - c_{n+1}$$

This proves induction hypothesis C and completes the induction.

We have thus divided α into a set of disjoint subpaths $\{\beta_i\}$ for each of which we have a lower bound. This allows us to compute a lower bound for the length of α:

$$|\alpha| \geq \sum_{i=1}^{r} |\beta_i| = \sum_{i=1}^{r} (R - c_i)$$

$$= rR - \sum_{i=0}^{r-1} c_{i+1}$$

$$= rR - 4r(K+d)\sum_{i=0}^{r-1} 2^i$$

$$= rR - 4r(K+d)(2^r - 1)$$

$$= R\log_2(\log_2(R)) - 4(K+d)\log_2(\log_2 R)\Big(\log_2 R - 1\Big)$$

There then exists a constant $m_2(k,d)$ such that for $R > m_2$,

$$|\alpha| \geq \frac{R\log_2(\log_2(R))}{2}$$

□

We are now in a position to prove Theorem 5.1.1.

Proof of Theorem 5.1.1 . The theorem follows from the above three lemmas. □

5.2.3 Symbolically Chain Separated Hyperplane HDiv Bounds

The proof of the following lemma is similar to that of Lemma 5.2.3. However, to get the quadratic bound on the HDiv function the proof requires a different counting technique.

Lemma 5.2.4. *Suppose X is essential, locally compact, has k–alternating geodesics and that Aut(X) acts cocompactly. Let $\mathcal{Y}$ and $\mathcal{Z}$ be symbolically k–chain separated. Set $d = d_X(\mathcal{Y}, \mathcal{Z})$. There exists a constant $R_0(d, k)$ such that for $R > R_0$, HDiv($\mathcal{Y}, \mathcal{Z}$) is bounded below by a quadratic function.*

Proof. Fix $R > 0$. Let g be a minimal geodesic from $\mathcal{Y}$ to $\mathcal{Z}$, $p = g \cap \mathcal{Y}$, and α a $B_p(R)$ avoidant combinatorial path from $\mathcal{Y}$ to $\mathcal{Z}$. Let M be the k–alternating constant from Definition 5.1.3. Set $c_1 = M + d$. Let c_2 be the number of different hyperplane types in X (this is finite since $Aut(X)$ acts cocompactly). Set $c = c_1(c_2)^k + 2d$. Set $r = \frac{R}{6c}$ and set $\mathcal{H}_0 = \mathcal{Y}$. Let D_0 be a combed diagram with boundary supported by $\bar{A}_0 = \{\mathcal{Y}, \alpha, \mathcal{Z}, g^{-1}\}$ and with boundary path $\bar{P}_0 = \{H_0, \alpha, Z, g^{-1}\}$. Orient H_0 from g to α.

Assume we have a combed disk diagram D_n supported by

$$\bar{A}_n = \{\mathcal{H}_0, \mathcal{H}_1, ..., \mathcal{H}_n, \alpha_n, \mathcal{Z}, g^{-1}\}$$

and with boundary path:

$$\bar{P}_n = \{H_0, H_1, ..., H_n, \alpha_n, Z, g^{-1}\}$$

where α_n is a subpath of α from $\mathcal{H}_n$ to Z. Assume H_n is oriented from H_{n-1} to α_n. Define β_n as the subpath of α from $\mathcal{H}_{n-1}$ to $\mathcal{H}_n$.

Let $Q = \{Q_1, Q_2, ..., Q_l\}$ be the set of dual curves to H_n labeled sequentially by the orientation on H_n. Define τ_n to be the largest integer such that Q_{τ_n} does not intersect α. Set $m_n = \tau_n + c$, and let $H_{n+1} = Q_{m_n}$. For $j \leq n$, define the paths in D_n:

$$T_j = H_{j+1} * H_{j+2} * ... * H_n * \alpha_n \text{ (set } T_n = \alpha_n)$$

$$B_j = Z^{-1} * g^{-1} * H_0 * H_1 * ... * H_{j-1}$$

$$B'_j = g^{-1} * H_0 * H_1 * ... * H_{j-1} \subset B_j$$

We assume by induction the following are true for $n < r$:

A D_n is well defined.

B In D_n, for $j < n$, exactly c dual curves to H_j intersect T_j and exactly τ_j dual curves to H_j intersect B_j.

C For $n > 0$, $|\beta_n| \geq |H_{n-1}| - \tau_{n-1} - c$

For the case when $n = 0$, A follows immediately and B, C are trivial. We now turn to the case $n = 1$. Note that in D_0, at most d dual curves can intersect g and M dual curves can intersect Z (since $\mathcal{Y}$ and $\mathcal{Z}$ are symbolically k–chain separated). Hence, $\tau_0 \leq d + M$, and Q_{m_0} does intersect α. Therefore, by using Lemma 3.3.5, we can define D_1. It is also clear that B holds in D_1. Furthermore, in D_0, it is clear that for $j > m_0$, Q_j intersects B_1. Hence, C is true as well. This shows the base cases $n = 0$ and $n = 1$ are true.

We now turn to the general case. Suppose the lemma is true for all integers n' such that $n' \leq n < r$. Consider the diagram D_n, and let $Q = \{Q_1, ..., Q_l\}$ be the set of dual curves to H_n labeled sequentially by the orientation on H_n.

Sub-Claim 5.2.5. *We will first show that we cannot have c_1 curves in Q intersect Z.*

Proof. We say two hyperplanes are almost symbolically k-chain connected, if they satisfy every condition of Definition 4.3.3 except maybe condition V . Suppose for a contradiction the claim is not true. It follows that $\mathcal{H}_n$ is almost symbolically k–chain connected to $\mathcal{Z}$. Assume for some $i \leq n$, $\mathcal{H}_i$ and $\mathcal{Z}$ are almost symbolically k–chain connected by sequences:

$$S_1 = \{\mathcal{P}_1^1, \mathcal{P}_2^1, ..., \mathcal{P}_k^1\}$$

$$S_2 = \{\mathcal{P}_1^2, \mathcal{P}_2^2, ..., \mathcal{P}_k^2\}$$

$$...$$

$$S_m = \{\mathcal{P}_1^m, \mathcal{P}_2^m, ..., \mathcal{P}_k^m\}$$

Additionally, we want this structure to be seen in the disk diagram D_n. So assume for every $\mathcal{P}_i^j$, there is a corresponding dual curve P_i^j in D_n. Also assume every dual curve corresponding to a hyperplane in S_1 intersects H_i, every dual curve corresponding to a hyperplane in S_m intersects Z, and $P_1^1 * P_1^2 * ... * P_1^m$ is well defined as a path in D_n from H_i to Z.

By the induction hypothesis, there are c dual cuves to H_{i-1} which intersect T_j. Let p be one such curve. Since D_n is combed, p cannot intersect T_{j+1}. Hence, p must intersect $P_1^1 * P_1^2 * ... * P_1^m$, and, consequently p must intersect P_1^j for some j. However, since the curves in S_j are pairwise disjoint, p must intersect every curve in S_j.

There are only c_2^k different possibilities for the tuple $Type(S_j)$. Hence, by the pigeonhole principle, there must be k dual curves to H_{n-1}, $S = \{P_1, ..., P_k\}$, such that the corresponding hyperplanes to the sequence:

$$S = \{P_1, P_2, ..., P_k\}$$

$$S_j = \{P_1^1, P_2^1, ..., P_k^1\}$$

$$S_{j+1} = \{P_1^2, P_2^2, ..., P_k^2\}$$

$$...$$

$$S_m = \{P_1^m, P_2^m, ..., P_k^m\}$$

almost symbolically k–chain connects $\mathcal{H}_{n-1}$ to $\mathcal{Z}$.

Proceeding in this manner, this would imply $\mathcal{Y}$ is symbolically k–chain connected to $\mathcal{Z}$, a contradiction. This finishes the proof of the subclaim. $\square$

We next want to show that m_n is well defined. By induction hypothesis B and the subclaim, at most $C = (n-2)c + d + c$ curves in Q can intersect B_n. Note that $\frac{R}{6} = rc \geq nc + c \geq C + c$. Hence, for m_n to be well defined, it is enough to know that $|H_n| \geq \frac{R}{6}$.

$\sum_{j=0}^{n-1} \tau_j$ is a count of how many dual curves have both endpoints on B_n. At most $(n-1)c$ such curves have endpoints on H_i and H_j for some $i < j < n$. At most d such curves have endpoints on g. Furthermore, by the subclaim at most $(n-1)c_1$ such curves have endpoints on H_i and Z for some $i < n$. Hence,

$$\sum_{j=0}^{n-1} \tau_j \leq (n-1)c + d + (n-1)c_1 \leq 2nc$$

Since α does not intersect $B_p(R)$, we have that:

$$\begin{aligned}
|H_n| &\geq R - \sum_{i=0}^{n-1} |H_i| \\
&= R - \sum_{i=0}^{n-1} (\tau_i + c) = R - nc - \sum_{i=0}^{n-1} \tau_i \\
&\geq R - 3nc \\
&\geq R - 3rc = \frac{R}{2}
\end{aligned}$$

Thus, m_n is well defined.

We are left to prove induction hypothesis C. For $j > m_n$, Q_j intersects β_{n+1}. Thus,

$$|\beta_n| \geq |H_{n-1}| - \tau_{n-1} - c$$

This finishes the induction.

We have broken α into a union of r subpaths $\{\beta_j\}$ for which we have a lower bound. We can then calculate a lower bound for α:

$$\begin{aligned}
|\alpha| &\geq \sum_{i=0}^{r-1} |\beta_{i+1}| \geq \sum_{i=0}^{r-1} |H_i| - \tau_i - c \\
&\geq \sum_{i=0}^{r-1} \left(\frac{R}{2} - \tau_i - c \right) \\
&\geq \frac{rR}{2} - \sum_{i=0}^{r-1} (\tau_i) - rc \\
&\geq \frac{rR}{2} - 3rc = \frac{R^2}{12c} - \frac{R}{2}
\end{aligned}$$

$\square$

5.2.4 From HDiv to Div

The following theorem allows us to deduce lower bounds for $\mathrm{Div}(X)$ through the existence of just two hyperplanes with strong enough separation properties. The proof of the theorem involves constructing an infinite sequence of nested hyperplanes. This is done primarily through the machinery developed in [CS11].

Theorem 5.2.6. *Let X be essential, locally compact and with cocompact automorphism group. Suppose HDiv$(\mathcal{Y}, \mathcal{Z}) \succeq F(r)$ for a pair of non-intersecting hyperplanes $\mathcal{Y}$ and $\mathcal{Z}$ in X. It then follows that Div$(X) \succeq rF(r)$.*

Proof. Let $\mathcal{Y}^+$ and $\mathcal{Z}^+$ be half-spaces associated to $\mathcal{Y}$ and $\mathcal{Z}$ such that $\mathcal{Y}^+ \subsetneq \mathcal{Z}^+$.

By the Double Skewering Lemma in [CS11], there exists a $\gamma \in G$ so that $\gamma\mathcal{Z}^+ \subsetneq \mathcal{Y}^+$. Note that $\mathrm{HDiv}(\gamma\mathcal{Z}, \mathcal{Z}) \geq F(r)$ since $\mathcal{Y}$ separates $\mathcal{Z}$ from $\gamma\mathcal{Z}$. By Lemma 2.3 in [CS11], γ is hyperbolic and its axis, l, intersects $\mathcal{Y}$ and $\mathcal{Z}$ (γ skewers both $\mathcal{Y}$ and

$\mathcal{Z}$).

Now, we have a chain of equally spaced pairs of hyperplanes $\{\gamma^n Z, \gamma^{n-1} Z\}$ along l (isometry moves hyperplanes through l). Hence, the divergence of the geodesic l is at least $rF(r)$. $\square$

Proof of Theorem 5.1.2 . The statements in Theorem 5.1.2 are now an easy consequence of Lemma 5.2.4, Theorem 5.1.1 and Theorem 5.2.6. $\square$

Remark 5.2.6.1. We note that Theorem 5.2.6, Theorem 5.1.2 and Theorem 6.0.3 all hold under the different assumption that X is essential, finite-dimensional and Aut(X) has no fixed point at infinity. This is true since the Double Skewering Lemma from [CS11] also holds under these assumptions.

Chapter 6

Higher Degree Polynomial Divergence of CAT(0) Cube Complexes

In this short chapter we provide an inductive definition for when a pair of hyperplanes are *degree d k–separated*. We show the divergence of two degree d k–separated hyperplanes is bounded below by a polynomial of degree d. Furthermore, under mild conditions, the existence of a pair of degree d k–separated hyperplanes in a CAT(0) cube complex X implies a degree $d+1$ polynomial lower bound on the divergence of X.

Definition 6.0.1. Hyperplanes $\mathcal{H}_1$ and $\mathcal{H}_2$ are degree 1 k–separated if $\mathcal{H}_1$ and $\mathcal{H}_2$ are k–separated. $\mathcal{H}_1$ and $\mathcal{H}_2$ are degree d k–separated if they are k–separated, and for either $i = 1$ or $i = 2$ every geodesic of length k contained in $N(\mathcal{H}_i)$ intersects a pair of degree $(d-1)$ k–separated hyperplanes.

The main theorem of this section is the following:

Theorem 6.0.2. *Suppose X is finite-dimensional. If $\mathcal{Y}$ and $\mathcal{Z}$ are degree d k–separated hyperplanes, then $Hdiv(\mathcal{Y}, \mathcal{Z})$ is bounded below by a polynomial of degree d.*

By combining Theorem 5.2.6 and 6.0.2 we immediately get the following:

Theorem 6.0.3. *Let X be an essential, locally compact CAT(0) cube complex with cocompact automorphism group. If X contains a pair of degree d k–separated hyperplanes, then $Div(X)$ is bounded below by a polynomial of degree $d+1$.*

Proof of Theorem 6.0.2 . The base case, $d = 1$, follows from 2 of Theorem 5.1.1. For the general case, assume the claim is true for degree $d-1$ k–separated hyperplanes. Suppose $\mathcal{Y}$ and $\mathcal{Z}$ are degree d k–separated. Let g be a minimal geodesic from $\mathcal{Y}$ to $\mathcal{Z}$ and let $p = g \cap \mathcal{Y}$. Fix $R > 0$, and let α be a path from $\mathcal{Y}$ to $\mathcal{Z}$ that avoids the ball $B_p(R)$. Let D be a combed disk diagram supported by $\{\mathcal{Y}, \alpha, \mathcal{Z}, g^{-1}\}$.

Orient $Y \subset D$ from g to α, and let $A = \{H_1, H_2, ..., H_n\}$ be dual curves to Y sequentially ordered by the orientation on Y. Since α does not intersect $B_p(R)$, we have that $n \geq r$. Since D is combed and since $\mathcal{Y}$ and $\mathcal{Z}$ are k–separated, it follows for $i > |g| + k$, H_i intersects α. By Definition 6.0.1, there is a subsequence $B = \{K_1, K_2, ..., K_m\} \subset A$ such that:

1. K_i intersects α

2. For i odd, K_i and K_{i+1} are degree $d-1$ k–separated.

3. $m \geq \frac{(r-k-|g|)}{k}$

4. $d(K_i, p) \leq ki + |g| + k$

For i odd, let α_i be the segment of α from $\mathcal{K}_i$ to $\mathcal{K}_{i+1}$. Note that α_i is a path from $\mathcal{K}_i$ to $\mathcal{K}_{i+1}$ which avoids the ball $B_{\mathcal{K}_i \cap \mathcal{Y}}(r - ki - |g| - k)$. By the induction hypothesis and Lemma 4.1.4, $|\alpha_i| \geq (r - ki - |g| - k)^{d-1}$. Since we have linearly many segments $\{\alpha_i\}$ whose length is bounded below by a degree $d-1$ polynomial, it follows the length of α is bounded below by a degree d polynomial. $\square$

Chapter 7

Right-angled Coxeter Group Divergence

Here we wish to apply the theorems from previous sections to the case of right-angled Coxeter groups (RACGs for short). In particular we will apply results from Section 5 to obtain a classification of RACGs of quadratic divergence and to show there are no RACGs exhibiting a divergence function strictly between quadratic and cubic. We also apply results from Section 6 to define a graph theoretic criteria that provides a degree d polynomial lower bound on divergence. Together with the thickness machinery, discussed in the next section, this allows the exact divergence of many RACGs to be known. Consequently, we can distinguish many distinct quasi-isometry classes of RACGs.

Let Γ be the graph associated to a RACG, W_Γ. Let Γ^c be the graph complement of Γ and let I be the set of isolated vertices in Γ^c. i.e.,

$$I = \{v \in V(\Gamma^c) \mid Link(v) = \emptyset\}$$

.

I forms a clique in Γ, and Γ is the graph join of the induced subgraph corresponding to I with the induced subgraph corresponding to $\Gamma - I$. Consequently, $W_{(\Gamma - I)}$ is finite index in W_Γ. Divergence is a quasi-isometry invariant, hence divergence results for $W_{(\Gamma - I)}$ apply to W_Γ.

We will from now on assume, without loss of generality, that Γ^c has no isolated vertices for all RACGs considered. By Lemma 2.4.6, the Davis complex for W_Γ under this assumption is essential.

7.1 CFS Graphs and Γ-Complete Words

The following definition is a construction used in [DT15].

Definition 7.1.1 (Γ-complete word)**.** Given a graph Γ which is not a join, let $w_0 = s_1 ... s_k$ be a word with the property that for every generator $s \in V(\Gamma)$, there exists an i such that $s_i = s$. Furthermore, $m(s_i, s_{i+1}) = \infty$ for all $1 \leq i < k$ and $m(s_1, s_k) = \infty$. Since Γ is not a join, it is always possible to define w_0, although w_0 is not unique. We call such a word a Γ-*complete word* and always denote it by w_0.

We use the definition of a CFS graph used in [BFRHS] and [DT15] (defined below). An *induced square* of a graph Γ is an embedded 4–cycle.

Definition 7.1.2. Given a graph Γ, define $\square(\Gamma)$ as the graph whose vertices are induced squares of Γ. Two vertices in $\square(\Gamma)$ are adjacent if and only if the corresponding induced squares in Γ have two non-adjacent vertices in common. For a set of induced squares $S \subset \square(\Gamma)$, define the *support* of S to be all vertices in Γ which are contained in some square in S. We say Γ is *CFS* if $\square(\Gamma)$ contains a component whose support is $V(\Gamma)$.

Remark 7.1.2.1. In [BFRHS], the graph join of a CFS graph with a clique graph is still CFS. With the assumption that Γ^c has no isolated vertices, such a graph is not possible and so we omit this from the definition. We note again, however, that the RACG corresponding to a graph that is a join with a clique is commensurable with the RACG corresponding to the graph. So the results in this section still hold in full generality.

7.2 Characterization of Linear Divergence

The authors of [DT15] characterize which 2–dimensional right-angled Coxeter groups exhibit linear divergence and the general case for arbitrary dimension is done [BFRHS]. For completeness and as a warm up for the quadratic case, we provide another proof here of this characterization.

Theorem 7.2.1. *If Γ is a join then $Div(W_\Gamma)$ is linear. Otherwise, $Div(W_\Gamma)$ is at least quadratic.*

Proof. Suppose Γ is a join. It follows $W_\Gamma = W_{\Gamma_1} \times W_{\Gamma_2}$. By the assumption that Γ^c has no isolated vertices, both W_{Γ_1} and W_{Γ_2} are infinite. Hence, $\text{Div}(W_\Gamma)$ is linear.

Now suppose Γ is not a join. Let $w_0 = s_1 s_2 ... s_k$ be a Γ–complete word. Let $\mathcal{Y}$ be the hyperplane dual to the letter s_1 in w_0 and $\mathcal{Z}$ the hyperplane dual to the letter s_k in w_0 in the Davis complex of W_Γ. Since, $m(s_1, s_k) = \infty$, it follows $\mathcal{Y}$ and $\mathcal{Z}$ do not intersect. Similarly, any hyperplane dual to the letter s_j in w_0 for $1 < j < k$, does not intersect $\mathcal{Y}$ or $\mathcal{Z}$.

We will show $\mathcal{Y}$ and $\mathcal{Z}$ are strongly separated. Suppose, for a contradiction, some hyperplane $\mathcal{H}$ intersects both $\mathcal{Y}$ and $\mathcal{Z}$. Let $\mathcal{H}$ be of type $s \in V(\Gamma)$. It follows that for every j such that $1 \leq j \leq k$, the hyperplane through the letter s_j in w_0 intersects $\mathcal{H}$. Hence, for every $t \in \Gamma$, $m(t, s) = 2$. But this implies that s is isolated in Γ^c, a contradiction.

Since $\mathcal{Y}$ and $\mathcal{Z}$ are strongly separated, by Theorem 5.1.2, $\text{Div}(W_\Gamma)$ is at least quadratic. $\square$

7.3 Characterization of Quadratic Divergence

We use results from Section 5 to characterize quadratic divergence in RACGs and show there is a gap between quadratic and cubic divergence in RACGs.

Theorem 7.3.1. *Suppose Γ is not CFS and is not a join, then W_Γ has divergence greater or equal to a cubic polynomial.*

The proof of the theorem will be given at the end of this subsection. We state the following corollaries which immediately follow.

Corollary 7.3.1.1. *W_Γ has quadratic divergence if and only if Γ is CFS and is not a join.*

Proof. If Γ is CFS and is not a join, it follows from [BFRHS] that it has quadratic divergence. The other direction follows from Theorem 7.3.1. $\square$

Corollary 7.3.1.2. *If W_Γ is strongly thick of order 2, then W_Γ has cubic divergence.*

Proof. By [BFRHS], W_Γ has at most cubic divergence. Hence, by Theorem 7.3.1, W_Γ has exactly cubic divergence. $\square$

The following lemma guarantees the existence of symbolically k–chain separated hyperplanes when Γ is not CFS.

Lemma 7.3.2. *Let M be the maximal clique size in Γ. Let $w_0 = s_1 s_2 ... s_k$ be a Γ-complete word and consider its image in the Davis complex X. Let $\mathcal{Y}$ be the hyperplane dual to w_0 which intersects s_1 and $\mathcal{Z}$ the hyperplane dual to w_0 intersecting s_k. If $\mathcal{Y}$ and $\mathcal{Z}$ are symbolically 2-chain connected then Γ is CFS.*

Proof. Assume $\mathcal{Y}$ and $\mathcal{Z}$ are symbolically 2-chain connected by sequences:

$$S_1 = \{\mathcal{H}^1, \mathcal{K}^1\}$$

$$S_2 = \{\mathcal{H}^2, \mathcal{K}^2\}$$

$$...$$

$$S_m = \{\mathcal{H}^m, \mathcal{K}^m\}$$

Let $a = \{a_1, a_2, ..., a_m\}$ and $b = \{b_1, b_2, ..., b_m\}$ be the letters in Γ corresponding respectively to the hyperplanes $\{\mathcal{H}^1, ..., \mathcal{H}^m\}$ and $\{\mathcal{K}_1, ..., \mathcal{K}^m\}$. It follows $a \cup b$ forms a CFS subgraph, Δ, of Γ.

Any hyperplane intersecting w_0 cannot intersect $\mathcal{Y}$ or $\mathcal{Z}$. Thus any such hyperplane separates $\mathcal{Y}$ from $\mathcal{Z}$. Consequently each hyperplane intersecting w_0 must intersect $\mathcal{H}_i$ and $\mathcal{K}_i$ for some i. Let $L(w_0)$ be the set of generators in the word w_0, namely $L(w_0) = \{s_1, s_2, ..., s_k\}$. It follows that for each $s \in L(w_0)$, there exists a j such that s commutes with both $h_j, k_j \in \Delta$.

Given $s \in L(w_0)$, assume $s \notin a \cup b$ as a vertex of Γ, and assume s does not commute with every generator in $a \cup b$. Let $t \in a \cup b$ be such that $m(s,t) = \infty$, $t \in \{a_r, b_r\}$ for some r, and $m(s, a_j) = m(s, b_j) = 2$ for some j with $|r - j| = 1$. This is possible by the above paragraph. It follows $\{s, t, a_j, b_j\}$ forms an induced square which shares two non-adjacent vertices with a square in Δ.

On the other hand, suppose $s_i \in L(w_0)$ commutes with every generator in $A \cup B$. s_{i+1} (if $i = k$ set $s_{i+1} = s_1$) commutes with a_j and b_j for some j. We then have that $\{s_i, s_{i+1}, a_j, b_j\}$ forms an induced square that shares two non-adjacent vertices with a square in Δ.

We have thus shown every generator in $L(w_0)$ is either contained in Δ or contained in an induced square C that shares two non-adjacent vertices with a square in Δ. Since $L(w_0)$ contains every generator in Γ, we have shown that Γ is CFS. $\square$

Lemma 7.3.3. *The Davis complex for W_Γ has 2-alternating geodesics.*

Proof. Choose M to be one larger than the maximal clique size in Γ. Let $g = s_1 s_2 ... s_M$ be a geodesic of length M with $s_i \in \Gamma$. By Tits' solution to the word problem (see [Dav08]), it follows that $m(s_i, s_j) = \infty$ for some $1 \leq i < j \leq k$. It follows the hyperplane intersecting s_i and the hyperplane intersecting s_j are of non-intersecting type. $\square$

Proof of Theorem 7.3.1. Theorem 7.3.1 now follows from the above two lemmas and Theorem 5.1.2. $\square$

7.4 Higher Degree Polynomial Divergence in RACGs

In this section, we apply results from Section 6 to give graph-theoretic criteria which imply lower bounds on the divergence of a RACG. Together with the machinery of thickness (see Sections 2.5 and 8) which provides upper bounds on divergence, these results allow one to compute the exact divergence of many RACGs.

Definition 7.4.1. Given distinct vertices $s, t \in \Gamma$, (s, t) is a *non-commuting pair* if s is not adjacent to t in Γ.

Definition 7.4.2. A non-commuting pair (s, t) is rank 1 if s, t are not contained in some induced square of Γ. Additionally (s, t) are rank n if either every non-commuting pair (s_1, s_2), with $s_1, s_2 \in Link(s)$, is rank $n-1$ or every non-commuting pair (t_1, t_2), with $t_1, t_2 \in Link(t)$, is rank $n-1$.

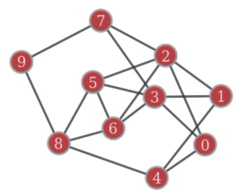

Figure 7.1: The non-commuting pairs (4,6), (4,5), (4,9), (5,9) and (6,9) for the example graph above are rank 1. It then follows that the non-commuting pair (7,8) is rank 2. Taking this one step further, we see that the non-commuting pair (9,0) is rank 3. By Theorem 7.4.3, the RACG associated to the above graph has divergence bounded below by a polynomial of degree 4. Furthermore, it can easily be checked using techniques from Section 8 that this RACG is thick of order 3 and so the divergence of this group is exactly a quartic polynomial.

Theorem 7.4.3. *Suppose Γ contains a rank n pair (s, t), then $Div(W_\Gamma)$ is bounded below by a polynomial of degree $n + 1$.*

Proof. Let M be the maximal clique size in Γ. We claim that hyperplanes of type s and t must be degree n M-separated, in the sense of Definition 6.0.1. If this claim is shown, the theorem follows from Theorem 6.0.3.

We first prove the base case when $n = 1$. Suppose, for a contradiction, $\mathcal{Y}$ and $\mathcal{Z}$ are of type s and t respectively and that $M+1$ hyperplanes intersect both $\mathcal{Y}$ and $\mathcal{Z}$. It follows from Lemma 7.3.3 that two such hyperplanes, $\mathcal{H}$ and $\mathcal{H}'$ are respectively of type $a, b \in \Gamma$ where (a, b) is a non-commuting pair. However, it then follows $\{s, a, b, t\}$ is an induced square in Γ, contradicting (s, t) being rank 1.

For the general case, suppose (s, t) are rank n and $\mathcal{Y}$ and $\mathcal{Z}$ are hyperplanes of type s and t respectively. Without loss of generality, assume every non-commuting pair (s_1, s_2), with $s_1, s_2 \in Link(s)$, are rank $n - 1$. By the induction hypothesis, hyperplanes of type s_1 and type s_2 are degree $n - 1$ M-separated.

Consider any geodesic, $g \subset N(\mathcal{Y})$, of length $M + 1$. By Lemma 7.3.3, g crosses two hyperplanes of non-commuting type, say of type s_1 and type s_2. By the above paragraph, (s_1, s_2) must be degree $n - 1$ $M - 1$-separated. The claim then follows.

□

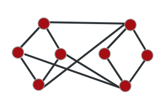

Figure 7.2: A graph that is not a join, is not CFS and only contains rank 0 and rank 1 pairs of vertices.

Remark 7.4.3.1. It is not true that the largest rank of a pair of vertices of a graph determines the corresponding RACG's divergence. The graph, Γ, in Figure 7.2 is not a join and is not CFS. Therefore, the divergence of W_Γ is at least cubic by Theorem 7.3.1. In fact, by applying Theorem 8.3.1 to obtain an upper bound, the divergence is determined to be exactly cubic. Furthermore, every non-adjacent pair of vertices in Γ is either rank 0 or rank 1. It follows we can only obtain a quadratic lower bound on the divergence of W_Γ through Theorem 7.4.3.

Chapter 8

Thick Structures on RACGs

In this chapter Γ will always denote a simplicial graph corresponding to a RACG W_Γ. **As in the previous chapter, without loss of generality, we assume that Γ^c has no isolated vertices.**

We define hypergraphs Λ_i derived from Γ for integers $i \geq 0$. Using this construction we define the hypergraph index of a right-angled Coxeter group. We show the hypergraph index gives an upper bound on the order of thickness and of algebraic thickness of a given non-relatively hyperbolic RACG. By Theorem 2.5.6 the order of thickness provides an upper bound on the divergence.

The thick structures defined here are similar to those defined in [BHS17]; however, we give lower order structures which provide good upper bounds on the divergence.

Furthermore, in Section 8.4 we provide examples of RACGs that are thick of order at most n but are algebraically thick of order at least $n+1$ and at most $2n-1$. On the other hand, the structures in [BHS17] are all algebraically thick.

Furthermore, in [Levb] it is proven that the hypergraph index is a quasi-isometry invariant of 2–dimensional RACGs.

8.1 Thickness of Order 0 and 1

The next two theorems summarize results in the literature that give many equivalent descriptions of thick of order 0 and thick of order 1 right-angled Coxeter groups. The proof follows from work in [BHS17], [BFRHS], [DT15] and this thesis.

Theorem 8.1.1 (Thick of order 0 classification)**.** *The following are equivalent:*

1. $\Gamma = A \star B$, *with* A *and* B *each containing a pair of non-adjacent vertices*
2. W_Γ *is algebraically thick of order 0*
3. W_Γ *is thick of order 0*
4. *The divergence of* W_Γ *is linear*
5. Γ *has hypergraph index 0*

Proof. The implications $1 \to 2 \to 3 \to 4 \to 1$ are either obvious or follow from [BHS17]. $5 \to 1 \to 5$ follows from the definition of hypergraph index. $\square$

Theorem 8.1.2 (Thick of order 1 classification)**.** *The following are equivalent:*

1. *Γ is CFS and $\Gamma \neq A \star B$, with A and B each containing a pair of non-adjacent vertices*

2. *W_Γ is algebraically thick of order 1*

3. *W_Γ is thick of order 1*

4. *The divergence of W_Γ is quadratic*

5. *Γ has hypergraph index 1*

Proof. The implication $1 \to 2 \to 3 \to 4$ are either obvious or follow from [BHS17] and [BFRHS]. $4 \to 1$ follows from Theorem 7.3.1. $5 \to 3$ follows from Theorem 8.3.1 below. $1 \to 5$ is an easy exercise. $\square$

8.2 Hypergraph Index Definition

8.2.1 Special Wide and Strip Subgroups

Definition 8.2.1 (Wide and strip subgraphs)**.** Let Γ be a simplicial graph. Let $\Omega = \Omega(\Gamma)$ denote the set of induced subgraphs of Γ such that given $L \in \Omega$, $L = A \star B$ where A and B are induced subgraphs which each contain a pair of non-adjacent vertices. Furthermore, L is maximal in Ω, i.e. if $L \subset L'$ for some $L' \in \Omega(\Gamma)$, then $L = L'$. The subgraphs in Ω are the *wide subgraphs* of Γ.

Let $\Psi = \Psi(\Gamma)$ denote the set of induced subgraphs of Γ such that given $L \in \Psi$, $L = A \star K$ where A is a set of two non-adjacent vertices and K is a non-empty clique. Furthermore, we require that if $L \subset L'$ for any $L' \in \Omega(\Gamma) \cup \Psi(\Gamma)$ then $L = L'$. The subgraphs in Ψ are the *strip subgraphs* of Γ.

Remark 8.2.1.1. By [BFRHS], Ω characterizes all maximal special subgroups of Γ which are wide (see Section 2.5 for the relevant definition). The term "strip subgraphs" is used since given $L = A \star K \in \Psi$, the Cayley graph of W_L is isometric to $\mathbb{Z} \times Q$, where Q is isometric to a cube of dimension $|W_K|$.

8.2.2 Hypergraph Index

We first recall the definition of a hypergraph. A *hypergraph* H consists of a set of vertices $V(H)$ and a set of hyperedges, $\mathcal{E}(H)$. An element of $\mathcal{E}(H)$ is a subset of $V(H)$ consisting of any number of vertices (edges in a standard graph only contain subsets of two elements).

Definition 8.2.2 (Lambda hypergraphs)**.** For each integer $i \geq 0$, we define the hypergraph $\Lambda_i = \Lambda_i(\Gamma)$ inductively. For each i, the vertex set of Λ_i is $V(\Gamma)$, the same as that of Γ.

1. For every $L \in \Omega(\Gamma) \cup \Psi(\Gamma)$, $V(L)$ is a hyperedge of Λ_0.

2. For $H, H' \in \Lambda_i$, set $H \equiv_i H'$ if there are hyperedges

$$H = H_0, H_1, ..., H_n = H' \in \mathcal{E}(\Lambda_i)$$

   such that for each j, $0 \leq j < n$, $H_j \cap H_{j+1}$ contains a pair of non-adjacent vertices. A hyperedge of Λ_{i+1} is the union of the vertices of a maximal set of pairwise $\equiv_i$-equivalent hyperedges of Λ_i.

For an example of these hypergraphs, see Figure 8.1 at the end of this section.

Given a hyperedge H of Λ_i, we define W_H as the special subgroup of W_Γ induced by the vertices of H.

Definition 8.2.3 (Hypergraph index)**.** Γ has *hypergraph index* $h \in \mathbb{N}$, if some hyperedge in $\Lambda_h(\Gamma)$ contains every vertex of Γ and no hyperedge of $\Lambda_{h-1}(\Gamma)$ contains every vertex of Γ. Additionally, it is required that the set of wide subgraphs, $\Omega(\Gamma)$, is not empty. If there is no such h or $\Omega(\Gamma)$ is empty, then we say Γ has infinite hypergraph index. The *hypergraph index of a right-angled Coxeter group,* W_Γ, is the hypergraph index of Γ.

Remark 8.2.3.1. It is not difficult to show, given the results of [BHS17], that Γ has hypergraph index $h = \infty$ if and only if W_Γ is relatively hyperbolic.

Remark 8.2.3.2. For $L = A \star K, L' = A' \star K' \in \Psi(\Gamma)$ distinct strip subgraphs, it follows that $A \neq A'$. For if $A = A'$, by the maximal property of strip subgraphs, there must be vertices $k \in K$ and $k' \in K'$ such that k and k' are not adjacent in Γ. Hence, $A \star (K \cup K')$ is contained in some subgraph of $\Omega(\Gamma)$, which is not allowed by the definition of strip subgraphs.

We define the *realization* of $\Lambda_i(\Gamma)$. These are cosets of special subgroups of W_Γ corresponding to hyperedges of $\Lambda_i(\Gamma)$, but excluding hyperedges corresponding to strip subgroups.

Definition 8.2.4. The realization $\mathcal{R}_i = \mathcal{R}_i(\Gamma)$ of a graph Γ is the set of cosets

$$\mathcal{R}_i = \{gW_H \subset W_\Gamma \mid H \text{ is a hyperedge of } \Lambda_i(\Gamma),\ H \notin \Psi(\Gamma),\ g \in W_\Gamma\}$$

Recall $\mathcal{H}(\Lambda_i(\Gamma))$ is the set of hyperedges of $\Lambda_i(\Gamma)$. By $H \notin \Psi(\Gamma)$, we mean that the subgraph of Γ induced by vertices of H is not in $\Psi(\Gamma)$. We often think of the cosets in $\mathcal{R}_i$ as geometric subsets of the Davis complex Σ_Γ.

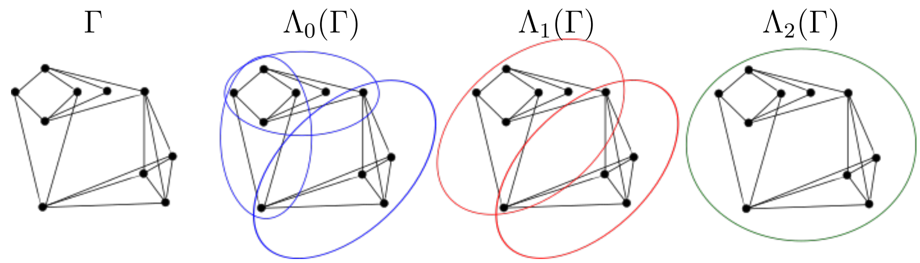

Figure 8.1: The hypergraphs $\{\Lambda_i(\Gamma)\}$ associated to the graph Γ. The hypergraph $\Lambda_0(\Gamma)$ has two hyperedges corresponding to wide subgraphs and several strip subgraph hyperedges (one is shown). As a hyperedge of $\Lambda_2(\Gamma)$ contains every vertex, the right-angled Coxeter group W_Γ has hypergraph index 2. For the relevant definitions, see definition 8.2.2 and 8.2.3.

8.3 The Hypergraph Index Bounds the Order of Thickness

We show the hypergraph index of the right-angled Coxeter group yields upper bounds on the group's order of thickness, order of algebraic thickness and divergence.

The hypergraph index yields an upper bound for the order of thickness:

Theorem 8.3.1. *If W_Γ has hypergraph index h, then W_Γ is thick of order at most h.*

Proof. We first note that the cosets in $\mathcal{R}_i$ are always convex by Lemma 2.4.2.

When $h = 0$, Λ_0 contains a hyperedge which contains every vertex. This hyperedge must be a set in Ω (by assumption Ω is non-empty). W_Γ is then wide by Theorem 8.1.1 and the base case follows.

Assume now we have hypergraph index $h = n$. By induction assume that every coset in the realization $\mathcal{R}_{n-1}$ is thick of order at most $n-1$. By definition, some hyperedge of Λ_n contains every vertex of Γ. We will show that $\mathcal{R}_{n-1}$ is a thick network of spaces and that a neighborhood of $\mathcal{R}_{n-1}$ covers W_Γ. Thus, this will show W_Γ is thick of order n.

Let M be a constant one larger than the maximal clique size of Γ. Let $\mathcal{R}'_i = \{N_M(L) | L \in \mathcal{R}_i\}$, the set of M neighborhoods of sets in $\mathcal{R}_i$. We first show that $\mathcal{R}'_i$ covers W_Γ. For this to fail, there would need to be some hyperedge corresponding to a strip subgroup $H \in \mathcal{H}(\Lambda_{n-1}) \cap \Psi$, some $g \in W_\Gamma$ and some $p \in gW_H$ such that p is not in $\mathcal{R}'_i$. However, since Γ has hypergraph index n, there must be some $H' \in \mathcal{H}(\Lambda_{n-1})$ such that $H' \cap H$ contain two non-adjacent vertices. Furthermore, $H' \notin \Psi$, since if it were $H \cup H' \in \Omega$, which would imply $p \in \mathcal{R}'_{n-1}$. It follows $H' \in \mathcal{R}_i$. However, it now follows that $p \in N_M(H')$, a contradiction. This establishes that $\mathcal{R}'_i$ covers W_Γ.

Let p, q be points in the Cayley graph of W_Γ. Let g be a geodesic connecting them. Each edge in g is contained in some coset of $\mathcal{R}'_i$. Furthermore, it follows for any two cosets corresponding to adjacent edges there are $g_0W_{H_0}, g_1W_{H_1}, ..., g_nW_{H_n} \in \mathcal{R}_{i-1}$ such that $N_M(g_iW_{H_i}) \cap N_M(g_{i+1}W_{H_{i+1}})$ is infinite. Furthermore, n is bounded by a constant only depending on Γ. This proves $\mathcal{R}_{n-1}$ is a thick network of spaces.

$\square$

The next corollary follows from the above theorem and Theorem 2.5.6.

Corollary 8.3.1.1. *If W_Γ has hypergraph index h, then the divergence of W_Γ is bounded above by a polynomial of degree $h+1$.*

The hypergraph index also provides an upper bound on the order of algebraic thickness:

Theorem 8.3.2. *If W_Γ has hypergraph index $h > 0$, then W_Γ is algebraically thick of order at most $2h-1$.*

Proof. The proof will be by induction. The base case when $h = 1$ follows from Theorem 8.1.2.

Assume the claim is true for graphs of hypergraph index h and suppose Γ has hypergraph index $h+1$. Let $\{E_1, ..., E_m\}$ be hyperedges of $\Lambda_h(\Gamma) = \Lambda_h$ which are not strip subgraphs (Λ_h is the hypergraph from Definition 8.2.2). By the induction hypothesis, the subgroups $\{W_{E_1}, ..., W_{E_m}\}$ are algebraically thick of order at most $2h-1$. Let $\{S_1, ..., S_r\}$ be hyperedges of Λ_h corresponding to strip subgraphs.

Since Λ has hypergraph index $h+1$ and by remark 8.2.3.2, for each S_i there is some $E_j \in \{E_1, ..., E_m\}$ such that $S_i \cap E_j$ contain two non-adjacent vertices. Set $\bar{S}_i = E_j$. By [BHS17, Proposition A.2], it follows that $W_{S_i \cup \bar{S}_i}$ is thick of order at most $(2h-1)+1 = 2h$. W_Γ is then algebraically thick of order at most $2h+1$ with respect to the special subgroups:

$$\{W_{E_1}, ..., W_{E_m}\} \cup \{W_{S_1 \cup \bar{S}_1}, ..., W_{S_r \cup \bar{S}_r}\}$$

□

8.4 Thickness $\neq$ Algebraic Thickness

As described in the introduction, there are known examples of groups which are thick of order 1, but are not algebraically thick of order 1. However, by Theorem 8.1.2 we know for the class of right-angled Coxeter groups thickness of order 1 is equivalent to algebraic thickness of order 1. The following natural question suggests itself: for at least the class of right-angled Coxeter groups, is algebraic thickness of order n equivalent to thickness of order n?

The goal of this section is to provide a negative answer to the above question: for each positive integer $n > 1$, there exists a right-angled Coxeter group that is thick of order n but is not algebraically thick of order n. This is the main content of Theorem 8.4.1 stated below.

Theorem 8.4.1. *Given an integer $n > 1$, let Γ be a graph satisfying the following* *hypotheses:*

1. *There is a subgraph, $B \subset \Gamma$ such that $V(\Gamma) \setminus V(B)$ is just two vertices: u and v.*

2. *B has hypergraph index $n-1$*

3. *$Link(u)$ is two non-adjacent vertices of B. Similarly, $Link(v)$ is two non-adjacent vertices of B.*

4. *For all $s \in \Gamma - Star(u)$, (u,s) is a rank n pair. Similarly, for all $s \in \Gamma - Star(v)$, (v,s) is a rank n pair.*

It follows the divergence of W_Γ is a polynomial of degree $n+1$, W_Γ is thick of order n, and W_Γ is algebraically thick of order d, where $n < d \leq 2n-1$. Furthermore, for every $n > 1$ such a graph, Γ, exists.

Figure 8.2 gives a family of graphs which can be readily checked to satisfy the hypotheses of Theorem 8.4.1. This family of graphs proves the last statement of the theorem, namely the existence of such graphs.

For the remainder of this section we fix an integer $n > 1$ and a graph Γ satisfying the hypotheses of Theorem 8.4.1. Furthermore we fix $u, v \in V(\Gamma)$ as in the statement of the theorem.

W_Γ decomposes as the amalgamated product $W_\Gamma = W_{star(u)} *_{link(u)} W_B *_{link(v)} W_{star(v)}$. It follows from Bass-Serre theory that W_Γ acts on a tree $\mathcal{T}$, with fundamental domain the graph of groups shown in figure 8.3. Fix this tree $\mathcal{T}$.

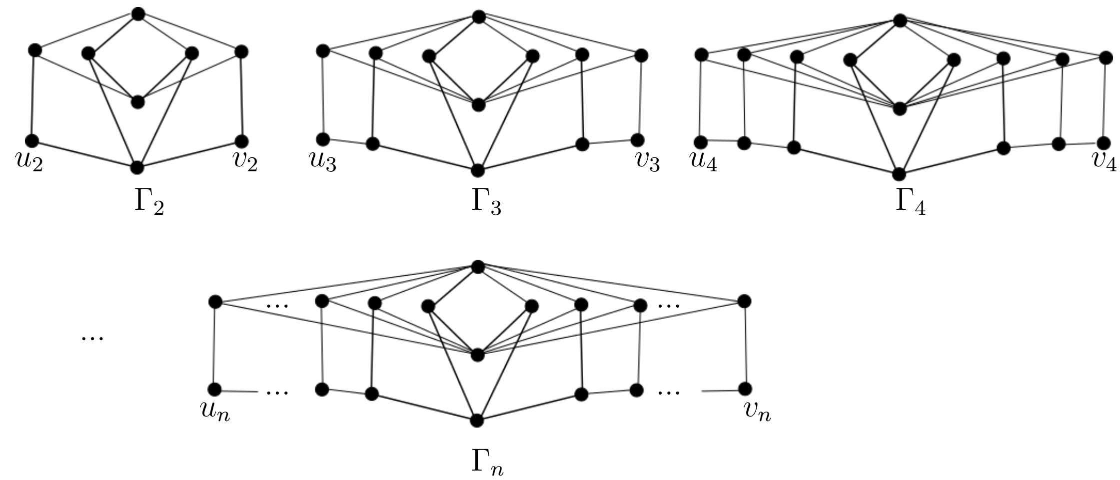

Figure 8.2: The given family of graphs provides a family of right-angled Coxeter groups, W_{Γ_n}, for $n > 1$. W_{Γ_n} is thick of order n but is algebraically thick of order strictly larger than n.

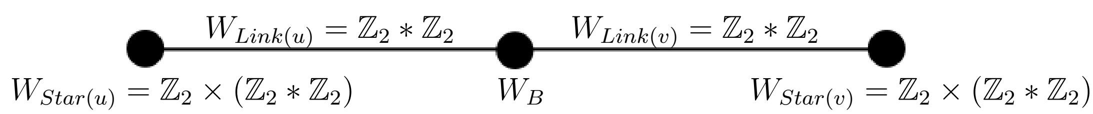

Figure 8.3: Graph of groups corresponding to W_Γ.

The proof of this theorem relies on some preliminary lemmas which we first prove.

Given w a minimal length expression of a word in W_Γ, let $L_s(w)$ be the number of occurrences of the generator s in w.

Lemma 8.4.2. *Let $w \in W_\Gamma$ be a hyperbolic isometry of the Bass-Serre tree $\mathcal{T}$. The bi-infinite geodesic ...www..., in the Davis complex Σ_Γ, has polynomial divergence of degree $n+1$.*

Proof. Since w is hyperbolic, by putting a reduced expression of w^i into normal form, we see that either $L_u(w^i)$ or $L_v(w^i)$ grows linearly with i. Without loss of generality, assume $L_u(w^i)$ grows linearly with i. Given a reduced expression, g, of w^i, it follows that given two occurrences of u in g there must exist some $s \in \Gamma$, which is not adjacent to u, between these occurrences (i.e. $g = ...u...s...u...$). Since for any such s, (s, u) forms a rank n pair by a hypothesis of Theorem 8.4.1, by a slight modification of the proof of Theorem 7.4.3, it follows that the bi-infinite geodesic $...www...$ has polynomial divergence of degree $n + 1$. □

Lemma 8.4.3. *Any quasi-isometrically embedded thick of order* $n - 1$ *subgroup is contained in a conjugate of* W_B*.*

Proof. Let G be such a thick of order $n-1$ quasi-isometrically embedded subgroup of W_Γ. Given $w \in G$, w cannot act as a hyperbolic isometry of the Bass-Serre tree $\mathcal{T}$, for then by Lemma 8.4.2, G would have divergence at least a polynomial of degree $n+1$ which is not possible since thick of order $n-1$ groups have divergence at most n by [BD14, Corollary 4.17]. It follows that any $w \in G$ acts elliptically on $\mathcal{T}$.

Since two elliptic isometries with disjoint fixed point sets generate a hyperbolic element (see [CM87, 1.5]), we have that every element of G is contained in some conjugate of W_B, some conjugate of $W_{Star(u)}$ or some conjugate of $W_{Star(v)}$. However, $W_{Star(u)}$ and $W_{Star(v)}$ are both virtually $\mathbb{Z}$ and so cannot contain a thick of order $n - 1$ subgroup. Thus, G must be contained in some conjugate of W_B. □

Lemma 8.4.4. *Let* $\{G_1, ..., G_m\}$ *be a finite set of subgroups contained in a conjugate of* W_B*. The subgroup,* G*, generated by* $\cup_{i=1}^n G_i$ *is infinite index in* W_Γ*.*

Proof. Let $H = \mathbb{Z}_2 * \mathbb{Z}_2$ and a, b the canonical generators of H. Define the homomorphism $\phi : W_\Gamma \to H$ by the map on generators: $\phi(u) = a$, $\phi(v) = b$, and $\phi(s) = 1$ for $s \neq u, v$.

Let $w \in G$. We can write $w = g_1 w_1 g_1^{-1} g_2 w_2 g_2^{-1} ... g_k w_k g_k^{-1}$ where $w_i \in W_B$ for each $1 \leq i \leq k$. It follows that $\phi(w) = 1$.

For a contradiction, suppose G is finite index in W_Γ. It follows for some $i > 0$ large enough, G must contain a word w of one of the following forms: $w = (uv)^i$, $w = (vu)^i$, $w = (uv)^i u$ or $w = (vu)^i v$, However, for each of these cases, $\phi(w) \neq 1$, a contradiction. $\square$

We are now in a position to prove Theorem 8.4.1:

Proof of Theorem 8.4.1. Given an integer $n > 1$, a graph Γ satisfying the hypotheses of the theorem exists by the family of examples given in Figure 8.2. In fact, one can construct many such families. Fix such a graph Γ.

It is immediate Γ has hypergraph index n, as B has hypergraph index $n - 1$ and Γ consists of the addition of two strip subgraphs to B. By Theorem 8.3.1, W_Γ is thick of order at most n. By Lemma 8.4.2, W_Γ has divergence a polynomial of degree $n+1$. By the lower bound on thickness provided by the divergence function, W_Γ is thick of order exactly n.

By Lemma 8.4.3 and Lemma 8.4.4, W_Γ cannot be algebraically thick of order n since no finite set of thick of order at most $n - 1$ subgroups generate a finite index subgroup of W_Γ. By Theorem 8.3.2, W_Γ is algebraically thick of order d where

$n < d \leq 2n - 1.$ □

8.5 A Conjecture

The following conjecture seems to hold for all examples we know:

Conjecture 8.5.1. *The following are equivalent:*

1. Γ *has hypergraph index* h.

2. W_Γ *is thick of order exactly* h

3. $Div(\Gamma) \asymp r^{h+1}$.

If this conjecture is true, there is then a systematic way to compute the order of thickness and divergence of RACGs using the hypergraph index construction.

Chapter 9

Coxeter Groups

This chapter explores lower bounds for divergence in Coxeter groups (not necessarily right-angled). We use similar arguments to those used in Section 6. We do not make use of a cube complex in this section. Instead, we use the construction of bands in Van-Kampen diagrams which behave similarly to dual curves in CAT(0) cube complex disk diagrams. We refer the reader to [Ol'91, Chapter 4] for a background on Van-Kampen diagrams and to [Bah05] for their application to Coxeter groups.

A characterization of thick Coxeter groups is given in [BHS17, Proposition A.2] by a class of edge-labelled graphs that can be constructed by an inductive procedure. Furthermore, the authors' proof of this proposition provides an upper bound on thickness, and hence divergence, at each step of the inductive construction. One can then carefully apply the results in this section, together with the work in [BHS17], and obtain the exact divergence for a large class of Coxeter groups.

In this section, we assume all Coxeter diagrams have at least one edge. Otherwise, W_Γ is virtually trivial or virtually free and exhibits either trivial or infinite divergence.

9.1 Higher Degree Polynomial Divergence

In this section we explore a graph theoretic criteria which implies a degree d polynomial lower bound on Coxeter groups.

The locally even and locally triangle free conditions are used in the hypotheses of results in this section.

Definition 9.1.1. Let Γ be a labeled graph. The vertex $v \in V(\Gamma)$ is *r-locally triangle-free* if for all $u \in V(\Gamma)$, such that $d_\Gamma(u, v) < r$, u is not in a triangle. We say v is *r-locally even* if for all $u \in V(\Gamma)$, such that $d_\Gamma(u, v) < r$, each edge adjacent to u is even labeled or not labeled.

Definition 9.1.2. For Γ a Coxeter diagram, let L_Γ be the largest integer edge label in Γ. If Γ contains no labeled edges, set $L_\Gamma = 2$.

9.1.1 Word Ordering Lemma

The following lemma allows us to choose boundedly spaced generators in a minimal expression for a word $w \in W_{Star(v)}$ such that these generators are not v and do not sequentially coincide.

Lemma 9.1.3. *Let Γ be a Coxeter diagram. For any $v \in \Gamma$ and any minimal expression, $w = s_1 s_2 ... s_n$, $s_i \in Star(v)$, for a word $w \in W_{Star(v)}$, there exists a subsequence $\{s_{i_1}, s_{i_2}, ..., s_{i_m}\}$ such that*

1. *$s_{i_j} \neq v$ for all j.*

2. *$i_1 \leq 2$*

3. *$i_{j+1} - i_j \leq L_\Gamma$*

4. *$s_{i_{j+1}} \neq s_{i_j}$ as vertices of Γ*

5. *$m \geq \frac{n}{L_\Gamma + 1}$.*

Proof. Since w is minimal length, there cannot be two v letters appearing consecutively. Hence either the first or second letter is not v. Set s_{i_1} to be this letter. Now note that for any letter $s \in Link(v)$ and $n > L_\Gamma$, we cannot have the expression $svsv...sv$ or $svsv...svs$ of length n appearing in w for this would contradict w being reduced. Hence, there is some letter, s_{i_2} not equal to s_{i_1} or v with $i_2 - i_1 \leq L_\Gamma$. We can keep repeating this process, and the lemma follows. □

9.1.2 Bands in Van-Kampen Diagrams

Let D be a Van-Kampen diagram for a Coxeter group. Each 2–cell in D has an even number of edges along its boundary path. For a given cell and a given edge along the cell's boundary path, there is a corresponding opposite edge. Furthermore, each edge in D is contained in exactly one cell if it is a boundary edge of D and in exactly two cells if it is not. Two edges e and e' in D are *opposite connected* if there is a sequence of edges $e = e_1, e_2, ..., e_n = e'$ such that e_i is opposite to e_{i+1} in some 2-cell of D. A *band* associated to an edge e in D is the set of all edges opposite connected to e and cells adjacent to these edges.

The construction of bands is utilized in [Bah05, Section 1.4]. There it is also shown that bands do not self-intersect and cannot intersect geodesics twice.

For $u, v \in V(\Gamma)$, an *odd path* from u to v is a path in Γ which only contains edges with odd labels. Let O_v consist of vertices $u \in V(\Gamma)$ for which there is an odd path from u to v. By definition $v \in O_v$.

Let e be an edge in D labeled by some $v \in V(\Gamma)$. It is easy to check the band corresponding to e only contains edges labeled by elements in O_v.

9.1.3 Higher Degree Divergence Theorem

Theorem 9.1.4. *Let Γ be a Coxeter graph. Suppose (u, v) is a rank n pair. Without loss of generality, we assume that for all distinct $u_1, u_2 \in Link(u)$, (u_1, u_2) is a rank $n - 1$ pair in Γ. Further assume that u is n-locally triangle free and $n + 1$-locally even and that v is 1-locally even. It follows that the divergence of the bi-infinite geodesic ...uvuv... is bounded below by a polynomial of degree $n + 1$.*

The following corollary is immediate:

Corollary 9.1.4.1. *Let W_Γ be an even Coxeter group such that Γ contains no* *triangles. If (u, v) is a rank n pair in Γ, then $Div(W_\Gamma)$ is bounded below by a polynomial of degree $n+1$.*

To prove the above theorem, we will first need the following technical lemma:

Lemma 9.1.5. *Let (u, v) be as in Theorem 9.1.4. Let $g \in W_{Star(u)}$ and $h \in W_{Star(v)}$ and $p \in W_\Gamma$ be words written in a minimal length expression. Suppose $|p| \leq L_\Gamma$, $|g| \geq r$ and $|ph| \geq r$. Let α be a shortest path from g to ph in the Cayley graph of W_Γ which does not intersect $B_{id}(r)$. It follows $|\alpha|$ is bounded below by a polynomial of degree n.*

Proof. The proof will follow by induction on n. We begin with the base case where the rank of (u, v) is $n = 1$. Suppose g is given by the following expression in generators, $g = s_1 s_2 ... s_l$. Note that $l \geq r$. Let D be a Van-Kampen with boundary path $g\alpha h^{-1}p^{-1}$.

Let $T = \{s_{i_1}, s_{i_2}, ..., s_{i_m}\}$ be a subsequence of $\{s_1, s_2, ..., s_l\}$ as in Lemma 9.1.3, and $B = \{B_1, B_2, ..., B_m\}$ bands in D corresponding to each letter in T. Since u is 2–locally even, each band B_j only contains edges labeled by s_{i_j}. Furthermore, since u is 1-locally triangle free, u is not contained in a triangle. It follows for $i \neq j$, B_i and B_j do not intersect.

At most L_Γ bands can intersect p. Additionally, u and v are rank 1, and so are not in a common square of Γ. It follows for $j > L_\Gamma$, B_j intersects α. Hence, $|\alpha|$ is linear in r, proving the base case.

Now assume the theorem is true for $n-1$ and that (u, v) are of rank n. The proof proceeds almost the same way as the base case. Consider all the same notation as the base case. For i, such that $L_\Gamma + 1 < i < m$, let α_i be the segment of α between B_i and B_{i+1}, and let p_i be the segment of g from B_i to B_{i+1}. Let g_i be the word along B_i from p_i to α_i, and let h_i be the word along B_{i+1} from p_i to α_i. By the induction hypothesis, $|\alpha_i|$ is bounded below by a polynomial of degree $n-1$ in $r-i$. Hence, $|\alpha|$ is bounded below by a polynomial of degree n. $\square$

Proof of Theorem 9.1.4 . Let α be a $B_{id}(r)$ avoidant path from $(uv)^r$ to $(vu)^r$ in the Cayley graph of W_Γ. Let D be a Van-Kampen diagram with boundary path $(uv)^r\alpha(vu)^{-r}$. Since (u, v) is a non-commuting pair in Γ, no pair of bands emanating from the words $(uv)^r$ or from $(vu)^r$ along the boundary path of D can intersect. Hence, each of these bands must intersect α.

Write $(uv)^r$ as $u_1v_1u_2v_2...u_rv_r$. Let U_i, V_i be bands corresponding respectively to u_i, v_i. Let D_i be the minimal connected subdiagram of D which includes U_i and V_i. Let α_i be the segment of α contained in D_i. By Lemma 9.1.5, $|\alpha_i|$ is bounded below by a polynomial of degree n in $r-i$. Hence, $|\alpha|$ is bounded below by a polynomial of degree $n+1$. $\square$

9.2 Quadratic Divergence Lower Bound

In this section, we provide a simple criteria for when a Coxeter group must have at least quadratic divergence.

Definition 9.2.1. Given an edge labeled Coxeter graph Γ, let $\hat{\Gamma}$ be the graph resulting from collapsing odd labeled edges in Γ to a point. For $v \in \Gamma$ we denote its image in $\hat{\Gamma}$ by $\pi(v)$. Each vertex $\hat{v} \in \hat{\Gamma}$ is labeled by a list, $\pi^{-1}(\hat{v})$. Each edge in $\hat{\Gamma}$ is labeled by the same integer as the corresponding edge in Γ. Note that this new graph can have multiple edges between two vertices.

Theorem 9.2.2. *Let Γ be a Coxeter graph. If the diameter of $\hat{\Gamma}$ is larger than 2, then W_Γ has at least quadratic divergence.*

Proof. Suppose $d_{\hat{\Gamma}}(\hat{u}, \hat{v}) > 2$ for some $\hat{u}, \hat{v} \in \hat{\Gamma}$. Choose $u \in \pi^{-1}(\hat{u})$ and $v \in \pi^{-1}(\hat{v})$. It follows that $m(u, v) = \infty$. We will show that the bi-infinite geodesic $...uvuv...$ exhibits at least quadratic divergence.

Let α be a shortest $B_{id}(2r)$–avoidant path from $(uv)^r$ to $(vu)^r$. Let D be a Van-Kampen diagram with boundary path $(uv)^r\alpha(vu)^{-r}$. Write $(uv)^r = u_1v_1u_2v_2...u_rv_r$. Let U_i denote the band emanating from u_i and V_i the band emanating from v_i. Note that for any i, j, U_i cannot intersect V_j. For then, there would be an odd path in Γ from u to some u', and an odd path from v to some v', so that $m(u', v') \neq \infty$. However, this would imply $d_{\hat{\Gamma}}(\hat{u}, \hat{v}) \leq 1$, a contradiction.

Fix i. Let D_i denote the minimal connected subdiagram of D containing both U_i and V_i, and let α_i be the subsegment of α contained in D_i. Let $g = s_1...s_m$ be the word along the boundary path of U_i from u_i to α_i. It follows $m \geq r - i$. Furthermore, $s_i \in A = \{Star(t) \mid t \in \pi^{-1}(\hat{u})\}$. Note that $d_{\hat{\Gamma}}(\hat{u}, \hat{t}) \leq 1$ for $t \in A$. Let S_j be the band in D_i emanating from s_j. It follows S_j cannot intersect V_i. For then $d_{\hat{\Gamma}}(\hat{u}, \hat{v}) \leq 2$. Hence S_j intersects α_i for each j. It follows, $|\alpha_i|$ is at least linear in $r - i$. Hence, $|\alpha|$ is at least quadratic in r. $\square$

Corollary 9.2.2.1. *Let W_Γ be an even Coxeter group. If $diam(\Gamma) > 2$, then $Div(W_\Gamma)$ is at least quadratic.*

Bibliography

[Ago13] Ian Agol, *The virtual Haken conjecture*, Doc. Math. **18** (2013), 1045–1087, With an appendix by Agol, Daniel Groves, and Jason Manning. MR 3104553

[Bah05] Patrick Bahls, *Isomorphism problem in Coxeter groups*, Imperial College Press, 2005.

[BB05] Anders Bjorner and Francesco Brenti, *Combinatorics of Coxeter groups*, Graduate Texts in Mathematics, vol. 32, Springer, 2005.

[BC12] Jason Behrstock and Ruth Charney, *Divergence and quasimorphisms of right-angled Artin groups*, Mathematische Annalen **352** (2012), 339–356.

[BD14] Jason Behrstock and Cornelia Druţu, *Divergence, thick groups, and short conjugators*, Illinois J. Math. **58** (2014), no. 4, 939–980. MR 3421592

[BDM09] Jason Behrstock, Cornelia Druţu, and Lee Mosher, *Thick metric spaces, relative hyperbolicity, and quasi-isometric rigidity*, Mathematische Annalen **344** (2009), 543–595.

[Beh06] Jason Behrstock, *Asymptotic geometry of the mapping class group and teichmuller space*, Geometry & Topology **10** (2006), 1523–1578.

[BFRHS] Jason Behrstock, Victor Falgas-Ravry, Mark F. Hagen, and Timothy Susse, *Global structural properties of random graphs*, arXiv:1505.01913, to appear in International Mathematics Research Notices.

[BH99] 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. MR 1744486 (2000k:53038)

[BH16] Jason Behrstock and Mark F. Hagen, *Cubulated groups: thickness, relative hyperbolicity, and simplicial boundaries*, Groups, Geometry, and Dynamics **10** (2016), no. 2, 649–707.

[BHS17] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto, *Thickness, relative hyperbolicity, and randomness in Coxeter groups*, Algebraic & Geometric Topology **17** (2017), 795–740, With an appendix written jointly with Pierre-Emmanuel Caprace.

[CM87] Marc Culler and John W. Morgan, *Group actions on* **R**-*trees*, Proc. London Math. Soc. (3) **55** (1987), no. 3, 571–604. MR 907233

[CS11] Pierre-Emmanuel Caprace and Michah Sageev, *Rank rigidity for CAT(0) cube complexes*, Geometric and Functional Analysis **21** (2011), no. 4, 851–891.

[Dav08] Michael W. Davis, *The geometry and topology of Coxeter groups*, London Mathematical Society Monographs Series, vol. 32, Princeton University Press, Princeton, NJ, 2008. MR 2360474 (2008k:20091)

[DJ00] Michael W. Davis and Tadeusz Januszkiewicz, *Right-angled Artin groups are commensurable with right-angled Coxeter groups*, Journal of Pure and Applied Algebra **153** (2000), 229–235.

[DMS10] 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. MR 2584607 (2011d:20084)

[DT15] Pallavi Dani and Anne Thomas, *Divergence in right-angled Coxeter groups*, Trans. Amer. Math. Soc. **367** (2015), no. 5, 3549–3577. MR 3314816

[Ger] S. M. Gersten, *Divergence in 3-manifold groups*, Geom. Funct. Anal. **4**.

[Ger94] ———, *Quadratic divergence of geodesics in CAT(0) spaces*, Geom. Funct. Anal. **4** (1994), 37–51.

[Gro93] Mikhail Gromov, *Asymptotic invariants of infinite groups*, London Math. Soc. Lecture Note Ser., vol. 182, Cambridge Univ. Press, Cambridge, 1993. MR 1253544 (95m:20041)

[GS] Dominik Gruber and Alessandro Sisto, *Infinitely presented graphical small cancellation groups are acylindrially hyperbolic*, Preprint, arXiv:1408.4488.

[Hag] M. F. Hagen, *The simplicial boundary of a CAT(0) cube complex*, Algebraic & Geometric Topology **13**.

[Hag13] ———, *Weak hyperbolicity of cube complexes and quasi-arboreal groups*, Journal of Topology **7** (2013).

[KL98] Michael Kapovich and Bernhard Leeb, *3-manifold groups and nonpositive curvature*, GAFA **8** (1998), 841–852.

[Leva] Ivan Levcovitz, *Divergence of CAT(0) cube complexes and Coxeter groups*, To appear in Algebraic & Geometric Topology.

[Levb] ______, *A quasi-isometry invariant and thickness bounds for right-angled Coxeter groups*, Preprint, arXiv:1705.06416.

[Mac13] Natasa Macura, *CAT(0) spaces with polynomial divergence of geodesics*, Geometriae Dedicata **163** (2013), 361–378.

[Ol'91] A. Yu. Ol'shanskii, *Geometry of defining relations in groups*, Springer, 1991.

[OOS09] Alexander Ol'shanskii, Denis Osin, and Mark Sapir, *Lacunary hyperbolic groups*, Geometry & Topology **13** (2009), 2051–2140.

[Wis11] Daniel T. Wise, *The structure of groups with a quasiconvex hierarchy*, 2011.