On Random Groups: the Square Model at Density $d < 1/3$ and as Quotients of Free Nilpotent Groups

BY

YEN DUONG
B.A., Yale University, New Haven, CT, 2010
M.A., University of California at Santa Barbara, 2012

THESIS
Submitted as partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Mathematics
in the Graduate College of the
University of Illinois at Chicago, 2017

Chicago, Illinois

Defense Committee:
Daniel Groves, Chair and Advisor
Marc Culler, Mathematics, Statistics, and Computer Science
David Dumas, Mathematics, Statistics, and Computer Science
Kevin Whyte, Mathematics, Statistics, and Computer Science
Benson Farb, University of Chicago

3 In loving memory of my father, Son Kim Duong. I still try to “take it easy” as much as I can.

Acknowledgments

First, a huge thank you to my advisor, Daniel Groves, who helped me with all of my math and also a good amount of my life throughout my years at UIC. His influence on nearly every aspect of my life is too large to describe–perhaps the most minor is the stroller caddy that I still use every day. Thanks to my committee members, especially to David Dumas and Marc Culler for their thoughtful and thorough feedback. Thanks to Piotr Przytycki who suggested my thesis problem, to Moon Duchin for running the program which led to the second half of my thesis, for coauthoring that paper with Ayla Sánchez, Turbo Ho, Matt Cordes, and me (thanks to those three too), and for being a good friend and role model. Thanks to my friend Teddy Einstein for offering helpful suggestions and bouncing ideas around with me. Thanks to the good people of UCSB and UT Austin, especially Medina Price and Alan Reid, for giving me a warm and welcoming environment to embrace math; same to the EDGE program and the Mellon-Mays fellowship. I could not have completed this without my mother, Lan Ho, helping during and after my pregnancies.

Finally, thank you to my partner, Mark Krasniewski, for everything. This thesis is no thanks to my kids Ian and Maeve, but they deserve acknowledgment anyway since they came along during the ride.

Y.D.

Contribution of Authors

Chapter 5 of this thesis is joint work with Matthew Cordes, Moon Duchin, Meng-Che (Turbo) Ho, and Ayla Sánchez. In addition, Duchin, Ho, and Sánchez wrote an additional appendix, which I reference but do not include in this document. I contributed to the group theory, linear algebra, and application of probability calculations in Chapter 5. I omitted the last section of the paper from which Chapter 5 is derived, which is a section of experiments conducted by Duchin and Sánchez, but for coherence I included Section 2 which contains the probability calculations.

Summary

We consider two fields of study–random groups and cubulations– from their definitions and history to current research. Motivated by group properties, we apply cubulation theory to a particular model of random groups, which are random quotients of free groups. We use a wall construction on the Cayley graph of a random group G in the square model to obtain a CAT(0) cube complex on which G acts, then show that this action results in the virtual specialness of G; in particular we show that random groups in the square model at density $d < 1/3$ are residually finite. In the style of random groups, we then explore random quotients of free nilpotent groups with results about distribution of rank, probability of being abelian, and a vanishing threshold.

Contents

List of Figures

CHAPTER 1

Introduction

Geometric group theory arose out of combinatorial group theory as its own field of study when mathematicians applied topological methods to understand aspects of groups. Just as the grouper is best studied in the context of its natural habitat (the ocean), a group is best studied in the context of its actions on topological or metric spaces. For instance, some of the first constructions of algebraic topology are the fundamental group and the presentation complex, establishing a correspondence between spaces and groups. The universal cover of the presentation complex is the Cayley complex, which we focus on for a major part of this document. One very nice property that a group might have is called *residual finiteness*. A group G is residually finite if for every nonidentity element x, there is a finite quotient of G in which the image of x is nontrivial. Examples of residually finite groups are free groups and finite groups. We apply geometric group theory techniques to show that certain groups are residually finite.

Broadly, one may want to say that "most" finitely presented groups have some property P–intuitively, that if we pick a finitely presented group "at random," it almost certainly has property P. The first and most intuitive approach is the *few-relators model of random groups*, where we fix a number of generators and randomly choose N strings as relators, where N is some number. In this model, random groups have property P asymptotically almost surely (a.a.s.) if the proportion of such groups approaches 1 as the length of the relators tends toward infinity. The most common and fruitful model, however, was created by Gromov as the *density model of random groups*, which fixes a number of generators and bound on the number of relators of a group, with an identical definition of having property P a.a.s. as the few-relators model. Random groups is a thriving field of research now, and researchers have created several different models to approach such an intuitive and appealing idea. In [**27**], Odrzygóźdź defines the square model of random groups, in which the length of relators is fixed at 4 and we say random groups in the square model at density d have property P asymptotically almost surely if the proportion of such groups approaches 1 as the *number of generators* tends toward infinity. In Chapter 4 we focus on this square model of random groups. As a different approach, we define a *few-relators model of random nilpotent groups*, which is defined identically to the few-relators model of random groups but taking quotients of free nilpotent groups instead of free groups. This model is the focus of Chapter 5.

Another exciting topic in mathematics today is *cubulation*, from which huge contributions to geometric group theory have arisen. A *cube complex* is formed by gluing Euclidean cubes together along faces, and it is *CAT(0)* if it is simply connected and non-positively curved. A group G is *cubulated* if there exists a CAT(0) cube complex X on which G acts cocompactly and properly by isometries. Cubulated groups have several nice properties; in particular, Wise showed that such groups are actually subgroups of right angled Artin groups (raAgs) and so inherit some of the properties of raAgs. These are described in Chapter 3. This subgroup property implies that cubulated groups are residually finite.

The goal of Chapter 4 is to show that random groups in the square model at density $d < 1/3$ are residually finite. One may want to cubulate such groups for a direct proof, but following the construction of Ollivier and Wise [**30**] which cubulates random groups in the density model at $d < 1/6$ as Odrzygóźdź does in [**27**] does not result in a proper action. So instead we focus on the action and use a workaround developed by Groves and Manning in [**13**] to conclude that the groups are still virtually special. The goal of Chapter 5 is to explore a new model of random groups, with a happy incidental result that random groups in the usual density model are almost surely perfect (equal to their commutator subgroup).

CHAPTER 2

Random Groups

1. History and Definitions

In this section we explore different models of random groups and explain the usefulness of each model. The motivating goal of this field is to say that "most" finitely presented groups, chosen "at random", will exhibit certain properties. The models make explicit definitions of "most" and "at random."

A first and intuitive definition of a random group is given by the *few-relators model*: we fix a number of generators m, a length l, and a number r of relators. Then we choose r of the $(2m)(2m-1)^{l-1}$ possible freely reduced words of length equal to l in m generators, assigning uniform probability to each word. We say that random groups in the few-relator models have property P with overwhelming probability (w.o.p.) if the proportion of groups constructed in such a fashion with P approaches 1 as l tends to ∞. Note: in the literature one will see both "with overwhelming probability" and "asymptotically almost surely," which technically cannot be used interchangeably, as probability theory says that w.o.p. is more likely that a.a.s., but in our document we use either as l tends to infinity throughout and renders the point moot. Gromov briefly alludes to the few-relators model in his 1987 paper on hyperbolic groups [**11**], in which he proves that w.o.p. groups in this model are hyperbolic. This intuitive model of random groups is subsumed by the *density model*, which we briefly described in the introduction and is the usual focus when studying random groups.

The main definition of random groups as random quotients of finitely generated free groups was introduced by Gromov in [**12**, §9.B]. Again we fix a number of generators m, choose a length l, and now add a density parameter d chosen such that $0 < d < 1$, getting rid of the number of relators r in the few-relators model. To construct a random group one chooses $\lfloor (2m-1)^{dl} \rfloor$ many words as relators from the same set as in the few-relators model. Again a random group in the density model has property P w.o.p. if the share of such groups with P approaches 1 as l tends to ∞. As a seminal result Gromov found that if $d > 1/2$, triviality or $\mathbb{Z}/2\mathbb{Z}$ occurs w.o.p, and if $d < 1/2$, hyperbolicity and torsion-freeness occur w.o.p.. Intuitively then, we can say that "most" finitely presented groups are hyperbolic. An excellent and thorough survey of random groups and early results in the field is available in Ollivier's book [**28**], and a more recent survey of current research in random groups is available in [**1**]. Also in Ollivier's book is an explanation of why we can switch between this "spherical" model of random groups to a "ball" model (also known as Arzhantseva-Olshanskii's *genericity* model), and choose relators of length less than or equal to l–we do not do so in this paper, but it is a useful technique in the literature. As a historical note, Kapovich and Schupp wrote a comparison of more early models of random groups in [**18**].

There are a number of limitations to the density model as it applies to our motivational goal. For instance, since the threshold is the same for hyperbolicity/torsion-freeness and for trivializing a group, we will never see any non-hyperbolic groups or groups with torsion appear in this model. Since $\mathbb{Z}^2$ can never be a subgroup of a hyperbolic group, we will never even have two commuting elements in a random group which are not powers of each other. Surely group theorists have some interest in abelian groups though, so this is a natural defect of the density model. We discuss this further in Section 3.

In order to prove a bound on Property (T) in the density model, Żuk introduced the *triangular model* of random groups [**40**], taking advantage of the fact that Property (T) passes through quotients. We need not worry about the full implications and definition of Property (T) in this paper, only that it implies that any action of the group on a finite-dimensional CAT(0) cube complex X must have a global fixed point, as proven in [**23**, Theorem B]. In particular, any such group action cannot be proper. In the triangular model one fixes all relators to length 3, and lets the number of generators tend to infinity rather than the length of the relators. By proving that random groups in this model have Property (T) when $d > 1/3$ and then passing to quotients, Żuk shows that random groups in the density model have Property (T) when $d > 1/3$. His original proof does not explicitly explain how to pass between models, and Kotowski and Kotowski fill this gap in [**19**].

The usefulness of bounding the length of relators as shown by the triangular model inspired Odrzygóźdź to define the *square model of random groups* [**27**, Definition 1.1]. Here all relators are fixed to length 4 and random groups have property P w.o.p. if the share of presentations with P tends to 1 as m approaches ∞. As in the density model and triangular model, Odrzygóźdź shows that random groups in the square model are almost surely trivial when $d > 1/2$ and almost surely hyperbolic when $d < 1/2$, and also proves bounds with Property (T) and freeness.

1.1. Van Kampen diagrams and Isoperimetric inequalities. The main method of studying random groups in any model is via their *isoperimetric inequalities*. To understand what these

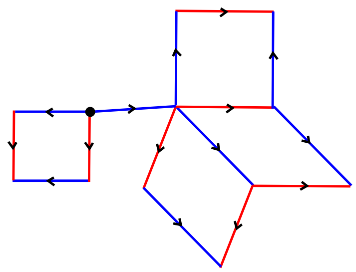

FIGURE 1. A Van Kampen diagram in $\mathbb{Z}^2$, where red edges are a and blue edges are b.

are, we first need to understand Van Kampen diagrams, which are visual means of understanding relations in a group. Given a group presentation $\langle S|R\rangle$, a *Van Kampen diagram* D for the group presentation is a planar 2-complex such that each edge is labeled by a letter in S, and the boundary of each 2-cell reads a freely reduced word in $\langle\langle R\rangle\rangle$, the normal closure of R. If D is a topological disk, then D is called a *disk diagram*. In this case, choosing a boundary vertex as a base point and reading the boundary word along ∂D gives a trivial word in the group presentation, formed as a product of conjugates of relators in R. See Figure 1 for an example of a Van Kampen diagram of $\mathbb{Z}^2 = \langle a, b|[a, b]\rangle$, where different generators are represented by different colored edges. Just as the boundary of a disk diagram reads off a trivial word in the presentation, any trivial word can be represented via a disk diagram by first writing it as a product of conjugates of elements in R. For example, Figure 1 may represent the element $r \cdot brb^{-1} \cdot brb^{-1} \cdot b^2rb^{-2} \cdot$, where $r = [a, b]$, the commutator of a and b–this can be seen by starting at the blue dot and reading each square clockwise, starting with the leftmost square as the initial r.

Isoperimetric inequalities offer bounds on the ratio between the length of the boundary of a Van Kampen diagram D and the number of 2-cells in D. When Gromov introduces random groups in [**12**], he includes a discussion of isoperimetric inequalities. In [**29**], Ollivier describes a sharp bound for the isoperimetric inequality of random groups.

THEOREM 2.1. [**29**, Theorem 2] *For random groups in the density model at $d < 1/2$, for any $\epsilon > 0$, w.o.p. any Van Kampen diagram D at length l satisfies $|\partial D| > (1 - 2d - \epsilon)l|D|$, where $|\partial D|$ indicates the length of the boundary of D and $|D|$ the number of 2-cells.*

Hyperbolicity of random groups follows directly from this isoperimetric inequality, as the definition of word-hyperbolicity for a group presentation is that any trivial element in the presentation satisfies a linear isoperimetric inequality. A full explanation of the proof of this theorem, which is beyond the scope of this document, is available in [**28**, §V].

2. The Square Model of Random Groups

As a reminder, in the square model, a random group at density d is a quotient of the free group on m letters $F_m/\langle\langle R\rangle\rangle$ such that relators in R are of length four with $\lfloor(2m-1)^{4d}\rfloor$ such relators chosen with uniform probability from all $2m(2m-1)^3$ such reduced 4-letter words.

We say a property P occurs *with overwhelming probability* in the square model at density d if the probability that a random group has P converges to 1 as $m \to \infty$. It is important to note here that unlike the density model, the relators do not lengthen, but the number of generators increases. Analogous to Gromov's result, Odrzygódźyź showed the same density threshold (1/2) for triviality and hyperbolicity in the square model [**27**, Theorem 2.8, Corollary 3.8].

One advantage of this model is that the Cayley complex Γ of a random group in the square model is a square complex. A *square complex* is a complex formed by gluing together Euclidean unit squares and unit intervals isometrically along vertices and edges. The first example of a square complex is the Euclidean plane. Square complexes are also seen as the 2-skeletons of cube complexes, which we explore in detail in Chapter 3. Square complexes come equipped with *hypergraphs*, graphs formed by connecting midpoints of opposite sides of squares in the complex with edges. More precisely, we build a graph whose vertices correspond to the 1-cells in the square complex, and add an edge between two vertices if there exists a 2-cell R such that the corresponding 1-cells in the square complex are opposite each other in R. A hypergraph h of Γ is a connected component of this newly formed graph. We say an edge $e \in \Gamma$ is *dual* to a hypergraph h if h passes through e. The *carrier* of a hypergraph h is the connected collection of 2-cells in Γ through which h passes. A *ladder* L of a hypergraph h is a segment of the carrier of h. Note that the term "hypergraph" means something different in graph theory; the two definitions have nothing to do with each other, but we use "hypergraph" as Ollivier and Wise do in [**30**]. Hypergraphs in the square model of random groups are embedded trees w.o.p. when $d < 1/3$ as proven in [**27**, Lemma 5.4], unlike the example in Figure 2.

2.1. Isoperimetric inequalities.

When Odrzygóźdź defines the square model of random groups in [**27**], he proves an analogous version of Ollivier's isoperimetric inequality. However, he improves this isoperimetric inequality in [**26**, Theorem 1.2] by extending to nonplanar diagrams. Nonplanar diagrams can be thought of as natural extensions of Van Kampen diagrams of group presentations $\langle S|R\rangle$, so they are still 2-complexes with each edge labeled by a letter in S, and any closed loop in the 1-skeleton reads a freely reduced word in $\langle\langle R\rangle\rangle$; we remove the requirement that the diagram be planar. Given a random group G defined by relators R with length l, we consider

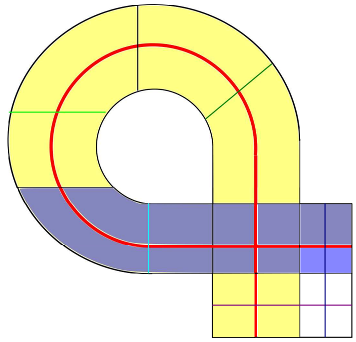

FIGURE 2. An example of a hypergraph (red) of a square complex. A ladder is indicated in purple; the rest of the carrier of the hypergraph is in yellow.

2-complexes Y which have a combinatorial map from Y to the presentation complex of G which is locally injective around edges, but not necessarily around vertices.

THEOREM 2.2. [**26**, Theorem 1.2] *At density d, for any $\epsilon > 0$ and for any 2-complex Y as described above, almost surely*

$$\sum_{e \in Y^{(1)}} (\deg(e) - 1) \leq (d + \epsilon)|Y|\ell$$

Odrzygóźdź defines $\mathrm{Cancel}(e) = \deg(e) - 1$ for any edge e in Y, and $\mathrm{Cancel}(Y) = \sum_{e \in Y^{(1)}} (\deg(e) - 1)$, and we use the Cancel notation throughout this paper. Intuitively, Cancel tells us how many times each edge is "double counted" when we consider $\ell \cdot |Y|$. For instance, any external edge adds 0 to Cancel, and internal edges contribute how many 2-cells are glued to that edge, less one. In the square model we let $\ell = 4$ and use this isoperimetric inequality. If Y is planar, then $\mathrm{Cancel}(e) = 1$ for internal edges and $\mathrm{Cancel}(e) = 0$ for boundary edges, so the nonplanar isoperimetric inequality simplifies to $1/2(4|D| - |\partial D|) \leq (d + \epsilon)4|D|$, or $|\partial D| > (1 - 2d - 2\epsilon)4|D|$, which is the same as Theorem 2.1 up to a change in choice of ϵ.

In Chapter 4 we extensively make use of Theorem 2.2 with $d = 1/3$.

3. Random Nilpotent Groups

As mentioned in Section 1, our definition of the density model of random groups completely neglects certain subclasses of groups, such as abelian groups. If one wants to study random abelian groups, one can replace the free group F_m that we quotient by a free abelian group $\mathbb{Z}^m$, and treat the random strings that form R as elements of $\mathbb{Z}_m$ instead. This is one model of random abelian groups, and there are many other ways to study them. Random abelian groups form a fairly well-studied area and results in it are beyond the scope of this thesis; for a useful survey of the literature, see [**22**] and its references.

One area that has *not* been studied before the paper from which Chapter 5 is derived, however, is random nilpotent groups, which are a natural extension of random abelian groups. Here we begin with a *free nilpotent group* $N_{s,m}$ *of step* s *and rank* m and add random relators as we do in other models. Rank m means that there are m generators, and step s means that all nested commutators of depth s are trivial: $[[\cdots[a_0, a_1], a_2], \dots, a_s] = 1$ for all elements a_i in the group, but not all nested commutators of depth $s-1$ are trivial. For example, abelian groups are 1–step nilpotent, so some results in random nilpotent groups can be applied to random abelian groups. As is common in nilpotent group literature, we write $[a_0, a_1, a_2] = [[a_0, a_1], a_2]$ for nested commutators. So we have an explicit presentation for $N_{s,m} = F_m / \langle\langle \{[a_{i_0}, \dots, a_{i_s}]\} : 1 \leq i_0, \dots, i_s \leq m\} \rangle\rangle = \langle a_1 \dots, a_m | [a_{i_0}, \dots, a_{i_s}] \text{ for all } i_j \rangle$, where $\langle\langle R \rangle\rangle$ indicates the normal closure of R in the ambient group.

Just as all finitely-generated groups are quotients of (finite-rank) free groups, all finitely-generated nilpotent groups are quotients of free nilpotent groups.

In Chapter 5 we begin to study the typical properties of random nilpotent groups. For instance, one would expect that the threshold for trivialization occurs with far fewer relators than for free groups, and also that nontrivial abelian quotients should occur with positive probability at some range of relator growth. In fact, rather than base the random nilpotent group model on the density model, we base it on the few-relators model because one of our results is that random nilpotent groups are a.a.s trivial whenever $|R|$ is unbounded as a function of l. Thus throughout Chapter 5, we explicitly define $|R|$, often as a function of m, the number of generators. Recall that Gromov's triviality threshold in the random density model implies that groups at $d > 1/2$ are either trivial or $\mathbb{Z}/2\mathbb{Z}$ due to a small parity issue. In Chapter 5, we choose our relators from those of length ℓ and $\ell - 1$ with equal probability in order to avoid the parity issue; with this convention, a random nilpotent group is trivial a.a.s. for $d > 1/2$.

Since the paper that Chapter 5 comes from has been published, a number of further studies have been made in this model, and more models have been developed. For instance, one of our results involves whether a random nilpotent group is abelian, based on if $|R|$ is larger than, equal to, or smaller than m. In [**9**], Garreta, Miasnikov, and Ovchinnikov expand on this result by further determining the structure of random nilpotent groups if $|R| \leq m-2, |R| = m-1, |R| = m$, and $|R| \geq m+1$. From this they extend to the decidability of Diophantine problems over random nilpotent groups, which is a discussion beyond this document. They continue their research by introducing new models of random nilpotent groups in [**10**], where they take advantage of the fact that finitely generated nilpotent groups are polycyclic. They develop a *polycyclic model of random nilpotent groups*,where presentations are written polycyclically; that is, every relator is a conjugate of an element in the form $x_i^{r_i} \prod_{k=i+1}^{n} a_k^{m_{i,k}}$, where the a_i are generators and x_i is a group element. We go through this in Section 5.1.1 when defining a *Mal'cev basis*, which these authors use extensively. Interestingly these authors restrict to 2-step random nilpotent groups, and explain that any higher step random nilpotent groups end up a.a.s. finite under their model.

Simultaneously developed alongside the paper from which Chapter 5 is derived, Delp, Dymarz, and Schaffer-Cohen explored yet another model of random nilpotent groups, viewing torsion-free random nilpotent groups as subgroups of $U_n(\mathbb{Z})$, the group of upper triangular matrices with integral entries and unit diagonal entries [**5**]. Rather than take random quotients, they randomly generate torsion free nilpotent groups with two generators of length l, and their results are based on how l increases as a function of n.

For more on random groups and the current state of research in the field, see the recent survey by Bassino, Nicaud, and Weil [**1**]. The survey extends beyond just random nilpotent groups and is meant to encompass all current research in random groups and random subgroups; Section 2.5 summarizes the results in random nilpotent groups mentioned above.

CHAPTER 3

Cubulation

1. History and Definitions

To understand free groups, Stallings developed the idea of *folding* graphs and adding edges to them without changing fundamental group, and analyzing the results in a seminal and beautiful paper [**35**]. He defines a canonical way to add edges to the graphs to prove Marshall Hall's theorem. Analogously, Wise developed a canonical completion of *special cube complexes* to understand subgroups of right angled Artin groups.

Right angled Artin groups (raAgs) are a particularly nice type of groups. Given a finite simplicial graph Γ, one forms A_Γ by assigning a generator v_i to each vertex of $V(\Gamma)$, and a commutator relator $[v_1, v_2]$ for each edge $(v_1, v_2) \in E(\Gamma)$. For example, the empty graph with n vertices creates F_n, the free group on n generators. On the other hand, the complete graph K_n creates $\mathbb{Z}_n$. For a more thorough and enjoyable introduction to right angled Artin groups, see Charney's survey paper [**3**]. RaAgs are linear and hence residually finite, which is explained in two different ways in [**3**] (one method shows that Artin groups are subgroups of Coxeter groups). Linearity is inherited by subgroups, so if one can show that a group is a subgroup of a raAg, one can show that such a group is residually finite. We do this via the machinery of cube complexes. For now we set aside raAgs and will return to them in Section 3.3.

A *cube complex* is a complex formed by gluing Euclidean unit cubes together isometrically along cells of the same dimension. One could glue two squares together along all four of their edges to form a sphere, but we do not want any such positive curvature. A cube complex is called *CAT(0)* if it is non-positively curved and simply connected. Gromov showed that one can determine non-positive curvature using a local condition, called *Gromov's link condition.* For any vertex in a cube complex, the *link* of that vertex is the complex which arises as the intersection of the cube complex with a sphere of radius 1/3 centered at the vertex. The link condition says that if every link is simplicial, and if every k−skeleton of a simplex is filled by a k−simplex in the link of every vertex in the cube complex, then the entire complex is non-positively curved. Intuitively this condition implies a "no empty triangles" rule for each link, which means if three 2-cells can form a corner of a 3-cube in the cube complex, then that 3-cube also exists. Otherwise, we would have a point of positive curvature.

We use cube complexes because they come equipped with *hyperplanes*, which we can build using *midcubes*. For any 1-simplex $[0,1]$ in the cube complex, one marks the midpoint $1/2$ in the interval. If the maximal dimension of cubes which contain the interval is 1, then the midcube is just the midpoint of the interval. If it is contained in a square, the midcube is a 1-simplex connecting midpoints on opposite sides of the square. Note any square contains two midcubes, and any d–dimensional cube contains d midcubes. One continues this process to find midcubes of every cube. If any midcubes share lower dimensional subcubes, then we connect them along these lower dimensional subcubes to form the hyperplanes. Note that in a square complex, hyperplanes are exactly the hypergraphs we defined in Section 2.2. Examples of hyperplanes are in Figure 3.

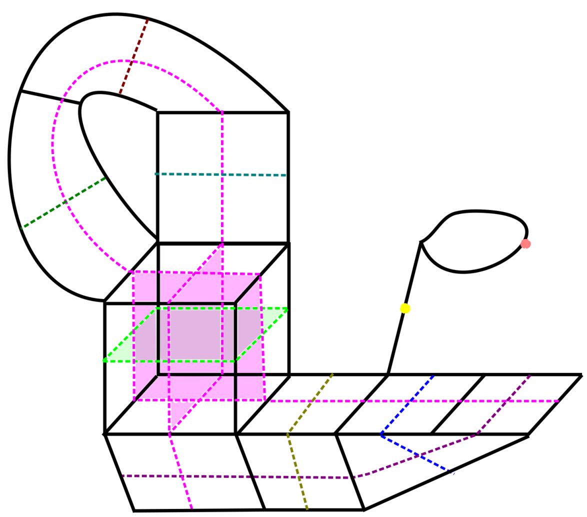

FIGURE 3. A non-$CAT(0)$ cube complex with its hyperplanes

Hyperplanes of $CAT(0)$ cube complexes have a number of nice properties: hyperplanes are themselves $CAT(0)$ cube complexes, and they are two-sided (so if $\hat{h}$ is a hyperplane, then $X - N(\hat{h})$ has two components). Here, $N(\hat{h})$ is the cubical neighborhood of $\hat{h}$, and is called the *carrier* of $\hat{h}$–as X is $CAT(0)$, $N(\hat{h}) \cong \hat{h} \times [-1,1]$. These properties are proven by Sageev in [**34**].

A group G is *cubulated* if G acts on a CAT(0) cube complex properly and cocompactly by cubical isometries. Warning: these properties are not consistent in the literature. In this paper we require our action to be proper, cocompact, and by (cubical) isometries, but in the literature this is sometimes referred to as *cocompactly cubulated* [for instance, in [**17**]]. By "cubical" we mean that n-cubes are sent to n-cubes and containment of cubes is preserved. RaAgs are cubulated, as each raAg defines a canonical *Salvetti complex* on which it acts, which is a cube complex with a CAT(0) universal cover. Briefly, one builds the Salvetti complex of a raAg by starting with one vertex, adding an edge for each generator of the raAg in the presentation which comes from the corresponding graph Γ, and adding n-tori for each n-cycle in Γ. Note that the fundamental group of a raAg's Salvetti complex is the raAg itself; the 2-skeleton of the Salvetti complex is the usual presentation 2-complex of the raAg. A major goal in current geometric group theory research is to cubulate groups.

2. Sageev's Wall Construction

In this section we carefully go through an example of Sageev's wall construction. In his thesis [**34**], Sageev described a way to construct a CAT(0) cube complex X out of a *space with walls.* A space W with walls is a set with a partially ordered set of subsets (called *halfspaces*) which is closed under complementation, so for any halfspace $A \subset W$, $A^C \subset W$ is also a halfspace, and if $A \subset B$, then $A^C \supset B^C$, such that any pair of points in W are separated by at most finitely many pairs of halfspaces and their complements (walls). Inclusion is the partial order, so if $A \subset B$ we have $A < B$. Given a locally finite space with walls W, we want to form the vertices of X using *ultrafilters.* An ultrafilter is a choice of a halfspace for each wall in W (so we pick either A or A^C for every halfspace in W) such that if we choose A and $A < B$, we also choose B: this is called the *consistency condition.*

To form the vertices of X, we choose every ultrafilter which satisfies a descending chain condition (so every descending chain $A_1 > A_2 > \dots$ terminates). Then we build an edge between two vertices if the corresponding ultrafilters differ in choice of exactly one halfspace. Note a square arises if we have four ultrafilters which agree in all but two halfspaces A and B, and they choose all four possible combinations of A, A^C, B, B^C. If such a square occurs in X, we fill it in with a 2-cell. In general, if an $n-1$-skeleton of an $n-$cube appears in X, we fill it with an n cube. The resulting cube complex is CAT(0) by construction.

EXAMPLE 3.1. *Consider the Cayley complex of $\mathbb{Z}^2$ arising from the presentation $\langle a, b, c \mid [a, b], ca^{-1}b^{-1}\rangle$ as our space with walls, where the bi-infinite geodesics formed by powers of generators are the walls (so, every vertical, horizontal, and diagonal line is a wall). We show that the CAT(0) cube complex from this presentation is $\mathbb{R}^3$ equipped with the Cayley complex of $\mathbb{Z}^3 = \langle a, b, c \mid [a, b], [b, c], [a, c]\rangle$ as its 2-skeleton.*

PROOF. For every triangular shaped region T in the Cayley complex, we can choose all halfspaces that contain it to form an ultrafilter that satisfies the descending chain condition. Note this is indeed an ultrafilter, for if $T \subset A$ and $A < B$, we know $T \subset B$ as well so we also choose B. Sageev defines ultrafilters that arise like this *principal ultrafilters*.

We can also form ultrafilters "at infinity" by picking an infinite geodesic ray from the origin and choosing halfspaces which contain infinitely much of the ray. Note these ultrafilters do not satisfy the descending chain condition, so we do not use them when constructing the cube complex.

Have we found all the ultrafilters? Any ultrafilter which does not satisfy the descending chain condition must be "at infinity" in some direction. Consider a vertex in the Cayley graph defined by three walls which pairwise intersect. Since each intersects the others, these three walls form incomparable halfspaces in the poset. We can cross any of them to find a different ultrafilter. In particular we have 2^3 ultrafilters involving these three walls, but only six principal ultrafilters which do so. We call the remaining two ultrafilters *nonprincipal ultrafilters*.

Consider one of the walls parallel to the diagonal wall. Since the halfspaces defined by this wall are incomparable in the poset to the four horizontal/vertical halfspaces, each halfspace of this diagonal wall can be chosen along with the four choices of horizontal/vertical halfspaces, and we choose halfspaces of all other walls to be consistent with this choice. Infinitely many of these choices result in principal ultrafilters, and infinitely many result in nonprincipal ultrafilters. So we have many nonprincipal ultrafilters.

Now let us find the CAT(0) cube complex formed by this space with walls. Every vertex is defined by just six halfspaces: three pairs of halfspaces formed by parallel walls which point "toward" each other. All other choices of halfspaces are determined by these six and the consistency condition. Switching over any one of the six walls leads to another vertex which still satisfies the consistency condition, so every vertex has six neighbors. Choosing any two generators (say, a and b) and moving along parallel walls, we form a copy of the Cayley graph of $\mathbb{Z}^2$ containing our vertex, where each vertex in this copy of $\mathbb{Z}^2$ has the same orientations in the c direction as our vertex. Swapping a single c orientation, we can do the same construction and form a parallel copy of the Cayley graph of $\mathbb{Z}^2$. We can do this infinitely many times, and so end up with a copy of the Cayley graph of $\mathbb{Z}^3$. Filling in the cubes gives the CAT(0) cube complex of $\mathbb{R}^3$ with its 2-skeleton formed by the Cayley complex of $\mathbb{Z}^3$.

□

Recall that a group is *cubulated* if it acts on a CAT(0) cube complex with certain properties. If a group G acts on a space with walls with no g sending a halfspace to its complementary halfspace,

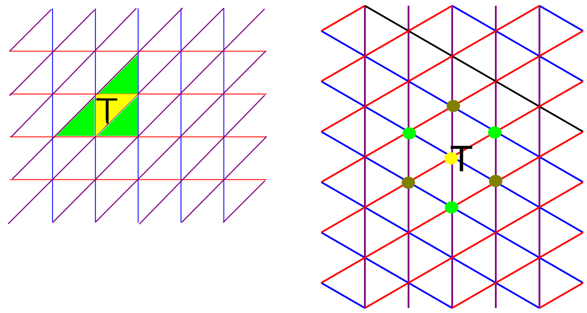

FIGURE 4. Left, the Cayley graph of $\mathbb{Z}^2$, and right, part of the ensuing $CAT(0)$ cube complex. Green marks principal vertices and corresponding regions; olive marks nonprincipal vertices.

then this naturally creates an action on the 0-skeleton of Sageev's CAT(0) cube complex as proven by Sageev in [**34**, §3]. As any $n-$cube in the cube complex is determined by its vertices and the action is by isometries, the action extends to each $n-$skeleton and so to the entire cube complex. Thus Sageev's construction gives an approach to cubulating groups.

For instance, Wise proves that small cancellation groups are cubulated in [**37**] by forming walls within the Cayley complexes of the small cancellation groups. In particular, he forms *hypergraphs* by connecting midpoints of edges of opposite sides of even-length relator polygons and gluing together edges that arise and share such midpoints (along with a small fix for odd-length relator polygons). From these hypergraphs he uses Sageev's construction to form a CAT(0) complex on which the small cancellation group acts. This approach is the basis of what we do in Chapter 4. Similarly, Ollivier and Wise cubulate random groups at density $d < 1/6$ using these same techniques in [**30**].

3. Special Cube Complexes

This section very briefly summarizes relevant results from the sprawling and active field of special cube complexes; for a more detailed survey with many examples see Wise's book [**38**]. Wise described a category of cube complexes, called *special cube complexes*, with a list of conditions on the hyperplanes of those cube complexes. Hyperplanes cannot self-intersect, must be two-sided (so there are no Möbius bands inside the cube complex), cannot *directly self-osculate*, and cannot *inter-osculate*. Two hyperplanes h and l *osculate* if there exists a vertex v adjacent to two edges e_1, e_2 which do not border a square and which are dual to h and l, respectively. If $h = l$ and we give h an orientation, h directly self-osculates if e_1 and e_2 cross h with opposite orientations (v is the terminus of both). If $h \neq l$ and they osculate and intersect, we say h and l inter-osculate. We illustrate these forbidden configurations in Figure 5. For more on this, see Haglund and Wise's paper [**14**].

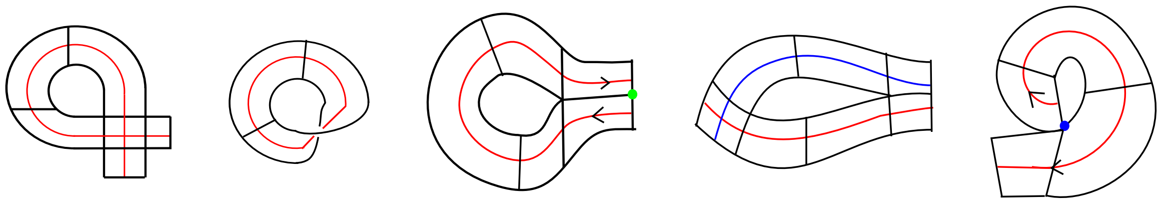

FIGURE 5. The four forbidden configurations for specialness. Fifth is indirect self-osculation and is allowed.

EXAMPLE 3.2. *Examples of special cube complexes.*

- Sageev proves all $CAT(0)$ cube complexes satisfy each condition for specialness in his thesis, though without the language of "special" [**34**].
- The Salvetti complex described earlier which arises from a raAg is special.
- Graphs are special, as their hypergraphs are just midpoints of edges.
- Any subcomplex of a product of graphs is also special. Any hypergraph in a product of graphs (which is a square complex) is a copy of one of the original graphs, and so does not self-intersect. The product structure guarantees that hypergraphs are 2-sided, as the carrier of any hypergraph is just a copy of one of the original graphs crossed with an interval. Any two intersecting hypergraphs are copies of the original graphs and diverge after meeting, so there is no interosculation. By this same reasoning there is no direct self-osculation.

A group is *special* if it is the fundamental group of a special cube complex. Immediately, raAgs are special. The major breakthrough by Haglund and Wise in [**14**, Theorem 4.2] is that special groups are subgroups of raAgs, and so inherit many nice properties of raAgs. In particular, special groups are residually finite. So we want to show that our groups of interest are special, or at least virtually special (that is, has a finite index subgroup which is itself special).

If we find a $CAT(0)$ cube complex X on which a group G acts properly and cocompactly by cubical isometries such that the quotient is special, then G is special. However, even if one has made a $CAT(0)$ cube complex on which a group G acts, it can be difficult to show that the action is proper and cocompact. Instead of requiring properness, a result by Groves and Manning gives us what we need. A subgroup $H < G$ is *quasiconvex* if there exists a K such that every geodesic in the Cayley graph of G with endpoints lying in H lies within a $K-$neighborhood of H, and is *convex* if $K = 0$. A trivial example of a non-convex but quasiconvex subgroup is $2\mathbb{Z}$ in $\mathbb{Z}$.

THEOREM 3.3. [**13**, Theorem D] *If G is a hyperbolic group acts cocompactly on a $CAT(0)$ cube complex with virtually special and quasiconvex cell stabilizers, then G is virtually special.*

A corollary of this theorem states that to show that G is virtually special, we need only prove vertex stabilizers are virtually special if the $CAT(0)$ cube complex in question arises via Sageev's construction using a collection of quasiconvex codimension 1 subgroups as the walls.

There is a major theorem by Haglund and Wise which is unrelated to this document but very worthy of mention. The theorem [**14**, Theorem 1.3] states that quasiconvex subgroups of hyperbolic virtually special groups are *separable*: for every element x not in the subgroup, there is a homomorphism f to a finite group such that $f(x) \notin f(H)$. And in fact, Haglund and Wise prove the converse as well: if all quasiconvex subgroups of a hyperbolic group are separable, then that group is virtually special. A group is residually finite if the trivial subgroup is separable.

CHAPTER 4

Cubulating Random Groups in the Square Model

1. Introduction

One advantage of the square model of random groups, as introduced in Chapter 2, is that the Cayley complex Γ of a random group in the square model is a square complex. Square complexes comes equipped with *hypergraphs*, graphs formed by connecting midpoints of opposite sides of squares in the complex with edges, as explained in Chapter 3. Note that the term "hypergraph" means something different in graph theory; the two definitions have nothing to do with each other. Lemma 5.4 in [**27**] states that these hypergraphs are embedded trees w.o.p. when $d < 1/3$.

The embeddedness of the hypergraphs suggests using the tools of special cube complex theory to explore random groups in the square model. There is some ambiguity in the literature between the terms "cubulated", "cocompactly cubulated", and "partially cubulated"; in this thesis we say a group G is *cubulated* if there exists a CAT(0) cube complex X on which G acts freely and cocompactly by isometries. Sageev's wall construction as explained in Section 3.2 offers a way to build a dual cube complex X on which a group G acts using codimension one subgroups; Haglund applied this construction to Cayley complexes of small cancellation groups in [**39**]. Throughout the chapter, we use Γ to denote the Cayley complex of our random group G, and X to denote the dual CAT(0) cube complex.

Some vertices of $X^{(0)}$ correspond to non-empty regions in the Cayley complex Γ: we call those *principal vertices*, in the same vein as our *principal ultrafilters* from Section 3.2. If a vertex is not principal, it is called *nonprincipal*. We use different techniques to address the stabilizers of each type. To cubulate groups, one shows that all of these stabilizers are trivial, as in [**30**] or [**39**] for instance. However, in the square model at $d > 1/4$ we do not have trivial stabilizers, so we must use different techniques to show:

THEOREM 4.1. *In the square model at $d < 1/3$, random groups act cocompactly on a CAT(0) cube complex with finitely generated free vertex stabilizers.*

By recent work of Groves-Manning [**13**], we immediately get the corollary:

COROLLARY 4.1. *In the square model at $d < 1/3$, random groups are virtually special and, in particular, are residually finite.*

PROOF. This follows immediately from Corollary D in [**13**], as finitely generated free groups are virtually special. □

While this manuscript was being prepared, Odrzygódźyź proved that random groups in the square model at $d < 3/10$ are virtually special [**25**] by modifying Sageev's wall construction and showing that the resulting action is proper. However, he uses completely different techniques from the ones used here.

2. Background

First, we remind ourselves of some important inequalities from Section 3.2, by Ollivier and Odrzygóźdź, respectively:

THEOREM 2.1. [**29**, Theorem 2] *At density d, for any $\epsilon > 0$ the following property occurs with overwhelming probability: every reduced Van Kampen diagram Δ satisfies*

$$|\partial\Delta| \geq (1 - 2d - \epsilon)\ell|\Delta|$$

.

THEOREM 2.2. [**26**, Theorem 1.2] *At density d, for any $\epsilon > 0$ and for any 2-complex Y, almost surely*

$$\sum_{e\in Y^{(1)}} (\deg(e) - 1) \leq (d + \epsilon)|Y|\ell$$

We choose $\epsilon < 1/3 - d$ and end up with an inequality we refer to throughout this chapter:

$$(*) \qquad \text{Cancel}(Y)/|Y| < 4/3$$

As with planar isoperimetric inequalities, $*$ bounds the ratio of internal gluing to the boundary of the diagram, which can be seen by considering $|Y| = \frac{|E(Y)|}{4}$, where we count the number of edges of Y. That is, the inequality says that less than a third of edges in Y are "double counted". In the planar case with Ollivier's inequality, we have $|\partial\Delta| > 4/3|\Delta|$, since $1/3 - d < 2(1/3 - d)$. Hence, as Ollivier and Wise do, we usually use the strict inequality. We say that a diagram is *valid* if all of its subdiagrams satisfies these inequalities. We make an observation about Cancel which follows directly from the definition:

REMARK 4.2. *If A, B are 2-complexes, then* $\mathrm{Cancel}(A \cup B) = \mathrm{Cancel}(A) + \mathrm{Cancel}(B) - \mathrm{Cancel}(A \cap B)$.

The following lemma is part of the proof of Theorem 5.14 in [**27**]. We include it because we refer to the diagrams in Figure 6 throughout the paper as diagrams (a), (b), and (c). This is identical to Figure 4 in [**27**].

LEMMA 4.3. *In the square model at $d < 1/3$, any valid diagram containing two hypergraphs which intersect twice must contain one of the three diagrams in Figure 6.*

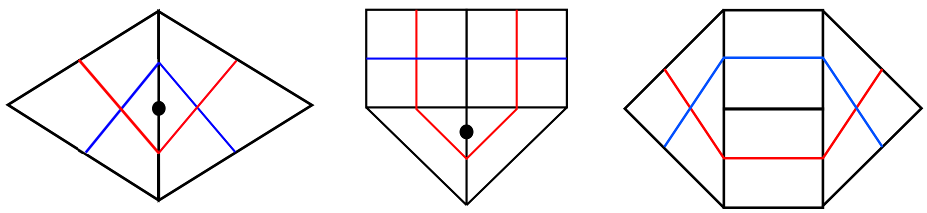

FIGURE 6. Diagrams (a), (b), and (c)

Figure 6 arises in [**27**, Theorem 5.14] as an application of [**30**, Theorem 3.12]. A key point is that hypergraphs in the square model are embedded trees w.o.p. when $d < 1/3$ [**27**].

A diagram Δ is *collared by two hypergraphs* if it is a disk diagram where every boundary cell of Δ contains one of the hypergraphs, neither hypergraph intersects an internal cell of Δ, and there exist two *corner* cells, which are boundary cells that contain intersections of the two hypergraphs. Theorem 3.12 in [**30**] states if two hypergraphs which are embedded trees intersect more than once, they form a reduced diagram collared by the two hypergraphs, and we can choose one of the corners to be one of our intersection points. This theorem allows us to choose just one of the corner cells, and we extend this in Lemma 4.4 to say we can in fact choose both.

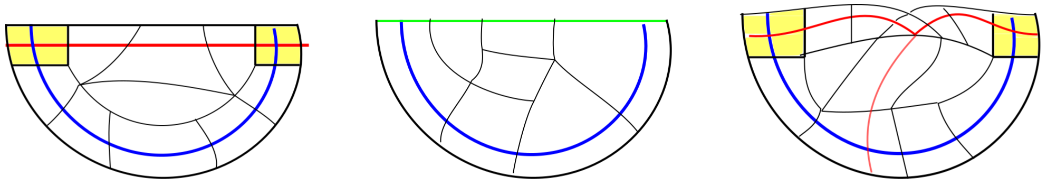

FIGURE 7. A diagram collared by two hypergraphs, one collared by a hypergraph and a path, and one quasi-collared by two hypergraphs, with corners indicated.

We also define *diagrams collared by hypergraphs and paths* as Ollivier and Wise do [**30**, Definition 3.11]. First, the *carrier* of a hypergraph h is the connected collection of 2-cells through which h

passes. A *ladder* L of a hypergraph h is a segment of the carrier of h. We say a disc diagram Δ is *collared* by a hypergraph h and a path p if the following conditions hold on a ladder L of h:

- There is a path l along L joining a point in the boundary of the first 2-cell of L with a point in the boundary of the last 2-cell of L such that $\partial\Delta = l \cup p$, so that l does not cross h.
- The carrier of h contains no internal 2-cells of Δ.

An example of a diagram collared by a hypergraph and a path is in the middle of Figure 7. A diagram Y is *quasi-collared* by two hypergraphs h and l if it is constructed in the following way. Ladders L_h and L_l of h and l intersect in two corner cells, and there exist paths p_h and p_l through the respective ladders which connect to a point in each of the corner cells. Then $p_h \cup p_l$ forms a cycle, and there exists a disk A with boundary $p_h \cup p_l$. As in the rightmost diagram in Figure 7, Δ is quasi-collared by h and l if it is equal to A glued to $L_h \cup L_l$ along $p_h \cup p_l$. Note if $p_h \cup p_l$ never crosses either ladder, Y is actually collared by h and l and Y is a disc diagram. Otherwise Y is not a topological disc, as one can see in Figure 7.

LEMMA 4.4. *Suppose h and l are two hypergraphs in a Cayley complex Γ which are embedded trees that intersect twice at points a and b. If h does not intersect l between points a and b and vice-versa, then h and l form a reduced diagram Δ collared by segments of h and l, and moreover the corners of this diagram include points a and b.*

PROOF. By Lemma 3.13 in [**30**], there exists a quasi-collared diagram Δ with a and b in its corners. By construction, within the diagram h and l intersect only in the corners of Δ. Since they are embedded trees, h and l cannot intersect themselves. So neither h nor l can appear inside an internal cell of Δ, so Δ is collared. By Lemma 3.15 in [**30**], reductions do not change the corners of Δ, so Δ remains collared after reduction. □

3. The Cayley Complex

The following lemma is used throughout the paper to say we cannot glue two of the allowable two-collared diagrams of Figure 6 together to form a larger diagram. We mention a few useful facts.

REMARK 4.5. *If the number of edges glued to the original diagram remains the same, the ratio $\frac{Cancel(\Delta)}{|\Delta|}$ is independent of how we glue additional 2-cells to Δ; that is, to which edges we choose to glue the 2-cells or which edges of the 2-cells we choose to glue.*

If a diagram Δ *fails* $(*)$*, then gluing two adjacent 2-cells to two edges of* Δ *results in a diagram that also fails the inequality, for the addition adds 2 to* $|\Delta|$ *and 2 to Cancel*(Δ)*.*

LEMMA 4.6. *There is no valid diagram* Y *which contains two distinct copies of diagram* (c) *from Figure 6 that share at least one 2-cell.*

PROOF. Suppose we have one copy of diagram (c) with hypergraphs h and l intersecting twice and a third hypergraph k, and we want to glue on a second copy to form a complex Y. Note Cancel$(c) = 5$ and $|c| = 4$. For all of the following cases refer to Figure 8.

(1) Suppose the two copies of (c) share one 2-cell, so Cancel of the overlap is 0 as there are no internal edges. By Remark 4.2, Cancel$(Y) = 5 + 5 - 0$, and $|Y| = 7$ as we share one two cell. This fails $(*)$.
(2) Suppose the two copies share two 2-cells. If the 2-cells are adjacent, then Cancel of the overlap is 1, and Cancel$(Y) = 5 + 5 - 1 = 9$, while $|Y| = 4 + 4 - 2 = 6$, failing the inequality. If the 2-cells are not adjacent, there is only one pair of possible 2-cells in both copies to glue together: the ones that include the two intersections of h and l. Such a gluing forms two loops in Y: one h loop and one l loop. As h and l are embedded trees, such loops cannot exist, so this cannot occur.
(3) Finally suppose the two copies share three 2-cells. If the excluded 2-cell is one of the middle cells, then Cancel of the overlap is 2 so Cancel$(Y) = 8$ while $|Y| = 5$. If the excluded 2-cell is an outside cell, then Cancel of the overlap is 3 so Cancel$(Y) = 5 + 5 - 3 = 7$ while $|Y| = 5$. Both fail the inequality.

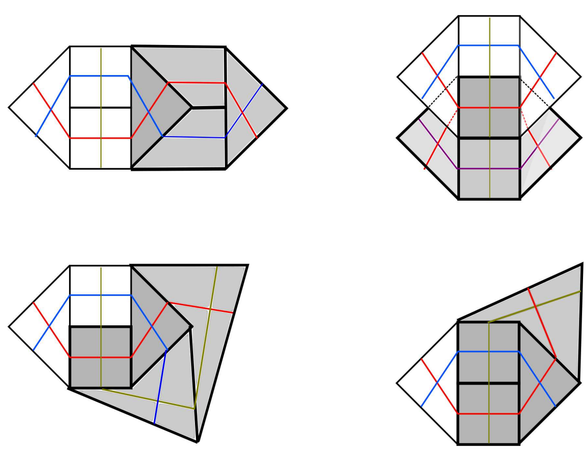

FIGURE 8. Top: sharing one dark gray cell, as in case (1). Bottom left: sharing two 2-cells as in case (2). Bottom right: Sharing three 2-cells, as in case (3).

□

LEMMA 4.7. *There is no valid planar diagram which glues a copy Δ of diagram (b) from Figure 6 to another diagram Δ' from Figure 6 along a vertex or an edge, such that the long hypergraph from Δ continues in Δ'.*

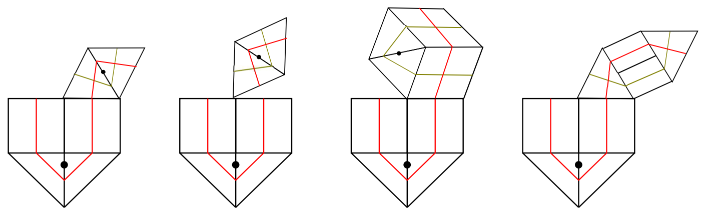

FIGURE 9. Left two: Case 1 in Lemma 4.7. Right: Cases 2 and 3. The red hypergraph is h.

PROOF. As in Remark 4.5, without loss of generality we select a particular edge in Δ, our copy of diagram (b), along which to glue the second diagram Δ'. As in Figure 9, this forms a diagram we call Y. Let h denote the long hypergraph from Δ.

(1) If Δ' is a diagram (a) and shares an edge with Δ, then Cancel of the overlap is -1 and we have $\text{Cancel}(Y) = 8$ and $|Y| = 6$, failing the inequality. If Δ' shares a vertex with Δ, the vertex must be on one of the boundary edges dual to h. In order for h to be dual to two edges which share a vertex in our diagram, it must pass through at least two 2-cells between these edges, forming a diagram (a). A diagram with two (a)s and a (b) sharing only edges has $\text{Cancel}(Y) = 11$ and $|Y| = 8$, which is not valid. Alternatively, h may pass through four 2-cells between Δ and Δ' and form a copy of diagram (b), which also fails the isoperimetric inequality.
(2) If Δ' is a diagram (b) and shares an edge with D, the overlap contributes 1 to Cancel and so $\text{Cancel}(Y)/|Y| = 11/8$, which is not valid. Then assume they share exactly one vertex. As above h must form a diagram (a) or (b) between Δ and Δ', and as above this fails the isoperimetric inequality.
(3) If Δ' is a diagram (c) and shares an edge with Δ, we have the same failure as in the previous case. If it shares a vertex then as in previous cases we fail the isoperimetric inequality.

□

4. The Cube Complex

We build a cube complex X from a set of hypergraphs of the Cayley complex Γ as we did in Example 3.1 in Chapter 3. Hypergraphs are 2-sided by Theorem 4.5 in [**27**], so we can pick an orientation for each hypergraph in the complex- that is, if h is a hypergraph, $\Gamma \setminus h$ has two components and we can choose one of these *halfspaces* to call h^+. Each vertex in $X^{(0)}$ is a choice of orientation for every hypergraph in Γ subject to the ultrafilter condition and the descending chain condition as described in Section 3.2.

Note that each vertex in Γ defines a *principal* vertex in X: one chooses all the halfspaces that contain that vertex. Multiple vertices in Γ may correspond to the same principal vertex in X. Vertices in X which do not correspond to a region in Γ are called *nonprincipal* vertices. If v is a vertex in Γ, we let x_v denote the principal vertex in X that corresponds to it. We say n-cells in X or Γ are *adjacent* if they share a common $(n-1)$ cell. We use the hypergraphs in the Cayley complex Γ as walls in Sageev's language. Figure 10 give an example of how a 3-cube appears in X as a result of a hypergraph arrangement in Γ.

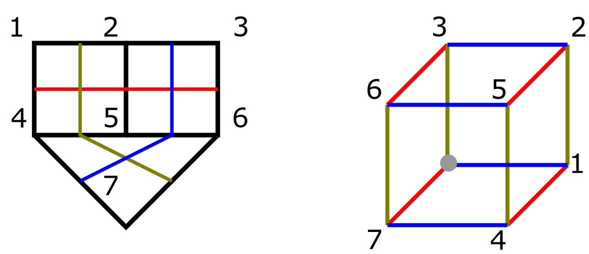

FIGURE 10. A portion of Γ and the corresponding 3-cube in X. Gray dot is the nonprincipal vertex. Edges in X correspond to the hypergraph with changing orientation.

LEMMA 4.8. *(3-cubes) Any 3-cube in X arises from one of the three diagrams in Figure 12 appearing in Γ.*

PROOF. A 3-cube C in X arises from three pairwise intersecting hypergraphs in Γ. These may form seven regions in Γ instead of the eight corners of the 3-cube. See Figure 10. They may also form eight regions, as in diagram (c) of Figure 4.3, but we focus on minimal diagrams which contain the three pairwise intersections.

Let us find all valid diagrams which minimally contain three pairwise intersecting hypergraphs. Such diagrams consist of three corners, ladders of the three hypergraphs, and any internal 2-cells contained in the disk formed by the three-hypergraphs. Let Δ indicate such a diagram. The boundaries of ladders of the three hypergraphs form $\partial\Delta$, for if any boundary cell is not from a ladder, then Δ is not minimal. To satisfy Theorem 2.1, we must have $\frac{4}{3} \leq \frac{6+E}{3+E+I}$, where E is the number of added boundary 2-cells and I is the number of added internal 2-cells, which implies $E + 4I \leq 6$. Thus we need $I \le 1$. If $I = 1$, $E = 0$. This leads to exactly one possible disk diagram, which contains the first diagram in Figure 12 as a subdiagram, hence is not minimal. See Figure 11.

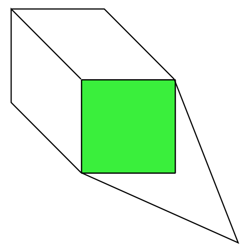

FIGURE 11. The only diagram with one internal cell and three corners

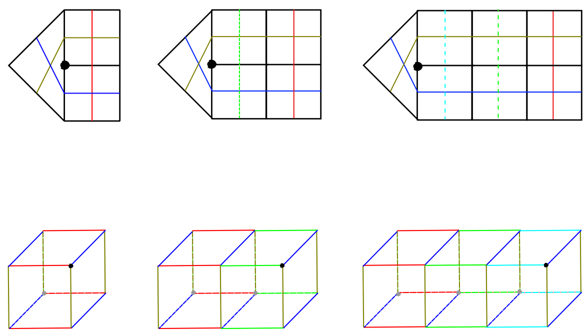

FIGURE 12. Top: (d), (e), and (f) are disk diagrams that lead to three cubes, below. Black dots mark the identity. Nonprincipal vertices are marked in gray.

So $I = 0$, and it follows that at least two of the corners are adjacent. With no internal cells, this constraint leads immediately to the three allowable diagrams which lead to 3-cubes: Figure 12.

Note that diagram (d) can be seen as the corner of a cube, and hence is actually symmetric. Note also that with the presence of other hypergraphs, the (e) and (f) 3-cube pictures actually create two or three adjacent 3-cubes in the cube complex, and hence more than one nonprincipal vertex in some 3-cubes. Finally, we see that diagram (d) is a subdiagram of diagram (c) from Figure 6. All eight corners of the cube are principal in diagram (c). □

COROLLARY 4.9. *Any 3-cube in the cube complex X has at least six principal vertices, and any nonprincipal vertices share an edge.*

PROOF. If a 3-cube C arises from diagram (d), (e), or (f), it has up to two nonprincipal vertices, marked as gray in Figure 12. If there are two nonprincipal vertices, they share an edge as seen in the Figure. □

The following lemma is used later in the chapter.

LEMMA 4.10. *No planar diagram Δ consists of a 3-cube diagram from Figure 12 sharing two 2-cells with a copy of diagram* (b) *from Figure 6.*

PROOF. By Remark 4.5, we need only show that it is impossible to glue a copy of diagram (b) to a copy of diagram (d) and satisfy the isoperimetric inequality; (e) and (f) follow. As Figure 13 shows, such a gluing fails the isoperimetric inequality with five 2-cells and boundary length six; alternatively one can use Remark 4.2 to see that Cancel of the resulting diagram is 7 while the size of the diagram is 5.

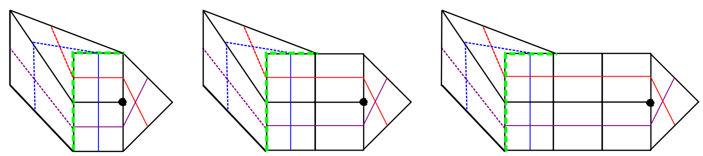

FIGURE 13. Gluing two of the cells of the original diagram (b) to the 3-cube diagrams of Figure 12.

$\square$

We use Lemma 4.8 to rule out cubes of dimension 4 or more in the cube complex.

LEMMA 4.11. *(4-cubes) There are no 4-cubes (and hence no n-cubes if $n \geq 4$) in the cube complex X.*

PROOF. We build all diagrams with four hypergraphs pairwise intersecting, starting with a 3-cube diagram from Figure 12, then show they fail our inequalities. Throughout we refer to Figure 14. First we use (e) as our base 3-cube, and note Cancel$(e) = 6$ and $|e| = 5$. There are already four hypergraphs present in (e), so to form a 4-cube diagram we need to add more cells and either make the non-intersecting hypergraphs intersect, or add a new fourth hypergraph. First we try adding no new hypergraphs. Forming a new diagram (e') by adding one more cell makes Cancel$(e') = 8$ and $|e'| = 6$, which violates the strictness of $(*)$. In Figure 14 we show two different gluings for e', and as mentioned in Remark 4.5 this does not change our failure of the inequality. We could also add two more 2-cells and form an (f) subdiagram, but by the second fact in Remark 4.5 this diagram (e'') will also fail the inequality.

Next we try to use a fourth hypergraph not already present in (e), which means we add at least three more 2-cells for the three intersections with existing hypergraphs. Adding exactly three 2-cells to form (d) and (e) subdiagrams in a non-planar diagram (e''') adds 5 to Cancel(e), which doesn't satisfy the inequality. As above, placement of the additional cells does not change the ratio

$\frac{\text{Cancel}(e''')}{|e'''|}$. Also as in e'', adding two more 2-cells to form (f) subdiagrams still fails the inequality (this case is not pictured).

Now we start with (f) as our base 3-cube. We need only consider diagrams that use (d) as the additional 3-cubes, for in the previous paragraphs we have dealt with all possible 4-cube diagrams that contain a copy of (e). Any such diagram will contain a copy of (e') as a subdiagram, so we are done.

Finally we try to build a 4-cube using only copies of (d). Adding squares to any boundary of (d) results in (e). If we add squares to interior edges of (d), we create (d') with $|d'| = 6$ and $\text{Cancel}(d') = 8$, which also fails the strict inequality. $\square$

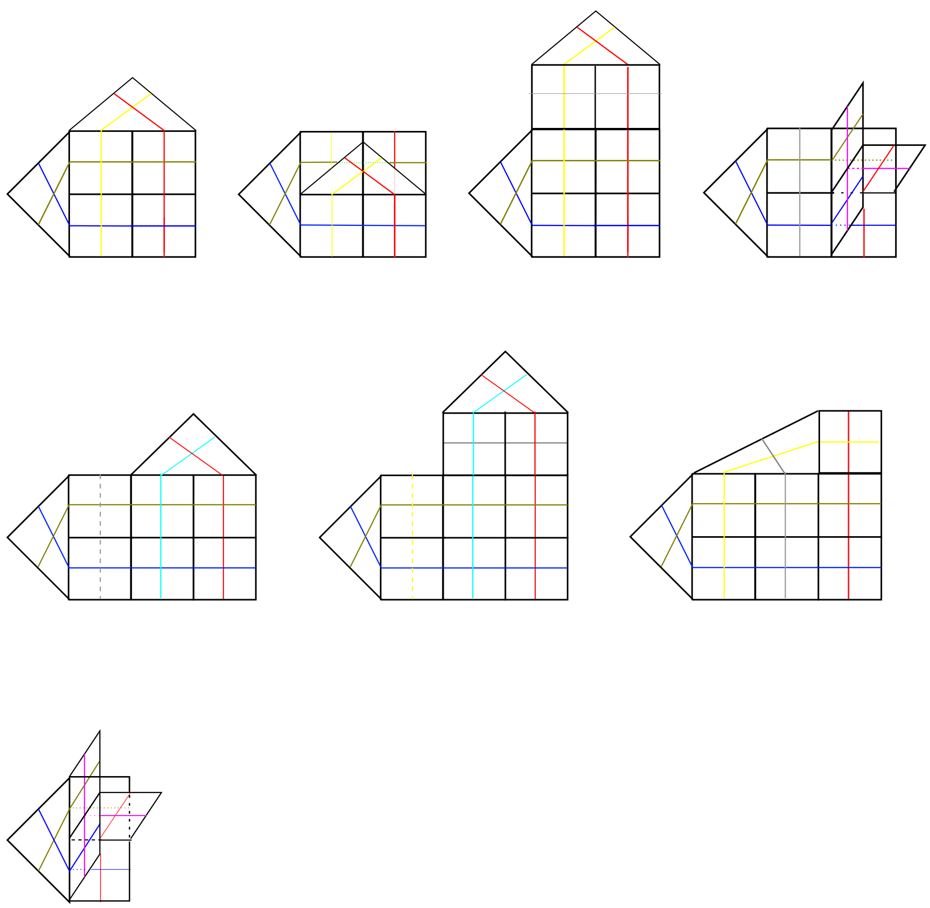

FIGURE 14. Cases from Lemma 4.11. Top: two different (e'), (e''), (e'''). Middle: (f'), (f''), (f'''). Bottom: (d').

We now turn to the G-action on X. If $g \in G, x_v, x_w \in X$ are principal and $g.v = w$ for the corresponding group elements v, w, then $g.x_v = x_w$. Throughout the paper, if g stabilizes a principal vertex

x_v, we assume that v is the identity element. For as $v.x_e = x_v$, we have $v^{-1}gv.x_e = x_e$ and $v^{-1}gv$ stabilizes the identity vertex e in Γ. This conjugation does not affect any of our arguments.

The following technical lemmas explore the action of G on X in more detail. For an illustrative example, Figure 15 shows some possible images of these isometries on an example 3-cube. We use these lemmas to prove Proposition 4.15 regarding stabilizers of vertices under this action.

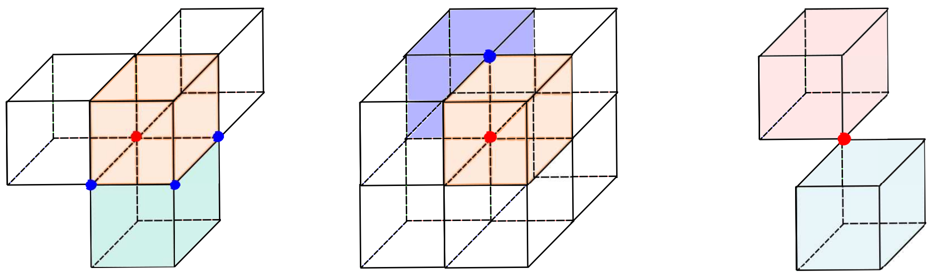

FIGURE 15. The left shows a reflection through a face with four fixed vertices and all x-stabilizing such cube images. Then a rotation about an edge with two fixed vertices, with the six different possibilities. Finally a rotation about a corner with one fixed vertex, rotated for perspective.

LEMMA 4.12. *If C is a 3-cube in X and there exists a $g \in G$ and vertex $x \in C$ such that $g.C \cap C = \{x\}$, then x is principal.*

PROOF. We need to address the cases where C arises from diagram (d), (e), or (f) in Figure 12. First, suppose C arises from diagram (d). As $g.C \cap C = \{x\}$, no hypergraph is sent to another of the three under the g action. Consider the principal vertex y which is antipodal to x in C, so $d(x, y) = 3$, where distance is measured in the 1-skeleton of the cube complex. Since C is non-positively curved, the combinatorial distance in the 1-skeleton equals the number of hypergraphs crossed by the geodesic path, per Theorem 14.3 in [**34**], and so $d(g.y, y) = 6$. Since y is principal by Corollary 4.9, we assume it is x_e, and so $g.x_e = x_g$ is principal. Then g is on the other side of six hypergraphs from e in Γ: the original three corresponding to the cube, and their images under the g action.

In the bottom of (d) in Figure 12, we can indicate where x_e is located in X as the black dot, if x is the antipodal vertex marked by the gray dot. Now g is somewhere in the Cayley complex on the other side of all three hypergraphs from e, as on the left hand side of Figure 16. Without loss of generality we choose a single two-sided hypergraph l to cross first. Crossing l puts us in a region l^+ from which we must cross the two other hypergraphs h and k. But h and k cross in l^-, so the only way we can cross both and end up with a principal vertex is if they cross again somewhere further in l^+. Two hypergraphs intersecting twice must satisfy one of the diagrams in Figure 6 per Lemma

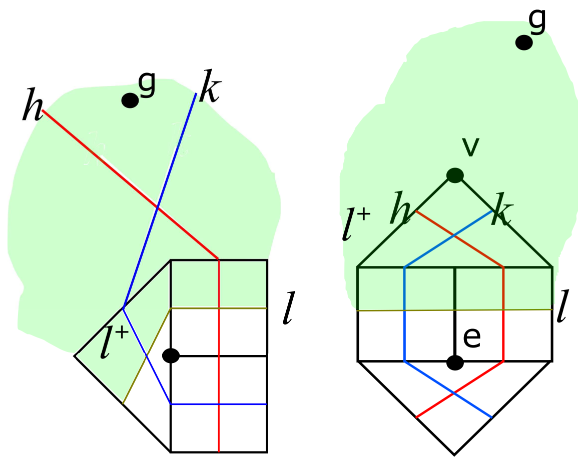

FIGURE 16. The hypergraphs from diagram (d) in Figure 12, with e and $g.e$ vertices marked.

4.3. Since l lies between the two crossings, h and k must form diagram (c), as in the right hand side of Figure 16. As no other hypergraphs have been involved, the new region includes a vertex v marked in the figure such that x_v has all the same orientations as x, so $x_v = x$. Thus x is principal.

Now suppose C is represented by diagram (e) and now has two nonprincipal vertices, and x is nonprincipal. If x is shared by C and its neighboring 3-cube which arises from the (d) subdiagram in (e), then x is the nonprincipal vertex in a (d)-diagram cube and the argument above holds. Then let x be the second nonprincipal vertex, and choose x_e to be antipodal to it in C. As with diagram (d), g must be on the other side of all three hypergraphs from e. By an identical argument to the (d) case, this requires two of the three to cross a second time, on the other side of the third. However, none of the diagrams in Figure 6 can accommodate diagram (e), so this is impossible.

The same argument as (e) holds for diagram (f). □

LEMMA 4.13. *If C is a 3-cube in X and $g.C \cap C$ is an edge with at least one fixed endpoint vertex, then the vertices in that edge are principal.*

PROOF. If $g.C \cap C$ is an edge including x and x_v, then exactly one of the three hyperplanes defining C, say h, is sent to one of this set. Suppose that C arises from diagram (d), and assume for contradiction that x is non-principal. Then x_v is a principal vertex associated with one of the corner vertices v in diagram (d). We glue on a second 3-cube diagram that shares the vertex v with the initial diagram or connects to the initial diagram through a series of copies of diagram (a), which form new vertices in Γ that correspond with the same vertex v in X. To share x, the second 3-cube diagram must contain the hyperplane h which runs through the two squares that do not include v (for v is adjacent to x through that hyperplane), as in Figure 17.

The carrier of h connects the two diagrams. Since Γ is simply connected and a path through v also connects the diagrams, the carrier of h and v form a loop which is filled by a disc in Γ. In particular, the hypergraph from our original diagram which enters this region next to v must exit the region at some point. Call this hypergraph l. Note l cannot exit through the original diagram nor through any copies of diagram (a). So l must cross h again. As in Lemma 4.12, l and h must form diagram (c), which forms a new region with all the same orientations as x, so x is principal.

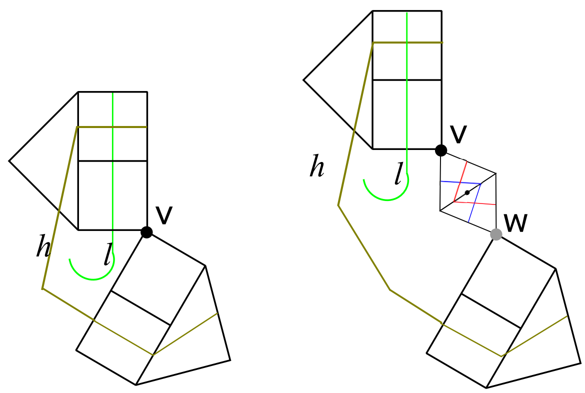

FIGURE 17. Gluing another 3-cube to the original (top left) one. They share vertex v and hypergraph h. On the right, an example of an included diagram (a), so v and w correspond to the same x_v in X.

Now suppose C arises from diagram (e), and again x is a nonprincipal vertex. If the shared edge is between x and a principal vertex, then the same argument as in diagram (d) holds. Hence assume the shared edge is between the two nonprincipal vertices in diagram (e). The image cube $g.C$ has six principal vertices by Corollary 4.9. There must exist an image hypergraph l' so that crossing l' from x results in a region containing a vertex w in the Cayley complex Γ, leading to a principal vertex x_w in $g.C$. Consider the orientations that result in x in diagram (e), as in Figure 18. We know x_w differs from x only in l', and so has the same orientations of the four hypergraphs in diagram (e) as x. Then, as in Lemma 4.12, two of the hypergraphs must cross again on the other side of the third, and hence form one of the diagrams from Figure 6. This is impossible, and so this case cannot occur. An identical argument holds for diagram (f).

□

LEMMA 4.14. *If C is a 3-cube in X, $g.C \cap C$ contains a face including a nonprincipal vertex x, and $g.x = x$, then g fixes more than one vertex in C.*

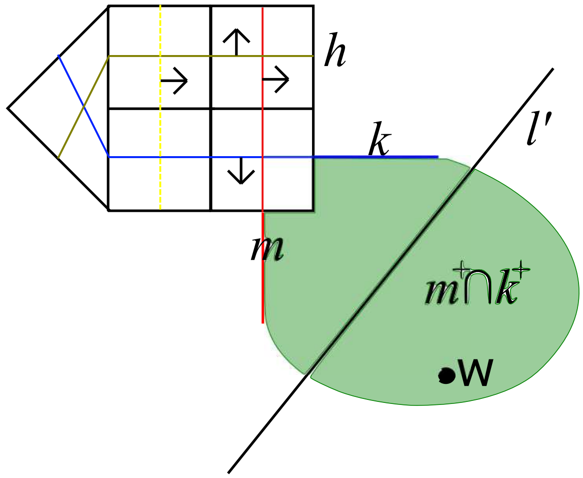

FIGURE 18. k and h need to cross in m^+, since w is on the positive side of h. This is impossible.

PROOF. Suppose $g.C \cap C$ contains a face and g fixes a vertex x in that face. Assume C arises from a (d) diagram. Choose three vertices as the rest of the shared face with x; these all lie on one boundary side of the (d) diagram. We need to find a diagram to represent the adjacent 3-cube which shares this face.

By Lemma 4.8, there are only three possible diagrams to represent the adjacent 3-cube. Since G acts on Γ by isometries, $g.C$ must also be represented by a (d) diagram. Since C and $g.C$ share x, which is defined by its orientations of three specific hypergraphs, $g.C$ must be represented by a (d) diagram with the same three hypergraphs as C. Note here that this may not be the exact same collection of 2-cells as C; for instance, C and $g.C$ could form a copy of diagram (c) from Figure 6. Nevertheless, any three pairwise intersecting hypergraphs uniquely determine a cube in X, so $g.C = C$ necessarily. Then the antipodal vertex to x is fixed and we are done.

An identical argument holds if C arises from diagrams (e) or (f). □

PROPOSITION 4.15. *If x is a vertex in the cube complex, then $Stab(x) \le Stab(x_w)$, where x_w is a principal vertex.*

PROOF. Let x be an nonprincipal vertex in X, and $g.x = x$. Let d be the maximal dimension of cubes containing x. We prove the claim for each d by either finding a principal vertex whose stabilizer contains $Stab(x)$ or showing that x itself is principal. Since the Cayley complex is a square complex where every hypergraph intersects at least one other hypergraph, every vertex of the cube complex is part of a square, so $d > 1$.

If $d = 2$, then x is the corner of a square, corresponding to an intersection of two hypergraphs in the Cayley complex. Any hypergraph intersection in the Cayley complex results in a collection of parallel squares in the cube complex. To have a nontrivial such parallelism between d-cubes, one needs $(d+1)$ cubes to support the parallel copies. But since $d = 2$ is maximal, there is only one such square, and x corresponds to one of the four regions formed by this intersection, so x is principal.

Now suppose $d = 3$, so we can apply the previous lemmas. Let C denote a 3-cube containing our nonprincipal vertex x. If there exists a $g \in Stab(x)$ that fixes only x and no other element in C, then $g.C \cap C$ is either x, an edge containing x, a face containing x, or all of C. But if $g.C \cap C = C$, then by the argument in Lemma 4.14 g fixes a second vertex, so this case cannot occur. The remaining cases are dealt with in Lemmas 4.12, 4.13, and 4.14, which say that x is actually principal, so we are done.

Otherwise every g in $Stab(x)$ fixes more than one element in C. By the previous three lemmas, g must send C to C. Since g fixes x, this means g must also fix the vertex x_w antipodal to x across C. Since x_w is antipodal to x, it cannot be nonprincipal by Corollary 4.9. Hence $Stab(x) \leq Stab(x_w)$, a principal vertex.

We cannot have any diagrams that satisfy $(*)$ with $d \geq 4$ by Lemma 4.11.

□

We use the lemmas from this section to prove freedom of vertex stabilizers in Proposition 4.18 and Corollary 4.19 in the next section.

5. Proof of Theorem 4.1

To prove Proposition 4.16, which shows finite generation, we build a specific diagram Δ out of diagrams collared by hypergraphs and paths. Our goal is make Δ a collection of disk diagrams glued together at cut points. We start with a principal vertex-stabilizing element $g \in G$, and assume as before that g stabilizes the identity vertex x_e in X. So we choose a geodesic path γ in Γ starting from e that reads off a word w that represents g.

Since $g.x = x$, any hypergraph that γ crosses is crossed again later in γ. Look at the hypergraph h_0 which crosses the first letter of w. Read along w until we cross h_0 again. The traversed segment of γ and part of h_0 form a diagram Δ collared by a path and a hypergraph, as guaranteed by Lemma 3.17 in [**30**]. Note h_0 does not intersect any interior cells of Δ, since we chose the first two (h_0, w)

intersections along γ, so we meet the requirement to be a collared rather than quasicollared diagram. Let σ indicate the non-γ part of $\partial\Delta$. We continue to build up Δ iteratively.

Look at the next hypergraph h dual to γ, reading along w as a word. We consider the parity of the number of times h has intersected γ, including this current intersection point.

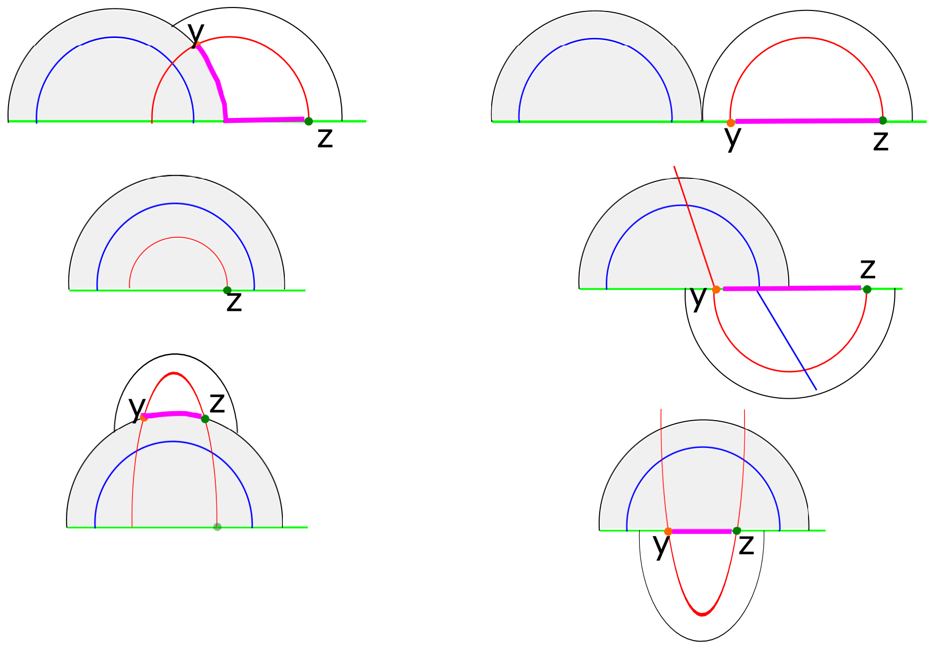

FIGURE 19. Building the diagram for Proposition 4.16. z is indicated by the green dots. On the right, y (orange dot) is our current point; on the left we choose y or it does not exist. In the bottom left, we also choose a new z; the old z is indicated in a lighter color. The path p is in fuschia.

Assume this parity is odd, so there is a next (h, γ) intersection point z as we read w which we match with our current intersection point. Starting from our current intersection point and traveling along h toward z, if h exits Δ, we label that point y. Note y may be our current intersection point. If no such y exists, the entire ladder of this segment of h between our current point and z is already in Δ and we do nothing. If the segment of h between z and y intersects Δ, we choose the intersection point of $\partial\Delta$ and h and rename it to z. These cases are illustrated in Figure 19.

Consider the segment of h between z and y, and the path on $\gamma \cup \partial\Delta$ that connects z and y consisting of subpaths of γ and possibly σ. These bound a disk, and since we chose z to be the next intersection point, this forms a diagram collared by a hypergraph h and a path p'. Here, p' is the path p between z and y extended to include the half-edges on either end which connect to the ladder of h.

By our choices of z and y, the interior of this new collared diagram is disjoint from our existing Δ and the next edge of γ which continues the ray from y to z. We glue the new diagram to Δ along

p', and call the resulting diagram Δ as well. We note by our construction that h may form disjoint components inside Δ, as in Figure 20. In this case when considering the hypergraphs of Δ we treat each of these components as a distinct hypergraph.

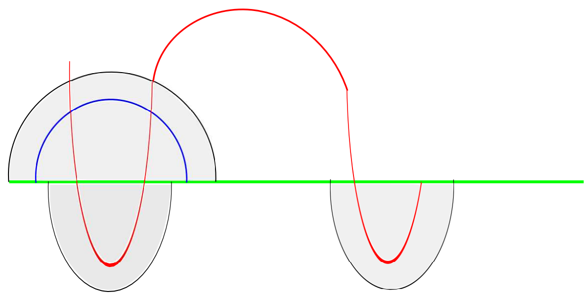

FIGURE 20. Same hypergraph in Γ may make separate components in Δ, we treat this as two different hypergraphs.

If the parity is even, then we have already added the relevant disk to the diagram from the previous time we encountered h, and we do nothing. Once we reach the end of γ, we have a Van Kampen diagram Δ of which part of the boundary may be part of γ, and the rest of the boundary is a concatenation of boundaries of carriers of $\gamma-$crossing hypergraphs.

Note that γ separates the planar diagram Δ into two sides, each of which may contain several or no components, which we can label g^+ and g^-. At least one of these sides is non-trivial if Δ is non-trivial, which is the case as long as γ is non-trivial. These sides depend on our choice of planar embedding of Δ, but once we have chosen such an embedding the sides are no longer ambiguous. Via this construction, each $\gamma-$crossing hypergraph h contributes a disk subdiagram to Δ, on one of these sides. We say h *belongs to* g^+ if its disk subdiagram appears in g^+, and the same for g^-. Note also that Δ can be seen as a collection of disks D_i connected by cut-points that lie along γ: every edge in γ borders a square, so there are no cut-edges, and by construction the only cut-points are on γ.

PROPOSITION 4.16. *(Finite generation) In the square model at $d < 1/3$, principal vertex stabilizers are generated by two letter words of a certain form, which can be seen as green paths in Figure 21.*

Note that as hypergraphs are embedded trees, the labels of the paths in ∂D complementary to the marked green paths are uniquely determined by the labels of the green paths in each diagram D in Figure 21. For the same reason, given a generating two letter word as in Proposition 4.16, there exists a unique diagram from Figure 21 where it appears as the label of a green path.

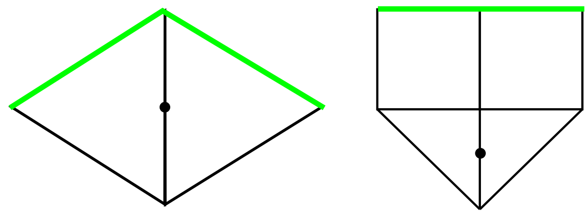

FIGURE 21. Options from Proposition 4.16: note the two letters of each green path are dual to the same hypergraph

COROLLARY 4.17. *Vertex stabilizers of principal vertices are finitely generated.*

PROOF. Random groups are finitely generated, so this follows immediately from Proposition 4.16. This also follows immediately from Corollary C in [**13**]. □

Throughout the proof of Proposition 4.16 we find two hypergraphs that intersect twice, and apply Lemma 4.4 to say they must form a diagram from Figure 6. We consider these requisite diagrams without explicitly referring to Lemma 4.4.

PROOF OF PROPOSITION 4.16. Suppose $g \in G$ stabilizes a principal vertex, which without loss of generality is x_e. Choose a word w of minimal length that represents g, and look at a geodesic path γ in Γ which reads off w. The path γ cannot essentially cross any wall, for then g would swap such a wall's orientation and move x_e to a different vertex in X under the action. We use the construction above to form the diagram Δ, consisting of disk diagrams D_i connected at cut points which lie along γ.

Each D_i contains a specified subpath $\gamma_i \subset \gamma$ representing a group element g_i so that $g = g_1 \cdots g_n$. Below, we inductively choose a w_i' equal in the group to g_i, and show w_i' can be written as a product of appropriate two-letter words by finding a path γ_i' that reads off w_i'. We then concatenate the w_i' appropriately to find a word w' equal in the group to g, and as each factor is a product of appropriate two-letter words, so too is w'. Then we need only consider a single factor, so in the following, we use Δ, g, w' for D_i, g_i, w_i' and assume we have no cut-points in Δ.

Starting from the beginning of w, we order the components of the boundary path σ (which is cut into components by γ) based on their first intersection points with γ, read as a word w. By construction, no two σ components intersect γ at the same initial point, and so we can index them into σ_j by these initial intersection points. We use these σ_j to build γ'.

Base Step. Begin with σ_1. Suppose the first letter of σ_1 is dual to a hypergraph l. Upon entry to Δ, l immediately intersects a γ–crossing hypergraph h. Note that h may be dual to the first or second letter of g, as in Figure 22. Say h belongs to g^+, and call this initial intersection 2-cell C. We consider the possibilities for l:

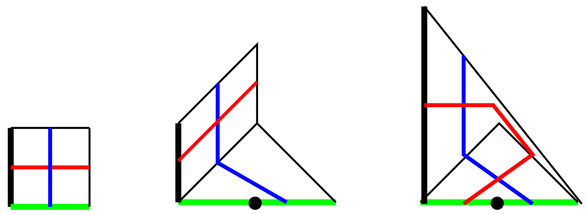

FIGURE 22. Left: h is dual to the first letter of g. Right: h is dual to the second letter of g. Far right: h is dual to the second letter of g and forms a diagram (a) with l. The blue hypergraph is h; the red is l. The thick line represents σ_1 and the green is γ.

(1) Suppose l does not cross γ. As l must exit Δ, it must re-cross h inside Δ and so form one of the diagrams from Figure 4.3 as a subdiagram of Δ. If they form diagram (c), then there is a third hypergraph k which intersects both.

(1a) Suppose k does not cross γ. Then as with l, k must intersect h again to exit the diagram. Consider the subdiagram of Δ formed by the two $k - h$ intersection cells and the ladders of k and h between these cells. As the immediate neighbors of k along h are both l, this subdiagram must be a diagram (c) with l as the third hypergraph. Thus we must glue together two copies of diagram (c) including the new intersection cell and the one included in our initial diagram (c), which is impossible by Lemma 4.6.

(1b) Now suppose k crosses γ. If k belongs to g^+, then l must cross k again to exit Δ. If k belongs to g^-, then h must cross k again to exit Δ. In either case, k again forms a second copy of diagram (c) sharing one of the cells from the initial diagram (c), impossible as in the previous case.

Thus l and h must form diagram (a) or (b). We choose the two letters dual to l along σ as the first two letters of w', which gives us a two letter factor of our final word with the property we desire.

(2A) Now suppose that l crosses γ and h is dual to the first letter of g. Since l is the first hypergraph dual to σ_1, l belongs to g^- and so must cross γ again. In order to exit Δ, l must intersect h twice and C must be a corner of the resulting diagram, which is one of diagrams (a), (b), or (c). By the above argument, they cannot form diagram (c). So the two intersections must form diagrams (a) or (b). Since h belongs to g^+ and l belongs to g^-, a subsection of γ must run through the diagram. By choice of C, γ intersects h in the diagram. We claim that γ also intersects l in the diagram: for if l does not intersect γ, l exits Δ after it intersects h for the second time. There is no possible other hypergraph belonging to g^+ to allow l to stay in Δ as h is the first hypergraph forming σ_1.

As γ runs along the boundary edge of C dual to h if h intersects γ on the first letter of g, the first edge must appear on the boundary of the resulting diagram. We deal with the case when h intersects γ on the second letter in Case 2B.The possibilities for the first few edges of γ appear in Figure 23.

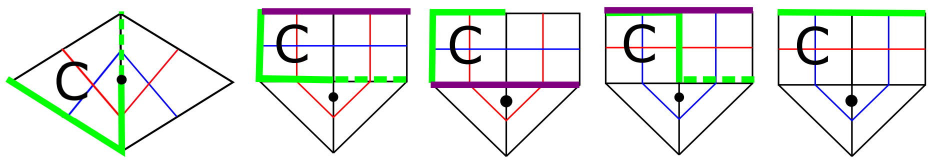

FIGURE 23. Possibilities for the first letters of w (hence of γ) in the base case are in green, and choices for w' are in purple if such a choice exists. The blue hypergraph is h; the red is l.

(1a) If l and h form diagram (a), the two 2-cells of diagram (a) must lie on opposite sides of γ. This forces three of the edges of C to lie along γ, which contradicts the geodesicity of γ.

(1b) If l and h form diagram (b), we have four possibilities for the γ path.

(1bi) If h is the shorter hypergraph in diagram (b) and γ begins in the top left corner, then we choose the first two letters of w' as the ones on the top of diagram (b), dual to l. Note that in this case that h cannot be the longer hypergraph, since the γ edge dual to h must lie on ∂C. We claim that this choice of two edges for γ' ends at a point on σ_1.

Suppose that we do not, so there is some other edge e which is the second edge of σ_1, and a corresponding hypergraph k dual to e. As h belongs to g^+, k cannot be h. Note the second edge of σ_1 shares a vertex with the second edge of our chosen γ'.

In order for l to be dual to two edges that share a common vertex, l must pass through at least two 2-cells between these edges. This may form a diagram (a), as on the left of Figure 24. Alternatively, l may pass through four 2-cells, which form a diagram (b), as in the next diagram of Figure 24. Any further nesting of copies of diagram a violates the isoperimetric inequality, so these are the only two possibilities. In either case, the resulting diagram violates the isoperimetric inequality by Lemma 4.7, so k cannot be l.

By the construction of Δ, after entering Δ k immediately crosses another hypergraph j which belongs to g^+. So j cannot be l. If two adjacent edges on σ_1 belong to the same γ–crossing hypergraph, then they must be consecutive in the ladder of that hypergraph, so j cannot be h either. See the right diagram in Figure 24. As j belongs to g^+, in order to exit Δ l must cross j twice. Then l and j form one of the diagrams from Figure 4.3 as a subdiagram of Δ. This subdiagram contains the entire ladder segments between intersections of both hypergraphs, so it must share at least a vertex with the original $l-h$ diagram b. By Lemma 4.7, this cannot occur.

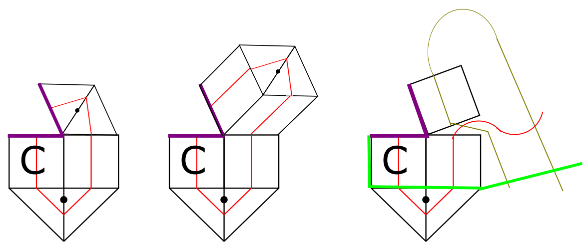

FIGURE 24. The argument that j is not l (left two diagrams) and intersects l twice (right diagram). In all, l is red hypergraph and j is olive.

Then l and j cannot intersect twice, so j cannot exist. Then k does not exist, and so both of the edges on the top of diagram (b) lie on σ_1, and so this choice of w' allows us to move on to the inductive step below.

(1bii) If h is the shorter hypergraph in diagram (b) and γ begins in the middle left of the diagram as in the third picture in Figure 23, then no part of C is on $\partial\Delta$, so l is not the first letter dual to σ_1 as we assumed. Then this case cannot occur.

(1biii) If h is the longer hypergraph in diagram (b), then we choose the first letter of w and the second letter dual to h as the first two letters of w'. Note that γ has only one possible third edge in the diagram, as marked in the fourth diagram of Figure 23, as γ must be a geodesic. So the top right cell of diagram (b) is the first cell in the carrier of l. Thus we end up on σ_2 and we can continue to the inductive step with this choice of w'.

(1biv) Suppose γ runs along the top two edges of diagram (b) and does not cross l in the diagram, as in the last diagram of Figure 23. In order for l to cross γ without exiting Δ, there must be a third hypergraph k which belongs to g^+ and defines some of σ_1. The only possibility for k is the third hypergraph in diagram

(b). To exit g^-, k must intersect l, which means that h, k, and l all pairwise intersect and form a 3-cube. This cannot occur by Lemma 4.10.

(1c) As in case 1 when l does not cross γ, l and h cannot form diagram (c).

(2B) Suppose l crosses γ and h intersects γ on the second letter of w, as in the right two diagrams of Figure 22. We show that this case cannot occur.

(1a) Suppose l and h form a copy of diagram (a) with γ on two edges, as in the right side of Figure 22. As l belongs to g^-, h must intersect l in g^- in order to exit Δ. Then h and l must form a second diagram from Figure 6 glued to this copy of diagram (a) sharing one 2-cell; any such possibility violates the isoperimetric inequality by Remark 4.2.

(1b) Suppose there exists a third hypergraph k which is dual to the first letter of w, and which intersects both h and γ, as in the middle diagram of Figure 22. Let C' denote the 2-cell whose boundary includes the first two edges of γ.

(1bi) Suppose k intersects h twice; this case is illustrated in Figure 25. If k intersects h again in g^-, then they may form a diagram (a) with C'. In this case k belongs to g^-, and intersects h again in g^+ to leave g^+. If they intersect in g^- but no diagram (a) occurs, then k belongs to g^+ and because of C along the ladder of h, k and h form a diagram (c). If k intersects h again in g^+, then due to the presence of C, k and h must form diagram (c) with l as the third hypergraph. We know l and h intersect twice, and as l is the third hypergraph in the $k-h$ diagram, l and h must form another diagram c, which is impossible by the argument in Case 1.

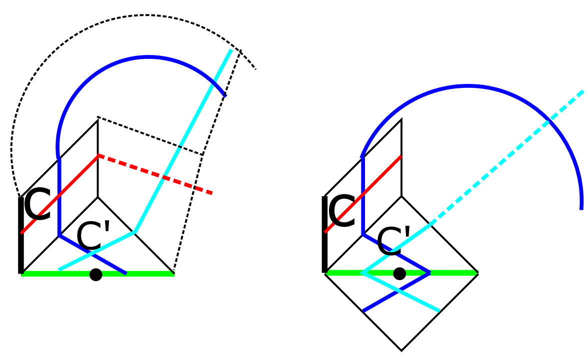

FIGURE 25. Left: presence of l (in red) forces a diagram c between k (in teal) and h (in blue). Right: If k belongs to g^-, same effect.

(1bii) Then assume k does not intersect h twice. Then k belongs to g^+, and the second $k-\gamma$ intersection point lies between the $h-\gamma$ intersection points along γ. The following argument is used often later (we refer to it as Case 1(b)ii), so we ignore the presence of C (which would just make the failure of the isoperimetric inequality worse).

(A) Suppose there exists a hypergraph j which intersects k twice in g^+. If j does not intersect h twice in g^+, then j belongs to g^+ and its γ intersection points lie between the $h-\gamma$ intersection points, and we apply this argument to j instead of to k. Then assume j intersects h twice in g^+. The only possibility is that the three hypergraphs form a subdiagram (b) from Figure 4.3. Note if this subdiagram b is glued to C' along two edges, we fail the isoperimetric inequality as we did in Case 1(b)i. So ladders of l and k must exist between C' and diagram b. Such a diagram with no internal 2-cells has $8+2n$ boundary length and $7+2n$ 2-cells as seen in the left side of Figure 26, which also fails the isoperimetric inequality. Any diagram with internal 2-cells will have a lower boundary-length-to-area ratio, and also fail the isoperimetric inequality. Then no j intersecting l twice can exist.

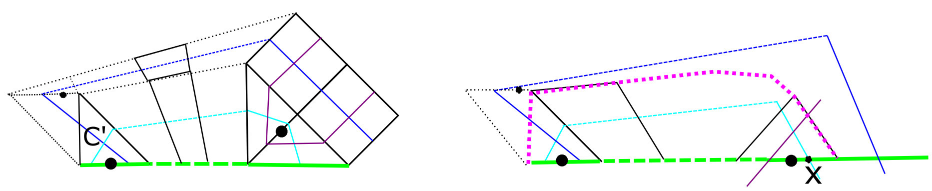

FIGURE 26. Failure of the isoperimetric inequality if the $k-\gamma$ and $h-\gamma$ intersection points stagger. The blue hypergraph is h, the teal is k, purple is a hypergraph j that intersects both twice in the left, and once on the right.

Note that this argument also holds if we assume γ runs along the other two edges of C', so that the $k-\gamma$ intersection points are nested between the $h-\gamma$ intersection points and k and h do not intersect. We use this argument in later cases.

(B) Every hypergraph that intersects k in g^+ only intersects it once, and so must intersect γ. Let x denote the second intersection point of k with γ. The hypergraph j that crosses k in the 2-cell containing x also intersects γ. Following the path from the first letter of γ to x along the top boundary of the ladder of k, as the pink line in Figure 26 indicates, gives a length $n+1.5$ path, where n indicates the number of hypergraphs that intersect k in g^+. However, the γ path between these points on the other side of the ladder of l has length $n+2.5$, so γ is not a geodesic, contradiction. So the $k-\gamma$ intersection points cannot stagger with the $j-\gamma$ intersection points. As above this argument holds if γ runs along the other two edges of C' as well and we cut out the $k-h$ intersection.

Inductive Step. We have established the base case of our inductive argument. Assume for our inductive hypothesis that we have just added a new pair of edges to γ', which has even length, and which ends at distance less than or equal to 1 from γ. Thus far γ' is a product of 2-letter words from the statement of Proposition 4.16, and we have ended at a vertex on σ_i.

Now we iterate the argument, and show that we end at a vertex on σ after choosing particular paths to add to γ'. Note if σ_i ends on γ and we are at that ending vertex before choosing our next edges for γ', then we have already written a subword of w as a product of desired 2-letter words. Then we could start at the next letter as the beginning of a new w, creating a new diagram for it, which would also create a cut-point in our diagram. So we have reached the end of w as we assumed there are no cut-points in the construction of the diagram, and the algorithm is complete. Hence we assume that we have not previously ended where σ_i and γ meet. Let l be the next hypergraph dual to σ_i, heading toward the end of w from the end of γ' as constructed so far.

As in Case 1 of the Base Step, if l does not cross γ, then l must immediately exit the diagram forming diagram (a) or (b).

So we assume l crosses γ, intersecting a γ-crossing hypergraph h as l enters Δ at a point p in a 2-cell C. Assume σ_i is in g^+, so h belongs to g^+ and l belongs to g^-. Then l must intersect h again, and form a subdiagram from Figure 4.3. As in the Base Step, l and h do not form a diagram (c), so l crosses h at consecutive edges along h. Since l enters Δ at p, the next edge of σ_i follows the ladder of h and not some other hypergraph that belongs to g^+. Thus if l does not cross γ between the two times that h crosses γ, then l exits σ_i without crossing γ at all.

Then we assume l crosses γ at least once between the two times that h crosses γ. Then at least one of the $l-\gamma$ intersection points lies between the two $l-h$ intersection points on l, so γ must intersect l in the subdiagram.

Each of the following cases offers a different configuration for l, h, the σ_i edge dual to l, and the portion of γ inside the subdiagram. In each we propose the two letters to add to γ'. We need to show two things: that we start at the appropriate vertex of the indicated σ_i edge, and that we end on σ so the algorithm can continue.

(1) If they form diagram (a), it may be contained in a (b) superdiagram.

(1a) If this (a) is contained in no (b) superdiagram, then there are two paths for γ to run through the diagram: either dual to both h and l, or along the other boundary edge dual to l, as illustrated in Figure 27. In either case, l belongs to g^- and we choose the two edges dual to l to add to γ'.

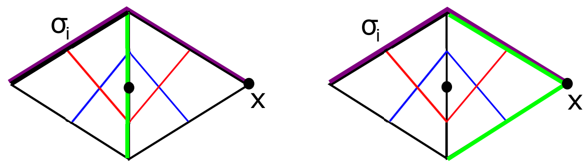

FIGURE 27. The possibilities for γ in diagram (a) are in green. The red is l, blue is h, top left edge is part of σ_i, and we choose the purple edges for γ'.

Our inductive hypothesis says that we do not start where σ_i meets γ, so we start at the desired left vertex of σ_i in the figure.

Suppose that we do not end on σ_{i+1}, so there is some hypergraph k which belongs to g^- that hits γ earlier than l does. To leave g^-, h must intersect k. Depending on whether h or k intersects γ earlier, h and k intersect a second time in g^- or g^+. So h and k intersect twice, so they form a subdiagram from Figure 4.3. As the $h-l$ intersections lie between the $h-k$ intersections as read along h, the $h-l$ diagram must be a subdiagram of the $h-k$ diagram. We assumed this diagram is contained in no (b) superdiagram, so k cannot exist. Thus the terminal vertex of the second edge dual to l edge lies on σ_{i+1}, so our algorithm can continue.

(1b) If this (a) is contained in a (b) superdiagram, then l is entirely contained in the (a) diagram as it must enter on a boundary, as in the left hand picture in Figure 23. Since γ must cross l inside the diagram, γ must intersect σ_i at the bottom of diagram (b). Let k denote the third hypergraph in the (b) superdiagram, as in Figure 28.

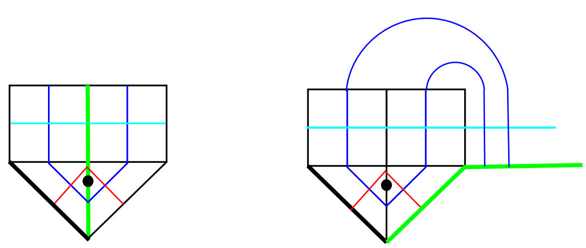

FIGURE 28. The possibilities for γ in diagram (a) with a (b) superdiagram are in green. As usual, red is l and blue is h. On the right, case where γ never intersects k, the hypergraph in teal.

If γ runs along the edge dual to l and never intersects k, then h necessarily intersects k two more times in order to cross γ, as on the right in Figure 28. There is no way to glue together three diagrams from Figure 4.3 that share appropriate 2-cells and satisfy the isoperimetric inequality, as can be readily seen using Remark 4.2.

Thus γ must run up the middle of diagram (b) by geodesicity as on the left in Figure 28, so γ intersects k.

Assume k belongs to g^+. If k intersects l, then we have three pairwise intersecting hypergraphs and need to glue a 3-cube diagram to this one, which is impossible by Lemma 4.6. Otherwise, either k intersects h a third time, which is impossible by Lemma 4.7, or its γ intersection points stagger with those of h which is impossible by an argument identical to Base Case 1(b)ii.

Thus k belongs to g^-. We have two possibilities.

(1bi) Suppose l intersects γ twice between the $k-\gamma$ intersection points. An argument identical to Base Case 1(b)ii shows this is impossible, where γ runs on the other two edges of C'. So the $l-\gamma$ intersection points cannot be nested inside the $k-\gamma$ intersection points.

(1bii) If l and k stagger their γ intersection points, then l and k intersect somewhere in g^-. Then h, l, k are three pairwise intersecting hypergraphs, and so form a 3–cube and a subdiagram from Figure 12. Gluing any of these three diagrams to the two 2-cells we know lie in g^+ fails the isoperimetric inequality as in Lemma 4.10, so this case also cannot occur.

(2) If l and h form diagram (b), our choice of γ' depends on how γ appears in diagram (b). Unlike in the Base Step, γ does not need to have an edge on the boundary of the diagram. The σ_1 case has already dealt with some of the possibilities for γ in Figure 23, but we need to address all of them. We consider all possible γ subpaths that intersect l, as in Figure 29.

(2a) Suppose γ runs along the vertical path of length 3 in diagram (b), so two cells of diagram (b) each lie in g^+ and g^-.

(2ai) Suppose l is the shorter hypergraph in diagram (b). This is the same as the earlier case with an (a) subdiagram contained in a (b) superdiagram (Inductive Case 1b), with l as k in that case and h remaining the same.

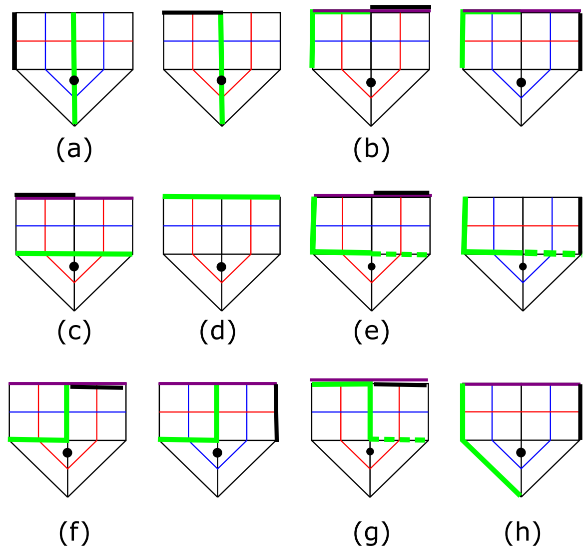

FIGURE 29. The other possibilities for γ in diagram (b), along with the choices for γ'. The red hypergraph is l, the blue is h, green is γ, purple is our choice of γ' if applicable, and the bolded edge is σ_i.

(2aii) Now suppose l is the longer hypergraph in diagram (b), and let k denote the third hypergraph in the diagram. By the argument in case 1b, k cannot belong to the same side of g as h does, so k belongs to g^-. To exit g^+, k must intersect h somewhere in g^+. Thus h, k, and l are three pairwise intersecting hyperplanes and must form a 3-cube diagram from Figure 12. By Lemma 4.10 this is impossible.

(2b) Suppose γ runs along two boundary edges of diagram (b), the outside of a corner. We choose the top two edges of diagram (b) as the next two edges of γ'.

(2bi) If l is the longer hypergraph in diagram (b), one of its edges is dual to γ and one edge is dual to σ_i by construction. If we start in the middle of the diagram, we either followed the edge dual to γ or a different edge to get here. If we followed the edge dual to γ, by the inductive hypothesis this edge is part of a diagram from Figure 19, which glues to our diagram b via an edge. Similarly, if we followed a different edge, we glue such a diagram to ours via a vertex. Both cases are impossible by Lemma 4.7.

(2bii) Suppose h is the longer hypergraph in diagram (b). Assume we start in the middle of the diagram, at the bottom of the edge of σ_i dual to l. If we have been following the ladder of h, then necessarily the bottom two edges of diagram (b) are also on σ_i. If h crosses l again in g^-, we have to glue a copy of diagram (a), (b), or (c) to this diagram (b) sharing a 2-cell, which fails the isoperimetric inequality by Remark 4.2. However, if h does not cross l again in g^-, by the inductive hypothesis the previous two letters we added to get to this point form one of the desired diagrams in Figure 19, which means we must have one of those diagrams glued to this one at a vertex. As in Lemma 4.7, this is impossible. So we start in the corner as desired.

Then we end in the middle of γ along the ladder of l. If this point is not the beginning of σ_{i+1}, then there exists a k belonging to g^- which hits γ earlier than l, and an argument almost identical to Inductive Case 1a holds, which would require a diagram from Figure 4.3 to contain this diagram (b) as a subdiagram. Thus we end at the beginning point of σ_{i+1} so our algorithm can continue.

(2c) Suppose γ runs along a path of length 2 in diagram (b) and separates an (a) subdiagram to g^-, so h is necessarily the shorter hypergraph . We choose the two letters dual to l along the top of diagram (b) for the next two letters of γ'.

If we are at the middle of the diagram, the argument for gluing another diagram to this one from Inductive Case 2(b)i holds. Thus we are at the corner of the diagram. We are either still on σ_i following the ladder of h and our algorithm can continue, or have hit γ and our algorithm terminates.

(2d) If γ runs along the top of diagram (b), then either l does not intersect γ (if it is the shorter hypergraph) or l does not intersect σ_i (if it is the longer hypergraph), contrary to our assumptions.

(2e) If γ runs along one side and then cuts across the diagram (see Figure 29), it necessarily separates a diagram (a) from a diagram (b) as in Inductive Case 2c. If l is the long hypergraph, then the argument from that case holds. So assume that l is the short hypergraph, and σ_i is dual to the other boundary edge dual to l. Then h belongs to g^-, contrary to our assumption.

(2f) Suppose γ runs along the non-boundary corner of one of the top 2-cells. We choose the top two boundary edges of the diagram for γ'.

(2fi) If l is the longer hypergraph, either boundary edge dual to l may be on σ_i, but in either case we want to show that we start at a corner of the diagram. The argument from Inductive Case 2(b)i shows we are not in the middle of the diagram.

Now we show that we end on σ_{i+1}. If not, then there is some other hypergraph k which belongs to g^- with the $k-\gamma$ intersection points surrounding the $l-\gamma$ intersection points. This is impossible by Base Case 1(b)ii.

(2fii) Assume h is the longer hypergraph. Consider the third hypergraph in the diagram, k. If k does not intersect γ, then it must intersect l twice to leave g^-, which makes a three-cube. We cannot glue a three cube diagram to this diagram by Lemma 4.10. Then k does intersect γ and does not intersect l, so k belongs to g^- and its γ intersection points are nested inside the $\gamma - l$ intersection points, which is impossible by Base Case 1(b)ii. Then h cannot be the longer hypergraph.

(2g) If γ runs along one top edge and turns down the middle of the diagram, then by geodesicity it necessarily turns the other direction to exit the diagram. If h is the longer hypergraph, this is identical to Base Case 1(b)iii. Then assume l is the longer hypergraph, so we know the non-γ edge dual to l corresponds to σ_i. By the argument in Case 2(b)i, we start in the corner we desire.

By the argument in Case 1(b)ii, no hypergraph belonging to g^- nests its γ-intersection points around l, so the entirety of σ_{i+1} is visible in the diagram, and we end on σ_{i+1}.

(2h) If γ runs down the side of the diagram, then l is necessarily the shorter hypergraph. The opposite edge dual to l is on σ_i, and we choose two edges that cut across the diagram for γ'.

(3) As in case 1 of the base step, l and h cannot form diagram (c).

Conclusion. The base step established that the first two letters of our chosen w' are in the form shown in Figure 4.16, and that we can continue on to the inductive step. The inductive step shows that either the next choice of two letters satisfies the proposition, or the algorithm terminates. Throughout we choose two particular edges to add to γ', and show that the end of these edges still lies on σ.

□

LEMMA 4.18. *(Freedom) Principal vertex stabilizers are free.*

PROOF. We choose from the collection of two-letter words from Proposition 4.16 for the basis of our free group, and show this group is isomorphic to a visible vertex stabilizer. Consider this collection of words as a set S generating a free group $F(S)$, and choose a minimally-sized generating set of two letter words. In S we label these as x_{ab} where ab is one of the two letter vertex stabilizing words. If any of our minimal generating set form diagram (a) with another word in the set, we consider those two generators equivalent. For example, if the words ab and cd form diagram (a), our group is $\langle x_{ab}, x_{cd}, \ldots \mid x_{ab} = x_{cd} \ldots \rangle$. Note S is finite as there are only finitely many options for such two letter words in the square model at $d < 1/3$. After adding the diagram (a) relations, our group is still finitely generated. The map from this group to our visible vertex stabilizer identifies each word in S with the corresponding two-letter word in G: so x_{ab} maps to the product ab. Certainly the map is surjective as we chose S from the full set of generators of the visible vertex stabilizers; we need now show the map is injective. We show if a relator r holds in the vertex stabilizer, it arises as a product of diagram (a) relators.

Suppose there is a relator r between vertex stabilizing elements with no ww^{-1} factors, where w is a two-letter word in our basis. Then r forms the boundary of a Van Kampen diagram Δ. Now freely reduce the diagram, so that the boundary is a reduced word r'. Note that reductions only occur between two different two-letter words, and form a new two-letter word which still stabilizes the vertex: for example, $ab^{-1} \cdot bc = ac$, so ac here still stabilizes the vertex.

Suppose these reductions result in a trivial r'. Then $|\partial\Delta| = 0$, so by $(*)$, we know $|\Delta| = 0$ and so Δ is trivial. Assume r contains no trivial subword which is a product of generators. Choose a particular algorithm for the free reduction to the trivial r', so each letter a is matched with a particular a^{-1} which cancels it out– call this instance of the inverse letter its *partner*. Consider the first 2-letter generator u that appears in r, and suppose $u = ab$. If the next letter in r is $c \neq b^{-1}$, then the partner c^{-1} occurs in r earlier than the partner b^{-1} does. An even number of letters must pass between c and its partner c^{-1}, which means the word that begins with c and ends with c^{-1} is of even length, freely reducible, and a product of generators, contrary to our assumption. Thus the next letter must be the partner b^{-1}, and so r is in the form $(ab)(b^{-1}a_3)(a_3^{-1}a_4)\ldots(a_n^{-1}a^{-1})$. That is, every letter is adjacent to its partner except for the first and last ones.

If any $a_i^{-1}a_{i+1}$ factor is equal to u besides the first, then the partner of the first letter is $a_i = a^{-1}$. Since every a_k is partnered with a_k^{-1} when $k \neq 1$, this means $ab \ldots a_{i-1}^{-1}a_i$ is a trivial subword which is a product of generators, contrary to our assumption. Similarly, if any $a_i^{-1}a_{i+1}$ is u^{-1}, then we again have a trivial subword in $b^{-1}a_3 \ldots a_{i-1}^{-1}a_i$. So u, u^{-1} do not appear again in r, so we can write u as a product of other 2-letter generators, which means we did not choose a minimal set for our generating set.

Thus we may assume that r' is not trivial, and consider the now-reduced Δ. Since hypergraphs are embedded, there must exist a corner 2-cell C where two hypergraphs h and l exit Δ along two adjacent edges on the boundary of Δ.

Now C contains letters from two stabilizing two-letter words, and by Proposition 4.16 each word forms one of two specified diagrams in Figure 21, with h and l dual to these words. Then h and l are each contained in one of these diagrams, either as a subdiagram of Δ or not.

Suppose h does not form a subdiagram of Δ, so two adjacent boundary edges of Δ are intersected by h. Since h is embedded, h must exit Δ at least two more times at other boundary edges. However, every two-letter word along $\partial\Delta$ forms such a diagram, so h must exit at adjacent edges and form another diagram outside Δ. While h may do this several times, it eventually makes a loop, as in Figure 30. But hypergraphs are embedded trees. So both h and l must form diagrams from Figure 21 as subdiagrams of Δ.

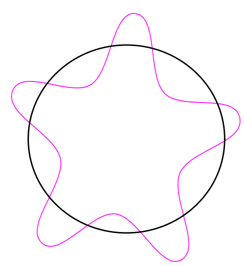

FIGURE 30. Forming a Figure 21 diagram outside Δ forces a loop in the hypergraph

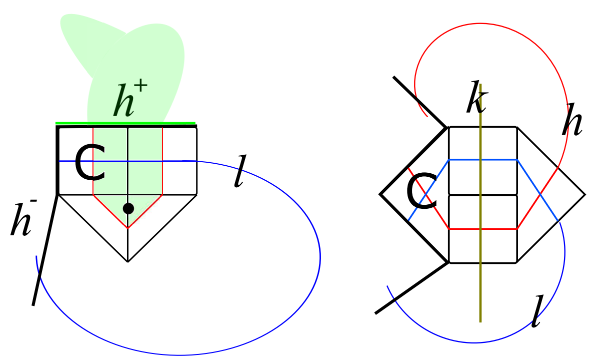

FIGURE 31. Last two cases in Lemma 4.18, with the heavy black lines indicating boundary edges of Δ. Left side indicates impossibility of including diagram (b) in a diagram from Figure 4.3 for l, right side indicates impossibility of including part of diagram (c) in one for k

We show that h and l must intersect twice in the diagram. As h separates Δ, we can call h^- the side that contains the boundary edge of C dual to l, and h^+ the other side, so l crosses from h^- to h^+ in C. We know h exits Δ at the next boundary cell of Δ, so there are no other boundary edges in h^+. As l needs to exit Δ again, we must have l exiting Δ in h^-, which means l crosses from h^+ back to h^- in Δ. Then h and l intersect twice, which mean that they too must form a diagram from Figure 6.

Suppose l and h form diagram (a). This gives a relation between the two words dual to h and l, and we chose only one of those two words as a basis element, so this is a trivial relation.

Suppose l and h form diagram (b), and h is the hypergraph that satisfies this diagram as in Proposition 4.16 (so the 2-letter word dual to h labels the green path in Figure 21). Then the diagram for l would necessarily contain the entire diagram (b) as a subdiagram, since it follows l and includes the disk collared by l and the two-letter word on the boundary dual to l. See the left side of Figure 31. None of the diagrams from Figure 21 contain diagram (b) as a subdiagram, so l cannot form one of the diagrams as a vertex-stabilizing word inside Δ, contradiction.

Suppose l and h form diagram (c). Then, as in Proposition 4.16 base case (1), the third hypergraph k in diagram c must intersect both l and h twice, as well as form a Figure 21 diagram along the boundary of Δ. See the right side of Figure 31. But none of the diagrams in Figure 21 contain diagram (c) as a subdiagram.

Then no such diagram can exist, and so no such r can exist. Thus any relator in the visible vertex stabilizer must include ww^{-1} factors, and so arise from a diagram (a) relator as in our group presentation. □

COROLLARY 4.19. *Vertex stabilizers are free.*

PROOF. By Proposition 4.15, nonprincipal vertex stabilizers are subgroups of principal vertex stabilizers. As those are free by Proposition 4.18, so are nonprincipal vertex stabilizers. □

PROOF OF THEOREM 4.1. By Corollary 3.15 in [**27**], these groups are hyperbolic. By Corollary 5.11 in [**27**], hypergraph stabilizers are quasiconvex. With these two facts we can apply Corollary C from [**13**] and say that nonprincipal vertex stabilizers are finitely generated. With Corollary 4.17 and Corollary 4.19, vertex stabilizers are free and finitely generated. □

This result is an application of Groves-Manning's result, which should prove fruitful in further research to show other groups are virtually special. For example, we hope that random groups in other models, namely the density model, may also be shown to be residually finite using this technique. There are two immediate potential paths to this method. Ollivier and Wise have shown that random groups at density $d < 1/6$ are residually finite, and we hope this technique can be applied to $d < 1/5$ as they show such hypergraphs are embedded but the resulting action on the cube complex is not proper [**30**]. By using a different method to create the cube complex, Mackay and Przytycki extend Ollivier and Wise's result to $d < 5/24$–it is possible this may be extended to $d < 1/4$ by applying this technique [**31**].

CHAPTER 5

Random Nilpotent Groups

Note: this chapter is excerpted from "Random Nilpotent Groups I" published in *International Mathematics Research Notices* as [**4**], joint with Matthew Cordes, Moon Duchin, Meng-Che (Turbo) Ho, and Ayla Sánchez. See Appendix for Publisher's Note. I added a few sentences throughout connecting it to the rest of this thesis.

1. Introduction and background

The results of this chapter are summarized as follows:

- In the remainder of this section, we establish a sequence of group theory and linear algebra lemmas for the following parts.
- In §2, the properties of non-backtracking random walk are described in order to deduce arithmetic statistics of Mal'cev coordinates—this is necessary to use the prior literature on random lattices and matrices.
- We survey the existing results from which ranks of random abelian groups can be calculated; a theorem of Magnus guarantees that the rank of a nilpotent group equals the rank of its abelianization. (§3)
- We give a complete description of one-relator quotients of the Heisenberg group, and compute the orders of finite quotients with any number of relators. (§4)
- Using a Freiheitssatz for nilpotent groups, we study the consequences of rank drop, and conclude that abelian groups occur with probability zero for $|R| \le m-2$, while they have positive probability for larger numbers of relators. Adding relators in a stochastic process drops the rank by at most one per new relator, with statistics for successive rank drop given by number-theoretic properties of the Mal'cev coordinates. (§5)
- We give a self-contained proof that a random nilpotent group is a.a.s. trivial exactly if $|R|$ is unbounded as a function of ℓ. We show how information about the nilpotent quotient lifts to information about the LCS of a standard (Gromov) random group and observe that standard random groups are *perfect* under the same conditions. (§6)

1.1. Nilpotent groups and Mal'cev coordinates. In the free group F_m of rank m, let

$$T_{j,m} = \left\{ \left[a_{i_1}, \dots, a_{i_j} \right] : 1 \le i_1, \dots, i_j \le m \right\}$$

be the set of all nested commutators with j arguments ranging over the generators. Then the *free s-step rank-m nilpotent group* is

$$N_{s,m} = F_m / \langle\!\langle T_{s+1,m} \rangle\!\rangle = \langle\, a_1, \dots, a_m \mid [a_{i_1}, \dots a_{i_{s+1}}] \text{ for all } i_j \,\rangle,$$

where $\langle\!\langle R \rangle\!\rangle$ denotes the normal closure of a set R when its ambient group is understood. This is the same expression we found in Section 2.3.

Just as all finitely-generated groups are quotients of (finite-rank) free groups, all finitely-generated nilpotent groups are quotients of free nilpotent groups. Note that the standard Heisenberg group $H(\mathbb{Z}) = \langle\, a, b \mid [a,b,a], [a,b,b] \,\rangle$ is realized as $N_{2,2}$. In the Heisenberg group, we will use the notation $c = [a,b]$, so that the center is $\langle c \rangle$.

The *lower central series* (LCS) for a s-step nilpotent group G is a sequence of subgroups inductively defined by $G_{k+1} = [G_k, G]$ which form a subnormal series

$$\{1\} = G_{s+1} \triangleleft \dots \triangleleft G_3 \triangleleft G_2 \triangleleft G_1 = G.$$

(The indexing is set up so that $[G_i, G_j] \subset G_{i+j}$.) For finitely generated nilpotent groups, this can always be refined to a *polycyclic series*

$$\{1\} = CG_{n+1} \triangleleft CG_n \triangleleft \dots \triangleleft CG_2 \triangleleft CG_1 = G$$

where each CG_i/CG_{i+1} is cyclic, so either $\mathbb{Z}$ or $\mathbb{Z}/n_i\mathbb{Z}$. The number of $\mathbb{Z}$ quotients in any polycyclic series for G is called the *Hirsch length* of G. From a polycyclic series we can form a generating set which supports a useful normal form for G. Make a choice of u_i in each CG_i so that $u_i CG_{i+1}$ generates CG_i/CG_{i+1}. An inductive argument shows that the set $\{u_1, \dots, u_n\}$ generates G. We call such a choice a *Mal'cev basis* for G, and we filter it as $\mathrm{MB}_1 \sqcup \cdots \sqcup \mathrm{MB}_s$, with MB_j consisting of basis elements belonging to $G_j \smallsetminus G_{j+1}$. Now if $u_i \in \mathrm{MB}_j$, let τ_i be the smallest value such that $u_i^{\tau_i} \in \mathrm{MB}_{j+1}$, putting $\tau_i = \infty$ if no such power exists. Then the Mal'cev normal form in G is as follows: every element $g \in G$ has a unique expression as $g = u_1^{t_1} \cdots u_n^{t_n}$, with integer exponents and $0 \le t_i \le \tau_i$ if $\tau_i < \infty$. Then the tuple of exponents $(t_1, \dots, t_n)$ gives a coordinate system on the group, called *Mal'cev coordinates*. We recall that $\mathrm{MB}_j \cup \cdots \cup \mathrm{MB}_s$ generates G_j for each j and that (by definition of s) the elements of MB_s are central.

We will construct a standard Mal'cev basis for free nilpotent groups $N_{s,m}$ as follows: let $\mathrm{MB}_1 = \{a_1, \dots, a_m\}$ be the basic generators, let $\mathrm{MB}_2 = \{b_{ij} := [a_i, a_j] : i < j\}$ be the basic commutators, and take each MB_j as a subset of $T_{j,m}$ consisting of some of the commutators from $[\mathrm{MB}_{j-1}, \mathrm{MB}_1]$. We note that $|\,\mathrm{MB}_2\,| = \binom{m}{2}$, and more generally the orders are given by the *necklace polynomials*

$$|\,\mathrm{MB}_j\,| = \frac{1}{j}\sum_{d|j} \mu(d) m^{j/d},$$

where μ is the Möbius function (see [**15**, Theorem 11.2.2]).

For example, the Heisenberg group $H(\mathbb{Z}) = N_{2,2}$ has the lower central series $\{1\} \lhd \mathbb{Z} \lhd H(\mathbb{Z})$, so its Hirsch length is 3. $H(\mathbb{Z})$ admits the Mal'cev basis a, b, c (with $a = a_1$, $b = a_2$, and c equal to their commutator), which supports a normal form $g = a^A b^B c^C$. The Mal'cev coordinates of a group element are the triple $(A, B, C) \in \mathbb{Z}^3$.

1.2. Group theory and linear algebra lemmas. In the free group $F_m = \langle a_1, \dots, a_m \rangle$, for any freely reduced $g \in F_m$, we define $A_i(g)$, called the *weight* of generator a_i in the word g, to be the exponent sum of a_i in g. Note that weights $A_1, \dots, A_m$ are well defined in the same way for the free nilpotent group $N_{s,m}$ for any s. We will let ab be the abelianization map of a group, so that $\mathrm{ab}(F_m) \cong \mathrm{ab}(N_{s,m}) \cong \mathbb{Z}^m$. Under this isomorphism, we can identify $\mathrm{ab}(g)$ with the vector $\mathbf{A}(g) := (A_1(g), \dots, A_m(g)) \in \mathbb{Z}^m$. If we have an automorphism ϕ on $N_{s,m}$, we write ϕ^{ab} for the induced map on $\mathbb{Z}^m$, which by construction satisfies $\mathrm{ab} \circ \phi = \phi^{\mathrm{ab}} \circ \mathrm{ab}$. Note that $\mathbf{A}(g)$ is also the MB_1 part of the Mal'cev coordinates for g, and we can similarly define a b-weight vector $\mathbf{B}(g)$ to be the MB_2 part, recording the exponents of the b_{ij} in the normal form.

To fix terminology: the *rank* of any finitely-generated group will be the minimum size of any generating set. Note this is different from the *dimension* of an abelian group, which we define by $\dim(\mathbb{Z}^d \times G_0) = d$ for any finite group G_0. (With this terminology, the Hirsch length of a nilpotent group G is the sum of the dimensions of its LCS quotients.) In any finitely-generated group, we say an element is *primitive* if it belongs to some basis (i.e., a generating set of minimum size). For a vector $w = (w_1, \dots, w_m) \in \mathbb{Z}^m$, we will write $\gcd(w)$ to denote the gcd of the entries. So a vector $w \in \mathbb{Z}^m$ is primitive iff $\gcd(w) = 1$. In this case we will say that the tuple $(w_1, \dots, w_m)$ has the *relatively prime property* or is RP. As we will see below, an element $g \in N_{s,m}$ is primitive in that nilpotent group if and only if its abelianization is primitive in $\mathbb{Z}^m$, i.e., if $\mathbf{A}(g)$ is RP. In free groups, there *exists* a primitive element with the same abelianization as g iff $\mathbf{A}(g)$ is RP.

The latter follows from a classic theorem of Nielsen [**24**].

THEOREM 5.1 (Nielsen primitivity theorem). *For every relatively prime pair of integers* (i,j), *there is a unique conjugacy class* $[g]$ *in the free group* $F_2 = \langle a,b\rangle$ *for which* $A(g) = i$, $B(g) = j$, *and* g *is primitive.*

COROLLARY 5.1 (Primitivity criterion in free groups). *There exists a primitive element* $g \in F_m$ *with* $A_i(g) = w_i$ *for* $i = 1,\dots,m$ *if and only if* $\gcd(w_1,\dots,w_m) = 1$.

PROOF. Let $w = (w_1,\dots,w_m)$. If $\gcd(w) \neq 1$, then the image of any g with those weights would not be primitive in the abelianization $\mathbb{Z}^m$, so no such g is primitive in F_m.

For the other direction we use induction on m, with the base case $m = 2$ established by Nielsen. Suppose there exists a primitive element of F_{m-1} with given weights $w_1,\dots,w_{m-1}$. For $\delta = \gcd(w_1,\dots,w_{m-1})$, we have $\gcd(\delta, w_m) = 1$. Let $\overline{w} = (\frac{w_1}{\delta},\cdots,\frac{w_{m-1}}{\delta})$. By the inductive hypothesis, there exists an element $\overline{g} \in F_{m-1}$ such that the weights of $\overline{g}$ are $\overline{w}$, and $\overline{g}$ can be extended to a basis $\{\overline{g}, h_2,\dots,h_{m-1}\}$ of F_{m-1}. Consider the free group $\langle \overline{g}, a_m\rangle \cong F_2$. Since $\gcd(\delta, w_m) = 1$, there exist $\hat{g},\hat{h}$ that generate this free group such that $\hat{g}$ has weights $A_{\overline{g}}(\hat{g}) = \delta$ and $A_m(\hat{g}) = w_m$ by Nielsen. Consequently, $A_i(\hat{g}) = w_i$. Then $\langle \hat{g},\hat{h},h_2,\cdots,h_{m-1}\rangle = \langle \overline{g},h_2,\dots,h_{m-1},a_m\rangle = F_m$, which shows that $\hat{g}$ is primitive, as desired. □

The criterion in free nilpotent groups easily follows from a powerful theorem due to Magnus.

THEOREM 5.2. [**21**, Lemma 5.9]*[Magnus lifting theorem] If* G *is nilpotent and* $S \subset G$ *is any set of elements such that* $\mathrm{ab}(S)$ *generates* $\mathrm{ab}(G)$, *then* S *generates* G.

Note that this implies that if G is nilpotent of rank m, then $G/\langle\langle g\rangle\rangle$ has rank at least $m-1$, because we can drop at most one dimension in the abelianization.

COROLLARY 5.2 (Primitivity criterion in free nilpotent groups). *An element* $g \in N_{s,m}$ *is primitive if and only if* $\mathbf{A}(g)$ *is primitive in* $\mathbb{Z}^m$.

Now we establish a sequence of lemmas for working with rank and primitivity. Recall that a,b are the basic generators of the Heisenberg group $H(\mathbb{Z})$ and that $c = [a,b]$ is the central letter.

LEMMA 5.3 (Heisenberg basis change). *For any integers* i,j, *there is an automorphism* ϕ *of* $H(\mathbb{Z}) = N_{2,2}$ *such that* $\phi(a^ib^jc^k) = b^dc^m$, *where* $d = \gcd(i,j)$ *and* $m = \frac{ij}{2d}(d-1)+k$.

In particular, if i,j *are relatively prime, then there is an automorphism* ϕ *of* $H(\mathbb{Z})$ *such that* $\phi(a^ib^j) = b$.

PROOF. Suppose $ri + sj = d = \gcd(i,j)$ for integers r, s and consider $\hat{a} = a^s b^{-r}$, $\hat{b} = a^{i/d} b^{j/d}$. We compute

$$[a^s b^{-r}, a^{i/d} b^{j/d}] = [a^s, b^{j/d}] \cdot [b^{-r}, a^{i/d}] = c^{(ri+sj)/d} = c.$$

If we set $\hat{c} = c$, we have $[\hat{a}, \hat{b}] = \hat{c}$ and $[\hat{c}, \hat{a}] = [\hat{c}, \hat{b}] = 1$, so $\langle \hat{a}, \hat{b} \rangle$ presents a quotient of the Heisenberg group. We need to check that it is the full group. Consider $h = (\hat{a})^{-i/d} (\hat{b})^s$. Writing h in terms of a, b, c, the a-weight of h is 0 and the b-weight is $(ri + sj)/d = 1$, so $h = bc^t$ for some t. But then $b = (\hat{a})^{-i/d} (\hat{b})^s (\hat{c})^{-t}$ and similarly $a = (\hat{a})^{j/d} (\hat{b})^r (\hat{c})^{-t'}$ for some t', so all of a, b, c can be expressed in terms of $\hat{a}, \hat{b}, \hat{c}$.

Finally,

$$(\hat{b})^d = (a^{i/d} b^{j/d})^d = a^i b^j c^{-\binom{d}{2} \frac{ij}{d^2}},$$

which gives the desired expression $a^i b^j c^k = (\hat{b})^d (\hat{c})^m$ from above. □

PROPOSITION 5.4 (General basis change). *Let $\delta = \gcd(A_1(g), \ldots, A_m(g))$ for any $g \in H = N_{s,m}$. Then there is an automorphism ϕ of H such that $\phi(g) = a_m^\delta \cdot h$ for some $h \in H_2$.*

PROOF. Let $w_i = A_i(g)$ for $i = 1, \ldots, m$ and let $r_i = w_i/\delta$, so that $\gcd(r_1, \ldots, r_m) = 1$. By Corollary 5.1, there exists a primitive element $x \in F_m$ with weights r_i. Let ϕ be a change of basis automorphism of F_m such that $\phi(x) = a_m$. This induces an automorphism of H, which we will also call ϕ.

By construction, x^δ and g have weight w. Since $\mathrm{ab}(x^\delta) = \mathrm{ab}(g) = w$, we must have $\phi^{\mathrm{ab}}(w) = \mathrm{ab}(\phi(x^\delta)) = \mathrm{ab}(\phi(g))$. Therefore $\phi(x^\delta)$ and $\phi(g)$ have the same weights.

Then $\mathrm{ab}(\phi(g)) = \mathrm{ab}(\phi(x^\delta)) = \mathrm{ab}(\phi(x)^\delta) = \mathrm{ab}(a_m^\delta)$, so $\phi(g)$ and a_m^δ only differ by commutators, i.e., $\phi(g) = a_m^\delta \cdot h$ for some $h \in H_2$. □

REMARK 5.5. *Given an abelian group $G = \mathbb{Z}^m / \langle R \rangle$, the classification of finitely-generated abelian groups provides that there are non-negative integers $d_1, \ldots, d_m$ with $d_m | d_{m-1} | \ldots | d_1$ such that $G \cong \bigoplus_{i=1}^m \mathbb{Z}/d_i\mathbb{Z}$. If G has dimension q and rank r, then $d_1 = \cdots = d_q = 0$, and $d_{r+1} = \cdots = d_m = 1$, so that*

$$G \cong \mathbb{Z}^q \times (\mathbb{Z}/d_{q+1}\mathbb{Z} \times \cdots \times \mathbb{Z}/d_r\mathbb{Z}).$$

Now consider a projection map $f : \mathbb{Z}^m \to \mathbb{Z}^m / K \cong \bigoplus_{i=1}^m \mathbb{Z}/d_i\mathbb{Z}$. We can choose a basis $e_1, \ldots, e_m$ of $\mathbb{Z}^m$ so that

$$K = \mathrm{span}\{d_1 e_1, \ldots, d_m e_m\} \cong \bigoplus_{i=1}^m d_i \mathbb{Z}.$$

Then since every element in K is a linear combination of $\{d_1 e_1, \dots, d_m e_m\}$ and $d_m | d_{m-1} | \dots | d_1$, we have that d_m divides all the coordinates of all the elements in K. Also $d_m e_m \in K$ with e_m being primitive.

LEMMA 5.6 (Criterion for existence of primitive vector). *Consider a set of r vectors in $\mathbb{Z}^m$, and let d be the gcd of the rm coordinate entries. Then there exists a vector in the span such that the gcd of its entries is d, and this is minimal among all vectors in the span.*

In particular, a set of r vectors in $\mathbb{Z}^m$ has a primitive vector in its span if and only if the gcd of the rm coordinate entries is 1.

PROOF. With d as above, let K be the $\mathbb{Z}$-span of the vectors and let

$$\gamma := \inf_{w \in K} \gcd(w).$$

One direction is clear: every vector in the span has every coordinate divisible by d, so $\gamma \geq d$. On the other hand $d_m e_m \in K$ and $\gcd(d_m e_m) = d_m$ because e_m is primitive. But d_m is a common divisor of all rm coordinates, and d is the greatest such, so $d_m \leq d$ and thus $\gamma \leq d$. □

LEMMA 5.7 (Killing a primitive element). *Let $H = N_{s,m}$ and let K be a normal subgroup of H. If* $\text{rank}(H/K) < m$ *then K contains a primitive element.*

PROOF. Since $\text{rank}(H/K) < m$, we also have $\text{rank}(\text{ab}(H/K)) < m$. Writing $\text{ab}(H/K) \cong \bigoplus_{i=1}^m \mathbb{Z}/d_i\mathbb{Z}$ as above, we have $d_m = 1$. By the previous lemma there is a primitive element in the kernel of the projection $\text{ab}(H) \to \text{ab}(H/K)$, and any preimage in K is still primitive (see Cor 5.2). □

LEMMA 5.8 (Linear algebra lemma). *Suppose $u_1, \dots, u_n \in \mathbb{Z}^m$ and suppose there exists a primitive vector v in their span. Then there exist $v_2, \dots, v_n$ such that* $\text{span}(v, v_2, \dots, v_n) = \text{span}(u_1, \dots, u_n)$.

PROOF. Since $v \in \text{span}(u_1, \dots, u_n)$, we can write $v = \alpha_1 u_1 + \dots + \alpha_n u_n$. Let $x \in \mathbb{Z}^n$ be the vector with coordinates α_i. Because $\gcd(v) = 1$, we have $\gcd(\alpha_i) = 1$, so x is primitive. Thus, we can complete x to a basis of $\mathbb{Z}^n$, say $\{x, x_2, \dots, x_n\}$. Then take $\begin{pmatrix} -v- \\ -v_2- \\ \vdots \\ -v_n- \end{pmatrix} = \begin{pmatrix} -x- \\ -x_2- \\ \vdots \\ -x_n- \end{pmatrix} \cdot \begin{pmatrix} -u_1- \\ -u_2- \\ \vdots \\ -u_n- \end{pmatrix}$.

Since $\begin{pmatrix} -x- \\ -x_2- \\ \vdots \\ -x_n- \end{pmatrix} \in SL_n(\mathbb{Z})$, it represents a change of basis matrix, so we have $\text{span}(v, v_2, \cdots, v_n) = \text{span}(u_1, \cdots, u_n)$, as needed. □

LEMMA 5.9 (String arithmetic). *Fix a free group $F = F_m$ on m generators and let R, S be arbitrary subsets, with normal closures $\langle\!\langle R\rangle\!\rangle, \langle\!\langle S\rangle\!\rangle$. Let $\phi : F \to F/\langle\!\langle R\rangle\!\rangle$ and $\psi : F \to F/\langle\!\langle S\rangle\!\rangle$ be the quotient homomorphisms. Then there exist canonical isomorphisms*

$$(F/\langle\!\langle R\rangle\!\rangle)\big/\langle\!\langle \phi(S)\rangle\!\rangle \cong F\big/\langle\!\langle R \cup S\rangle\!\rangle \cong (F/\langle\!\langle S\rangle\!\rangle)\big/\langle\!\langle \psi(R)\rangle\!\rangle$$

that are compatible with the underlying presentation (i.e., the projections from F commute with these isomorphisms).

PROOF. We will abuse notation by writing strings from F and interpreting them in the various quotients we are considering. Then if $G = \langle F \mid T\rangle \cong F/\langle\!\langle T\rangle\!\rangle$ is a quotient of F and U is a subset of F, we can write $\langle G \mid U\rangle$ to mean $F\big/\langle\!\langle T \cup U\rangle\!\rangle$ and can equally well write $\langle F \mid T, U\rangle$. Then the isomorphisms we need just record the fact that

$$\langle\, F \mid R, S\,\rangle = \langle\, F/\langle\!\langle R\rangle\!\rangle \mid S\,\rangle = \langle\, F/\langle\!\langle S\rangle\!\rangle \mid R\,\rangle. \quad \square$$

Because of this standard abuse of notation where we will variously interpret a string in $\{a_1, \ldots, a_m\}^{\pm}$ as belonging to F_m, $N_{s,m}$, or some other quotient group, we will use the symbol $=_G$ to denote equality in the group G when trying to emphasize the appropriate ambient group.

2. Random walk and arithmetic uniformity

In this section we survey properties of the simple nearest-neighbor random walk (SRW) and the non-backtracking random walk (NBSRW) on the integer lattice $\mathbb{Z}^m$, then deduce consequences for the distribution of Mal'cev coordinates for random relators in free nilpotent groups. For the standard basis $\{e_i\}$ of $\mathbb{Z}^m$, SRW is defined by giving the steps $\pm e_i$ equal probability $1/2m$, and NBSRW is similarly defined but with the added condition that the step $\pm e_i$ cannot be immediately followed by the step $\mp e_i$ (that is, a step can't undo the immediately previous step; equivalently, the position after k steps cannot equal the position after $k+2$ steps). Then for a random string w_ℓ of ℓ letters from $\{a_1, \ldots, a_m\}^{\pm}$, we have $w_\ell = \alpha_1\alpha_2\cdots\alpha_\ell$, where the α_i are i.i.d. random variables which equal each basic generator or its inverse with equal probability $1/2m$. The abelianization $X_\ell = \mathbf{A}(w_\ell)$ is a $\mathbb{Z}^m$-valued random variable corresponding to ℓ-step SRW. A random freely reduced string does not have an expression as a product of variables identically distributed under the same law, but if v_ℓ is such a string, its weight vector $Y_\ell = \mathbf{A}(v_\ell)$ is another $\mathbb{Z}^m$-valued random variable, this time corresponding to NBSRW.

It is well known that the distribution of endpoints for a simple random walk in $\mathbb{Z}^m$ converges to a multivariate Gaussian: if X_ℓ is again the random variable recording the endpoint after ℓ steps of simple random walk on $\mathbb{Z}^m$, and δ_t is the dilation in $\mathbb{R}^m$ sending $v \mapsto tv$, we have the central limit theorem:

$$\delta_{\frac{1}{\sqrt{\ell}}} X_\ell \longrightarrow \mathcal{N}(\mathbf{0}, \tfrac{1}{m} I).$$

This convergence notation for a vector-valued random variable V_ℓ and a multivariate normal $\mathcal{N}(\mu, \Sigma)$ means that V_ℓ converges in distribution to $AW + \mu$, where the vector μ is the mean, $\Sigma = AA^T$ is the covariance matrix, and W is a vector-valued random variable with i.i.d. entries drawn from a standard (univariate) Gaussian distribution $\mathcal{N}(0,1)$. In other words, this central limit theorem tells us that the individual entries of X_ℓ are asymptotically independent, Gaussian random variables with mean zero and expected magnitude $\sqrt{\ell}/m$. This is a special case of a much more general result of Wehn for Lie groups and can be found for instance in [**2**, Thm 1.3]. Fitzner and van der Hofstad derived a corresponding central limit theorem for NBSRW in [**8**]. Letting Y_ℓ be the $\mathbb{Z}^m$-valued random variable for ℓ-step NBSRW as before, they find that for $m \geq 2$,

$$\delta_{\frac{1}{\sqrt{\ell}}} Y_\ell \longrightarrow \mathcal{N}(\mathbf{0}, \tfrac{1}{m-1} I).$$

Note that the difference between the two statements records something intuitive: the non-backtracking walk still has mean zero, but the rule causes the expected size of the coordinates to be slightly higher than in the simple case; also, it blows up (as it should) in the case $m = 1$.

The setting of nilpotent groups is also well studied. To state the central limit theorem for free nilpotent groups, we take δ_t to be the similarity which scales each coordinate from MB_j by t^j, so that for instance in the Heisenberg group, $\delta_t(x,y,z) = (tx, ty, t^2 z)$.

PROPOSITION 5.10 (Distribution of Mal'cev coordinates). *Suppose* NB_ℓ *is an* $N_{s,m}$*-valued random variable chosen by non-backtracking simple random walk (NBSRW) on* $\{a_1, \ldots, a_m\}^\pm$ *for* ℓ *steps. Then the distribution on the Mal'cev coordinates is asymptotically normal:*

$$\delta_{\frac{1}{\sqrt{\ell}}} \mathrm{NB}_\ell \sim \mathcal{N}(\mathbf{0}, \Sigma).$$

For SRW, this is called a "simple corollary" of Wehn's theorem, [**2**, Theorem 3.11], where the only hypotheses are that the steps of the random walk are i.i.d. under a probability measure on $N_{s,m}$ that is centered, with finite second moment (in this case, the measure has finite support, so all moments are finite). Each Mal'cev coordinate is given by a polynomial formula in the a-weights of the step elements α_i (the polynomial for an MB_j coordinate has degree j), where the number of summands gets large as $\ell \to \infty$. Switching to NBSRW, it is still the case that NB_ℓ is a product of group elements whose a-weight vectors are independent and normally distributed, so their images under the same polynomials will be normally distributed as well, with only the covariance differing from the SRW case. We sketch a simple and self-contained argument for this in the $N_{2,2}$ non-backtracking case—that the third Mal'cev coordinate in $H(\mathbb{Z})$ is normally distributed— which we note is easily generalizable to the other $N_{s,m}$ with (only) considerable notational suffering. Without loss of generality, the sample path of the random walk is

$$g = a^{i_1} b^{j_1} a^{i_2} b^{j_2} \ldots a^{i_r} b^{j_r}$$

for some integers i_s, j_t summing to ℓ or $\ell - 1$, with all but possibly i_1 and j_r nonzero. After a certain number of steps, suppose the last letter so far was a. Then the next letter is either a, b, or b^{-1} with equal probability, so there is a 1/3 chance of repeating the same letter and a 2/3 chance of switching. This means that the i_s and j_t are (asymptotically independent) run-lengths of heads for a biased coin (Bernoulli trial) which lands heads with probability 1/3. On the other hand, r is half the number of tails flipped by that coin in ℓ (or $\ell - 1$) trials. In Mal'cev normal form,

$$g = a^{\sum i_s} b^{\sum j_t} c^{\sum_{t<s} i_s j_t}.$$

Thus the exponent of c is obtained by adding products of run-lengths together $\binom{r}{2}$ times, and general central limit theorems ensure that adding many independent and identically distributed (i.i.d.) random variables together tends to a normal distribution.

Our distribution statement has a particularly nice formulation in this Heisenberg case, where the third Mal'cev coordinate records the *signed area* enclosed between the x-axis and the path traced out by a word in $\{a, b\}^{\pm}$.

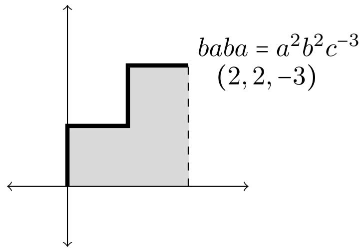

Corollary 5.11 (Area interpretation for Heisenberg case). *For the simple random walk on the plane, the signed area enclosed by the path is a normally distributed random variable.*

Next, we want to describe the effect of a group automorphism on the distribution of coordinates. Then we conclude this section by considering the distribution of coordinates in various $\mathbb{Z}/p\mathbb{Z}$.

Corollary 5.12 (Distributions induced by automorphisms). *If ϕ is an automorphism of $N_{s,m}$ and g is a random freely reduced word of length ℓ in $\{a_1,\dots,a_m\}^{\pm}$, then the Mal'cev coordinates of* $\mathrm{ab}(\phi(g))$ *are also normally distributed.*

Proof. The automorphism ϕ induces a change of basis on the copy of $\mathbb{Z}^m$ in the MB_1 coordinates, which is given by left-multiplication by a matrix $B \in SL_m(\mathbb{Z})$. Then $\phi_*(Y_\ell) \to \mathcal{N}(\mathbf{0}, B\Sigma B^T)$. □

Note that normality of the MB_j coordinates follows as well, as before: they are still described by sums of statistics coming from asymptotically independent Bernoulli trials, and only the coin bias has changed.

Relative primality of MB_1 coefficients turns out to be the key to studying the rank of quotient groups, so we will need some arithmetic lemmas.

LEMMA 5.13 (Arithmetic uniformity). *Let $A_{i,\ell}$ be the $\mathbb{Z}$-valued random variable given by the a_i-weight of a random freely reduced word of length ℓ in $\{a_1,\dots,a_m\}^{\pm}$, for $1 \le i \le m$. Let $\hat{A}_{i,\ell}$ equal $A_{i,\ell}$ with probability $\frac{1}{2}$ and $A_{i,\ell-1}$ with probability $\frac{1}{2}$. Then for fixed $m \ge 2$, fixed i, and $\ell \to \infty$,*

$$\forall k, n, \qquad \Pr\Big(\hat{A}_{i,\ell} \equiv k \mod n\Big) = \frac{1}{n} + o(1).$$

Furthermore, the distributions are independent in different coordinates:

$$\Pr\Big(\hat{A}_{i_1,\ell} \equiv k_1, \cdots \hat{A}_{i_s,\ell} \equiv k_s \mod n\Big) = \frac{1}{n^s} + o(1) \quad \textit{for } i_j \textit{ distinct}, s \le m.$$

In other words, the $\mathbb{Z}/n\mathbb{Z}$-valued random variables induced by the coordinate projections from NBSRW on MB_1 approach independent uniform distributions.

PROOF. First, consider SRW on $\mathbb{Z}^m$, which induces a lazy random walk (i.e., stays still with some probability, and moves forward or back with equal probabilities) on each coordinate. For odd n, the random walk is a Markov process on the finite graph given by the torus $(\mathbb{Z}/n\mathbb{Z})^m$, so $A_{i,\ell}$ approaches a uniform distribution (see [**6**, Chapter 3C]), and the result follows for $\hat{A}_{i,\ell}$. In fact, in that case the error term decays exponentially fast in ℓ:

$$\forall k, \forall n \le \sqrt{\ell}, \qquad \left|\Pr\Big(A_{i,\ell} \equiv k \pmod n\Big) - \frac{1}{n}\right| \le e^{-\pi^2 \ell/n^2}.$$

For $n = 2$ (and likewise for other even n) the construction of $\hat{A}$ corrects the parity bias, since $\hat{A}_{i,\ell}$ is now equally likely to have same parity as ℓ or the opposite parity. To make NBSRW into a Markov process, we must create a new state space where the states correspond to directed edges on the discrete torus, which encodes the one step of memory required to avoid backtracking. This new state space can itself be rendered as a homogeneous finite graph, and the result follows. Since the ith coordinate of the torus position corresponds to the $\hat{A}_{i,\ell}$ value, uniformity over the torus implies the independence and uniformity we need. □

COROLLARY 5.14 (Uniformity mod p). *The abelianization of a random freely reduced word in F_m has entries that are asymptotically uniformly distributed in $\mathbb{Z}/p\mathbb{Z}$ for each prime p, and the distribution mod p is independent of the distribution mod q for any distinct primes p, q.*

PROOF. For independence, consider $n = pq$ in the previous Lemma. □

COROLLARY 5.15 (Probability of primitivity). *For a random freely reduced word in F_m, the probability that it is primitive in abelianization tends to $1/\zeta(m)$, where ζ is the Riemann zeta function. In particular, for $m = 2$, the probability is $6/\pi^2$.*

PROOF. Using arithmetic uniformity, one derives a probability expression that agrees with $1/\zeta(m)$ by Euler's product formula for the zeta function; see [**16**]. □

REMARK 5.16 (Comparison of random models). *As we have seen, abelianizations of Gromov random groups are computed as cokernels of random matrices M whose columns are given by non-backtracking simple random walk on $\mathbb{Z}^m$.*

*Other models in the random abelian groups literature have somewhat different setup. Dunfield and Thurston [**7**] use a lazy random walk: ℓ letters are chosen uniformly from the $(2m+1)$ possibilities* *of $a_i^{\pm}$ and the identity letter, creating a word of length $\le \ell$, whose abelianization becomes a column of M.*

Kravchenko–Mazur–Petrenko [**20**] *and Wang–Stanley* [**36**] *use the standard "box" model: integer entries are drawn uniformly at random from $[-\ell, \ell]$, and asymptotics are calculated as $\ell \to \infty$. (This is the most classical way to randomize integers in number theory; see* [**16**].*)*

However, all of the arguments in all of these settings rely on arithmetic uniformity of coordinates mod p to calculate probabilities of relative primality, which is why the Riemann zeta function comes up repeatedly in the calculations. Since we have also established arithmetic uniformity for our setting in Corollary 5.14, the results achieved in these other models will carry over to our groups directly.

3. Preliminary facts about random nilpotent groups via abelianization

In this section we make a few observations relevant to the model of random nilpotent groups we study below. In particular, there has been substantial work on quotients of free abelian groups $\mathbb{Z}^m$ by random lattices, so it is important to understand the relationship between a random nilpotent group and its abelianization. Below, and throughout the paper, recall that probabilities are asymptotic as $\ell \to \infty$.

First, we record the simple observation that depth in the LCS is respected by homomorphisms.

LEMMA 5.17. *Let $\phi : G \to H$ be a surjective group homomorphism. Then $\phi(G_k) = H_k$ where G_k, H_k are the level-k subgroups in the respective lower central series.*

PROOF. Since ϕ is a homomorphism, depth-k commutators are mapped to depth-k commutators, i.e., $\phi(G_k) \subseteq H_k$. Let $h \in H_k$. Without loss of generality we can assume h is a single nested commutator $h = [w_1, \dots, w_k]$. By surjectivity of ϕ we can choose lifts $\overline{w_1}, \dots, \overline{w_k}$ of $w_1, \dots, w_k$. We see $[\overline{w_1}, \dots, \overline{w_k}] \in G_k$ and $\phi(G_k) \supseteq H_k$. □

To begin the consideration of ranks of random nilpotent groups, note that the Magnus lifting theorem (Theorem 5.2) tells us the rank of $N_{s,m}/\langle\langle R \rangle\rangle$ equals the rank of its abelianization $\mathbb{Z}^m/\langle R \rangle$, so we quickly deduce the probability of rank drop.

PROPOSITION 5.18 (Rank drop). *For a random r-relator nilpotent group $G = N_{s,m}/\langle\langle g_1, \dots, g_r \rangle\rangle$,*

$$\Pr(\text{rank}(G) < m) = \frac{1}{\zeta(rm)}.$$

PROOF. This follows directly from considering the existence of a primitive element in $\langle \mathrm{ab}(R) \rangle$. By Lemma 5.6, this occurs if and only if the rm entries are relatively prime, and by arithmetic uniformity (Lemma 5.13), this is computed by the Riemann zeta function, as in Corollary 5.15. □

Next we observe that a nilpotent group is trivial if and only if its abelianization (i.e., the corresponding $\mathbb{Z}^m$ quotient) is trivial, and more generally it is finite if and only if the abelianization is finite. Equivalence of triviality follows directly from the Magnus lifting theorem (Theorem 5.2). For the other claim, suppose the abelianization is finite. Then powers of all the images of a_i are trivial in the abelianization, so in the nilpotent group G there are finite powers $a_i^{r_i}$ in the commutator subgroup G_2. A simple inductive argument shows that every element of G_j has a finite power in G_{j+1}; for example, consider $b_{ij} \in G_2$. Since $[a_i^{r_i}, a_j] = b_{ij}^{r_i}$ is a commutator of elements from G_2 and G_1, it must be in G_3, as claimed. But then we can see that there are only finitely many distinct elements in the group by considering the Mal'cev normal form

$$g = u_1^* u_2^* \dots u_r^*$$

and noting that each exponent can take only finitely many values. Since the rank of a nilpotent group equals that of its abelianization (by Theorem 5.2 again), it is also true that a nilpotent group is cyclic if and only if its abelianization is cyclic.

We introduce the term *balanced* for groups presented with the number of relators equal to the number of generators, so that it applies to models of random groups $F_m/\langle\langle R \rangle\rangle$, random nilpotent groups $N_{s,m}/\langle\langle R \rangle\rangle$, or random abelian groups $\mathbb{Z}^m/\langle R \rangle$, where $|R| = m$, the rank of the seed group. We will correspondingly use the terms *nearly-balanced* for $|R| = m-1$, and *underbalanced* or *overbalanced* in the $|R| < m-1$ and $|R| > m$ cases, respectively.

Then it is very easy to see that nearly-balanced (and thus underbalanced) groups are a.a.s. infinite, while balanced (and thus overbalanced) groups are a.a.s. finite because m random integer vectors are $\mathbb{R}$-linearly independent with probability one. However, is also easy to see that if $|R|$ is held constant, no matter how large, then there is a nonzero probability that the group is nontrivial (because, for example, all the a-weights could be even).

To set up the statement of the next lemma, let $Z(m) := \zeta(2) \dots \zeta(m)$ and

$$P(m) := \prod_{\text{primes } p} \left(1 + \frac{1/p - 1/p^m}{p-1}\right).$$

Recall from Remark 5.16 that we can quote the distribution results of [**7**],[**20**],[**36**] because of the common feature of arithmetic uniformity.

LEMMA 5.19 (Cyclic quotients of abelian groups). *The probability that the quotient of* $\mathbb{Z}^m$ *by* $m-1$ *random vectors is cyclic is* $1/Z(m)$. *With* m *random vectors, the probability is* $P(m)/Z(m)$.

These facts, particularly the first, can readily be derived "by hand," but can also be computed using Dunfield–Thurston [**7**] as follows: their generating functions give expressions for the probability that i random vectors with $\mathbb{Z}/p\mathbb{Z}$ entries generate a subgroup of rank j, and the product over primes of the probability that the $\mathbb{Z}/p\mathbb{Z}$ reduction has rank $\geq m-1$ produces the probability of a cyclic quotient over $\mathbb{Z}$.

The latter fact appears directly in Wang–Stanley [**36**] as Theorem 4.9(i). We note that corresponding facts for higher-rank quotients could also be derived from either of these two papers, but the expressions have successively less succinct forms.

COROLLARY 5.20 (Explicit probabilities for cyclic quotients). *For balanced and nearly-balanced presentations, the probability that a random abelian group or a random nilpotent group is cyclic is a strictly decreasing function of* m *which converges as* $m \to \infty$.

In the balanced case, the limiting value is a well-known number-theoretic invariant. Values are estimated in the table below.

The convergence for both cases is proved in [**36**, Theorem 4.9] as a corollary of the more general statement about the Smith normal form of a random not-necessarily-square matrix M, which is an expression $A = SMT$ for invertible S,T in which A has all zero entries except possibly its diagonal entries $a_{ii} = \alpha_i$. These α_i are then the abelian invariants for the quotient of $\mathbb{Z}^m$ by the column span of M (that is, they are the d_i from Remark 5.5 but with opposite indexing, $d_i = \alpha_{m+1-i}$). The rank of the quotient is the number of these that are not equal to 1.

The probabilities of cyclic groups among balanced and nearly-balanced quotients of free abelian groups and therefore also for random nilpotent groups are approximated below. Values in the table are truncated (not rounded) at four digits.

Pr(cyclic)	$m=2$	$m=3$	$m=4$	$m=10$	$m=100$	$m=1000$	$m\to\infty$
$\lvert R\rvert = m-1$	.6079	.5057	.4672	.4361	.4357	.4357	.4357
$\lvert R\rvert = m$	.9239	.8842	.8651	.8469	.8469	.8469	.8469

Computing the probability of a trivial quotient with r relators is equivalent to the probability that r random vectors generate $\mathbb{Z}^m$.

LEMMA 5.21 (Explicit probability of trivial quotients). *For $r > m$,*

$$\Pr\left(\mathbb{Z}^m/\langle v_1,\dots,v_r\rangle = 0\right) = \frac{1}{\zeta(r-m+1)\cdots\zeta(r)}.$$

This is a rephrasing of [**20**, Corollary 3.6] and [**36**, Theorem 4.8].

REMARK 5.22. *From the description of Smith normal form, we get a symmetry in r and m, namely*

$$\Pr\left(\operatorname{rank}\left(\mathbb{Z}^m/\langle v_1,\dots,v_r\rangle\right) = m-k\right) = \Pr\left(\operatorname{rank}\left(\mathbb{Z}^r/\langle v_1,\dots,v_m\rangle\right) = r-k\right) \quad \forall 1 \le k \le \min(r,m)$$

just by the observation that the transpose of the normal form expression has the same invariants. For example, applying duality to Lemma 5.19 and reindexing, we immediately obtain, as in Lemma 5.21,

$$\Pr\left(\mathbb{Z}^m/\langle v_1,\dots,v_{m+1}\rangle = 0\right) = \frac{1}{Z(m+1)} = \frac{1}{\zeta(2)\cdots\zeta(m+1)}.$$

4. Quotients of the Heisenberg group

We will classify all $G := H(\mathbb{Z})/\langle\langle g\rangle\rangle$ for single relators g, up to isomorphism. As above, we write a,b for the generators of $H(\mathbb{Z})$, and $c = [a,b]$. With this notation, $H(\mathbb{Z})$ can be written as a semidirect product $\mathbb{Z}^2 \rtimes \mathbb{Z}$ via $\langle b,c\rangle \rtimes \langle a\rangle$ with the action of $\mathbb{Z}$ on $\mathbb{Z}^2$ given by $ba = abc^{-1}$, $ca = ac$.

THEOREM 5.3 (Classification of one-relator Heisenberg quotients). *Suppose $g = a^i b^j c^k \neq 1$. Let $d = \gcd(i,j)$, let $m = \frac{ij}{2d}(d-1)+k$ as in Lemma 5.3, and let $D = \gcd(d,m)$. Then*

$$G := H(\mathbb{Z})/\langle\langle g\rangle\rangle \cong \begin{cases} (\mathbb{Z}\times\mathbb{Z}/k\mathbb{Z}) \rtimes \mathbb{Z}, & \text{if } i=j=0; \\ (\mathbb{Z}/\frac{d^2}{D}\mathbb{Z} \times \mathbb{Z}/D\mathbb{Z}) \rtimes \mathbb{Z}, & \text{else,} \end{cases}$$

with the convention that $\mathbb{Z}/0\mathbb{Z} = \mathbb{Z}$ and $\mathbb{Z}/1\mathbb{Z} = \{1\}$. In particular, G is abelian if and only if $g = c^{\pm1}$ or $\gcd(i,j) = 1$; otherwise, it has step two. Furthermore, unless g is a power of c (the $i=j=0$ case), the quotient group is virtually cyclic.

Note that this theorem is exact, not probabilistic.

REMARK 5.23 (Baumslag-Solitar case). *The* Baumslag-Solitar groups *are a famous class of groups given by the presentations $BS(p,q) = \langle a,b \mid ab^p a^{-1} = b^q\rangle$ for various p,q. For the Heisenberg quotients* *as described above, we will refer to* $D = 1$ *as the Baumslag-Solitar case, because in that case* $sd - tm = 1$ *has solutions in* s, t*, and one easily checks that the group is presented as*

$$G = \langle a, b \mid [a, b] = b^{td}, \quad b^{d^2} = 1 \rangle \cong BS(1, 1 + td) / \langle\!\langle b^{d^2} \rangle\!\rangle,$$

a 1-relator quotient of a solvable Baumslag-Solitar group $BS(1, q)$*.*

Examples:

(1) if $g = a$, then $G = \mathbb{Z}$.
(2) if $g = c$, then $G = \mathbb{Z}^2$.
(3) if $g = c^2$, then $G = (\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}) \rtimes \mathbb{Z}$.
(4) if $g = a^{20}b^{28}c^{16}$, we have $d = 4$, $m = 226$, $D = 2$, so we get

$$G = \left(\mathbb{Z}^2 / \langle \left(\begin{smallmatrix} 4 \\ 226 \end{smallmatrix}\right), \left(\begin{smallmatrix} 0 \\ 4 \end{smallmatrix}\right) \rangle\right) \rtimes \mathbb{Z} \cong \left(\mathbb{Z}^2 / \langle \left(\begin{smallmatrix} 4 \\ 2 \end{smallmatrix}\right), \left(\begin{smallmatrix} 0 \\ 4 \end{smallmatrix}\right) \rangle\right) \rtimes \mathbb{Z} \cong (\mathbb{Z}/8\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}) \rtimes \mathbb{Z}.$$

(5) if $g = a^2b^2c^2$, we have $d = 2$, $m = 3$, $D = 1$. In this case, $b^4 =_G c^2 =_G 1$ and the quotient group is isomorphic to $\mathbb{Z}/4\mathbb{Z} \rtimes \mathbb{Z}$ with the action given by $aba^{-1} = b^3$. This is a two-step-nilpotent quotient of the Baumslag-Solitar group $BS(1, 3)$ by introducing the relation $b^4 = 1$.

We see that the quotient group G collapses down to $\mathbb{Z}$ precisely if $\gcd(i, j) = 1$. Namely, $c =_G 1$ in that case, so we have a quotient of $\mathbb{Z}^2$ by a primitive vector.

Corollary 5.24. *For one-relator quotients of the Heisenberg group,* $G = N_{2,2}/\langle\!\langle g \rangle\!\rangle$*,*

$$\Pr(G \cong \mathbb{Z}) = \frac{6}{\pi^2} \approx 60.8\% \ ; \qquad \Pr(G \text{ step } 2, \text{ rank } 2) = 1 - \frac{6}{\pi^2}.$$

Of course, if $g = c$, we have $\mathbb{Z}^2$, but this event occurs with probability zero. If $\gcd(i, j) \neq 1$, then G is two-step (thus non-abelian) and has torsion.

Proof of theorem. First, the $(i, j) = (0, 0)$ case is very straightforward: then $g = c^k$ and the desired expression for G follows.

Below, we assume $(i, j) \neq (0, 0)$, and by Lemma 5.3, without loss of generality, we will write $g = b^d c^m$.

Consider the normal closure of b, which is $\langle\!\langle b \rangle\!\rangle = \langle b, c \rangle$. This intersects trivially with $\langle a \rangle$, and $G = \langle\!\langle b \rangle\!\rangle \langle a \rangle$. Thus $G = \langle b, c \rangle \rtimes \langle a \rangle$.

Now in $H(\mathbb{Z})$, we compute $\langle\!\langle g \rangle\!\rangle = \langle b^d c^m, c^d \rangle \subset \langle b, c \rangle$. Thus

$$\langle b, c \rangle \cong \mathbb{Z}^2 \big/ \langle \left(\begin{smallmatrix} d \\ m \end{smallmatrix}\right), \left(\begin{smallmatrix} 0 \\ d \end{smallmatrix}\right) \rangle .$$

We have the semidirect product structure $G \cong \mathbb{Z}^2 \big/ \langle \left(\begin{smallmatrix} d \\ m \end{smallmatrix}\right), \left(\begin{smallmatrix} 0 \\ d \end{smallmatrix}\right) \rangle \rtimes \mathbb{Z}$, where the action sends $\left(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right) \mapsto \left(\begin{smallmatrix} 1 \\ 1 \end{smallmatrix}\right)$ and fixes $\left(\begin{smallmatrix} 0 \\ 1 \end{smallmatrix}\right)$. Note that c has order d in G, and a simple calculation verifies that b has order d^2/D, where $D = \gcd(d, m)$. If we are willing to lose track of the action and just write the group up to isomorphism, then we can perform both row and column operations on $\left[\begin{smallmatrix} d & 0 \\ m & d \end{smallmatrix}\right]$ to get $\left[\begin{smallmatrix} d^2/D & 0 \\ 0 & D \end{smallmatrix}\right]$, which produces the desired expression. □

In fact, we can say something about quotients of $H(\mathbb{Z})$ with arbitrary numbers of relators. First let us define the *K-factor* $K(R)$ of a relator set $R = \{g_1, \ldots, g_r\}$, where relator g_1 has the Mal'cev coordinates (i_1, j_1, k_1), and similarly for $g_2, \ldots, g_r$. Let $M = \begin{pmatrix} i_1 & i_2 & \ldots & i_r \\ j_1 & j_2 & \ldots & j_r \end{pmatrix}$ and suppose its nullity (the dimension of its kernel) is q. Then let W be a kernel matrix of M, i.e., an $r \times q$ matrix with rank q such that $MW = \mathbf{0}$. (Note that if R is a random relator set, then $q = r - 2$, since the rank of M is 2 with probability one.) Let $k = (k_1, \ldots, k_r)$ be the vector of c-coordinates of relators, so that $kW \in \mathbb{Z}^q$. Then $K(R) := \gcd(kW)$ is defined to be the gcd of those q integers.

THEOREM 5.4 (Orders of Heisenberg quotients). *Consider the group $G = H(\mathbb{Z})/\langle\!\langle g_1, \ldots, g_r \rangle\!\rangle$, where relator g_1 has the Mal'cev coordinates (i_1, j_1, k_1), and similarly for $g_2, \ldots, g_r$. Let $d = \gcd(i_1, j_1, \ldots, i_r, j_r)$; let Δ be the co-area of the lattice spanned by the $\left(\begin{smallmatrix} i_\alpha \\ j_\alpha \end{smallmatrix}\right)$ in $\mathbb{Z}^2$; and let $K = K(R)$ be the K-factor defined above. Then c has order $\gamma = \gcd(d, K)$ in G and $|G| = \Delta \cdot \gamma$.*

PROOF. Clearly Δ is the order of $\mathrm{ab}(G) = G/\langle c \rangle$. So to compute the order of G, we just need to show that the order of c in G is γ. Consider for which n we can have $c^n \in \langle\!\langle g_1, \ldots, g_r \rangle\!\rangle$, i.e.,

$$c^n = \prod_{\alpha=1}^{N} w_\alpha \; g_\alpha^{\epsilon_\alpha} \; w_\alpha^{-1}$$

for arbitrary words w_α and integers ϵ_α. First note that all commutators $[w, g_\alpha]$ are of this form, and that by letting $w = a$ or b, these commutators can equal c^{i_α} or c^{j_α} for any α, so n can be an arbitrary multiple of d.

Next, consider the expression in full generality and note that $\mathbf{A}(c^n) = \left(\begin{smallmatrix} 0 \\ 0 \end{smallmatrix}\right)$. Conjugation preserves weights, so $\mathbf{A}(w_\alpha g_\alpha^{\epsilon_\alpha} w_\alpha^{-1}) = \mathbf{A}(g_\alpha^{\epsilon_\alpha}) = \epsilon_\alpha \mathbf{A}(g_\alpha) = \epsilon_\alpha \left(\begin{smallmatrix} i_\alpha \\ j_\alpha \end{smallmatrix}\right)$. To get the two sides to be equal in abelianization, the ϵ_α must record a linear dependency in the $\binom{i_\alpha}{j_\alpha}$. Finally we compute

$$n = \sum_\alpha \epsilon_\alpha(x_\alpha j_\alpha - y_\alpha i_\alpha) + \sum_{\alpha<\beta} \epsilon_\alpha \epsilon_\beta i_\beta j_\alpha - \sum_\alpha i_\alpha j_\alpha \frac{\epsilon_\alpha(\epsilon_\alpha - 1)}{2} + \sum_\alpha \epsilon_\alpha k_\alpha,$$

where $\binom{x_\alpha}{y_\alpha} = \mathbf{A}(w_\alpha)$. We can observe that each of the first three terms is a multiple of d and the fourth term is an arbitrary integer multiple of K. (To see this, note that the column span of W is exactly the space of linear dependencies in the $\mathbf{A}(g_\alpha)$, so $\sum \epsilon_\alpha k_\alpha$ is a scalar product of the k vector with something in that column span, and is therefore a multiple of K.) Thus n can be any integer combination of d and K, as we needed to prove. □

5. Rank drop

First, we establish that adding a single relator to a (sufficiently complicated) free nilpotent group does not drop the nilpotency class; the rank drops by one if the relator is primitive in abelianization and it stays the same otherwise. Furthermore, a single relator never drops the step unless the starting rank was two. This is a nilpotent version of Magnus's famous *Freiheitssatz* (freeness theorem) for free groups [**21**, Theorem 4.10].

THEOREM 5.5 (Nilpotent *Freiheitssatz*). *For any $g \in N_{s,m}$ with $s \geq 2, m \geq 3$, there is an injective homomorphism*

$$N_{s,m-1} \hookrightarrow N_{s,m}/\langle\langle g \rangle\rangle.$$

This is an isomorphism if and only if $\gcd(A_1(g), \ldots, A_m(g)) = 1$.

If $m = 2$ the result holds with $\mathbb{Z} \hookrightarrow N_{s,2}/\langle\langle g \rangle\rangle$.

PROOF. Romanovskii's 1971 theorem [**33**, Theorem 1] does most of this. In our language, the theorem says that if $A_m(g) \neq 0$, then $\langle a_1, \ldots, a_{m-1} \rangle$ is a copy of $N_{s,m-1}$. This establishes the needed injection except in the case $g \in [N_{s,m}, N_{s,m}]$, where $\mathbf{A}(g)$ is the zero vector. In the $m = 2$ case, any such $N_{s,2}/\langle\langle g \rangle\rangle$ has abelianization $\mathbb{Z}^2$, so the statement holds. For $m > 2$, one can apply an automorphism so that g is spelled with only commutators involving a_m. Even killing all such commutators does not drop the nilpotency class because $m > 2$ ensures that there are some Mal'cev generators spelled without a_m in each level. Thus in this case $\langle a_1, \ldots, a_{m-1} \rangle \cong N_{s,m-1}$ still embeds.

It is easy to see that if g is non-primitive in abelianization, then the rank of $\mathrm{ab}(N_{s,m}/\langle\langle g \rangle\rangle)$ is m, and so the quotient nilpotent group has rank m as well. However, the image of Romanovskii's map has rank $m - 1$, so it is not a surjection.

On the other hand, suppose $\mathrm{ab}(g)$ is a primitive vector. Then the rank of the abelianized quotient is $m-1$, and by Magnus's theorem (Theorem 5.2) the rank of the nilpotent quotient is the same. The group $G = N_{s,m}/\langle\langle g\rangle\rangle$ is therefore realizable as a quotient of that copy of $N_{s,m-1}$. Since the lower central series of $N_{s,m-1}$ has all free abelian quotients, any proper quotient would have smaller Hirsch length, and this contradicts Romanovskii's injection. Thus relative primality implies that the injection is an isomorphism. □

Now we can use rank drop to analyze the probability of an abelian quotient for a free nilpotent group in the underbalanced, nearly balanced, and balanced cases (i.e., cases with the number of relators at most the rank).

LEMMA 5.25 (Abelian implies rank drop for up to m relators). *Let $G = N_{s,m}/\langle\langle R\rangle\rangle$, where $R = \{g_1, \dots, g_r\}$ is a set of $r \le m$ random relators. Suppose $s \ge 2$ and $m \ge 2$. Then*

$$\Pr(G \text{ abelian } \mid \mathrm{rank}(G) = m) = 0.$$

PROOF. Suppose that $\mathrm{rank}(G) = m$ and G is abelian. We use the form of the classification of abelian groups (Remark 5.5) in which $G \cong \oplus_{i=1}^m \mathbb{Z}/d_i\mathbb{Z}$, where $d_m \mid \dots \mid d_1$ so that $d_1 = \dots = d_q = 0$ for $q = \dim(G)$, and we write $\langle\langle \mathrm{ab}(R)\rangle\rangle = \langle d_1e_1, \dots, d_me_m\rangle$ for a basis $\{e_i\}$ of $\mathbb{Z}^m$. Since $\mathrm{rank}(G) = m$, we can assume no $d_i = 1$. We can lift the basis $\{e_i\}$ of $\mathbb{Z}^m$ to a generating set $\{a_i\}$ of $N_{s,m}$ by Magnus (Theorem 5.2). Note that the exponent of each generator in each relator is a multiple of d_m.

Next we show that we cannot kill a commutator in G without dropping rank. Let $b_1 = [a_1, a_m]$. We claim that $b_1 \notin \langle\langle g_1, \dots, g_r\rangle\rangle$. To do so, we compute an arbitrary element

$$\prod_{\alpha}^{n} w_\alpha g_\alpha^{\epsilon_\alpha} w_\alpha^{-1} \in \langle\langle g_1, \dots, g_r\rangle\rangle.$$

Conjugation preserves weights, so $\mathbf{A}(w_\alpha g_{i_\alpha}^{\epsilon_\alpha} w_\alpha^{-1}) = \mathbf{A}(g_{i_\alpha}^{\epsilon_\alpha}) = \epsilon_\alpha \mathbf{A}(g_{i_\alpha})$. If the product is equal to b_1, then its a-weights are all zero. Now consider the b-weights. For the product, the b-weights are the combination of the b-weights of the g_α, modified by amounts created by commutation. However, since all the a-exponents of all the g_α are multiples of d_m, we get

$$\sum \epsilon_i \mathbf{A}(g_i) = \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \qquad \sum \epsilon_i \mathbf{B}(g_i) \equiv \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix} \pmod{d_m},$$

where each ϵ_i is the sum of the ϵ_α corresponding to g_i. The second expression ensures that the ϵ_i are not all zero, so the first equality is a linear dependence in the $\mathbf{A}(g_i)$, which has probability zero since $r \le m$. □

THEOREM 5.6. *(Underbalanced quotients are not abelian) Let* $G = N_{s,m}/\langle\langle R\rangle\rangle$*, where* $R = \{g_1, \dots, g_r\}$ *is a set of* $r \leq m-2$ *random relators* g_i*. Then*

$$\Pr(G \text{ abelian}) = 0.$$

PROOF. Suppose that G is abelian, and consider elements of G as vectors in $\mathbb{Z}^m$ via the abelianization map on $N_{s,m}$; in this way we get vectors $v_1 = \mathbf{A}(g_1), \dots, v_r = \mathbf{A}(g_r)$. From the previous result we may assume $\operatorname{rank}(G) < m$. By Lemma 5.6, we can find a primitive vector w as a linear combination of the v_i. Then we apply the linear algebra lemma (Lemma 5.8) to extend w appropriately so that $\operatorname{span}(v_1, \dots, v_r) = \operatorname{span}(w, w_2, \dots, w_r)$. We can find a series of elementary row operations (switching, multiplication by -1, or addition) to get $(w, w_2, \dots, w_r)$ from $(v_1, \dots, v_r)$, and we lift these operations to elementary Nielsen transformations (switching, inverse, or multiplication, respectively) in $N_{s,m}$ to get $(g', g_2', \dots, g_r')$ from $(g_1, \dots, g_r)$. Note that Nielsen transformations on a set of group elements preserve the subgroup they generate, so also preserve normal closure. This lets us define $R' = \{g', g_2', \dots, g_r'\}$ with $\langle\langle R'\rangle\rangle = \langle\langle R\rangle\rangle$. Since g' has a weight vector w whose coordinates are relatively prime, the Freiheitssatz (Theorem 5.5) ensures that $N_{s,m}/\langle\langle g'\rangle\rangle \cong N_{s,m-1}$. Thus we have $G = N_{s,m-1}/\langle\langle g_2', \dots, g_r'\rangle\rangle$.

If $r \leq m-2$, then iterating this argument $r-1$ times gives $G \cong N_{s,m-r+1}/\langle\langle g_r\rangle\rangle$ for some new g_r, and $m - r + 1 \geq 3$. Then we can apply Theorem 5.5 to conclude that this quotient is not abelian, because its nilpotency class is $s > 1$. ☐

PROPOSITION 5.26 (Cyclic quotients). *If* $|R| = m-1$ *or* $|R| = m$*, then abelian implies cyclic:*

$$\Pr(G \text{ cyclic} \mid G \text{ abelian}) = 1.$$

PROOF. Running the proof as above, we iterate the reduction $m-2$ times to obtain $G \cong N_{s,2}/\langle\langle g\rangle\rangle$ or $N_{s,2}/\langle\langle g, g'\rangle\rangle$.

If g (or any element of $\langle\langle g, g'\rangle\rangle$) is primitive, then G is isomorphic to $\mathbb{Z}$ or a quotient of $\mathbb{Z}$, i.e., G is cyclic.

Otherwise, note that $N := N_{s,2}$ has the Heisenberg group as a quotient ($H(\mathbb{Z}) = N_1/N_3$). If G is abelian, then the corresponding quotient of $H(\mathbb{Z})$ is abelian. In the non-primitive case, this can only occur if $c \in \langle\langle g, g'\rangle\rangle$, which (as in the proof of Lemma 5.25) implies $\mathbf{A}(g) = (0,0)$ (or a linear dependency between $\mathbf{A}(g)$ and $\mathbf{A}(g')$). But by Corollary 5.12, the changes of basis do not affect the probability of linear dependency, so this has probability zero. ☐

Corollary 5.27. *For nearly-balanced and balanced models, the probability that a random nilpotent group is abelian equals the probability that it is cyclic.*

We reprise the table from §3, recalling that values are truncated at four digits.

Pr(abelian)	$m=2$	$m=3$	$m=4$	$m=10$	$m=100$	$m=1000$	$m\to\infty$
$\lvert R\rvert = m-1$	.6079	.5057	.4672	.4361	.4357	.4357	.4357
$\lvert R\rvert = m$	.9239	.8842	.8651	.8469	.8469	.8469	.8469

Corollary 5.28 (Abelian one-relator). *For any step $s\ge 2$,*

$$\Pr(N_{s,m}/\langle\!\langle g\rangle\!\rangle \text{ is abelian}) = \begin{cases} 6/\pi^2, & m=2 \\ 0, & m\ge 3. \end{cases}$$

Note that these last two statements agree for $m=2$, $|R| = m-1 = 1$.

6. Trivializing and perfecting random groups

In this final section, we first find the threshold for collapse of a random nilpotent group, using the abelianization. Then we will prove a statement lifting facts about random nilpotent group to facts about the LCS of classical random groups, deducing that *random groups are perfect* with exactly the same threshold again.

Recall that $T_{j,m} = \{[a_{i_1},\dots,a_{i_j}] : 1\le i_1,\dots,i_j\le m\}$ contains the basic nested commutators with j arguments. In this section we fix m and write F for the free group, so we can write F_i for the groups in its lower central series. Similarly we write N for $N_{s,m}$ (when s is understood), and T_j for $T_{j,m}$. Note that $\langle\!\langle T_j\rangle\!\rangle = F_j$, so $N = F/F_{s+1}$.

For a random relator set $R\subset F$, we write $\Gamma = F/\langle\!\langle R\rangle\!\rangle$, $G = N/\langle\!\langle R\rangle\!\rangle$, and $H = \mathbb{Z}^m/\langle R\rangle = \mathrm{ab}(\Gamma) = \mathrm{ab}(G)$, using the abuse of notation from Lemma 5.9 and treating R as a set of strings from F to be identified with its image in N or $\mathbb{Z}^m$. In all cases, R is chosen uniformly from freely reduced words of length ℓ or $\ell-1$ in F.

First we need a result describing the divisibility properties of the determinants of matrices whose columns record the coordinates of random relators.

LEMMA 5.29 (Arithmetic distribution of determinants). *For a fixed rank $m\ge 1$ and prime p, let M be an $m\times m$ random matrix whose entries are independently uniformly distributed in $\mathbb{Z}/p\mathbb{Z}$ and* *let* $\Delta = \det M$*. Then*

$$\Pr(\Delta \equiv 0 \mod p) = 1 - \left[\left(1 - \frac{1}{p}\right)\left(1 - \frac{1}{p^2}\right)\dots\left(1 - \frac{1}{p^m}\right)\right],$$

and the remaining probability is uniformly distributed over the nonzero residues.

Proof. The number of nonsingular matrices with $\mathbb{F}_q$ entries is

$$\left|GL_m(\mathbb{F}_q)\right| = (q^m - 1)(q^m - q)\dots\left(q^m - q^{m-1}\right)$$

out of q^{m^2} total matrices, where q is any prime power [**32**]. This establishes the probability that $p \mid \Delta$. On the other hand, it is a classical fact due to Gauss that every prime modulus has a primitive root, or a generator for its multiplicative group of nonzero elements. Suppose α is a primitive root mod p. If x is the random variable that is uniformly distributed in $\mathbb{Z}/p\mathbb{Z}$, then αx is as well. For any $m \times m$ matrix A, let $f(A)$ be the matrix whose entries are identical to A but $(f(A))_{1j} = \alpha a_{1j}$ for the first-row entries. Then $\det(f^k(A)) \equiv \alpha^k \Delta$. If $\Delta \not\equiv 0$, then this takes all nonzero values modulo p for $k = 0, 1, \dots, m - 1$. But since all of the matrix entries are distributed by the same law for each random matrix $f^k(M)$, it follows that Δ gives equal probability to each nonzero value mod p. □

The following theorem tells us that in sharp contrast to Gromov random groups, where the number of relators required to trivialize the group is exponential in ℓ, even the slowest-growing unbounded functions, like $\log\log\log \ell$ or an inverse Ackermann function, suffice to collapse random abelian groups and random nilpotent groups.

THEOREM 5.7 (Collapsing abelian quotients). *For random abelian groups* $H = \mathbb{Z}^m/\langle R\rangle$ *with* $|R|$ *random relators, if* $|R| \to \infty$ *as a function of* ℓ*, then* $H = \{0\}$ *with probability one (a.a.s.). If* $|R|$ *is bounded as a function of* ℓ*, then there is a positive probability of a nontrivial quotient, both for each* ℓ *and asymptotically.*

Proof. For a relator g, its image in $\mathbb{Z}^m$ is the random vector $\mathbf{A}(g)$, which converges in distribution to a multivariate normal, as described in §2. Furthermore, the image of this vector in projection to $\mathbb{Z}/p\mathbb{Z}$ has entries independently and uniformly distributed. We will consider adding vectors to this collection R until they span $\mathbb{Z}^m$, which suffices to get $H = \{0\}$.

Choose m vectors $v_1, \dots v_m$ in $\mathbb{Z}^m$ at random. These vectors are a.a.s. $\mathbb{R}$-linearly independent, because their distribution is normal and linear dependence is a codimension-one condition. Therefore they span a sublattice $L_1 \subset \mathbb{Z}^m$. The covolume of L_1 (i.e., the volume of the fundamental domain) is $\Delta_1 = \det(v_1, \dots, v_m)$. As we add more vectors, we refine the lattice. Note that $\Delta_1 = 1$ if and only if $L_1 = \mathbb{Z}^m$. Similarly define L_k to be spanned by $v_{(k-1)m+1}, \dots, v_{km}$ for $k = 2, 3, \dots$, and define Δ_k to be the corresponding covolumes.

Note that for two lattices L, L', the covolume of the lattice $L \cup L'$ is always a common divisor of the respective covolumes Δ, Δ'. Therefore, the lattice $L_1 \cup \dots \cup L_k$ has covolume $\le \gcd(\Delta_1, \dots, \Delta_k)$. Here, the Δ_k are identically and independently distributed with the probabilities described in the previous result (Lemma 5.29), and divisibility by different primes is independent, and therefore the probability of having $\gcd(\Delta_1, \dots, \Delta_k) = 1$ is

$$\prod_{\text{primes } p} 1 - \left[1 - \left(1 - \frac{1}{p}\right)\left(1 - \frac{1}{p^2}\right) \dots \left(1 - \frac{1}{p^m}\right)\right]^k,$$

which goes to 1 as $k \to \infty$. (To see this, note that first applying a logarithm, then exchanging the sum and the limit, gives an absolutely convergent sequence.)

On the other hand, it is immediate that for any finite $|R|$ there is a small but nonzero chance that all entries are even, say, which would produce a nontrivial quotient group. ☐

Of course this also follows immediately from the statement in Lemma 5.21, because

$$\Pr(\text{span}\{v_1, \dots, v_r\} = \mathbb{Z}^m) = \frac{1}{\zeta(r-m+1)\cdots\zeta(r)} \longrightarrow 1$$

for any fixed m as $r \to \infty$, but the above argument is appealingly self-contained.

We immediately get corresponding statements for random nilpotent groups and standard random groups. Recall that a group Γ is called *perfect* if $\Gamma = [\Gamma, \Gamma]$; equivalently, if $\text{ab}(\Gamma) = \Gamma/[\Gamma, \Gamma] = \{0\}$.

Corollary 5.30 (Threshold for collapsing random nilpotent groups). *A random nilpotent group $G = N_{s,m}/\langle\langle R \rangle\rangle$ is a.a.s. trivial precisely in those models for which $|R| \to \infty$ as a function of ℓ.*

Corollary 5.31 (Random groups are perfect). *Random groups $\Gamma = F_m/\langle\langle R \rangle\rangle$ are a.a.s. perfect precisely in those models for which $|R| \to \infty$ as a function of ℓ.*

Proof. $\mathbb{Z}^m/\langle R \rangle = \{0\} \iff \text{ab}(\Gamma) = \{0\} \iff \text{ab}(G) = \{0\} \iff G = \{1\}$, with the last equivalence from Theorem 5.2. ☐

We have established that the collapse to triviality of a random nilpotent group G corresponds to the immediate stabilization of the lower central series of the corresponding standard random group:

$\Gamma_1 = \Gamma_2 = \dots$ In fact, we can be somewhat more detailed about the relationship between G and the LCS of Γ.

THEOREM 5.8 (Lifting to random groups). *For $\Gamma = F_m/\langle\langle R\rangle\rangle$ and $G = N_{s,m}/\langle\langle R\rangle\rangle$, they are related by the isomorphism $\Gamma/\Gamma_{s+1} \cong G$. Furthermore, the first s of the successive LCS quotients of Γ are the same as those in the LCS of G, i.e.,*

$$\Gamma_i/\Gamma_{i+1} \cong G_i/G_{i+1} \qquad \textit{for } 1 \le i \le s.$$

PROOF. Since homomorphisms respect LCS depth (Lemma 5.17), the quotient map $\phi : F \to \Gamma$ gives $\phi(F_j) = \Gamma_j$ for all j. We have

$$\Gamma/\Gamma_{s+1} \cong F/\langle\langle R, F_{s+1}\rangle\rangle \cong N/\langle\langle R\rangle\rangle = G$$

by Lemma 5.9 (string arithmetic).

From the quotient map $\psi : \Gamma \to G$, we get $\Gamma_i/\Gamma_{s+1} = \psi(\Gamma_i) = G_i$. Thus

$$G_i\big/_{G_{i+1}} \cong \Gamma_i/\Gamma_{s+1}\big/_{\Gamma_{i+1}/\Gamma_{s+1}} \cong \Gamma_i\big/_{\Gamma_{i+1}}. \quad \square$$

COROLLARY 5.32 (Step drop implies LCS stabilization). *For $G = N_{s,m}/\langle\langle R\rangle\rangle$, if $\mathrm{step}(G) = k < s = \mathrm{step}(N_{s,m})$, then the LCS of the random group Γ stabilizes: $\Gamma_{k+1} = \Gamma_{k+2} = \dots$.*

PROOF. This follows directly from the previous result, since $\mathrm{step}(G) = k$ implies that $G_{k+1} = G_{k+2} = 1$, which means $G_{k+1}/G_{k+2} = 1$. Since $k+1 \le s$, we conclude that $\Gamma_{k+1}/\Gamma_{k+2} = 1$. Thus $\Gamma_{k+2} = \Gamma_{k+1}$, and it follows by the definition of LCS that these also equal Γ_i for all $i \ge k+1$. $\square$

Thus, in particular, when a random nilpotent group (with $m \ge 2$) is abelian but not trivial, the corresponding standard random group has its lower central series stabilize after one proper step:

$$\dots \Gamma_4 = \Gamma_3 = \Gamma_2 \triangleleft \Gamma_1 = \Gamma$$

For instance, with balanced quotients of F_2 this happens about 92% of the time.

In future work, we hope to further study the distribution of steps for random nilpotent groups.

CHAPTER 6

Conclusion

The first few chapters of this thesis focused on the history, background, and necessary definitions to understand Chapters 4 and 5. We learned about random groups and various models of random groups along with a few preliminary results in the field, then switched to cube complexes and various results in that field. In Chapter 4 we married the two fields and used cubulation techniques to prove that random groups in the square model at a certain density are residually finite. In Chapter 5 we focused on a new model of random groups using a different seed group–rather than quotients of free groups, we used quotients of free nilpotent groups–and explored a few results there. Together this document offers a light introduction to cube complex theory and explores two newer models of random groups: the square model and random nilpotent groups.

Appendix: Publisher's Note from *International Mathematics Research Notices*

Cited Literature

[1] F. Bassino, C. Nicaud, and P. Weil. Random presentations and random subgroups: a survey. *ArXiv e-prints*, February 2017.

[2] E.l Breuillard. Random walks on Lie groups. Technical report, 2004. preprint.

[3] Ruth Charney. An introduction to right-angled Artin groups. *Geometriae Dedicata*, 125(1):141–158, 2007.

[4] Matt Cordes, Moon Duchin, Yen Duong, Meng-Che Ho, and Andrew Sánchez. Random nilpotent groups I. *Inter. Math. Res. Notices*, February 2017.

[5] Kelly Delp, Tullia Dymarz, and Anschel Schaffer-Cohen. A matrix model for random nilpotent groups. *ArXiv e-prints*, February 2016.

[6] Persi Diaconis. *Group representations in probability and statistics*. Lecture Notes–Monograph Series. Institute of Mathematical Statistics, 1988.

[7] Nathan M. Dunfield and William P. Thurston. Finite covers of random 3-manifolds. *Inventiones mathematicae*, 166(3):457–521, 2006.

[8] Robert Fitzner and Remco van der Hofstad. Non-backtracking random walk. *Journal of Statistical Physics*, 150(2):264–284, 2013.

[9] A. Garreta, A. Miasnikov, and D. Ovchinnikov. Properties of random nilpotent groups. *ArXiv e-prints*, December 2016.

[10] A. Garreta, A. Miasnikov, and D. Ovchinnikov. Random nilpotent groups, polycyclic presentations, and diophantine problems. *ArXiv e-prints*, December 2016.

[11] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *MSRI Publications*, pages 75–265. Springer, 1987.

[12] M. Gromov. Asymptotic invariants of infinite groups. In *Geometric group theory, Vol. 2 (Sussex, 1991)*, volume 182 of *London Math. Soc. Lecture Note Ser.*, pages 1–295. Cambridge Univ. Press, Cambridge, 1993.

[13] Daniel Groves and Jason Fox Manning. Hyperbolic groups acting improperly. 2017. Preprint.

[14] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *GAFA*, 17:1551–1620, 2008.

[15] Marshall Hall Jr. *The Theory of Groups*. AMS Chelsea Publishing Series. AMS Chelsea Pub., 1976.

[16] G.H. Hardy, E.M. Wright, R. Heath-Brown, and J. Silverman. *An Introduction to the Theory of Numbers*. Oxford mathematics. OUP Oxford, 2008.

[17] Jingyin Huang, Kasia Jankiewicz, and Piotr Przytycki. Cocompactly cubulated 2-dimensional Artin groups. *Comment. Math. Helv.*, 91(3):519–542, 2015.

[18] I. Kapovich and P. Schupp. On group-theoretic models of randomness and genericity. *GROUPS GEOMETRY AND DYNAMICS*, 2(3):383–404, 2008.

[19] Marcin Kotowski and Michał Kotowski. Random groups and Property (T): Żuk's theorem revisited. *Journal of the London Mathematical Society*, 2013.

[20] Rotstyslav V. Kravchenko, Marcin Mazur, and Bogdan V. Petrenko. On the smallest number of generators and the probability of generating an algebra. *Algebra & Number Theory*, 6(2):243–291, 2012.

[21] Wilhelm Magnus, Abraham Karrass, and Donald Solitar. *Combinatorial Group Theory: Presentations of Groups in Terms of Generators and Relations*. Dover Books on Mathematics. Dover Publications, 2004.

[22] Melanie Matchett Wood. Random integral matrices and the Cohen Lenstra Heuristics. *ArXiv e-prints*, April 2015.

[23] Graham Niblo and Lawrence Reeves. Groups acting on cat(0) cube complexes. *Geometry and Topology Monographs: The Epstein Birthday Schrift*, 1:1–7, 1997.

[24] J. Nielsen. Die Isomorphismen der allgemeinen, unendlichen Gruppe mit zwei Erzeugenden. (Mit 3 Figuren im Text) (German). *Mathematische Annalen*, 78:385–397, 1917.

[25] Tomasz Odrzygóźdź. Cubulating random groups in the square model. *ArXiv e-prints*. Preprint.

[26] Tomasz Odrzygóźdź. Nonplanar isoperimetric inequality for random groups, 2014. Available at `http://students.mimuw.edu.pl/~to277393/web/files/nonplanar.pdf`.

[27] Tomasz Odrzygóźdź. The square model for random groups. *Colloq. Math.*, 142(2):227–254, 2016.

[28] Yann Ollivier. A January 2005 invitation to random groups. In *Ensaios Matemáticos (Mathematical Surveys)*, volume 10. Sociedade Brasiliera de Matematica, Rio de Janiero, 2005.

[29] Yann Ollivier. Some small cancellation properties of random groups. *Internat. J. Algebra Comput.*, 17(1):37–51, 2007.

[30] Yann Ollivier and Daniel T. Wise. Cubulating random groups at density less than 1/6. *Trans. Amer. Math. Soc.*, 363(9):4701–4733, 2011.

[31] Piotr Przytycki and John M. Mackay. Balanced walls for random groups. *The Michigan Mathematical Journal*, 2015:–, 2015.

[32] Derek J.S. Robinson. *A Course in the Theory of Groups*. Graduate Texts in Mathematics. Springer New York, 2012.

[33] N S Romanovski. A freedom theorem for groups with one defining relation in the varieties of solvable and nilpotent groups of given lengths. *Mathematics of the USSR-Sbornik*, 18(1):93, 1972.

[34] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.

[35] John R. Stallings. Topology of finite graphs. *Inventiones mathematicae*, 71:551–566, 1983.

[36] Yinghui Wang and Richard P. Stanley. The Smith Normal Form Distribution of a Random Integer Matrix. *ArXiv e-prints*, May 2015.

[37] Daniel T. Wise. Cubulating small cancellation groups. *Geometric & Functional Analysis GAFA*, 14(1):150–214, 2004.

[38] Daniel T. Wise. *From Riches to Raags: 3-manifolds, Right-angled Artin Groups, and Cubical Geometry*. Conference board of the Mathematical Sciences. Regional conference series in mathematics. American Mathematical Society, 2012.

[39] D.T. Wise. Cubulating small cancellation groups. *Geom. Funct. Anal.*, 14:150–214, 2004.

[40] Andrzej Żuk. Property (T) and Kazhdan constants for discrete groups. *GAFA*, 13:643–670, 2003.

Vita

EDUCATION

Ph.D., Mathematics
University of Illinois at Chicago. Expected graduation: August 2017.
M.A., Mathematics
University of California, Santa Barbara. June 2012.
B.A., Mathematics and Philosophy
Yale University, New Haven, CT. May 2010.

PUBLICATIONS

M. Cordes, M. Duchin, **Y. Duong**, A. Sánchez, T. Ho. "Random Nilpotent Groups I," *International Mathematics Research Notices*. February 2017.

Y. Duong, K. Meehan, L. Merrill, L. Snyder. "Intrinsically Linked Signed Graphs in Projective Space," *Discrete Mathematics*. 312 (2012), no. 12-13, 2009-2022.

AWARDS AND FELLOWSHIPS

Research and Training Grant Pre-Doctoral Fellow — Fall 2014
One semester research fellowship

University Fellow — Fall 2012-present
Two-year fellowship at University of Illinois at Chicago

Graduate Opportunity Fellow — 2010-2011
One-year fellowship at University of California, Santa Barbara

Mellon Mays Undergraduate Fellow — Spring 2008-present
National fellowship supporting minorities in academia

Barge and Runk Prize — Spring 2007
Yale University: underclassmen competitive examination in mathematics

TEACHING ASSISTANT EXPERIENCE

- Math 3C. Differential and Integral Calculus I. — UCSB, Summer 2011
- Math 3A. Calculus with Applications, First Course. — UCSB, Fall 2011
- Math 34A. Calculus for Social and Life Sciences. — UCSB, Winter 2012
- Math 3B. Differential and Integral Calculus II. — UCSB, Spring 2012
- Math 181. Calculus II. — UIC, Fall 2012
- Math 180. Calculus I. — UIC, Spring 2013
- Math 179. Emerging Scholars Workshop for Calculus. — UIC, 2013-2014

SEMINARS, COLLOQUIA, SERVICE

Texas Women in Mathematics Symposium Fall 2016

- Founded annual traveling conference for women in mathematics in Texas area

Graduate Student Seminar, Temple University. October 2015

- Talk: "An introduction to cube complexes and random groups"

Junior Topology Seminar, University of Texas at Austin. Fall 2015-2016

- Talk: "The contracting boundary of CAT(0) spaces is great!"
- Talk: "One weird trick to tell if RAAGS split over Abelian groups"
- Talk: "From folds to Cubulations: an introduction to special cube complexes"

Mellon Mays Undergraduate Fellowship Professional Network 2013-2014

- Vice-Chair of Development. Organize Chicago area networking for MMUF alumni
- Talk: "Hyperbolic Groups Rule the World!" Yale March 26, 2014
- Talk: "Special Cube Complexes: The future of hyperbolic studies" MMFPN Chicago event May 10, 2014

UIC Recruitment Day Mar 2014

- Special lecture: "A problem in geometric group theory"

Midwest Women in Mathematics Symposium

- UIC. Founder and Organizer. April 20, 2013
- Notre Dame. Talk: "Small Cancellation Theory and Hyperbolicity" April 5, 2014

Graduate Student Colloquium, UIC.

- Talk: "Introduction to Group Cohomology" Nov 2012
- Talk: "Introduction to Teichmuller space, the mapping class group, and the curve complex" Mar 2014

Graduate Geometry, Topology and Dynamics Seminar, UIC.

- Talks: "Fundamental Theorem of Bass Serre Theory," "Schreier's Theorem," "Trees" Spring 2013
- Organizer. Talks: "What is a Stallings Fold?," "What is small cancellation theory?," "Random groups are perfect" 2013-2014

= *Association for Women in Mathematics, UIC.* 2012-2014

- Helped organize various AWM activities, served on a panel for undergraduates applying to graduate school, moderated a Work-Life balance panel.

M&M Seminar, UCSB. 2011-2012

- Established semiweekly interdisciplinary mathematics and philosophy seminar
- Grant Received: Graduate Student Scholarly Collaborations
- Some talks: "Why Should I Believe You? Mathematics and Epistemological Privilege", "Everything", "On Infinity"

Hypatian Seminar, UCSB. 2010-2012

- Seminar for underrepresented groups in academia, general student guidance.
- Attendee and talk given, presented at JMM Outreach Poster Session

WORKSHOPS AND SUMMER SCHOOLS

IAS Women And Mathematics Program: Geometry and Randomness in Group Theory. May 15-26 2016

- TA for Advanced Course: Random Groups.

IAS Women And Mathematics Program: Words, Curves and Loops. May 9-20 2016

- Talk: "Working with the Square Model of Random Groups." Available online.

MSRI Summer Graduate School: Geometric Group Theory. June 15-26 2015

Research Cluster in Random Groups. Tufts University, Boston. June 25- Aug 5 2014

Summer School on the Geometry of Outer Space. Aix Marseille Université June 24-29 2013

Topology Students Workshop. Georgia Tech University, Atlanta. June 11-15 2012

Geometric Group Theory Student RTG. UC Berkeley. June 18-22 2012

Enhancing Diversity in Graduate Education. NC State. June 7-July 10 2010

CONFERENCES ATTENDED

Tech Topology Conference Georgia Tech, Atlanta. Dec 9-11, 2016

- Lightning Talk: "Random Groups and Cubulations."

Graduate Student Conference in Algebra, Geometry and Topology. Temple University, Philadelphia. May 14-15 2016

- Talk: "Working with the Square Model of Random Groups."

Underrepresented Students in Topology and Algebra Research Symposium. Sam Houston State University, Huntsville. April 14-15 2016

- Distinguished Graduate Speaker: "Random Groups and Cubulations."

Graduate Student Combinatorics Conference. Clemson University. April 1-3 2016

- Talk: "The Square Model of Random Groups."

Workshop on Mapping Class Groups and $Out(F_N)$. University of Texas, Austin. May 25-30 2015

Cornell Topology Festival. May 8-10 2015

Graduate Student Topology and Geometry Conference. UIUC. Mar 28-29 2015

- Talk: "Random Groups are Perfect."

GEAR Junior Retreat. University of Michigan, Ann Arbor. May 23-June 1 2014

Young Geometric Group Theory III. CIRM, Luminy. Jan 20-24 2014

Geometry and Topology in New York. Columbia University . Aug 12-16 2013

Hot Topics: Surface subgroups and cube complexes. MSRI. March 18 - 22, 2013

Young Geometric Group Theory II. Technion, Haifa, Israel. Feb 3-7, 2013

Women in Mathematics Symposium.

- Pomona College, CA. Talk: "Expander Exposition" Nov 20, 2010

- Loyola Marymount University, CA. Talk: "Some Philosophies of Mathematics." Jan 21, 2011

Joint Mathematics Meetings.

- San Francisco, CA. Jan 13-16, 2010. Undergraduate research poster session.
- Boston, MA. Jan 5-7, 2012. Outreach poster session.

OTHER

Science Fair Judging:

- California State Science Fair, Los Angeles. Springs 2011 and 2012
- QED: Chicago's Youth Math Symposium. April 27, 2013
- Austin Energy Regional Science Festival. February 2016 and 2017

Languages: Intermediate French, intermediate Vietnamese

Media: Blog: Baking and Math, `http://bakingandmath.com`

Joint work with Mark Krasniewski:

- Ian Minh Krasniewski October 7, 2014
- Maeve Kim Duong March 30,2017