

RESIDUAL FINITENESS GROWTHS OF VIRTUALLY SPECIAL GROUPS

KHALID BOU-RABEE, MARK F. HAGEN, AND PRIYAM PATEL

ABSTRACT. Let G be a virtually special group. Then the residual finiteness growth of G is at most linear. This result cannot be found by embedding G into a special linear group. Indeed, the special linear group $\mathrm{SL}_k(\mathbb{Z})$, for $k > 2$, has residual finiteness growth n^{k-1} .

CONTENTS

1. Introduction	1
Acknowledgments	2
2. Background	3
2.1. Quantifying residual finiteness	3
2.2. Right-angled Artin groups and virtually special groups	4
3. Virtually special groups	8
4. Special linear groups	11
4.1. The upper bound	11
4.2. The lower bound	11
References	12

1. INTRODUCTION

This paper demonstrates that Stallings’s topological proof of the residual finiteness of free groups [Sta83, Theorem 6.1] extends efficiently to the class of right-angled Artin groups (and, more generally, to virtually special groups). To state our results, we use notation developed in [BRM10], [BRM], [BR11a] for quantifying residual finiteness. Let A be a residually finite group with generating set X . The *divisibility function* $D_A : A \setminus \{1\} \rightarrow \mathbb{N}$ is defined by

$$D_A(g) = \min\{[A : B] : g \notin B \wedge B \leq A\}.$$

Define $F_{A,X}(n)$ to be the maximal value of D_A on the set

$$\{g : \|g\|_X \leq n \wedge g \neq 1\},$$

where $\|\cdot\|_X$ is the word-length norm with respect to X . The growth of $F_{A,X}$ is called the *residual finiteness growth function*. The growth of $F_{A,X}(n)$ is, up to a natural equivalence, independent of the choice of generating set (see §2.1 for this and generalizations of $F_{A,X}$). Loosely speaking, residual finiteness growth measures how well a group is approximated by its profinite completion.

Date: August 11, 2014.

2010 Mathematics Subject Classification. Primary: 20E26; Secondary: 20F65, 20F36.

Key words and phrases. Residual finiteness growth, special cube complex, right-angled Artin group.

Our first result, proved in §3, gives bounds on the residual finiteness growth of any right-angled Artin group:

Theorem 1.1. *Let Γ be a simplicial graph. Let A_Γ be the corresponding right-angled Artin group with standard generating set X . Then $F_{A_\Gamma, X}(n) \leq n + 1$.*

The canonical completion [HW08], a pillar from the structure theory of special groups, plays an integral role in our proof of Theorem 1.1. Thus, we include a complete proof of a quantified version of the canonical completion construction in §2.2.4.

Our residual finiteness bound for right-angled Artin groups, hereafter known as *raAgs*, is an *efficient* extension of Stallings' proof of [Sta83, Theorem 6.1]. Indeed, the bound achieved over the class of all raAgs exactly coincides with that given by a direct application of Stallings' proof in the case $A_\Gamma = \mathbb{Z} * \mathbb{Z}$. Further, Theorem 1.1 immediately extends to many other groups. That is, since bounds on residual finiteness growth are inherited by passing to subgroups and to super groups of finite-index (see §2.1), Theorem 1.1 gives bounds on residual finiteness growth for any group that virtually embeds into a raAg, the so-called *virtually special groups* (defined in Section 2.2.3). This class includes, for example, Coxeter groups [NR03, HW10], free-by- \mathbb{Z} groups with irreducible atoroidal monodromy [HW13], hyperbolic 3-manifold groups [Wis, Ago12], fundamental groups of many aspherical graph manifolds [PW11, Liu], fundamental groups of mixed 3-manifolds [PW], and random groups at sufficiently low density [Ago12, OW11]. See also [AFW13] for more explicit results in the 3-manifold case.

Our second result gives a sense of the topological nature of Theorem 1.1. It is known that any finitely generated raAg embeds in $SL_k(\mathbb{Z})$ for appropriately-chosen k , and that $k > 2$ unless the raAg in question is free [Big01], [Kra02], [CW02]. However, Theorem 1.2 shows that the best upper bound on residual finiteness growth of the raAg that can be inferred in this manner is superlinear (c.f. [BRK12], where the normal residual finiteness growth for $SL_k(\mathbb{Z}), k > 2$, is shown to be n^{k^2-1}).

Theorem 1.2. *The residual finiteness growth of $SL_k(\mathbb{Z}), k > 2$, is bounded above by Cn^{k-1} and below by $1/Cn^{k-1}$ for some fixed $C > 0$. That is, $F_{SL_k(\mathbb{Z})}(n) \simeq n^{k-1}$.*

Since residual finiteness growth is a commensurability invariant (see §2.1), this theorem computes the residual finiteness growth for any S arithmetic subgroup of $SL_k(\mathbb{C})$ for $k > 2$ (in the sense of [BRK12]). Our proof of Theorem 1.2 relies on a result of Lubotzky, Mozes, and Raghunathan [LMR00] and the structure theory of finite-index subgroups of unipotent groups [GSS88].

This article is organized as follows. In §2, we give some necessary background on quantifying residual finiteness, raAgs, and cubical geometry. In the interest of self-containment, we also provide the construction of the canonical completion from [HW08] for a special case that is relevant to our framework. In §3, we generalize Stallings' proof to raAgs to give a proof of Theorem 1.1. In §4, we conclude with a proof of Theorem 1.2.

Acknowledgments. The authors are grateful to Ian Agol, Benson Farb, Michael Larsen, Feng Luo, and Ben McReynolds for useful and stimulating conversations. K.B. and M.F.H. gratefully acknowledge the hospitality and support given to them from the Ventotene 2013 conference for a week while they worked on some of the material in this paper. K.B. was partially supported by the NSF grant DMS-1405609. The work of M.F.H. was supported by the National Science Foundation under Grant Number NSF 1045119. Finally, the authors would like to thank the referee for several excellent suggestions and comments that have improved the quality of this paper.

2. BACKGROUND

2.1. Quantifying residual finiteness. For previous works on quantifying this basic property see, for instance, [Bus09], [Pat12], [KM11], [BRM], [BRK12], [BRM11], [BR11b], [BRM10], and [Riv12]. Here we unify the notation used in previous papers and demonstrate that this unification preserves basic properties of residual finiteness growth. Further, we will see that this unification also elucidates the behavior of residual finiteness functions under commensurability.

Let A be a group with generating set X . Given a class of subgroups Ω of A , set $\Lambda_\Omega = \bigcap_{\Delta \in \Omega} \Delta$ (c.f. [IBM]). The *divisibility function* (c.f. [BRM10], [BRM], [BR11a]), $D_A^\Omega : A \setminus \Lambda_\Omega \rightarrow \mathbb{N}$ is defined by

$$D_A^\Omega(g) = \min\{[A : B] : g \notin B \wedge B \in \Omega\}.$$

Define $\text{RF}_{A,X}^\Omega(n)$ to be the maximal value of D_A^Ω on the set

$$\{g : \|g\|_X \leq n \wedge g \notin \Lambda_\Omega\},$$

where $\|\cdot\|_X$ is the word-length norm with respect to X . The growth of $\text{RF}_{A,X}^\Omega$ is called the *residual Ω growth function*. The growth of $\text{RF}_{A,X}^\Omega(n)$ is, up to a natural equivalence, independent of the choice of generating set (see Lemmas 2.1 and 2.2 below).

When Ω is the class of all subgroups of a residually finite group, A , we have $\Lambda_\Omega = \{1\}$ and the growth of $\text{RF}_{A,X}^\Omega(n)$ is the *residual finiteness growth* of A . In the case when Ω consists of all normal subgroups of A and $\Lambda_\Omega = \{1\}$, the function $\text{RF}_{A,X}^\Omega$ is the *normal residual finiteness growth function*.

Our first result demonstrates that the residual Ω growth of a group is well-behaved under passing to subgroups.

Lemma 2.1. *Let G be generated by a set S , and let $H \leq G$ be generated by a finite set $L \subset G$. Let Ω be a class of subgroups of G . Then there exists $C > 0$ such that $\text{RF}_{H,L}^{\Omega \cap H}(n) \leq \text{RF}_{G,S}^\Omega(Cn)$ for all $n \geq 1$.*

Proof. By definition of Ω and $\Omega \cap H$ we have $D_H^{\Omega \cap H}(h) \leq D_G^\Omega(h)$ for all $h \in H$ and $h \notin \Lambda_{\Omega \cap H} = \Lambda_\Omega \cap H$. Hence,

$$(1) \quad \begin{aligned} \text{RF}_{H,L}^{\Omega \cap H}(n) &= \max\{D_H^{\Omega \cap H}(h) : \|h\|_L \leq n \wedge h \notin \Lambda_{\Omega \cap H}\} \\ &\leq \max\{D_G^\Omega(h) : \|h\|_L \leq n \wedge h \notin \Lambda_\Omega\}. \end{aligned}$$

Further, there exists a $C > 0$ such that any element in L can be written in terms of at most C elements of S . Thus,

$$(2) \quad \{h \in H : \|h\|_L \leq n\} \subseteq \{g \in G : \|g\|_S \leq Cn\}.$$

So by (1) and (2), we have that

$$\text{RF}_{H,L}^{\Omega \cap H}(n) \leq \max\{D_G^\Omega(h) : \|h\|_L \leq n\} \leq \max\{D_G^\Omega(g) : \|g\|_S \leq Cn\} = \text{RF}_{G,S}^\Omega(Cn),$$

as desired. \square

Lemma 2.1 shows that the residual Ω growth of a group does not depend heavily on the choice of generating set. To this end, we write $f \preceq g$ if there exists $C > 0$ such that $f(n) \leq Cg(Cn)$. Further, we write $f \simeq g$ if $f \preceq g$ and $g \preceq f$. If $f \preceq g$, we say that g *dominates* f . So, in particular, Lemma 2.1 gives that, up to \simeq equivalence, the growth of

$\text{RF}_{A,X}^\Omega$ does not depend on the choice of generating set. We can and will often, therefore, drop the decoration indicating the generating set X when dealing with growth functions.

The next result, coupled with Lemma 2.1, demonstrates that residual Ω growth is also well-behaved under passing to super groups of finite-index.

Lemma 2.2. *Let H be a finite-index subgroup of a finitely generated group G . Let H be generated by L and G by S . Let Ω be a class of subgroups of G with $H \in \Omega$. Then $\text{RF}_G^\Omega(n) \preceq \text{RF}_H^{\Omega \cap H}(n)$.*

Proof. $D_G^\Omega(g) \leq [G : H]D_H^{\Omega \cap H}(g)$ if $g \in H$ and because $H \in \Omega$, $D_G^\Omega(g) \leq [G : H]$ for $g \notin H$. Bringing these facts together gives

$$(3) \quad \text{RF}_{G,S}^\Omega(n) = \max\{D_G^\Omega(g) : \|g\|_S \leq n\} \leq [G : H] \max\{D_H^{\Omega \cap H}(g) : \|g\|_S \leq n \wedge g \in H \setminus \Lambda_\Omega\}.$$

Fix word metrics d_H and d_G for H and G with respect to L and S respectively. The identity map $H \rightarrow G$ is a $(K, 0)$ quasi-isometry by the Milnor-Schwarz lemma. For any element $g \in G$ with $d_G(g, 1) \leq n$, either $g \notin H$ or $d_H(g, 1) \leq Kd_G(g, 1) \leq Kn$. Thus, there exists a natural number C such that for all n ,

$$(4) \quad \{h \in H : \|h\|_S \leq n\} \subseteq \{g \in H : \|g\|_L \leq Cn\}.$$

So by (3) and (4), we have that

$$\begin{aligned} \text{RF}_{G,S}^\Omega(n) &\leq [G : H] \max\{D_H^{\Omega \cap H}(g) : \|g\|_S \leq n \wedge g \in H \setminus \Lambda_\Omega\} \\ &\leq [G : H] \max\{D_H^{\Omega \cap H}(g) : \|g\|_L \leq Cn \wedge g \in H \setminus \Lambda_\Omega\}. \end{aligned}$$

Thus, $\text{RF}_{G,S}^\Omega(n) \leq [G : H] \text{RF}_{H,L}^{\Omega \cap H}(Cn)$, as desired. \square

Recall that subgroups G and H of A are *commensurable* if $G \cap H$ is finite-index in both G and H . Lemma 2.2 and Lemma 2.1 demonstrate that residual Ω growth, and residual finiteness growth, behave well under commensurability as noted in the following Proposition.

Proposition 2.3. *Let G be a finitely generated subgroup of A . Let Ω be a class of subgroups of G . Let H be commensurable with G , and let $G \cap H \in \Omega$. Then $\text{RF}_H^{\Omega \cap H}(n) \simeq \text{RF}_G^\Omega(n)$. In particular, $\text{F}_G(n) \simeq \text{F}_H(n)$.*

Proof. Since $G \cap H$ is a finite-index subgroup of both G and H , Lemma 2.2 gives that $\text{RF}_{G \cap H}^{\Omega \cap H}(n)$ dominates both $\text{RF}_H^{\Omega \cap H}(n)$ and $\text{RF}_G^\Omega(n)$. Further, Lemma 2.1, demonstrates that both $\text{RF}_H^{\Omega \cap H}(n)$ and $\text{RF}_G^\Omega(n)$ dominate $\text{RF}_{G \cap H}^{\Omega \cap H}(n)$. Thus, $\text{RF}_H^{\Omega \cap H}(n) \simeq \text{RF}_G^\Omega(n)$, as desired.

In the case when Ω is the set of all finite-index subgroups of G , we see that $\Omega \cap G \cap H$ is precisely the set of all finite-index subgroups of $G \cap H$. Thus, we have $\text{F}_G(n) \preceq \text{F}_{G \cap H}(n)$. So as $G \cap H$ is finite-index in H , we similarly achieve $\text{F}_H(n) \preceq \text{F}_{G \cap H}(n)$. So by Lemma 2.1, we are done. \square

2.2. Right-angled Artin groups and virtually special groups.

2.2.1. *Right-angled Artin groups.* Right-angled Artin groups (raAgs) are a widely-studied class of groups (see [Cha07] for a comprehensive survey). These groups have great utility in geometric group theory both because the class of subgroups of raAgs is very rich and because raAgs are fundamental groups of particularly nice nonpositively-curved cube complexes. For each finite simplicial graph Γ , the associated finitely generated raAg A_Γ is given by the presentation

$$\left\langle x_i \in \text{Vertices}(\Gamma) \mid [x_j, x_k] = x_j x_k x_j^{-1} x_k^{-1}, \{x_j, x_k\} \in \text{Edges}(\Gamma) \right\rangle.$$

For example, if Γ has no edges, then A_Γ is the free group, freely generated by $\text{Vertices}(\Gamma)$, while $A_\Gamma \cong \mathbb{Z}^{|\text{Vertices}(\Gamma)|}$ when Γ is a clique. More generally, A_Γ decomposes as the free product of the raAgs associated to the various components of Γ , and if Γ is the join of subgraphs Γ_1, Γ_2 , then $A_\Gamma \cong A_{\Gamma_1} \times A_{\Gamma_2}$.

2.2.2. Nonpositively-curved cube complexes. We recall some basic notions about nonpositively-curved cube complexes that will be required below. More comprehensive discussions of $\text{CAT}(0)$ and nonpositively-curved cube complexes can be found in, e.g., [Che00, Hag12, Sag95, Wis, Wis12]. We largely follow the discussion in [Wis12].

For $d \geq 0$, a d -cube is a metric space isometric to $[-\frac{1}{2}, \frac{1}{2}]^d$ with its ℓ^1 metric. A d' -face of the d -cube C is a subspace obtained by restricting $d - d'$ of the coordinates to $\pm\frac{1}{2}$, while a *midcube* of C is a subspace obtained by restricting exactly one coordinate to 0. A *cube complex* is a CW-complex whose cells are cubes of various dimensions and whose attaching maps restrict to combinatorial isometries on faces.

Let X be a cube complex and let $x \in X$ be a 0-cube. The *link* $\text{lk}(x)$ of x is the simplicial complex with an n -simplex σ_c for each $(n + 1)$ -cube c containing x , with the property that $\sigma_c \cap \sigma_{c'} = \cup_{c''} \sigma_{c''}$, where c'' varies over the constituent cubes of $c \cap c'$. A simplicial complex is *flag* if each $(n + 1)$ -clique in the 1-skeleton spans an n -simplex, and X is *nonpositively-curved* if $\text{lk}(x)$ is flag for each $x \in X^{(0)}$. If the nonpositively-curved cube complex X is simply connected, then X is a $\text{CAT}(0)$ cube complex, so named for the existence of a natural piecewise-Euclidean $\text{CAT}(0)$ metric [Gro87, Bri91, Lea10].

All maps of nonpositively-curved cube complexes in this paper are, unless stated otherwise, *cubical maps*, i.e. they send open d -cubes isomorphically to open d -cubes for $d > 0$ and send 0-cubes to 0-cubes.

2.2.3. Special cube complexes. Special cube complexes were defined in [HW08] in terms of the absence of certain pathological configurations of immersed hyperplanes. Here, we are interested in the characterization of special cube complexes in terms of raAgs, established in the same paper.

Let Γ be a simplicial graph and A_Γ the associated raAg. The *Salvetti complex* S_Γ associated to Γ is a $K(A_\Gamma, 1)$ cube complex, first constructed in [CD95], that we now describe. S_Γ has a single 0-cube v and a 1-cube e_{x_i} for each $x_i \in \text{Vertices}(\Gamma)$. For each relation $[x_j, x_k]$ in the above presentation of A_Γ , we add a 2-cube with attaching map $e_{x_j} e_{x_k} e_{x_j}^{-1} e_{x_k}^{-1}$. Finally, we add an n -cube for each size- n set of pairwise-commuting generators. Note that the image in S_Γ of each n -cube is an embedded n -torus and S_Γ is a nonpositively-curved cube complex.

The cubical map $f : Y \rightarrow X$ of nonpositively-curved cube complexes is a *local isometry* if the following conditions are satisfied:

- (1) f is locally injective; equivalently, the induced map $\text{lk}(x) \rightarrow \text{lk}(f(x))$ is injective for each $x \in X^{(0)}$, and
- (2) for each $x \in X^{(0)}$, the subcomplex $\text{lk}(x) \subseteq \text{lk}(f(x))$ is an *induced* subcomplex in the sense that $n + 1$ vertices of $\text{lk}(x)$ span an n -simplex whenever their images in $\text{lk}(f(x))$ span an n -simplex.

If X, Y are $\text{CAT}(0)$ cube complexes, and there is an injective local isometry $Y \rightarrow X$, then Y is *convex* in X . More generally, if X, Y are nonpositively-curved and there is a local isometry $Y \rightarrow X$, then the image of Y is *locally convex* in X . It should be noted that covering maps of nonpositively-curved cube complexes are local isometries.

Remark 2.4 (Cubical convexity). The term “convex” is justified by the fact that if Y is a convex subcomplex of the CAT(0) cube complex X , in the preceding sense, then $Y^{(1)}$, with the usual graph metric, is metrically convex in $X^{(1)}$. When working with a CAT(0) cube complex X , we will only use the usual graph metric on $X^{(1)}$, ignoring the CAT(0) metric. In particular, a (*combinatorial*) *geodesic* in X shall be understood to be a path in $X^{(1)}$ with a minimal number of edges among all paths with the given endpoints. Equivalently, a combinatorial path $\gamma \rightarrow X$ is a geodesic if and only if each hyperplane of X intersects at most one 1-cube of γ , and a connected subcomplex Y of X has isometrically embedded 1-skeleton if and only if it has connected intersection with each hyperplane. We will refer to a connected subcomplex of X as *isometrically embedded* if it has the latter property.

The notion of a cubical local isometry yields an elegant characterization of special cube complexes (see [HW08]) which we shall take to be the definition:

Definition 2.5 (Special cube complex, virtually special group). The nonpositively-curved cube complex X is *special* if there exists a local isometry $X \rightarrow S_\Gamma$ for some simplicial graph Γ . The group G is [*virtually*] *special* if there exists a special cube complex X having [a finite-index subgroup of] G as its fundamental group. If this cube complex can be chosen to be compact, then G is *virtually compact special*.

2.2.4. *Canonical completion.* A substantial part of the utility of special cube complexes is the fact that they behave in several important ways like graphs. Chief among the graph-like features of special cube complexes is the ability to extend compact local isometries to covers, generalizing the fact that finite immersions of graphs extend to covering maps [Sta83]. This procedure, introduced in [HW08] and outlined presently, is called “canonical completion”. Since it is more directly suited to our situation, we will follow the discussion in [HW12]; in the interest of a relatively self-contained exposition, we now sketch the special case of the construction in [HW12, Definition 3.2] that we will later require.

Theorem 2.6. [*Canonical completion for Salvetti complexes*] *Let Y be a compact cube complex, and let $f : Y \rightarrow S_\Gamma$ be a local isometry, where S_Γ is the Salvetti complex of a raAg A_Γ . Then there exists a finite-sheeted cover $\widehat{S}_\Gamma \rightarrow S_\Gamma$ such that f lifts to an embedding $\widehat{f} : Y \rightarrow \widehat{S}_\Gamma$.*

The space \widehat{S}_Γ is called the *canonical completion* of f and will be denoted by $\mathfrak{C}(Y \rightarrow S_\Gamma)$.

Proof of Thm. 2.6. Let e be a (closed) oriented 1-cube of S_Γ . Each component of the preimage of e in Y is either a cycle, an interval, or a 0-cube mapping to the base-point, since $Y \rightarrow S_\Gamma$ is locally injective. For each non-cycle component, we add an appropriately oriented open 1-cube to Y to form an oriented cycle covering e . The map $f : Y \rightarrow S_\Gamma$ extends by declaring the new open 1-cube to map by an orientation-preserving homeomorphism to $\text{Int}(e)$. Let Y° be the union of Y and all of these new 1-cubes. We thus have a map $\widehat{f} : Y^\circ \rightarrow S_\Gamma$ that extends f and is a covering map on 1-skeleta. The 1-skeleton of Y° will be the 1-skeleton of $\mathfrak{C}(Y \rightarrow S_\Gamma)$. By Lemma 2.7 below, for each 2-cube c of S_Γ , the boundary path of c lifts to a closed path in Y° , and hence we can attach 2-cubes to Y to form a complex Y^\bullet equipped with a cubical map $Y^\bullet \rightarrow S_\Gamma$ that extends $Y^\circ \rightarrow S_\Gamma$ and restricts to a covering map on 2-skeleta. For each higher-dimensional cube c of S_Γ , the 2-skeleton of c lifts to Y^\bullet , and we form $\mathfrak{C}(Y \rightarrow S_\Gamma)$ by adding to Y^\bullet each cube whose 2-skeleton appears. By construction, $\mathfrak{C}(Y \rightarrow S_\Gamma)$ covers S_Γ and is thus non-positively curved. \square

Lemma 2.7. *For each 2-cube c of S_Γ , the boundary path of c lifts to a closed path in Y° .*

Proof. Let $f : Y \rightarrow S_\Gamma$ be the local isometry from the proof of Theorem 2.6, and also denote by f its extension to Y° . Let $\bar{\gamma} : [0, 4] \rightarrow S_\Gamma$ be the boundary path of c and let γ be a lift of $\bar{\gamma}$ through f .

There exists, by construction, a finite sequence $(C_i)_{i=1}^4$ of (not necessarily distinct) cycles in Y° such that

- (1) $\forall i$, $f(C_i)$ is a 1-cube in S_Γ ;
- (2) $(C_1 \cap \gamma) \cdot (C_2 \cap \gamma) \cdot (C_3 \cap \gamma) \cdot (C_4 \cap \gamma) = \gamma$; and
- (3) each C_i has at most one 1-cube not in Y .

Set $\gamma_i = \gamma([i-1, i]) = C_i \cap \gamma$ and set $C_i^c := C_i \setminus \text{int}(\gamma_i)$. We note that when γ_i is a loop, C_i^c consists of the single vertex $\gamma(i)$.




FIGURE 1

Case 1a: Suppose that γ_1 is not a loop in Y° so that C_1^c is not a vertex. Suppose that γ_2 is also not a loop. Condition (3) above then ensures that either $\gamma_1 \subset Y$, in which case we set $D_1 = \gamma_1$, or $\gamma_1 \subset Y^\circ \setminus Y$ so that $C_1^c \subset Y$ and we set $D_1 = C_1^c$. Similarly, either $\gamma_2 \subset Y$ and $\gamma_2 = D_2$, or $D_2 = C_2^c$. Since $f : Y \rightarrow S_\Gamma$ is a local isometry and $D_1 \cup D_2 \subset Y$, the map $D_1 \cup D_2 \hookrightarrow Y$ extends to a map $D_1 \times D_2 \hookrightarrow Y$. Let ℓ_i denote the length of the path D_i for $i = 1, 2$. The map $D_1 \times D_2 \rightarrow Y$ is a cubical embedding of the standard tiling of $[0, \ell_1] \times [0, \ell_2]$ by 2-cubes. Since $f| : (Y^\circ)^{(1)} \rightarrow S_\Gamma^{(1)}$ is an immersion, the third side of the rectangle $D_1 \times D_2$ must coincide with either γ_3 or C_3^c (depending on the orientation of the third side of $D_1 \times D_2$) as shown by Figure 1. If $D_1 \times D_2 \cap \gamma_3 \neq \emptyset$ we set $D_3 = \gamma_3$. Otherwise, we set $D_3 = C_3^c$.

The fact that f is a covering map on the 1-skeleta of Y° and S_Γ implies that $\ell_1 = \ell_3$, where ℓ_3 is the length of D_3 . Indeed, Figure 2 shows that if $\ell_1 > \ell_3$, the covering map condition fails at the vertex $v_0 \in f^{-1}(v)$. Similarly, Figure 3 shows that if $\ell_1 < \ell_3$, then the covering criterion would fail at the vertex $v_1 \in f^{-1}(v)$. Thus D_3 is precisely the third side of $D_1 \times D_2$.

A similar argument shows that the fourth side of $D_1 \times D_2$ must coincide with either γ_4 , in which case we set $D_4 = \gamma_4$, or C_4^c so that we set $D_4 = C_4^c$. Again, we can show that $\ell_2 = \ell_4$ so that D_4 forms the fourth side of $D_1 \times D_2$.

By construction, each D_i shares endpoints with the γ_i for $i = 1, 2, 3, 4$. Thus, $\gamma = \gamma_1\gamma_2\gamma_3\gamma_4$ forms a closed path since $D_1D_2D_3D_4$ forms the closed boundary path of the rectangle $D_1 \times D_2$.

Case 1b: Suppose that γ_1 is not a loop and that γ_2 is a loop. In this case, γ_3 necessarily equals γ_1^{-1} since γ is a lift of a path in S_Γ representing a commutator relation in A_Γ .

If γ_4 is also a loop, then γ is a closed path. Hence suppose that γ_4 is not a loop. Then either $\gamma_i \subset Y$ and we set $D_i = \gamma_i$, or $D_i = C_i^c$ for $i = 1, 4$. Thus the map $D_1 \cup D_4 \hookrightarrow Y$ extends to a map $D_1 \times D_4 \hookrightarrow Y$, so that we have an embedded rectangle $D_1 \times D_4$ in Y . Now, a similar argument as in the previous case shows that γ_2 must be the third side of $D_1 \times D_4$, contradicting the fact that $D_1 \times D_4$ is embedded. Therefore, γ_4 must be a loop and γ is a closed path in Y° .




FIGURE 2. The map f fails to be locally injective at v_0 .




FIGURE 3. The map f fails to be locally injective at v_1 .

Case 2: Suppose that γ_1 is a loop in Y° . If γ_2 is also a loop, γ is of course a closed path in Y° . Therefore, we assume that γ_2 is not a loop. Note that if γ_3 is a loop, then γ_4 necessarily equals γ_2^{-1} and γ is a closed path.

Assume for contradiction that γ_3 is not a loop. Then either $\gamma_i \subset Y$ and we set $D_i = \gamma_i$, or $D_i = C_i^c$ for $i = 2, 3$, and thus, we have an embedded rectangle $D_2 \times D_3$ in Y . As in the previous case, γ_1 must form another side of $D_2 \times D_3$, contradicting the fact that $D_2 \times D_3$ is embedded. Thus, γ_3 must be a loop and γ is a closed path in Y° . \square

The following simple observation plays a crucial role in the proof of Theorem 1.1.

Lemma 2.8. *The nonpositively-curved cube complex $\mathfrak{C}(Y \rightarrow S_\Gamma)$ is connected when Y is connected, and $|\mathfrak{C}(Y \rightarrow S_\Gamma)^{(0)}| = |Y^{(0)}|$. Hence $\deg(\mathfrak{C}(Y \rightarrow S_\Gamma) \rightarrow S_\Gamma) = |Y^{(0)}|$.*

Proof. The first assertion is immediate from the construction of $\mathfrak{C}(Y \rightarrow S_\Gamma)$. The second follows from the fact that $\mathfrak{C}(Y \rightarrow S_\Gamma)$ contains Y and does not contain any 0-cube not in Y . This, together with the fact that S_Γ has a single 0-cube, implies the third assertion. \square

2.2.5. *Structure of \tilde{S}_Γ .* Let Γ be a finite simplicial graph, and let S_Γ be the Salvetti complex of A_Γ . Recall that for each cube $c \rightarrow S_\Gamma$, the attaching map identifies opposite faces of c , so that the image of c is an embedded $\dim c$ -torus. Such a torus is a *standard torus* of S_Γ , and a standard torus $T \subseteq S_\Gamma$ is *maximal* if it is not properly contained in a standard torus. (We emphasize that 0-cubes and 1-cubes in S_Γ are also standard tori.) The inclusion $T_n \rightarrow S_\Gamma$ of the standard n -torus T_n lifts to an isometric embedding $\tilde{T}_n \rightarrow \tilde{S}_\Gamma$ of universal covers. In fact, \tilde{T}_n has a natural $\text{CAT}(0)$ cubical structure obtained by pulling back the cell structure on T_n : as a cube complex, \tilde{T}_n is the standard tiling of \mathbb{E}^n by unit n -cubes. Such a subcomplex $\tilde{T}_n \subseteq \tilde{S}_\Gamma$ is a *standard flat* (and a *maximal standard flat* if T_n is a maximal standard torus). Since the inclusion $T_n \hookrightarrow S_\Gamma$ is easily seen to be a local isometry, $\tilde{T}_n \subseteq \tilde{S}_\Gamma$ is a convex subcomplex.

3. VIRTUALLY SPECIAL GROUPS

This section presents a proof of Theorem 1.1. To this end, let Γ be a simplicial graph, let A_Γ be the corresponding raAg, and let S_Γ be the corresponding Salvetti complex. The *label* of a 1-cube e of \tilde{S}_Γ is the 1-cube of S_Γ to which e maps. For each hyperplane H of \tilde{S}_Γ , the 1-cubes dual to H all have the same label, which we call the *label* of H . Let H be labeled by

a. Then \tilde{S}_Γ has a convex subcomplex $P(H) = H \times L_a$, where L_a is a convex combinatorial line all of whose 1-cubes are labelled by a .

Lemma 3.1. *Let $K \subset \tilde{S}_\Gamma$ be a convex subcomplex and let H be a hyperplane such that $H \cap K \neq \emptyset$. Then $P(H) \cap K = (H \cap K) \times L'_a$, where L'_a is a combinatorial subinterval of L_a of length at least one.*

Proof. This follows from convexity of K and Lemma 2.5 of [CS11]. \square

The complex $\mathcal{F}_K(H) = P(H) \cap K$ is the *frame* of H in K and is shown in Figure 4.




FIGURE 4. A frame showing **two translates** of the hyperplane H .

Proof of Theorem 1.1. Let \tilde{v} be a lift of v to \tilde{S}_Γ and let $g \in A_\Gamma \setminus \{1\}$ and let $d(\tilde{v}, g\tilde{v}) = n \geq 1$. Let K be the convex hull of $\{\tilde{v}, g\tilde{v}\}$. There exist a set of hyperplanes H_1, \dots, H_k with the property that each H_i separates $\tilde{v}, g\tilde{v}$, such that $H_{i\pm 1} \cap K$, for $2 \leq i \leq k-1$, lie in two distinct connected components of $K \setminus H_i \cap K$. By passing to a subset if necessary, we can assume that $\mathcal{F}_K(H_i) \neq \mathcal{F}_K(H_j)$ for $i \neq j$. For each i , we have $\mathcal{F}_K(H_i) \cong (H_i \cap K) \times L'_i$, where $L'_i \cong [0, \ell_i]$ with $\ell_i \geq 1$.

By definition of a frame, the fact that $\mathcal{F}_K(H_i) \neq \mathcal{F}_K(H_{i+1})$ and that $\text{Stab}(H_i)$ is the centralizer of the generator labeling L_i , the labels of L'_i and $L'_{i\pm 1}$ are distinct for all i . Moreover, since $H_i \cap H_{i\pm 1} = \emptyset$, no 1-cube of $L'_{i\pm 1}$ lies in the $(H_i \cap K)$ -factor of $\mathcal{F}_K(H_i)$ or vice versa. This fact together with the fact that each hyperplane intersecting K must separate \tilde{v} from $g\tilde{v}$, implies that $\mathcal{F}_K(H_i) \cap \mathcal{F}_K(H_{i+1}) = (H_i \cap K) \times \{\ell_i\} \cap (H_{i+1} \cap K) \times \{0\}$. Indeed, the intersection involves the 0 and ℓ_i factors only since for any three pairwise non-intersecting hyperplanes of K , some pair is separated by the third. Since H_i separates $H_{i'}$ from $H_{i''}$ when $i' < i < i''$, we have that $\mathcal{F}_K(H_i) \cap \mathcal{F}_K(H_{i'}) = \emptyset$ if $|i - i'| > 1$.

Finally, we can make the above choices so that $\mathcal{F}_K(H_i) \cap \mathcal{F}_K(H_{i+1}) \neq \emptyset$. Indeed, were the intersection empty, then by convexity of frames, there would be a hyperplane H separating $\mathcal{F}_K(H_i)$ from $\mathcal{F}_K(H_{i+1})$ and hence separating H_i from H_{i+1} ; such an H could be included in our original sequence and its frame in K in our original sequence of frames. Moreover, by a similar argument, \tilde{v} is in the $\mathcal{F}_K(H_1)$ and $g\tilde{v}$ is in $\mathcal{F}_K(H_k)$. Hence, without loss of generality, there is an embedded piecewise-geodesic combinatorial path $\gamma = Q_1 L'_1 \cdots Q_k L'_k Q_{k+1}$ in K joining \tilde{v} to $g\tilde{v}$, where $Q_i \subset (H_i \cap K) \times \{0\}$ for $i \leq k$, $Q_{k+1} \subset (H_k \cap K) \times \{\ell_k\}$ and L'_i is chosen within its parallelism class so that the above concatenation exists. Note that $\sum_i \ell_i \leq n$.

Let $P = \cup_{i=1}^k \mathcal{F}_K(H_i)$, so P is connected and contains γ . For each i , let $\rho_i : \mathcal{F}_K(H_i) \rightarrow \overline{H_i \cap K} \times L'_i$ be the cubical quotient induced by identifying the endpoints of each 1-cube of H_i

and folding as necessary. More precisely, for each i , we identify the endpoints of each 1-cube of $H_i \cap K$. This induces a cubical quotient $H_i \rightarrow \overline{H}_i$. We then *fold*, i.e. identify cubes c_1, c_2 for which $c_1 \cup c_2 \rightarrow S_\Gamma$ is not locally injective. (This straightforwardly generalizes Stallings folding for maps of graphs.) The resulting (folded) quotient is $\overline{H}_i \cap \overline{K}$, and ρ_i is the induced map acting as the identity on L'_i .

Since ρ_i and ρ_{i+1} agree on $\mathcal{F}_K(H_i) \cap \mathcal{F}_K(H_{i+1})$, these maps can be pasted together to form a quotient $\rho : P \rightarrow \overline{P}$ with \overline{P} a nonpositively-curved cube complex. Note that the restriction of $\tilde{S}_\Gamma \rightarrow S_\Gamma$ descends to a locally injective cubical map $\overline{P} \rightarrow S_\Gamma$.

We claim that $\rho \circ \gamma$ is a path in \overline{P} that contains every 0-cube and has distinct endpoints. This follows from the fact that $\rho(L'_i) \cap \rho(L'_j)$ is a single 0-cube if $|i - j| = 1$ and is otherwise empty if $i \neq j$, and ρ is injective on each L'_i . Since γ passes through each 1-cube of $\cup_i L'_i$ exactly once, and the image of each Q_i maps to a wedge of circles in \overline{P} , it follows that γ has the desired properties. Hence, $|\overline{P}^{(0)}| \leq n + 1$.

We would like to finish by applying Lemma 2.8 to \overline{P} . However, the constructed cube-complex, \overline{P} , is not necessarily locally convex in S_Γ . To fix this, let $s = c_i \times c_{i+1}$ be a 2-cube of K such that c_i is a 1-cube in $\mathcal{F}_K(H_i)$ and c_{i+1} is a 1-cube in $\mathcal{F}_K(H_{i+1})$, as in Figure 5.




FIGURE 5. Two frames and a missing 2-cube.

Let $\bar{s} = \rho(c_i) \times \rho(c_{i+1})$. Without loss of generality, $c_{i+1} \subset H_{i+1} \times \{0\}$. Indeed, the generators labeling L'_i and L'_{i+1} do not commute, so at most one of c_i, c_{i+1} is in L'_i, L'_{i+1} . Hence \bar{s} is either a cylinder or a torus. In the latter case, glue \bar{s} to \overline{P} along $\rho(c_i) \cup \rho(c_{i+1})$, noting that we do not add 0-cubes in so doing and moreover, we do not add 1-cubes. Hence no missing corners are introduced.

In the former case, the label of c_{i+1} corresponds to a generator of A_Γ that commutes with the generator labeling L'_i , and hence $c_{i+1} \subset (H_i \cap K) \times \{\ell_i\} \cap (H_{i+1} \cap K) \times \{0\}$. Therefore, $s \subset \mathcal{F}_K(H_i)$. We conclude that the quotient ρ extends to a quotient $K \rightarrow \overline{K}$ such that $\overline{P} \subseteq \overline{K}$ and the restriction of $\tilde{S}_\Gamma \rightarrow S_\Gamma$ to K descends to a local isometry $\overline{K} \rightarrow S_\Gamma$. Moreover, since \overline{K} is formed from \overline{P} by attaching 2-tori as above, and adding higher-dimensional tori when their 2-skeleta appear, we see that $|\overline{K}^{(0)}| = |\overline{P}^{(0)}| \leq n + 1$. Hence, $\mathcal{C}(\overline{K} \rightarrow S_\Gamma)$ is a cover of S_Γ of degree at most $n + 1$, by Lemma 2.8, such that $\gamma : [0, n] \rightarrow \tilde{S}_\Gamma \rightarrow S_\Gamma$ lifts to a non-closed path in $\mathcal{C}(\overline{K} \rightarrow S_\Gamma)$, and the proof is complete. \square

4. SPECIAL LINEAR GROUPS

4.1. The upper bound. Fix a generating set for $\mathrm{SL}_k(\mathbb{Z})$. Let g be a nontrivial element in the word-metric ball of radius n in $\mathrm{SL}_k(\mathbb{Z})$. Since $Z(\mathrm{SL}_k(\mathbb{Z}))$ is finite, we may assume that $g \notin Z(\mathrm{SL}_k(\mathbb{Z}))$. Thus there exists an off-diagonal entry of g that is not zero or two diagonal entries with non-zero difference. Select α so that it is one of these values and non-zero. By [BRK12, Proposition 4.1], there exists some prime p with

$$p \leq Cn$$

for some fixed constant C , where the image of g in $\mathrm{SL}_k(\mathbb{Z}/p\mathbb{Z})$ is not central (that is, α does not vanish in $\mathbb{Z}/p\mathbb{Z}$). The group $\mathrm{SL}_k(\mathbb{Z}/p\mathbb{Z})$ has subgroup

$$\Delta := \left\{ \begin{pmatrix} * & \cdots & * & * \\ \vdots & & \vdots & \vdots \\ * & \cdots & * & * \\ 0 & \cdots & 0 & * \end{pmatrix} \in \mathrm{SL}_k(\mathbb{Z}/p\mathbb{Z}) : * \text{ entries are arbitrary} \right\}.$$

Using a dimension counting argument, it is straightforward to see that the index of Δ in $\mathrm{SL}_k(\mathbb{Z}/p\mathbb{Z})$ is bounded above by $C'p^{k-1}$ where C' depends only on k . Since $\mathrm{SL}_k(\mathbb{Z}/p\mathbb{Z})$ maps onto a simple group with kernel $Z(\mathrm{SL}_k(\mathbb{Z}/p\mathbb{Z}))$, it follows that the intersection of all conjugates of Δ is contained in $Z(\mathrm{SL}_k(\mathbb{Z}/p\mathbb{Z}))$ (note that Δ contains $Z(\mathrm{SL}_k(\mathbb{Z}/p\mathbb{Z}))$). Thus, there exists some conjugate of Δ that misses the image of g , which is not central, in $\mathrm{SL}_k(\mathbb{Z}/p\mathbb{Z})$. Thus, we get

$$\mathrm{RF}_{\mathrm{SL}_k(\mathbb{Z})}(n) \preceq n^{k-1}.$$

4.2. The lower bound. Here, we show that the residual finiteness growth of $\mathrm{SL}_k(\mathbb{Z})$, $k > 2$, is bounded below by n^{k-1} . Before we get into the proof, we need a lemma involving unipotent subgroups of $\mathrm{SL}_k(\mathbb{Z})$. Let $E_{i,j}(\alpha)$ be the elementary matrix with α in the i th row and j th column.

Lemma 4.1. *Let H be the subgroup of $\mathrm{SL}_k(\mathbb{Z})$ that is the $2k - 1$ dimensional generalized Heisenberg group. Set $g_n = E_{1,k}(\mathrm{lcm}(1, \dots, n))$. Then $D_H(g_n) \geq n^{k-1}$.*

Proof. Let Δ be a finite-index subgroup of H that does not contain g_n . Set $d = 2k - 3$. By using $E_{i,j}(1)$ as a Mal'cev basis we may associate to Δ a matrix $\{m_{i,j}\}$ (see [GSS88, Lemma 2.3]) with $[\Gamma : \Delta] = \prod_{i=1}^d m_{i,i}$ where $(E_{1,k}(1))^{m_{1,1}} \in \Delta$ and, in fact, we have $k - 2$ conditions:

$$\begin{aligned} m_{1,1} & \text{ divides } m_{2,2}m_{k,k}, \\ m_{1,1} & \text{ divides } m_{3,3}m_{k+1,k+1}, \\ & \vdots \\ m_{1,1} & \text{ divides } m_{k-1,k-1}m_{2k-3,2k-3}. \end{aligned}$$

Thus, $\prod_{i=1}^d m_{i,i} \geq m_{1,1}^{k-1}$. As $g_n \notin \Delta$, we have that $m_{1,1}$ does not divide $\mathrm{lcm}(1, \dots, n)$, i.e. $m_{1,1} > n$, so

$$D_H(g_n) \geq n^{k-1},$$

as desired. \square

We can now prove the lower bound. We begin by following the first part of the proof of [BR10, Theorem 2.6]. By [LMR00, Theorem A], there exists a finite generating set, S , for $\mathrm{SL}_k(\mathbb{Z})$ (see also Riley [Ril05]) and a $C > 0$ satisfying

$$\| - \|_S \leq C \log(\| - \|_1),$$

where $\| - \|_1$ is the 1-operator norm for matrices. Thus, as $\log(\|E_{1,k}(\mathrm{lcm}(1, \dots, n))\|_1) = \log(\mathrm{lcm}(1, \dots, n)) + 1 \approx n$ by the prime number theorem, the elementary matrix may be written in terms of at most Cn elements from S . The matrix $g_n := E_{1,k}(\mathrm{lcm}(1, \dots, n))$ in $\mathrm{SL}_k(\mathbb{Z})$ is our candidate.

g_n is contained in $H \leq \mathrm{SL}_k(\mathbb{Z})$ as in Lemma 4.1. It follows then that

$$D_{\mathrm{SL}_k(\mathbb{Z})}(g_n) \geq D_H(g_n) \geq n^{k-1}.$$

Thus, $\mathrm{RF}_{\mathrm{SL}_k(\mathbb{Z})}(n) \succeq n^{k-1}$, as desired.

REFERENCES

- [AFW13] M. Aschenbrenner, S. Friedl, and H. Wilton. 3-manifold groups. *ArXiv 1205.0202*, pages 1–149, 2013.
- [Ago12] Ian Agol. The virtual Haken conjecture. *arXiv:1204.2810*, 2012. Primary article by Ian Agol, with an appendix by Ian Agol, Daniel Groves, and Jason Manning.
- [Big01] Stephen J. Bigelow. Braid groups are linear. *J. Amer. Math. Soc.*, 14(2):471–486 (electronic), 2001.
- [BR10] Khalid Bou-Rabee. Quantifying residual finiteness. *J. Algebra*, 323(3):729–737, 2010.
- [BR11a] Khalid Bou-Rabee. Approximating a group by its solvable quotients. *New York J. Math.*, 17:699–712, 2011.
- [BR11b] Khalid Bou-Rabee. Approximating a group by its solvable quotients. *New York J. Math.*, 17:699–712, 2011.
- [Bri91] M.R. Bridson. Geodesics and curvature in metric simplicial complexes. In E. Ghys, A. Haefliger, and A. Verjovsky, editors, *Group theory from a geometrical viewpoint*, Proc. ICTP, Trieste, Italy, pages 373–463. World Scientific, Singapore, 1991.
- [BRK12] Khalid Bou-Rabee and Tasho Kaletha. Quantifying residual finiteness of arithmetic groups. *Compos. Math.*, 148(3):907–920, 2012.
- [BRM] K. Bou-Rabee and D. B. McReynolds. Extremal behavior of divisibility functions. (*submitted*).
- [BRM10] K. Bou-Rabee and D. B. McReynolds. Bertrand’s postulate and subgroup growth. *J. Algebra*, 324(4):793–819, 2010.
- [BRM11] K. Bou-Rabee and D. B. McReynolds. Asymptotic growth and least common multiples in groups. *Bull. Lond. Math. Soc.*, 43(6):1059–1068, 2011.
- [Bus09] N. V. Buskin. Efficient separability in free groups. *Sibirsk. Mat. Zh.*, 50(4):765–771, 2009.
- [CD95] Ruth Charney and Michael W. Davis. Finite $K(\pi, 1)$ s for Artin groups. In *Prospects in topology (Princeton, NJ, 1994)*, volume 138 of *Ann. of Math. Stud.*, pages 110–124. Princeton Univ. Press, Princeton, NJ, 1995.
- [Cha07] Ruth Charney. An introduction to right-angled Artin groups. *Geometriae Dedicata*, 125:141–158, 2007.
- [Che00] Victor Chepoi. Graphs of some CAT(0) complexes. *Adv. in Appl. Math.*, 24(2):125–179, 2000.
- [CS11] Pierre-Emmanuel Caprace and Michah Sageev. Rank rigidity for CAT(0) cube complexes. *Geom. Funct. Anal.*, 21:851–891, 2011.
- [CW02] Arjeh M. Cohen and David B. Wales. Linearity of Artin groups of finite type. *Israel J. Math.*, 131:101–123, 2002.
- [Gro87] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.
- [GSS88] F. J. Grunewald, D. Segal, and G. C. Smith. Subgroups of finite index in nilpotent groups. *Invent. Math.*, 93(1):185–223, 1988.
- [Hag12] M.F. Hagen. *Geometry and combinatorics of cube complexes*. PhD thesis, McGill University, 2012.
- [HW08] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *Geom. Funct. Anal.*, 17(5):1 551–1620, 2008.

- [HW10] Frédéric Haglund and Daniel T. Wise. Coxeter groups are virtually special. *Adv. Math.*, 224(5):1890–1903, 2010.
- [HW12] Frédéric Haglund and Daniel T. Wise. A combination theorem for special cube complexes. *Ann. of Math. (2)*, 176(3):1427–1482, 2012.
- [HW13] M.F. Hagen and D.T. Wise. Cubulating hyperbolic free-by-cyclic groups: the irreducible case. *arXiv 1311.2084*, pages 1–39, 2013.
- [IBM] M. Kassabov I. Biringer, K. Bou-Rabee and F. Matucci. Intersection growth in groups. in preparation.
- [KM11] Martin Kassabov and Francesco Matucci. Bounding the residual finiteness of free groups. *Proc. Amer. Math. Soc.*, 139(7):2281–2286, 2011.
- [Kra02] Daan Krammer. Braid groups are linear. *Ann. of Math. (2)*, 155(1):131–156, 2002.
- [Lea10] Ian Leary. A metric Kan-Thurston theorem. *Preprint*, 2010.
- [Liu] Yi Liu. Virtual cubulation of nonpositively curved graph manifolds.
- [LMR00] Alexander Lubotzky, Shahar Mozes, and M. S. Raghunathan. The word and Riemannian metrics on lattices of semisimple groups. *Inst. Hautes Études Sci. Publ. Math.*, (91):5–53 (2001), 2000.
- [NR03] G. A. Niblo and L. D. Reeves. Coxeter groups act on CAT(0) cube complexes. *J. Group Theory*, 6(3):399–413, 2003.
- [OW11] Yann Ollivier and Daniel T. Wise. Cubulating random groups at density $< \frac{1}{6}$. *Trans. Amer. Math. Soc.*, 363:4701–4733, 2011.
- [Pat12] Priyam Patel. On a theorem of Peter Scott. *Proc. Amer. Math. Soc.*, 2012. To appear.
- [PW] Piotr Przytycki and Daniel T. Wise. Mixed 3-manifolds are virtually special. pages 1–24. Available at arXiv:1205.6742.
- [PW11] Piotr Przytycki and Daniel T. Wise. Graph manifolds with boundary are virtually special. *ArXiv eprint 1110.3513*, page 24, 2011.
- [Ril05] T. R. Riley. Navigating in the Cayley graphs of $SL_N(\mathbb{Z})$ and $SL_N(\mathbb{F}_p)$. *Geom. Dedicata*, 113:215–229, 2005.
- [Riv12] Igor Rivin. Geodesics with one self-intersection, and other stories. *Adv. Math.*, 231(5):2391–2412, 2012.
- [Sag95] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.
- [Sta83] John R. Stallings. Topology of finite graphs. *Inventiones mathematicae*, 71(3):551–565, 1983.
- [Wis] Daniel T. Wise. The structure of groups with a quasiconvex hierarchy. 205 pp. Preprint 2011.
- [Wis12] Daniel T. Wise. From riches to RAAGs: 3-manifolds, right-angled Artin groups, and cubical geometry. In *Lecture notes*, NSF-CBMS Conference, CUNY Graduate Center, New York, 2012.

SCHOOL OF MATHEMATICS, UNIVERSITY OF MINNESOTA—TWIN CITIES, MINNEAPOLIS, MINNESOTA, USA
E-mail address: khalid.math@gmail.com

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF MICHIGAN, ANN ARBOR, MICHIGAN, USA
E-mail address: markfhagen@gmail.com

DEPARTMENT OF MATHEMATICS, PURDUE UNIVERSITY, WEST LAFAYETTE, INDIANA, USA
E-mail address: patel1376@math.purdue.edu