The asymptotic density of finite-order elements in virtually nilpotent groups

Pallavi Dani

Department of Mathematics, University of Oklahoma, Norman, OK 73019-0315, USA

Abstract

Let Γ be a finitely generated group with a given word metric. The *asymptotic density* of elements in Γ that have a particular property P is the limit, as $r \to \infty$, of the proportion of elements in the ball of radius r which have the property P. We obtain a formula to compute the asymptotic density of finite-order elements in any virtually nilpotent group. Further, we show that the spectrum of numbers that occur as such asymptotic densities consists of exactly the rational numbers in $[0, 1)$.
© 2007 Elsevier Inc. All rights reserved.

Keywords: Asymptotic density; Virtually nilpotent groups; Finite-order

1. Introduction

Let Γ be a finitely generated infinite group. If P is a property that elements of Γ may have, such as having finite order, having cyclic centraliser or having a root, it is natural to ask: What is the density of elements of Γ that have the property P?

To make this more precise, fix a finite set S of generators for Γ. Given two elements g and h in Γ, set $d(g, h)$ to be the length of the shortest word in S representing $g^{-1}h$. This defines the *word metric* on Γ, which makes Γ into a discrete, proper metric space. For $r \geqslant 1$, let $B_S(r)$ denote the ball of radius r centred at the identity of Γ with respect to this metric. Let $E_S(r)$ denote the set of elements with property P in the ball of radius r.

E-mail address: pdani@math.ou.edu.

General Problem. *Compute the asymptotics of $|E_S(r)|$. In particular, find*

$$\mathcal{D}(\Gamma, S) = \lim_{r \to \infty} \frac{|E_S(r)|}{|B_S(r)|}$$

if this limit exists.

$\mathcal{D}(\Gamma, S)$ is the *asymptotic density* of elements in Γ which have the property P. In this paper we study the asymptotic density of finite-order elements in the class of virtually nilpotent groups (i.e. groups containing a nilpotent subgroup of finite index). In Theorem 1.1 and Corollary 1.2 we obtain a formula to compute $\mathcal{D}(\Gamma, S)$ for any virtually nilpotent group Γ.

It is worth pointing out that if Γ is actually a nilpotent group, the finite-order elements of Γ form a finite subgroup, so that $\mathcal{D}(\Gamma, S) = 0$ for any generating set S. However, the situation is very different when one passes to *virtually* nilpotent groups. For example, Theorem 1.1 can be used to show that the densities of finite-order elements in the square and triangle reflection groups in the Euclidean plane are 1/4 and 1/3, respectively. In fact, we prove in Theorem 1.3 that every rational number in $[0, 1)$ occurs as the density of finite-order elements in some virtually nilpotent group. This is noteworthy in light of the fact that in many results of this nature in the literature the limit is always either 0 or 1. A number of such examples is listed in [12]. The authors themselves give an example exhibiting "intermediate" density; they show that the union of all proper retracts in the free group on two generators has asymptotic density $6/\pi^2$.

The phenomenon of positivity of $\mathcal{D}(\Gamma, S)$ is not restricted to groups of polynomial growth. In fact there exist infinite torsion groups with intermediate [8] and even exponential [1,15] growth. (For these $\mathcal{D}(\Gamma, S) = 1$.)

The quantity $\mathcal{D}(\Gamma, S)$ is not a *geometric* property; it may change drastically under quasi-isometry. For example, every virtually nilpotent group contains a nilpotent subgroup of finite index (for which $\mathcal{D} = 0$). However, the large-scale geometry of nilpotent Lie groups plays an important role in the methods used to study $|E_S(r)|$.

The idea of studying groups from a statistical viewpoint was introduced by Gromov, when he indicated that "almost every" group is word-hyperbolic. Since then the notions of generic group theoretic properties and generic-case behavior have been extensively studied by Arzhantseva, Champetier, Kapovich, Myasnikov, Ollivier, Ol'shanskii, Rivin, Schupp, Schpilrain, Zuk and others (see [12] and the references therein).

1.1. Virtually nilpotent groups

Our goal is to compute $\mathcal{D}(\Gamma, S)$ for virtually nilpotent groups Γ. First consider the following geometric case.

Let G be a connected, simply connected nilpotent Lie group endowed with a left-invariant Riemannian metric. Its group of isometries is given by $\operatorname{Isom} G = G \rtimes C$, where G acts by left multiplication and C is the group of automorphisms of G which preserve the metric. Let Γ be a discrete, cocompact subgroup of $\operatorname{Isom} G$. Auslander [4] generalised Bieberbach's First Theorem to show that Γ has a unique maximal normal nilpotent subgroup Λ, which is torsion-free, and that the quotient $F = \Gamma/\Lambda$ is finite. (In particular Γ is virtually nilpotent.)

This information determines a representation $\rho : F \to \operatorname{Aut}(\mathfrak{g})$, where $\mathfrak{g}$ is the Lie algebra of G. (See Section 3.) If A is an element of F, then the automorphism $\rho(A)$ has eigenvalues for its

action on $\mathfrak{g}$. The eigenvalues determined by elements of F in this way depend only on the isomorphism type of Γ. The following theorem computes the asymptotics of $E_S(r)$ and gives a formula for $\mathcal{D}(\Gamma, S)$ in terms of the above eigenvalues.

Theorem 1.1. *Retaining the above notation, let S be a finite set of generators for Γ and let $E_S(r)$ denote the set of finite-order elements in the ball of radius r in the word metric. Let $\mathfrak{g} = \mathfrak{g}^1 \supset \mathfrak{g}^2 \supset \cdots \supset \mathfrak{g}^{k+1} = 0$ be the lower central series of $\mathfrak{g}$ and let $\pi : \Gamma \to F$ denote the projection map.*

Then there exists $c > 0$ such that for any $A \in F$, if $\mathfrak{h}$ denotes the 1-eigenspace of $\rho(A)$, then

$$\left| \pi^{-1}(A) \cap E_S(r) \right| \leqslant c r^{d-p} \tag{1}$$

where

$$d = \sum_{i=1}^{k} i \cdot \text{rank}\left(\mathfrak{g}^i / \mathfrak{g}^{i+1}\right) \quad and \quad p = \sum_{i=1}^{k} i \cdot \text{rank}\left(\mathfrak{h} \cap \mathfrak{g}^i / \mathfrak{h} \cap \mathfrak{g}^{i+1}\right).$$

Further,

$$\mathcal{D}(\Gamma, S) = \frac{m}{|F|} \tag{2}$$

where m is the number of elements of $\rho(F)$ that do not have 1 as an eigenvalue.

In particular, $\mathcal{D}(\Gamma, S)$ is independent of the generating set S, so we may write $\mathcal{D}(\Gamma)$ instead of $\mathcal{D}(\Gamma, S)$.

Dekimpe and Igodt [6] show that every finitely generated virtually nilpotent group has a subgroup of finite index that arises as in the geometric case. In particular, they show (see Section 3) that any virtually nilpotent group Γ has a unique maximal finite normal subgroup, say Q, and Γ/Q acts geometrically on a connected, simply connected nilpotent Lie group. So $\mathcal{D}(\Gamma/Q)$ can be computed using Theorem 1.1. Further, we have the following result.

Corollary 1.2. *Let Γ be an arbitrary finitely generated virtually nilpotent group with maximal finite normal subgroup Q. Then $\mathcal{D}(\Gamma, S) = \mathcal{D}(\Gamma/Q)$ for any generating set S of Γ.*

The formula in Theorem 1.1 makes it very easy to compute $\mathcal{D}(\Gamma)$ using algebraic data associated with Γ. A large class of examples is provided by *crystallographic groups*, i.e. groups acting properly discontinuously and cocompactly on Euclidean space. These groups are virtually abelian (by Bieberbach's First Theorem), and hence virtually nilpotent. There are 17, 230, and 4783 crystallographic groups in dimensions 2, 3, and 4, respectively. These are available as libraries designed for use with the computer algebra software GAP [19]. The results of the computation of $\mathcal{D}(\Gamma)$ for these groups (obtained using GAP) are summarised in Appendix A.

Theorem 1.1 shows that $\mathcal{D}(\Gamma)$ is always a rational number. In the following theorem we address the question of which rational numbers in $[0, 1]$ can occur.

Theorem 1.3. *Given any rational number p/q with $0 \leqslant p/q < 1$, there exists a crystallographic group Γ such that $\mathcal{D}(\Gamma) = p/q$.*

This is proved in Section 11 by explicitly constructing finite subgroups of $Gl(n, \mathbb{Z})$ in which exactly $(q - p)/q$ of the elements have eigenvalue 1.

The paper is organised as follows. Sections 2–6 contain definitions and background on nilpotent Lie groups. In particular, Section 4 describes certain useful "polynomial" coordinate systems for nilpotent Lie groups. Section 6 contains some technical lemmas about polynomial coordinates.

The proof of Theorem 1.1 is contained in Sections 7–9. In Section 7 we show that a finite-order element of length r in Γ fixes a point in a certain ball centered at the identity in G. The key is to use now the geometry of G to estimate the number of fixed sets of torsion elements that intersect this ball. In Section 8, an argument about volumes of balls in G yields the upper bound (1) in Theorem 1.1 for the number of torsion elements in any coset $\pi^{-1}(A)$ of Λ. From this bound it follows that if 1 is an eigenvalue of $\rho(A)$, the torsion in $\pi^{-1}(A)$ does not contribute to $\mathcal{D}(\Gamma, S)$. In Section 9 an inductive argument shows that if 1 is not an eigenvalue of $\rho(A)$, then the coset $\pi^{-1}(A)$ consists entirely of torsion elements. Theorem 1.1 then follows from the fact that the asymptotic density of a coset of Λ in Γ is $1/|F|$.

The proof of Corollary 1.2 appears in Section 10. Finally, in Section 11 we construct examples to prove Theorem 1.3 and also investigate $\mathcal{D}(\Gamma)$ for some virtually nilpotent groups which are not virtually abelian.

2. Definitions and basic facts

2.1. Nilpotent Lie groups and Lie algebras

In this section we recall some background material, which can be found, for example, in [5] or [7]. Recall that the *lower central series* for a Lie algebra $\mathfrak{g}$ is defined by

$$\mathfrak{g}^1 = \mathfrak{g}, \qquad \mathfrak{g}^{i+1} = [\mathfrak{g}, \mathfrak{g}^i] = \mathbb{R}\text{-span}\{[X, Y]: X \in \mathfrak{g}, \ Y \in \mathfrak{g}^i\} \quad \text{for } i \geqslant 1.$$

Then $\mathfrak{g}$ is said to be *nilpotent* if $\mathfrak{g}^{k+1} = \{0\}$ for some k. If, in addition, $\mathfrak{g}^k$ is non-trivial, then $\mathfrak{g}$ is called a k-*step* nilpotent Lie algebra.

The lower central series for a group G is given by $G^1 = G$, and $G^{i+1} = [G, G^i]$ and G is nilpotent if its lower central series is finite. If $G^{k+1} = \{1\}$, with G^k non-trivial, then G is called a k-*step* nilpotent group. The Lie algebra of a connected nilpotent Lie group is nilpotent.

A *Lie subgroup* of G is a subgroup which is a submanifold of the underlying manifold of G. If G is connected, the subgroups G^i are Lie subgroups and the Lie algebra of G^i is $\mathfrak{g}^i$. Thus G is k-step nilpotent if and only if $\mathfrak{g}$ is. For each i, the subgroup G^{i+1} is normal in G^i and the quotients G^i/G^{i+1} are abelian.

If G is a connected, simply connected nilpotent Lie group, the exponential map, $\exp: \mathfrak{g} \to G$, is an analytic diffeomorphism. Denote its inverse by log. Define a map $*: \mathfrak{g} \times \mathfrak{g} \to \mathfrak{g}$ by

$$X * Y = \log(\exp X \exp Y). \tag{3}$$

The *Baker–Campbell–Hausdorff* formula expresses $X * Y$ as a universal power series which involves commutators in X and Y. While the general term cannot be expressed in closed form, the low-order terms in the formula are well known:

$$X * Y = X + Y + \frac{1}{2}[X, Y] + \frac{1}{12}[X, [X, Y]] - \frac{1}{12}[Y, [X, Y]]$$

$$- \frac{1}{48}\big[Y, [X, [X, Y]]\big] - \frac{1}{48}\big[X, [Y, [X, Y]]\big] + (\text{commutators in} \geqslant 5 \text{ terms}). \quad (4)$$

If G is k-step nilpotent, then commutators in more than k terms are trivial, which makes this a finite sum.

2.1.1. Automorphisms and isometries

An automorphism A of G leaves invariant the groups G^i. Further, A satisfies the relation $A \circ \exp = \exp \circ dA$. The fixed set of A is the image in G of the 1-eigenspace of dA under the exponential map. It is a Lie subgroup of G.

Let G be endowed with a left-invariant Riemannian metric. Its group of isometries is given by $\operatorname{Isom} G = G \rtimes C$, where G acts by left multiplication and C is the group of automorphisms of G which preserve the inner product at the identity. We will write the action of an element $(g, A) \in \operatorname{Isom} G$ on $t \in G$ as $(g, A)(t) = gA(t)$. Any isometry fixing the identity is also an automorphism of G.

If G is abelian, then $G = \mathbb{R}^n$ with the standard inner product, where n is the dimension of G. In this case $\operatorname{Isom} G = \mathbb{R}^n \rtimes O(n)$.

The identity elements of G and $\operatorname{Aut}(G)$ will be denoted by 1 and I, respectively. We will freely make use of the identifications $(g, I) = g$ and $(1, A) = A$.

Any finite-order isometry of G has a fixed point. This follows from a more general result of Auslander in [3]. If (g, A) is a finite-order isometry with fixed point p (so that $gA(p) = p$), then

$$\big(p^{-1}, I\big)(g, A)(p, I) = \big(p^{-1}gA(p), A\big) = \big(p^{-1}p, A\big) = (1, A).$$

Thus (g, A) is conjugate to $(1, A)$ in $\operatorname{Isom} G$ and hence $\operatorname{Fix}((g, A)) = p\operatorname{Fix}(A)$.

Lemma 2.1. *Let $A = (1, A)$ be a finite-order isometry of G fixing the identity. Let K be a normal, A-invariant Lie subgroup of G with projection map $\pi : G \to G/K$. If $\bar{A}$ is the automorphism of G/K induced by A, then $\operatorname{Fix}(\bar{A}) = \pi(\operatorname{Fix}(A))$.*

Proof. Clearly $\pi(\operatorname{Fix}(A)) \subseteq \operatorname{Fix}(\bar{A})$. Now if gK is fixed by $\bar{A}$, then A leaves gK invariant. Thus $(g^{-1}, I)(1, A)(g, I)$ is a finite-order isometry (equal to $(g^{-1}A(g), A)$) leaving K invariant, which means it has a fixed point in K, say b. We now have

$$\big(g^{-1}A(g), A\big)(b) = b \quad \Rightarrow \quad g^{-1}A(g)A(b) = b \quad \Rightarrow \quad A(gb) = gb.$$

So gb is fixed by A and $gK = \pi(gb)$. Thus $\operatorname{Fix}(\bar{A}) \subseteq \pi(\operatorname{Fix}(A))$. $\square$

2.2. Quasi-isometries

A map $\phi : X \to X'$ between two metric spaces (X, d) and (X', d') is a *quasi-isometry* if there exist constants $\lambda \geqslant 1$, and $C, D \geqslant 0$, such that

$$\frac{1}{\lambda}d(x, y) - C \leqslant d'\big(\phi(x), \phi(y)\big) \leqslant \lambda d(x, y) + C$$

for all $x, y \in X$ and every point of X' is in a D-neighbourhood of $\phi(X)$.

The following classical result can be found, for example, in [11].

Theorem 2.2 *(Milnor, Efremovich, Švarc). If Γ is a group acting properly discontinuously and cocompactly by isometries on a proper geodesic metric space X, then Γ is quasi-isometric to X. More precisely, for any $x_0 \in X$, the mapping $\Gamma \to X$ given by $\gamma \mapsto \gamma(x_0)$ is a quasi-isometry.*

3. Virtually nilpotent groups

A finitely generated group is said to be *virtually nilpotent* if it has a nilpotent subgroup of finite index. *Almost crystallographic* groups, i.e. groups acting properly discontinuously and cocompactly by isometries on a connected, simply connected nilpotent Lie group, are examples of virtually nilpotent groups. This follows from the following theorem of Auslander.

Theorem 3.1. *(See [4].) If Γ is a discrete, cocompact subgroup of $\operatorname{Isom} G = G \rtimes C$, where G is a connected, simply connected nilpotent Lie group, then $\Lambda = \Gamma \cap G$ is cocompact in G and $F = \Gamma/\Lambda$ is a finite group. Further, Λ is the unique maximal normal nilpotent subgroup of Γ and it is torsion-free.*

Using the work of Kamishima, Lee, and Raymond, Dekimpe and Igodt gave an algebraic condition for a virtually nilpotent group to be almost crystallographic. It is proved in [6] that every virtually nilpotent group has a unique maximal finite normal subgroup. Further, they prove the following:
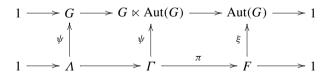
Theorem 3.2. *If Γ' is a virtually nilpotent group with maximal finite normal subgroup Q, then $\Gamma = \Gamma'/Q$ is almost crystallographic.*

This is a generalisation of Malcev's result [16] that any finitely generated, torsion-free nilpotent group can be embedded as a discrete subgroup of a nilpotent Lie group, which is unique up to isomorphism. Theorem 3.2 allows us to focus on almost crystallographic groups.

3.1. Eigenvalues

Let Γ be an almost crystallographic group acting on G, i.e. there is an injection $\psi : \Gamma \to G \ltimes \operatorname{Aut}(G)$. By Theorem 3.1, Γ has a unique maximal normal nilpotent subgroup Λ with $\psi(\Lambda) = G \cap \psi(\Gamma)$, such that $F = \Gamma/\Lambda$ is finite.

Let $\pi : \Gamma \to F$ be the projection map. There is a unique homomorphism $\xi : F \to \operatorname{Aut}(G)$ which makes the following diagram commute:

$$
\begin{array}{ccccccccc}
1 & \longrightarrow & G & \longrightarrow & G \ltimes \operatorname{Aut}(G) & \longrightarrow & \operatorname{Aut}(G) & \longrightarrow & 1 \\
 & & \uparrow{\psi} & & \uparrow{\psi} & & \uparrow{\xi} & & \\
1 & \longrightarrow & \Lambda & \longrightarrow & \Gamma & \xrightarrow{\ \pi\ } & F & \longrightarrow & 1
\end{array}
$$

A diagram-chase shows that ξ is injective. In other words, F can be realised as a group of automorphisms of G. We obtain an injective homomorphism $\rho : F \to \operatorname{Aut}(\mathfrak{g})$ by composing ξ with the map that assigns to each automorphism in $\operatorname{Aut}(G)$ its derivative. If $A \in F$, the *eigenvalues* of A are the eigenvalues of the automorphism $\rho(A)$ for its action on $\mathfrak{g}$.

The fact that these eigenvalues are well defined follows from a theorem of Lee and Raymond in [14] which says that any two isomorphic almost crystallographic groups acting on G are conjugate by an element of $G \ltimes \text{Aut}(G)$. Indeed if $\psi' : \Gamma \to G \ltimes \text{Aut}\, G$ is another injection, giving rise to the homomorphism $\xi' : F \to \text{Aut}(G)$, and the element $(g, B) \in G \ltimes \text{Aut}(G)$ conjugates $\psi(\Gamma)$ to $\psi'(\Gamma)$, then we also have $B\xi(F)B^{-1} = \xi'(F)$, which implies that the eigenvalues assigned to elements of F via ξ are the same as those assigned via ξ'.

4. Polynomial coordinates on G

For the rest of the paper, G will denote a connected, simply connected nilpotent Lie group. In this section we describe how G can be naturally identified with $\mathbb{R}^n$, where n is the dimension of G, so that the group structure is "polynomial" relative to the linear coordinates on $\mathbb{R}^n$. Such polynomial coordinate systems are treated, for example, in [5,10] or [20].

A map $f : V \to W$ between two vector spaces is *polynomial* if it is described by polynomials in the coordinates for some (and hence any) pair of bases. A *polynomial coordinate map* for G is a diffeomorphism $\phi : \mathbb{R}^n \to G$, such that $\log \circ \phi$ and $\phi^{-1} \circ \exp$ are polynomial maps. We start by defining a useful polynomial coordinate map on G.

Let $\mathfrak{g}$ be a nilpotent Lie algebra and let $\mathfrak{g} = \mathfrak{g}^1 \supset \mathfrak{g}^2 \supset \cdots \supset \mathfrak{g}^{k+1} = 0$ be its lower central series. We define a basis which respects this filtration of $\mathfrak{g}$.

Definition 4.1 *(Triangular basis).* Let $\{X_1, \ldots, X_n\}$ be an ordered basis for $\mathfrak{g}$ with $[X_i, X_j] = \sum_{l=1}^{n} \alpha_{ijl} X_l$. The basis is *triangular* if $\alpha_{ijl} = 0$ when $l \leqslant \max\{i, j\}$.

Example. For the three-dimensional Heisenberg Lie algebra (generated by X, Y and Z, where $[X, Y] = Z$ and all other brackets are trivial), the ordered sets $\{X, Y, Z\}$ and $\{X + Z, Y, Z\}$ are triangular bases, while the sets $\{Y, Z, X\}$ and $\{X, Y, X + Z\}$ are not.

A triangular basis can be constructed by starting with an ordered basis for $\mathfrak{g}^k$ and then successively pulling back ordered bases for the factors $\mathfrak{g}^i / \mathfrak{g}^{i+1}$, for $i < k$. If $\mathfrak{g}$ has an inner product, then the triangular basis can be chosen to be orthonormal.

Definition 4.2 *(Coordinate map on G).* Let $\{X_1, \ldots, X_n\}$ be a triangular basis for $\mathfrak{g}$. Define a map $\phi : \mathbb{R}^n \to G$ by

$$\phi(s_1, \ldots, s_n) = (\exp s_n X_n) \cdots (\exp s_1 X_1) = \exp(s_n X_n * \cdots * s_1 X_1).$$

See [5, Proposition 1.2.7] for a proof of the fact that ϕ defines a polynomial coordinate map on G.

Each vector V in $\mathfrak{g}$ is assigned a weight $\mathcal{W}$, which specifies the smallest group in the lower central series which contains V:

$$\mathcal{W}(V) = \max\{i \mid V \in \mathfrak{g}^i\}.$$

Example. In the Heisenberg Lie algebra, with triangular basis $\{X, Y, Z\}$, we have $\mathcal{W}(X) = \mathcal{W}(Y) = 1$ and $\mathcal{W}(Z) = 2$.

In a triangular basis for $\mathfrak{g}$, there are exactly $\text{rank}\,(\mathfrak{g}^i / \mathfrak{g}^{i+1})$ vectors which have weight i. With this in mind we fix the following notation.

Notation. Let $\mathfrak{g}$ be k-step nilpotent and let $\rho_i = \text{rank}\,(\mathfrak{g}^i/\mathfrak{g}^{i+1})$. A triangular basis for $\mathfrak{g}$ will be written as $\{X_{ij}\} = \{X_{ij} \mid 1 \leqslant i \leqslant k;\ 1 \leqslant j \leqslant \rho_i\}$, where $\mathcal{W}(X_{ij}) = i$. We will assume that $\{X_{ij}\}$ has the "dictionary order." Sometimes we will write the basis as $\{X_1, \ldots, X_k\}$, where X_i will mean $X_{i1}, \ldots, X_{i\rho_i}$.

We will identify G with its preimage under the polynomial coordinate map ϕ from Definition 4.2. Thus the element $s = \exp(s_{k\rho_k} X_{k\rho_k} * \cdots * s_{11} X_{11})$ of G will be written either as (s_{ij}), where it is assumed that $1 \leqslant i \leqslant k$ and $1 \leqslant j \leqslant \rho_i$, or as $(s_1, \ldots, s_k)$, where $s_i = s_{i1}, \ldots, s_{i\rho_i}$ for all i.

Finally, we will use $s_i \cdot X_i$ to denote $s_{i\rho_i} X_{i\rho_i} * \cdots * s_{i1} X_{i1}$.

5. Geometry of nilpotent Lie groups

Let G be endowed with a left-invariant Riemannian metric. The Ball–Box Theorem of Gromov and Karidi (Theorem 5.2) says that in certain polynomial coordinates, the ball of radius r about the identity in G is bounded by certain boxes with sides parallel to the coordinate axes.

Let $\{X_{ij} \mid 1 \leqslant i \leqslant k;\ 1 \leqslant j \leqslant \rho_i\}$ be an orthonormal triangular basis for the Lie algebra $\mathfrak{g}$ of G, where $\rho_i = \text{rank}\,(\mathfrak{g}^i/\mathfrak{g}^{i+1})$. Identify G with its preimage under the corresponding polynomial coordinate map, and let $B_G(1, r)$ denote the ball of radius r about the identity in G.

Definition 5.1. In the above coordinates, for any $l > 0$, define $\text{Box}(l) \subset G$ by

$$\text{Box}(l) = \big\{ (s_{ij}) \,\big|\, |s_{ij}| \leqslant (l)^i \text{ for } 1 \leqslant i \leqslant k;\ 1 \leqslant j \leqslant \rho_i \big\}.$$

This is a box in G with sides parallel to the coordinate axes. For each i, it has ρ_i sides of length $2l^i$. Note that the Lebesgue measure of this box is $2^n l^d$, where $n = \sum_{1 \leqslant i \leqslant k} \rho_i$ is the dimension of G, and $d = \sum_{1 \leqslant i \leqslant k} i\rho_i$.

Theorem 5.2 *(Ball–Box Comparison Theorem [9,13]). There exists $a > 1$, which depends only on G, such that for every $r > 1$,*

$$\text{Box}(r/a) \subset B_G(1, r) \subset \text{Box}(ra).$$

The Ball–Box Theorem can be used to estimate the volume of $B_G(1, r)$ and the distances of elements of G from the identity. First we make the following definition.

Definition 5.3. Two functions f_1 and f_2, from a set S to $\mathbb{R}$ are said to be *comparable*, denoted by $f_1(x) \sim f_2(x)$, if there exists $M > 1$ such that for all $x \in S$,

$$\frac{1}{M} f_2(x) < f_1(x) < M f_2(x).$$

There is a unique left-invariant volume form on G, up to a scalar multiple. Also, the left-invariant measure on G pulls back to Lebesgue measure on $\mathbb{R}^n$ under the polynomial coordinate map. (See [5].) This yields the following corollary.

Corollary 5.4 *(Polynomial growth [9,13]). Retaining the above notation, if* vol_G *denotes the left-invariant volume on* G, *we have*

$$\mathrm{vol}_G\big[B_G(1,r)\big] \sim r^d.$$

Let $\|s\|_G$ denote the distance of $s \in G$ from the identity, in the left-invariant metric on G. The following corollary is proved in [2].

Corollary 5.5 *(Distances in nilpotent groups [2]). Let* $s \in G$, *with* $s = (s_{ij})$ *in polynomial coordinates. Then*

$$\|s\|_G \sim \max_{i,j}\big\{|s_{ij}|^{1/i}\big\}.$$

If G^l is a group in the lower central series of G, the metric on G induces a left-invariant metric on G/G^l. (The inner product at the identity is obtained by identifying $\mathfrak{g}/\mathfrak{g}^l$ with $\mathfrak{g}^\perp$). The corresponding distances are related as follows:

Corollary 5.6 *(Distances in quotients). Let* $\pi_l : G \to G/G^l$ *be the projection map. Then there exists a constant* $\delta = \delta(G,l)$, *such that for any* $s \in G$,

$$\big\|\pi_l(s)\big\|_{G/G^l} \leqslant \delta\|s\|_G.$$

Proof. If $\{X_{ij} \mid 1 \leqslant i \leqslant k; \ 1 \leqslant j \leqslant \rho_i\}$ is an orthonormal triangular basis for $\mathfrak{g}$ then $\{d\pi_l(X_{ij}) \mid 1 \leqslant i \leqslant l-1; \ 1 \leqslant j \leqslant \rho_i\}$ is an orthonormal triangular basis for $\mathfrak{g}/\mathfrak{g}^l$, in the induced left-invariant metric on G/G^l. In the corresponding polynomial coordinates, π_l is given by $(s_1, \ldots, s_k) \mapsto (s_1, \ldots, s_{l-1})$, and the result follows from Corollary 5.5. $\square$

6. More on polynomial coordinates

In this section we show that various functions associated with G, in particular, the group operations and automorphisms, are polynomial maps which preserve certain suitably defined weights. (See Proposition 6.3.) These results are well known (cf. [5,10] or [20]) but proofs are provided here for completeness. In Lemma 6.4 we obtain a bound on the amount that such weight preserving polynomial maps can stretch distances. These results are used in the proof of Proposition 7.2.

We start with an example:

Example. Consider the Heisenberg group with polynomial coordinates associated to the triangular basis $\{X, Y, Z\}$. Group multiplication and inversion expressed in these coordinates are given by

$$(x, y, z)(x_1, y_1, z_1) = (x + x_1, y + y_1, z + z_1 + xy_1), \tag{5}$$

$$(x, y, z)^{-1} = (-x, -y, -z + xy). \tag{6}$$

Recall that $\mathcal{W}(X) = \mathcal{W}(Y) = 1$ and $\mathcal{W}(Z) = 2$. If we assign the weight 1 to the variables x, y, x_1, and y_1 and the weight 2 to z and z_1, then on the right-hand side of both (5) and (6), the X- and Y-coordinates are sums of terms of weight 1, and the Z-coordinates are sums of terms such that the total weight of each term is 2.

Motivated by this example, we make the following definition. Let $y = \{y_i\}$ be a set of variables and let $\mathcal{W}$ be a function assigning a weight to each y_i. Then polynomials in $\{y_i\}$ can be assigned weights as follows:

$$\mathcal{W}(\alpha y_{i_1} \cdots y_{i_s}) = \mathcal{W}(y_{i_1}) + \cdots + \mathcal{W}(y_{i_s}), \quad \text{where } \alpha \text{ is any constant,}$$

$$\mathcal{W}\big(P(y)\big) = \max\big\{\mathcal{W}(\alpha y_{i_1} \cdots y_{i_s}) \,\big|\, \alpha y_{i_1} \cdots y_{i_s} \text{ is a term of } P(y)\big\}.$$

Observe that $\mathcal{W}(P + Q) \leqslant \max\{\mathcal{W}(P), \mathcal{W}(Q)\}$ and $\mathcal{W}(PQ) \leqslant \mathcal{W}(P) + \mathcal{W}(Q)$.

Definition 6.1 *(Weight-preserving map).* Let V and V' be vector spaces with bases $\mathcal{B} = \{X_1, \ldots, X_s\}$ and $\mathcal{B}' = \{X'_1, \ldots, X'_{s'}\}$, respectively. A polynomial map $f : V \to V'$ can be written, with respect to these bases, as $f(v) = (P_1(v), \ldots, P_{s'}(v))$, where $v = (v_1, \ldots, v_s) = \sum_{i=1}^{s} v_i X_i \in V$, and the P_i's are polynomials. Let $\mathcal{W}$ (respectively $\mathcal{W}'$) be a function assigning weights to the X_i's (respectively X'_i's) and define $\mathcal{W}(v_i) = \mathcal{W}(X_i)$. As described above, this induces a weight function $\mathcal{W}$ on the polynomials P_i. Then f is *weight-preserving* if $\mathcal{W}(P_l) \leqslant \mathcal{W}'(X'_l)$ for all l.

Observation 6.2. Finite sums and composites of weight-preserving polynomial maps are weight-preserving polynomial maps.

Proposition 6.3. *Let G be endowed with a polynomial coordinate system corresponding to the triangular basis $\{X_{ij} \mid 1 \leqslant i \leqslant k, \ 1 \leqslant j \leqslant \rho_i\}$ of its Lie algebra $\mathfrak{g}$, where $\mathcal{W}(X_{ij}) = i$. Then the bracket, $*$, exp, multiplication and inversion in G, and all automorphisms of G, when expressed in these coordinates, are weight-preserving polynomial maps.*

Proof. It is easy to see that the bracket is a polynomial map. To prove that it is weight-preserving, it is enough to show that for given i, j, l, and m, if α and β are polynomials with $\mathcal{W}(\alpha) \leqslant i$ and $\mathcal{W}(\beta) \leqslant l$, then $[\alpha X_{ij}, \beta X_{lm}]$ is weight-preserving. Observe that $[X_{ij}, X_{lm}] \in \mathfrak{g}^{i+l}$, so that

$$[\alpha X_{ij}, \beta X_{lm}] = \sum_{s \geqslant i+l} \alpha \beta a_{st} X_{st}.$$

where the a_{st} are structure constants which depend on X_{ij} and X_{lm}. This is weight-preserving, since $\mathcal{W}(\alpha \beta a_{st}) \leqslant i + l \leqslant s$ for all s and t.

Now, the Baker–Campbell–Hausdorff Formula (4) expresses $*$ as a finite sum involving brackets. So $*$ is a weight-preserving polynomial map as well, by Observation 6.2.

To prove that exp is a weight-preserving polynomial map, we produce polynomials Q_{ij}, with $\mathcal{W}(Q_{ij}) \leqslant i$, such that when expressed in coordinates,

$$\exp(v) = \big(Q_{11}(v), \ldots, Q_{k\rho_k}(v)\big) \tag{7}$$

for all $v \in \mathfrak{g}$. Recall that $Q_l(v) \cdot X_l$ denotes $Q_{l\rho_l}(v) X_{l\rho_l} * \cdots * Q_{l1}(v) X_{l1}$. Define

$$\psi_l(v) = Q_l(v) \cdot X_l * \cdots * Q_1(v) \cdot X_1$$

for all l. We prove that $v = \psi_k(v)$ (where $\mathfrak{g}^{k+1}$ is trivial), i.e. $\exp v = \exp \psi_k(v)$. By Definition 4.2, this is equivalent to Eq. (7). The Q_l's are chosen inductively so that $\mathcal{W}(Q_{lj}) \leqslant l$ and $\psi_l(v) - v \in \mathfrak{g}^{l+1}$, for all v.

Let $v = \sum v_{ij} X_{ij}$ be an element of $\mathfrak{g}$. Set $Q_{1j}(v) = v_{1j}$, for $1 \leqslant j \leqslant \rho_1$. Clearly, $\mathcal{W}(Q_{1j}) = 1$ and $\psi_1(v) - v = \sum_{i>1} v_{ij} X_{ij} \in \mathfrak{g}^2$.

Now assume the Q_{ij}'s for $i < l$ have been chosen, with $\psi_{l-1}(v) - v \in \mathfrak{g}^l$, say

$$\psi_{l-1}(v) - v = q_{l1}(v) X_{l1} + \cdots + q_{l\rho_l}(v) X_{l\rho_l} + \text{ an element of } \mathfrak{g}^{l+1}. \tag{8}$$

Equivalently, $\psi_{l-1}(v) - v = (0, \ldots, 0, q_{l1}, \ldots, q_{l\rho_l}, \ldots)$ (this follows from the Baker–Campbell–Hausdorff formula). Observe that $\psi_{l-1}(v) - v$ is a weight-preserving polynomial map, as it is defined in terms of $*$. Thus $\mathcal{W}(q_{lj}) \leqslant l$. Choose $Q_{lj} = -q_{lj}$ for $1 \leqslant j \leqslant \rho_l$. Then $\mathcal{W}(Q_{lj}) \leqslant l$. Moreover, using the Baker–Campbell–Hausdorff formula again, we have

$$\begin{aligned}
\psi_l(v) &= Q_{l\rho_l}(v) X_{l\rho_l} * \cdots * Q_{l1}(v) X_{l1} * \psi_{l-1}(v) \\
&= Q_{l1}(v) X_{l1} + \cdots + Q_{l\rho_l}(v) X_{l\rho_l} + \psi_{l-1}(v) + \text{ an element of } \mathfrak{g}^{l+1} \\
&= -q_{l1}(v) X_{l1} + \cdots - q_{l\rho_l}(v) X_{l\rho_l} + \psi_{l-1}(v) + \text{ an element of } \mathfrak{g}^{l+1}. \tag{9}
\end{aligned}$$

Equations (8) and (9) imply that $\psi_l(v) - v \in \mathfrak{g}^{l+1}$, completing the induction. Since $\mathfrak{g}^{k+1}$ is trivial, we have $v = \psi_k(v)$ as required.

Now let $s = (s_{ij})$ and $t = (t_{ij}) \in G$. Using Definition 4.2, multiplication and inversion can be written in terms of $*$ and $\exp$ as follows:

$$\begin{aligned}
st &= \exp(s_k \cdot X_k * \cdots * s_1 \cdot X_1 * t_k \cdot X_k * \cdots * t_1 \cdot X_1), \\
s^{-1} &= \exp\big[-(s_k \cdot X_k * \cdots * s_1 \cdot X_1)\big].
\end{aligned}$$

If A is an automorphism, dA preserves the bracket, and hence $*$. Thus

$$\begin{aligned}
A(s) &= A\big(\exp(s_{k\rho_k} X_{k\rho_k} * \cdots * s_{k1} X_{k1} * \cdots * s_{1\rho_1} X_{1\rho_1} * \cdots * s_{11} X_{11})\big) \\
&= \exp dA(s_{k\rho_k} X_{k\rho_k} * \cdots * s_{k1} X_{k1} * \cdots * s_{1\rho_1} X_{1\rho_1} * \cdots * s_{11} X_{11}) \\
&= \exp(s_{k\rho_k} dA X_{k\rho_k} * \cdots * s_{k1} dA X_{k1} * \cdots * s_{1\rho_1} dA X_{1\rho_1} * \cdots * s_{11} dA X_{11}).
\end{aligned}$$

Now for each i and j, we have $dA(X_{ij}) \in \mathfrak{g}^i$, so that

$$s_{ij} dA X_{ij} = \sum_{\substack{l \geqslant i \\ 1 \leqslant m \leqslant \rho_l}} s_{ij} \alpha_{lm} X_{lm},$$

where the α_{lm}'s are constants depending on A. Thus $s \mapsto s_{ij} dA X_{ij}$ is a weight-preserving polynomial map.

It follows from Observation 6.2 that multiplication inversion and all automorphisms are weight-preserving polynomial maps. $\square$

We now obtain a bound on the amount that a weight-preserving polynomial map can stretch distances.

Lemma 6.4. *If $P : G \to G$ is a weight-preserving polynomial map, then there exists a constant $\lambda = \lambda(P) > 0$ such that for all $y \in G$,*

$$\|P(y)\|_G \leqslant \lambda \|y\|_G.$$

Proof. By Corollary 5.5 we know that there exists $\mu > 1$ such that

$$\frac{1}{\mu} \|y\|_G \leqslant \max_{i,j}\{|y_{ij}|^{1/i}\} \leqslant \mu \|y\|_G \quad \text{and} \quad \frac{1}{\mu} \|P(y)\|_G \leqslant \max_{i,j}\{|P_{ij}(y)|^{1/i}\} \leqslant \mu \|P(y)\|_G.$$

Let $\alpha y_{i_1 j_1} \cdots y_{i_l j_l}$ be a term occurring in P_{ij} for some i and j. We omit the second subscript for convenience. The weight-preserving condition, $\mathcal{W}(P_{ij}) \leqslant i$, implies that $i_1 + \cdots + i_l \leqslant i$. Let s be such that $|y_{i_s}|^{1/i_s} = \max\{|y_{i_1}|^{1/i_1}, \ldots, |y_{i_l}|^{1/i_l}\}$. Then

$$
\begin{aligned}
|\alpha y_{i_1} \cdots y_{i_l}|^{1/i} &= |\alpha|^{1/i} \big[\big(|y_{i_1}|^{1/i_1} \big)^{i_1} \cdots \big(|y_{i_l}|^{1/i_l} \big)^{i_l} \big]^{1/i} \\
&\leqslant |\alpha|^{1/i} \big[|y_{i_s}|^{i_1/i_s} \cdots |y_{i_s}|^{i_l/i_s} \big]^{1/i} \\
&= |\alpha|^{1/i} \big[|y_{i_s}|^{(i_1 + \cdots + i_l)/i_s} \big]^{1/i} \leqslant |\alpha|^{1/i} |y_{i_s}|^{1/i_s} \leqslant |\alpha|^{1/i} \mu \|y\|_G.
\end{aligned}
$$

This, combined with the fact that $|P_{ij}(y)|^{1/i} \leqslant \sum_{\text{terms of } P_{ij}(y)} |\alpha y_{i_1} \cdots y_{i_l}|^{1/i}$, enables us to choose a constant ν such that $|P_{ij}(y)|^{1/i} \leqslant \nu \mu \|y\|_G$ for all i and j. Thus $\|P\|_G \leqslant \mu^2 \nu \|y\|_G$, and we can take $\lambda = \mu^2 \nu$. □

7. The relation between Γ and G

We now return to the set-up in Theorem 1.1. Let Γ be a discrete, cocompact subgroup of $\text{Isom} \, G = G \rtimes C$, with maximal normal nilpotent subgroup Λ and finite quotient $F = \Gamma/\Lambda$, which we identify with its image in $\text{Aut}(G)$ under the map ξ defined in Section 3.

Let S be a finite generating set for Γ and let $E_S(r)$ be the set of finite-order elements of length less than or equal to r in the word metric.

Every finite-order isometry of G has a fixed point (see Section 2.1.1). The following lemma will allow us to estimate the cardinality of $E_S(r)$ by counting fixed sets of elements in $E_S(r)$ that intersect a certain ball in G.

Lemma 7.1. *Let G, Γ, S and $E_S(r)$ be as above. Then there exists a positive constant κ such that if $(g, A) \in E_S(r)$, then the fixed set of (g, A) in G intersects $B_G(1, \kappa r)$.*

The proof relies on the following proposition, which establishes a relation between the displacement of a point under the action of a finite-order isometry fixing the identity of G, and the distance of that point from the fixed set of the isometry.

Proposition 7.2. *Let A be a finite-order isometry of G fixing the identity, with fixed subgroup H. Then there exists $K = K(A, G)$, such that for all $r > 0$, if $t \in G$ satisfies $\|t A(t^{-1})\|_G < r$, then there exists $h \in H$ such that $\|th\|_G < Kr$.*

Proof of Lemma 7.1. Let ℓ_S denote the distance from the identity in Γ. Since the map $\Gamma \to G$ given by $\gamma \mapsto \gamma(1)$ is a quasi-isometry (Theorem 2.2), there exist positive constants λ and C, such that for all $(g, A) \in \Gamma$,

$$\frac{1}{\lambda}\ell_S\big((g, A)\big) - C \leqslant \|g\|_G \leqslant \lambda\ell_S\big((g, A)\big) + C.$$

So if (g, A) is an element of $E_S(r)$, then $\|g\|_G \leqslant \lambda r + C \leqslant (\lambda + C)r$ (since $r \geqslant 1$).

Let H be the fixed subgroup of A. As discussed in Section 2.1.1, if (g, A) fixes t_g, then its fixed set is $t_g H$, and $g = t_g A(t_g^{-1})$. Thus $\|t_g A(t_g^{-1})\|_G < (\lambda + C)r$. So by Proposition 7.2 there exists $K = K(A, G)$, and an element $h \in H$ such that

$$\|t_g h\|_G < K(\lambda + C)r.$$

In other words, the fixed set of (g, A) intersects $B_G(1, \kappa r)$, where $\kappa = K(\lambda + C)$. Since F is finite, K can be chosen to work simultaneously for all $A \in F$. $\square$

Proof of Proposition 7.2. The proof is by induction on the lower central series of G. Firstly, if G is abelian, we may write the condition on t as $\|t - A(t)\|_G < r$, where $A \in O(n)$. In this case, let $H^\perp$ be the orthogonal complement of H in G. There exists $h \in H$ such that $t + h \in H^\perp$. Since h is a fixed point of A, we have

$$\big\|(t + h) - A(t + h)\big\|_G = \big\|t - A(t)\big\|_G < r. \tag{10}$$

The map A leaves $H^\perp$ invariant and has no fixed points on the compact set

$$\big\{x \in H^\perp \mid \|x\|_G = 1\big\}.$$

Thus the function $\|x - A(x)\|_G$ attains a positive minimum, say m, on this set. Now,

$$\big\|(t + h) - A(t + h)\big\|_G = \|t + h\|_G \left\|\frac{t + h}{\|t + h\|_G} - A\left(\frac{t + h}{\|t + h\|_G}\right)\right\|_G \geqslant m\|t + h\|_G.$$

Inequality (10) now implies that $\|t + h\|_G \leqslant (\frac{1}{m})r$.

Now let G be k-step nilpotent with an orthonormal polynomial coordinate system as in Section 5. Let $\pi : G \to G/G^k$ be the canonical projection, i.e. $\pi(x_1, \dots, x_k) = (x_1, \dots, x_{k-1})$, and let $\bar{A}$ be the automorphism of G/G^k induced by A.

Let $t \in G$ with $\|t A(t^{-1})\|_G < r$. For the rest of the proof, we use $\prec$ to mean *less than, up to a constant factor that depends only on G and A*. We produce $h \in H$ such that $\|t h\|_G \prec r$.

Corollary 5.6 on distances in quotients implies that

$$\big\|\pi(t)\bar{A}\big(\pi(t)^{-1}\big)\big\|_{G/G^k} = \big\|\pi\big(t A(t^{-1})\big)\big\|_{G/G^k} \prec \big\|t A(t^{-1})\big\|_G < r.$$

By Lemma 2.1 the fixed set of $\bar{A}$ is $\pi(H)$. So by the induction hypothesis, there exists $h_1 \in H$ such that $\|\pi(t)\pi(h_1)\|_{G/G^k} \prec r$.

We may write $th_1 = yz$, where $y = (y_1, \ldots, y_{k-1}, 0)$ and $z = (0, \ldots, 0, z) \in G^k$. Note that $\pi(y) = \pi(th_1) = \pi(t)\pi(h_1)$, so that $\|\pi(y)\|_{G/G^k} \prec r$. Further, Corollary 5.5 implies that

$$\|y\|_G \sim \|\pi(y)\|_{G/G^k} \prec r.$$

However, $\|z\|_G$, and hence $\|th_1\|_G$, may be arbitrarily large. This will be fixed by correcting th_1 by an element of $H \cap G^k$.

We first show that $\|zA(z^{-1})\|_G \prec r$. Note that z is in the centre of G, which is preserved by A, so that

$$yzA\big([yz]^{-1}\big) = yzA\big(z^{-1}y^{-1}\big) = yA\big(y^{-1}\big)zA\big(z^{-1}\big). \tag{11}$$

Proposition 6.3 and Observation 6.2 imply that $x \mapsto xA(x^{-1})$ is a weight-preserving polynomial map. Lemma 6.4 now implies that $\|yA(y^{-1})\|_G \prec \|y\|_G \prec r$. Further, $yzA([yz]^{-1}) = th_1A([th_1]^{-1}) = tA(t^{-1})$, so that $\|yzA([yz]^{-1})\|_G \prec r$. Equation (11) now implies that $\|zA(z^{-1})\|_G \prec r$.

Corollary 5.5 implies that $\|x\|_{G^k} \sim \|x\|_G^k$ for all $x \in G^k$. Thus $\|zA(z^{-1})\|_{G^k} \prec r^k$. By the first step of the induction, (since G^k is abelian) there exists $h_2 \in H \cap G^k$, such that $\|zh_2\|_{G^k} \prec r^k$, which means $\|zh_2\|_G \prec r$. Setting $h = h_1h_2$ completes the inductive step:

$$\|th\|_G = \|th_1h_2\|_G = \|yzh_2\|_G \leqslant \|y\|_G + \|zh_2\|_G \prec r. \qquad \square$$

8. Volume estimates in G

In this section we fix an element A in the finite quotient F, and obtain an upper bound for $|\pi^{-1}(A) \cap E_S(r)|$. Let $\mathcal{H}$ be the collection of fixed sets in G of finite-order elements in $\pi^{-1}(A)$. Then Lemma 7.1 says that every element of $\mathcal{H}$ intersects $B_G(1, R)$, where $R = \kappa r$.

If H is the fixed subgroup of A, then $\mathcal{H}$ consists of cosets of H in G. We will use a volume argument to count the number of elements of $\mathcal{H}$ intersecting $B_G(1, R)$. We first obtain disjoint neighbourhoods of the submanifolds in $\mathcal{H}$ and intersect them with $B_G(1, R)$. Then we use the fact that the volume of $B_G(1, R)$ is greater than the sum of the volumes of these disjoint pieces contained in it.

Lemma 8.1. *Let H and $\mathcal{H}$ be as above. Then there exists $\epsilon > 0$ such that the ϵ-neighbourhoods of cosets in $\mathcal{H}$ are pairwise disjoint.*

Proof. We first show that if $t \notin H$, then H and tH have disjoint δ-neighbourhoods for some $\delta > 0$. If there is no such δ, we can find sequences $\{h_i\}$ and $\{tk_i\}$, with $h_i, k_i \in H$, such that $d(h_i, tk_i) \to 0$, which means $\{h_i^{-1}tk_i\}$ is a sequence converging to 1. We now have

$$
\begin{aligned}
A\big(h_i^{-1}tk_i\big) \to A(1) &\Rightarrow& h_i^{-1}A(t)k_i \to 1 \\
&\Rightarrow& k_i^{-1}A\big(t^{-1}\big)h_i \to 1 \\
&\Rightarrow& \big(h_i^{-1}tk_i\big)\big(k_i^{-1}A\big(t^{-1}\big)h_i\big) \to 1 \\
&\Rightarrow& h_i^{-1}tA\big(t^{-1}\big)h_i \to 1.
\end{aligned}
$$

Set $g = tA(t^{-1})$. Suppose $g \in G^j$. Then $\{h_i^{-1}gh_i\} = \{g(g^{-1}h_i^{-1}gh_i)\}$ is a sequence in gG^{j+1} converging to 1. Since gG^{j+1} is closed set, the limit, 1, is in gG^{j+1}. Thus $g \in G^{j+1}$. This inductive argument shows that $g = 1$. This means that t is a fixed point of A, contradicting the fact that $t \notin H$.

Since $\mathcal{H}$ is a discrete collection of cosets of H, the above implies the existence of $\epsilon > 0$ such that for any $tH \in \mathcal{H}$, the ϵ-neighbourhoods of H and tH are disjoint. It follows that the ϵ-neighbourhoods of any two cosets in $\mathcal{H}$ are disjoint. $\quad\square$

We denote the ϵ-neighbourhood of a set Y by $\mathrm{Nbd}_\epsilon(Y)$. We will need to estimate the volumes of intersections of ϵ-neighbourhoods of elements of $\mathcal{H}$ with $B_G(1, R)$. Since $\mathrm{Nbd}_\epsilon(tH) = t\,\mathrm{Nbd}_\epsilon(H)$, we will focus on estimating the volume of $\mathrm{Nbd}_\epsilon(H) \cap B_G(1, R)$.

The following lemma relates the volume of $\mathrm{Nbd}_\epsilon(H) \cap B_G(1, R)$ to the volume of $H \cap B_G(1, R)$ with respect to the left-invariant measure on H.

Lemma 8.2. *Let* vol_G *and* vol_H *denote the left-invariant volumes in G and H, respectively. Let ϵ be the constant obtained in Lemma 8.1. There exists a constant $V_\epsilon > 0$, which is independent of R, such that*

$$\mathrm{vol}_G\big[\mathrm{Nbd}_\epsilon(H) \cap B_G(1, R)\big] > V_\epsilon\, \mathrm{vol}_H\big[H \cap B_G(1, R)\big].$$

Proof. For any R, there exists a finite set $\mathcal{A}_R$ of points in $H \cap B_G(1, R - \epsilon)$ which satisfies the following two conditions:

(1) Balls of radius ϵ in G, centred at points in $\mathcal{A}_R$ are disjoint.
(2) Balls of radius 3ϵ in G, centred at points in $\mathcal{A}_R$ cover $H \cap B_G(1, R)$.

Note that each ϵ-ball as in (1) is contained in $\mathrm{Nbd}_\epsilon(H) \cap B_G(1, R)$ and has volume equal to $V_1^\epsilon = \mathrm{vol}_G[B_G(1, \epsilon)]$. Thus we have

$$\mathrm{vol}_G\big[\mathrm{Nbd}_\epsilon(H) \cap B_G(1, R)\big] > V_1^\epsilon |\mathcal{A}_R|. \tag{12}$$

If $h \in H$, then $B_G(h, 3\epsilon) \cap H = h(B_G(1, 3\epsilon) \cap H)$. Thus, the volume in H of the intersection of H with a 3ϵ-ball as in (2) is a constant, $V_2^\epsilon = \mathrm{vol}_H[B_G(1, 3\epsilon) \cap H]$. Since the collection of balls in (2) covers $H \cap B_G(1, R)$, we have

$$V_2^\epsilon |\mathcal{A}_R| > \mathrm{vol}_H\big[H \cap B_G(1, R)\big]. \tag{13}$$

Set $V_\epsilon = \frac{V_1^\epsilon}{V_2^\epsilon}$. Combining inequalities (12) and (13) yields the result. $\quad\square$

The next step is to estimate $\mathrm{vol}_H[H \cap B_G(1, R)]$. Note that the distance between two points in H measured in the metric on G may be less then their distance in the induced metric on H. If $B_H(1, R)$ denotes the ball of radius R in the induced metric on H, then $B_H(1, R)$ is, in general, a subset of $H \cap B_G(1, R)$.

8.1. Polynomial coordinates compatible with H

We will define a new polynomial coordinate system on G, such that the preimage of H under the polynomial coordinate map is a subspace of $\mathbb{R}^n$ parallel to the coordinate axes. We will then be able to use the ball–box technique from Theorem 5.2 to estimate volumes in H and G simultaneously.

Let $\mathfrak{h}$ be the Lie algebra of H. We choose a triangular basis for $\mathfrak{g}$, such that a subset of the basis is a triangular basis for $\mathfrak{h}$. This can be done as follows.

Let $\mathfrak{g} = \mathfrak{g}^1 \supset \mathfrak{g}^2 \supset \cdots \supset \mathfrak{g}^{k+1} = 0$ be the lower central series of $\mathfrak{g}$. Let

$$\rho_i = \text{rank}\big(\mathfrak{g}^i/\mathfrak{g}^{i+1}\big) \quad \text{and} \quad \eta_i = \text{rank}\big(\mathfrak{h} \cap \mathfrak{g}^i/\mathfrak{h} \cap \mathfrak{g}^{i+1}\big).$$

For each i, pick $X_{i1}, \ldots, X_{i\rho_i}$ to be a pullback of a basis for $\mathfrak{g}^i/\mathfrak{g}^{i+1}$, such that $X_{i1}, \ldots, X_{i\eta_i}$ projects to a basis for $\mathfrak{h} \cap \mathfrak{g}^i/\mathfrak{h} \cap \mathfrak{g}^{i+1}$. Give $\{X_{ij} \mid 1 \leqslant i \leqslant k; \ 1 \leqslant j \leqslant \rho_i\}$ the dictionary-order.

It is easy to see that this gives a triangular basis for $\mathfrak{g}$. Since H is a Lie subgroup, $\mathfrak{h}$ is a subalgebra. In particular, it is closed under the bracket, so that $\{X_{ij} \mid 1 \leqslant i \leqslant k; \ 1 \leqslant j \leqslant \eta_i\}$ is a triangular basis for $\mathfrak{h}$.

Now define a polynomial coordinate map $\phi: \mathbb{R}^n \to G$ as in Definition 4.2. Observe that $\phi^{-1}(H)$ is the set of points $\{(s_{ij}) \in \mathbb{R}^n \mid s_{ij} = 0 \text{ if } \eta_i < j \leqslant \rho_i\}$, which is a plane spanned by a subset of the coordinate axes for G.

We now endow G with a new left-invariant metric that makes the above basis orthonormal. Note that, up to a constant factor, there is only one left-invariant volume form on a Lie group. Since we are only interested in the degree of growth, we may use this new metric to estimate volume.

Recall that the symbol $\sim$ denotes comparable functions. (See Definition 5.3.)

Lemma 8.3. *Retaining the above notation, let $p = \sum_{i=1}^k i\eta_i$. Then*

$$\text{vol}_H\big[H \cap B_G(1, R)\big] \sim R^p.$$

Proof. Theorem 5.2 tells us that in the coordinate system defined above, $B_G(1, R)$ can be bounded by two boxes (one contained in it and one containing it) which have sides parallel to the coordinate axes. In particular, there exists $a > 1$ such that for $R > 1$,

$$\big\{(s_{ij}) \mid |s_{ij}| \leqslant (R/a)^i \text{ for all } i\big\} \subset B_G(1, R) \subset \big\{(s_{ij}) \mid |s_{ij}| \leqslant (aR)^i \text{ for all } i\big\}.$$

For each i, the outer box has ρ_i sides of length $2(aR)^i$. The intersection of this box with H is a box parallel to the coordinate axes in H, with η_i sides of length $2(aR)^i$, for each i. The Lebesgue measure of this intersection is therefore a constant multiple of R^p, where $p = \sum_{i=1}^k i\eta_i$. A similar statement holds for the inner box. Moreover, $H \cap B_G(1, R)$ is contained in the outer box and contains the inner box. This proves the lemma, since the Lebesgue measure on $\phi^{-1}(H)$ is comparable to the left-invariant measure on H. $\quad\square$

We can now prove inequality (1) in Theorem 1.1.

Lemma 8.4. *There exists $c > 0$ such that $|\pi^{-1}(A) \cap E_S(r)| \leqslant cr^{d-p}$, where $d = \sum_{i=1}^k i\rho_i$ and $p = \sum_{i=1}^k i\eta_i$.*

Proof. In Lemma 8.1 we obtained disjoint ϵ-neighbourhoods of the cosets in $\mathcal{H}$. For every $tH \in \mathcal{H}$ which intersects $B_G(1, R)$, choose an element p_t in the intersection. Observe that $\mathrm{Nbd}_\epsilon(tH) = p_t(\mathrm{Nbd}_\epsilon(H))$, so that the sets $p_t(\mathrm{Nbd}_\epsilon(H) \cap B_G(1, R))$ corresponding to distinct cosets in $\mathcal{H}$ are disjoint. Since $\|p_t\|_G \leqslant R$, we have

$$p_t\big(\mathrm{Nbd}_\epsilon(H) \cap B_G(1, R)\big) \subseteq B_G(1, 2R).$$

Let M be the number of elements of $\mathcal{H}$ intersecting $B_G(1, R)$. Then

$$
\begin{aligned}
\mathrm{vol}_G\big[B_G(1, 2R)\big] &> \mathrm{vol}_G\Bigg[\bigcup_{t\mathcal{H}\cap B_G(1,R)\neq\emptyset} p_t\big(\mathrm{Nbd}_\epsilon(H) \cap B_G(1, R)\big)\Bigg] \\
&= M\,\mathrm{vol}_G\big[\mathrm{Nbd}_\epsilon(H) \cap B_G(1, R)\big] \\
&> M V_\epsilon\,\mathrm{vol}_H\big[H \cap B_G(1, R)\big] \quad \text{(Lemma 8.2)},
\end{aligned}
$$

so that

$$M < \left(\frac{1}{V_\epsilon}\right)\frac{\mathrm{vol}_G[B_G(1, 2R)]}{\mathrm{vol}_H[H \cap B_G(1, R)]}.$$

Corollary 5.4 and Lemma 8.3 now imply the existence of a constant $c' > 0$ such that $M < c'R^{d-p}$. Now by Lemma 7.1, M is an upper bound for $|\pi^{-1}(A) \cap E_S(r)|$, so that $|\pi^{-1}(A) \cap E_S(r)| \leqslant cr^{d-p}$, where $c = \kappa c'$. $\square$

9. Finishing the proof

To complete the proof we will need the following theorem of Pansu on the growth of balls in virtually nilpotent groups.

Theorem 9.1. *(See [18].) Let Γ be a finitely generated, virtually nilpotent group with finite generating set S. Let $d = \sum_{i=1}^{\infty} i\,\mathrm{rank}(\Gamma^i/\Gamma^{i+1})$. Then $\lim_{r\to\infty} \frac{|B_S(r)|}{r^d}$ exists.*

In particular, $|B_S(r)| \sim r^d$. Together with Lemma 8.4, this implies that if 1 is an eigenvalue of dA, then

$$\lim_{r\to\infty} \frac{|\pi^{-1}(A) \cap E_S(r)|}{|B_S(r)|} = 0. \tag{14}$$

This is because in this case, the fixed set H of A has dimension at least 1, which implies that $p \geqslant 1$ and hence $d - p < d$.

On the other hand if 1 is not an eigenvalue of dA, we have the following.

Lemma 9.2. *Let A be a finite-order isometry fixing the identity, such that 1 is not an eigenvalue of dA. Then (g, A) has finite-order for every $g \in G$.*

Proof. We prove in Lemma 9.3 below, that for every $g \in G$, there exists $t \in G$ with $g = tA(t^{-1})$. Now $(t, I)(1, A)(t^{-1}, I) = (tA(t^{-1}), A) = (g, A)$. In other words, (g, A) is conjugate in $\mathrm{Isom}\,G$ to $(1, A)$, and hence has finite-order. $\square$

Lemma 9.3. *Let A be a finite-order isometry fixing the identity, such that 1 is not an eigenvalue of dA. The map $\psi : G \to G$ defined by $\psi(t) = t A(t^{-1})$ is surjective.*

Proof. The proof is by induction on the lower central series. If G is abelian, we can write $\psi(t) = t - A(t) = (I - A)t$, where $A = dA$ is linear. Since 1 is not an eigenvalue of A, there is no non-zero v with $(I - A)v = 0$. Thus $I - A$ is invertible and the equation $\psi(t) = b$ has a solution for every b.

Now let $G = G^1 \supset G^2 \supset \cdots \supset G^{k+1} = 1_G$ be the lower central series for G. Then ψ leaves G^i invariant for all i, since A does. Assume $\psi|_{G^i} : G^i \to G^i$ is surjective.

The automorphism A induces an automorphism A_i on G^{i-1}/G^i. It follows from Lemma 2.1 that dA_i does not have 1 as an eigenvalue. The map ψ_i, induced by ψ on G^{i-1}/G^i, is given by $\psi_i(tG^i) = t A(t^{-1})G^i = (tG^i)A_i([tG^i]^{-1})$. Since G^{i-1}/G^i is abelian, ψ_i is surjective.

To prove the surjectivity of $\psi|_{G^{i-1}}$, let $b \in G^{i-1}$. Then there exists $w \in G^{i-1}$ with $bG^i = \psi_i(wG^i) = wA(w^{-1})G^i$. This means $A(w)w^{-1}b$, and hence $w^{-1}bA(w)$ is an element of G^i. Now the surjectivity of $\psi|_{G^i}$ implies that there exists $y \in G^i$ such that $\psi(y) = yA(y^{-1}) = w^{-1}bA(w)$. Then we have

$$\psi(wy) = wy A(y^{-1}) A(w^{-1}) = ww^{-1} bA(w) A(w^{-1}) = b. \qquad \square$$

Thus every element of the coset $\pi^{-1}(A)$ has finite order if 1 is not an eigenvalue of dA. The asymptotic density of a coset is computed in the following corollary.

Corollary 9.4. *If $\pi^{-1}(A)$ is any coset of Λ in Γ, then*

$$\lim_{r \to \infty} \frac{|\pi^{-1}(A) \cap B_S(r)|}{|B_S(r)|} = \frac{1}{|F|}.$$

Proof. Pick a set of coset representatives $\{\gamma_B \mid B \in F;\ \gamma_B \in \pi^{-1}(B)\}$ and let $L = \max\{\ell_S(\gamma_B) \mid B \in F\}$. For any $A \in F$, there is a bijective map $\pi^{-1}(A) \to \Lambda$ given by $x \mapsto x\gamma_A^{-1}$. Then for any $r > 0$, we have

$$\big|\Lambda \cap B_S(r - L)\big| \leqslant \big|\pi^{-1}(A) \cap B_S(r)\big| \leqslant \big|\Lambda \cap B_S(r + L)\big|. \tag{15}$$

Since $|B_S(r)| = \sum_{A \in F} |\pi^{-1}(A) \cap B_S(r)|$, it follows that

$$\big|B_S(r - L)\big| \leqslant |F|\big|\Lambda \cap B_S(r)\big| \leqslant \big|B_S(r + L)\big|. \tag{16}$$

A simple consequence of Theorem 9.1 is that $\lim_{r \to \infty} |B_S(r + N)|/|B_S(r)| = 1$, for any $N \in \mathbb{Z}$. Now Eq. (16) implies that $\lim_{r \to \infty} |\Lambda \cap B_S(r)|/|B_S(r)| = 1/|F|$ and the result follows from Eq. (15). $\quad \square$

Putting together the different pieces yields the formula for $\mathcal{D}(\Gamma, S)$:

End of proof of Theorem 1.1. The inequality (1) was proved in Lemma 8.4. Now let m be the number of elements of $\rho(F)$ which do not have 1 as an eigenvalue. Combining Eq. (14), Lemma 9.2, and Corollary 9.4 we have

$$\mathcal{D}(\Gamma, S) = \lim_{r \to \infty} \sum_{\substack{1 \text{ not an} \\ \text{eigenvalue of } A}} \frac{|\pi^{-1}(A) \cap E_S(r)|}{|B_S(r)|} = m \lim_{r \to \infty} \frac{|\pi^{-1}(A) \cap B_S(r)|}{|B_S(r)|} = \frac{m}{|F|}. \qquad \square$$

10. Arbitrary virtually nilpotent groups

In this section we prove Corollary 1.2. Let Γ be any finitely generated virtually nilpotent group. As discussed in Section 3, Γ has a unique maximal finite normal subgroup, say Q, and Γ/Q is almost crystallographic. We wish to show that $\mathcal{D}(\Gamma, S) = \mathcal{D}(\Gamma/Q)$ for any generating set S of Γ.

Proof of Corollary 1.2. Let $S = \{\gamma_1, \dots, \gamma_l\}$ be a generating set for Γ. Then $\bar{S} = \{\gamma_1 Q, \dots, \gamma_l Q\}$ generates Γ/Q. Let ℓ_S and $\ell_{\bar{S}}$ denote the corresponding length functions on Γ and Γ/Q, respectively.

Let $g \in \Gamma$. Clearly, $\ell_{\bar{S}}(gQ) \leqslant \ell_S(g)$. Moreover, if $\gamma_{i_1} Q \cdots \gamma_{i_n} Q$ is a geodesic word representing gQ, then $g = \gamma_{i_1} \cdots \gamma_{i_n} q'$ for some $q' \in Q$. If $M = \max\{\ell_S(q) \mid q \in Q\}$, then $\ell_S(g) \leqslant \ell_{\bar{S}}(gQ) + M$.

Let $B_\Gamma(r)$ and $B_{\Gamma/Q}(r)$ denote the balls of radius r in Γ and Γ/Q, respectively. Let $E_\Gamma(r)$ and $E_{\Gamma/Q}(r)$ represent the corresponding sets of finite-order elements. The above inequalities yield:

$$\frac{|B_\Gamma(r)|}{|Q|} \leqslant |B_{\Gamma/Q}(r)| \quad \text{and} \quad |B_{\Gamma/Q}(r)| \leqslant \frac{|B_\Gamma(r + M)|}{|Q|}.$$

Since Q is finite, an element of Γ has finite order if and only if its projection in Γ/Q has finite order. Thus we have

$$\frac{|E_\Gamma(r)|}{|Q|} \leqslant |E_{\Gamma/Q}(r)| \quad \text{and} \quad |E_{\Gamma/Q}(r)| \leqslant \frac{|E_\Gamma(r + M)|}{|Q|}.$$

Putting together the above information, we have

$$\frac{|E_{\Gamma/Q}(r - M)|}{|B_{\Gamma/Q}(r)|} \leqslant \frac{|E_\Gamma(r)|}{|B_\Gamma(r)|} \leqslant \frac{|E_{\Gamma/Q}(r)|}{|B_{\Gamma/Q}(r - M)|}.$$

Theorems 1.1 and 9.1 can now be used to conclude that

$$\mathcal{D}(\Gamma, S) = \lim_{r \to \infty} \frac{|E_\Gamma(r)|}{|B_\Gamma(r)|} = \mathcal{D}(\Gamma/Q). \qquad \square$$

11. Examples

Crystallographic groups are groups which act properly discontinuously and cocompactly on Euclidean space. They are virtually abelian, and hence virtually nilpotent. The results of the

computation of $\mathcal{D}(\Gamma)$ for the crystallographic groups in dimensions 2, 3, and 4, computed using GAP, are summarised in Appendix A.

The study of $\mathcal{D}(\Gamma)$ for crystallographic groups leads to a number of questions:

- Given a rational number $r \in [0, 1)$, is there a crystallographic group Γ with $\mathcal{D}(\Gamma) = r$?
- More generally, for every k, is there a k-step nilpotent group Γ with $\mathcal{D}(\Gamma) = r$?
- What is the highest density that can occur in crystallographic groups of a given dimension?
- What is the smallest dimension that a given density occurs in?
- Is there an interesting explanation for the spectrum of densities in a given dimension?

In this section we answer the first of these by constructing examples to show that in fact, every rational number in $[0, 1)$ occurs as $\mathcal{D}(\Gamma)$ for some crystallographic group Γ. We also give a partial answer to the third question, and finally we investigate $\mathcal{D}(\Gamma)$ for some virtually nilpotent groups which are not virtually abelian.

11.1. Constructing examples

The finite quotient $F = \Gamma/\Lambda$ is called the *holonomy group* of Γ. The holonomy group of a crystallographic group can be realised as a finite subgroup of $Gl(n, \mathbb{Z})$. On the other hand, if F is a finite subgroup of $Gl(n, \mathbb{Z})$, an averaging argument can be used to show that F preserves an inner product on $\mathbb{R}^n$. Equivalently, there exists $M \in Gl(n, \mathbb{R})$ such that $F' = MFM^{-1} \subset O(n)$. Then the lattice $\Lambda = M\mathbb{Z}^n$ is preserved by F' and $\Gamma = \Lambda \rtimes F'$ defines a crystallographic group.

The following Lemma is useful for constructing many examples.

Lemma 11.1. *Let Γ_1 and Γ_2 be virtually nilpotent groups. Then*

$$\mathcal{D}(\Gamma_1 \times \Gamma_2) = \mathcal{D}(\Gamma_1)\mathcal{D}(\Gamma_2).$$

Proof. In light of Corollary 1.2, we may assume Γ_i acts geometrically on a nilpotent Lie group G_i, for $i = 1, 2$. In this case $\Gamma_1 \times \Gamma_2$ acts geometrically on $G_1 \times G_2$. If Γ_i fits into

$$0 \to \Lambda_i \to \Gamma_i \to F_i \to 1,$$

where Λ_i is maximal normal nilpotent, then we have

$$0 \to \Lambda_1 \times \Lambda_2 \to \Gamma_1 \times \Gamma_2 \to F_1 \times F_2 \to 1,$$

and $\Lambda_1 \times \Lambda_2$ is the maximal normal nilpotent subgroup of $\Gamma_1 \times \Gamma_2$.

If $A = (A_1, A_2) \in F_1 \times F_2$, then the set of eigenvalues of dA is the union of the eigenvalues of dA_1 and dA_2. In particular, 1 is not an eigenvalue for dA if and only if neither dA_1 nor dA_2 has 1 as an eigenvalue. Thus $\mathcal{D}(\Gamma_1 \times \Gamma_2) = \mathcal{D}(\Gamma_1)\mathcal{D}(\Gamma_2)$. $\quad\square$

Proof of Theorem 1.3. We start by constructing, for any $m \in \mathbb{Z}$, a crystallographic group Γ_m such that $\mathcal{D}(\Gamma_m) = \frac{m-1}{m}$. Let ζ be a primitive mth root of unity. If Φ denotes the Euler function,

then $\{1, \zeta, \zeta^2, \ldots, \zeta^{\Phi(m)-1}\}$ is a basis for $\mathbb{Z}[\zeta]$, and we have $\zeta^{\Phi(m)} = \sum_{i=0}^{\Phi(m)-1} a_i \zeta^i$, where $a_i \in \mathbb{Z}$. The matrix T, representing multiplication by ζ on $\mathbb{Z}[\zeta]$ is given as follows:

$$T = \begin{pmatrix} 0 & & & & & & a_0 \\ 1 & 0 & & & & & a_2 \\ & 1 & 0 & & & & a_3 \\ & & 1 & \ddots & & & \vdots \\ & & & \ddots & 0 & & a_{\Phi(m)-2} \\ & & & & 1 & & a_{\Phi(m)-1} \end{pmatrix}.$$

The characteristic polynomial of T is $x^{\Phi(m)} - \sum_{i=0}^{\Phi(m)-1} a_i x^i$, which is also the minimal polynomial of ζ. The eigenvalues of T are exactly the $\Phi(m)$ primitive mth roots of unity. Thus T has order m and the matrices T^i do not have 1 as an eigenvalue, for $i < m$.

Let $F \in O(n)$ be conjugate to $\langle T \rangle$ and let $\Lambda \simeq \mathbb{Z}^n$ be the lattice preserved by F. Then $\Gamma_m = \Lambda \rtimes F$ is the desired group.

Now let $\frac{p}{q} \in [0, 1)$ and note that $\frac{p}{q} = \frac{p}{p+1}\frac{p+1}{p+2}\cdots\frac{q-1}{q}$. Appealing to Lemma 11.1 we can construct an example of a group Γ with $\mathcal{D}(\Gamma) = \frac{p}{q}$. $\quad\square$

The above construction gives a very high-dimensional crystallographic group if p is much smaller than q. It would be nice to obtain a more efficient example.

11.2. Highest densities

As seen in the tables in Appendix A, the highest values of $\mathcal{D}(\Gamma)$ in two-, three- and four-dimensional crystallographic groups are 5/6, 1/2, and 23/24, respectively. The fact that the highest density in three dimensions is 1/2 is part of a more general phenomenon:

Proposition 11.2. *If Γ is an odd-dimensional crystallographic group, $\mathcal{D}(\Gamma) \leqslant 1/2$.*

Proof. The holonomy group F of Γ can be realised as a finite subgroup of $Gl(n, \mathbb{Z})$. Since elements of F have finite order, all their eigenvalues are roots of unity. Thus if n is odd and $A \in F \subset Gl(n, \mathbb{Z})$ is orientation preserving (i.e. A has determinant 1), then 1 is necessarily an eigenvalue of A. Thus at least half the elements of F have 1 as an eigenvalue, proving the result. $\quad\square$

The upper bound is attained, for example by $\mathbb{Z}^n \rtimes \mathbb{Z}_2$, where the non-trivial element of $\mathbb{Z}_2$ is the automorphism T of $\mathbb{Z}^n$ defined by $T(v) = -v$ for all v.

11.3. Almost crystallographic groups

We now investigate $\mathcal{D}(\Gamma)$ for some virtually nilpotent (but not virtually abelian) groups which act geometrically on nilpotent Lie groups. For these, the holonomy group can be realised as a finite group of automorphisms of the associated Lie algebra. We first show that in three and four dimensions, this turns out to be too restrictive, and $\mathcal{D}(\Gamma)$ is always 0.

Recall that the *complexification* of $\mathfrak{g}$, denoted by $\mathfrak{g}_{\mathbb{C}}$, is $\mathfrak{g} \otimes_{\mathbb{R}} \mathbb{C}$. Any inner product on $\mathfrak{g}$ extends to a positive definite hermitian form on $\mathfrak{g}_{\mathbb{C}}$. Any inner-product-preserving automorphism T of $\mathfrak{g}$ extends to a unitary operator with the same eigenvalues. Further, if λ is an eigenvalue of T, then there is an eigenvector (or a generalised eigenspace of the appropriate dimension) corresponding to λ in $\mathfrak{g}_{\mathbb{C}}$.

Lemma 11.3. *Let T be an automorphism of a 3- or 4-dimensional nilpotent, non-abelian Lie algebra $\mathfrak{g}$. If all eigenvalues of T have absolute value 1, then 1 is an eigenvalue of T.*

Proof. In each case below, we pass to the complexification of $\mathfrak{g}$ to ensure the existence of eigenvectors or generalised eigenspaces.

If $\mathfrak{g}$ is 3-dimensional, it is isomorphic to the Heisenberg Lie algebra. If Z generates $\mathfrak{g}^2$, then $T(Z) = \pm Z$. Thus we may assume that $T(Z) = -Z$ and that the eigenvalues of T are -1, λ_1 and λ_2, where λ_1 and λ_2 are either both real, or complex conjugates of each other.

If λ_1 and λ_2 are distinct, there exist linearly independent eigenvectors v_1 and v_2 in $\mathfrak{g}_{\mathbb{C}}$. Then $[v_1, v_2] = cZ$ for some $c \in \mathbb{C}$. Since T preserves the bracket, $[Tv_1, Tv_2] = T(cZ) = -cZ$. On the other hand, $[Tv_1, Tv_2] = [\lambda_1 v_1, \lambda_2 v_2] = \lambda_1 \lambda_2 [v_1, v_2] = \lambda_1 \lambda_2 cZ$. We conclude that $\lambda_1 \lambda_2 = -1$. This cannot happen if λ_1 and λ_2 are complex conjugates, so the two eigenvalues have to be 1 and -1.

If $\lambda_1 = \lambda_2 = \lambda$, then there exist vectors v_1 and v_2 in $\mathfrak{g}_{\mathbb{C}}$ such that $T(v_1) = \lambda v_1$ and $Tv_2 = \lambda v_2 + \alpha v_1$, for some α. Let $[v_1, v_2] = cZ$ for some $c \in \mathbb{C}$. Then $-cZ = [Tv_1, Tv_2] = [\lambda v_1, \lambda v_2 + \alpha v_1] = \lambda^2 [v_1, v_2] = \lambda^2 cZ$, which is impossible, since λ is real.

If $\mathfrak{g}$ is 4-dimensional and 2-step nilpotent, then $\mathfrak{g}^2$ is necessarily 1-dimensional. If Z generates $\mathfrak{g}^1$, then $T(Z) = \pm Z$. Assuming $T(Z) = -Z$, so that -1 is an eigenvalue, at least one of the other eigenvalues must be real. Thus we may assume the set of eigenvalues is $\{-1, -1, \lambda_1, \lambda_2\}$ and proceed as above.

If $\mathfrak{g}$ is 3-step nilpotent, $\mathfrak{g}^3$ and $\mathfrak{g}^2/\mathfrak{g}^3$ are 1-dimensional. Let $\mathfrak{g}^3 = \langle Z \rangle$ and $\mathfrak{g}^2 = \langle W, Z \rangle$. Then $T(Z) = \pm Z$ and $T(W) = \pm W + aZ$, for some $a \in \mathbb{R}$. We may assume the set of eigenvalues is $\{-1, -1, \lambda_1, \lambda_2\}$. An argument similar to the above completes the proof. $\square$

Remark 11.4. If T is an inner-product-preserving automorphism of $\mathfrak{g}$, then every eigenvalue of T has absolute value 1.

Theorem 1.1, Lemma 11.3, and Remark 11.4 imply the following corollary.

Corollary 11.5. *If Γ is a group acting geometrically on a 3- or 4-dimensional nilpotent (non-abelian) Lie group, then $\mathcal{D}(\Gamma) = 0$.*

We now construct a class of almost crystallographic groups Γ, such that $\mathcal{D}(\Gamma)$ is non-zero.

Definition 11.6 (*Generalised Heisenberg Lie algebras*)*.* Define $\mathfrak{h}_n$ to be the Lie algebra generated by $\{X_1, \ldots, X_n, Y_1, \ldots, Y_n, Z\}$ such that $[X_i, Y_i] = Z$ for $1 \leqslant i \leqslant n$, and all other brackets are 0.

The following construction can be done for any generalised Heisenberg Lie algebra $\mathfrak{h}_{2n}$, of dimension $4n + 1$. We give the construction for $\mathfrak{h}_2$:

Define an automorphism T on $\mathfrak{h}_2$ by

$$X_1 \mapsto X_2, \qquad Y_1 \mapsto -Y_2, \qquad Z \mapsto -Z, \qquad X_2 \mapsto -X_1, \qquad Y_2 \mapsto Y_1.$$

It is easy to check that T preserves the bracket. The matrix of T with respect to the basis $\{X_1, X_2, Y_1, Y_2, Z\}$ is given by

$$\begin{pmatrix} 0 & -1 & & & \\ 1 & 0 & & & \\ & & 0 & 1 & \\ & & -1 & 0 & \\ & & & & -1 \end{pmatrix}.$$

Thus T is an automorphism of order 4 whose eigenvalues are $\pm i$ and -1.

Let H_2 be the connected, simply connected nilpotent Lie group corresponding to $\mathfrak{h}_2$. Let $\widetilde{T}$ be the automorphism of H_2 with $d\widetilde{T} = T$ and $N = \exp(\mathbb{Z}^5)$ be the lattice in H_2 preserved by $\widetilde{T}$. Then $\Gamma = N \rtimes \langle \widetilde{T} \rangle$ is an almost-crystallographic group with $\mathcal{D}(\Gamma) = \frac{1}{2}$.

Note that for any group acting on a Lie group with 1-dimensional centre, the maximum value of $\mathcal{D}$ is $\frac{1}{2}$, since the square of any automorphism fixes the central direction. The groups Γ defined above attain this maximum value.

Acknowledgments

Appendix A

The computations summarised below were done using the computer algebra software GAP [19] and the software package "Cryst," which contains libraries of 2-, 3- and 4-dimensional crystallographic groups.

Table 1 gives the values of $\mathcal{D}(\Gamma)$ for all the 2-dimensional crystallographic groups. See [17] for a description of the notation. In Tables 1 and 2, $\mathcal{N}(q)$ denotes the number of groups Γ for which $\mathcal{D}(\Gamma) = q$.

Table 1

Dimension two					Dimension three	

Γ	$\mathcal{D}(\Gamma)$		Γ	$\mathcal{D}(\Gamma)$	q	$\mathcal{N}(q)$
$\mathcal{W}_1$	0		$\mathcal{W}_4$	3/4	0	113
$\mathcal{W}_1^1$	0		$\mathcal{W}_4^1$	3/8	1/8	28
$\mathcal{W}_1^2$	0		$\mathcal{W}_4^2$	3/8	1/6	4
$\mathcal{W}_1^3$	0				3/16	20
			$\mathcal{W}_2$	1/2	5/24	4
$\mathcal{W}_6$	5/6		$\mathcal{W}_2^1$	1/4	1/4	30
$\mathcal{W}_6^1$	5/12		$\mathcal{W}_2^2$	1/4	5/16	10
			$\mathcal{W}_2^3$	1/4	1/3	1
$\mathcal{W}_3$	2/3		$\mathcal{W}_2^4$	1/4	3/8	13
$\mathcal{W}_3^1$	2/3				5/12	2
$\mathcal{W}_3^2$	1/3				1/2	5

Table 2

Dimension four							
q	$\mathcal{N}(q)$	q	$\mathcal{N}(q)$	q	$\mathcal{N}(q)$	q	$\mathcal{N}(q)$
0	1875	37/144	3	9/20	1	11/16	4
1/16	605	33/128	16	11/24	1	25/36	3
1/12	64	35/128	8	15/32	4	17/24	1
3/32	426	5/18	9	1/2	6	205/288	1
5/48	48	9/32	34	37/72	2	137/192	2
1/9	25	23/80	2	33/64	10	13/18	2
1/8	558	85/288	1	25/48	1	35/48	3
5/36	5	5/16	91	17/32	7	3/4	3
9/64	50	21/64	20	35/64	5	55/72	2
5/32	193	1/3	20	5/9	7	19/24	1
1/6	38	385/1152	1	9/16	17	4/5	1
25/144	2	49/144	2	23/40	2	77/96	4
13/72	7	25/72	2	85/144	1	13/16	4
3/16	229	205/576	2	43/72	1	5/6	1
1/5	10	13/36	6	5/8	14	41/48	3
13/64	23	3/8	27	91/144	1	31/36	1
5/24	31	2/5	6	21/32	13	7/8	4
2/9	31	13/32	9	2/3	4	9/10	1
9/40	2	5/12	4	385/576	1	11/12	2
15/64	11	7/16	2	65/96	2	23/24	4
1/4	125	4/9	12	49/72	1		

References

[1] S.I. Adyan, The Burnside Problem and Identities in Groups, Nauka, Moscow, 1975.
[2] A. Ahlin, The large scale geometry of nilpotent-by-cyclic groups, PhD thesis, University of Chicago, 2001.
[3] L. Auslander, A fixed point theorem for nilpotent Lie groups, Proc. Amer. Math. Soc. 9 (1958) 822–823.
[4] L. Auslander, Bieberbach's theorems on space groups and discrete uniform subgroups of Lie groups, Ann. of Math. (2) 71 (1960) 579–590.
[5] L. Corwin, F. Greenleaf, Representations of Nilpotent Lie Groups and Their Applications. Part I, Cambridge Stud. Adv. Math., vol. 18, Cambridge Univ. Press, Cambridge, 1990.
[6] K. Dekimpe, P. Igodt, The structure and topological meaning of almost-torsion free groups, Comm. Algebra 22 (1994) 2547–2558.
[7] V.V. Gorbatsevich, A.L. Onishchik, E.B. Vinberg, Foundations of Lie Theory and Lie Transformation Groups, Springer-Verlag, Berlin, 1997 (translated from the Russian by A. Kozlowski). Reprint of Lie Groups and Lie Algebras. I, Encyclopaedia Math. Sci., vol. 20, Springer-Verlag, Berlin, 1993.
[8] R.I. Grigorchuk, On Burnside's problem on periodic groups, Funktsional. Anal. i Prilozhen. 14 (1980) 53–54 (in Russian).
[9] M. Gromov, Carnot–Carathéodory spaces seen from within, in: Sub-Riemannian Geometry, in: Progr. Math., vol. 144, Birkhäuser, Basel, 1996, pp. 79–323.
[10] P. Hall, The Edmonton Notes on Nilpotent Groups, Queen Mary College Mathematics Notes, Mathematics Department, Queen Mary College, London, 1969.
[11] P. de la Harpe, Topics in Geometric Group Theory, Chicago Lectures in Math., Univ. of Chicago Press, Chicago, IL, 2000.
[12] I. Kapovich, I. Rivin, P. Schupp, V. Schpilrain, Asymptotic density in free groups and $\mathbb{Z}^k$, visible points and test elements, math.GR/0507573.
[13] R. Karidi, Geometry of balls in nilpotent Lie groups, Duke Math. J. 74 (1994) 301–317.
[14] K.B. Lee, F. Raymond, Rigidity of almost-crystallographic groups, in: Combinatorial Methods in Topology and Algebraic Geometry, Rochester, NY, 1982, in: Contemp. Math., vol. 44, Amer. Math. Soc., Providence, RI, 1985, pp. 73–78.
[15] I.G. Lysionok, Infinite Burnside groups of even period, Izv. Ross. Akad. Nauk Ser. Mat. 60 (1996) 3–224 (in Russian); translation in Izv. Math. 60 (1996) 453–654.
[16] A.I. Malcev, On a Class of Homogeneous Spaces, Amer. Math. Soc. Transl., vol. 39, Amer. Math. Soc., Providence, RI, 1951.
[17] G. Martin, Transformation Geometry, Undergrad. Texts in Math., Springer-Verlag, New York, 1982.
[18] P. Pansu, Croissance des boules et des géodésiques fermées dans les nilvariétés, Ergodic Theory Dynam. Systems 3 (3) (1983) 415–445.
[19] M. Schönert, et al., GAP—Groups, Algorithms, and Programming, version 3, release 4, patchlevel 4, share package CRYST, Lehrstuhl D für Mathematik, Rheinisch Westfälische Techn. Hochsch., Aachen, Germany, 1997.
[20] D. Segal, Polycyclic Groups, Cambridge Tracts in Math., vol. 82, Cambridge Univ. Press, Cambridge, 1983.