

THE STRUCTURE OF APPROXIMATE GROUPS

by EMMANUEL BREUILLARD, BEN GREEN, and TERENCE TAO

ABSTRACT

Let $K \geq 1$ be a parameter. A K -approximate group is a finite set A in a (local) group which contains the identity, is symmetric, and such that $A \cdot A$ is covered by K left translates of A .

The main result of this paper is a qualitative description of approximate groups as being essentially finite-by-nilpotent, answering a conjecture of H. Helfgott and E. Lindenstrauss. This may be viewed as a generalisation of the Freiman-Ruzsa theorem on sets of small doubling in the integers to arbitrary groups.

We begin by establishing a correspondence principle between approximate groups and locally compact (local) groups that allows us to recover many results recently established in a fundamental paper of Hrushovski. In particular we establish that approximate groups can be approximately modeled by Lie groups.

To prove our main theorem we apply some additional arguments essentially due to Gleason. These arose in the solution of Hilbert's fifth problem in the 1950s.

Applications of our main theorem include a finitary refinement of Gromov's theorem, as well as a generalized Margulis lemma conjectured by Gromov and a result on the virtual nilpotence of the fundamental group of Ricci almost nonnegatively curved manifolds.

CONTENTS

1. Introduction	115
2. Coset nilprogressions and a more detailed version of the Main Theorem	124
3. Ultra approximate groups and Hrushovski's Lie Model Theorem	130
4. An outline of the argument	140
5. Sanders-Croot-Sisask theory	142
6. Proof of the Hrushovski Lie model theorem	148
7. Strong approximate groups	160
8. The escape norm and a Gleason type theorem	163
9. Proof of the main theorem	172
10. A dimension bound	183
11. Applications to growth in groups and geometry	190
Acknowledgements	200
Appendix A: Basic theory of ultralimits and ultraproducts	200
Appendix B: Local groups	205
Appendix C: Nilprogressions and related objects	213
References	219

1. Introduction

Approximate groups. — A fair proportion of the subject of additive combinatorics is concerned with approximate analogues of exact algebraic properties, and the extent to which they resemble those algebraic properties. In this paper we are concerned with sets that are approximately closed under multiplication, which we do not necessarily assume to be commutative, and more specifically with *approximate groups*. These are finite non-empty sets A with group-like properties which we shall state precisely later. First we will motivate the definition of an approximate group with some discussion and examples.

Suppose first of all that A is a finite subset of some ambient group $G = (G, \cdot)$. This is the setting considered in essentially all of the existing literature, and the one of

importance in applications. However, as we shall see later, our method of proof is in fact more naturally adapted to a more general setting, in which A lies in a *local* group rather than a global one.

It is easy to see that a finite non-empty subset A of G is a genuine subgroup if, and only if, we have $xy^{-1} \in A$ whenever $x, y \in A$. Perhaps the most natural way in which a set A may be *approximately* a subgroup, then, is if the set $A \cdot A^{-1} := \{xy^{-1} : x, y \in A\}$ has cardinality not much bigger than the cardinality of A : for example, we might ask that $|A \cdot A^{-1}| \leq K|A|$ for some constant K .

Sets with this property or with the closely related property $|A^2| \leq K|A|$, where $A^2 := A \cdot A = \{xy : x, y \in A\}$, are said to have *small doubling*, and this is indeed a commonly encountered condition in various fields of mathematics, in particular in additive combinatorics. It is a perfectly workable notion of approximate group in the abelian setting and the celebrated Freiman-Ruzsa theorem, Theorem 2.1 below, describes subsets of \mathbf{Z} with this property. However in [52] it was noted that in non-commutative settings a somewhat different, though closely related, notion of approximate group is more natural: A is an approximate group if it is symmetric in the sense that the identity id lies in A , if $a^{-1} \in A$ whenever $a \in A$, and if $A \cdot A$ is covered by K left-translates of A .

As suggested above we consider in this paper a slightly more general (and perhaps more natural, in retrospect) “local” definition of approximate group in which there is no ambient global group G . It will be convenient to introduce the following definition. This requires the concept of a *local group*, which is discussed at some length in Appendix B.

Definition 1.1 (Multiplicative set). — *A multiplicative set is a finite non-empty set A contained in a (symmetric) local group $G = (G, \cdot)$, such that the product set $(A \cup A^{-1})^{200}$ is well-defined, where $A^{-1} := \{a^{-1} : a \in A\}$ is the inverse of A . Strictly speaking, one should refer to the pair (A, G) as the multiplicative set rather than just A , but we will usually abuse notation and omit the ambient local group G .*

In some (abelian) examples, we will use additive group notation $G = (G, +)$ rather than multiplicative notation $G = (G, \cdot)$. In such cases, we will refer to multiplicative sets as additive sets instead.

Clearly, any finite non-empty subset of a (global) group G is a multiplicative set. The reader should probably keep this model case in mind throughout a first reading of this paper. Indeed the additional generality afforded by the local setting is only needed at a single, albeit critical, place in the argument in Section 9. One should informally think of a multiplicative set A as a set that behaves “as if” it were in a global group, so long as one only works “locally” in the sense that one only considers products of up to 200 elements of A and their inverses. The exponent 200 in Definition 1.1 is somewhat arbitrary, but for the purposes of studying approximate groups, the exact choice of this exponent is not important in practice, so long as it is at least 8 (see Theorem 5.3 for a precise formalisation of this assertion). For the reader familiar with Freiman homomorphisms (cf. [54, §5.3]), we remark that these are essentially the morphisms in the category of multiplicative sets.

Definition 1.2 (*Approximate groups*). — Let $K \geq 1$. A K -approximate group is a multiplicative set A with the following properties:

- (i) the set A is symmetric in the sense that $\text{id} \in A$ and $a^{-1} \in A$ if $a \in A$;
- (ii) there is a symmetric subset $X \subset A^3$ with $|X| \leq K$ such that $A \cdot A \subseteq X \cdot A$.

We will sometimes refer to actual (global) groups as genuine groups, in order to distinguish them from approximate groups. We define a global K -approximate group to be a K -approximate group A that lies inside a global group G . We refer to K as the covering parameter of the approximate group A .

Remark 1.3. — We will also have occasion to deal with infinite K -approximate groups, which are defined exactly as ordinary K -approximate groups, except that they are no longer required to be finite sets. A convex body in a Euclidean space, or a small ball in a Lie group, are examples of infinite approximate groups. Later we will introduce the important notion of an *ultra approximate group*, which is another example. However, by default, approximate groups in this paper will be understood to be finite unless otherwise stated.

The connection between sets with small doubling and the apparently stronger property of being an approximate group was worked out in [52], building on work of Ruzsa [45]; see Remark 1.5 below.

When we speak of an “approximate group” we shall generally imagine that K is fixed (e.g. $K = 10$) and that $|A|$ is large. Let us give some examples.

Example 1 (*Finite group*). — A 1-approximate group is the same thing as a finite group.

Example 2 (*Arithmetic/geometric progression*). — If $N \in \mathbf{N}$ is a natural number, then the arithmetic progression $P(1; N) := \{-N, \dots, N\}$ (which one can view inside the (additive) global group \mathbf{Z} , or the local group $\{-200N, \dots, 200N\}$) is a 2-approximate group. More generally, if $G = (G, \cdot)$ is any (global) group and $g \in G$ then the geometric progression $P(g, N) := \{g^{-N}, \dots, g^N\}$ is a 2-approximate group.

Example 3 (*Generalised arithmetic progression*). — Let $G = (G, +)$ be an abelian group, let $u_1, \dots, u_r \in G$ for some $r \geq 0$, and let $N_1, \dots, N_r > 0$ be real numbers. We refer to the set

$$\begin{aligned} &P(u_1, \dots, u_r; N_1, \dots, N_r) \\ &:= \{n_1 u_1 + \dots + n_r u_r : n_1, \dots, n_r \in \mathbf{Z}; |n_1| \leq N_1, \dots, |n_r| \leq N_r\} \end{aligned}$$

as a *generalised arithmetic progression* of rank r . One easily verifies that this is a 2^r -approximate group.

Example 4 (Homomorphic images). — Let $\phi : G \rightarrow H$ be a homomorphism between local or global groups. If A is a K -approximate subgroup of G , then $\phi(A)$ is a K -approximate subgroup of H . This observation can be generalised to the case when ϕ is a *Freiman homomorphism* (of order 3) rather than a group homomorphism; see [54, §5.3] for more discussion. Indeed, Freiman homomorphisms are very similar to homomorphisms of local groups, although for technical reasons we will rely on the latter concept rather than the former.

Conversely, if B is a K -approximate subgroup of H , ϕ is surjective, and $\ker(\phi)$ is finite, then $\phi^{-1}(B)$ is a K -approximate subgroup of G . In the latter case one can view the K -approximate group $\phi^{-1}(B)$ as a “finite extension” of the K -approximate group B by the genuine group $\ker(\phi)$.

Example 5 (Large subsets). — Let A be a K -approximate group, and let A' be a symmetric neighbourhood of the identity in A such that A is covered by K' left-translates of A' . Then A' is a KK' -approximate group. This hints that approximate groups are considerably more numerous than genuine groups, because the latter property is preserved under passage to “large” subsets, whereas the former is not.

Example 6 (Heisenberg example). — Let G be the free nilpotent group of step 2 generated by two generators u_1, u_2 . More concretely, one can take G to be the *Heisenberg group*

$$(1.1) \quad G := \begin{pmatrix} 1 & \mathbf{Z} & \mathbf{Z} \\ 0 & 1 & \mathbf{Z} \\ 0 & 0 & 1 \end{pmatrix}$$

with generators

$$u_1 := \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix} \quad \text{and} \quad u_2 := \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Consider also the commutator

$$[u_2, u_1] := u_2^{-1} u_1^{-1} u_2 u_1 = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix};$$

one has

$$\begin{pmatrix} 1 & n_1 & n_{12} \\ 0 & 1 & n_2 \\ 0 & 0 & 1 \end{pmatrix} = u_1^{n_1} u_2^{n_2} [u_2, u_1]^{n_{12}}$$

for all integers n_1, n_2, n_{12} .

Let $N_1, N_2 \geq 10$ be real numbers. Define the *nilprogression* $P(u_1, u_2; N_1, N_2)$ to be the set of all words in $u_1, u_1^{-1}, u_2, u_2^{-1}$ that involve at most N_1 occurrences of u_1, u_1^{-1} and at most N_2 occurrences of u_2, u_2^{-1} . It is not difficult to verify that $P(u_1, u_2; N_1, N_2)$ is a symmetric neighbourhood of the identity which contains the set

$$\{u_1^{n_1} u_2^{n_2} [u_2, u_1]^{n_{12}} : |n_1| \leq N_1/10, |n_2| \leq N_2/10, |n_{12}| \leq N_1 N_2/10\}$$

and is contained in the set

$$\{u_1^{n_1} u_2^{n_2} [u_2, u_1]^{n_{12}} : |n_1| \leq 10N_1, |n_2| \leq 10N_2, |n_{12}| \leq 10N_1 N_2\}.$$

One can easily verify that $P(u_1, u_2; N_1, N_2)$ is a K -approximate group for some absolute constant K (for instance, one could take $K = 100$).

Remark 1.4. — The above example was constructed inside the Heisenberg group. Later on we will discuss a generalisation of this example to arbitrary nilpotent groups. These examples, which we will call *nilprogressions*, will be needed to state the precise version of our main theorem (Theorem 2.10) below. We will define them later in this introduction.

Example 7 (Direct products). — The direct product of a K_1 -approximate group and a K_2 -approximate group is a $K_1 K_2$ -approximate group, and so one may build up examples of approximate groups using both subgroups and nilprogressions.

Example 8 (Helfgott's example). — The following example of Helfgott¹ is a less obvious way of combining a subgroup and a nilprogression.

Let $A \subseteq \mathrm{GL}_3(\mathbf{F}_p)$ be the following set of 3×3 matrices:

$$A := \left\{ \begin{pmatrix} r^n & x & z \\ 0 & s^n & y \\ 0 & 0 & (rs)^{-n} \end{pmatrix} : x, y, z \in \mathbf{F}_p, -N \leq n \leq N \right\}.$$

Here, $r, s \in \mathbf{F}_p^\times$ are fixed and N is large yet much smaller than p . Then A is a $O(1)$ -approximate group.

Note that A has the following form: it admits a subgroup H , normalised by A , such that A/H is a geometric progression. Indeed

$$H = \left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} : x, y, z \in \mathbf{F}_p \right\}.$$

In the language of Example 4, A is a finite extension of a geometric progression by the finite group H .

¹ See terrytao.wordpress.com/2009/06/21/freimans-theorem-for-solvable-groups/#comment-39705.

Each of the above examples was rather “algebraic” in nature, whereas the definition of approximate group is somewhat combinatorial. We also have some more combinatorial criteria for generating approximate groups using sets of small doubling or tripling.

Remark 1.5 (Relationship between small doubling and approximate groups). — Let A be a non-empty finite subset of a global group G . If $|A^3| \leq K|A|$, then the set $H := (A \cup \{\text{id}\} \cup A^{-1})^2$ is² a $O(K^{O(1)})$ -approximate group that contains A ; see [52, Theorem 3.9]. In a similar vein if $|A^2| \leq K|A|$ or $|A \cdot A^{-1}| \leq K|A|$, then there exists a $O(K^{O(1)})$ -approximate group H of size $|H| = O(K^{O(1)}|A|)$ such that A can be covered by $O(K^{O(1)})$ left-translates gH of H ; see [52, Theorem 4.6].

Our aim in this paper is to “describe” the structure of approximate subgroups in an arbitrary ambient group in terms of more explicit algebraic objects such as those listed in the examples. Here is one form of our main result in this regard.

Theorem 1.6 (Main theorem, simple form). — *Let A be a global K -approximate group, thus it is contained in a (global) group G . Then there exists a subgroup G_0 of G and a finite normal subgroup H of G_0 with the following properties:*

- (i) A can be covered by $O_K(1)$ left-translates of G_0 ;
- (ii) G_0/H is nilpotent and finitely generated of rank³ and step at most $O_K(1)$;
- (iii) A^4 contains H and a generating set of G_0 .

In particular, the group G_0 is finite-by-nilpotent, and hence also virtually nilpotent. Indeed, the stabiliser in G_0 of the conjugation action on H has finite index in G_0 and is a central extension of a finite index subgroup of G_0/H , and therefore is also nilpotent.

By specialising Theorem 1.6 to the combinatorial examples in Remark 1.5 we obtain an analogous structure theorem for sets of small doubling.

Corollary 1.7 (Freiman-type theorem). — *Let A and B be finite non-empty subsets in a (global) group G such that $|AB| \leq K|A|^{\frac{1}{2}}|B|^{\frac{1}{2}}$. Then there exists a subgroup G_0 of G and a finite normal subgroup H of G_0 with the following properties:*

- (i) A can be covered by at most $O_K(1)$ right translates of G_0 ;
- (ii) G_0/H is nilpotent and finitely generated of rank and step $O_K(1)$. In particular, G_0 is finite-by-nilpotent and hence also virtually nilpotent.

² Here and in the rest of the paper we use $X = O_K(Y)$, $X \ll_K Y$, or $Y \gg_K X$ for two (standard) quantities X, Y and a (standard) parameter K to denote the assertion that $|X| \leq C_K Y$ for some (standard) quantity $C_K > 0$ depending only on K , and similarly for other choices of subscripted parameters. We also adopt an analogous notation for nonstandard quantities; see Appendix A.

³ The *rank* of a finitely generated group is the least number of generators required to generate the group. The *step* is the length of the lower central series, minus 1.

Proof. — By [52, Theorem 4.6], there exists a $O(K^{O(1)})$ -approximate group A' of size $O(K^{O(1)}|A|)$ such that A can be covered by $O(K^{O(1)})$ right translates of A' and B can be covered by $O(K^{O(1)})$ left translates of A' . We may thus apply Theorem 1.6 to A' . \square

Theorem 1.6 (or Corollary 1.7) answers in the affirmative a conjecture that we have been referring to as the *Helgott-Lindenstrauss Conjecture*, on account of its having been raised independently in private communications by both Harald Helgott and Elon Lindenstrauss. In fact, the conjecture is reasonably explicit in the comments surrounding [31, Theorem 1.1].

Remark 1.8 (*The linear case*). — Various forms of the main theorem are also known in groups of Lie type of bounded dimension, as a consequence of results of many authors [6–8, 18, 19, 30, 31, 33, 43]. For instance, in [18] an analogue of Theorem 1.6 was established in the case when G is a solvable algebraic group of bounded dimension over a finite field of prime order. In that case, the group G_0/H has bounded rank, and the number of cosets of G_0 needed to cover A is *polynomial* in K . We have no examples to rule out the possibility that this polynomiality in K holds in all groups G , perhaps at the cost of weakening the rank and step bounds on G_0/H . Unfortunately our methods, which rely on ultrafilter arguments, give no quantitative bounds on the covering number whatsoever.

Remark 1.9 (*Bounds on the nilpotent group*). — Our method allows us to give an explicit bound on the dimension (rank and step) of the nilpotent group G_0/H in Theorem 1.6 at the expense of replacing A^4 in item (iii) by a larger power of A . Namely, if we allow for H and the generating set of G_0 to be contained in A^{12} , then we may ensure that the nilpotent group G_0/H is ℓ -nilpotent with $\ell = O(K^2 \log K)$. If we are happy to go as far as $A^{O(K^1)}$, then this may be further reduced to $\ell \leq 6 \log_2 K$. Here we say that a group is ℓ -nilpotent if it admits a generating set u_1, \dots, u_ℓ such that $[u_i, u_j] \in \langle u_{j+1}, \dots, u_\ell \rangle$ for all $i < j$. In particular such a group admits a normal series with cyclic factors of length at most ℓ , and so is also nilpotent of step at most ℓ . We refer the reader to Theorem 2.12 and to Section 10 for a detailed statement and proof.

Remark 1.10. — Note that no bound is provided on the size of the finite group H in Theorem 1.6, other than that it is finite. Indeed, by considering A to be a large finite simple group it is not difficult to see that H can be arbitrarily large.

We will in fact prove a much more precise version of Theorem 1.6 involving a slightly complicated type of approximate group which we call a *coset nilprogression*. We discuss this concept in some detail in the next section. For many applications, however, Theorem 1.6 is quite sufficient.

Applications. — We now give a small selection of applications to growth in groups and to Riemannian geometry; a greater variety is assembled in Section 11, which also contains proofs of these statements.

Polynomial growth conditions and Gromov's theorem. — Firstly, Theorem 1.6 yields a quick proof of Gromov's theorem [27] on groups of polynomial growth.

Theorem 1.11 (*Gromov's theorem*). — *Let G be a group of polynomial growth. That is, G is generated by a finite symmetric set S , and there are constants C and d such that $|S^n| \leq Cn^d$ for all $n \in \mathbf{N}$. Then G is virtually nilpotent.*

Remark 1.12. — In fact, our arguments show that there is some function $f : \mathbf{N} \rightarrow \mathbf{N}, f(n) \rightarrow \infty$, such that, if G does not have polynomial growth, then $|S^n| \geq n^{f(n)}$ for all n . We do not get an explicit function f . However, if the control parameter $O_K(1)$ in Theorem 2.10 were known to be polynomial in K , we could take $f(n) = c \log n$. The best (in fact only) lower bound known for this function at present is $(\log \log n)^c$, due to Shalom and the third author [51]. It is conjectured by some, in the absence of any examples to the contrary, that $f(n) > n^c$, and possibly even that $|S^n| \geq e^{c\sqrt{n}}$.

In [33] Hrushovski also gave a derivation of Gromov's theorem from his Lie model theorem (see Theorem 3.10 below). He in fact proved a strengthening of Gromov's theorem (see [33, Theorem 7.1] or Theorem 11.1 below). We will be able to recover Hrushovski's result more directly (see Corollary 11.2 below). In fact, our approach can also yield the following other strengthening of Gromov's theorem, which is uniform in the size of the generating set S and appears to be new. Recall that if $\ell \in \mathbf{N}$ then we say that a group is ℓ -nilpotent if it admits a generating set u_1, \dots, u_ℓ such that $[u_i, u_j] \in \langle u_{j+1}, \dots, u_\ell \rangle$ for all $i < j$.

Theorem 1.13. — *Let $d > 0$. Then there is $n_0 = n_0(d) > 0$ such that if G is a group generated by a finite symmetric set S with $1 \in S$ for which $|S^n| \leq n^d |S|$ for some $n \geq n_0(d)$, then G is virtually nilpotent. In fact G has a normal subgroup of index at most $O_d(1)$ which is finite-by- $(O(d)$ -nilpotent).*

Proof. — The proof of this (and hence of Theorem 1.11) is a short enough deduction that we can give it here in the introduction. We refer the reader to Section 11 for more details. Let $N = N(d)$ be a large quantity to be specified later, and let n_0 be sufficiently large depending on N and d . By the pigeonhole principle and the hypothesis $|S^n| \leq n^d |S|$ we see that if n_0 is sufficiently large depending on N then there exists $n', N \leq n' \leq n_0/100$, such that $|S^{100n'}| \leq (200)^d |S^{n'}|$. By Corollary 5.2 (which is quite easy) this implies that $S^{2n'}$ is a $e^{O(d)}$ -approximate group. By our main theorem, Theorem 1.6 (and Remark 1.9), we can thus find a finite-by- $(O(d)$ -nilpotent) and hence virtually nilpotent group G_0 such that $S^{2n'}$ is covered by $O_d(1)$ left-translates of G_0 . By the pigeonhole principle, if N is large enough, we can find a nonnegative $m < 2n'$ such that $S^{m+1}G_0 = S^mG_0$. Multiplying on the left by S repeatedly we conclude that $S^{m+k}G_0 = S^mG_0$ for all $k \geq 0$. Since S generates G , we conclude that $G = S^mG_0 = S^{2n'}G_0$. Since $S^{2n'}$ was covered by $O_d(1)$ left-translates of G , G_0 has index $O_d(1)$ in G , and so G is also virtually nilpotent. \square

Riemannian manifolds. — A. Petrunin suggested to us some years ago⁴ that a result such as Corollary 1.13 would give a purely group-theoretical proof of a theorem of Fukaya and Yamaguchi [16] according to which fundamental groups of almost non-negatively curved manifolds are virtually nilpotent. Recall that a closed manifold M is said to be almost non-negatively curved if one can find a sequence g_n of Riemannian metrics on it for which $\text{diam}(M, g_n) \leq 1$ while $K_{M, g_n} \geq -1/n$ where K_{M, g_n} is the sectional curvature. Indeed, a simple application of the Bishop-Gromov inequalities combined with Corollary 11.5 yields the following improvement assuming only a lower bound on the Ricci curvature and an upper bound on the diameter.

Corollary 1.14 (Ricci gap). — *Given $d \in \mathbf{N}$, there is $\varepsilon(d) > 0$ such that the following holds. Let $M = (M, g)$ be an d -dimensional compact Riemannian manifold with Ricci curvature bounded below by $-\varepsilon$ and diameter $\text{diam}(M) \leq 1$. Then $\pi(M)$ is virtually nilpotent.*

This result is known to differential geometers and follows from the works of Cheeger-Colding [11] and Kapovitch and Wilking [36]. We refer the reader to Section 11.1 for more discussion and references concerning the above result. We only note that Corollary 11.5 yields in fact an explicit bound on the nilpotency class, namely that after passing to a subgroup of $\pi_1(M)$ with index $O_d(1)$ and quotienting by a finite normal subgroup, we obtain a $O(d)$ -nilpotent group.

Generalised Margulis lemma. — Another corollary of Theorem 1.6 is a “generalised Margulis lemma” for metric spaces of a type conjectured by Gromov in [28, §5.F]. A metric space X is said to have *bounded packing* with packing constant K if there is $K > 0$ such that every ball of radius 4 in X can be covered by at most K balls of radius 1. Say that a subgroup Γ of isometries of X acts *discretely* on X if every orbit is discrete in the sense that $\{\gamma \in \Gamma : \gamma \cdot x \in \Sigma\}$ is finite for every $x \in X$ and for every bounded set $\Sigma \subseteq X$.

Corollary 1.15 (Generalized Margulis Lemma). — *Let $K \geq 1$ be a parameter. Then there is some $\varepsilon(K) > 0$ such that the following is true. Suppose that X is a metric space with packing constant K , and that Γ is a subgroup of isometries of X which acts discretely. Then for every $x \in X$ the “almost stabiliser” $\Gamma_\varepsilon(x) = \langle S_\varepsilon(x) \rangle$, where $S_\varepsilon(x) := \{\gamma \in \Gamma : d(\gamma \cdot x, x) < \varepsilon\}$, is virtually nilpotent.*

Note that the space X is not assumed to be a manifold. The traditional Margulis lemma establishes a similar statement for subgroups of isometries of pinched negatively curved manifolds, or more generally under a curvature lower bound.

Approximate groups and polynomial growth. — Finally we remark on an additive-combinatorial application, which asserts that approximate groups have large subsets with “polynomial growth”.

⁴ See also <http://mathoverflow.net/questions/11091>.

Theorem 1.16 (*Approximate groups are locally of polynomial growth*). — Suppose that A is a global K -approximate group. Then A^4 contains a $O_K(1)$ -approximate group A' with $(A')^4 \subset A^4$ and $|A'| \gg_K |A|$ such that $|(A')^m| \ll_K m^{O_K(1)} |A'|$ for all $m \geq 1$.

This theorem is an immediate consequence of Theorem 2.10 and Proposition C.5 below.

Remark 1.17. — The above argument converted nilpotent structure (or more precisely, coset nilprogression structure, see below) to polynomial growth. In the reverse direction, there is the result of Sanders [46] in certain monomial groups, in which polynomial growth is shown to imply a metric ball type structure, at least under the (rather strong) restriction that the approximate group A is normal in the ambient group G .

2. Coset nilprogressions and a more detailed version of the Main Theorem

This section concerns the more precise variants of our main theorem, whose existence we hinted at in the first introductory section. Let us first recall the fundamental inverse sumset theorem for abelian approximate groups. This was first introduced by Freiman [15], and a simplified argument was subsequently given in the paper [44] of Ruzsa. Here is the theorem in the torsion-free setting. Recall the notion of a generalised arithmetic progression, defined in Example 3 above.

Theorem 2.1 (*Freiman-Ruzsa theorem*). — Let $G = (G, +)$ be a torsion-free (global) abelian group, and let $K \geq 2$ be a parameter. Suppose that $A \subseteq G$ is a K -approximate group. Then $4A = A + A + A + A$ contains a generalised arithmetic progression

$$P = P(u_1, \dots, u_r; N_1, \dots, N_r)$$

with $r \leq \log^{O(1)} K$ and $|P| \gg e^{-\log^{O(1)} K} |A|$. In particular A can be covered by $O(e^{\log^{O(1)} K})$ translates of P .

Proof. — See [48] for the main part of this; the final assertion is then a consequence of Ruzsa’s covering lemma, Lemma 5.1. For earlier results of this type with weaker bounds on r and P , see [10, 44]. In [26] it was noted that one can take r as small as $\lfloor \log_2 K + \varepsilon \rfloor$ for any $\varepsilon > 0$, at the cost of decreasing the size of $|P|$ somewhat; see also [3, 4] for prior results along these lines. \square

Roughly speaking, Theorem 2.1 asserts that, in a global torsion-free abelian group such as the integers \mathbf{Z} , approximate groups are “controlled” by generalised arithmetic progressions of bounded rank. In the case of abelian groups with torsion, the class of generalised arithmetic progressions is not sufficient, as one must also now deal with the

example of finite genuine groups (Example 1). It is thus natural to introduce the concept of a *coset progression* $H + P$: the sum of a finite genuine group H and a generalised arithmetic progression $P = P(u_1, \dots, u_r; N_1, \dots, N_r)$. This concept is sufficient for the formulation of a Freiman type theorem in an arbitrary abelian group.

Theorem 2.2 (*Abelian Freiman-Ruzsa theorem*). — *Let $G = (G, +)$ be a (global) abelian group, and let $K \geq 2$ be a parameter. Suppose that $A \subseteq G$ is a K -approximate group. Then $4A$ contains a coset progression $H + P$, where*

$$P = P(u_1, \dots, u_r; N_1, \dots, N_r)$$

is a generalised arithmetic progression with $r \leq \log^{O(1)} K$, H is a finite abelian subgroup disjoint from P , and $|H + P| = |H||P| \gg e^{-\log^{O(1)} K} |A|$. In particular, A can be covered by $O(e^{\log^{O(1)} K})$ translates of $H + P$.

Proof. — Again, see [48]; see also [25] for an earlier result in this direction. \square

We turn now to the business of dropping the commutativity assumption. We will also drop the assumption that A is contained in a *global* group and merely assume that A is a subset of a *local* group G . Informally, this means that we will not require the multiplication law to be defined everywhere in G , but only in a certain neighborhood of id . We refer the reader to Appendix B for a precise definition and basic properties; see also [50, IV.3] for a discussion of the closely related notion of *group chunk*. We generalise the concept of a generalised arithmetic progression to this setting as follows.

Definition 2.3 (*Non-commutative progression*). — *Let u_1, \dots, u_r be r elements in a local group $G = (G, \cdot)$, and let N_1, \dots, N_r be r positive real numbers. If all products $g_1 \dots g_n$ are well-defined in G , where each g_i is equal to one of u_j or u_j^{-1} and, for each $j = 1, \dots, r$, the formal expression⁵ u_j and its inverse u_j^{-1} appear at most N_j times, then we call the set of such products a non-commutative progression of rank r and side lengths N_1, \dots, N_r and we denote it by $P(u_1, \dots, u_r; N_1, \dots, N_r)$. We refer to r as the rank of the non-commutative progression.*

Remark 2.4. — One can view non-commutative progressions as multiparameter variants of balls in a word metric. For instance when all N_j take the same value N and one is working in a global group, the progression $P(u_1, \dots, u_r; N, \dots, N)$ is comparable with the word ball $B(N)$ of radius N in the group $\langle u_1, \dots, u_r \rangle$ for the word metric with generating set $\{u_1, \dots, u_r\}$ in the sense that $B(N) \subseteq P(u_1, \dots, u_r; N, \dots, N) \subseteq B(rN)$.

In the global abelian setting, all generalised arithmetic progressions of bounded rank are automatically approximate groups with a bounded covering parameter K . This

⁵ For this definition, we consider u_i and u_j to be distinct formal expressions when $i \neq j$, even if u_i and u_j take the same value in G , and similarly for u_i^{-1} , u_j^{-1} . Thus, for instance, $P(u_1, u_2; 1, 1)$ contains $u_1 u_2$ even if u_1, u_2 are equal.

is not the case in general non-abelian groups, even in the global setting. For instance, if F is the free non-abelian group on two generators e_1, e_2 , then the non-commutative progression $P(e_1, e_2; \mathbb{N}, \mathbb{N})$ (which, as remarked earlier, is essentially the ball of radius N in F) grows exponentially in N , and one can easily verify that $P(e_1, e_2; \mathbb{N}, \mathbb{N})$ is only a K -approximate group for K growing exponentially in N . However, the situation is much closer to the abelian case if the ambient group G is *nilpotent*. Given the link between progressions and balls, the reader familiar with Gromov's theorem on groups of polynomial growth [27] (to be discussed later on) will not find this surprising. Indeed, it can be shown (though we will not do so here) that if G is a global nilpotent group of step s , a non-commutative progression $P(u_1, \dots, u_r; N_1, \dots, N_r)$ in G will be a $O_{r,s}(1)$ -approximate group if N_1, \dots, N_r are sufficiently large depending on r and s .

This motivates the following definition. Given some generators u_1, \dots, u_r , let us recursively define an *iterated commutator of degree k* involving these generators for a natural number $k \geq 1$ by declaring $u_1^{\pm 1}, \dots, u_r^{\pm 1}$ to be the iterated commutators of degree 1, and $[g, h]$ to be a iterated commutator of degree $j + k$ whenever g, h are iterated commutators of weight j, k respectively for some $j, k \geq 1$. Thus for instance $[[u_2, u_3^{-1}], [u_2^{-1}, u_4]]$ is an iterated commutator of u_1, u_2, u_3, u_4 of degree 4.

Definition 2.5 (*Nilprogression*). — Suppose that G is a local group and that $s \geq 0$ is an integer. A nilprogression of rank r and step s is a non-commutative progression $P(u_1, \dots, u_r; N_1, \dots, N_r)$ with the property that every iterated commutator of degree $s + 1$ in the generators u_1, \dots, u_r is well-defined and equals the identity id .

Example 9. — The generalised arithmetic progressions $P(u_1, \dots, u_r; N_1, \dots, N_r)$ in Example 3 is a nilprogression (in additive notation) of rank r and step 1. The set $P(u_1, u_2; N_1, N_2)$ in Example 6 is a nilprogression of rank 2 and step 2.

It can be shown (though we shall not do so here) that if N_1, \dots, N_r are sufficiently large depending on r, s , and $P(u_1, \dots, u_r; CN_1, \dots, CN_r)$ is a well-defined nilprogression of step s for some sufficiently large C depending on r, s , then $P(u_1, \dots, u_r; N_1, \dots, N_r)$ is a $O_{r,s}(1)$ -approximate group.

The concept of a nilprogression as defined above is related to, though not quite identical with, the one given in [5]. As a byproduct of our proof methods, we will be able to work with a more tractable subclass of nilprogressions, which we will call *nilprogressions in C -normal form*. These generalise the notion of a *proper* generalised arithmetic progression in the additive combinatorics literature, and are also close in spirit to the nilprogressions introduced in [53].

Definition 2.6 (*C -normal form*). — Let $C \geq 1$. A non-commutative progression

$$P(u_1, \dots, u_r; N_1, \dots, N_r)$$

is said to be in C -normal form if the following axioms are obeyed.

- (i) (*Upper-triangular form*) For every i, j with $1 \leq i < j \leq r$ and for all four choices of signs \pm one has

$$(2.1) \quad [u_i^{\pm 1}, u_j^{\pm 1}] \in P\left(u_{j+1}, \dots, u_r; \frac{CN_{j+1}}{N_i N_j}, \dots, \frac{CN_r}{N_i N_j}\right).$$

In particular, $[u_i, u_r] = \text{id}$ whenever $1 \leq i < r$.

- (ii) (*Local properness*) The expressions $u_1^{n_1} \dots u_r^{n_r}$ are distinct as n_1, \dots, n_r range over integers with $|n_i| \leq \frac{1}{C}N_i$, $i = 1, \dots, r$.
- (iii) (*Volume bound*) One has

$$(2.2) \quad \frac{1}{C}(2\lfloor N_1 \rfloor + 1) \cdots (2\lfloor N_r \rfloor + 1) \leq |\mathbf{P}| \leq C(2\lfloor N_1 \rfloor + 1) \cdots (2\lfloor N_r \rfloor + 1).$$

The somewhat ugly expression $(2\lfloor N_1 \rfloor + 1) \cdots (2\lfloor N_r \rfloor + 1)$ is convenient to have in (2.2) for some minor technical reasons, but it would not do much harm for the reader to mentally substitute $N_1 \dots N_r$ for this expression instead if desired. The volume bound (2.2) is morally (up to some degradation in the constants C) implied by the other axioms of a nilprogression in C -normal form, when the N_1, \dots, N_r are sufficiently large, and one is working in a global group (or at least if one assumes $P(u_1, \dots, u_r; DN_1, \dots, DN_r)$ to be well-defined for some sufficiently large $D = D_{r,s}$), but for some further minor technical reasons it is convenient to state this bound explicitly in the definition.

Example 10. — The generalised arithmetic progressions $P(u_1, \dots, u_r; N_1, \dots, N_r)$ in Example 3 will be in 1-normal form if it is proper, i.e. if all the expressions $n_1 u_1 + \dots + n_r u_r$ for $|n_i| \leq N_i$ are distinct.

Example 11. — The set $P(u_1, u_2; N_1, N_2)$ in Example 6 is not in C -normal form for any bounded C , because $[u_1, u_2]$ is non-trivial. However, the closely related nilprogression

$$P(u_1, u_2, [u_1, u_2]; N_1, N_2, N_1 N_2)$$

of rank 3 and step 2 is in 1-normal form. The two sets are “comparable” in a number of ways; for instance, one can easily verify that

$$\begin{aligned} P\left(u_1, u_2; \frac{1}{C}N_1, \frac{1}{C}N_2\right) &\subset P(u_1, u_2, [u_1, u_2]; N_1, N_2, N_1 N_2) \\ &\subset P(u_1, u_2; CN_1, CN_2) \end{aligned}$$

for some absolute constant C (e.g. one can take $C = 100$).

Remark 2.7. — Note that in the global group case, the step of a nilprogression in C -normal form is less or equal to its rank.

In Lemma C.1 we will show that any non-commutative progression $P(u_1, \dots, u_r; N_1, \dots, N_r)$ in C-normal form is “essentially” a $O_{r,C}(1)$ -approximate group. More precisely, we will show that $P(u_1, \dots, u_r; \varepsilon N_1, \dots, \varepsilon N_r)$ is a $O_{r,C,\varepsilon}(1)$ -approximate group whenever $\varepsilon > 0$ is sufficiently small and the N_i ’s are sufficiently large depending on C, r . We will also show that every element of $P(u_1, \dots, u_r; \varepsilon N_1, \dots, \varepsilon N_r)$ can be rewritten in the form $u_1^{n_1} \dots u_r^{n_r} h$, where $h \in H$ and $|n_i| = O_{r,s}(\varepsilon N_i)$, while conversely every such product with $|n_i| \leq \varepsilon N_i$ obviously belongs to $P(u_1, \dots, u_r; \varepsilon N_1, \dots, \varepsilon N_r)$.

Just as in the abelian case, we need to account for genuine subgroups. The analogue of coset progression is a coset nilprogression, a concept we first define in the simpler setting of global groups.

Definition 2.8 (*Global coset nilprogression*). — Let G be a (global) group. By a coset nilprogression of rank r and step s in G , we mean a set P of the form $\pi^{-1}(Q)$, where G_0 is a subgroup of G , H is a finite normal subgroup of G_0 , $\pi : G_0 \rightarrow G_0/H$ is the quotient map, and Q is a nilprogression of rank r and step s in G_0/H .

We say that P is in C-normal form if Q is in C-normal form.

We can extend this definition to local groups, using the local notion of quotient group reviewed in Lemma B.12.

Definition 2.9 (*Local coset nilprogression*). — Let G be a (local) group, which we endow with the discrete topology. By a coset nilprogression of rank r and step s in G , we mean a set P of the form $\pi^{-1}(Q)$, where H is a finite genuine subgroup of G with a cancellative normalising neighbourhood G_0 , W is a neighbourhood of H in G_0 with $W^6 \subset G_0$, $WH = HW = W$, $\pi : W \rightarrow W/H$ is the quotient map defined in Lemma B.12, and Q is a nilprogression of rank r and step s in W/H .

We say that P is in C-normal form if Q is in C-normal form.

We call H the finite group associated with P , and Q the nilprogression associated with P . If $Q = P(u_1, \dots, u_r; N_1, \dots, N_r)$, then we write $P = P_H(u_1, \dots, u_r; N_1, \dots, N_r)$.

Example 12. — A subgroup is a coset nilprogression of rank 0 and step 0. More generally, the direct product of a subgroup with a nilprogression of rank r and step s is a coset nilprogression of rank r and step s . The coset nilprogression will be in C-normal form if the associated nilprogression is.

Example 13. — The set A constructed in Example 8 is a coset nilprogression of rank 1 and step 1, and is also in 1-normal form as long as $N < \frac{p-1}{2}$.

Again, coset nilprogressions in normal form are essentially approximate groups; see Lemma C.1 for a precise version of this statement.

We are now ready to state our main technical theorem, which among other things implies Theorem 1.6, and whose proof will occupy the bulk of this paper.

Theorem 2.10 (Main theorem). — Let A be a K -approximate group. Then A^4 contains a coset nilprogression P of rank and step $O_K(1)$ and $|P| \gg_K |A|$. Furthermore, P can be taken to be in $O_K(1)$ -normal form.

We remark that precursor results to this theorem in the case of nilpotent or solvable groups were obtained in [5, 6, 14, 18, 52, 53]. Theorem 2.10 also provides an independent proof of a qualitative version of the abelian results of Theorem 2.1 and Theorem 2.2, which, in contrast to the other known proofs of these results, manages to almost completely avoid the use of Fourier analysis.⁶

It is easy to see that Theorem 2.10 implies Theorem 1.6, by taking G_0 to be the global group generated by P . The key point here is that a group generated by a set u_1, \dots, u_r is nilpotent of step at most s if every iterated commutator of the u_1, \dots, u_r of degree $s + 1$ is trivial. A proof of this assertion may be found in Hall's book [29].

By standard non-commutative product estimates, we can also establish the following Freiman-type theorem for sets of bounded doubling.

Corollary 2.11 (Freiman-type theorem). — Let $K \geq 1$. Let G be a (global) group and A, B be finite non-empty subsets of G such that $|AB| \leq K|A|^{1/2}|B|^{1/2}$. Then there exists a coset nilprogression P of rank and step $O_K(1)$ with $|P| \gg_K |A|$ which is in $O_K(1)$ -normal form, such that A can be covered by $O_K(1)$ left-translates of P , and B can be covered by $O_K(1)$ right-translates of P .

Proof. — This follows immediately from combining Theorem 2.10 with [52, Theorem 4.6]. \square

In Section 10, we will show the following explicit bounds on the rank and step of P .

Theorem 2.12 (Bounds on the rank and step of the nilprogression). — In Corollary 2.11 (and in Theorem 2.10 if A is assumed to be a global K -approximate group), at the expense of replacing the conclusion $P \subseteq A^4$ with the weaker statement that $P \subseteq A^{12}$, the coset nilprogression P can be taken to have rank and step at most $O(K^2 \log K)$ while remaining in $O_K(1)$ -normal form. Moreover, if we settle for the even weaker inclusion $P \subset A^{O_K(1)}$, one can ensure that P has rank and step at most $6 \log_2 K$ (while still remaining in $O_K(1)$ -normal form).

It is likely that the numerical constants 6 and 12 here can be improved, but we will not pursue such improvements here.

Local approximate groups can be embedded in global groups. — As we have remarked above, the approximate groups A considered in this paper are *local* in the sense that we do not

⁶ However, our argument still uses results relating to Hilbert's fifth problem which require Fourier-analytic tools, such as Pontryagin duality, even in the abelian setting.

need to assume that A lies in a global group G . However as a consequence of Theorem 2.10, the more detailed version of our main theorem, we have the following statement. It asserts that, at least at the qualitative level, there is in fact no loss of generality in dealing with the global case.

Theorem 2.13. — *Suppose that A is a K -approximate group. Then A^4 contains a $O_K(1)$ -approximate group A' with $(A')^4 \subset A^4$ and $|A'| \gg_K |A|$ which is isomorphic to a subset of a global group G .*

This theorem follows from Theorem 2.10 and the fact (which we prove in Lemma C.3) that a large portion of a coset nilprogression in normal form can be embedded in a global group. This theorem can be viewed as a discrete analogue to a recent result of Goldbring and van den Dries [56], who established that every locally compact local group is locally isomorphic near the identity to a locally compact *global* group (thus there is a neighbourhood of the identity in the former group that is isomorphic to a neighbourhood of the identity in the latter group). One should also compare this result with Lie’s third theorem that every local Lie group is locally isomorphic to a global Lie group (see Theorem B.16 and the discussion in Serre’s book [50]).

3. Ultra approximate groups and Hrushovski’s Lie Model Theorem

In the next section we will give an outline of the argument we shall use to prove Theorem 2.10. An extremely important component of it will be a *Lie Model Theorem* that implicitly appears in a remarkable paper of Hrushovski [33, Theorem 4.2], which provided the foundation for much of the work here, and for which we will give a self-contained proof later in this paper. We can state this theorem very informally as follows:

Theorem 3.1 (Hrushovski’s Lie Model Theorem, informal version). — *In a suitable limit, an approximate group is virtually modelled by a precompact neighbourhood of the identity in a Lie group.*

Of course, to make this theorem more precise, one has to formalise terms such as “suitable limit”, “virtually”, and “modelled”. We shall do so presently, but first we point out that Theorem 3.1 is very similar in spirit to a key step [27, §7] in Gromov’s proof of his celebrated theorem on groups of polynomial growth, which we state informally as follows.

Theorem 3.2 (Gromov’s Lie Model Theorem, informal version). — *In a suitable limit, a group of polynomial growth can be modeled by a finite-dimensional locally compact space with a transitive isometric action of a Lie group.*

To deepen the analogy between the two results, we note that Theorem 3.1 and Theorem 3.2 both require the deep body of results surrounding the solution to Hilbert’s

fifth problem on the topological description of the category of Lie groups (see [40]) in order to bring into view the Lie structure, which is not manifestly present when one first takes a limit. There are however some technical differences between the precise formulations of Theorem 3.1 and Theorem 3.2. In the latter theorem, one has a group G (of polynomial growth) generated by a finite set S . This gives a metric on G , the word metric given by the generating set S . Gromov then looks at the discrete balls S^n , $n = 1, 2, 3, \dots$ “from a distance” to get some continuous limit metric space X . For example if $G = \mathbf{Z}$ and $S = \{-1, 0, 1\}$, then $S^n = \{-n, \dots, n\}$, and it is heuristically clear that these discrete intervals S^n , after rescaling by n , “converge” in a suitable sense to the continuous interval $[-1, 1] \subseteq \mathbf{R}$.

To effect this limit, Gromov introduced what is now known as *Gromov-Hausdorff convergence* of a sequence of metric spaces. In subsequent work of van der Dries and Wilkie [58] a slightly different approach, using ultralimits (or non-standard analysis) was pioneered. This construction is now known, in the geometric group theory literature, as the *asymptotic cone*.

The asymptotic cone, then, is (a quotient of) an ultraproduct of the sequence of balls $(S^n)_{n \in \mathbf{N}}$. We will use a similar limit⁷ in order to formalise Theorem 3.1, namely an *ultraproduct* \mathbf{A} of an arbitrary sequence $(A_n)_{n \in \mathbf{N}}$ of \mathbf{K} -approximate groups, an object we call an *ultra approximate group*. We now define this term more precisely.

Definition 3.3 (*Ultra approximate group*). — *Throughout this paper, we fix a non-principal ultrafilter $\alpha \in \beta\mathbf{N} \setminus \mathbf{N}$ (see Lemma A.1 for a definition of this concept). If $\mathbf{K} > 0$ is a real number then an ultra \mathbf{K} -approximate group is an ultraproduct $\mathbf{A} := \prod_{n \rightarrow \alpha} A_n$, where each A_n is a (standard) \mathbf{K} -approximate group. Thus, \mathbf{A} is the space of all formal limits $\lim_{n \rightarrow \alpha} a_n$ with $a_n \in A_n$, where two formal limits $\lim_{n \rightarrow \alpha} a_n$ and $\lim_{n \rightarrow \alpha} a'_n$ are considered equal if $a_n = a'_n$ for all n sufficiently close to α (i.e. for all n in an α -large subset of \mathbf{N}). See Appendix A for more discussion on ultraproducts. Often we will not need to refer to \mathbf{K} explicitly, in which case we speak simply of an ultra approximate group.*

Note that we allow the approximate groups A_n to lie in different ambient groups G_n (much as the notion of Gromov-Hausdorff convergence also does not require the spaces X_n involved to all live in a common ambient space). Ultraproducts are a model-theoretic limit, in contrast to the more geometric notion of a limit defined by Gromov-Hausdorff convergence. There are two key properties of these model-theoretic limits that make them convenient to use for our purposes. The first is *Łoś’s theorem*, which roughly speaking asserts that any property that can be stated in the language of first-order logic holds for an ultraproduct $\mathbf{A} = \prod_{n \rightarrow \alpha} A_n$ if and only if it holds for those A_n with n sufficiently close to α ; see Theorem A.6. The second is *countable saturation*, which we will use to establish the completeness of a certain (pseudo)metric space associated to an ultra approximate group; see Proposition 6.1.

⁷ In [33], more saturated limits (not necessarily constructed using ultrafilters) were also considered, but we will not need such constructions here.

Next, we discuss what it would mean to⁸ “model” an ultra approximate group \mathbf{A} . Informally, a model would seek to describe the “coarse-scale” behaviour of \mathbf{A} , and in particular be able to predict when an orbit $\text{id}, a, a^2, a^3, \dots$ of an element a of \mathbf{A} will “escape” \mathbf{A} , while ignoring the “fine-scale” behaviour of \mathbf{A} . Such a model will be formalised by a homomorphism $\phi : \mathbf{A} \rightarrow \mathbf{L}$ of local groups that obey certain good properties (see Definition 3.5 below). Before we present this formal definition, though, we first discuss some key examples of ultra approximate groups and their models.

Example 14 (Nonstandard finite groups). — Suppose that A_n is a sequence of (standard) finite groups; then the ultraproduct $\mathbf{A} := \prod_{n \rightarrow \alpha} A_n$ is an ultra approximate group. In this case, \mathbf{A} is in fact a genuine group, with group operation given by the law

$$\left(\lim_{n \rightarrow \alpha} a_n \right) \cdot \left(\lim_{n \rightarrow \alpha} b_n \right) := \lim_{n \rightarrow \alpha} (a_n b_n).$$

We will refer to such groups as *nonstandard finite groups*. A typical example of a nonstandard finite group is the nonstandard cyclic group⁹

$$\mathbf{Z}/\mathbf{NZ} := \prod_{n \rightarrow \alpha} \mathbf{Z}/n\mathbf{Z},$$

where $\mathbf{N} \in {}^*\mathbf{N}$ is the nonstandard natural number

$$(3.1) \quad \mathbf{N} := \lim_{n \rightarrow \alpha} n.$$

In a nonstandard finite group \mathbf{A} , there are no elements that ever escape \mathbf{A} : if $a \in \mathbf{A}$, then one has $a^n \in \mathbf{A}$ for all $n \in \mathbf{N}$. As such, it will turn out that \mathbf{A} can be modeled by a trivial homomorphism $\phi : \mathbf{A} \rightarrow \{\text{id}\}$ to the trivial group.

Example 15 (Nonstandard intervals). — Now consider the sequence $A_n := P(1; n) = \{-n, \dots, n\}$ of (standard) arithmetic progressions in \mathbf{Z} . The ultraproduct $\mathbf{A} := \prod_{n \rightarrow \alpha} A_n$ can be viewed as the nonstandard arithmetic progression $\mathbf{A} = P(1; \mathbf{N}) = \{-\mathbf{N}, \dots, \mathbf{N}\}$ in the nonstandard integers ${}^*\mathbf{Z} := \prod_{n \rightarrow \alpha} \mathbf{Z}$, where \mathbf{N} was defined in (3.1). Then \mathbf{A} is an ultra approximate group, and it can also be viewed as a local group inside the nonstandard integers ${}^*\mathbf{Z}$.

Consider now the map $\pi : \mathbf{A} \rightarrow \mathbf{R}$ defined by

$$\pi \left(\lim_{n \rightarrow \alpha} a_n \right) := \text{st} \lim_{n \rightarrow \alpha} \frac{a_n}{n},$$

⁸ Our use of the term “model” here is not, strictly speaking, the precise notion that is used in model theory, but is closer to the notion of a “Freiman model” from additive combinatorics, as used for instance in [25, 44].

⁹ This group is the analogue of the profinite completion $\hat{\mathbf{Z}} = \varprojlim \mathbf{Z}/n\mathbf{Z}$ of the integers, but is built using the machinery of ultralimits rather than inverse limits. The two groups are however not identical. For instance, $\hat{\mathbf{Z}}$ is torsion-free, whereas \mathbf{Z}/\mathbf{NZ} can contain torsion; for example if \mathbf{N} is even, or equivalently if the set of even natural numbers is α -large, then \mathbf{Z}/\mathbf{NZ} contains the element $\mathbf{N}/2 \pmod{\mathbf{N}}$, which has order 2. But see Remark 3.4 below for a link between ultraproducts and inverse limits.

where $\text{st}x$ is the standard part of a nonstandard real x (see Appendix A). Thus, for every standard $\varepsilon > 0$, one has

$$\pi\left(\lim_{\mathbf{n} \rightarrow \alpha} a_{\mathbf{n}}\right) - \varepsilon \leq \frac{a_{\mathbf{n}}}{\mathbf{n}} \leq \pi\left(\lim_{\mathbf{n} \rightarrow \alpha} a_{\mathbf{n}}\right) + \varepsilon$$

for all \mathbf{n} sufficiently close to α . One may also write

$$\pi(a) = \text{st} \frac{a}{\mathbf{N}}$$

for all $a \in \mathbf{A}$. The map π is a homomorphism of local groups from \mathbf{A} into $[-1, 1]$. It is surjective since, for any $\gamma \in [-1, 1]$, the nonstandard integer

$$x := \lfloor \gamma \mathbf{N} \rfloor = \lim_{\mathbf{n} \rightarrow \infty} \lfloor \gamma \mathbf{n} \rfloor,$$

where $\lfloor \cdot \rfloor$ is the integer part function, has image $\pi(x) = \gamma$. The kernel $\ker(\pi)$ is the set of $x \in \mathbf{A}$ with $x = o(\mathbf{N})$ (thus if $x = \lim_{\mathbf{n} \rightarrow \alpha} x_{\mathbf{n}}$ and $\varepsilon > 0$ is standard, then $|x_{\mathbf{n}}| \leq \varepsilon \mathbf{n}$ an α -large set of \mathbf{n}). For instance, every standard integer lies in $\ker(\pi)$, as do some nonstandard integers such as $\lfloor \sqrt{\mathbf{N}} \rfloor = \lim_{\mathbf{n} \rightarrow \infty} \lfloor \sqrt{\mathbf{n}} \rfloor$.

There are similar maps from¹⁰ \mathbf{A}^m to $[-m, m]$ for any fixed natural number m , which by abuse of notation we also call π . Informally, these maps *model* \mathbf{A} by the interval $[-1, 1]$, and more generally *model* \mathbf{A}^m by $[-m, m]$. In this particular case, the model $\pi : \mathbf{A} \rightarrow [-1, 1]$ of the ultraproduct \mathbf{A} can be viewed as a limiting object for models $\pi_{\mathbf{n}} : \mathbf{A}_{\mathbf{n}} \rightarrow [-1, 1]$ of the individual factors $\mathbf{A}_{\mathbf{n}}$, by defining $\pi_{\mathbf{n}}(a) := \frac{a}{\mathbf{n}}$. However, in more general situations, the model for the ultraproduct is only a limit for *approximate* models of the factors, and this is one reason why we need to work in the ultraproduct setting as much as we do.

The model $\pi : \mathbf{A}^m \rightarrow [-m, m]$ is not injective: if $\pi(a)$ is trivial, this does not imply that a is trivial. However, π does have an injectivity-like property which will be important later, which roughly speaking asserts that if $\pi(a)$ is *small*, then a is *small*. For instance, observe that if $a \in \mathbf{A}^{1000}$ is such that¹¹ $\pi(a) \in (-1, 1)$, then $a \in \mathbf{A}$. This property on the model π can be used to derive some important facts about the ultraproduct \mathbf{A} ; for instance it implies the *escape property* that if a, a^2, \dots, a^{100} all lie in \mathbf{A}^{10} , then a lies in \mathbf{A} . These sorts of escape properties will play a major role in our arguments in later sections.

Example 16 (Generalised arithmetic progression). — We still work in the integers \mathbf{Z} , but now take $\mathbf{A}_{\mathbf{n}}$ to be the rank two generalised arithmetic progression

$$\mathbf{A}_{\mathbf{n}} := \text{P}(1, \mathbf{n}^{10}; \mathbf{n}, \mathbf{n}) := \{a + b\mathbf{n}^{10} : a, b \in \{-\mathbf{n}, \dots, \mathbf{n}\}\}.$$

¹⁰ Strictly speaking, as we are currently in an additive setting, one should write $m\mathbf{A} = \mathbf{A} + \dots + \mathbf{A}$ rather than $\mathbf{A}^m = \mathbf{A} \cdot \dots \cdot \mathbf{A}$ here.

¹¹ This claim is not quite true when $\pi(a)$ is -1 or $+1$, as can be seen for instance by considering $a = \mathbf{N} + 1 = \lim_{\mathbf{n} \rightarrow \alpha} \mathbf{n} + 1$.

Then the ultraproduct $\mathbf{A} := \prod_{n \rightarrow \alpha} A_n$ is the subset of the nonstandard integers ${}^*\mathbf{Z}$ of the form

$$\mathbf{A} = P(1, N^{10}; N, N) = \{a + bN^{10} : a, b \in \{-N, \dots, N\}\}.$$

This is an ultra approximate group which can be modeled by the Euclidean plane \mathbf{R}^2 , using the model maps $\pi : \mathbf{A}^m \rightarrow \mathbf{R}^2$ defined for each standard m by the formula

$$\pi(a + bN^{10}) := \left(\text{st}_{\mathbf{N}} \frac{a}{N}, \text{st}_{\mathbf{N}} \frac{b}{N} \right)$$

whenever $a, b = O(N)$. The image $\pi(\mathbf{A}^m)$ is then the square $[-m, m]^2$. As before, if $a \in \mathbf{A}^{1000}$ is such that $\pi(a) \in (-1, 1)^2$, then $a \in \mathbf{A}$; this can be used to conclude that if $a, a^2, \dots, a^{100} \in \mathbf{A}^{10}$, then $a \in \mathbf{A}$. Note here that while \mathbf{A} lives in a “one-dimensional” group ${}^*\mathbf{Z}$, the model \mathbf{R}^2 is “two-dimensional”. This is also reflected in the volume growth of the powers A_n^m of A_n for small m and large n , which grow quadratically rather than linearly in m .

Example 17 (Heisenberg box, I). — This example is related to the Heisenberg example in Example 6. We take each A_n to be the “nilbox”

$$(3.2) \quad A_n := \left\{ \begin{pmatrix} 1 & x_n & z_n \\ 0 & 1 & y_n \\ 0 & 0 & 1 \end{pmatrix} \in \begin{pmatrix} 1 & \mathbf{Z} & \mathbf{Z} \\ 0 & 1 & \mathbf{Z} \\ 0 & 0 & 1 \end{pmatrix} : |x_n|, |y_n| \leq n, |z_n| \leq n^2 \right\}.$$

This is not quite an approximate group because it is not quite symmetric (cf. Example 6), but we will ignore this technicality for sake of exposition. In any case it can be repaired in a number of ways, for instance by replacing A_n with $A_n \cup A_n^{-1}$. Once again we consider the ultraproduct $\mathbf{A} := \prod_{n \rightarrow \alpha} A_n$; this is a subset of the nilpotent (nonstandard) group $\begin{pmatrix} 1 & {}^*\mathbf{Z} & {}^*\mathbf{Z} \\ 0 & 1 & {}^*\mathbf{Z} \\ 0 & 0 & 1 \end{pmatrix}$, consisting of all elements $\begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix}$ with $|x|, |y| \leq N$ and $|z| \leq N^2$; again, this is a (discrete) local group.

Consider now the map

$$\pi : \mathbf{A}^3 \rightarrow \begin{pmatrix} 1 & \mathbf{R} & \mathbf{R} \\ 0 & 1 & \mathbf{R} \\ 0 & 0 & 1 \end{pmatrix}$$

defined by

$$(3.3) \quad \pi \left(\begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \right) := \begin{pmatrix} 1 & \text{st}_{\mathbf{N}} \frac{x}{N} & \text{st}_{\mathbf{N}^2} \frac{z}{N^2} \\ 0 & 1 & \text{st}_{\mathbf{N}} \frac{y}{N} \\ 0 & 0 & 1 \end{pmatrix}.$$

This is easily seen to be a homomorphism (of local groups) to the Heisenberg group, whose image is the compact set

$$(3.4) \quad \left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \in \begin{pmatrix} 1 & \mathbf{R} & \mathbf{R} \\ 0 & 1 & \mathbf{R} \\ 0 & 0 & 1 \end{pmatrix} : |x|, |y|, |z| \leq 1 \right\}.$$

Informally, π models \mathbf{A} (or \mathbf{A}^8) by what is essentially a unit ball in this Lie group. As before, we have the injectivity-like property that if $a \in \mathbf{A}^{1000}$ is such that $\pi(a)$ is sufficiently close to the identity, then $a \in \mathbf{A}$; as such, one can again establish the *escape property* that if a, a^2, \dots, a^{100} all lie in \mathbf{A}^{10} , then a lies in \mathbf{A} .

Example 18 (Heisenberg box, II). — This is a variant of the preceding example, in which the (not quite) approximate groups \mathbf{A}_n now take the form

$$(3.5) \quad \mathbf{A}_n := \left\{ \begin{pmatrix} 1 & x_n & z_n \\ 0 & 1 & y_n \\ 0 & 0 & 1 \end{pmatrix} : |x_n|, |y_n| \leq n, |z_n| \leq n^{10} \right\}$$

so that the ultralimit $\mathbf{A} := \prod_{n \rightarrow \alpha} \mathbf{A}_n$ takes the form

$$\mathbf{A} := \left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \in \begin{pmatrix} 1 & *Z & *Z \\ 0 & 1 & *Z \\ 0 & 0 & 1 \end{pmatrix} : |x|, |y| \leq N, |z| \leq N^{10} \right\}.$$

Now consider the map

$$\pi : \mathbf{A}^8 \rightarrow \mathbf{R}^3$$

defined by

$$\pi \left(\begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \right) = \left(\text{st} \frac{x}{N}, \text{st} \frac{y}{N}, \text{st} \frac{z}{N^{10}} \right).$$

The image of this map is the unit cube $[-1, 1]^3$, and is in particular compact. It is also a homomorphism of local groups, since

$$\pi \left(\begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & x' & z' \\ 0 & 1 & y' \\ 0 & 0 & 1 \end{pmatrix} \right) = \left(\text{st} \frac{x+x'}{N}, \text{st} \frac{y+y'}{N}, \text{st} \frac{z+z'+xy'}{N^{10}} \right),$$

but the nonstandard real $xy'/N^{10} = O(N^2/N^{10})$ is infinitesimal, and so the previous expression is equal to

$$\left(\text{st} \frac{x+x'}{N}, \text{st} \frac{y+y'}{N}, \text{st} \frac{z+z'}{N^{10}} \right)$$

which establishes the homomorphic nature of π .

Here we note that the homomorphism $\pi : \mathbf{A}^8 \rightarrow \mathbf{R}^3$ is not associated to any exact homomorphisms π_n from \mathbf{A}_n^8 to \mathbf{R}^3 . Instead, it is only associated to *approximate* homomorphisms

$$\pi_n \left(\begin{pmatrix} 1 & x_n & z_n \\ 0 & 1 & y_n \\ 0 & 0 & 1 \end{pmatrix} \right) := \left(\frac{x_n}{n}, \frac{y_n}{n}, \frac{z_n}{n^{10}} \right)$$

into \mathbf{R}^3 . Such approximate homomorphisms are somewhat less pleasant to work with than genuine homomorphisms; one of the main reasons why we work in the ultraproduct setting is so that we can use genuine group homomorphisms, or at least local group homomorphisms, throughout the paper.

Note that the preceding example (3.2) admits a homomorphism $\tilde{\pi}$ onto the abelian group \mathbf{R}^2 by composing the map (3.3) with the natural map from $\begin{pmatrix} \mathbf{R} & \mathbf{R} \\ 0 & \mathbf{R} \\ 0 & 0 & 1 \end{pmatrix}$ to its abelianisation \mathbf{R}^2 . However the kernel of $\tilde{\pi}$ is, for us, too “big”. In particular it contains every $\begin{pmatrix} 1 & 0 & z \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$, and in particular contains elements of \mathbf{A}^8 not in \mathbf{A} . By contrast there are no such elements in the example (3.5). In particular, we can still use the model π to establish the same escape property for \mathbf{A} as before, namely that whenever $a, \dots, a^{100} \in \mathbf{A}^{10}$, one has $a \in \mathbf{A}$.

We also note the sets \mathbf{A}_n^m for small m and large n grow cubically in m in this example, and quartically in m in the previous example. This is consistent with the model groups having homogeneous dimension 3 in the current example and 4 in the previous example.

In all the above examples, the model group \mathbf{L} was a Lie group. We now give some examples to show that the model need not *initially* be of Lie type, but can then be replaced with a Lie model after some modification.

Example 19 (Nonstandard cyclic group, revisited). — The first example is the nonstandard cyclic group $\mathbf{A} := \mathbf{Z}/2^{\mathbf{N}}\mathbf{Z} = \prod_{n \rightarrow \alpha} \mathbf{Z}/2^n\mathbf{Z}$. This is a nonstandard finite group and can thus be modeled by the trivial group $\{\text{id}\}$ as discussed in Example 14. However, it can also be modeled by the compact abelian group \mathbf{Z}_2 of 2-adic integers using the model $\pi : \mathbf{A} \rightarrow \mathbf{Z}_2$ defined by the formula

$$\pi(a) := \lim_{n \rightarrow \infty} a(\text{mod } 2^n)$$

where for each standard natural number n , $a(\text{mod } 2^n) \in \{0, \dots, 2^{n-1}\}$ is the remainder of a modulo 2^n (this is well-defined in \mathbf{A}) and the limit is in the 2-adic metric. Note that the image $\pi(\mathbf{A})$ of \mathbf{A} is the entire group \mathbf{Z}_2 , and conversely the preimage of \mathbf{Z}_2 in $\mathbf{A}^8 = \mathbf{A}$ is trivially all of $\mathbf{Z}/2^{\mathbf{N}}\mathbf{Z}$; as such, one can quotient out \mathbf{Z}_2 in this model and recover the trivial model of \mathbf{A} .

Example 20 (Nonstandard abelian 2-torsion group). — In a similar spirit to the preceding example, the nonstandard 2-torsion group $\mathbf{A} := (\mathbf{Z}/2\mathbf{Z})^{\mathbf{N}} = \prod_{n \rightarrow \alpha} (\mathbf{Z}/2\mathbf{Z})^n$ can be modeled by the compact abelian group $(\mathbf{Z}/2\mathbf{Z})^{\mathbf{N}}$ by the formula

$$\pi(a) := \lim_{n \rightarrow \infty} \pi_n(a)$$

where $\pi_n : \mathbf{A} \rightarrow (\mathbf{Z}/2\mathbf{Z})^n$ is the obvious projection, and the limit is in the product topology of $(\mathbf{Z}/2\mathbf{Z})^{\mathbf{N}}$. As before, we can quotient out $(\mathbf{Z}/2\mathbf{Z})^{\mathbf{N}}$ and model \mathbf{A} instead by the trivial group.

Remark 3.4. — The above two examples can be generalised to model any nonstandard finite group $G = \prod_{n \rightarrow \alpha} G_n$ equipped with surjective homomorphisms from G_{n+1} to G_n by the inverse limit of the G_n .

Example 21 (Lamplighter group). — Let F_2 be the field of two elements. Let G be the lamplighter group $\mathbf{Z} \ltimes F_2^{\mathbf{Z}}$, where \mathbf{Z} acts on $F_2^{\mathbf{Z}}$ by the shift $T : F_2^{\mathbf{Z}}$ defined by $T(a_n)_{n \in \mathbf{Z}} := (a_{n-1})_{n \in \mathbf{Z}}$. Thus the group law in G is given by

$$(i, x)(j, y) := (i + j, x + T^i y).$$

For each n , we then set $A_n \subseteq G$ to be the set

$$A_n := \{(i, x) \in G : i \in \{-1, 0, +1\}; x \in F_2^n\},$$

where we identify F_2^n with the space of elements $(a_n)_{n \in \mathbf{Z}}$ of $F_2^{\mathbf{Z}}$ such that $a_n \neq 0$ only for $n \in \{1, \dots, n\}$. These sets A_n are not quite approximate groups because they are not symmetric, but they are close enough to approximate groups for this discussion. For instance, they have bounded doubling or bounded tripling, and $A_n \cup A_n^{-1}$ is an approximate group. We model the ultraproduct $\mathbf{A} := \prod_{n \rightarrow \alpha} A_n \subset \mathbf{Z} \ltimes {}^*F_2^{\mathbf{Z}}$ by the group

$$G_0 \times_{\mathbf{Z}} G_0 := \{((i, x), (j, y)) \in G_0 \times G_0 : i = j\},$$

where G_0 is the modified lamplighter group $\mathbf{Z} \ltimes F_2((t))$, where $F_2((t))$ is the ring of formal Laurent series $\sum_{n \in \mathbf{Z}} a_n t^n$ over F_2 with only finitely many non-zero a_n for n negative, and the shift given by the multiplication map $T : f \mapsto tf$. We give $F_2((t))$ (and hence G_0 and $G_0 \times_{\mathbf{Z}} G_0$) a topology by declaring the norm of a non-zero element $\sum_{n \in \mathbf{Z}} a_n t^n$ of $F_2((t))$ to be 2^{-n} , where n is the least integer for which a_n is non-zero. The model map $\pi : \langle \mathbf{A} \rangle \rightarrow G_0 \times_{\mathbf{Z}} G_0$ is then given by the formula

$$\pi\left(\left(i, \lim_{n \rightarrow \alpha} (a_n^{(n)})_{n \in \mathbf{Z}}\right)\right) := \left(\left(i, \sum_{n \in \mathbf{Z}} \lim_{n \rightarrow \alpha} a_n^{(n)} t^n\right), \left(i, \sum_{n \in \mathbf{Z}} \lim_{n \rightarrow \alpha} a_{n-n}^{(n)} t^n\right)\right).$$

Roughly speaking, $\pi(a)$ captures the behaviour of a at the two “ends” of $F_2^{\mathbf{N}}$. The image $\pi(\mathbf{A})$ of \mathbf{A} under this model is then the compact neighbourhood of the identity

$$\pi(\mathbf{A}) = \{((i, x), (i, y)) \in G_0 : i \in \{-1, 0, +1\}, x, y \in F_2[[t]]\}$$

where $F_2[[t]] \subset F_2((t))$ is the ring of formal power series $\sum_{n=0}^{\infty} a_n t^n$ over F_2 . One can also compute the images $\pi(\mathbf{A}^m)$ for larger values of m , although they are a bit more complicated. One can verify the escape property that if $g, g^2, \dots, g^{100} \in \pi(\mathbf{A}^{10})$ for some $g \in G_0 \times_{\mathbf{Z}} G_0$, then $g \in \pi(\mathbf{A})$; here it is essential that we use both of the two factors of $G_0 \times_{\mathbf{Z}} G_0$, as the claim is false if we project π to just one of the two factors G , or to the base group \mathbf{Z} . So, in this case, one needs a moderately complicated (though still locally compact) group $G_0 \times_{\mathbf{Z}} G_0$ to properly¹² model \mathbf{A} and its powers \mathbf{A}^m . However, if we pass to the large subset \mathbf{A}' of \mathbf{A} defined by $\mathbf{A}' := \prod_{n \rightarrow \alpha} A'_n$, where

$$A'_n := \{(i, x) \in G : i = 0; x \in F_2^n\}$$

then \mathbf{A}' is now a nonstandard finite group (isomorphic to the group $F_2^{\mathbb{N}}$ considered in Example 20) and can be modeled simply by the trivial group $\{\text{id}\}$. Thus we see that we can sometimes greatly simplify the modeling of an ultra approximate group by passing to a large ultra approximate subgroup.

Let us formalise the properties enjoyed by the above examples in the following definition, which will play a key role in this paper.

Definition 3.5 (Good models). — *Let \mathbf{A} be an ultra approximate group. A good model for \mathbf{A} is a symmetric local topological group \mathbf{L} (see Definition B.1), together with a homomorphism $\pi : \mathbf{A}^8 \rightarrow \mathbf{L}$ of local groups with the following properties:*

- (i) (Thick image) *There exists an open neighbourhood of the identity U_0 in \mathbf{L} such that $\pi^{-1}(U_0) \subseteq \mathbf{A}$ and $U_0 \subseteq \pi(\mathbf{A})$. In particular $\ker \pi \subseteq \mathbf{A}$;*
- (ii) (Compact image) *$\pi(\mathbf{A})$ is contained in a compact set.*
- (iii) (Approximation by “internal” sets) *Suppose that $F \subseteq U \subseteq U_0$, where F is compact and U is open. Then there is an ultraproduct $\mathbf{A}' = \prod_{n \rightarrow \alpha} A'_n$ of finite sets $A'_n \subseteq A_n$ such that $\pi^{-1}(F) \subseteq \mathbf{A}' \subseteq \pi^{-1}(U)$.*

We will often abuse notation and refer to just \mathbf{L} or π as the good model for \mathbf{A} , rather than the pair (\mathbf{L}, π) .

Remark 3.6. — Properties (i) and (ii) together imply that \mathbf{L} is locally compact. We leave it to the reader to check that the examples given above have all of the properties of this definition. One can think of a good model as accurately describing the “coarse-scale” structure of the ultra approximate group \mathbf{A} , without directly controlling the “fine-scale” structure. For instance, in the example (3.5) which is “abelian at coarse scales” but “2-step nilpotent at fine scales”, the model π only detects the abelian structure and not the 2-step nilpotent structure.

¹² This can also be seen from volume growth considerations: A_n^m grows like 4^m , which is also the rate of volume growth of $\pi(\mathbf{A})$ in $G_0 \times_{\mathbf{Z}} G_0$, whereas the volume growth in a single factor G_0 would only grow like 2^m , and the volume growth in \mathbf{Z} is only linear in m .

Remark 3.7. — In (iii), if F and U are symmetric neighbourhoods of the identity, then \mathbf{A}' can be chosen to be symmetric (since one can replace \mathbf{A}' with $\mathbf{A}' \cap (\mathbf{A}')^{-1}$). As L is locally compact, we may shrink U to be precompact; then U^2 can be covered by finitely many translates of F , and thus \mathbf{A}' is then an ultra approximate group.

Finally, we need to explain the adjective “virtually” in Theorem 3.1. In group theory, “virtually” means “after passing to a finite index subgroup”. Note that a subgroup G' of a group G has finite index if and only if G can be covered by finitely many left-translates—or, equivalently, right-translates—of G' . This motivates the following definition.

Definition 3.8 (Large approximate subgroups). — *Let \mathbf{A}, \mathbf{A}' be ultra approximate groups. We say that \mathbf{A}' is a large ultra approximate subgroup of \mathbf{A} if one has $(\mathbf{A}')^4 \subset \mathbf{A}^4$, and \mathbf{A} can be covered by finitely many left-translates of \mathbf{A}' (by elements of $\mathbf{A} \cdot (\mathbf{A}')^{-1}$, of course).*

Remark 3.9. — It would be more aesthetically pleasing to have $\mathbf{A}' \subset \mathbf{A}$ instead of $(\mathbf{A}')^4 \subset \mathbf{A}^4$, but we need the exponent 4 in the inclusion for some minor technical reasons. Note that the property of being a large ultra approximate subgroup is transitive.

We are now in a position to state Hrushovski’s Lie Model Theorem.

Theorem 3.10 (Hrushovski Lie Model Theorem). — *Let \mathbf{A} be an ultra approximate group. Then there is a large ultra approximate subgroup \mathbf{A}' of \mathbf{A} such that \mathbf{A}' admits a local Lie group as a good model.*

We will prove this theorem in Section 6. As stated above, the basic idea of the proof is to first establish that \mathbf{A} itself admits a locally compact local group as a good model. Here results of multiplicative combinatorics, and in particular a lemma of Sanders [47] (see also [13]), are critical. Once this is done, Theorem 3.10 follows relatively quickly from the deep results in the literature on Hilbert’s fifth problem. This theorem will then play a key role in the proof of Theorem 2.10 in two ways: firstly by allowing us to establish certain “escape” properties on (ultra) approximate groups that will be used to build useful metric structures on these groups; and secondly by giving a natural notion of the “dimension” of an (ultra) approximate group which we will need to induct on. Note that one can invoke Lie’s third theorem (Theorem B.16) to upgrade the local Lie group in Theorem 3.10 to a connected, simply connected, global Lie group, but for technical inductive reasons it will be more convenient to keep the model in the category of local Lie groups for now.

Theorem 3.10 will be proven in Section 6. We will also establish a “global” variant of this theorem later, first in a weak form as Proposition 6.12 and then in a stronger form as Theorem 10.10.

4. An outline of the argument

In the previous section we introduced the notion of a (Hrushovski) Lie model, one of the key technical tools we will use to prove Theorem 2.10. In this section we outline the argument for this proof as a whole.

Our aim is to show that every \mathbf{K} -approximate group is controlled in some sense by a coset nilprogression of rank and density $O_{\mathbf{K}}(1)$. We shall prove this by contradiction, assuming that there is a sequence $(A_n)_{n \in \mathbf{N}}$ of \mathbf{K} -approximate groups for which the statement fails in the limit for any given choice of implied constant in the $O_{\mathbf{K}}(1)$ notation. In particular, the cardinality $|A_n|$ will go to infinity as $n \rightarrow \infty$. We assemble these approximate groups into an ultra approximate group $\mathbf{A} := \prod_{n \rightarrow \alpha} A_n$. Our assumption implies that \mathbf{A} is not “controlled” in a certain sense by what we call an ultra coset nilprogression, which we now define.

Definition 4.1 (Ultra coset nilprogression). — An ultra coset nilprogression is an ultraproduct $\mathbf{P} = \prod_{n \rightarrow \alpha} \mathbf{P}_n$ of coset nilprogressions $\mathbf{P}_n = \mathbf{P}(u_{1,n}, \dots, u_{r,n}; N_{1,n}, \dots, N_{r,n})$ of fixed (standard) rank r and step s . We then say that \mathbf{P} has rank r and step s . If the \mathbf{P}_n are also all in \mathbf{C} -normal form for some (standard) \mathbf{C} independent of n , we say that the ultra coset nilprogression is in normal form. We call $N_i := \lim_{n \rightarrow \alpha} N_{i,n}$ for $i = 1, \dots, r$ the lengths of the ultra coset nilprogression, and say that the nilprogression is nondegenerate if all the N_i are unbounded.

We define the concept of an ultra nilprogression similarly, but replacing “coset nilprogression” by “nilprogression” throughout.

As with all ultraproducts, it suffices to have the \mathbf{P}_n obey the stated properties for all n sufficiently close to α , as one can redefine \mathbf{P}_n arbitrarily on the remaining values of n without affecting the ultraproduct \mathbf{P} . Note that an ultra nilprogression \mathbf{P} can be expressed as

$$\mathbf{P} = \mathbf{P}(u_1, \dots, u_r; N_1, \dots, N_r)$$

where r is the rank, u_1, \dots, u_r are elements of the ambient nonstandard local group, and N_1, \dots, N_r are nonstandard positive reals.

To obtain the contradiction, then, it is sufficient to establish the following ultraproduct version of our main theorem.

Theorem 4.2. — Suppose that \mathbf{A} is an ultra approximate group. Then \mathbf{A}^4 contains a nondegenerate ultra coset nilprogression \mathbf{P} in normal form with $|\mathbf{P}| \gg |\mathbf{A}|$.

Here $|\mathbf{P}| \gg |\mathbf{A}|$ means that the non-standard numbers $|\mathbf{A}|$ and $|\mathbf{P}|$ satisfy $|\mathbf{A}| = O(|\mathbf{P}|)$, or in other words that there is a (standard) number $C > 0$ such that $|A_n| \leq C|P_n|$ for an α -large set of $n \in \mathbf{N}$. See the end of Appendix A for more information.

The Hrushovski Lie model theorem, Theorem 3.10, will be a key tool in establishing this, as we discuss below. In addition to this theorem, a further fundamental additional concept in our argument will be the notion of an *escape norm*.

Definition 4.3 (*Escape norm*). — Let \mathbf{A} be a multiplicative set. For a group element $g \in \mathbf{A}^{10}$, we define the escape norm $\|g\|_{e,\mathbf{A}} \in [0, 1]$ to be the quantity

$$\|g\|_{e,\mathbf{A}} := \inf \left\{ \frac{1}{n+1} : n \in \mathbf{N}; g^i \in \mathbf{A} \text{ for all } 0 \leq i \leq n \right\}.$$

Recall that by convention, the statement $g^i \in \mathbf{A}$ is false if g^i is not well-defined. Now suppose that \mathbf{A} is a nonstandard multiplicative set, i.e. an ultraproduct $\mathbf{A} = \prod_{n \rightarrow \alpha} \mathbf{A}_n$ of standard multiplicative sets \mathbf{A}_n . If $g = \lim_{n \rightarrow \alpha} g_n$ is an element of \mathbf{A}^{10} , we define the escape norm $\|g\|_{e,\mathbf{A}} \in [0, 1]$ to be the quantity

$$\|g\|_{e,\mathbf{A}} := \lim_{n \rightarrow \alpha} \|g_n\|_{e,\mathbf{A}_n}.$$

The escape norm can always be defined, but there are some remarkable lemmas essentially due to Gleason [21] concerning its properties when \mathbf{A} is an approximate group. Specifically we will show in Section 8 that there is a set A' controlling \mathbf{A} for which the escape norms satisfy (precise versions of) the following estimates:

- (i) (Product property) If $g_1, \dots, g_n \in A'$ then $\|g_1 \dots g_n\|_{e,A'} \ll \|g_1\|_{e,A'} + \dots + \|g_n\|_{e,A'}$;
- (ii) (Conjugation property) If $g, h \in A'$ then $\|h^{-1}gh\|_{e,A'} \ll \|g\|_{e,A'}$;
- (iii) (Commutator property) If $g, h \in A'$ then $\|[g, h]\|_{e,A'} \ll \|g\|_{e,A'} \|h\|_{e,A'}$.

These estimates, which we shall informally term “Gleason’s lemmas”, will be proven in Section 8. They are valid in both the finitary and the ultralimit settings; the latter will be deduced, quite straightforwardly, from the former.

The remarks in the following paragraph pertain to the finitary situation. To prove the Gleason lemmas, the set A' must be what we call a *strong approximate group*. The precise definition of this is Definition 7.1. It is by no means obvious that there is a large strong approximate group A' contained in A^4 , but this will follow from the Hrushovski Lie model theorem (Theorem 3.10), basically because small balls in a Lie group are automatically strong approximate groups, and can then be pulled back by the model map.

Once A' has been defined, Gleason’s lemmas are proven by an argument closely analogous to that of Gleason himself [21]. We shall say nothing further about the details here; the argument is self-contained and is discussed in Section 8.

With Gleason’s lemmas in hand, let us describe the rest of the argument.

Firstly, the set $H = \{g : \|g\|_{e,A'} = 0\}$ of elements which do not escape is a normal (genuine) subgroup of A' ; this follows from (i) and (ii). We may quotient by H to get an ultra approximate group $A_0 := A'/H$, all of whose non-identity elements have nonzero

escape norm. We shall call such approximate groups *NSS approximate groups*, in analogy with the *no small subgroups* property in the theory of locally compact groups.

Now, if $g_1 \in A_0^\dagger$ is an element other than the identity with smallest (nonzero) $\|\cdot\|_{e,A_0}$ -escape norm then we shall see that in fact, if A' is chosen appropriately, $g_1 \in A_0$. Item (iii) then implies that for any $h \in A_0$, $[g_1, h] \in A_0^\dagger$ has smaller escape norm than g_1 , and hence must be the identity. In other words, g_1 is central in A_0 and we may quotient again to get a new approximate group $A_1 := A_0/\langle g_1 \rangle$. We are being quite fuzzy at this point; in fact, the quotienting takes place in the category of *local* groups and one is quotienting not by the entire group $\langle g_1 \rangle$ but by an appropriate geometric progression within it.

Continuing in this vein we pick $g_2 \in A_1^\dagger$ other than the identity with smallest $\|\cdot\|_{e,A_1}$ -norm. We shall see that this norm is automatically nonzero, a consequence of the *local* nature of the quotienting operation.

Continuing further, we pick g_3, g_4, \dots

All of this makes sense at the level of ultralimits as well, and in this setting one can show that \mathbf{A}_i has a Hrushovski Lie model L_i with $\dim L_i < \dim L_{i-1}$ for all i . Because of this, the quotienting procedure terminates in finite time with an element g_k and one concludes by reversing these finitely many quotienting operations that \mathbf{A} is controlled by an ultra coset nilprogression with “generators” H, g_1, \dots, g_k , thereby leading to a proof of Theorem 4.2.

This concludes our brief summary of the argument. Let us summarise the content of the remaining core sections of the paper.

- In Section 5, we discuss results from multiplicative combinatorics, essentially due to Sanders and Croot-Sisask, which are relevant to the proof of Hrushovski’s Lie model theorem.
- In Section 6, we prove the Hrushovski Lie model theorem.
- In Section 7, we use the Hrushovski Lie model theorem to construct strong approximate groups.
- In Section 8, we state and prove Gleason’s lemmas.
- In Section 9, we give details of the inductive strategy outlined above for constructing H and g_1, \dots, g_k , and conclude the proof of Theorem 2.10 (except for the rank bound).
- In Section 10 we show that the rank and step of the coset nilprogression can be bounded by $6 \log_2 K$ in the global case.
- Section 11 is devoted to various applications to the growth of groups and to Riemannian geometry. We prove there the corollaries stated in the introduction.

5. Sanders-Croot-Sisask theory

In the next section we will establish Hrushovski’s Lie Model Theorem (Theorem 3.10), in which an ultra approximate group is related first to a locally compact

metrisable local group and then, via Goldbring’s solution [24] of the local Hilbert’s Fifth problem, to a local Lie group. In locally compact metrisable local groups we have total boundedness, which means that the unit (say) ball $B(\text{id}, 1) := \{x \in G : d(x, \text{id}) \leq 1\}$ may be covered by $O_\varepsilon(1)$ smaller balls $B(x_i, \varepsilon) := \{x \in G : d(x, x_i) \leq \varepsilon\}$. On the other hand, by continuity of the group operation, $B(\text{id}, 1)$ will contain high powers like $B(\text{id}, \varepsilon)^{100}$ for suitably small ε .

It is not surprising, then, that we need tools for showing (roughly speaking) that approximate groups A contain high powers of somewhat smaller, but still quite large, approximate subgroups A' , which do not immediately escape A in the sense that $(A')^m$ is contained inside A (or perhaps a slightly larger set such as A^4) for a reasonably large value of m . Such a tool is provided by a result from multiplicative combinatorics due to Sanders [47] and to Croot-Sisask [13, Theorem 1.6], namely Theorem 5.3 below. We shall also need a “normal” variant of this result, which essentially follows by combining Theorem 5.3 with [49, Lemma 13.1]. Our version of this is Theorem 5.6 below, and once again we provide a self-contained proof.

Let us remark that by appealing to these results from multiplicative combinatorics we differ fairly substantially from the approach taken by Hrushovski [33], although one may perceive structural similarities in the model-theoretic arguments he uses.

All of the results below are essentially already in the literature, but always for subsets A of some ambient (global) group G . As it turns out, though, the proofs of these results end up being equally valid for the more local setting of multiplicative sets. Indeed, most of the tools used in multiplicative combinatorics (with the notable exception of the Fourier transform) are already “local” in nature in that they only require one to do $O(1)$ multiplications.

Our first such tool is Ruzsa’s covering lemma, which essentially allows one to select a “complete set of coset representatives” in the approximate group setting.

Lemma 5.1 (*Local Ruzsa covering lemma*). — *Let A, B be finite sets, and suppose that $A \cup B$ is a multiplicative set. Then there exists a finite set $X \subseteq B$ with $|X| \leq |AB|/|A|$ and $B \subseteq A^2X$. Similarly there exists a finite set $Y \subseteq B$ such that $|Y| \leq |BA|/|A|$ and $B \subseteq YA^2$.*

Proof. — Let X be a subset of B such that the sets $A \cdot x$ for $x \in X$ are disjoint, and such that X is maximal with respect to set inclusion; then we have $|X| \leq |AB|/|B|$. If $b \in B$, then $A \cdot b$ and $A \cdot x$ must intersect for some x , thus $a \cdot b = a' \cdot x$ for some $a, a' \in A$. Multiplying on the left by a^{-1} , we conclude that $b = a^{-1} \cdot a' \cdot x$, and the claim follows. \square

A corollary of this is the following result, which allows one to produce an approximate group from a set with small growth.

Corollary 5.2. — *Let A be a symmetric multiplicative set, and suppose that $|A^5| \leq K|A|$. Then A^2 is a $2K$ -approximate group.*

Proof. — Clearly A^2 is a symmetric set containing the identity. Since $|A^5| \leq K|A| \leq K|A^2|$, we see from Lemma 5.1 that there exists $X \subseteq A^4$ with $|X| \leq K$ such that $A^4 \subseteq A^2X$, and there similarly exists $Y \subseteq A^4$ with $|Y| \leq K$ such that $A^4 \subseteq YA^2$. Taking the union of X and Y we obtain the claim. \square

We turn now to the result of Sanders [47] that drives our whole approach.

Theorem 5.3 (*Small neighbourhoods*). — *Suppose that A is a K -approximate group, and let $m \geq 1$ be an integer. Then there is a $O_{K,m}(1)$ -approximate group S with $|S| \gg_{K,m} |A|$ such that $S^m \subseteq A^4$.*

Remark 5.4. — Explicit bounds for the implied constants are given in, for example, [13, Theorem 1.6]. As much of the remainder of the argument is not explicitly effective with respect to bounds, we do not worry about such quantitative issues here. Similar remarks can be made in connection with the normal variant, Theorem 5.6 below.

Proof. — We use the argument from [47], generalised to the setting of multiplicative sets. For the convenience of the reader, we reproduce it here. A somewhat different proof of Theorem 5.3 can also be obtained by using the techniques of [13].

For each $0 < t < 1$, let $f(t)$ denote the quantity

$$f(t) := \inf \left\{ \frac{|AB|}{|A|} : B \subseteq A; |B| \geq t|A| \right\}.$$

Since $|A^2| \leq K|A|$, we have $1 \leq f(t) \leq K$ for all $0 < t < 1$. By the pigeonhole principle, we can thus find $t \gg_{K,m} 1$ such that

$$(5.1) \quad f\left(\frac{t^2}{2K^2}\right) \geq \left(1 - \frac{1}{100m}\right)f(t).$$

Fix this t . As there are only finitely many sets B that make up the infimum for f , we can find a $B \subset A$ with $|B| \geq t|A|$ such that

$$(5.2) \quad |AB| = f(t)|A|.$$

For each $a \in A$, the set Ba has cardinality $|B|$ and is contained in A^2 .

$$\sum_{x \in A^2} \sum_{a \in A} 1_{Ba}(x) = |A||B|$$

and hence by Cauchy-Schwarz we obtain

$$\sum_{x \in A^2} \left(\sum_{a \in A} 1_{Ba}(x) \right)^2 = \frac{|A|^2|B|^2}{|A^2|}.$$

The left-hand side can be rewritten as

$$\sum_{a \in A} \sum_{a' \in A} |\mathbf{B}a \cap \mathbf{B}a'|,$$

and so by the pigeonhole principle, there exists $a_0 \in A$ such that

$$\sum_{a \in A} |\mathbf{B}a \cap \mathbf{B}a_0| \geq \frac{|A||\mathbf{B}|^2}{|A^2|}.$$

Since $|\mathbf{B}| \geq t|A|$ and $|A^2| \leq K|A|$, we thus have

$$\sum_{a \in A} |\mathbf{B}a \cap \mathbf{B}a_0| \geq \frac{t^2}{K^2}|A|^2,$$

and hence we can find a subset C of A of cardinality

$$(5.3) \quad |C| \geq t^2/2K^2|A|$$

such that $|\mathbf{B}a \cap \mathbf{B}a_0| \geq t^2|A|/2K^2$ for all $a \in C$. Multiplying by a_0^{-1} and by a^{-1} , we see that $|\mathbf{B}h \cap \mathbf{B}| \geq t^2|A|/2K^2$ for all $h \in S_0$, where $S_0 := a_0^{-1}C \cup C^{-1}a_0 \cup \{\text{id}\}$ is a symmetric subset in A^2 containing the identity. From (5.1), we conclude that

$$|A(\mathbf{B}h \cap \mathbf{B})| \geq \left(1 - \frac{1}{100m}\right) f(t)|A|.$$

From (5.2), we conclude that

$$|\mathbf{A}Bh \cap \mathbf{A}B| \geq \left(1 - \frac{1}{100m}\right) |\mathbf{A}B|.$$

Using induction (and the hypothesis that A^8 is well-defined, noting that $\mathbf{B} \subset A$ and $S_0 \subset A^2$) we then see that for any $1 \leq m' < 100m$, the set $S_0^{m'}$ is well-defined and

$$|\mathbf{A}Bh \cap \mathbf{A}B| \geq \left(1 - \frac{m'}{100m}\right) |\mathbf{A}B|$$

for all $h \in S_0^{m'}$, which in particular implies that $S_0^{m'} \subset A^4$. On the other hand, from (5.3) we have $|S_0| \gg_{K,m} |A|$. From Corollary 5.2 we see that $S := S_0^2$ is a $O_{K,m}(1)$ -approximate group. Since $S^m = S_0^{2m} \subset A^4$, we obtain the first claim of the lemma. The second claim follows by applying the Ruzsa covering lemma (with $\mathbf{B} := S_0$). \square

Remark 5.5. — Let us pause to note a consequence of this result. We defined multiplicative sets to be ones in which one was at liberty to take up to 100 multiplications (i.e. A^{100} is well-defined), and the associative law would hold to this extent. Theorem 5.3, or

more accurately a close examination of the proof of it, says that if A is an approximate group and a multiplicative set in which merely 8 multiplications are allowed (i.e. A^8 is well-defined) then A is $O_{m,K}(1)$ -controlled by an $O_{m,K}(1)$ -approximate group $A' = S$ in which up to m multiplications are defined an associative. For this reason Theorem 2.10 holds if only 8 multiplications are allowed. We shall not dwell on such details further in this paper, allowing ourselves the luxury of 100 multiplications.

We turn now to proving a “normal” variant of Theorem 5.3. Here, we use the notation

$$a^b := b^{-1}ab$$

and

$$A^B := \{a^b : a \in A, b \in B\}$$

for elements a, b and subsets A, B of a local group.

Theorem 5.6 (*Small normal neighbourhoods*). — *Suppose that A is a K -approximate group, and let $m \geq 1$ be an integer. Let $S \subseteq A^4$ be a K' -approximate group with $|S| = \delta|A|$. Then there is an $O_{m,K,K',\delta}(1)$ -approximate group \tilde{S} with $|\tilde{S}| \gg_{K,K',m,\delta} |A|$ such that $(\tilde{S}^m)^{A^4} \subseteq S^4$.*

Theorem 5.6 will be deduced from Theorem 5.3. To motivate the argument, let us first recall a standard lemma from group theory.

Lemma 5.7. — *Let A be a finite group, and let S be a subgroup of A with $|S| \geq |A|/K$. Then there exists a further subgroup $\tilde{S} \subset S$ of A with $|\tilde{S}| \gg_K |A|$ which is normal in A .*

Note that this lemma would easily yield Theorem 5.6 from Theorem 5.3 in the special case when A and \tilde{S} are genuine groups and not merely approximate groups.

Proof. — Let x_1, \dots, x_k be a complete set of right coset representatives for S in A , and set

$$\tilde{S} = \bigcap x_i^{-1}Sx_i = \bigcap_{x \in A} x^{-1}Sx.$$

All the claims of the lemma are immediate, except for the claim that $|\tilde{S}| \gg_K |A|$. However, this follows from iterating the fact that if $H_1, H_2 \leq G$ are subgroups of small index in a group G then so is $H_1 \cap H_2$; in fact we have the well-known inequality

$$(5.4) \quad [G : H_1 \cap H_2] \leq [G : H_1][G : H_2]. \quad \square$$

To adapt this argument to the approximate setting we need an analogue of (5.4) for approximate groups. This is provided by the following lemma.

Lemma 5.8. — *Suppose that A is a K -approximate group and that $A_1, A_2 \subseteq A$ are sets with $|A_i| = \delta_i |A|$. Then $A_1 A_1^{-1} \cap A_2 A_2^{-1}$ contains a set BB^{-1} with $B \subseteq A$ and $|B| \geq \delta_1 \delta_2 |A|/K$.*

Proof. — Since $A_1^{-1} A_2 \subseteq A^2$, we have $|A_1^{-1} A_2| \leq K|A|$. It follows that there is some x with at least $\delta_1 \delta_2 |A|/K$ representations as $a_1^{-1} a_2$. Let B be the set of all values of a_2 that appear. Obviously $BB^{-1} \subseteq A_2 A_2^{-1}$. Suppose that $a_2, a'_2 \in B$. Then there are a_1, a'_1 such that $x = a_1^{-1} a_2 = (a'_1)^{-1} a'_2$, and so $a'_1 a_1^{-1} = a'_2 a_2^{-1}$. Thus BB^{-1} lies in $A_1 A_1^{-1}$ as well. \square

By iterating the above lemma we obtain the following corollary.

Corollary 5.9. — *Suppose that A is a K -approximate group and that $A_1, \dots, A_k \subseteq A$ are sets with $|A_i| \geq \delta |A|$ for each i . Then $|\bigcap_{i=1}^k A_i A_i^{-1}| \gg_{\delta, k, K} |A|$.*

Now we can prove Theorem 5.6.

Proof of Theorem 5.6. — By Theorem 5.3, there is an $O_{l, K, K'}(1)$ -approximate subgroup $S_0 \subseteq S^4$,

$$|S_0| \gg_{m, K, K', \delta} |A|,$$

such that

$$(5.5) \quad S_0^{4m+4} \subseteq S^4.$$

The Ruzsa covering lemma allows us to do the analogue of picking a complete set of coset representatives in the approximate group setting. Specifically, there are x_1, \dots, x_k , $k = O_{m, \delta, K}(1)$, such that

$$(5.6) \quad A^4 \subseteq \bigcup_{i=1}^k S_0^2 x_i.$$

Let us assume without loss of generality that $x_1 = \text{id}$.

By Corollary 5.9, the set

$$T := \bigcap_{i=1}^k x_i S_0^2 x_i^{-1}$$

has cardinality

$$|T| \gg_{m, K, K', \delta} |A|.$$

We claim that the set $\tilde{S} := T^2$ has the required properties. First of all note that, by Corollary 5.2, \tilde{S} is indeed an $O_{m, K, K', \delta}(1)$ -approximate group.

Next observe that, since $x_1 = \text{id}$,

$$(5.7) \quad x_i^{-1}Tx_i \subseteq S_0^2$$

for each i .

Suppose that $x \in A^4$. Then, by (5.6), there is some i with $1 \leq i \leq k$ and some $s \in S_0^2$ such that $x = sx_i$. It follows from this, (5.7) and (5.5) that

$$x^{-1}\tilde{S}^m x = x^{-1}T^{2m}x = s^{-1}x_i^{-1}T^{2m}x_i s = s^{-1}(x_i^{-1}Tx_i)^{2m}s \subseteq S_0^{4m+4} \subseteq S^4.$$

This concludes the proof. \square

6. Proof of the Hrushovski Lie model theorem

In this section we establish Theorem 3.10. The reader may wish to reread Section 3, which gave an overview of this theorem. We will deduce this theorem from the following two propositions.

Proposition 6.1 (Locally compact model). — *Let \mathbf{A} be an ultra approximate group. Then \mathbf{A}^4 admits a model $\pi : \mathbf{A}^{32} \rightarrow G$ by a metrisable locally compact local group G .*

Proposition 6.2 (From locally compact models to Lie models). — *Let \mathbf{A} be an ultra approximate group and suppose that \mathbf{A}^4 admits a model $\pi : \mathbf{A}^{32} \rightarrow G$ into a locally compact local group G . Then there is a large ultra approximate group $\tilde{\mathbf{A}}$ of \mathbf{A} (thus $\tilde{\mathbf{A}}^4 \subset \mathbf{A}^4$) which admits a model $\tilde{\pi} : \tilde{\mathbf{A}}^8 \rightarrow L$ into a connected, simply-connected Lie group L .*

It is clear that the above two propositions together imply Theorem 3.10.

We will give a self-contained proof of Proposition 6.1, using the multiplicative combinatorics results of the previous section, together with the countable saturation property of ultraproducts. In contrast, the proof of Proposition 6.2 requires deep material related to (the local version of) Hilbert's fifth problem, for which we provide suitable references.

Building metrics on local groups. — We now begin the proof of Proposition 6.1. Suppose that we have a pseudometric $d : G \times G \rightarrow [0, \infty)$ on some local group G , that is to say d satisfies the axioms of a metric, except that we may have $d(x, y) = 0$ when $x \neq y$. Then we may of course define the balls $B(\text{id}, \varepsilon) := \{x \in G : d(x, \text{id}) < \varepsilon\}$, and these will be nested in the sense that $B(\text{id}, \varepsilon) \subseteq B(\text{id}, \varepsilon')$ if $\varepsilon < \varepsilon'$. We now examine ways to reverse this construction, beginning with a quite general way to construct pseudometrics on symmetric local groups; this will be needed to prove Proposition 6.1.

Let G be a symmetric local group. For any function $\psi : G \rightarrow \mathbf{R}$ and $g \in G$, we define the shift $T_g\psi : G \rightarrow \mathbf{R}$ by setting

$$T_g\psi(x) := \psi(g^{-1}x)$$

if $g^{-1}x$ is well-defined in G , and $T_g\psi(x) = 0$ otherwise. We then define the “derivative” operator

$$\partial_g\psi := \psi - T_g\psi.$$

The expression

$$\|\partial_g\psi\|_{\ell^\infty(G)} := \sup_{x \in G} |\partial_g\psi(x)|$$

can be viewed heuristically as a “norm” of g relative to ψ , and this makes it natural to consider the function

$$(6.1) \quad d(g, h) := \|T_g\psi - T_h\psi\|_{\ell^\infty(G)} = \|\partial_{h^{-1}g}\psi\|_{\ell^\infty(G)}.$$

One can view d as the pullback of the metric on $\ell^\infty(G)$ to G using the translation action $g \mapsto T_g\psi$ of G on ψ .

Lemma 6.3 (Using functions to build (pseudo-)metrics). — *Let G be a local group, and let A be a symmetric neighbourhood of the identity such that A^{128} is well-defined in G . Let $\psi : G \rightarrow \mathbf{R}$ be non-negative and supported on A .*

- (i) *We have $\|\partial_g\psi\|_{\ell^\infty(G)} \leq \|\psi\|_{\ell^\infty(G)}$ for all $g \in A^{128}$, with equality holding when $g \notin A^2$.*
- (ii) *Whenever $g, h \in A^{128}$, one has*

$$(6.2) \quad \|\partial_{gh}\psi\|_{\ell^\infty(G)} \leq \|\partial_g\psi\|_{\ell^\infty(G)} + \|\partial_h\psi\|_{\ell^\infty(G)}.$$

- (iii) *For any $g \in A^{128}$, we have*

$$(6.3) \quad \|\partial_{g^{-1}}\psi\|_{\ell^\infty(G)} = \|\partial_g\psi\|_{\ell^\infty(G)}.$$

- (iv) *The function $d : A^{64} \times A^{64} \rightarrow \mathbf{R}^+$ defined by the formula (6.1) is a left-invariant pseudo-metric on A^{64} .*

Remark 6.4. — To spell out what we mean in (iv), we are asserting that $d(g, g) = 0$, that $d(g, h) = d(h, g)$, and that $d(g, k) \leq d(g, h) + d(h, k)$ for all $g, h, k \in A^{64}$. Furthermore it has the left-invariance property $d(gh, gk) = d(h, k)$ whenever $h, k \in A^{64}$, $g \in A^{128}$, and $gh, gk \in A^{64}$. Later on, when proving Gleason’s lemmas, we shall require some slightly more exotic properties of these cocycle “norms”, related to commutation and a certain “Taylor expansion”.

Proof. — The property (i) is clear from construction. For $g, h \in A^{128}$ we have the representation property $T_g T_h\psi = T_{gh}\psi$ and hence the *cocycle identity*

$$\partial_{gh}\psi = \partial_g\psi + T_g\partial_h\psi$$

which gives (6.2).

Similarly, for $g \in A^{128}$ we have the inverse identity

$$\partial_{g^{-1}} \psi = -T_{g^{-1}} \partial_g \psi$$

which gives (6.3).

The claims in (iv) follow easily from (ii) and (iii). \square

In the next lemma we give a variant of the Birkhoff-Kakutani construction [40, §1.22], in which a function ψ is constructed so that the pseudometric $d(g, h) = \|\partial_{h^{-1}g} \psi\|_{\ell^\infty(G)}$ is adapted to a given nested sequence of symmetric sets which are supposed to resemble “balls” in this pseudometric.

Lemma 6.5 (Birkhoff-Kakutani construction). — *Suppose that G is a local group and that we have a sequence of symmetric neighbourhoods A_0, A_1, \dots of the identity in G with the nesting property that $A_{i+1}^2 \subseteq A_i$ for $i = 0, 1, 2, \dots$, and with A_0^{200} well-defined. Then there is a pseudometric*

$$d : A_0^{64} \times A_0^{64} \rightarrow [0, 1]$$

such that we have the inclusions

$$(6.4) \quad \{g \in A_0^{64} : d(g, \text{id}) < 2^{-k}\} \subseteq A_k \subseteq \{g \in A_0^{64} : d(g, \text{id}) \leq 2 \cdot 2^{-k}\}$$

for all nonnegative integers k . In particular $x_n \rightarrow x$ in the pseudometric d if and only if, for each $k \in \mathbf{N}$, we have $x^{-1}x_n \in A_k$ for all sufficiently large n .

Proof. — Suppose that $q = 2^{-i_1} + \dots + 2^{-i_k}$, $0 < q < 1$, is a dyadic rational, and define

$$B_q := A_{i_k} A_{i_{k-1}} \dots A_{i_1}.$$

Even though the definition uses a potentially large number k of multiplications, the nesting property of the A_i means that these sets B_q are well-defined in the local group G .

We claim that $B_q \subseteq B_{q+2^{-k}}$ whenever q is a dyadic rational with denominator dividing 2^k ; this easily implies that

$$(6.5) \quad B_q \subseteq B_{q'} \quad \text{whenever } 0 < q < q' < 1.$$

The claim follows by repeated use of the nesting $A_{i+1}^2 \subseteq A_i$ (the number of times it will be required is the number of carries when 2^{-k} is added to q in binary).

In particular, $B_q \subseteq A_{i_1-1} \subset A_0$.

Define $\psi : A_0^{64} \rightarrow [0, 1]$ by

$$\psi(x) := \sup\{1 - q : 0 < q < 1; x \in B_q\} \cup \{0\},$$

and consider the pseudometric $d(g, h) := \|\partial_{h^{-1}g}\psi\|_{\ell^\infty(G)}$ as discussed in Lemma 6.3. Note that for $g, h \in A_0^{64}$, $\partial_{h^{-1}g}\psi$ is supported in A_0^{192} , and so one can replace $\ell^\infty(G)$ here with A_0^{192} if desired.

If $d(g, \text{id}) < 2^{-k}$ then $|\partial_g\psi(\text{id})| < 2^k$, which implies that $\psi(g) > 1 - 2^{-k}$ and therefore $g \in B_{2^{-k}}$ and hence $g \in A_k$.

Conversely, suppose that $g \in A_k$: we are to show that $d(g, \text{id}) \leq 2 \cdot 2^{-k}$. To show this we must confirm that $|\partial_g\psi(h)| < 2^{1-k}$ for all $h \in G$. As discussed before, we may assume that $h \in A_0^{192}$. Suppose that $h \in B_q$, where $0 < q < 1 - 2^{-k}$ is an integer multiple of 2^{-k} , but that $h \notin B_{q-2^{-k}}$. Then $\psi(h) \leq 1 - q + 2^{-k}$. On the other hand, $g^{-1}h \in A_k B_q \subseteq B_{q+2^{-k}}$, by the claim established above, and therefore $\psi(g^{-1}h) \geq 1 - q - 2^{-k}$. It follows that $\partial_g\psi(h) = \psi(g^{-1}h) - \psi(h) \geq -2 \cdot 2^{-k}$. Similarly, $\partial_g\psi(h) \leq 2 \cdot 2^{-k}$. Since h was arbitrary it follows that $d(g, \text{id}) = \|\partial_g\psi\|_{\ell^\infty(G)} \leq 2 \cdot 2^{-k}$, and the claim follows. \square

If the sets A_i satisfy a certain normality condition, the group operations are continuous with respect to the pseudometric d :

Lemma 6.6 (Normal Birkhoff-Kakutani construction). — *Suppose that G is a local group and that we have a sequence of symmetric sets A_0, A_1, \dots in G with A_0^{200} well-defined and with the nesting property that $(A_{i+1}^2)^{A_0^{100}} \subseteq A_i$ for $i = 0, 1, \dots$ (and so, in particular, we certainly have the weaker nesting property $A_{i+1}^2 \subseteq A_i$ required by the preceding lemma). Consider the pseudometric $d : A_0^{64} \times A_0^{64} \rightarrow [0, 1]$ defined in the preceding lemma. Then the product map $\cdot : A_0^{32} \times A_0^{32} \rightarrow A_0^{64}$ and the inversion map $^{-1} : A_0^{32} \rightarrow A_0^{32}$ are both continuous with respect to d .*

Proof. — Suppose that $g_n \rightarrow g$ and that $h_n \rightarrow h$. We wish to show that $g_n h_n \rightarrow gh$, to which end it suffices to establish that $(gh)^{-1} g_n h_n \in A_k$ for all sufficiently large n . However, for n sufficiently large in terms of k we have $g^{-1} g_n \in A_{k+2}$, and hence

$$h_n^{-1} g^{-1} g_n h_n \in A_{k+2}^{h_n} \subseteq A_{k+2}^{A_0^{100}} \subseteq A_{k+1}.$$

Furthermore, $h^{-1} h_n \in A_{k+1}$ for n sufficiently large, and so

$$(gh)^{-1} g_n h_n = (h^{-1} h_n) (h_n^{-1} g^{-1} g_n h_n) \in A_{k+1}^2 \in A_k,$$

as required. The statement about the inverse map is easier. Suppose that $g_n \rightarrow g$. Then $g^{-1} g_n \in A_{k+1}$ for n sufficiently large, and so

$$g_n g^{-1} = g(g^{-1} g_n) g^{-1} \in A_{k+1}^{g^{-1}} \subseteq A_{k+1}^{A_0^{100}} \in A_k.$$

But this means that $g_n^{-1} \rightarrow g$ as $n \rightarrow \infty$. \square

The previous lemma showed how to get a local topological group given a sequence of balls satisfying a suitable normalisation condition. The normal variant of the Croot-Sisask-Sanders lemma, Theorem 5.6, allows us to find precisely such a sequence of balls

given any \mathbf{K} -approximate group \mathbf{A} . Of course, these balls are just finite sets and, for sufficiently large i , \mathbf{A}_i may well consist only of the identity element e . This will be the case, for example, when $\mathbf{A} = [-N, N]$. However when transferred to the setting of an ultra approximate group $\mathbf{A} = \prod_{n \rightarrow \alpha} \mathbf{A}_n$, these balls have “finite index” in \mathbf{A} , and this ultimately leads to the important conclusion that the metric d gives \mathbf{A} the structure of a locally compact local group.

Lemma 6.7. — *Let \mathbf{A} be an ultra approximate group. Then there is a sequence of ultra approximate groups $\mathbf{A}_0, \mathbf{A}_1, \dots$ such that $\mathbf{A}_0 = \mathbf{A}^4$, we have the nesting property that $(\mathbf{A}_{i+1}^2)^{\mathbf{A}_0^{100}} \subseteq \mathbf{A}_i$ for $i = 0, 1, \dots$, and each \mathbf{A}_i is large in the sense that \mathbf{A} can be covered by finitely many left-translates of \mathbf{A}_i .*

Proof. — By definition, one has $\mathbf{A} = \prod_{n \rightarrow \alpha} \mathbf{A}_n$ for some \mathbf{K} -approximate groups \mathbf{A}_n and some fixed \mathbf{K} . Applying Theorem 5.6 repeatedly we see that there are, for each n , $\mathbf{O}_{\mathbf{K},i}(1)$ -approximate groups $\mathbf{S}_{n,i}$, $i = 1, 2, 3, \dots$, such that $\mathbf{S}_{n,0} := \mathbf{A}_n$ and $(\mathbf{S}_{n,i+1}^8)^{\mathbf{S}_{n,0}^{400}} \subseteq \mathbf{S}_{n,i}^4$ for each i . Furthermore we have $\mathbf{S}_{n,i} \subseteq \mathbf{A}_n^4$ and $|\mathbf{S}_{n,i}| \gg_{\mathbf{K},i} |\mathbf{A}_n|$ for each i . Setting

$$\mathbf{A}_i := \prod_{n \rightarrow \alpha} \mathbf{S}_{n,i}^4,$$

all of the properties except the assertion about covering are immediate. To check that each \mathbf{A}_i is large, we need only check that $\mathbf{S}_{n,0}$ is covered by $\mathbf{O}_{\mathbf{K},i}(1)$ left-translates of $\mathbf{S}_{n,i}^4$, for each i . This, however, is an immediate consequence of Lemma 5.1 and the lower bound on $|\mathbf{S}_{n,i}|$. \square

Lemma 6.8. — *Let \mathbf{A} be an ultra approximate group. Consider a sequence of ultra approximate groups $\mathbf{A}_0, \mathbf{A}_1, \dots$ as found in the preceding lemma, and let $d : \mathbf{A}^{32} \times \mathbf{A}^{32} \rightarrow [0, 1]$ be the pseudometric associated to these sets as in Lemma 6.5. Then \mathbf{A}^{32} is locally compact with respect to the topology generated by d .*

Proof. — By the Heine-Borel theorem (which is usually stated for metrics, but which extends without difficulty¹³ to pseudometrics) it suffices to show that \mathbf{A}^{32} is complete and totally bounded. We deal with the latter task first. From the inclusion $\mathbf{A}_k \subseteq \{x : d(x, \text{id}) \leq 2 \cdot 2^{-k}\}$ and the left-invariance of d , this follows from the fact that \mathbf{A}^{32} is covered by finitely many left-translates of \mathbf{A}_k .

We turn now to completeness. Suppose that $(x_n)_{n \in \mathbf{N}}$ is a Cauchy sequence. By refining the sequence if necessary we may assume that it is rapidly Cauchy in the sense that $d(x_n, x_m) \leq 2^{-n-1}$.

We claim that the sets $x_n \mathbf{A}_n$ are nested in the sense that $x_m \mathbf{A}_m \subseteq x_n \mathbf{A}_n$ whenever $m > n$. To see this note that by left-invariance we have $d(\text{id}, x_n^{-1} x_m) \leq 2^{-n-1}$ and hence,

¹³ Indeed, one can deduce the pseudometric case from the metric case by quotienting out by the equivalence relation $x \sim y$ defined by the equation $d(x, y) = 0$.

by the inclusions of Lemma 6.5, $x_n^{-1}x_m \in \mathbf{A}_{n+1}$. Since $\mathbf{A}_{n+1}\mathbf{A}_m \subseteq \mathbf{A}_{n+1}^2 \subseteq \mathbf{A}_n$, it follows that $x_n^{-1}x_m\mathbf{A}_m \subseteq \mathbf{A}_n$, thereby confirming the claim.

Now each set $x_m\mathbf{A}_m$ is an ultraproduct $\prod_{n \rightarrow \alpha} S_{m,n}$, by construction. The nesting property just established of course implies that, for any positive integer M , $\bigcap_{m \leq M} x_m\mathbf{A}_m \neq \emptyset$. Let y_M be an element of this intersection; this means that there is a set $\Sigma_M \in \alpha$ such that $(y_M)_n \in \bigcap_{m \leq M} S_{m,n}$ for all $n \in \Sigma_M$. By replacing Σ_2 with $\Sigma_1 \cap \Sigma_2$ if necessary, and so on, and using the basic properties of ultrafilters, we may assume that $\Sigma_1 \supseteq \Sigma_2 \supseteq \Sigma_3 \supseteq \dots$. By removing 2 from Σ_2 , 3 from Σ_3 and so on, if necessary, we may also assume that no integer lies in infinitely many Σ_M .

Now define a sequence x by setting $x_n = (y_M)_n$, where M is the largest integer for which $n \in \Sigma_M$. Then, by construction, $x_n \in \bigcap_{m \leq M} S_{m,n}$ for all $n \in \Sigma_M$, that is to say for a set of n tending to α . This means that $x \in \bigcap_{m \leq M} x_m\mathbf{A}_m$ for every M , and hence $x \in \bigcap x_m\mathbf{A}_m$. In particular we have $x_m^{-1}x \in \mathbf{A}_m$ for every m and hence $d(x, x_m) \leq 2 \cdot 2^{-m}$. It follows that $x_m \rightarrow x$, thereby confirming that \mathbf{A} is complete with the metric d . \square

Remark 6.9. — The last part of this argument, in which an element is found in the infinite intersection $\bigcap_m x_m\mathbf{A}_m$ given that each finite intersection $\bigcap_{m \leq M} x_m\mathbf{A}_m$ is nonempty, is an instance of the *countable saturation* property of the ultraproduct construction. The completeness that is afforded by the countable saturation property is one of the main reasons why we work in the ultraproduct setting. Note that a similar completeness also appears in the ultralimit $(X, d)/\sim$ of bounded metric spaces (X_n, d_n) , where $X := \prod_{n \rightarrow \alpha} X_n$, $d := \text{stlim}_{n \rightarrow \alpha} d_n$, and \sim is the equivalence relation defined by setting $x \sim y$ whenever $d(x, y) = 0$. Indeed, it is not difficult to use countable saturation to verify that such ultralimits are automatically complete, even if the original spaces X_n are not.

Proof of Proposition 6.1. — We have shown that \mathbf{A}^{32} has the structure of a locally compact local group with respect to the metric d . To complete the proof of Proposition, we need only quotient by the equivalence relation \sim on \mathbf{A}^{32} , defined by $x \sim y$ if and only if $d(x, y) = 0$. The quotient $L := \mathbf{A}^{32}/\sim$ is then a metrisable, locally compact, local group and there is a natural map $\pi : \mathbf{A}^{32} \rightarrow L$. We must check that L is a good model for \mathbf{A}^4 in the sense of Definition 3.5.

Property (i) requires us to show that there is some open neighbourhood U_0 of the identity in L such that $\pi^{-1}(U_0) \subseteq \mathbf{A}^4$ and $U_0 \subseteq \pi(\mathbf{A}^4)$, or in other words some ball $\{x \in \mathbf{A}^{32} : d(x, \text{id}) < \varepsilon\}$ lies in \mathbf{A}^4 . This again follows from (6.4) and the fact that each of the sets \mathbf{A}_k constructed in Lemma 6.7 lies in \mathbf{A}^4 .

Finally, property (ii) in the definition of good model requires us to show that $\pi(\mathbf{A}^4)$ is compact. This is immediate.

To prove property (iii), we first establish the following weaker property:

(iii)': for any open neighbourhood U of the identity in L there is some $U' \subseteq U$ and some ultra finite set $\mathbf{A}' = \prod_{n \rightarrow \alpha} A'_n$ with $\pi^{-1}(U') \subseteq \mathbf{A}' \subseteq \pi^{-1}(U)$.

This is quite easily established: suppose that U contains the ball $B(\text{id}, 2^{-k})$. Then it follows immediately from the inclusions of Lemma 6.5 that we may take $\mathbf{A}' := A_{k+1}$ and then $U' := B(\text{id}, 2^{-k-1})$.

We now upgrade this to property (iii) in the definition of good model. Suppose that $F \subseteq U \subseteq U_0$ with F compact and U open. Then there is some open neighbourhood of the identity U' such that $FU' \subseteq U$. Applying (iii)', we may locate a further open set $U'' \subseteq U'$ and an ultra finite set \mathbf{A}' such that $\pi^{-1}(U'') \subseteq \mathbf{A}' \subseteq \pi^{-1}(U')$. By compactness there are elements x_1, \dots, x_M such that $F \subseteq \bigcup_{m=1}^M x_m U''$; we may assume that these elements lie in $F(U'')^{-1} = FU'' \subseteq U \subseteq U_0$, and hence each is of the form $x_i = \pi(a_i)$ with $a_i \in A$. To conclude the proof of property (iii) simply take $\mathbf{A}'' := \bigcup_{m=1}^M a_m \mathbf{A}'$. This completes the proof of Proposition 6.1. \square

To complete the proof of Theorem 3.10, we invoke results about Hilbert's fifth problem, and specifically the structural theorem of Goldbring [23] describing locally compact local groups, which we state as Theorem B.18 in Appendix B.

Proof of Proposition 6.2. — Suppose that we have a model $\pi : \mathbf{A}^{32} \rightarrow G$ from \mathbf{A}^{32} to a locally compact local group G , and let U_0 be the open neighbourhood of the identity featuring in the definition of good model (Definition 3.5), thus $\pi^{-1}(U_0) \subset \mathbf{A}^4$ and $U_0 \subset \pi(\mathbf{A}^4)$. By Theorem B.18, there are symmetric neighbourhoods $U_2 \subseteq U_1 \subseteq U_0 \subseteq G$ with $U_2^4 \subseteq U_1$ (say) and a compact normal subgroup H of U_2 such that U_1/H is isomorphic to a local Lie group L . Let $\phi : U_1 \rightarrow U_1/H$ be the projection map.

By property (iii) of Definition 3.5 (applied to $\pi : \mathbf{A} \rightarrow G$) there is a symmetric ultra finite set $\tilde{\mathbf{A}} \subseteq \mathbf{A}^4$ with $\pi^{-1}(U_2^2) \subseteq \tilde{\mathbf{A}} \subseteq \pi^{-1}(U_2^3)$. Certainly, the map $\tilde{\pi} := \phi \circ \pi$ is well-defined and gives a homomorphism from $\tilde{\mathbf{A}}^8$ to L ; since $\pi^{-1}(U_2^3)^4 \subset \pi^{-1}(U_2^{12}) \subset \mathbf{A}^4$, we have $\tilde{\mathbf{A}}^4 \subseteq \mathbf{A}^4$, and by Remark 3.7, $\tilde{\mathbf{A}}$ is an ultra approximate group. We verify that this is a good model by checking (i), (ii) and (iii) of Definition 3.5 in turn. For (i), first note that $\tilde{\pi}(\tilde{\mathbf{A}})$ contains $\tilde{U}_0 := \phi(U_2) = U_2H/H \subseteq L$, which is an open neighbourhood of the identity in L since $U_2H \subseteq G$ is open.

Furthermore we have

$$\tilde{\pi}^{-1}(\tilde{U}_0) = \pi^{-1}\phi^{-1}\phi(U_2) \subseteq \pi^{-1}(U_2H) \subseteq \pi^{-1}(U_2^2) \subseteq \tilde{\mathbf{A}}.$$

Turning to (ii), $\tilde{\pi}(\tilde{\mathbf{A}})$ is contained in the compact set $\phi(\overline{U_2^2})$.

Finally, we check the ‘‘approximation by internal sets’’ property, which is (iii) in Definition 3.5. Suppose that $\tilde{F} \subseteq \tilde{U} \subseteq \tilde{U}_0$, with \tilde{F} compact and \tilde{U} open. Then $\phi^{-1}(\tilde{F}) = \tilde{F}H$ is compact, whilst $\phi^{-1}(\tilde{U}) = \tilde{U}H$ is open. The approximation by internal sets property then follows from that fact that $\pi : \mathbf{A} \rightarrow G$ is a good model.

Finally, we check that $\tilde{\mathbf{A}}$ is a large ultra approximate group. To see this note that $\tilde{\pi}(\tilde{\mathbf{A}}^2)$ is contained in a compact subset of L ; therefore there are finitely many elements

x_1, \dots, x_k such that $\tilde{\pi}(\tilde{\mathbf{A}}^2) \subseteq \bigcup_{i=1}^k \tilde{\pi}(x_k)\tilde{U}_0$. It follows that

$$\tilde{\mathbf{A}}^2 \subseteq \bigcup_{i=1}^k x_k \tilde{\pi}^{-1} U_0 \subseteq \bigcup_{i=1}^k x_k \tilde{\mathbf{A}},$$

thereby confirming that $\tilde{\mathbf{A}}$ is an ultra approximate group. By essentially the same argument, \mathbf{A} may be covered by finitely many translates of $\tilde{\mathbf{A}}$; thus $\tilde{\mathbf{A}}$ is indeed large. \square

We now record some analogues of the above results in the setting of *global* ultra approximate groups (i.e. ultraproducts of global \mathbf{K} -approximate groups for some fixed \mathbf{K}), which are closer to the results of Hrushovski [33]. Define a *global model* $\pi : \langle \mathbf{A} \rangle \rightarrow \mathbf{G}$ to be the same notion as a good model $\pi : \mathbf{A}^8 \rightarrow \mathbf{G}$ from Definition 3.5, except that \mathbf{A}^8 is replaced by the whole group $\langle \mathbf{A} \rangle$ generated by \mathbf{A} , and \mathbf{G} is now required to be a global group rather than a local group.

Proposition 6.10 (Global locally compact model). — *Let \mathbf{A} be a global ultra approximate group. Then \mathbf{A}^4 admits a global model $\pi : \langle \mathbf{A} \rangle \rightarrow \mathbf{G}$ by a metrisable locally compact global group \mathbf{G} .*

Proof. — This is obtained by a modification of the proof of Proposition 6.1. The one main change is that the nesting condition $(\mathbf{A}_{i+1}^2)^{\mathbf{A}_0^{100}} \subseteq \mathbf{A}_i$ appearing in Lemma 6.7 needs to be strengthened to $(\mathbf{A}_{i+1}^2)^{\mathbf{A}_0^{100(i+1)}} \subseteq \mathbf{A}_i$, but this is easily accomplished. \square

Proposition 6.11 (From locally compact models to Lie models). — *Let \mathbf{A} be a global ultra approximate group and suppose that \mathbf{A}^4 admits a global model $\pi : \langle \mathbf{A} \rangle \rightarrow \mathbf{G}$ into a locally compact global group \mathbf{G} . Then there is a large ultra approximate group $\tilde{\mathbf{A}}$ of \mathbf{A} which admits a global model $\tilde{\pi} : \langle \mathbf{A} \rangle \rightarrow \mathbf{L}$ into a connected Lie group \mathbf{L} .*

Proof. — This is obtained by a modification of the proof of Proposition 6.2. The one main change is that one needs to replace Theorem B.18 with Theorem B.17. \square

Note that in contrast to Proposition 6.2 that we do *not* assert that the global Lie group \mathbf{L} is simply connected (as this is not provided by the global Gleason-Yamabe theorem (Theorem B.17), which only promises connectedness). And indeed, in general we do not have simple connectedness of the model. For instance, if $\mathbf{A} = \{-N, \dots, N\} \subset \mathbf{Z}/100N\mathbf{Z}$ for some unbounded nonstandard natural number N , then the obvious global model here is the map $\pi : \mathbf{Z}/100N\mathbf{Z} \rightarrow \mathbf{R}/\mathbf{Z}$ defined by $\pi(x) = \text{st}(\frac{x}{100N}) \bmod 1$, and of course the unit circle \mathbf{R}/\mathbf{Z} is not simply connected. On the other hand, $\mathbf{A}^{100} = \mathbf{Z}/100N\mathbf{Z}$ is globally modeled by the trivial group; and so one can still recover simple connectedness by passing from \mathbf{A} to a suitably large power. See [33, Remark 4.11] for some further discussion of this point, as well as Theorem 10.10 below.

Combining Proposition 6.10 and Proposition 6.11 we obtain the following result, originally due to Hrushovski [33].

Proposition 6.12 (Weak global Lie model theorem). — Suppose that \mathbf{A} is a global ultra approximate group. Then there is a large ultra approximate group $\tilde{\mathbf{A}}$ of \mathbf{A} which admits a global model $\tilde{\pi} : \langle \mathbf{A} \rangle \rightarrow \mathbf{L}$ into a connected Lie group \mathbf{L} .

We will strengthen this proposition in Theorem 10.10 below.

Remark 6.13. — Let $\pi : \mathbf{A}^{\delta} \rightarrow \mathbf{L}$ be a good model for an ultra approximate group $\mathbf{A} = \prod_{n \rightarrow \alpha} A_n$ by a locally compact local group \mathbf{L} , and let U_0 be the neighbourhood in Definition 3.5. Let U_1 be a symmetric neighbourhood of the identity such that $U_1^{100} \subset U_0$. For any continuous function $f : \mathbf{L} \rightarrow \mathbf{R}$ with compact support in U_1 , we can define a functional $I(f)$ by the formula

$$I(f) = \inf_{\text{st}} \frac{\sum_{a \in \mathbf{A}} F^+(a)}{|\mathbf{A}|}$$

where $F^+ = \lim_{n \rightarrow \alpha} F_n^+$ is the ultralimit of functions $F_n^+ : A_n \rightarrow \mathbf{R}$, with the nonstandard real $\sum_{a \in \mathbf{A}} F^+$ and nonstandard natural number $|\mathbf{A}|$ defined in the usual fashion as

$$\sum_{a \in \mathbf{A}} F^+(a) := \lim_{n \rightarrow \alpha} \sum_{a_n \in A_n} F_n^+(a_n)$$

and

$$|\mathbf{A}| := \lim_{n \rightarrow \alpha} |A_n|,$$

and the infimum is over all F^+ for which $F^+(a) \geq f(\pi(a))$ for all $a \in \mathbf{A}$. Using Definition 3.5(iii) it is not difficult to also obtain the equivalent formula

$$I(f) = \sup_{\text{st}} \frac{\sum_{a \in \mathbf{A}} F^-(a)}{\sum_{a \in \mathbf{A}} 1}$$

where the supremum is over all F^- for which $F^-(a) \leq f(\pi(a))$ for all $a \in \mathbf{A}$. From these two definitions we see that $I(f)$ is both super-linear and sub-linear, and is thus a continuous linear functional on the space $C_c(U_1)$ of continuous compactly supported functions in U_1 . By the Riesz representation theorem, there thus exists a Radon measure μ on U_1 such that $I(f) = \int_{U_1} f d\mu$ for all $f \in C_c(U_1)$. From the translation invariant properties of $I(f)$, we see that $\mu(gE) = \mu(E)$ for any measurable subset E of U_1 , and any $g \in \mathbf{L}$ such that gE are defined in U_1 , and similarly for gE replaced by Eg . Thus μ is a bi-invariant Haar measure on U_1 ; since \mathbf{A} can be covered by finitely many left-translates of $\pi^{-1}(F)$ for any compact neighbourhood F of the identity, we see that μ is non-trivial (which implies in particular by bi-invariance that the locally compact local group \mathbf{L} is unimodular). This Haar measure can then be used to estimate the (nonstandard) cardinality of various

nonstandard finite sets that are “close” to \mathbf{A} in some sense. Indeed, from the definitions (and the regular nature of Radon measures) we see that

$$\mu(F)|\mathbf{A}| \leq |\mathbf{A}'| \leq \mu(U)|\mathbf{A}|$$

whenever $F \subseteq U \subseteq U_1$, F is compact, U is open, and \mathbf{A}' is a nonstandard set with

$$\pi^{-1}(F) \subset \mathbf{A}' \subseteq \pi^{-1}(U).$$

We will not use this measure μ in this paper, but see [33] for some further discussion of this measure and its relationship to *Kiesler measures* from model theory. One can also use μ to relate the volume growth of \mathbf{A}^m to the volume growth of the model group L , giving some rigorous substance to some of the volume growth heuristics invoked in the examples in Section 3, but we will not formalise this relationship here.

Remark 6.14. — As remarked in [33], the Lie Model theorem is not only valid in the context of nonstandard finite ultra approximate groups, i.e. the ultraproduct of finite K -approximate groups for a fixed K , but also for “continuous” ultra approximate groups, that is to say the ultraproduct of precompact open subsets of a locally compact local group that obey all of the approximate group axioms other than finiteness. See [52] for the basic theory of such continuous approximate groups. Indeed, one can check that the machinery in Section 5 can be adapted to this setting by replacing the cardinality of finite sets with the Haar measure of various precompact open subsets of a locally compact local group, as in [52]. Some other components of this paper, such as the construction of strong approximate groups and Gleason metrics, can also be extended to this setting after some minor notational changes. However, there will be a key place in the argument¹⁴ in Section 9 in which the (nonstandard) finiteness of the ultra approximate groups is used in an absolutely crucial way, namely to locate an element in such a group element of minimal non-zero “escape norm”. As such, the main result of this paper, Theorem 2.10, does not immediately extend to the continuous setting. Indeed, the basic example of a small ball in a Lie group shows that continuous approximate groups need not resemble coset nilprogressions at all. We will not pursue this matter further here.

Some finitary consequences of the Lie Model Theorem. — To illustrate the power of the Lie Model Theorem in the analysis of approximate groups, we offer two fairly quick applications. The reader interested in the proof of our main results may skip ahead to the next section.

The first application is a special case of our main theorem (Theorem 2.10), following Hrushovski [33, Corollary 4.18].

¹⁴ Another, much more minor, place where ultra finiteness is used in Remark 6.13 above, as we implicitly used the trivial fact that counting measure is bi-invariant. In general, one can only conclude that the measure associated to a good model is bi-invariant if each of the individual approximate groups in the ultraproduct is also equipped with a finite bi-invariant measure.

Theorem 6.15 (Hrushovski). — Suppose that G be a group of exponent m and suppose that $A \subseteq G$ is a K -approximate group. Then A^4 contains a genuine subgroup H of G with $|H| \gg_{K,m} |A|$. In particular, by Lemma 5.1, A is covered by $O_{K,m}(1)$ left-translates of H .

Remark 6.16. — When $m = 2$ the group G must be abelian, and in this case the theorem is due to Imre Ruzsa [45].

Proof. — Suppose for sake of contradiction that the claim failed. Then we may find fixed K, m and a sequence of K -approximate groups $A_n \subseteq G_n$ in groups G_n of exponent m , such that for each n , A_n^4 does *not* contain a genuine subgroup H_n of cardinality $|H_n| \geq |A_n|/n$. As usual we form the ultra approximate group $\mathbf{A} := \prod_{n \rightarrow \alpha} A_n$. The ultra-product group $\mathbf{G} := \prod_{n \rightarrow \alpha} G_n$ also has exponent m , and by Hrushovski's Lie model theorem we can find a large approximate group $\mathbf{A}' \subseteq \mathbf{A}^4$ with a local Lie model $\pi : (\mathbf{A}')^{\delta} \rightarrow L$. By Definition 3.5(i), we may find a neighbourhood U_0 of the identity in L such that $\pi^{-1}(U_0) \subseteq \mathbf{A}'$ and $U_0 \subseteq \pi(\mathbf{A}')$. Using the fact that the exponential map is a homeomorphism near the identity of L , we may then find a neighbourhood U_1 of the identity with $U_1^m \subseteq U_0$ such that U_1 contains no elements of order m other than the identity. If $a \in \pi^{-1}(U_1)$, then we conclude that a^m is well-defined in \mathbf{A}' with $\pi(a)^m = \pi(a^m) = \text{id}$, and so $\pi(a)$ is trivial, which means that $\pi^{-1}(U_1) = \ker(\pi)$. As $\pi(\mathbf{A}')$ is precompact, we conclude that \mathbf{A}' is covered by a finite number of translates of $\ker(\pi)$; as \mathbf{A}' is large, \mathbf{A} is also covered by M such translates for some (standard) finite M .

From Definition 3.5(iii), we see that the set $\pi^{-1}(U_1) = \ker(\pi)$ is a nonstandard finite set, and so $\ker(\pi) = \prod_{n \rightarrow \alpha} H_n$ for some finite subsets H_n of G_n . Since $\ker(\pi) \subseteq \mathbf{A}^4$ is a group and \mathbf{A} is covered by M translates of $\ker(\pi)$, we see from Łoś's theorem (Theorem A.6) that for all n sufficiently close to α , $H_n \subseteq A_n^4$ is a group and A_n is covered by M translates of H_n . However if one takes n larger than M then this contradicts the construction of A_n , and the claim follows. \square

Remark 6.17. — The astute reader will notice that the only properties of the local Lie group L that were really used in the above argument were that L was locally compact and had the NSS (no small subgroups property). Thus, one could prove Theorem 6.15 using a weaker form of the Gleason-Yamabe theorem (Theorem B.17), in which the model group is merely locally compact NSS rather than Lie. (The machinery of Hilbert's fifth problem implies that these two concepts coincide, but this is considerably deeper.) However, we do not know of a proof of Theorem 6.15 that avoids the machinery of Hilbert's fifth problem completely, and in particular some variant of the Gleason lemmas is required.

Next, we prove (a slight variant of) the main theorem from Hrushovski's paper [33, Theorem 1.1], which uses the Lie structure (via the Baker-Campbell-Hausdorff formula) more thoroughly than the preceding application.

Theorem 6.18 (*Hrushovski's structure theorem*). — *Let A be a K -approximate group, and let $F : \mathbf{N} \times \mathbf{N} \rightarrow \mathbf{N}$ be a function. Then there exist natural numbers L, M, N with $N \geq F(L, M)$ and $L, M \ll_{K,F} 1$, and nested sets*

$$\{\text{id}\} \subset A_N \subseteq \cdots \subseteq A_1 \subseteq A^4$$

with the following properties:

- (i) *For each $1 \leq n \leq N$, A_n is symmetric;*
- (ii) *For each $1 \leq n < N$, $A_{n+1}^2 \subseteq A_n$;*
- (iii) *For each $1 \leq n \leq N$, A_n is contained in M left-translates of A_{n+1} ;*
- (iv) *For $1 \leq n, m, k \leq N$ with $k < n + m$, the set $[A_n, A_m] := \{[g, h] : g \in A_n, h \in A_m\}$ is contained in A_k ;*
- (v) *A can be covered by L left-translates of A_1 .*

Proof. — Suppose this is not the case. Carefully negating all the quantifiers, we conclude that there exist K, F and a sequence $A^{(n)}$ of K -approximate groups, such that for each n and each $L, M \leq n$, there does not exist $N \geq F(L, M)$ and $A_1^{(n)}, \dots, A_N^{(n)}$ obeying the conclusions of the theorem.

As usual, we form the ultraproduct $\mathbf{A} := \prod_{n \rightarrow \alpha} A^{(n)}$, which is an ultra approximate group. By Theorem 3.10, we may find a large ultra approximate subgroup $\tilde{\mathbf{A}} = \prod_{n \rightarrow \alpha} \tilde{A}^{(n)}$ which has a good model $\phi : \tilde{\mathbf{A}}^8 \rightarrow L$ by a local Lie group.

Let \mathfrak{l} be the Lie algebra of L , and fix an open bounded convex symmetric body B in \mathfrak{L} . Let $\varepsilon > 0$ be a sufficiently small (standard) real number depending on B, L to be chosen later; in particular we may assume that the exponential map is a homeomorphism from εB to $\exp(\varepsilon B)$, and that $\exp(\varepsilon B)$ is contained in the neighbourhood U_0 appearing in Definition 3.5. For each standard natural number $n \geq 1$, we apply Definition 3.5 and Remark 3.7 to find an ultra approximate group \mathbf{A}_n with

$$\pi^{-1}(\exp(10^{-n}\varepsilon B)) \subseteq \mathbf{A}_n \subseteq \pi^{-1}(\exp(2 \times 10^{-n}\varepsilon B));$$

In particular we have the nesting

$$\cdots \subseteq \mathbf{A}_2 \subseteq \mathbf{A}_1 \subseteq \mathbf{A}^4.$$

From the Baker-Campbell-Hausdorff formula, we have

$$\exp(2 \times 10^{-n-1}\varepsilon B)^2 \subseteq \exp(10^{-n}\varepsilon B)$$

if ε is small enough, and thus $\mathbf{A}_{n+1}^2 \subseteq \mathbf{A}_n$. In a similar spirit, we can find an M depending only on the dimension of L or \mathfrak{l} such that each ball $10^{-n}\varepsilon B$ is covered by at most M translates of $4 \times 10^{-n-1}\varepsilon B$, which by the Baker-Campbell-Hausdorff formula again implies, for small enough ε , that each \mathbf{A}_n is covered by at most M left-translates of \mathbf{A}_{n+1} . Finally, another application of the Baker-Campbell-Hausdorff formula reveals that

$$[\exp(2 \times 10^{-n}\varepsilon B), \exp(2 \times 10^{-m}\varepsilon B)] \subseteq \exp(10^{-k}\varepsilon B)$$

whenever $k < n + m$, and hence $[\mathbf{A}_n, \mathbf{A}_m] \subseteq \mathbf{A}_k$.

Finally, since one can cover $\pi(\mathbf{A})$ by a finite number of translates of $\exp(\varepsilon B)$, we see that \mathbf{A} can be covered by at most L left-translates of \mathbf{A}_1 for some standard $L \in \mathbf{N}$.

Now set $\mathbf{A}_n = \prod_{n \rightarrow \infty} A_n^{(n)}$ for some finite sets $A^{(n)}$, and set $N := F(L, M)$. Applying Łos's theorem (Theorem A.6) repeatedly (but only finitely many times), we see that for \mathbf{n} sufficiently close to α the sets $A_1^{(n)}, \dots, A_N^{(n)}, A^{(n)}$ obey all the properties in the conclusion of the theorem. This contradicts the construction of the $A^{(n)}$ for \mathbf{n} larger than L, M , and the claim follows. \square

Remark 6.19. — One can also use the Lie Model Theorem to establish a stronger statement than Theorem 6.18, which roughly speaking asserts that given a (finite) K -approximate group A , one can find a large sub-approximate group A' which has an approximate homomorphism $\pi : (A')^8 \rightarrow L$ into a local Lie group L with bounded range that obeys an approximate version of Property (i) in Definition 3.5, where the accuracy of these approximations exceeds the “complexity” of the model¹⁵ by any given function F . The precise formulation of this statement, which is in fact a logically equivalent “finitisation” of Theorem 3.10, is somewhat complicated. We will not need it elsewhere in the paper, and so we leave it as an exercise to the reader.

7. Strong approximate groups

We now give a combinatorial consequence of the Lie Model Theorem (Theorem 3.10) which will be important later, involving a concept which we will call a *strong approximate group*.

Definition 7.1 (Strong Approximate Group). — Let A be a K -approximate group for some $K \geq 1$. We say that A is a strong K -approximate group if it admits a symmetric subset S such that

$$(7.1) \quad (S^{A^4})^{1000K^3} \subseteq A$$

and for which the following two trapping conditions are satisfied:

- (i) (First trapping condition) If $g, g^2, g^3, \dots, g^{1000} \in A^{100}$ then $g \in A$;
- (ii) (Second trapping condition) If $g, g^2, \dots, g^{10^6 K^3} \in A$ then $g \in S$.

An ultra strong approximate group is an ultraproduct $\mathbf{A} = \prod_{n \rightarrow \infty} A_n$ of strong K -approximate groups A_n , for some $K \geq 1$ independent of \mathbf{n} .

¹⁵ The complexity, which we do not define here, would be some quantity taking account of the dimension and structure constants on the Lie algebra \mathfrak{l} of L , the diameter of the range of π and the inradius of the neighbourhood U_0 appearing in Definition 3.5(i).

At present this definition will seem somewhat unmotivated, although it can be demystified to some extent by remarking that these properties suggest that S and A are behaving like very small neighbourhoods of the identity in a Lie group L , with S much smaller than A . This point should become clearer shortly. The reader should not pay too much attention to exponents such as $1000K^3$ or 10^6K^3 in the definition; they are chosen for the sake of concreteness.

The main reason for introducing this concept is that we will be able to show, in the next section, that the escape norm $\|g\|_{e,\mathbf{A}}$ (defined in Definition 4.3) for an ultra strong approximate group \mathbf{A} has the pleasant properties outlined in Section 4. There is scope for varying the parameters in the definition of strong approximate group, but the ones we have given here are strong enough to prove the desired properties of the escape norm.

It is easy to give examples of strong approximate groups. For instance, if $A = \{-N, \dots, N\}$ (and $K = 3$) then we may take $S = \{-N', \dots, N'\}$ with $N' \sim N/1000K^3$. If A is a subgroup, then we may simply take $S = A$. On the other hand, if one randomly removes a small number (e.g. $N^{0.01}$) of elements symmetrically from $\{-N, \dots, N\}$, the resulting set is likely to remain a $O(1)$ -approximate subgroup, but not a strong $O(1)$ -approximate subgroup.

The main result of this section implies the following.

Proposition 7.2 (Finding a ultra strong approximate group). — Let \mathbf{A} be an ultra approximate group. Then there is a large ultra approximate subgroup $\tilde{\mathbf{A}}$ of \mathbf{A} which is a strong ultra approximate group.

For use in Section 9 we will require the following somewhat more precise result.

Proposition 7.3 (Balls are ultra strong approximate groups). — Let \mathbf{A} be an ultra approximate group with a good model $\pi : \mathbf{A} \rightarrow L$ to a local Lie group L . Let B be an open bounded convex symmetric subset of the Lie algebra \mathfrak{L} of L . Then there exists a standard radius $r_0 > 0$ such that for all $0 < r < r_0$, any symmetric nonstandard finite set $\tilde{\mathbf{A}}$ with

$$(7.2) \quad \pi^{-1}(\exp(rB)) \subseteq \tilde{\mathbf{A}} \subseteq \pi^{-1}(\exp(2rB))$$

is a large strong ultra approximate subgroup of \mathbf{A} .

It is clear that Proposition 7.2 follows from Proposition 7.3, Theorem 3.10, and Definition 3.5(iii); we will, however, only need Proposition 7.3 in the sequel.

We now prove Proposition 7.3. Let $r_0 > 0$ be a sufficiently small quantity depending on \mathbf{A} , π , L , B to be chosen later; in particular, we take r_0 so small so that the exponential map is a homeomorphism from $2r_0B$ to $\exp(2r_0B)$, and $\exp(2r_0B)^{100}$ is contained inside the open neighbourhood U_0 of L_0 from Definition 3.5.

Let $\tilde{\mathbf{A}}$ be as in the proposition. In particular $\tilde{\mathbf{A}}^{100} \subset \mathbf{A}$. By Remark 3.7 we may take $\tilde{\mathbf{A}}$ to be a ultra \mathbf{K} -approximate group for some¹⁶ \mathbf{K} , and therefore an ultra approximate subgroup of \mathbf{A} . Since $\pi(\mathbf{A})$ is precompact, it may be covered by finitely many left-translates of $\exp(r\mathbf{B})$, and so \mathbf{A} can be covered by finitely many left-translates of $\tilde{\mathbf{A}}$. Thus $\tilde{\mathbf{A}}$ is a large ultra approximate subgroup of \mathbf{A} .

It remains to establish that $\tilde{\mathbf{A}}$ is a strong ultra approximate subgroup. Suppose that $g \in \tilde{\mathbf{A}}^{100}$ is such that $g, \dots, g^{1000} \in \tilde{\mathbf{A}}^{100}$. Applying π , we see that

$$\pi(g), \dots, \pi(g)^{1000} \in \exp(2r\mathbf{B})^{100}.$$

Working in exponential coordinates and using the Baker-Campbell-Hausdorff formula we conclude, if r_0 is small enough, that $\pi(g) \in \exp(r\mathbf{B})$ and thus $g \in \tilde{\mathbf{A}}$. We have thus shown the first trapping condition for $\tilde{\mathbf{A}}$.

Next, we use Definition 3.5 to find a symmetric nonstandard finite set \mathbf{S} with

$$\exp(10^{-5}\mathbf{K}^{-3}r\mathbf{B}) \subseteq \mathbf{S} \subseteq \exp(10^{-4}\mathbf{K}^{-3}r\mathbf{B}).$$

From the Baker-Campbell-Hausdorff formula we see that

$$\left(\exp(10^{-4}\mathbf{K}^{-3}r\mathbf{B})^{\exp(2r\mathbf{B})^4}\right)^{1000\mathbf{K}^3} \subseteq \exp(r\mathbf{B})$$

and thus

$$(7.3) \quad ((\mathbf{S})^{\tilde{\mathbf{A}}^4})^{1000\mathbf{K}^3} \subseteq \tilde{\mathbf{A}}.$$

Finally, suppose that $g \in \tilde{\mathbf{A}}$ is such that $g, \dots, g^{10^6\mathbf{K}^3} \in \tilde{\mathbf{A}}$. Applying π , we conclude that

$$\pi(g), \dots, \pi(g)^{10^6\mathbf{K}^3} \in \exp(2r\mathbf{B}).$$

Working in exponential coordinates, we conclude that $\pi(g) \in \exp(10^{-5}\mathbf{K}^{-3}r\mathbf{B})$ and hence $g \in \mathbf{S}$. Thus we have verified the second trapping condition for $\tilde{\mathbf{A}}$.

Finally, we need to push the trapping conditions from the ultraproduct $\tilde{\mathbf{A}}$ back to the finitary setting. Write $\mathbf{A} = \prod_{\mathbf{n} \rightarrow \alpha} \mathbf{A}_{\mathbf{n}}$, $\tilde{\mathbf{A}} = \prod_{\mathbf{n} \rightarrow \alpha} \tilde{\mathbf{A}}_{\mathbf{n}}$ and $\mathbf{S} = \prod_{\mathbf{n} \rightarrow \alpha} \mathbf{S}_{\mathbf{n}}$ for some finite sets $\mathbf{A}_{\mathbf{n}}, \tilde{\mathbf{A}}_{\mathbf{n}}, \mathbf{S}_{\mathbf{n}}$. By Łos's theorem (Theorem A.6), we see that for \mathbf{n} sufficiently close to α , $\tilde{\mathbf{A}}_{\mathbf{n}}$ is symmetric and contains the identity with $\tilde{\mathbf{A}}_{\mathbf{n}}^4 \subset \mathbf{A}_{\mathbf{n}}^4$, with $\tilde{\mathbf{A}}_{\mathbf{n}}^2$ covered by \mathbf{K} left-translates of $\tilde{\mathbf{A}}_{\mathbf{n}}$, that

$$((\mathbf{S}_{\mathbf{n}})^{\tilde{\mathbf{A}}_{\mathbf{n}}^4})^{1000\mathbf{K}^3} \subseteq \tilde{\mathbf{A}}_{\mathbf{n}},$$

and that the first and second trapping properties hold for $\tilde{\mathbf{A}}_{\mathbf{n}}$ and $\mathbf{S}_{\mathbf{n}}$. Thus we see that $\tilde{\mathbf{A}}_{\mathbf{n}}$ is a strong \mathbf{K} -approximate group for \mathbf{n} sufficiently close to α . After redefining $\tilde{\mathbf{A}}_{\mathbf{n}}$ suitably for all other values of \mathbf{n} , we conclude that $\tilde{\mathbf{A}}$ is an ultra strong approximate group as required. This concludes the proof of Proposition 7.3 and hence Proposition 7.2.

¹⁶ Indeed, by using the Baker-Campbell-Hausdorff formula one can take \mathbf{K} to depend only on the dimension of \mathbf{L} , if r_0 is small enough, but we will not need this fact here.

Remark 7.4. — This proposition represents by far the most serious use of Hrushovski’s Lie Model Theorem in our paper. Although we use that theorem elsewhere in the paper, it is only for this proposition that we do not currently have a plausible alternative approach.

8. The escape norm and a Gleason type theorem

In this section we prove a variant of “Gleason’s lemmas” in the setting of approximate groups. These show that if A is a strong approximate group then the escape norm has pleasant properties with respect to product, conjugation and commutation. The role of these lemmas was briefly discussed in Section 4.

Here is a precise statement.

Theorem 8.1 (Gleason-type theorem). — *Suppose that A is a strong \mathbf{K} -approximate group. Consider the escape norm*

$$\|g\|_{e,A} := \inf \left\{ \frac{1}{n+1} : n \in \mathbf{N}; g^i \in A \text{ for all } 0 \leq i \leq n \right\},$$

with the convention that $\|g\|_{e,A} = 1$ when g is undefined. This has the following properties:

- (i) (Conjugation) *If $g, h \in A^{10}$ then $\|g^h\|_{e,A} \leq 1000\|g\|_{e,A}$;*
- (ii) (Product) *We have $\|g_1 \dots g_n\|_{e,A} \leq \mathbf{K}^{O(1)}(\|g_1\|_{e,A} + \dots + \|g_n\|_{e,A})$ if $g_1, \dots, g_n \in A^{10}$;*
- (iii) (Commutators) *If $g, h \in A^{10}$ then we have $\|[g, h]\|_{e,A} \leq \mathbf{K}^{O(1)}\|g\|_{e,A}\|h\|_{e,A}$.*

Note that, as a consequence of (i) and (ii), the set of $g \in A$ with $\|g\|_{e,A} = 0$ is a subgroup normalised by A^{10} .

Remark 8.2. — Note that this lemma is trivial when the ambient local group is abelian. For that reason, this section can be ignored by those readers interested in seeing our alternative proof of the (abelian) Freiman’s theorem.

Motivation and heuristic discussion. — We will shortly give a self-contained proof of Theorem 8.1, but as motivation we first offer some comments and discussion of the context in which these ideas were first invented: the solution of Hilbert’s fifth problem by Gleason, Montgomery-Zippin and Yamabe [20, 21, 39, 40, 60, 61] (see also [23] for the local group analogue of these lemmas).

In that context, the Gleason lemmas show the existence, in an arbitrary locally compact group G , of arbitrarily small compact neighborhoods A of the identity whose associated escape norm satisfies properties (i) to (iii) as above. The Gleason lemmas lie at the heart of Hilbert’s fifth problem and are used at several places in its proof, both in

the reduction step from general locally compact groups to NSS (No Small Subgroups) groups, and in order to deal with NSS groups.

For example, if G is NSS, the Gleason lemmas are needed in order to establish that the set of one-parameter subgroups of G forms a vector space. If $X(t)$ and $Y(t)$ are two one-parameter subgroups, then a natural candidate for $X + Y$ is $\lim_{n \rightarrow +\infty} (X(t/n)Y(t/n))^n$. In order to show that such a limit does exist, the bound (ii) on the escape norm of a product is precisely what is needed. For the full story, the reader may wish to consult the classical references [34, 40], the more recent non-standard treatments of the Gleason lemmas by Hirschfeld [32] and by Goldbring and van den Dries [57], or the blog posts of the third author.¹⁷

To give a flavour of how the Gleason lemmas are proven, let us discuss a simple case of the product estimate, namely

$$(8.1) \quad \|uv\|_{e,A} \leq C(\|u\|_{e,A} + \|v\|_{e,A}).$$

Here, A is a ball $B(\text{id}, 1)$ about the identity in a locally compact group G with the NSS property, where the ball is with respect to some left-invariant distance d , and C is some finite quantity depending on A . In the discussion below we will make use of the following points concerning this situation:

- (i) We may construct a distance d with the additional property that $d(\text{id}, x^g) \leq Cd(\text{id}, x)$ for $g, x \in B(\text{id}, 2)$ (for example by the Birkhoff-Kakutani construction [40, §1.22]).
- (ii) The balls in G enjoy an escape property quite similar to that in the definition of a strong approximate group. More precisely, given $\varepsilon > 0$ there is an $M_\varepsilon \in \mathbf{N}$ such that if $g, g^2, \dots, g^{M_\varepsilon} \in B(\text{id}, 1)$ then $g \in B(\text{id}, \varepsilon)$. The proof of this is by contradiction—taking a limit of putative “bad” g s, one can contradict the NSS property.

The key idea behind the proof of the product estimate (8.1) is to relate the escape norm $\|g\|_{e,A}$ to the auxiliary quantity $\|\partial_g \Psi\|_\infty$, where $\partial_g \Psi(x) = \Psi(g^{-1}x) - \Psi(x)$ and Ψ is a non-negative “bump” function supported on $B(\text{id}, 1)$, let us say with $\|\Psi\|_\infty = \Psi(\text{id}) = 1$. As noted in Lemma 6.3, such a “norm” automatically satisfies the product inequality (with $C = 1$), and so we need only show that $\|g\|_{e,A} \sim \|\partial_g \Psi\|_\infty$ in a suitable sense, and for a suitable Ψ .

In one direction, it is easy to link the two quantities. Indeed if $\|\partial_g \Psi\|_\infty \leq \delta$ for some $\delta > 0$, then a simple telescoping sum argument confirms that $\Psi(g^i) > 0$, and hence $g^i \in A$ whenever $i < 1/\delta$. Therefore

$$(8.2) \quad \|g\|_{e,A} \leq \|\partial_g \Psi\|_\infty.$$

¹⁷ <http://terrytao.wordpress.com/tag/hilberts-fifth-problem/>.

Suppose, conversely, that we know that $g, g^2, \dots, g^n \in A = B(\text{id}, 1)$. Then certainly, by the escape property, we have $g, g^2, \dots, g^{n'} \in B(\text{id}, \varepsilon)$ for some $n' \gg_\varepsilon n$. Now if G were a Lie group, and if Ψ were smooth with bounded derivatives, we would have

$$(8.3) \quad \partial_{g^{n'}} \Psi \approx n' \partial_g \Psi,$$

the approximation being better as ε gets smaller. This immediately gives the bound $\|\partial_g \Psi\|_\infty \lesssim_\varepsilon 1/n$, and thus we have linked the escape norm and the auxiliary norm $\|\partial_g \Psi\|_\infty$ in both directions.

Now unfortunately (8.3) is only an approximate identity and, more seriously, G is not known to be a Lie group. In fact, as noted above, these Gleason lemmas are required to prove statements of that form. On a more positive note, observe that we only need to bound $\|\partial_g \Psi\|_\infty$ above in terms of $\|g\|_{e,A}$ when $g = u$ or $g = v$, and not for all g . We are at liberty to design the auxiliary function Ψ with this in mind.

Now the exact version of (8.3) is basically Taylor's formula, and it reads

$$(8.4) \quad \partial_{g^n} \Psi = n \partial_g \Psi + \sum_{i=0}^{n-1} \partial_g \partial_{g^i} \Psi.$$

(We replace n' by n for ease of notation.) This makes it desirable to bound the second derivatives $\partial_g \partial_{g^i} \Psi$. At this point another key idea enters: it is possible to get good control on these second derivatives when $\Psi = \phi * \psi$ is the convolution of two ‘‘Lipschitz’’ functions, that is

$$\Psi(x) = \phi * \psi(x) = \int \phi(xz^{-1}) \psi(z) dz,$$

the integral being with respect to Haar measure on G . This is because of the formula

$$(8.5) \quad \partial_g \partial_h (\phi * \psi) = \int \partial_g \phi(z) \partial_{h^z} \psi(z^{-1}x) dz.$$

To make this useful, ϕ is chosen to be somewhat Lipschitz with respect to shifts by $g = u$ and $g = v$, and ψ is chosen to be Lipschitz with respect to the distance d . We omit the details.

Rigorous argument. — We turn now to the details of such a strategy in the discrete setting, that is to say a rigorous proof of Theorem 8.1.

Proof of Theorem 8.1. — To simplify the notation, we will abbreviate $\|\cdot\|_{e,A}$ in this proof as $\|\cdot\|_e$.

We start with (i), which is a relatively easy consequence of the first trapping property in the definition of strong approximate group (Definition 7.1). Indeed suppose that $g, g^2, \dots, g^n \in A$ for some n ; then certainly $g^h, (g^h)^2, \dots, (g^h)^n \in A^h \subseteq A^{12}$. By the first trapping property this implies that $g^h, (g^h)^2, \dots, (g^h)^{n'} \in A$ for any $n' \leq n/1000$, and this confirms (i).

The proof of (ii) is significantly trickier and is based on the construction of Gleason that was briefly discussed earlier. In order to facilitate a certain technical “bootstrap argument”, it will be convenient to temporarily replace the escape norm $\|g\|_e$ by the regularised version $\|g\|_e^{(\varepsilon)} := \|g\|_e + \varepsilon$, where $\varepsilon > 0$ is a small quantity. We shall obtain estimates uniform in ε , and then let $\varepsilon \rightarrow 0$.

It is natural to introduce the norm-like quantity

$$d^{(\varepsilon)}(g) := \inf \left\{ \sum_{i=1}^n \|g_i\|_e^{(\varepsilon)} : g = g_1 \dots g_n, n \geq 1 \right\}.$$

It is clear that

$$(8.6) \quad d^{(\varepsilon)}(g) \leq \|g\|_e^{(\varepsilon)}.$$

We shall prove an estimate in the opposite direction, namely

$$(8.7) \quad \|g\|_e \leq \mathbf{K}^{O(1)} d^{(\varepsilon)}(g).$$

The exponent $O(1)$ will be independent of ε . This implies that, for each positive integer n and all g_1, \dots, g_n ,

$$\|g_1 \dots g_n\|_e \leq \mathbf{K}^{O(1)} (\|g_1\|_e^{(\varepsilon)} + \dots + \|g_n\|_e^{(\varepsilon)}).$$

Letting $\varepsilon \rightarrow 0$, we recover the product estimate (ii).

In order to establish this we shall, as in Gleason’s argument, relate $\|g\|_e$ and $d^{(\varepsilon)}(g)$ to an auxiliary quantity $\|\partial_g \Psi\|_\infty$, where $\Psi : A^4 \rightarrow [0, \infty)$ is a certain “smooth” function supported on A^4 . We will specify Ψ shortly; as in Gleason’s argument it will be constructed as a convolution of two functions ϕ and ψ . The former is taken to be a kind of smoothed version of 1_A defined using the metric $d^{(\varepsilon)}$ and Lipschitz for this metric, and the latter constructed using the set S appearing in the definition of strong approximate group (Definition 7.1) and Lipschitz with respect to the word metric on S .

One link between these quantities is relatively easy to establish for any function Ψ with $\Psi(\text{id}) \geq 1$. Indeed suppose that $\|\partial_g \Psi\|_\infty = \delta$ for some $g \in A^{100}$. Then certainly $|\Psi(g^i) - \Psi(g^{i+1})| \leq \delta$ for all i with $g^i \in A^{100}$, which implies by an easy telescoping sum argument that $\Psi(g^i) \geq 1 - \delta i$ for all i . In particular g^i lies in the support of Ψ , and hence in A^4 , for $i < 1/\delta$; note that the hypothesis $g^i \in A^{100}$ can be removed by induction. By the first trapping condition in Definition 7.1 this implies that $g^i \in A$ for $i < 1/1000\delta$, and hence $\|g\|_e \leq 1000\delta$. Thus

$$(8.8) \quad \|g\|_e \leq 1000 \|\partial_g \Psi\|_\infty$$

whenever $g \in A^{100}$.

To establish (8.7) and hence the product estimate (ii) it therefore suffices to prove a bound

$$(8.9) \quad \|\partial_g \Psi\|_\infty \ll K^{O(1)} d^{(\varepsilon)}(g)$$

in the opposite direction for all $g \in A^{100}$ (the claim for $g \notin A^{100}$ being an easy consequence). This argument will depend crucially on the specific form of Ψ . The following two lemmas describe the construction of the functions ϕ and ψ .

Lemma 8.3 (Properties of ϕ). — *There is a function $\phi^{(\varepsilon)} : A^{1000} \rightarrow [0, 1]$ such that*

- (i) $\phi^{(\varepsilon)}(x) = 1$ for $x \in A$;
- (ii) $\phi^{(\varepsilon)}(x) = 0$ if $x \notin A^2$;
- (iii) (Lipschitz bound) For all $g \in A^{1000}$, one has

$$\|\partial_g \phi^{(\varepsilon)}\|_\infty \leq \frac{d^{(\varepsilon)}(g)}{d^{(\varepsilon)}(\text{id}, A^c)}.$$

Here $d^{(\varepsilon)}(y, B) := \inf\{d^{(\varepsilon)}(b^{-1}y) : b \in B\}$, and A^c is the complement of A in G .

Proof. — Define

$$\phi^{(\varepsilon)}(x) := \left(1 - \frac{d^{(\varepsilon)}(x, A)}{d^{(\varepsilon)}(\text{id}, A^c)}\right)_+.$$

Note that this is well-defined since $d^{(\varepsilon)}(\text{id}, A^c) \neq 0$; this would be an issue without the fudge factor of ε that we have introduced.

Obviously $\phi^{(\varepsilon)}(x) = 1$ for $x \in A$. If $\phi^{(\varepsilon)}(x) \neq 0$ then $d^{(\varepsilon)}(\text{id}, x^{-1}A) = d^{(\varepsilon)}(x, A) < d^{(\varepsilon)}(\text{id}, A^c)$, and so $x^{-1}A$ contains a point outside of A^c . This implies that $x \in A^2$.

The Lipschitz bound is easily established. \square

Lemma 8.4 (Properties of ψ). — *There is a function $\psi : A^{1000} \rightarrow [0, 1]$ such that*

- (i) $\psi(x) = 1$ for $x \in A$;
- (ii) $\psi(x) = 0$ if $x \notin A^2$;
- (iii) (Lipschitz bound) $\|\partial_{h^i} \psi\|_\infty \leq 1/10^4 K^3$ for $h \in S$ and $y \in A^4$.

Proof. — Let $Q := S^{A^4}$; recall from the definition of strong approximate group that $Q^N \subseteq A$, where $N := 10^4 K^3$. Define $\psi(g) = 0$ if $g \notin Q^N A$, $\psi(g) = 1$ if $g \in A$ and $\psi(g) = 1 - i/N$ if $g \in Q^{i+1}A \setminus Q^i A$ for $i = 0, 1, \dots, N-1$. The claimed properties of ψ are easily checked. \square

We now define Ψ to be the convolution

$$\Psi(x) := \frac{1}{|\mathbf{A}|} \sum_{y \in \mathbf{A}^2} \phi^{(\varepsilon)}(y) \psi(y^{-1}x)$$

for all $x \in \mathbf{A}^{100}$, with the convention that $\Psi(x) = 0$ for x outside \mathbf{A}^{100} .

We note that

$$\Psi(\text{id}) = \frac{1}{|\mathbf{A}|} \sum_x \phi^{(\varepsilon)}(x) \psi(x^{-1}) \geq \frac{1}{|\mathbf{A}|} \sum_{x \in \mathbf{A}} \phi^{(\varepsilon)}(x) \psi(x^{-1}) = 1,$$

a property required in the proof of (8.8). Note also that since ϕ and ψ are both at most 1 pointwise and are supported on \mathbf{A}^2 we have, for all x such that $\Psi(x) \neq 0$,

$$\Psi(x) = \frac{1}{|\mathbf{A}|} \sum_y \phi^{(\varepsilon)}(y) \psi(y^{-1}x) = \frac{1}{|\mathbf{A}|} \sum_{y \in \mathbf{A}^4} \phi^{(\varepsilon)}(y) \psi(y^{-1}x) \leq \frac{|\mathbf{A}^4|}{|\mathbf{A}|} \leq \mathbf{K}^3,$$

that is to say

$$(8.10) \quad \|\Psi\|_\infty \leq \mathbf{K}^3.$$

Let $g \in \mathbf{A}^{100}$. Now since $\text{id} \in \mathbf{A}$ we have the crude bound $d^{(\varepsilon)}(\text{id}, \mathbf{A}^c) \geq \varepsilon$. It follows from Lemma 8.3 that $\|\partial_g \phi^{(\varepsilon)}\|_\infty \leq d^{(\varepsilon)}(g)/\varepsilon$. From the identity

$$\partial_g \Psi(x) = \frac{1}{|\mathbf{A}|} \sum_{y \in \mathbf{A}^{200}} \partial_g \phi^{(\varepsilon)}(y) \psi(y^{-1}x),$$

we have that

$$|\partial_g \Psi(x)| \leq \|\partial_g \phi\|_\infty \frac{1}{|\mathbf{A}|} \sum_{y \in \mathbf{A}^4} \psi(y^{-1}x) \leq \frac{\mathbf{K}^3}{\varepsilon} d^{(\varepsilon)}(g).$$

This immediately yields the crude bound

$$(8.11) \quad \|\partial_g \Psi\|_\infty \leq \frac{\mathbf{K}^3}{\varepsilon} d^{(\varepsilon)}(g)$$

in the direction of (8.9), the statement we are trying to prove.

Denote by $\mathbf{P}(\mathbf{X})$ the bound

$$(8.12) \quad \|\partial_g \Psi\|_\infty \leq \mathbf{X}d(g)$$

for all $g \in \mathbf{A}^{100}$. We have just demonstrated $\mathbf{P}(\mathbf{K}^3/\varepsilon)$, and we wish to prove $\mathbf{P}(\mathbf{K}^{O(1)})$, which is (8.9). To this end we will implement a bootstrapping argument, showing that

$P(X)$ implies a stronger version of itself, namely $P(X')$ with some $X' < X$, under appropriate conditions.

The hypothesis $P(X)$ (cf. (8.12)) implies an improved Lipschitz bound on ϕ . To see this note that if $d^{(\varepsilon)}(g) < 1/1000X$ then from assumption $P(X)$ we have $\|\partial_g \Psi\|_\infty < 1/1000$ and hence, from (8.8), that $\|g\|_\varepsilon < 1$. By definition of the escape norm this implies that $g \in A$. Phrased in the contrapositive, it follows that $d^{(\varepsilon)}(\text{id}, A^c) \geq 1/1000X$, and therefore the Lipschitz bound in Lemma 8.3 implies that

$$(8.13) \quad \|\partial_g \phi\|_\infty \leq 1000X d^{(\varepsilon)}(g).$$

The bootstrapping argument hinges on the Taylor expansion identity

$$\partial_{g^n} \Psi = n \partial_g \Psi + \sum_{i=0}^{n-1} \partial_{g^i} \partial_g \Psi,$$

valid whenever $g, \dots, g^n \in A^{200}$ (say). This identity implies, using the triangle inequality and (8.10), that

$$(8.14) \quad \|\partial_g \Psi\|_\infty \leq \frac{1}{n} \|\partial_{g^n} \Psi\|_\infty + \frac{1}{n} \sum_{i=0}^{n-1} \|\partial_{g^i} \partial_g \Psi\|_\infty \leq \frac{2K^3}{n} + \frac{1}{n} \sum_{i=0}^{n-1} \|\partial_{g^i} \partial_g \Psi\|_\infty.$$

To use this, we need to focus attention on the first and second derivatives of Ψ . To bound the first derivative we use the identity

$$\partial_h \Psi(x) = \frac{1}{|A|} \sum_y \phi(y) \partial_{h'} \psi(y^{-1}x),$$

valid for $h \in A^{100}$. Since $\phi \leq 1$, this and the Lipschitz bound on ψ given in Lemma 8.4 imply that

$$(8.15) \quad \|\partial_h \Psi\|_\infty \leq \frac{1}{|A|} \sum_{y \in A^2} \|\partial_{h'} \psi\|_\infty \leq 1/10^4 K^2$$

if $h \in S$.

We turn to the second derivative $\partial_h \partial_g \Psi$ for $g \in A$ and $h \in S$. Here we use the identity

$$\partial_h \partial_g \Psi(x) = \frac{1}{|A|} \sum_y (\partial_g \phi)(y) \partial_{h'} \psi(y^{-1}x).$$

Recalling that ϕ, ψ are supported on A^2 and using the Lipschitz bound (8.13) on ϕ together with the Lipschitz bound on ψ given in Lemma 8.4, we obtain the bound

$$(8.16) \quad \|\partial_h \partial_g \Psi\|_\infty \leq \frac{1}{|A|} \sum_{j \in A^4} \|\partial_g \phi\|_\infty \|\partial_{h^j} \psi\|_\infty \leq \frac{1}{10} X d^{(\varepsilon)}(g)$$

if $g \in A$ and $h \in S$.

These bounds are useful in (8.14) provided that n is such that $g, g^2, \dots, g^n \in S$. However, the second trapping property in the definition of strong approximate group ensures that this is so for a reasonably large value of n , indeed for n as large as $\frac{1}{10^6 K^3 \|g\|_\varepsilon}$. Taking n this large and substituting into (8.14) yields

$$\|\partial_g \Psi\|_\infty \leq 10^7 K^6 \|g\|_\varepsilon + \frac{1}{10} X d^{(\varepsilon)}(g) \leq X' \|g\|_\varepsilon^{(\varepsilon)},$$

where $X' = 10^7 K^6 + \frac{1}{10} X$ and $g \in S$. The claim also trivially holds when $g \notin S$.

It is easy to improve this to the stronger statement $P(X')$ using the triangle inequality $\|\partial_{gh} \Psi\|_\infty \leq \|\partial_g \Psi\|_\infty + \|\partial_h \Psi\|_\infty$, already observed in (6.2). Indeed for every $\eta > 0$ there are, by the definition of d, g_1, \dots, g_n such that $g = g_1 \dots g_n$ and

$$d^{(\varepsilon)}(g) > \|g_1\|_\varepsilon^{(\varepsilon)} + \dots + \|g_n\|_\varepsilon^{(\varepsilon)} - \eta.$$

Therefore

$$\begin{aligned} \|\partial_g \Psi\|_\infty &\leq \|\partial_{g_1} \Psi\|_\infty + \dots + \|\partial_{g_n} \Psi\|_\infty \leq X' (\|g_1\|_\varepsilon^{(\varepsilon)} + \dots + \|g_n\|_\varepsilon^{(\varepsilon)}) \\ &\leq X' (\eta + d^{(\varepsilon)}(g)). \end{aligned}$$

Since η was arbitrary, we do indeed obtain the bound $\|\partial_g \Psi\|_\infty \leq X' d(g)$, which is the statement $P(X')$.

By repeating this deduction of $P(X')$ from $P(X)$ many times, we see that the crude bound $P(K^3/\varepsilon)$, established in (8.11), eventually implies $P(10^9 K^6)$, and hence (8.9). By earlier remarks, this concludes the proof of (ii), the inequality for products.

Finally, we turn to the commutator bound (iii). Now that we have the product inequality (ii), we may define a function ϕ obeying the properties in Lemma 8.3 but using $\|g\|_\varepsilon$ instead of the fudged quantity $\|g\|_\varepsilon^{(\varepsilon)} = \|g\|_\varepsilon + \varepsilon$, that is to say with

$$d(g) := \inf \left\{ \sum_{i=1}^n \|g_i\|_\varepsilon; g = g_1 \dots g_n, n \geq 1 \right\}.$$

This is because (ii) implies the lower bound $d(\text{id}, A^\varepsilon) \geq K^{-O(1)}$, and in particular $d(\text{id}, A^\varepsilon) \neq 0$. Moreover we have the Lipschitz bound

$$(8.17) \quad \|\partial_g \phi\|_\infty \leq K^{O(1)} d(g).$$

We will use this function ϕ in establishing (iii), the bound for commutators. Once again we consider an auxillary function Φ , defined now to be the convolution

$$\Phi(x) := \frac{1}{|\mathbf{A}|} \sum_{y \in \mathbf{A}^2} \phi(y) \phi(y^{-1}x)$$

again with the convention that Φ vanishes outside of \mathbf{A}^{100} . We observe the identity

$$\partial_g \partial_h \Phi - \partial_h \partial_g \Phi = -\mathbf{T}_{hg} \partial_{[g,h]} \Phi,$$

for $g, h \in \mathbf{A}^{10}$, where \mathbf{T}_g denotes the shift defined by $\mathbf{T}_g f(x) := f(g^{-1}x)$ if $g^{-1}x$ is well-defined, and 0 otherwise. It follows that

$$\|\partial_{[g,h]} \Phi\|_\infty \leq \|\partial_h \partial_g \Phi\|_\infty + \|\partial_g \partial_h \Phi\|_\infty.$$

By the first bound in (8.16) (which holds equally well for this Φ) we have

$$\|\partial_{[g,h]} \Phi\|_\infty \leq \frac{2}{|\mathbf{A}|} \sum_{y \in \mathbf{A}^4} \|\partial_g \phi\|_\infty \|\partial_{h^y} \phi\|_\infty.$$

From (8.17) we obtain

$$\|\partial_{[g,h]} \Phi\|_\infty \leq \mathbf{K}^{\mathbf{O}(1)} d(g) \sup_{y \in \mathbf{A}^4} d(h^y) \leq \mathbf{K}^{\mathbf{O}(1)} \|g\|_e \sup_{y \in \mathbf{A}^4} \|h^y\|_e.$$

By part (i), this implies

$$\|\partial_{[g,h]} \Phi\|_\infty \leq \mathbf{K}^{\mathbf{O}(1)} \|g\|_e \|h\|_e.$$

To conclude, we note that (8.8) holds for this new auxillary function Φ as well, since the only fact we used in establishing that other than trapping properties of \mathbf{A} was the lower bound $\Phi(\text{id}) \geq 1$.

This, at last, concludes the proof of Theorem 8.1. \square

To conclude this section we assemble the main results of it and the previous section in a portable form. The following is the only result we shall need from Section 7 and the present section going forward to the next (and final) part of the paper.

Proposition 8.5. — *Suppose that \mathbf{A} is an ultra approximate group and that $\pi : \mathbf{A}^8 \rightarrow \mathbf{L}$ is a good model for \mathbf{A} into a connected Lie group \mathbf{L} with Lie algebra \mathfrak{l} . Let \mathbf{B} be an arbitrary compact convex neighbourhood of 0 in \mathfrak{l} . Then, for sufficiently small r, r' with $2r > r' > r > 0$, we may find a large strong ultra approximate subgroup \mathbf{A}' of \mathbf{A} such that*

- (i) $\pi^{-1}(\exp(r\mathbf{B})) \subset \mathbf{A}' \subset \pi^{-1}(\exp(r'\mathbf{B}))$;
- (ii) $(\mathbf{A}')^{10^4}$ is well defined;

- (iii) *The escape norm $\|g\|_{e,\mathbf{A}'}$ satisfies*
 - (a) (Conjugation) *If $g, h \in (\mathbf{A}')^{10}$ then $\|h^{-1}gh\|_e = O(\|g\|_e)$;*
 - (b) (Product) *If n is a nonstandard natural number and $g_1, \dots, g_n \in (\mathbf{A}')^{10}$ is a non-standard finite sequence of elements of $(\mathbf{A}')^{10}$ (i.e. an ultraproduct of standard finite sequences, see Section A) then $\|g_1 \dots g_n\|_e = O(\sum_{i=1}^n \|g_i\|_e)$;*
 - (c) (Commutators) *If $g, h \in (\mathbf{A}')^{10}$ then we have $\|[g, h]\|_e = O(\|g\|_e \|h\|_e)$.*
- (iv) *The set $\mathbf{H} := \{g \in \mathbf{A}'; \|g\|_e = 0\}$ is a global internal subgroup, that is to say it is of the form $\mathbf{H} = \prod_{n \rightarrow \alpha} \mathbf{H}_n$, where $\mathbf{H}_n \subset \mathbf{A}_n$ contains id and is stable under multiplication and inverse, which is contained in \mathbf{A}' and is normalised by \mathbf{A}' .*

Proof. — The existence of \mathbf{A}' satisfying (i) and (ii) follows from part (iii) of Definition 3.5. If r, r' are small enough then Proposition 7.3 ensures that \mathbf{A}' is a ultra strong approximate group in the sense of Definition 7.1. Properties (iii)(a), (b) and (c) then follow immediately from Theorem 8.1 and taking ultraproducts, and (iv) then follows from (iii). \square

Remark 8.6. — Observe that if \mathbf{A} is a strong ultra approximate group, that is to say an ultraproduct of \mathbf{K} -strong finite approximate groups, and if \mathbf{L} is a locally compact model of \mathbf{A} as given for example by Proposition 6.1, then from the strong approximate group hypothesis made on \mathbf{A} we see that the standard part of the escape norm $\text{st}(\|g\|_{e,\mathbf{A}})$ and the escape norm of $\pi(g) \in \mathbf{L}$ with respect to the neighborhood of the identity $\pi(\mathbf{A})$ of \mathbf{L} are comparable. Namely $\|\pi(g)\|_{e,\pi(\mathbf{A})} \ll \text{st}(\|g\|_{e,\mathbf{A}}) \leq \|\pi(g)\|_{e,\pi(\mathbf{A})}$. As a consequence, if we take the standard parts of the escape norm in properties (i) to (iii), then what we obtain is precisely the analogous properties for the escape norm in \mathbf{L} with respect to $\pi(\mathbf{A})$. In that case, the three properties are essentially equivalent to the original Gleason lemmas in the literature on Hilbert's fifth problem, applied to the locally compact (local) group \mathbf{L} . In the sequel however, it will be very important that the three bounds (i) to (iii) obtained in Proposition 8.5 hold at the ultra level in ${}^*\mathbf{R}$ and not only at the level of standard parts.

9. Proof of the main theorem

In this section, we complete the proof of our main theorem, Theorem 4.2. We will do so by first reducing to the case when \mathbf{A} has no global internal subgroup. For convenience, we introduce the following definition.

Definition 9.1 (No small subgroups). — *An ultra approximate group \mathbf{A} has the NSS property if \mathbf{A} does not contain any non-trivial global internal subgroup.*

By a *global internal subgroup* of $\mathbf{A} = \prod_{n \rightarrow \alpha} \mathbf{A}_n$, we mean a subset of the form $\prod_{n \rightarrow \alpha} \mathbf{H}_n$, where $\mathbf{H}_n \subseteq \mathbf{A}_n$ is a genuine subgroup. Note that \mathbf{A} is NSS if and only if, for any $g \in \mathbf{A} \setminus \text{id}$, the escape norm $\|g\|_{e,\mathbf{A}}$ is non-zero (though it may be infinitesimal). We remark that an analogous NSS condition for locally compact groups plays a key role in the theory of Hilbert's fifth problem.

Example 22. — Let $N \in {}^*\mathbf{N}$ be an unbounded (nonstandard) integer. Then the interval $\mathbf{A} := [-N, N]$ (in the nonstandard integers ${}^*\mathbf{Z}$) is NSS. Note that while \mathbf{A} contains global subgroups such as \mathbf{Z} or $\{x \in {}^*\mathbf{Z} : x = o(N)\}$, such subgroups are not internal (they are not the ultralimits of standard sets).

Clearly, any ultra approximate subgroup of an NSS ultra approximate group is also an NSS ultra approximate group. Using the Gleason lemmas from Section 8 we can reduce the proof of our main theorem to consideration of the NSS case.

Proposition 9.2 (NSS reduction). — *Let \mathbf{A} be an ultra approximate group. Then there exists a large ultra approximate subgroup \mathbf{A}' of \mathbf{A} , with $(\mathbf{A}')^{1000}$ well-defined and contained in \mathbf{A}^4 , and a global internal subgroup \mathbf{H} contained in \mathbf{A}' and normalised by $(\mathbf{A}')^{100}$, such that \mathbf{A}'/\mathbf{H} is an NSS ultra approximate subgroup, which admits a connected Lie group as a good model.*

We refer the reader to Definition 3.5 for the definition of a good model. Here \mathbf{A}'/\mathbf{H} denotes the quotient local group as defined in Lemma B.12.

Proof. — By Proposition 7.2 there is a ultra (strong) approximate group $\mathbf{A}' \subseteq \mathbf{A}^4$ which is large relative to \mathbf{A} , for which $(\mathbf{A}')^{10^4}$ is well-defined and contained in \mathbf{A}^4 , and a good model $\pi : (\mathbf{A}')^8 \rightarrow \mathbf{L}$, where \mathbf{L} is a connected Lie group. Let \mathbf{B} be an open bounded convex symmetric neighbourhood of the identity in the Lie algebra of \mathbf{L} . Then for sufficiently small $r > 0$, $\exp(r\mathbf{B})$ contains no non-trivial subgroups of \mathbf{L} .

Let \mathbf{H} denote the global internal subgroup $\mathbf{H} = \{g \in \mathbf{A}' : \|g\|_e = 0\}$ given by Proposition 8.5. Since \mathbf{H} is normalised by \mathbf{A}' , it is also normalised by $(\mathbf{A}')^{1000}$. We may then apply Lemma B.12 and consider the quotient local group $(\mathbf{A}')^{100}/\mathbf{H}$. Then $(\mathbf{A}')^8/\mathbf{H} = (\mathbf{A}'/\mathbf{H})^8$ is well-defined. Since \mathbf{A}' , \mathbf{H} are nonstandard finite symmetric sets, \mathbf{A}'/\mathbf{H} is also; since $(\mathbf{A}')^2$ can be covered by finitely many left-translates of \mathbf{A} , $(\mathbf{A}'/\mathbf{H})^2$ can be covered by finitely many left-translates of \mathbf{A}'/\mathbf{H} . We conclude that \mathbf{A}'/\mathbf{H} is an ultra approximate group.

Since $\exp(r\mathbf{B})$ contains no non-trivial subgroups, the image of \mathbf{H} under π has to be trivial, thus the homomorphism π descends to a homomorphism of \mathbf{A}'/\mathbf{H} to \mathbf{L} , which satisfies the conditions for a good model (see Definition 3.5). By construction, every element $g \in \mathbf{A}'$ that is not in \mathbf{H} has positive (but nonstandard) escape norm $\|g\|_{e,\mathbf{A}'}$. If $g \in \mathbf{A}'$ and $\langle [g] \rangle \subseteq \mathbf{A}'/\mathbf{H}$, where $[g]$ is the class of g in \mathbf{A}'/\mathbf{H} , then $\langle g \rangle \subseteq \mathbf{A}'^2$. On the other hand \mathbf{A}' is a strong ultra approximate group, and thus $\|g\|_{e,\mathbf{A}'}$ is non-zero if and only if $\|g\|_{e,\mathbf{A}'^2}$ is non-zero. This implies that every non-identity element $[g]$ in \mathbf{A}'/\mathbf{H} also has positive escape norm $\|[g]\|_{e,\mathbf{A}'/\mathbf{H}}$. Thus \mathbf{A}'/\mathbf{H} is NSS and the claim follows. \square

Let us now state Theorem 4.2 in the special case of NSS groups, and show how the general case of Theorem 4.2 follows from it.

Theorem 9.3 (NSS approximate groups contain large nilprogressions). — *Let \mathbf{A} be an NSS ultra approximate group which admits a connected Lie group \mathbf{L} as a good model. Then \mathbf{A}^4 contains a*

nondegenerate ultra nilprogression P in normal form, which is large relative to \mathbf{A} . Furthermore, the rank and step of P are no greater than the dimension of L .

Proof that Theorem 9.3 implies Theorem 4.2. — Let \mathbf{A} be an ultra approximate group. We may find a large ultra approximate subgroup \mathbf{A}' of \mathbf{A} which satisfies the conclusions of Proposition 9.2. We may then apply Theorem 9.3 to \mathbf{A}'/H and find in $(\mathbf{A}')^4/H$ a nondegenerate ultra nilprogression P_0 in normal form with $|P_0| \gg |\mathbf{A}/H|$. We can write $P_0 = P_s(\bar{u}_1, \dots, \bar{u}_r; N_1, \dots, N_r)$, where $N_i \in {}^*\mathbf{N}$ are unbounded and $\bar{u}_i \in \mathbf{A}'/H$. We may then pick arbitrary lifts $u_i \in \mathbf{A}'$ and set $P = P(u_1, \dots, u_r; N_1, \dots, N_r)$. Then HP is a nondegenerate ultra coset progression in normal form contained in $(\mathbf{A}')^5 \subseteq \mathbf{A}^4$, and $|HP| \geq |H||P_0| \gg |\mathbf{A}|$ as desired. \square

We turn now to the proof of Theorem 9.3, which will occupy the remainder of this section. We begin with a brief sketch, fleshing out a little more the overview given in Section 4. The proof will proceed by induction on the dimension of the connected Lie group L . The base case of the induction, when $\dim L = 0$, is trivial as in this case the NSS ultra approximate group \mathbf{A} is also trivial. To treat the induction step we will consider an element u of \mathbf{A} with smallest possible escape norm. The existence of such an element is guaranteed by our standing hypothesis that approximate groups are finite objects, i.e. that each A_n in $\mathbf{A} = \prod_{n \rightarrow \alpha} A_n$ is finite. Then we will mod out \mathbf{A} by the geometric $P := \{u^n, |n| \leq 1/\|u\|_{e,\mathbf{A}}\}$, where u is an element of \mathbf{A} which the smallest possible escape norm $\|u\|_{e,\mathbf{A}}$. The quotient local group \mathbf{A}/P (in the sense of Lemma B.12) will be shown to be both NSS and to admit a Lie group with dimension at most $\dim L - 1$ as a good model. It is at this step that we crucially rely on the fact that we are only quotienting out by a *local* group, the progression P , rather than a *global* one such as the group $\langle u \rangle$ generated by u . We do this in order to avoid accidentally creating torsion with an excessively large quotient. Indeed, it is because of this component of the induction that it was necessary to cast the entire argument in the setting of local groups rather than global groups, even if one had been willing to restrict the main results of the paper to the global group case. Finally, making key use of the properties of the escape norm given by the Gleason lemmas, we will lift the nilprogression from \mathbf{A}/P to \mathbf{A} .

Let us turn to the details.

Proof of Theorem 9.3. — Let \mathbf{A} be an NSS ultra approximate group which admits a connected Lie group L as a good model $\pi : \mathbf{A}^8 \rightarrow L$. We proceed by induction on $\dim L$ and first dispose of the trivial case when L has dimension zero. As L is connected, it must thus be trivial. Applying Definition 3.5(iii), we conclude that \mathbf{A}^8 is a large global internal subgroup of \mathbf{A} . Since \mathbf{A} is NSS, this kernel must therefore be trivial. Therefore \mathbf{A} is trivial.

Now suppose that $\dim L \geq 1$, and that the claim has already been proven for connected Lie groups of smaller dimension. To complete the proof of Theorem 9.3 it suffices to establish the following lemma.

Lemma 9.4 (*Induction step*). — Suppose that \mathbf{A} is an ultra approximate group admitting a connected Lie group \mathbf{L} of positive dimension as a good model. Then \mathbf{A} contains large ultra approximate subgroups $\mathbf{A}''' \subseteq \mathbf{A}'' \subseteq \mathbf{A}' \subseteq \mathbf{A}$ with the following properties. Let $u \in \mathbf{A}'$ be such that $\|u\|_{e,\mathbf{A}'}$ is minimal and non zero, and set $\mathbf{P} := \{u^n : |n| < 1/\|u\|_{e,\mathbf{A}'}\}$. Then \mathbf{P} commutes with $(\mathbf{A}')^{10}$ and obeys the following properties:

- (i) the quotient \mathbf{A}'''/\mathbf{P} is an ultra approximate group which admits a connected Lie group of dimension $\dim \mathbf{L} - 1$ as a good model, whose Lie algebra is formed from the Lie algebra of \mathbf{L} by quotienting out by a one-dimensional central subalgebra;
- (ii) if \mathbf{A} is NSS, so is \mathbf{A}'''/\mathbf{P} ;
- (iii) to any large ultra nilprogression $\overline{\mathbf{Q}}$ in \mathbf{A}'''/\mathbf{P} in normal form, one can associate a large ultra nilprogression \mathbf{Q} in \mathbf{A}''' in normal form, whose rank exceeds the rank of $\overline{\mathbf{Q}}$ by at most one, and similarly for the step; and
- (iv) $(\mathbf{A}''')^4 \subseteq \mathbf{A}''$.

Proof of Theorem 9.3. — Indeed apply the induction hypothesis to \mathbf{A}'''/\mathbf{P} , which we can do by (i) and (ii). We may then conclude, using (iv), that \mathbf{A}'''/\mathbf{P} contains a large ultra nilprogression. Finally, apply (iii) to conclude. \square

Proof of Lemma 9.4. — Take \mathbf{B} to be some small convex neighbourhood of 0 in the Lie algebra \mathfrak{l} of \mathbf{L} . We shall take $\mathbf{A}', \mathbf{A}'', \mathbf{A}'''$ to be such that

$$(9.1) \quad \pi^{-1}(\exp(\mathbf{B})) \subseteq \mathbf{A}' \subseteq \pi^{-1}(\exp(1.001\mathbf{B}))$$

and

$$(9.2) \quad \pi^{-1}(\exp(\delta\mathbf{B})) \subseteq \mathbf{A}'' \subseteq \pi^{-1}(\exp(1.001\delta\mathbf{B}))$$

and

$$(9.3) \quad \pi^{-1}\left(\exp\left(\frac{\delta}{10}\mathbf{B}\right)\right) \subseteq \mathbf{A}''' \subseteq \pi^{-1}\left(\exp\left(1.001\frac{\delta}{10}\mathbf{B}\right)\right),$$

where $\delta > 0$ is a small (standard) real number to be specified later.

It follows from Proposition 8.5 that large ultra approximate subgroups of \mathbf{A} exist with these properties, and furthermore that, if \mathbf{B} is small enough, the escape norm $\|\cdot\|_{e,\mathbf{A}'}$ satisfies the conjugation, product and commutator inequalities laid out in (iii) of that proposition. Note also that $\mathbf{A}', \mathbf{A}''$ admit \mathbf{L} as a good model and have the NSS property.

Property (iv) of the lemma is essentially immediate; we turn to the more substantial (i), (ii) and (iii).

We begin with the proof of (i). Recall that $u \in \mathbf{A}'$ is chosen so that $\|u\|_{e,\mathbf{A}'}$ is minimal and nonzero. Observing that $\|x\|_{e,\mathbf{A}'} \leq 100\delta$ for $x \in (\mathbf{A}')^{10}$, it follows from the commutator estimate of Proposition 8.5 (iii)(b) that for such x we have

$$\|[x, u]\|_{e,\mathbf{A}'} = \mathcal{O}(\|x\|_{e,\mathbf{A}'}\|u\|_{e,\mathbf{A}'}) < \|u\|_{e,\mathbf{A}'}$$

provided that δ is chosen sufficiently small in terms of the implied constant $\mathcal{O}(\cdot)$.

Note that $[x, u]$ lies in \mathbf{A}' , rather than merely $(\mathbf{A}')^4$, since its escape norm is less than 1. From the extremal property of u , it follows that $[x, u] = \text{id}$, that is to say x commutes with u , whenever $x \in (\mathbf{A}'')^{10}$.

Recall that we are taking $\mathbf{P} := \{u^n : n \leq 1/\|u\|_{e, \mathbf{A}'}\}$.

Since $(\mathbf{A}'')^6$ is well defined, we may apply Lemma B.12 and form the quotient local group $\mathbf{A}''/\mathbf{P} \simeq \prod_{n \rightarrow \alpha} \mathbf{A}''_n/\mathbf{P}_n$, which is clearly also an ultra approximate group. We now show that \mathbf{A}''/\mathbf{P} admits a proper quotient of \mathbf{L} as a good model. To do this, we first verify that $\pi(\mathbf{P})$ is a non-trivial central one-parameter local subgroup of \mathbf{L} .

Since $\dim \mathbf{L} \geq 1$, the groups \mathbf{A}' , \mathbf{A}'' are non trivial, and this implies that $\|u\|_{e, \mathbf{A}'}$ is infinitesimal, i.e. that $M'_0 := 1/\|u\|_{e, \mathbf{A}'}$ is unbounded. Let $n \in {}^*\mathbf{N}$ be such that $n = o(M'_0)$. We must have $u^n \in \ker \pi$, because $u^{kn} \in \mathbf{A}'$ for all $k \in \mathbf{N}$. Define a map $\phi : [-100, 100] \rightarrow \mathbf{L}$ by setting

$$(9.4) \quad \phi(t) := \pi(u^{\lfloor tM'_0 \rfloor}),$$

where $\lfloor \cdot \rfloor$ is the (nonstandard) greatest integer function. Then $\pi(u^n) = \phi(\text{st}(n/M'_0))$ for all $n \in [-100M'_0, 100M'_0]$, and ϕ is a local homomorphism in the sense that $\phi(t)\phi(s) = \phi(t+s)$ whenever $t, s, t+s \in [-100, 100]$. Also $\pi(\mathbf{P}) = \phi([-1, 1])$. Finally, we verify that ϕ is continuous. Because of the local homomorphism property, it is enough to check this at 0. If t is small, then $(\phi(t))^k = \phi(tk) \in \exp(1.001\mathbf{B})$ for every integer $k \in [0, \frac{1}{t}]$, hence $\phi(t) \in \exp(1.001\mathbf{B}/k)$ is close to the identity in \mathbf{L} , which gives the desired continuity.

As ϕ is a continuous homomorphism from $[-1, 1]$ to the Lie group \mathbf{L} , there exists an element \mathbf{X} of the Lie algebra \mathfrak{l} such that $\phi(t) = \exp(t\mathbf{X})$ for all $t \in [-1, 1]$. Moreover $\mathbf{X} \in 1.001\mathbf{B}$. On the other hand by the definition of the escape norm we have $u^{M'_0} \notin \mathbf{A}'$, and hence $\phi(1) = \exp(\mathbf{X}) \notin \exp(\mathbf{B})$ and thus $\mathbf{X} \notin \mathbf{B}$. In particular \mathbf{X} is non-zero, and it follows that $\phi([-1, 1])$ is a non-trivial local one-parameter subgroup of \mathbf{L} . Finally it is central in \mathbf{L} , because \mathbf{L} is connected and $\phi(t) = \pi(u^{\lfloor tM'_0 \rfloor})$ commutes with the neighbourhood of identity $\pi(\mathbf{A}'')$ as shown above. Thus \mathbf{X} lies in the centre of the Lie algebra \mathfrak{l} .

If we choose a neighbourhood \mathbf{U} of the identity in \mathbf{L} small enough, then by Lemma B.12 we may form the quotient space $\mathbf{U}/\phi([-1, 1])$, which one easily verifies to be a local Lie group of dimension $\dim \mathbf{L} - 1$, whose Lie algebra is obtained from the Lie algebra of \mathbf{L} by quotienting out by a one-dimensional central subalgebra. By Lie's third theorem every local Lie group is locally identifiable with an open neighbourhood of a global connected Lie group \mathbf{L}' , which in our case still has dimension $\dim \mathbf{L} - 1$. Thus, by shrinking \mathbf{U} if necessary, we may find a local homomorphism $\eta : \mathbf{U} \rightarrow \mathbf{L}'$ whose kernel lies in $\phi([-1, 1])$. The local homomorphism $\eta \circ \pi : (\mathbf{A}'')^8 \rightarrow \mathbf{L}'$ then pushes down to a local homomorphism $\psi : (\mathbf{A}'')^8/\mathbf{P} \rightarrow \mathbf{L}'$. Choosing δ smaller if necessary, we may assume that ψ is defined on all of $(\mathbf{A}''/\mathbf{P})^8$, thus making \mathbf{L}' a good model for \mathbf{A}''/\mathbf{P} . Note we may also ensure that $\pi(\mathbf{A}'')$ contains no non-trivial subgroup of \mathbf{L}' , a property that will be needed in Lemma 9.5 below. This completes the proof of (i).

We turn now to (ii), which asserted that \mathbf{A}'''/\mathbf{P} is NSS. In fact we shall prove the same statement for \mathbf{A}''/\mathbf{P} , from which the statement for \mathbf{A}'''/\mathbf{P} follows (or note that an

identical proof works). Key to this endeavour is the following lifting lemma, which we will require again in the proof of (iii).

Lemma 9.5 (Lifting lemma). — *Let $g \in \mathbf{A}''/\mathbf{P}$, and let $\kappa : (\mathbf{A}'')^8 \rightarrow (\mathbf{A}'')^8/\mathbf{P}$ be the projection map. Then there exists $\tilde{g} \in \mathbf{A}''$ such that $\kappa(\tilde{g}) = g$ and $\|\tilde{g}\|_{e,\mathbf{A}''} = O(\|g\|_{e,\mathbf{A}''/\mathbf{P}})$.*

Let us first remark on why the NSS property of \mathbf{A}''/\mathbf{P} follows quickly from this. Indeed suppose that $g \in \mathbf{A}''/\mathbf{P}$ is not the identity. Then the element \tilde{g} generated by the above lemma is not the identity either, and hence has positive escape norm since \mathbf{A}'' is NSS. By the lemma, g also has positive escape norm. Since $g \neq \text{id}$ was arbitrary, this establishes the NSS property for \mathbf{A}''/\mathbf{P} .

Proof of Lemma 9.5. — Fix $g \in \mathbf{A}''/\mathbf{P}$. Let \tilde{g} be a lift of g in \mathbf{A}'' which minimizes the escape norm $\|\tilde{g}\|_{e,\mathbf{A}''}$ among all possible lifts of g . If \tilde{g} is trivial, then so is g and there is nothing to prove. Therefore we may assume that \tilde{g} is not the identity and hence, since \mathbf{A}'' is NSS, that it has positive escape norm. Suppose, by way of contradiction, that $\|g\|_{e,\mathbf{A}''/\mathbf{P}} = o(\|\tilde{g}\|_{e,\mathbf{A}''})$. Our goal will be to reach a contradiction by finding another lift of g with strictly smaller escape norm than \tilde{g} .

Set $M_1'' := 1/\|\tilde{g}\|_{e,\mathbf{A}''} \in {}^*\mathbf{N}$.

We now make an important deduction from our hypothesis. For every $n \in {}^*\mathbf{N}$ such that $n = O(M_1'')$, we have $g^n \in \mathbf{A}''/\mathbf{P}$. In particular, for every (standard) integer $k \in \mathbf{N}$, $g^{kM_1''} \in \mathbf{A}''/\mathbf{P}$. This implies that the group generated by $g^{M_1''}$ lies in \mathbf{A}''/\mathbf{P} . However, in projection to the Lie model, \mathbf{A}''/\mathbf{P} gets mapped into a neighbourhood of the identity in L' , which we chose small enough so as not to contain any non-trivial subgroup. We thus conclude that $g^{M_1''}$ maps to the identity in L' , and therefore $\tilde{g}^{M_1''}$ maps into the local one-parameter subgroup $\phi([-1, 1])$.

Now there is another element which maps to $\phi([-1, 1])$, namely u , the element for which $\|u\|_{e,\mathbf{A}'}$ is minimal.

In order to motivate the rest of the argument, let us temporarily work in a heuristic setting (using informal notation such as \approx), returning to tighten the argument rigorously later. Since

$$(9.5) \quad \mathbf{A}'' \approx \pi^{-1}(\exp(\delta B)),$$

and since M_1'' is the least n for which \tilde{g}^n escapes \mathbf{A}'' , we have

$$(9.6) \quad \pi(\tilde{g}^{M_1''}) \approx \phi(\delta) = \exp(\delta X).$$

Similarly

$$(9.7) \quad \pi(u^{M_1''}) \approx \phi(\delta).$$

Now u^n takes at least as long as \tilde{g}^n to escape from $\mathbf{A}' \approx \pi^{-1}(\exp \mathbf{B})$. Hence (roughly) it takes at least as long to escape from $\mathbf{A}'' \approx \pi^{-1}(\exp \delta \mathbf{B})$ as well, which means that

$$(9.8) \quad M_1'' \lesssim M_0''.$$

We are trying to find a lift of g with smaller escape norm than that of \tilde{g} . To do this it is sensible to look for elements of the form $h := \tilde{g}u^{-m}$, $m \in {}^*\mathbf{N}$. Provided that m is chosen judiciously, h will also be a lift of g since (by definition) u lies in \mathbf{P} . Since, measured in $\phi([-1, 1])$ by applying π , the element u is “shorter” than \tilde{g} , it seems reasonable that by an appropriate choice of m we can make h shorter than \tilde{g} as well.

Being a little more precise, suppose that $m, n \in {}^*\mathbf{N}$. Since u is central in \mathbf{A}'' we have $h^n = \tilde{g}^n u^{-mn}$ whenever these expressions are well-defined, and hence $\pi(h^n) = \pi(\tilde{g}^n)\pi(u^{-mn})$. From (9.6) and (9.7) we have

$$(9.9) \quad \pi(\tilde{g}^n) \approx \phi(\text{st}(\delta n/M_1'')) \quad \text{and} \quad \pi(\tilde{u}^{-mn}) \approx \phi(\text{st}(-\delta mn/M_0'')).$$

These expressions will be legitimate if m, n are chosen so that the arguments of the ϕ 's always lie in $[-50, 50]$ (say). It follows that

$$(9.10) \quad \pi(h^n) \approx \phi\left(\text{st}\left(\frac{1}{M_1''} - \frac{m}{M_0''}\right)\delta n\right).$$

However (by the Euclidean algorithm) there is a choice of $m \in {}^*\mathbf{N}$ such that $|1/M_1'' - m/M_0''| \leq 1/2M_0''$. Comparing with (9.10) we see that for $n = 1, \dots, 2M_0''$ we have $\pi(h^n) = \phi(\delta')$ with $\delta' \leq \delta$. Since $\pi^{-1}(\phi([0, \delta])) \subseteq \mathbf{A}''$, we must raise h to at least the power $2M_0''$ before it escapes \mathbf{A}'' . Since $2M_0'' > M_0'' \gtrsim M_1''$, this h is a lift of g with smaller escape norm than \tilde{g} . Note that the computations (9.10) are legitimate for this choice of m and for $n \leq 2M_0''$.

We now perform the above argument rigorously. Instead of the heuristic statement (9.5), we must work with the inclusions

$$(9.11) \quad \pi^{-1}(\exp(\delta \mathbf{B})) \subseteq \mathbf{A}'' \subseteq \pi^{-1}(\exp(1.001\delta \mathbf{B})).$$

To get a precise form of (9.6), note that by definition of the escape norm we have $\tilde{g}^{M_1''-1} \in \mathbf{A}''$, whilst $\tilde{g}^{M_1''} \notin \mathbf{A}''$. In particular, as a consequence of (9.11), $\pi(\tilde{g}^{M_1''-1}) \in \exp(1.001\delta \mathbf{B})$, whilst $\pi(\tilde{g}^{M_1''}) \notin \exp(\delta \mathbf{B})$. Since M_1'' is unbounded, the first of these actually implies that $\pi(\tilde{g}^{M_1''}) \in \exp(1.001\delta \mathbf{B})$.

Similarly $\pi(u^{M_0''-1}) \in \exp(1.01\delta \mathbf{B})$, whilst $\pi(u^{M_0''}) \notin \exp(\delta \mathbf{B})$. Once again, the first of these implies that $\pi(u^{M_0''}) \in \exp(1.001\delta \mathbf{B})$.

Since \mathbf{B} is convex, comparison of these facts shows that $\pi(\tilde{g}^{M_1''}) = \phi(t)$ and $\pi(u^{M_0''}) = \phi(t')$ with

$$(9.12) \quad t, t' \in [0.9\delta, 1.1\delta].$$

Suppose that M'_1 and M'_0 are the escape times of \tilde{g} and u from \mathbf{A}' , respectively. Since $u \in \mathbf{A}'$ was assumed to have minimal escape norm, $M'_1 \leq M'_0$. On the other hand (9.11) implies that $M''_0/M'_0, M''_1/M'_1 \in [0.99\delta, 1.01\delta]$, and so

$$(9.13) \quad M''_1 \leq 1.1M''_0.$$

As in the heuristic discussion above, take $h := \tilde{g}u^{-m}$, for some $m \in {}^*\mathbf{N}$. Let $n \in {}^*\mathbf{N}$. Then we have

$$\pi(\tilde{g}^n) = \phi(\text{st}(tn/M''_1)) \quad \text{and} \quad \pi(u^{-mn}) = \phi(-\text{st}(t'mn/M''_0))$$

provided that the arguments of the ϕ 's are in $[-50, 50]$, which will always be the case later on in the argument. Since u is central we have

$$(9.14) \quad \pi(h^n) = \phi\left(\text{st}\left(\frac{t}{\delta M''_1} - \frac{m t'}{\delta M''_0}\right)\delta n\right).$$

Roughly as before, we use the Euclidean algorithm to find $m \in {}^*\mathbf{N}$ such that $|1/M''_1 - m t'/\delta M''_0| \leq t'/2\delta M''_0$. By (9.12) and (9.13) it follows that

$$\left|\frac{t}{\delta M''_1} - \frac{m t'}{\delta M''_0}\right| \leq \frac{t'}{2\delta M''_0} + \left(1 - \frac{t}{\delta}\right)\frac{1}{M''_1} < \frac{0.9}{M''_1}.$$

It follows from this and (9.14) that $\pi(h^n) \in \phi([0, \delta])$ for $n \leq M''_1$, and hence h^n lies in \mathbf{A}'' for these same values of n . As a consequence, h has smaller $\|\cdot\|_{e, \mathbf{A}''}$ escape norm than \tilde{g} , contrary to assumption. \square

Finally we prove item (iii) of Lemma 9.4. Suppose then that \overline{Q} is a nondegenerate large ultra nilprogression in \mathbf{A}''/\mathbf{P} in normal form; we wish to lift this to a large ultra nilprogression Q in \mathbf{A}'' of at most one higher rank and step, while preserving the nondegeneracy and normal form properties. The main difficulty is that if one lifts the generators of \overline{Q} arbitrarily then there is no guarantee that the progression they generate, or even a significant part of it, will be contained in \mathbf{A}'' . The key to ensuring that we do achieve this lies in making judicious use of the lifting lemma (Lemma 9.5) and the product and commutator properties of the escape norm (Proposition 8.5(iii) (b) and (c)). At this point we advise the reader to quickly review Definition 2.6 and Appendix C, where nilprogressions in C-normal form are discussed.

We may write the non-degenerate ultra nilprogression in normal form as

$$\overline{Q} = \mathbf{P}(\overline{u}_1, \dots, \overline{u}_r; N_1, \dots, N_r),$$

where the \overline{u}_i are in \mathbf{A}''/\mathbf{P} , the $N_i \in {}^*\mathbf{N}$ are unbounded, and r is the rank of \overline{Q} , and some standard step s . From the normal form hypothesis (and taking ultraproducts), we have the following properties:

(i) (Upper-triangular form) For every $1 \leq i < j \leq r$ and $\epsilon_i, \epsilon_j \in \{-1, +1\}$, one has

$$(9.15) \quad [u_i^{\epsilon_i}, u_j^{\epsilon_j}] \in P\left(u_{j+1}, \dots, u_r; O\left(\frac{N_{j+1}}{N_i N_j}\right), \dots, O\left(\frac{N_r}{N_i N_j}\right)\right).$$

(ii) (Local properness) The expressions $u_1^{n_1} \dots u_r^{n_r}$ for nonstandard integers n_1, \dots, n_r with $|n_i| \leq \frac{1}{C} N_i$ for all $1 \leq i \leq r$ are all well-defined and distinct, if C is a sufficiently large standard real.

(iii) (Volume bound) One has

$$N_1 \dots N_r \ll |\mathbf{Q}| \ll N_1 \dots N_r.$$

(Note that as the N_i are unbounded, $2\lfloor N_i \rfloor + 1$ and N_i are comparable.)

Also, since $u_i^{n_i} \in \overline{\mathbf{Q}} \subseteq \mathbf{A}''/\mathbf{P}$ for all $1 \leq i \leq r$ and $|n_i| \leq N_i$, we have

$$\|u_i\|_{e, \mathbf{A}''/\mathbf{P}} \leq \frac{1}{N_i}.$$

By Lemma 9.5, we may find lifts $u_i \in \mathbf{A}''$ which project to \bar{u}_i in the quotient local group \mathbf{A}''/\mathbf{P} , and are such that

$$\|u_i\|_{e, \mathbf{A}''} = O(\|\bar{u}_i\|_{e, \mathbf{A}''/\mathbf{P}})$$

and thus

$$(9.16) \quad \|u_i\|_{e, \mathbf{A}''} \ll \frac{1}{N_i}.$$

In order to include \mathbf{P} in the lifted progression, we set $u_{r+1} := u$, the generator of \mathbf{P} , and $N_{r+1} := 1/\|u\|_{e, \mathbf{A}''}$. From (9.4) we see that

$$(9.17) \quad M'_0 \ll N_{r+1} \ll M'_0.$$

We then define

$$\mathbf{Q} := P(u_1, \dots, u_r, u_{r+1}; \varepsilon N_1, \dots, \varepsilon N_{r+1})$$

for some sufficiently small standard $\varepsilon > 0$. We claim that \mathbf{Q} is well-defined in \mathbf{A}'' as a nondegenerate ultra nilprogression in normal form, of rank $(r+1)$ and step at most $s+1$.

We begin with the claim that \mathbf{Q} is well-defined in \mathbf{A}'' . From (9.16) and Proposition 8.5 one has

$$\|g\|_{e, \mathbf{A}''} \ll \varepsilon$$

for all $g \in \mathbf{Q}$, and in particular every product in \mathbf{Q} lies in \mathbf{A}'' as required.

It is clear that \mathbf{Q} is a nondegenerate ultra non-commutative progression of rank $(r + 1)$. To show that it is a nilprogression of step at most $s + 1$, it suffices to show that any iterated commutator g of length $s + 2$ in the generators $u_1^{\pm 1}, \dots, u_{r+1}^{\pm 1}$ is trivial. Using commutator identities such as the Hall-Witt identity

$$(9.18) \quad [z, [x, y]] = [[y^{-1}, z^{-1}], x]^{\bar{y}} [[z, x^{-1}], y^{-1}]^{\bar{y}}$$

where $x^{\bar{y}} := y^{-1}xy$ (using the unbounded nature of the N_i to justify all operations) we may restrict attention to iterated commutators g of the form $g = [h, u_i^{\pm 1}]$ where h is an iterated commutator of length $s + 1$ and $1 \leq i \leq r + 1$. But by projecting down to \mathbf{A}''/\mathbf{P} , we know that the image of h vanishes and thus $h \in \mathbf{P}$. Since \mathbf{P} is central in \mathbf{A}'' , the claim follows.

Finally, we need to show that \mathbf{Q} is in normal form. We begin by establishing the upper triangular form (2.1), i.e. that

$$[u_i^{\epsilon_i}, u_j^{\epsilon_j}] \in \mathbf{P} \left(u_{j+1}, \dots, u_r; \mathbf{O} \left(\frac{N_{j+1}}{N_i N_j} \right), \dots, \mathbf{O} \left(\frac{N_r}{N_i N_j} \right) \right)$$

whenever $1 \leq i < j \leq r + 1$ and $\epsilon_i, \epsilon_j \in \{-1, +1\}$.

If $j = r + 1$, then $u_j = u$ commutes with every element of \mathbf{A}'' , and in particular with u_i , so the claim follows in this case. Now suppose that $j \leq r$. From (9.15) we then have

$$[\bar{u}_i^{\epsilon_i}, \bar{u}_j^{\epsilon_j}] \in \mathbf{P} \left(\bar{u}_{j+1}, \dots, \bar{u}_r; \mathbf{O} \left(\frac{N_{j+1}}{N_i N_j} \right), \dots, \mathbf{O} \left(\frac{N_r}{N_i N_j} \right) \right)$$

which lifts to

$$[u_i^{\epsilon_i}, u_j^{\epsilon_j}] \in \mathbf{P} \left(u_{j+1}, \dots, u_r; \mathbf{O} \left(\frac{N_{j+1}}{N_i N_j} \right), \dots, \mathbf{O} \left(\frac{N_r}{N_i N_j} \right) \right) \cdot \mathbf{P}.$$

Thus we may write $[u_i^{\epsilon_i}, u_j^{\epsilon_j}] = gu^n$, where

$$g \in \mathbf{P} \left(u_{j+1}, \dots, u_r; \mathbf{O} \left(\frac{N_{j+1}}{N_i N_j} \right), \dots, \mathbf{O} \left(\frac{N_r}{N_i N_j} \right) \right)$$

and $n \leq 1/\|u\|_{e, \mathbf{A}'} = M'_0$. From (9.16) and Proposition 8.5 one has

$$\|[u_i^{\epsilon_i}, u_j^{\epsilon_j}]\|_{e, \mathbf{A}''} \ll \frac{1}{N_i N_j},$$

and therefore

$$\|g\|_{e, \mathbf{A}''} \ll \frac{1}{N_i N_j}$$

and hence

$$(9.19) \quad \|u^n\|_{e, \mathbf{A}''} \ll \frac{1}{N_i N_j}.$$

In particular, $\|u^n\|_{e, \mathbf{A}''}$ is infinitesimal, which implies that $\pi(u^n) = \text{id}$ and hence $n = o(M'_0)$ by (9.4). Since $\|u\|_{e, \mathbf{A}''} = 1/N_{r+1}$, we conclude that

$$|n| \ll \frac{N_{r+1}}{N_i N_j}$$

and thus

$$[u_i^{\varepsilon_i}, u_j^{\varepsilon_j}] \in \mathbf{P}\left(u_{j+1}, \dots, u_{r+1}; \mathbf{O}\left(\frac{N_{j+1}}{N_i N_j}\right), \dots, \mathbf{O}\left(\frac{N_{r+1}}{N_i N_j}\right)\right).$$

Noting that $\varepsilon > 0$ is standard and thus can be absorbed into the $\mathbf{O}()$ notation, this gives the desired upper triangular property.

Next, we establish the local properness. Suppose that

$$u_1^{n_1} \dots u_{r+1}^{n_{r+1}} = u_1^{n'_1} \dots u_{r+1}^{n'_{r+1}}$$

for some $|n_i|, |n'_i| \leq \varepsilon N_i$. Quotienting by \mathbf{P} , we conclude that

$$\overline{u_1^{n_1}} \dots \overline{u_r^{n_r}} = \overline{u_1^{n'_1}} \dots \overline{u_r^{n'_r}}.$$

This quotienting can be justified because all products here lie in \mathbf{Q} and hence in \mathbf{A}'' . By the local properness of $\overline{\mathbf{Q}}$, we conclude if ε is small enough that $n_i = n'_i$ for all $1 \leq i \leq r$; we may then cancel and conclude that

$$u^{n_{r+1} - n'_{r+1}} = \text{id}.$$

Since $\|u\|_{e, \mathbf{A}''} = 1/N_{r+1}$ and $|n_{r+1} - n'_{r+1}| < N_{r+1}$, this implies that $n_{r+1} = n'_{r+1}$, giving the desired local properness.

From local properness one immediately has the lower bound

$$|\mathbf{Q}| \gg N_1 \dots N_r N_{r+1}.$$

Now we establish the matching upper bound

$$|\mathbf{Q}| \ll N_1 \dots N_r N_{r+1}.$$

We first recall from the normal form of $\overline{\mathbf{Q}}$ that

$$|\overline{\mathbf{Q}}| \ll N_1 \dots N_r.$$

From construction it is also clear that the image of \mathbf{Q} under projection by \mathbf{P} lies in $\overline{\mathbf{Q}}$. It therefore suffices to show that the preimage of any element in $\overline{\mathbf{Q}}$ contains at most $\mathbf{O}(N_r)$ elements of \mathbf{Q} . By construction of the quotient map, we see that the preimage is contained in a translate of \mathbf{P} , and thus has cardinality $\mathbf{O}(M'_0)$; the claim then follows from (9.17). This concludes the proof of Lemma 9.4 and thus Theorem 9.3. \square

We can now conclude the proof of Theorem 2.10, the most basic form of our main theorem.

Proof of the first part of Theorem 2.10. — We argue by contradiction. Negating the quantifiers, we see that there exists some $K \geq 1$ and an infinite sequence of local groups G_n and finite K -approximate groups $A_n \subseteq G_n$, $n \in \mathbf{N}$, for which the conclusion of the theorem fails, namely for which A_n^4 does not contain any coset nilprogression of rank and step at most n in n -normal form and of cardinality at least $\frac{1}{n}|A_n|$.

Now form the ultraproduct $\mathbf{A} = \prod_{n \rightarrow \alpha} A_n$ inside $\mathbf{G} = \prod_{n \rightarrow \alpha} G_n$. By Łos's Theorem (Theorem A.6), \mathbf{G} is a local group and \mathbf{A} an ultra approximate subgroup. We can now apply Theorem 4.2, whose proof we just completed, to conclude that \mathbf{A}^4 contains an ultra coset nilprogression P in normal form with $|P| \gg |A|$. Using Łos's theorem again we conclude that $P = \prod_{n \rightarrow \alpha} P_n$, where for an α -large set of n , P_n is a $1/c$ -proper coset nilprogression contained A_n^4 of rank and step at most $1/c$ and of size at least $c|A_n|$ for some standard positive number $c > 0$. But this contradicts the construction of the A_n , thereby yielding the claim. \square

To conclude this section we record another useful conclusion from the above analysis: Hrushovski's Lie model is nilpotent.

Proposition 9.6 (Nonstandard finite approximate groups have nilpotent Lie models). — *Suppose that \mathbf{A} is an ultra approximate group and that $\pi : \mathbf{A}^{\delta} \rightarrow \mathbf{L}$ is a good model for \mathbf{A} into a connected Lie group \mathbf{L} with Lie algebra \mathfrak{L} . Then \mathfrak{L} and \mathbf{L} are nilpotent.*

Proof. — By Proposition 8.5 we may find a large strong ultra approximate subgroup \mathbf{A}' of \mathbf{A} obeying the conclusion of that proposition. By quotienting out the elements H of \mathbf{A}' of zero escape norm as in the proof of Proposition 9.2, we obtain an NSS ultra approximate group \mathbf{A}'/H . Now one runs the argument in Theorem 9.3. An inspection of this argument shows that if one unfolds the induction from Lemma 9.4, the Lie algebra \mathfrak{L} of \mathbf{L} is repeatedly quotiented out by central algebras until it becomes trivial. Thus, \mathfrak{L} can be obtained from the trivial Lie algebra by a finite tower of central extensions and is therefore nilpotent as required. The nilpotence of \mathbf{L} is an immediate consequence of this and basic Lie theory. \square

10. A dimension bound

In this section we prove Theorem 2.12, in which it is shown that the rank of the nilprogression P in the main theorem may be taken to be $O(K^{O(1)})$. We will also show that so long as we work in a global group G , and replace A^4 with $A^{O_K(1)}$, it may be taken to be $O(\log K)$. By the usual ultraproduct argument, it will suffice to establish the following nonstandard analysis formulation of the theorem.

Theorem 10.1. — *Suppose that \mathbf{A} is an ultra global \mathbf{K} -approximate group, thus $\mathbf{A} = \prod_{n \rightarrow \alpha} A_n$ for some finite \mathbf{K} -approximate groups, each contained in a global group G_n . Then \mathbf{A}^{12} contains an ultra coset nilprogression P in normal form with $|P| \gg |A|$ and rank at most $O(\mathbf{K}^2 \log \mathbf{K})$. Moreover, there exists a standard natural number m such that \mathbf{A}^m contains an ultra coset nilprogression P in normal form with $|P| \gg |\mathbf{A}|$ and rank at most $6 \log_2 \mathbf{K}$.*

Recall that the step of a nilprogression in normal form is always less or equal to its rank. The derivation of Theorem 2.12 from Theorem 10.1 proceeds analogously to the derivation of Theorem 2.10 from Theorem 4.2 and is omitted.

It remains to establish Theorem 10.1. The arguments here are inspired by some remarks of Hrushovski in [33, §4]. In particular, a key tool will be the following lemma from [33, Lemma 4.9].

Lemma 10.2 (*Doubling in a simply connected nilpotent Lie group*). — *Let G be a connected, simply connected nilpotent Lie group of dimension d , and let A be a measurable subset of G . Let μ be a Haar measure on G (note that nilpotent groups are automatically unimodular, and so there is no distinction between left and right Haar measure). Then $\mu(A^2) \geq 2^d \mu(A)$.*

Proof. — We use an argument of Gelfand from [33, Lemma 4.9]. As is well known (e.g. see [12]), in a simply connected nilpotent Lie group, the exponential map $\exp : \mathfrak{g} \rightarrow G$ is a diffeomorphism, which pushes forward the Lebesgue measure $\mu_{\mathfrak{g}}$ on the d -dimensional vector space \mathfrak{g} , the Lie algebra of G , to the Haar measure μ on G . Thus it will suffice to show that $\mu_{\mathfrak{g}}(\log(A^2)) \geq 2^d \mu_{\mathfrak{g}}(\log A)$, where \log is the inverse of \exp . But as A^2 contains $\{a^2 : a \in A\}$, $\log(A^2)$ contains the dilate $2 \cdot \log A$ of $\log A$, and the claim follows. \square

One is tempted to combine this theorem with the Hrushovski Lie Model Theorem directly (i.e. Theorem 3.10), to get some dimensional control on the Lie group L . However, there is a technical obstruction; the Lie model is only available for an ultra approximate subgroup \mathbf{A}' of \mathbf{A} , and the covering parameter \mathbf{K}' of this subgroup \mathbf{A}' may be much worse than the covering parameter \mathbf{K} of the original ultra approximate subgroup \mathbf{A} . To get around this problem, we need to choose the subgroup \mathbf{A}' more carefully. A clue as to how to proceed is provided by the following basic observation (cf. [31, Lemma 7.3]).

Lemma 10.3 (*Slicing approximate groups by genuine subgroups*). — *Let A be a (possibly infinite) \mathbf{K} -approximate group in a global group G , and let G' be a genuine subgroup of G . Then $A' := A^2 \cap G'$ is a \mathbf{K}^3 -approximate subgroup and $A^4 \cap G'$ can be covered by at most \mathbf{K}^3 left translates of A' .*

Proof. — Since $(A')^2 \subseteq A^4 \cap G'$, it suffices to show that $A^4 \cap G'$ can be covered by \mathbf{K}^3 left-translates of A' . But A^4 can be covered by \mathbf{K}^3 left-translates of A since A is a \mathbf{K} -approximate group. Next, observe that if a left-translate gA of A intersects $A^4 \cap G'$ in

at least one point g' , then

$$gA \cap A^4 \cap G' \subseteq gA \cap G' \subseteq g'A'.$$

Thus $A^4 \cap G'$ can be covered by K^3 left-translates of A' , as required. \square

Lemma 10.3 suggests that we should look for Lie models of approximate groups \mathbf{A}' that are formed by slicing \mathbf{A}^2 with a genuine subgroup G' of G .

We turn to the details. Let A_n be a sequence of K -approximate groups in global groups G_n , and let $\mathbf{A} = \prod_{n \rightarrow \alpha} A_n$ be their ultraproduct; thus \mathbf{A} is a ultra K -approximate group that lies inside an ultra genuine group $\prod_{n \rightarrow \alpha} G_n$. By Proposition 6.10, we may find a model $\pi : \langle \mathbf{A} \rangle \rightarrow G$ of \mathbf{A}^4 by a locally compact group G .

Let U_0 be the neighbourhood in Definition 3.5. We have $\pi^{-1}(U_0) \subseteq \mathbf{A}^4$ and $U_0 \subseteq \pi(\mathbf{A}^4)$. By Theorem B.17, there is an open subgroup G' of G and a closed subgroup H of G contained in U_0 and normalized by G' such that $L := G'/H$ is a connected Lie group. Let $U_1 \subseteq G'$ be an open subset such that $H \subseteq U_1 \subseteq U_1^2 \subseteq U_0$ and let $\phi : G' \rightarrow L$ denote the quotient map.

Now set $\mathbf{A}' := \mathbf{A}^4 \cap \pi^{-1}(G')$. From Lemma 10.3 applied to \mathbf{A}^2 , we see that \mathbf{A}' is a K^6 -approximate group. We now also claim that \mathbf{A}' is a nonstandard finite set, which would make \mathbf{A}' an ultra K^6 -approximate group. To see this, observe first from Definition 3.5(ii) that $\pi(\mathbf{A}^4)$ is contained in some compact set F . As G' is an open subgroup of G , it is also closed and $F \cap G'$ is compact. We then see from Definition 3.5(iii) that we can find a nonstandard finite set \mathbf{A}_* such that $\pi^{-1}(F \cap G') \subseteq \mathbf{A}_* \subseteq \pi^{-1}(G')$. Thus $\mathbf{A}' = \mathbf{A}^4 \cap \mathbf{A}_*$, and so \mathbf{A} is a nonstandard finite set as required.

Note that $\pi(\mathbf{A}')$ contains the open set $U_0 \cap G'$ and is itself contained in a compact subset of G' . Hence the set $E := \phi \circ \pi(\mathbf{A}')$ is precompact and contains a neighbourhood of the identity in L . Moreover $(\phi \circ \pi)^{-1}(\phi(U_1)) \subseteq \mathbf{A}'$, hence it follows that $\phi \circ \pi : \langle \mathbf{A}' \rangle \rightarrow L$ is a good model for \mathbf{A}' . From Lemma 9.6, we conclude that L is nilpotent. Every connected nilpotent Lie group admits a unique maximal compact subgroup which, moreover, is central. Let N be the maximal compact subgroup of L and $\theta : L \rightarrow L/N$ be the quotient map.

We claim that $\dim(L/N) \leq 6 \log_2 K$. To see this note that, as \mathbf{A}' is a K^6 -approximate group, we see that E^2 is covered by at most K^6 left-translates of E . Therefore $\theta(E)^2$ can be covered by at most K^6 left-translates of $\theta(E)$, and hence $\overline{\theta(E)^2}$ can be covered by at most K^6 translates of $\overline{\theta(E)}$, where $\overline{\theta(E)}$ is the topological closure of $\theta(E)$, a compact set with non-empty interior. If we let μ be a Haar measure on L/N , it follows that

$$\mu(\overline{\theta(E)^2}) \leq K^6 \mu(\overline{\theta(E)}).$$

On the other hand, from Lemma 10.2 one has

$$\mu(\overline{\theta(E)^2}) \geq 2^{\dim(L/N)} \mu(\overline{\theta(E)}).$$

Since $\overline{\theta(\mathbf{E})}$ has non-empty interior, $\mu(\overline{\theta(\mathbf{E})}) \neq 0$ and so comparison of these two inequalities implies that

$$(10.1) \quad \dim(L/N) \leq 6 \log_2 K.$$

We now explain how to derive the second part of Theorem 10.1 from the above; we will turn to the first part later. We consider $\phi^{-1}(N)$, which is the kernel of the projection map from G' to $L/N (= (G'/H)/N)$. This is a compact subgroup of G' . Since $\pi(\mathbf{A}')$ contains an open neighbourhood of the identity, we conclude that there exists a standard natural number m such that $\phi^{-1}(N) \subseteq \pi(\mathbf{A}'^{m-1})$, and this implies that \mathbf{A}'^m contains the kernel of $\theta \circ \phi \circ \pi$, which implies that $\theta \circ \phi \circ \pi : \mathbf{A}'^{3m} \rightarrow L/N$ is a good model. By Proposition 9.2 we conclude that \mathbf{A}'^{4m} contains a large approximate subgroup \mathbf{A}'' with $(\mathbf{A}'')^{1000}$ well-defined and contained in \mathbf{A}'^{4m} , and a global internal subgroup H' of \mathbf{A}'' such that \mathbf{A}''/H' is an NSS approximate subgroup with a connected Lie group as a good model. An inspection of the proof of that proposition reveals that we may take $(\mathbf{A}'')^{1000}$ inside \mathbf{A}^m and that we may take the connected Lie group to be L/N , which we have shown to have dimension at most $6 \log_2 K$. Applying Theorem 9.3, we see that $(\mathbf{A}''/H')^4$ contains a large ultra nilprogression in normal form of rank at most $6 \log_2 K$, and thus $(\mathbf{A}'')^4$ contains a large ultra coset nilprogression in normal form of rank at most $6 \log_2 K$. As $(\mathbf{A}'')^4$ is contained in \mathbf{A}'^m , which is in turn contained in \mathbf{A}^{4m} , the second part of Theorem 10.1 follows (after redefining m).

We now turn to the first part of Theorem 10.1. As we see from the last paragraph, the difficulty here is that $\phi^{-1}(N)$ may not be contained in $\pi(\mathbf{A}')$. We will show that nevertheless $\pi(\mathbf{A}')$ still contains a subgroup $\phi^{-1}(N_0)$, where N_0 is a closed subgroup of N with small codimension. For this the key is the following lemma, which is potentially of interest its own right. Here, and below, we write $\mathbf{T}^d := \mathbf{R}^d/\mathbf{Z}^d$ for the d -dimensional torus. By a *subtorus* we mean a closed connected subgroup of \mathbf{T}^d .

Lemma 10.4. — *Let $K, d \geq 1$ and A be a closed K -approximate group in \mathbf{T}^d containing a neighbourhood of 0. That is, A is closed, contains a neighbourhood of 0, is centrally symmetric, and there is a finite set $X \subseteq \mathbf{T}^d$, $|X| \leq K$, such that $A + A \subseteq A + X$. Then $4A := A + A + A + A$ contains a subtorus $T \subseteq \mathbf{T}^d$ of codimension at most $O(K^2 \log K)$.*

Before proving Lemma 10.4, we explain how to conclude the proof of Theorem 10.1 with Lemma 10.4 in hand. First we observe that setting $\mathbf{A}'_1 := \mathbf{A}^2 \cap \pi^{-1}(G')$, $\pi(\mathbf{A}'_1)$ is a neighbourhood of id in G . Indeed \mathbf{A}^4 is and $\mathbf{A}^4 \subseteq X\mathbf{A}$ for some finite X , so that $\pi(\mathbf{A})$ has non-empty interior, and hence $\pi(\mathbf{A}^2)$ is a neighborhood of id . We now apply the lemma to $A = \phi \circ \pi(\mathbf{A}'_1)$, and conclude that $\phi^{-1}(T) \subseteq \pi(\mathbf{A}'_1^4) \subseteq \pi(\mathbf{A}^2)$. Writing $\theta' : L \rightarrow L/T$ for the projection map, we see that \mathbf{A}'^3 contains the kernel of $\theta' \circ \phi \circ \pi$ and this implies that \mathbf{A}'^3 admits the connected nilpotent Lie group L/T as a good model. Moreover $\dim L/T = \dim L/\mathbf{T}^d + \dim \mathbf{T}^d/T = O(K^2 \log K)$ by (10.1) and by Lemma 10.4. The rest of the proof is then identical to the previous case: by Proposition

9.2 and it's proof we conclude that \mathbf{A}^3 contains a large approximate subgroup \mathbf{A}'' with $(\mathbf{A}'')^{1000}$ well-defined and contained in \mathbf{A}^3 , and a global internal subgroup H' of \mathbf{A}'' such that \mathbf{A}''/H' is an NSS ultra approximate subgroup admitting L/N_0 as a good model. By Theorem 9.3, we see that $(\mathbf{A}''/H')^4$ contains a large ultra nilprogression in normal form of rank at most $O(K^2 \log K)$, and thus $(\mathbf{A}'')^4$ contains a large ultra coset nilprogression in normal form with rank at most $O(K^2 \log K)$. As $(\mathbf{A}'')^4 \subseteq \mathbf{A}^3 \subseteq \mathbf{A}^{12}$, the first part of Theorem 10.1 follows.

We now turn to the proof of Lemma 10.4. Let μ be the normalized Haar measure on \mathbf{T}^d . Note that the group $\widehat{\mathbf{T}}^d$ of characters of \mathbf{T}^d identifies with \mathbf{Z}^d . Our main tool is the notion of the (α) -large spectrum of an additive set A , defined by

$$\text{Spec}_\alpha(A) := \{\xi \in \mathbf{Z}^d : |\widehat{1_A}(\xi)| \geq \alpha \mu(A)\}.$$

See [54, Definition 4.34] for this definition and a further discussion.

If $S \subseteq \mathbf{Z}^d$ is a set of characters, we write

$$S^\perp := \bigcap_{\xi \in S} \ker \xi.$$

Note that S^\perp is a closed subgroup of \mathbf{T}^d with codimension the rank of the subgroup of \mathbf{Z}^d generated by S .

To prove Lemma 10.4 we first reduce to the case in which $\mu(A)$ is somewhat large by establishing the following lemma.

Lemma 10.5. — *Suppose that $A \subseteq \mathbf{T}^d$ is a closed K -approximate group containing a neighbourhood of 0. Then there is a subtorus $T_0 \subseteq \mathbf{T}^d$ with $\dim(T_0) \geq d - O(\log K)$ and some $x_0 \in \mathbf{T}^d$ such that, writing μ_0 for the Haar measure on T_0 , we have $\mu_0((A + x_0) \cap T_0) \gg e^{-O(K \log K)}$.*

Then we handle the case in which $\mu(A)$ is somewhat large by proving the following, which is a straightforward continuous analogue of the so-called Bogolyubov-Chang lemma [10].

Lemma 10.6. — *Suppose that $A \subseteq \mathbf{T}^d$ is measurable, that $\mu(A) \geq \alpha$, and that $\mu(2A) \leq K\mu(A)$. Then $2A - 2A$ contains a subtorus $T \subseteq \mathbf{T}^d$ of codimension at most $O(K \log(1/\alpha))$.*

To deduce Lemma 10.4 from Lemmas 10.5 and 10.6, we proceed as follows. Locate an x_0 as in Lemma 10.5, and suppose furthermore that for this x_0 the measure of $(A + x_0) \cap T_0$ is close to maximal in the sense that

$$(10.2) \quad \mu_0((A + x_0) \cap T_0) \geq \frac{1}{2} \mu_0((A + x) \cap T_0)$$

for all $x \in \mathbf{T}^d$. Set $A_1 := (A + x_0) \cap T_0$. Then, since $A + A \subseteq A + X$, we have

$$A_1 + A_1 \subseteq (A + X + 2x_0) \cap T_0.$$

By (10.2) it follows that $\mu(2A_1) \leq 2K\mu(A_1)$. We are now in a position to apply Lemma 10.6 to A_1 , with $\alpha = e^{-O(K \log K)}$. We conclude that there is a further subtorus $T \subseteq T_0$ of codimension $O(K^2 \log K)$ inside $2A_1 - 2A_1$. Since $2A_1 - 2A_1 \subseteq 4A$, this concludes the proof of Lemma 10.4.

For the proofs of both Lemmas 10.5 and 10.6 we will require the following lemma of Bogolyubov type.

Lemma 10.7 (Bogolyubov-type lemma). — *Let $A \subseteq \mathbf{T}^d$ have positive measure and let $k \geq 2$ be a natural number. Suppose that*

$$\delta \leq (\mu(A)/2\mu(kA))^{1/(2k-2)}.$$

Then $kA - kA$ contains $(\text{Spec}_\delta(A))^\perp$.

Proof. — It suffices (in fact, it is equivalent) to show that if $x \in (\text{Spec}_\delta(A))^\perp$ then $f(x) > 0$, where $f = 1_A^{*k} * 1_{-A}^{*k} = 1_A * \dots * 1_A * 1_{-A} * \dots * 1_{-A}$ is the convolution of k copies of 1_A and k copies of 1_{-A} . Now by the Fourier inversion formula we have

$$\begin{aligned} (10.3) \quad f(x) &= \sum_{\xi \in \mathbf{Z}^d} |\widehat{1}_A(\xi)|^{2k} \overline{\xi(x)} \geq \sum_{\xi \in \text{Spec}_\delta(A)} |\widehat{1}_A(\xi)|^{2k} - \sum_{\xi \notin \text{Spec}_\delta(A)} |\widehat{1}_A(\xi)|^{2k} \\ &\geq \sum_{\xi \in \mathbf{Z}^d} |\widehat{1}_A(\xi)|^{2k} - 2 \sum_{\xi \notin \text{Spec}_\delta(A)} |\widehat{1}_A(\xi)|^{2k}, \end{aligned}$$

where we have used the fact that $\xi(x) = 1$ if $\xi \in \text{Spec}_\delta(A)$ and $x \in (\text{Spec}_\delta(A))^\perp$. Now Parseval's identity and the Cauchy-Schwarz inequality imply that

$$\begin{aligned} \sum_{\xi \in \mathbf{Z}^d} |\widehat{1}_A(\xi)|^{2k} &= \int_{x \in \mathbf{T}^d} 1_A^{*k}(x)^2 d\mu(x) \geq \frac{1}{\mu(kA)} \left(\int_{x \in \mathbf{T}^d} 1_A^{*k}(x) d\mu(x) \right)^2 \\ &= \frac{\mu(A)^{2k}}{\mu(kA)}. \end{aligned}$$

On the other hand, by a second application of Parseval's identity, we have

$$\sum_{\xi \notin \text{Spec}_\delta(A)} |\widehat{1}_A(\xi)|^{2k} < \delta^{2k-2} \mu(A)^{2k-2} \sum_{\xi \in \mathbf{Z}^d} |\widehat{1}_A(\xi)|^2 = \delta^{2k-2} \mu(A)^{2k-1}.$$

Substituting these inequalities into (10.3) yields

$$f(x) \geq \frac{\mu(A)^{2k}}{\mu(kA)} - 2\delta^{2k-2} \mu(A)^{2k-1} = \frac{\mu(A)^{2k-1}}{\mu(kA)} (\mu(A) - 2\delta^{2k-2} \mu(kA)).$$

The lemma follows immediately. \square

Lemma 10.6 is an immediate consequence of the case $k = 2$ of this lemma and (the continuous variant of) “Chang’s lemma” [10], which is the following statement. For a proof, see [54, Lemma 4.36].

Lemma 10.8 (Chang’s lemma). — Suppose that $\alpha < 1/2$ and that $A \subseteq \mathbf{T}^d$ is a measurable set with $\mu(A) \geq \alpha$. Then $\text{Spec}_\delta(A)$ generates a subgroup of \mathbf{Z}^d of rank at most $O(\delta^{-2} \log(1/\alpha))$.

Proof of Lemma 10.6. — Noting that $\mu(A)/\mu(2A) \leq 1/K$, the lemma follows from Lemma 10.7 with $k = 2$ and $\delta := 1/2\sqrt{K}$ followed by an application of Lemma 10.8. \square

To prove Lemma 10.5 we will apply Lemma 10.7 with a much larger value of k , as well as the following result.

Lemma 10.9. — There is an absolute constant $c > 0$ with the following property. Suppose that $A \subseteq \mathbf{T}^d$ is a closed K -approximate group containing a neighbourhood of 0. Then $\text{Spec}_{1-c/\log K}(A)$ generates a subgroup of \mathbf{Z}^d of rank $O(\log K)$.

Proof. — Let $\varepsilon = c/\log K$, where $c > 0$ is to be chosen later. Suppose that $\xi \in \text{Spec}_{1-\varepsilon}(A)$ and that $\xi \neq 0$. Let $\eta : \mathbf{T}^d \mapsto \mathbf{R}/\mathbf{Z}$ be such that $\xi(x) = e^{2i\pi\eta(x)}$. Then

$$\int_{\mathbf{T}^d} 1_A(x) \xi(x) d\mu(x)$$

is real, since A is symmetric, and at least $(1 - \varepsilon)\mu(A)$. Thus

$$\int_{\mathbf{T}^d} 1_A(x) \cos(2\pi\eta(x)) d\mu(x) \geq (1 - \varepsilon)\mu(A),$$

and hence the (symmetric) subset $A'(\xi) \subseteq A$, consisting of those x for which $\cos(2\pi\eta(x)) \geq 99/100$, has measure $\mu(A'(\xi)) \geq (1 - 100\varepsilon)\mu(A)$. In particular $\|\eta(x)\| < \frac{1}{10}$ whenever $x \in A'(\xi)$, where $\|\theta\| := \inf_{z \in \mathbf{Z}} |\theta - z|$.

Suppose now that $\text{Spec}_{1-\varepsilon}(A)$ contains elements ξ_1, \dots, ξ_m which are linearly independent over \mathbf{Q} , and let η_1, \dots, η_m be the corresponding \mathbf{R}/\mathbf{Z} -valued characters. Consider the set $A'' := \bigcap_{i=1}^m A'(\xi_i)$; provided that $m < 1/100\varepsilon$, this will have $\mu(A'') \geq \frac{1}{2}\mu(A)$. Note that $A'' = -A''$.

Consider now the homomorphism $\psi : \mathbf{T}^d \rightarrow \mathbf{T}^m$ given by $x \mapsto (\eta_1(x), \dots, \eta_m(x))$. The image of A'' under ψ lies in a box of diameter $1/5$.

Now any subset U of $\mathbf{T}^m = (\mathbf{R}/\mathbf{Z})^m$ which lies in a box of diameter $< \frac{1}{2}$ is Freiman 2-isomorphic to an open subset of \mathbf{R}^m , and thus by the abelian case of Gelfander’s Lemma 10.2 (which, in this case, is just a very simple case of the Brunn-Minkowski inequality), we have

$$\mu_m(2U) \geq 2^m \mu_m(U),$$

where μ_m is the normalized Haar measure on $\mathbf{T}^m = (\mathbf{R}/\mathbf{Z})^m$. However, we have $\mu(5A'') \leq \mu(5A) \leq K^4 \mu(A) \leq 2K^4 \mu(A'')$, and an application of the Ruzsa covering lemma (here our Lemma 5.2) shows that $2A''$ is a $4K^4$ -approximate group, and consequently so is $U := \psi(2A'')$.

Therefore, noting that $\mu_m(U) \neq 0$ since $\mu(A'') \geq \frac{1}{2}\mu(A) > 0$, we obtain

$$2^m \leq (4K^4)^2 = (2K)^8,$$

a contradiction if $m > 8 \log_2 2K$. Such a choice of m is acceptable if $\varepsilon < c/\log K$ with c sufficiently small, and so we are forced to conclude that ξ_1, \dots, ξ_m cannot exist. The lemma follows. \square

Proof of Lemma 10.5. — Note that $kA \subseteq (k-1)X + A$, and that $|mX| \leq m^{|X|} = m^K$ for all natural numbers m . It follows that $\mu(kA) \leq k^K \mu(A)$, and so Lemma 10.7 is applicable with $\delta = 1 - c/\log K$ for some $k \leq K^C$. The conclusion is that $2kA$ contains the subtorus $T_0 := (\text{Spec}_{1-c/\log K}(A))^\perp$ which, by Lemma 10.9, has codimension $O(\log K)$. However $2kA$ is covered by at most $(2k)^K = e^{O(K \log K)}$ translates of A , and so one of these translates has $\mu_0(A + x_0) \gg e^{-O(K \log K)}$, which was precisely what we claimed. \square

To conclude this section we record the observation that the above arguments also yield the following more precise version of Proposition 6.12, the weak global Lie model theorem. This builds upon a previous result in this direction by Hrushovski: see [33, Theorem 4.2] and the discussion before [33, Lemma 4.9].

Theorem 10.10 (*Strong global Lie Model Theorem*). — *Suppose that \mathbf{A} is a global ultra K -approximate group. Then there is a large ultra approximate subgroup \mathbf{A}' of \mathbf{A} for some standard $m \geq 1$ which admits a global model $\tilde{\pi} : \langle \mathbf{A}' \rangle \rightarrow \mathbf{L}$ into a connected, simply connected nilpotent Lie group \mathbf{L} of dimension at most $6 \log_2 K$.*

Furthermore, there exists a large ultra approximate group \mathbf{A}'' of \mathbf{A} which admits a global model $\pi' : \langle \mathbf{A}'' \rangle \rightarrow \mathbf{L}'$, a connected nilpotent Lie group, whose maximal (central) compact subgroup \mathbf{N} verifies $\mathbf{L}'/\mathbf{N} \simeq \mathbf{L}$.

11. Applications to growth in groups and geometry

In this section we collect a variety of applications of our main results, in particular proving the various results stated in the introduction.

As an application of his method Hrushovski [33] established the following strengthening of Gromov's theorem on groups with polynomial growth.

Theorem 11.1. — *Let G be a finitely generated group and let $K \geq 1$. Suppose $G = \bigcup_{n \geq 1} A_n$, where A_n is an increasing union of finite subsets of G such that $|A_n^2| \leq K|A_n|$ for all $n \geq 1$. Then G is virtually nilpotent.*

This is indeed a strengthening of Gromov's theorem because if G has polynomial growth with respect to some generating set S then the A_n may be taken to be some subsequence of the word metric balls relative to S .

Unsurprisingly, our main theorem also admits an application of this kind. The following is a corollary of Theorem 2.10 and subsumes Theorem 11.1 above.

Corollary 11.2 (Gromov-type theorem). — *Let $K \geq 1$. Then there is some K' , depending on K , such that the following holds. Assume G is a group generated by a finite symmetric set S containing the identity. Let A be a finite subset of G such that $|A^2| \leq K|A|$ and $S^{K'} \subseteq A$. Then there is a finite normal subgroup $N \triangleleft G$ and a subgroup $G_1 \leq G$ containing N such that*

- (i) G_1 has index $O_K(1)$ in G ;
- (ii) G_1/N has step and rank $O_K(1)$.

In particular G is virtually nilpotent.

Proof of Corollary 11.2. — First we make the following simple observation. Suppose G is a group generated by a finite symmetric set S and let G_0 be a subgroup of index $n = [G : G_0]$. Then for every $k < n$ the ball S^k meets at least $k + 1$ different left cosets of G_0 in G . Indeed if not then by the pigeonhole principle we have $S^i G_0 = S^{i+1} G_0$ for some $i < k$, and so by multiplying on the left with S it follows that $S^k G_0 = S^{k+1} G_0$. Multiplying on the left by further copies of S implies that $S^k G_0 = \langle S \rangle G_0 = G$, and so G_0 has index at most k in G , contrary to assumption.

Now, we apply Corollary 1.7. Thus there exists a subgroup G_0 of G and a normal subgroup H of G_0 such that A may be covered by K' left-translates G_0 for some $K' = O_K(1)$ depending only on K' , and G_0/H is nilpotent of step and rank $O_K(1)$. In particular, G_0 is finite-by-nilpotent.

Using this value of K' , we see by assumption that $S^{K'}$ is contained in A and thus $S^{K'}$ is covered by at most K' cosets of G_0 . From our initial observation we conclude that $[G : G_0] \leq K'$.

Note that for some $s = O_K(1)$ the s -th term of the central descending series $C^s(G_0)$ is contained in H . Moreover, $G_1 := \bigcap_{g \in G} g G_0 g^{-1}$ is a normal subgroup of G with index at most $O_K(1)$ contained in G_0 . Hence $N := C^s(G_1)$ is a normal subgroup of G contained in H . On the other hand, G_1/N is nilpotent of complexity bounded in terms of K only and it has index $O_K(1)$ in G/N .

To conclude from this that G is virtually nilpotent, it suffices to show that G_1 is. However G_1 is actually finite-by-nilpotent (the finite group being N) and any such group is virtually nilpotent. To see this note that the kernel of the action by conjugation on N is a nilpotent subgroup of finite index. \square

Remark 11.3. — Recall that the condition $|A^2| \leq K|A|$ implies the existence of an approximate group Z of size $O(K^{O(1)}|A|)$ and of $O(K^{O(1)})$ left translates of Z which cover

A (see [52, Theorem 4.6]). Using Remark 1.9 and Theorem 2.12, we then see that G_1 can be taken so that G_1/N is $O(\log K)$ -nilpotent in the sense that it admits a generating set u_1, \dots, u_ℓ with $\ell = O(\log K)$ such that $[u_i, u_j] \in \langle u_{j+1}, \dots, u_\ell \rangle$ for all $i < j$. In particular such a group admits a normal series with cyclic factors of length at most $O(\log K)$.

Remark 11.4. — If one assumes that A is a K -approximate group instead of the doubling condition $|A^2| \leq K|A|$ in Corollary 11.2, then we may also conclude from Theorem 1.6 that N and a generating set of G_1 are contained in A^4 . Using Theorem 2.12 we see that if one additionally wishes to ensure the logarithmic bound as in the previous remark, then one can only guarantee that N lies inside $A^{O_K(1)}$.

The following corollary is reminiscent of Gromov's theorem but it involves a weaker type of polynomial growth condition in which the generating set may be arbitrarily large. Furthermore it only requires that at one scale.

Corollary 11.5. — *Let $d > 0$. Then there is $R(d) > 0$ such that the following holds. Suppose that G is generated by a finite symmetric set S and that there is some scale $r > R(d)$ such that $|S^r| \leq r^d|S|$. Then there is a finite normal subgroup $N \triangleleft G$ and a subgroup $G_1 \leq G$ containing N such that*

- (i) $N \subseteq S^r$;
- (ii) G_1 has index $O_d(1)$ in G ;
- (iii) G_1/N is $O(d)$ -nilpotent (see Remark 1.9 for a definition).

Proof. — Our assumption is that $|S^r| \leq r^d|S|$. Let $K = 2 \cdot 10^d$ and C_K be such that, in the last part of Remark 11.4, N lies in A^{C_K} . We claim that there is some $r_0, \sqrt{r} \leq r_0 \leq r/2C_K$, such that $A := S^{r_0}$ has $|A^5| \leq 10^d|A|$. Note that A^2 is then a K -approximate group with $K = 2 \cdot 10^d$ (see Lemma 5.2). Applying Corollary 11.2 and Remark 11.4 and ensuring that $R(d)$ is so large that $R(d) > (K')^2$ (K' being the quantity in Corollary 11.2), we obtain a finite normal subgroup $N \triangleleft G$ and a subgroup $G_1 \leq G$ containing N such that G_1 has index $O_d(1)$ in G and G_1/N is $O(d)$ -nilpotent. Furthermore N and a set of generators for G_1 are contained in $A^{2C_K} = S^{2r_0C_K} \subseteq S^r$.

It remains to justify the claim. If it is false then $|S^{5^{i+1}\sqrt{r}}| > 10^d|S^{5^i\sqrt{r}}|$ whenever $5^i\sqrt{r} < r/10C_K$, and in particular $|S^r| \geq (10^d)^{\log_5(\sqrt{r}/10C_K)-1}|S^{\sqrt{r}}|$. If r is greater than some absolute constant, this is greater than $r^d|S|$, contrary to assumption. \square

Remark 11.6. — Note that there is no bound on the size of N . Indeed, if G is a large finite simple group and $S = G$ then N must equal G , which shows that $|N|$ can be arbitrarily large compared to d, r .

In [51] Y. Shalom and the third author gave a quantitative refinement of Gromov's theorem inspired by Kleiner's recent new proof (see also [38, Corollary 4.2] for an earlier

result in that direction). A consequence of their result is that a polynomial growth condition at one large scale is enough to guarantee virtual nilpotence. We take the opportunity to record that this follows easily from Corollary 11.2.

Corollary 11.7. — *Let $d > 0$. Then there is $R(d) > 0$ such that the following holds. Suppose that G is generated by a finite symmetric set S containing the identity and that there is some scale $r > R(d)$ such that $|S^r| \leq r^d$. Then G contains G' , where*

- (i) G' has index $O_d((r^d)!) in G ;$
- (ii) G' is nilpotent with step $O_d(1)$.

Proof. — We apply Corollary 11.5 to obtain groups N, G_1 with the properties stated there. As N is contained in S^r , it has cardinality at most r^d . The group G_1 acts on N by conjugation; since the permutation group of N has cardinality at most $(r^d)!$, we conclude that the stabiliser G' of this action has index at most $(r^d)!$ in G_1 . As G_1/N is nilpotent of step $O_d(1)$, we conclude that G' is nilpotent of step $O_d(1) + 1 = O_d(1)$, and the claim follows. \square

Remark 11.8. — As observed in the last section of Gromov's original paper [27], Gromov's theorem on polynomial growth already easily implies a weaker result of this kind in which the hypothesis is that $|S^r| \leq r^d$ for all $r = 1, 2, \dots, R(d)$. Note that this result of Gromov (and, *a fortiori*, Corollary 11.7) have content even when the group G is finite. Another weakening of the above result appears in [58], where $|S^r| \leq r^d$ is assumed for infinitely many r rather than for all r .

Corollary 11.7 is stronger than the results in [38, 51] in the sense that the bounds do not depend on the cardinality $|S|$ of S . On the other hand, the results in [38, 51], which follow a strategy close to that of Kleiner's work [37], yield more effective quantitative control on the index and step of G' , especially in the case when S is of bounded cardinality.

Another consequence of our main theorem is that polynomial growth in the sense of Corollary 11.5 at one large scale implies polynomial growth at all subsequent scales.

Corollary 11.9. — *Let $d > 0$. Then there is $R'(d) > 0$ such that the following holds. Suppose that G is generated by a finite symmetric set S and that $|S^r| \leq r^d |S|$ for some $r \geq R'(d)$. Then $|S^{r'}| \leq (r')^{O_d(1)} |S|$ for all $r' \geq r$.*

Proof. — A simple modification of the proof of Corollary 11.5 shows that there is some $r_0, \sqrt{r} \leq r_0 \leq r/6$, such that $|S^{5r_0}| \leq K |S^{r_0}|$ where $K = 100^d$ (say). Applying Corollary 11.2 with $A := S^{r_0}$ (as before) we obtain a normal subgroup $H \subseteq S^{4r_0}$ such that G/H is virtually nilpotent with the index, step and number of generators of the nilpotent subgroup G_1/H all being $O_d(1)$.

Now by Corollary 5.2, $A^2 = S^{2r_0}$ is a $2K$ -approximate group. This means that there is some set X , $|X| \leq 2K$, such that $S^{4r_0} \subseteq XS^{2r_0}$. From this it follows that

$$(11.1) \quad S^{2mr_0} \subseteq X^m S^{2r_0}$$

for every positive integer m .

Let $\pi : G \rightarrow G/H$ be the quotient homomorphism. We have

$$(11.2) \quad |S^{2mr_0}| \leq |H| |\pi(S^{2mr_0})|,$$

since the cardinality of any fibre is at most $|H|$.

From (11.1) and the fact that π is a homomorphism we have

$$(11.3) \quad \pi(S^{2mr_0}) \subseteq \pi(X)^m \pi(S^{2r_0}).$$

On the other hand, since $H \subseteq S^{4r_0}$, we have

$$(11.4) \quad |H| |\pi(S^{2r_0})| \leq |S^{6r_0}|.$$

Moreover, since $\sqrt{r} \leq r_0 \leq r/6$, we have

$$(11.5) \quad |S^{6r_0}| \leq |S^r| \leq r^d |S| \leq r_0^{2d} |S|.$$

Putting (11.2), (11.3), (11.4) and (11.5) together gives

$$(11.6) \quad |S^{2mr_0}| \leq |\pi(X)^m| r_0^{2d} |S|.$$

Now $\pi(X)$ is a set of size $O_d(1)$, contained in a virtually nilpotent group in which the index and step of the nilpotent subgroup are $O_d(1)$. Every such group is a quotient of one fixed virtually nilpotent group with number of generators, index and step of the nilpotent subgroup also $O_d(1)$ and whose generators are lifts of the elements in $\pi(X)$. Hence there is a bound of the form

$$|\pi(X)^m| \leq m^{O_d(1)}$$

for all $m > 1$. Comparing this with (11.6) confirms that

$$|S^{r'}| \leq r'^{O_d(1)} |S|$$

whenever r' is a multiple $2mr_0$ with $m > 1$. It is not hard to see that the same estimate therefore holds for all r' , at the expense of increasing the exponent $O_d(1)$ if necessary. \square

A consequence/reformulation of the preceding result is the following.

Corollary 11.10. — *Let $\alpha > 0$. Then there are $r_0 \in \mathbf{N}$ and $\beta > 0$ with $\lim_{\alpha \rightarrow 0} \beta(\alpha) = 0$ such that the following holds. Let G be a finite group generated by a symmetric set S and assume that the diameter of the associated Cayley graph satisfies $\text{diam}_S(G) \leq (|G|/|S|)^\alpha$. Then $|S^r| \geq \min\{r^{1/\beta}|S|, |G|\}$ if $r \geq r_0(\alpha)$.*

Proof. — If this does not hold for some r, r_0 and β then as soon as r_0 is large enough (in terms of β) Corollary 11.9 applies and yields $|S^n| \leq n^e|S|$ for all $n \geq r$ and some $e = e(\beta) > 0$. In particular, when n reaches the diameter of G , we obtain $S^n = G$ so $|G| \leq (\text{diam}_S(G))^e|S|$. This contradicts our hypothesis if $e < 1/\alpha$. \square

We shall apply Corollary 11.10 later on to deduce an isoperimetric inequality; see Corollary 11.15.

Finally we show that by repeatedly applying Corollary 11.2 we can obtain the following more precise result, which says something non trivial for finite groups as well. We say that a polycyclic group has length at most L if it is obtained from the trivial group by at most L successive extensions by a cyclic group.

Corollary 11.11. — *Let G be a group which has a left-invariant metric $d : G \times G \rightarrow [0, \infty)$ satisfying the following conditions for some $K \geq 1$:*

- (i) (Uniform doubling property) *We have $|B(2r)| \leq K|B(r)|$ for every $r > 0$;*
- (ii) (Finiteness condition) *There are at most K different subgroups of the form $\langle B(r) \rangle$ as r ranges over $(0, \infty)$.*

Then G has a subgroup of index at most $O_K(1)$ which is polycyclic of length $O_K(1)$.

Proof. — Given $d \in \mathbf{N}$ and $R \geq 0$ we claim that if there are at most d groups of the form $\langle B(r) \rangle$ for $r \leq R$, then $\langle B(R) \rangle$ contains a polycyclic subgroup of index $O_{K,d}(1)$. This is clearly enough to establish the corollary. To prove the claim, we proceed by induction on d . It is clear for $d = 1$, since $\langle B(R) \rangle$ is then the trivial group.

Let R_0 be the upper bound of those $R' \geq 0$ such that there are at most $d - 1$ groups of the form $\langle B(r) \rangle$ for $r \leq R'$. Without loss of generality¹⁸ $0 < R_0 \leq R$. Then $\langle B(r) \rangle = \langle B(R) \rangle$ whenever $R_0 \leq r \leq R$. By the induction hypothesis, $\langle B(R_0/2) \rangle$ contains a polycyclic subgroup P of index $O_{K,d}(1)$ and length $O_{K,d}(1)$.

Let $K' = O_K(1)$ be the constant obtained in Corollary 11.2. Setting $S = B(R_0)$ and $A = B(K'R_0)$, we may apply Corollary 11.2 and conclude that $G = \langle B(R_0) \rangle$ contains a subgroup G_1 of index $O_K(1)$ such that G_1 has a normal subgroup $N \subset B(4K'R_0)$ with G_1/N nilpotent with step and number of generators $O_K(1)$. It is enough to show that G_1 has a polycyclic subgroup of index $O_{K,d}(1)$, because then so will $G = \langle B(R) \rangle$.

By the uniform doubling assumption and a covering argument, $B(4K'R_0)$ can be covered by $O_K(1)$ translates of $B(R_0/2)$. It follows that N can be covered by $O_{K,d}(1)$

¹⁸ Recall that $B(R)$ is the *closed* ball $\{g \in G; d(1, g) \leq R\}$.

translates of P , and in particular $[N : N \cap P] = O_{K,d}(1)$. Now $N \cap P$ is a subgroup of P and hence is also polycyclic of length $O_{K,d}(1)$; in particular, it is generated by $O_{K,d}(1)$ elements. Therefore so is N , and hence N_0 , the intersection of all subgroups of N of index at most $[N : N \cap P]$, has index $O_{K,d}(1)$ in N . (To see this recall Schreier's theorem that if S is a symmetric generating set for a group Γ , and if $\Gamma' \leq \Gamma$ has index k , then S^{2k-1} contains a set of generators for Γ' .)

The group N_0 , being a subgroup of $N \cap P$, is polycyclic. It is also characteristic in N and hence, since N is normal in G_1 , N_0 is also normal in G_1 .

However G_1 acts by conjugation N/N_0 , and the kernel of this action is a subgroup G'_1 of G_1 with index $O_{K,d}(1)$. Now $(N \cap G'_1)/N_0$ is central in G'_1/N_0 and of size $O_{K,d}(1)$. We thus have $N_0 \leq N \cap G'_1 \leq G'_1$, where each successive quotient is polycyclic of length $O_{K,d}(1)$. It follows that G'_1 is polycyclic of length $O_{K,d}(1)$, which is what we wanted to establish. \square

Remark 11.12. — There are examples of groups which satisfy the assumptions of Corollary 11.11 yet have no nilpotent subgroup of index $O_K(1)$. For instance, let p be a large prime and set $G := (\mathbf{Z}/p\mathbf{Z})^2 \rtimes \mathbf{Z}$, where the action is by an element of $SL_2(\mathbf{Z}/p\mathbf{Z})$ which is a diagonal matrix γ of the form $\gamma := \text{diag}(x, x^{-1})$, where $x \in \mathbf{F}_p^*$ is a generator of the multiplicative group of \mathbf{F}_p . Then no subgroup of G of index less than $p-1$ is nilpotent (note that such a subgroup must contain $(\mathbf{Z}/p\mathbf{Z})^2$ and be the preimage of the subgroup of \mathbf{Z} with that index). However we can endow G with a uniformly doubling weighted word metric (with 3 generators) by letting the two standard generators of $(\mathbf{Z}/p\mathbf{Z})^2$ each have weight $\frac{1}{p}$ and γ have weight 1.

We turn now to some geometric applications of the above results.

Manifolds with a lower bound on Ricci Curvature. — A. Petrunin suggested to us some years ago¹⁹ that a result such as Corollary 11.5 would give a purely group-theoretical proof of a theorem of Fukaya and Yamaguchi [16] according to which fundamental groups of almost non-negatively curved manifolds are virtually nilpotent. Recall that a closed manifold M is said to be almost non-negatively curved if one can find a sequence of Riemannian metrics on it for which $\text{diam}(M) \leq 1$ while $K_M \geq -1/n$ where K_M is the sectional curvature. Indeed, a simple application of the Bishop-Gromov inequalities combined with Corollary 11.5 yields the following improvement assuming only a condition on the Ricci curvature.

Corollary 11.13 (Ricci gap). — Given $d \in \mathbf{N}$, there is $\varepsilon(d) > 0$ such that the following holds. Let M be an d -dimensional compact Riemannian manifold with Ricci curvature bounded below by $-\varepsilon$ and diameter at most 1. Then $\pi_1(M)$ has normal subgroup of index $O_d(1)$, which is finite-by- $(O(d)$ -nilpotent). In particular $\pi_1(M)$ is virtually nilpotent.

¹⁹ See also <http://mathoverflow.net/questions/11091>.

Proof. — Fix a base point x_0 on the universal cover \tilde{M} and let \mathcal{F} be a Dirichlet fundamental domain based at x_0 for the action of $\Gamma := \pi_1(M)$: that is,

$$\mathcal{F} := \{p \in \tilde{M} : d(x_0, p) \leq d(\gamma \cdot x_0, p) \text{ for all } \gamma \in \Gamma\}.$$

Set $S := \{\gamma \in \Gamma : d(\gamma \cdot x_0, x_0) \leq 3\}$. Note that $\text{diam}(\mathcal{F}) \leq 1$ and that S is symmetric and contains 1. Observe further that S generates Γ and that for every integer $r \geq 1$ we have $B(x_0, r) \subset S^r \cdot \mathcal{F} \subset B(x_0, 3r + 1)$, where $B(x_0, r)$ is the ball of radius r on \tilde{M} for the Riemannian metric lifted from M . It follows that

$$(11.7) \quad \frac{|S^r|}{|S|} \leq \frac{|B(x_0, 3r + 1)|}{|B(x_0, 1)|}.$$

From the assumed Ricci curvature bound and the Bishop-Gromov volume comparison estimates (see [17, Theorem 4.19]) we have the bound

$$\frac{|B(x_0, r)|}{|B(x_0, 1)|} \leq \frac{|B_{-\varepsilon}(r)|}{|B_{-\varepsilon}(1)|},$$

where $B_{-\varepsilon}(r)$ is a metric ball in the comparison model space with constant curvature $-\varepsilon$ and dimension d . The volume of this ball is $|B_{-\varepsilon}(r)| = \frac{1}{(\sqrt{\varepsilon})^d} |B_{-1}(r/\sqrt{\varepsilon})| = c_d \int_0^r \left(\frac{\sinh(\sqrt{\varepsilon}t)}{\sqrt{\varepsilon}}\right)^{d-1} dt$, where $c_d > 0$ is the volume of the $d - 1$ -dimensional unit sphere (see [17, p. 138] for this volume computation). As ε tends to 0, this tends to $c_d r^d / d$. Combining this with (11.7) we obtain that for every $R_0 \geq 1$ there is some $\varepsilon_0 = \varepsilon_0(d, R_0)$ such that

$$\frac{|S^r|}{|S|} \leq 2(3r + 1)^d$$

for all $r \leq R_0$ provided that $0 < \varepsilon < \varepsilon_0$. Letting $R_0 = R_0(2d)$ be as in Corollary 11.5, we obtain the existence of some $\varepsilon = \varepsilon(d) > 0$ for which the conclusion of that statement holds. This completes the proof. \square

Remark 11.14. — The fact that $\pi_1(M)$ is virtually nilpotent under the above Ricci bounds assumptions was obtained by Cheeger and Colding in [11] (and had been conjectured earlier by Gromov) and their proof was recently completed and extended by Kapovitch and Wilking [36], who also established that the index of the nilpotent subgroup is uniformly bounded by a constant depending on the dimension d only, an improvement which seems beyond the scope of our methods. This extended earlier work of Kapovitch, Petrunin and Tuschmann in [35] which proved the same result under sectional curvature bounds instead of Ricci. The work of these authors is, unlike our work, differential-geometric in nature. The linear dependence in d of the nilpotency length proven in our Corollary 11.13 seems new however.

An isoperimetric inequality. — It has been well-known since the work of Varopoulos on Kesten’s conjecture ([42, 59]) that isoperimetric inequalities on Cayley graphs are closely related to lower bounds on the volume growth. Using this idea and Corollary 11.10, we can derive the following property of finite Cayley graphs with a polynomial upper bound on the diameter.

Corollary 11.15 (*Isoperimetric inequality on finite groups*). — *Let $\alpha > 0$. Then there are $r_0 \in \mathbf{N}$ and $\beta > 0$ with $\lim_{\alpha \rightarrow 0} \beta(\alpha) = 0$ such that the following holds. Let G be a finite group generated by a symmetric set S and assume that the diameter of the associated Cayley graph satisfies $\text{diam}_S(G) \leq (|G|/|S|)^\alpha$. Then for every subset E in G with $\frac{1}{2}r_0 \leq |E| \leq \frac{1}{2}|G|$, $|\partial E| \geq \frac{1}{8}|S|^\beta |E|^{1-\beta}$.*

Proof. — This follows almost immediately from Corollary 11.10 and the following well-known lemma, which may be found in [28, Chapter 5] or [42] and references therein. For the convenience of the reader we offer a self-contained proof. \square

Lemma 11.16 (*Isoperimetry versus growth*). — *Let G be group and S some finite symmetric generating set containing 1. Let $B(r) = S^r$ be the word ball of radius r in the word metric. Let $\partial E = SE \setminus E$ be the boundary of a subset $E \subset G$. If $E \subseteq G$ is a set, write $r(E)$ for the infimum of those r for which $|B(r)| \geq 2|E|$. Then for all E with $|E| < |G|/2$ we have $|E| \leq 4r(E)|\partial E|$.*

Proof. — We give a proof for the reader’s convenience. Let $f = 1_E$ the indicator function of the set E , and $f_r := \frac{1}{|B(r)|} \sum_{g \in B(r)} g \cdot f$ be the average of f over balls of radius r . By the triangle inequality we have $\|g \cdot f - f\|_1 \leq |g| \cdot \max_{s \in S} \|s \cdot f - f\|_1$, where $|g|$ is the distance to the identity in the word metric. Moreover $\|s \cdot f - f\|_1 = |sE \Delta E| \leq 2|\partial E|$ for every $s \in S$. Hence $\|f_r - f\|_1 \leq 2r|\partial E|$. On the other hand for every $x \in E$, there are at most $|E|$ elements $g \in B(r)$ such that $g \cdot 1_E(x) \neq 0$. Therefore if $|B(r)| \geq 2|E|$ then $f_r(x) \leq \frac{1}{2}$ and hence $\|f_r - f\|_1 \geq \frac{1}{2}|E|$. The claim follows. \square

In [1], Benjamini and Kozma conjecture that one can take $\beta = \alpha$ in the Corollary 11.15 (at the expense of introducing a possible multiplicative constant c_α in place of $|S|^\beta/8$ in (ii)). This, however, is beyond the scope of our method. We would like to thank Itai Benjamini for drawing our attention to their work and its connection to Gromov-type theorems.

A generalized Margulis lemma. — In hyperbolic geometry, the Margulis lemma asserts that there is a constant $\varepsilon = \varepsilon(n) > 0$, the Margulis constant, such for any discrete subgroup Γ of isometries of the hyperbolic n -space \mathbf{H}^n , and any point $x \in \mathbf{H}^n$, the almost stabiliser $\Gamma_\varepsilon(x) := \{\gamma \in \Gamma : d(\gamma \cdot x, x) < \varepsilon\}$ is virtually cyclic. This lemma is important for describing the geometry of cusps in hyperbolic manifolds, or for establishing volume lower bounds (see e.g. [55]). Various generalisations of this lemma have been established in the past for more general Riemannian manifolds under curvature upper and lower bounds (e.g. [9, Chapter 6]). Typically in these results, unless the manifold has strictly

negative curvature, “virtually cyclic” in the conclusion of the lemma must be replaced by “virtually nilpotent”.

In [28, §5.F] Gromov raises the issue of establishing a *generalized Margulis lemma* under very weak assumptions on the metric space and he proposes a conjectural statement in this direction. Below we answer Gromov’s question affirmatively.

A metric space X is said to have *bounded packing* with packing constant K if there is $K > 0$ such that every ball of radius 4 in X can be covered by at most K balls of radius 1. Say that a subgroup Γ of isometries of X acts *discretely* on X if every orbit is discrete in the sense that $\{\gamma \in \Gamma : \gamma \cdot x \in \Sigma\}$ is finite for every $x \in X$ and for every bounded set $\Sigma \subseteq X$.

Corollary 11.17 (Generalized Margulis Lemma). — *Let $K \geq 1$ be a parameter. Then there is some $\varepsilon(K) > 0$ such that the following is true. Suppose that X is a metric space with packing constant K , and that Γ is a subgroup of isometries of X which acts discretely. Then for every $x \in X$ the “almost stabiliser” $\Gamma_\varepsilon(x) = \langle S_\varepsilon(x) \rangle$, where $S_\varepsilon(x) := \{\gamma \in \Gamma : d(\gamma \cdot x, x) < \varepsilon\}$, is virtually nilpotent.*

Proof. — Each set $S_r(x)$ is symmetric and contains the identity. Now by the assumption on X the ball $B(x, 4)$ can be covered by collection of balls $B(x_i, 1)$, $i = 1, 2, \dots, K$. Suppose that for $i = 1, 2, \dots, k$ there is at least one element $\gamma_i \in S_4(x)$ with $\gamma_i \cdot x \in B(x_i, 1)$. Suppose now that $\gamma \in S_4(x)$ is arbitrary; then there is some $i \in \{1, 2, \dots, k\}$ such that $\gamma \cdot x \in B(x_i, 1)$. But this means that $d(\gamma \cdot x, \gamma_i \cdot x) < 2$, and therefore $\gamma_i^{-1}\gamma \in S_2(x)$. This implies that $S_4(x) \subseteq \bigcup_{i=1}^k \gamma_i S_2(x)$, which yields (since $S_2(x)^2 \subseteq S_4(x)$) the doubling estimate $|S_2(x)^2| \leq K|S_2(x)|$.

Let $K' = K'(K) > 0$ be the constant from Corollary 11.2. Set $\varepsilon := 2/K'$, $S = S_\varepsilon(x)$ and $A = S_2(x)$. A direct application of Corollary 11.2 shows that $\Gamma_\varepsilon(x) = \langle S \rangle$ is virtually nilpotent. \square

Remark 11.18. — This confirms Gromov’s conjecture, which suggested the same conclusion under the slightly stronger hypotheses that every ball of radius R in X can be covered by at most $C(R/r)^m$ balls of radius r for all $0 < r < R \leq 1$ and some fixed constants $C, m > 0$.

The assumptions of this generalized Margulis lemma are satisfied for example if X is a complete Riemannian manifold with a lower bound on its Ricci curvature, by an immediate application of the Bishop-Gromov volume comparison estimates. In this case, the result was proved by Cheeger-Colding [11] and Kapovitch-Wilking [36], namely:

Corollary 11.19. — *Let $d \geq 1$ be an integer. Then there is $\varepsilon = \varepsilon(d) > 0$ with the following property. Suppose that M is a d -dimensional complete Riemannian manifold with a Ricci curvature lower bound $\text{Ric} \geq -(d-1)$ and that Γ is a subgroup of $\text{Isom}(M)$ which acts properly discontinuously by isometries on M . Then for every $x \in M$ the “almost stabiliser” $\Gamma_\varepsilon(x) := \{\gamma \in \Gamma : d(\gamma \cdot x, x) < \varepsilon\}$ is virtually nilpotent.*

In fact the result in [36, Theorem 1] is a stronger version of Corollary 11.19, establishing that $\Gamma_\varepsilon(x)$ has a nilpotent subgroup of index $O_d(1)$. This stronger result seems to be beyond the scope of our method.

We also note that Corollary 11.11 applies to the Margulis lemma in the context of Riemannian d -manifolds with a lower bound on sectional curvature, because then the Gromov short basis has bounded cardinality from Toponogov’s theorem (see for instance [9, 37.3]). We thus get this way an alternate proof of the Fukaya-Yamaguchi theorem [16] according to which almost non-negatively curved n -manifolds have $O_n(1)$ -virtually polycyclic fundamental group. Again, by [35] we know better, namely that they are $O_n(1)$ -virtually nilpotent, but once again this seems beyond the scope of our method.

Finally we would like to remark that the usual proofs of the classical Margulis lemma bear some resemblance to the proof of our main theorem in as much as they use a similar “shrinking commutator trick” to establish nilpotence. While we proved this shrinking commutator estimate for the escape norm associated to an approximate group as part of the Gleason lemmas (Theorem 8.1), in the Margulis lemma, one proves a similar estimate for the norm $\|\gamma\|_x = d(\gamma \cdot x, x)$ by a Riemannian geometric argument using the assumed curvature bounds. This “shrinking commutator trick” dates back at least to Bieberbach [2] in his proof of Jordan’s theorem on finite linear groups.

Acknowledgements

EB is supported in part by the ERC starting grant 208091-GADA. He also acknowledges support from MSRI where part of this work was finalized. TT is supported by a grant from the MacArthur Foundation, by NSF grant DMS-0649473, and by the NSF Waterman award.

The first author would like to thank E. Lindenstrauss from whom he first learned about these questions and for several related discussions. We also acknowledge the huge intellectual debt we owe to prior work of Hrushovski [33], without which we would not have started this project. We are grateful to him for several enlightening discussions regarding his work and the subject matter of the present paper. We also thank I. Goldbring and L. van den Dries for showing us a preliminary version of their notes on Hilbert’s fifth problem and its local versions, B. Hayes for corrections, and J. Lott for help with the references. Finally, all three authors would like to thank T. Sanders for a number of valuable discussions concerning this work.

Appendix A: Basic theory of ultralimits and ultraproducts

In this appendix we review the machinery of ultralimits and ultraproducts. We will borrow some terminology from nonstandard analysis in order to do this, although we will not rely too heavily on nonstandard machinery in this paper.

We will assume the existence of a *standard universe* \mathfrak{U} which contains all the objects and spaces that one is interested in (such as the natural numbers \mathbf{N} , the real numbers \mathbf{R} , the classical Lie groups, etc.). The precise construction of this universe is not particularly important for our purposes, so long as it forms a set. We refer to objects and spaces inside the standard universe as *standard objects* and *standard spaces*, with the latter being sets whose elements are in the former category.

We will rely heavily on the existence of a *nonprincipal ultrafilter*.

Lemma A.1 (*Ultrafilter lemma*). — *There exists a collection α of subsets of the natural numbers \mathbf{N} with the following properties:*

- (i) (Monotonicity) *If $A \in \alpha$ and $B \supseteq A$, then $B \in \alpha$.*
- (ii) (Closure under intersection) *If $A, B \in \alpha$, then $A \cap B \in \alpha$.*
- (iii) (Maximality) *If $A \subseteq \mathbf{N}$, then either $A \in \alpha$ or $\mathbf{N} \setminus A \in \alpha$, but not both.*
- (iv) (Non-principality) *If $A \in \alpha$, and A' is formed from A by adding or deleting finitely many elements to or from A , then $A' \in \alpha$.*

We refer to a collection α obeying the above axioms as a nonprincipal ultrafilter.

Proof. — The collection of cofinite subsets of \mathbf{N} already obeys the monotonicity, closure under intersection, and non-principality properties. Using Zorn’s lemma,²⁰ one can enlarge this collection to a maximal collection which, it may be verified, has all the required properties. \square

Throughout the paper, we fix a non-principal ultrafilter α . A property $P(\mathfrak{n})$ depending on a natural number \mathfrak{n} is said to hold *for \mathfrak{n} sufficiently close to α* if the set of \mathfrak{n} for which $P(\mathfrak{n})$ holds lies in α . A set of natural numbers lying in α will also be called an *α -large set*.

Once we have fixed this ultrafilter, we can define nonstandard objects and spaces.

Definition A.2 (*Nonstandard objects*). — *Given a sequence $(x_{\mathfrak{n}})_{\mathfrak{n} \in \mathbf{N}}$ of standard objects in \mathfrak{U} , we define their ultralimit $\lim_{\mathfrak{n} \rightarrow \alpha} x_{\mathfrak{n}}$ to be the equivalence class of all sequences $(y_{\mathfrak{n}})_{\mathfrak{n} \in \mathbf{N}}$ of standard objects in \mathfrak{U} such that $x_{\mathfrak{n}} = y_{\mathfrak{n}}$ for \mathfrak{n} sufficiently close to α . Note that the ultralimit $\lim_{\mathfrak{n} \rightarrow \alpha} x_{\mathfrak{n}}$ can also be defined even if $x_{\mathfrak{n}}$ is only defined for \mathfrak{n} sufficiently close to α .*

An ultralimit of standard natural numbers is known as a nonstandard natural number, an ultralimit of standard real numbers is known as a nonstandard real number, and so on.

²⁰ By using this lemma, our results thus rely on the axiom of choice, which we will of course assume throughout this paper. On the other hand, it is possible to rephrase the purely combinatorial results in this paper, such as Theorem 2.10, in the language of Peano arithmetic. Applying a famous theorem of Gödel [22], we then conclude that Theorem 2.10 is provable in ZFC if and only if it is provable in ZF. In fact it is possible, with significant effort, to directly translate these ultrafilter arguments to a much lengthier argument in which neither ultrafilters nor the axiom of choice are used. However, this would require one to “finitise” or “proof-mine” such infinitary results as the Heine-Borel theorem or Theorem B.18, and this in turn would require finitisations of the construction of Haar measure and the Peter-Weyl theorem. This would lead to a vastly messier argument.

For any standard object x , we identify x with its own ultralimit $\lim_{\mathbf{n} \rightarrow \alpha} x$. Thus, every standard natural number is a nonstandard natural number, etc.

Any operation or relation on standard objects can be extended to nonstandard objects in the obvious manner. Indeed, if \mathbf{O} is a k -ary operation, we define

$$\mathbf{O}\left(\lim_{\mathbf{n} \rightarrow \alpha} x_{\mathbf{n}}^1, \dots, \lim_{\mathbf{n} \rightarrow \alpha} x_{\mathbf{n}}^k\right) := \lim_{\mathbf{n} \rightarrow \alpha} \mathbf{O}(x_{\mathbf{n}}^1, \dots, x_{\mathbf{n}}^k)$$

and if \mathbf{R} is a k -ary relation, we define $\mathbf{R}(\lim_{\mathbf{n} \rightarrow \alpha} x_{\mathbf{n}}^1, \dots, \lim_{\mathbf{n} \rightarrow \alpha} x_{\mathbf{n}}^k)$ to be true iff $\mathbf{R}(x_{\mathbf{n}}^1, \dots, x_{\mathbf{n}}^k)$ is true for all \mathbf{n} sufficiently close to α . One easily verifies that these nonstandard extensions of \mathbf{O} and \mathbf{R} are well-defined.

Example 23. — The sum of two nonstandard real numbers $\lim_{\mathbf{n} \rightarrow \alpha} x_{\mathbf{n}}$, $\lim_{\mathbf{n} \rightarrow \alpha} y_{\mathbf{n}}$ is the nonstandard real number

$$\lim_{\mathbf{n} \rightarrow \alpha} x_{\mathbf{n}} + \lim_{\mathbf{n} \rightarrow \alpha} y_{\mathbf{n}} = \lim_{\mathbf{n} \rightarrow \alpha} (x_{\mathbf{n}} + y_{\mathbf{n}}),$$

and the statement $\lim_{\mathbf{n} \rightarrow \alpha} x_{\mathbf{n}} < \lim_{\mathbf{n} \rightarrow \alpha} y_{\mathbf{n}}$ means that $x_{\mathbf{n}} < y_{\mathbf{n}}$ for all \mathbf{n} sufficiently close to α .

Definition A.3 (Ultraproducts). — Let $(\mathbf{X}_{\mathbf{n}})_{\mathbf{n} \in \mathbf{N}}$ be a sequence of standard spaces $\mathbf{X}_{\mathbf{n}}$ in \mathfrak{U} indexed by the natural numbers. The ultraproduct $\prod_{\mathbf{n} \rightarrow \alpha} \mathbf{X}_{\mathbf{n}}$ of the $\mathbf{X}_{\mathbf{n}}$ is defined to be the space of all ultralimits $\lim_{\mathbf{n} \rightarrow \alpha} x_{\mathbf{n}}$, where $x_{\mathbf{n}} \in \mathbf{X}_{\mathbf{n}}$ for all \mathbf{n} . We refer to the ultraproduct of standard sets as a nonstandard set; in a similar vein, an ultraproduct of standard groups is a nonstandard group, and an ultraproduct of standard finite sets is a nonstandard finite set. We refer to ${}^*\mathbf{X} := \prod_{\mathbf{n} \rightarrow \alpha} \mathbf{X}$ as the ultrapower of a standard set \mathbf{X} ; the identification of x with $\lim_{\mathbf{n} \rightarrow \alpha} x$ causes \mathbf{X} to be identified with a subset of ${}^*\mathbf{X}$. We will refer to the ultrapower ${}^*\mathfrak{U}$ of the standard universe \mathfrak{U} as the nonstandard universe.

Remark A.4. — Nonstandard sets in nonstandard analysis behave analogously in some ways to measurable sets²¹ in measure theory; for instance, the union or intersection of two nonstandard sets is again a nonstandard set. Also, just as a subset of a measurable set need not be measurable, a subset of a nonstandard set need not be another nonstandard set. For instance, the nonstandard natural numbers ${}^*\mathbf{N}$ is a nonstandard set (being the ultraproduct of the sequence $\mathbf{N}, \mathbf{N}, \dots$), but the standard natural numbers \mathbf{N} , despite being a subset of ${}^*\mathbf{N}$, is not a nonstandard set.

A fundamental property of ultralimits is that they preserve first-order statements and predicates, a fact known as *Łos's theorem*. Here is one formalisation of this theorem.

²¹ Actually, the notion of an *elementary set* (e.g. a finite union of intervals) would be an even closer analogy here than the notion of a measurable set.

Theorem A.5 (Łos's theorem with parameters). — Let m be a standard natural number, and for each $1 \leq i \leq m$, let $x_i = \lim_{\mathbf{n} \rightarrow \alpha} x_{i,\mathbf{n}}$ be a nonstandard object. If $P(y_1, \dots, y_m)$ is a predicate, then $P(x_1, \dots, x_m)$ is true (as quantified over the nonstandard universe ${}^*\mathfrak{U}$) if and only if $P(x_{1,\mathbf{n}}, \dots, x_{m,\mathbf{n}})$ is true for all \mathbf{n} sufficiently close to α (as quantified over the standard universe \mathfrak{U}).

Proof. — (Sketch) By definition, Łos's theorem is true for “primitive” predicates which take the form $R(x_1, \dots, x_k)$ for some primitive k -ary relation R and objects x_1, \dots, x_k , or of the form $x_{k+1} = O(x_1, \dots, x_k)$ for some primitive k -ary operator O . From the ultrafilter axioms, we also see that Łos's theorem is closed with respect to boolean operations; for instance, if Theorem A.6 holds for $P(x_1, \dots, x_m)$ and $Q(x_1, \dots, x_m)$, then it also holds for $\neg P$ or $P \wedge Q$.

Now, we claim that if Łos's theorem holds for the predicate $P(x_1, \dots, x_m)$, then it also holds for the quantified predicates $\exists x_m : P(x_1, \dots, x_m)$ and $\forall x_m : P(x_1, \dots, x_m)$ (where now there are only $m - 1$ free variables x_1, \dots, x_{m-1} , with x_m being bound). We show this just for the existential quantifier \exists , as the case of the universal quantifier \forall is similar (and can be deduced from the existential case by negation). Suppose first that $\exists x_m : P(x_1, \dots, x_m)$ is true in ${}^*\mathfrak{U}$. Then there exists $x_m = \lim_{\mathbf{n} \rightarrow \alpha} x_{m,\mathbf{n}}$ such that $P(x_1, \dots, x_m)$ holds; by hypothesis, this implies that $P(x_{1,\mathbf{n}}, \dots, x_{m,\mathbf{n}})$ holds for \mathbf{n} sufficiently close to α , and thus $\exists x_m : P(x_{1,\mathbf{n}}, \dots, x_{m-1,\mathbf{n}}, x_m)$ holds for \mathbf{n} in \mathfrak{U} sufficiently close to α as desired. Conversely, if $\exists x_m : P(x_{1,\mathbf{n}}, \dots, x_{m-1,\mathbf{n}}, x_m)$ holds in \mathfrak{U} for \mathbf{n} sufficiently close to α , then by the axiom of (countable) choice, we may find $x_{m,\mathbf{n}} \in \mathfrak{U}$ for such \mathbf{n} such that $P(x_{1,\mathbf{n}}, \dots, x_{m-1,\mathbf{n}}, x_{m,\mathbf{n}})$ holds. Setting $x_m := \lim_{\mathbf{n} \rightarrow \alpha} x_{m,\mathbf{n}}$, we conclude that $P(x_1, \dots, x_m)$ holds, and the claim follows.

The above discussion yields Łos's theorem for any predicate that can be built out of primitive predicates by a finite number of boolean operations and quantifications. However, it is easy to see that all predicates are logically equivalent to a predicate of this form. For instance, $\forall a \forall b \forall c : (a + b) + c = a + (b + c)$ is equivalent to

$$\forall a \forall b \forall c \exists d \exists e \exists f : (d = a + b) \wedge (e = b + c) \wedge (f = d + c) \wedge (f = a + e).$$

This completes the proof. □

In applications, we will actually use a slight generalisation of Łos's theorem.

Theorem A.6 (Łos's theorem with parameters and ultraproducts). — Let m, k be standard natural numbers. For each $1 \leq i \leq m$, let $x_i = \lim_{\mathbf{n} \rightarrow \alpha} x_{i,\mathbf{n}}$ be a nonstandard object, and for each $1 \leq j \leq k$, let $A_j = \prod_{\mathbf{n} \rightarrow \alpha} A_{j,\mathbf{n}}$ be a nonstandard set. If $P(y_1, \dots, y_m; B_1, \dots, B_k)$ is a predicate over m objects and k sets, with the sets A_1, \dots, A_k only appearing in P through the membership predicate $x \in B_j$ for various j and various objects B_j , then $P(x_1, \dots, x_m; A_1, \dots, A_k)$ is true (as quantified over the nonstandard universe ${}^*\mathfrak{U}$) if and only if $P(x_{1,\mathbf{n}}, \dots, x_{m,\mathbf{n}}; A_{1,\mathbf{n}}, \dots, A_{k,\mathbf{n}})$ is true for all \mathbf{n} sufficiently close to α (as quantified over the standard universe \mathfrak{U}).

Proof. — We replace each appearance of $x \in B_j$ in P with a new primitive relation $R_j(x, \mathbf{n})$, which is interpreted in \mathfrak{U} as $x \in A_{j, \mathbf{n}}$. This replaces the predicate $P(y_1, \dots, y_m; B_1, \dots, B_k)$ by a predicate $Q(y_1, \dots, y_m, \mathbf{n})$, with $P(x_{1, \mathbf{n}}, \dots, x_{m, \mathbf{n}}; A_{1, \mathbf{n}}, \dots, A_{k, \mathbf{n}})$ logically equivalent to $Q(x_{1, \mathbf{n}}, \dots, x_{m, \mathbf{n}}, \mathbf{n})$. One easily verifies that $P(x_1, \dots, x_m; A_1, \dots, A_k)$ is logically equivalent to $Q(x_1, \dots, x_m, \lim_{\mathbf{n} \rightarrow \alpha} \mathbf{n})$, and the claim now follows from Theorem A.5. \square

Example 24. — Any ultraproduct $G := \prod_{\mathbf{n} \rightarrow \alpha} G_{\mathbf{n}}$ of groups $G_{\mathbf{n}}$ is again a group, because one can write the property of G being a group as a predicate $P(G)$ that involves membership in G (as well as the constant id and the group operations $\cdot, ()^{-1}$, of course). Conversely, if $G = \prod_{\mathbf{n} \rightarrow \alpha} G_{\mathbf{n}}$ is a group, then $G_{\mathbf{n}}$ is a group for all \mathbf{n} sufficiently close to α .

Example 25. — Let $G = \prod_{\mathbf{n} \rightarrow \alpha} G_{\mathbf{n}}$ be an ultraproduct of groups (and thus also a group), and let $A = \prod_{\mathbf{n} \rightarrow \alpha} A_{\mathbf{n}}$ and $B = \prod_{\mathbf{n} \rightarrow \alpha} B_{\mathbf{n}}$ be subsets of G that are nonstandard sets. Then, for \mathbf{n} sufficiently close to α , $A_{\mathbf{n}}$ and $B_{\mathbf{n}}$ are subsets of $G_{\mathbf{n}}$ and $B_{\mathbf{n}}$ (because this statement can be written as a predicate involving membership in $A_{\mathbf{n}}, B_{\mathbf{n}}, G_{\mathbf{n}}$). In a similar (but more complicated) spirit, for any standard $K \in \mathbf{N}$, A can be covered by K left-translates of B if and only if, for \mathbf{n} sufficiently close to α , $A_{\mathbf{n}}$ can be covered by K left-translates of $B_{\mathbf{n}}$.

A nonstandard real number $x \in {}^*\mathbf{R}$ is said to be *bounded* if one has $|x| \leq C$ for some standard $C > 0$, and *unbounded* otherwise. Similarly, we say that x is *infinitesimal* if $|x| \leq c$ for all standard $c > 0$; in the former case we write $x = O(1)$, and in the latter $x = o(1)$. For every bounded real number $x \in {}^*\mathbf{R}$ there is a unique standard real number $\text{st}(x) \in \mathbf{R}$, called the *standard part* of \mathbf{R} , such that $x = \text{st}(x) + o(1)$, or equivalently that $\text{st}(x) - \varepsilon \leq x \leq \text{st}(x) + \varepsilon$ for all standard $\varepsilon > 0$. Indeed, one can set $\text{st}(x)$ to be the supremum of all the real numbers y such that $x > y$ (or equivalently, the infimum of all the real numbers y such that $x < y$). We write $X = O(Y)$, $X \ll Y$, or $Y \gg X$ if we have $X \leq CY$ for some standard C .

Given a sequence $f_{\mathbf{n}} : X_{\mathbf{n}} \rightarrow Y_{\mathbf{n}}$ of standard functions between standard sets $X_{\mathbf{n}}, Y_{\mathbf{n}}$, one can form the *ultralimit* $f := \lim_{\mathbf{n} \rightarrow \alpha} f_{\mathbf{n}}$, which is a function from the ultraproduct $X := \prod_{\mathbf{n} \rightarrow \alpha} X_{\mathbf{n}}$ to the ultraproduct $Y := \prod_{\mathbf{n} \rightarrow \alpha} Y_{\mathbf{n}}$ defined by the formula

$$f\left(\lim_{\mathbf{n} \rightarrow \alpha} x_{\mathbf{n}}\right) := \lim_{\mathbf{n} \rightarrow \alpha} f_{\mathbf{n}}(x_{\mathbf{n}}).$$

Such ultralimits will be called *nonstandard functions* (and are also known as *internal functions* in the nonstandard analysis literature). In particular, since standard finite sequences $(a_n)_{n=1}^N$ of standard reals $a_n \in \mathbf{R}$ with some standard length $N \in \mathbf{N}$ can be viewed as a function $n \mapsto a_n$ from $\{1, \dots, N\}$ to \mathbf{R} , one can thus define *nonstandard finite sequences* $(a_n)_{n=1}^N$ of nonstandard reals $a_n \in {}^*\mathbf{R}$ with some nonstandard length $N \in {}^*\mathbf{N}$ as an ultralimit of standard finite sequences $(a_{n, \mathbf{n}})_{n=1}^{N_{\mathbf{n}}}$, thus $N = \lim_{\mathbf{n} \rightarrow \alpha} N_{\mathbf{n}}$ and

$$a_{\lim_{\mathbf{n} \rightarrow \alpha} n_{\mathbf{n}}} = \lim_{\mathbf{n} \rightarrow \alpha} a_{n_{\mathbf{n}}, \mathbf{n}}.$$

One can then transplant various operations on standard finite sequences to their non-standard counterparts, and can in particular define the sum

$$\sum_{n=1}^N a_n \in {}^*\mathbf{R}$$

of a nonstandard finite sequence $(a_n)_{n=1}^N = \lim_{n \rightarrow \alpha} (a_{n,n})_{n_n=1}^{N_n}$ by the formula

$$\sum_{n=1}^N a_n := \lim_{n \rightarrow \alpha} \sum_{n_n=1}^{N_n} a_{n,n}.$$

Appendix B: Local groups

In this appendix we recall the basic definitions and notations of (symmetric) local group theory, following Goldbring [24].

Definition B.1 (Local group). — A symmetric local group $G = (G, \text{id}, \cdot, ()^{-1})$ is a topological space G with a distinguished element $\text{id} \in G$ (the identity element), together with a globally defined inversion map $()^{-1} : G \rightarrow G$ and a partially defined product map $\cdot : \Omega \rightarrow G$, obeying the following axioms:

- (i) (Partial closure) Ω is an open neighbourhood of $(G \times \{1\}) \cup (\{1\} \times G)$ in $G \times G$.
- (ii) (Continuity) The maps $()^{-1} : x \mapsto x^{-1}$ and $\cdot : (x, y) \mapsto x \cdot y$ are continuous on G and Ω respectively.
- (iii) (Local associativity) If $g, h, k \in G$ are such that $(g \cdot h) \cdot k$ and $g \cdot (h \cdot k)$ are well-defined (thus $(g, h), (g \cdot h, k), (h, k), (g, h \cdot k)$ all lie in Ω), then $(g \cdot h) \cdot k = g \cdot (h \cdot k)$.
- (iv) (Identity) For any $g \in G$, one has $\text{id} \cdot g = g \cdot \text{id} = g$.
- (v) (Invertibility) If $g \in G$, then $g \cdot g^{-1}$ and $g^{-1} \cdot g$ are well-defined (i.e. $(g, g^{-1}), (g^{-1}, g) \in \Omega$) and are equal to id .

If necessary, we will write id, Ω as id_G, Ω_G to reduce confusion. If $\Omega = G \times G$, we call G a global group or a topological group.

If G has the structure of a smooth finite-dimensional real manifold, and the inversion map $()^{-1}$ and product map \cdot are smooth maps, we say that G is a local Lie group.

Remark B.2. — One can also consider non-symmetric local groups, in which the inversion map $()^{-1}$ is only defined on an open neighbourhood Λ of the identity. However, the theory of non-symmetric local groups contains some minor additional technicalities caused by the existence of non-invertible elements which we wish to avoid here. As we will not consider non-symmetric local groups anywhere in this paper, we will often omit the adjective “symmetric” from the term “local group” when there is no chance of confusion.

Following [24], we do not explicitly assume that G is Hausdorff. In practice, though, one can reduce to the Hausdorff case because the closure of the identity element will turn out to be a closed normal subgroup that one can quotient out by.

Example 26. — If G is a symmetric local group and U is a symmetric open neighbourhood of the identity (thus $g^{-1} \in U$ whenever $g \in U$), then U can also be viewed as a symmetric local group, by restricting the domain Ω of the product maps to $\{(g, h) \in \Omega \cap (U \times U) : g \cdot h \in U\}$ (and also restricting the topological structure of G to U). We will sometimes write this symmetric local group as $G \downarrow_U$ to emphasise that it is the restriction of G to U . In particular, an important source of local groups comes from restricting a global group to an open symmetric neighbourhood of the identity.

One can also restrict G to non-open symmetric neighbourhoods of the identity, but the resulting object obtained is not necessarily a symmetric local group (see e.g. Example 28 below).

We say that two symmetric local groups G, G' are *locally identical* if they have a common restriction, thus there exists a U which is an open symmetric neighbourhood of the identity $1_G = 1_{G'}$ in both G and G' for which the group operations on G and G' , when restricted to U , agree completely (in particular, they have the same domain and range). This is an equivalence relation, and we will focus on those properties of symmetric local groups that are preserved up to local identity.

In a similar spirit, we say that two subsets A, B of a symmetric local group in G are *locally identical* if there exists an open neighbourhood U of the identity in G such that $A \cap U = B \cap U$. For instance, all neighbourhoods of the identity are locally identical. Note that every open neighbourhood of the identity contains an open symmetric neighbourhood, so we can assume here that U is symmetric without loss of generality.

Remark B.3. — Symmetric local groups are defined as topological groups, but if one wishes, one can restrict attention to *discrete* symmetric local groups, in which every set is open. In this case, all references to continuity, openness, and the Hausdorff property in Definition B.1 can be omitted as being automatically satisfied. On the other hand, all discrete local groups are locally equivalent to the trivial local group $\{\text{id}\}$.

Example 27. — If \mathfrak{g} is a (finite-dimensional) Lie algebra, and B is a sufficiently small symmetric open neighbourhood of the identity in \mathfrak{g} , then $\exp(B)$ is a symmetric local group, with the multiplication law given by the Baker-Campbell-Hausdorff formula.

Example 28. — The closed interval $[-1, 1]$ in \mathbf{R} with the addition operation is *not* a symmetric local group, because the set $\{(x, y) \in [-1, 1] \times [-1, 1] : x + y \in [-1, 1]\}$ is not open in $[-1, 1] \times [-1, 1]$. However, the *open* interval $(-1, 1)$ is a symmetric local group.

Given any finite number of elements g_1, \dots, g_m in a global group G , one can use the associativity axiom to unambiguously define the product $g_1 \dots g_m$. In a symmetric local group, one can only define this product $g_1 \dots g_m$ locally. We formalise this as a definition:

Definition B.4 (*Finite products*). — Let g_1, \dots, g_m be a finite number of elements in a symmetric local group G . We say that the product $g_1 \dots g_m$ is well-defined in G (or well-defined for short) if, for each $1 \leq i \leq j \leq m$, we can find a group element $g_{[i,j]} \in G$ with the following properties:

- For each $1 \leq i \leq m$, we have $g_{[i,i]} = g_i$.
- If $1 \leq i \leq j < k \leq m$, the product $g_{[i,j]} \cdot g_{[j+1,k]}$ is well-defined (i.e. $(g_{[i,j]}, g_{[j+1,k]}) \in \Omega$) and equal to $g_{[i,k]}$.

By induction we see that if these group elements $g_{[i,j]}$ exist, then they are unique. We then define $g_1 \dots g_k := g_{[1,k]}$. If $g_1 = \dots = g_k = g$, we abbreviate $g_1 \dots g_k$ as g^k . By abuse of notation, we also write $g_1 \dots g_m \in G$ to denote the assertion that $g_1 \dots g_m$ is defined in G .

We adopt the convention that $g_1 \dots g_m = \text{id}$ when $m = 0$.

An easy induction using the local associativity axiom shows that if $g_1, \dots, g_m \in G$ is such that $g_i \dots g_j$ is well-defined whenever $1 \leq i < j \leq m$ with $(i, j) \neq (1, m)$, and $(g_i \dots g_j) \cdot (g_{j+1} \dots g_k)$ is well-defined whenever $1 \leq i \leq j < k \leq m$, then $g_1 \dots g_m$ is well-defined, and we have

$$(g_i \dots g_k) = (g_i \dots g_j) \cdot (g_{j+1} \dots g_k)$$

for all $1 \leq i \leq j < k \leq m$.

Remark B.5. — It is worth pointing out one subtlety here: in order for $g_1 \dots g_m$ to be well-defined, it is necessary that *all* possible ways of decomposing this m -fold product into pairwise products be well-defined. For instance, for $g_1 g_2 g_3$ to be well-defined, both $(g_1 \cdot g_2) \cdot g_3$ and $g_1 \cdot (g_2 \cdot g_3)$ need to be well-defined. Similarly, if g_1, g_2, g_3, g_4 are such that $g_1 g_2 g_3$, $(g_1 g_2 g_3) \cdot g_4$, $g_2 g_3 g_4$, and $g_1 \cdot (g_2 g_3 g_4)$ are well-defined, this is not yet sufficient to deduce that $g_1 g_2 g_3 g_4$ is well-defined, because $(g_1 g_2) \cdot (g_3 g_4)$ need not be well-defined. For instance, in the (additive) local group $\{-1, 0, +1\}$, the expression $(+1) + (-1) + (-1) + (+1)$ is not well-defined, because $(-1) + (-1)$ is not well-defined.

Related to this is the well-known fact that local associativity does not imply global associativity: it is possible for two different ways of decomposing an m -fold product into pairwise products to both exist, but give distinct values; see [41] for further discussion. For instance, there exists a local group G and elements $g_1, g_2, g_3, g_4 \in G$ such that $((g_1 \cdot g_2) \cdot g_3) \cdot g_4$ and $g_1 \cdot (g_2 \cdot (g_3 \cdot g_4))$ both exist, but are not equal to one another. Of course, in this case, we do not consider $g_1 g_2 g_3 g_4$ to be well-defined.

Another easy induction also shows that for each $m \geq 1$, the set of tuples $(g_1, \dots, g_m) \in G^m$ for which $g_1 \dots g_m$ is well-defined is an open subset of G^m .

Now we extend the notion of products and inverses from individual group elements to sets of such elements.

Definition B.6. — *Let G be a symmetric local group. A subset A of G is said to be symmetric if the set $A^{-1} := \{g^{-1} : g \in A\}$ is contained in A . If A_1, \dots, A_m are subsets of G , we say that $A_1 \dots A_m$ is well-defined in G (or well-defined for short) if $g_1 \dots g_m$ is well-defined for all $g_1 \in A_1, \dots, g_m \in A_m$, in which case we write $A_1 \dots A_m := \{g_1 \dots g_m : g_1 \in A_1, \dots, g_m \in A_m\}$. If $A_1 = \dots = A_m = A$, we abbreviate $A_1 \dots A_m$ as A^m . By abuse of notation, we write $A_1 \dots A_m \subset G$ for the assertion that $A_1 \dots A_m$ is well-defined in G . We adopt the convention that $A_1 \dots A_m = \{\text{id}\}$ when $m = 0$. In particular, $A^0 = \{\text{id}\}$ for any $A \subset G$.*

An easy induction (see [24, Lemma 2.5]) shows that for any local group G and any open neighbourhood U_0 of the identity, there exists a nested sequence $U_0 \supset U_1 \supset U_2 \supset \dots$ of symmetric open neighbourhoods of the identity such that $U_{m+1}^2 \subset U_m$ for every $m \geq 0$, which in particular implies that U_m^m is well-defined in U_0 , and thus $A_1 \dots A_m$ is well-defined in U_0 whenever $A_1, \dots, A_m \subset U_m$.

We make the trivial remark that multiplication of sets is associative: if $A_1 \dots A_m$ is well-defined, then for any $1 \leq i \leq j < k \leq m$, $(A_i \cdot A_j) \cdot (A_{j+1} \dots A_k)$ and $A_i \dots A_k$ are well-defined and equal to each other.

By passing to neighbourhoods such as U_m , one can improve the group-like properties of a local group. To illustrate this principle, let us first introduce the following definition.

Definition B.7 (Cancellative local groups). — *A symmetric local group G is said to be cancellative if the following assertions hold:*

- (i) *Whenever $g, h, k \in G$ are such that gh and gk are well-defined and equal to each other, then $h = k$. (Note that this implies in particular that $(g^{-1})^{-1} = g$.)*
- (ii) *Whenever $g, h, k \in G$ are such that hg and kg are well-defined and equal to each other, then $h = k$.*
- (iii) *Whenever $g, h \in G$ are such that gh and $h^{-1}g^{-1}$ are well-defined, then $(gh)^{-1} = h^{-1}g^{-1}$. (In particular, if $U \subset G$ is symmetric and U^m is well-defined in G for some $m \geq 1$, then U^m is also symmetric.)*

Clearly all global groups are cancellative. A local group need not be cancellative everywhere; however, we can restrict to a large subset on which it is cancellative, by using the following proposition.

Proposition B.8. — *Let G be a symmetric local group, and let U be an open symmetric neighbourhood of the identity in G such that U^6 is well-defined. Then the restriction of G to U is cancellative.*

In particular, the restriction of G to the open symmetric neighbourhood U_6 discussed earlier is cancellative. We shall see later that the property of being cancellative is

hereditary in that it is inherited by passing to subgroups and quotients, and because of this we will be able to easily restrict attention to the cancellative case in our arguments.

Proof. — If $g, h \in U$, then $(gh)^{-1}ghh^{-1}g^{-1}$ is well-defined in G . By evaluating this well-defined expression in two different ways we conclude property (iii). In a similar spirit, by evaluating $g^{-1}gh$ and $g^{-1}gk$ for $g, h, k \in U$ in two different ways, we obtain (i); and similarly for (ii). \square

Lemma B.9. — *Let G be a symmetric local group, and let U, V be open sets with $\text{id} \in V$. Then $\overline{U} \subset U \cdot V$ if $U \cdot V$ is well-defined, and similarly $\overline{U} \subset V \cdot U$ if $V \cdot U$ is well-defined.*

Proof. — We prove the first claim only, as the second is similar. Suppose that g is an adherent point of U . By continuity, we can find an open neighbourhood W of g and an open neighbourhood Y of the identity such that $g \cdot g^{-1} \cdot W \cdot Y^{-1}$ is well-defined and $Y^{-1} \subset V$. By continuity, the set $\{h \in W : g^{-1}h \in Y\}$ is an open neighbourhood of g , and thus contains an element h of U . Writing $v := g^{-1}h$ and expanding out $g \cdot g^{-1} \cdot h \cdot v^{-1}$ in two different ways, we conclude that $g = hv^{-1}$, and thus $g \in U \cdot V$ as required. \square

We can give the class of local groups the structure of a category by defining the notion of a (continuous) homomorphism.

Definition B.10 (Homomorphisms). — *Let G, H, K be symmetric local groups. A continuous homomorphism $\phi : G \rightarrow H$ is a continuous map from G to H with the following properties:*

- (i) ϕ maps the identity of G to the identity of H : $\phi(\text{id}_G) = \text{id}_H$.
- (ii) For every $g \in G$, we have $\phi(g)^{-1} = \phi(g^{-1})$.
- (iii) If $g, h \in G$ are such that $g \cdot h$ is well-defined, then $\phi(g) \cdot \phi(h)$ is well-defined and is equal to $\phi(g \cdot h)$.

We will often omit the adjective “continuous” when G is discrete.

A local homomorphism from G to H is a continuous homomorphism $\phi : U \rightarrow H$ from a symmetric open neighbourhood U of the identity of G to H , where of course we give U the structure of the restricted local group $G \downarrow_U$ from Example 26. Two local homomorphisms $\phi : U \rightarrow H$, $\phi' : U' \rightarrow H$ are equivalent if there exists a neighbourhood V of the identity contained in both U and U' such that ϕ and ϕ' agree on V ; this is an equivalence relation. A local morphism is an equivalence class of local homomorphisms.

Given two local homomorphisms $\phi : U \rightarrow H$ and $\psi : V \rightarrow K$ from G to H and H to K respectively, we define the composition map $\psi \circ \phi : U' \rightarrow K$ by $\psi \circ \phi(g) := \psi(\phi(g))$, where $U' := \{g \in U : \phi(g) \in V\}$. This allows one to define a composition of two local morphisms in the obvious manner.

Example 29. — There are no non-trivial global morphisms from the unit circle \mathbf{R}/\mathbf{Z} to \mathbf{R} . However, there do exist non-trivial local morphisms, such as (the equivalence

class of) the map ϕ from $(-1/4, 1/4) \bmod 1$ to \mathbf{R} defined by setting $\phi(x \bmod 1) := x$ for all $x \in (-1/4, 1/4)$. The concept of a local homomorphism is closely related to that of a *Freiman homomorphism* in additive combinatorics, as discussed for example in [54].

One easily verifies that continuous homomorphisms and local morphisms both obey the axioms of a category; in particular, the composition of two continuous homomorphisms is a continuous homomorphism, and the composition of two local morphisms is again a local morphism. As usual in category theory, we can now say that two local groups G, G' are *locally isomorphic* if there exists a local morphism ϕ from G to G' with an inverse ϕ' from G' to G which is also a local morphism, such that the compositions $\phi \circ \phi'$ or $\phi' \circ \phi$ are equivalent to the identity. Thus, for instance, the unit circle \mathbf{R}/\mathbf{Z} and the line \mathbf{R} are locally isomorphic. This notion of local isomorphism generalises the notion of local identity from Remark 26.

Definition B.11 (Sub-local groups [24]). — Given two symmetric local groups G' and G , we say that G' is a sub-local group of G if G' is the restriction of G to a symmetric neighbourhood of the identity, and there exists an open neighbourhood V of G' with the property that whenever $g, h \in G'$ are such that gh is defined in V , then $gh \in G'$; we refer to V as an associated neighbourhood for G' . If G' is also a global group, we say that G' is a subgroup of G .

If G' is a sub-local group of G , we say that G' is normal if there exists an associated neighbourhood V for G' with the additional property that whenever $g' \in G', h \in V$ are such that $hg'h^{-1}$ is well-defined and lies in V , then $hg'h^{-1} \in G'$. We call V a normalising neighbourhood of G' .

Example 30. — If G, G' are the (additive) local groups $G := \{-2, -1, 0, +1, +2\}$ and $G' := \{-1, 0, +1\}$, then G' is a sub-local group of G (with associated neighbourhood $V = G'$). Note that this is despite G' not being closed with respect to addition in G ; thus we see why it is necessary to allow the associated neighbourhood V to be strictly smaller than G . In a similar vein, the open interval $(-1, 1)$ is a sub-local group of $(-2, 2)$.

The interval $(-1, 1) \times \{0\}$ is also a sub-local group of \mathbf{R}^2 ; here, one can take for instance $(-1, 1)^2$ as the associated neighbourhood. As all these examples are abelian, they are clearly normal.

Example 31. — Let $T : V \rightarrow V$ be a linear transformation on a finite-dimensional vector space V , and let $G := \mathbf{Z} \rtimes_T V$ be the associated semi-direct product. Let $G' := \{0\} \times W$, where W is a subspace of V that is not preserved by T . Then G' is not a normal subgroup of G , but it is a normal sub-local group of G , where one can take $\{0\} \times V$ as a normalising neighbourhood of G' .

Observe that any sub-local group of a cancellative local group is again a cancellative local group.

One also easily verifies that if $\phi : U \rightarrow H$ is a local homomorphism from G to H for some open neighbourhood U of the identity in G , then $\ker(\phi)$ is a normal sub-local

group of U , and hence of G . Note that the kernel of a local morphism is well-defined up to local identity. If H is Hausdorff, then the kernel $\ker(\phi)$ will also be closed.

Conversely, normal sub-local groups give rise to local homomorphisms into quotient spaces.

Lemma B.12 (*Quotient spaces [24]*). — *Let G be a cancellative local group, and let H be a normal sub-local group with normalising neighbourhood V . Let W be a symmetric open neighbourhood of the identity such that $W^6 \subset V$. Then there exists a cancellative local group W/H and a surjective continuous homomorphism $\phi : W \rightarrow W/H$ such that, for any $g, h \in W$, one has $\phi(g) = \phi(h)$ if and only if $gh^{-1} \in H$, and for any $E \subset W/H$, one has E open if and only if $\phi^{-1}(E)$ is open.*

Proof. — We define an equivalence relation on W by declaring $g \sim h$ if $gh^{-1} \in H$. Using the cancellative properties of V (and hence of W^6) we see that this is indeed an equivalence relation. We let $W/H := \{[g]_{\sim} : g \in W\}$ be the set of equivalence classes $[g]_{\sim} := \{h \in W : g \sim h\}$, with the obvious projection map $\pi : W \rightarrow W/H$. We define an inversion relation on W/H by setting $[g]_{\sim}^{-1} := [g^{-1}]_{\sim}$, and a product operation by setting $[g]_{\sim}[h]_{\sim}$ to equal $[g'h']_{\sim}$ if $g'h' \in W$ for at least one representative g', h' of $[g]_{\sim}, [h]_{\sim}$ respectively.

We now verify that these relations are well-defined. To make the inversion relation well-defined, we need to verify that if $g \sim h$, then $g^{-1} \sim h^{-1}$. But from the cancellative properties of W^6 , we have $g^{-1}(h^{-1})^{-1} = g^{-1}(gh^{-1})^{-1}g$, and the claim follows as W^6 is a normalising neighbourhood for H . Similarly, to make the multiplication relation well-defined, we need to verify that if g, g', h, h' are such that $g \sim g', h \sim h'$, and $gh, g'h' \in W$, then $gh \sim g'h'$. But $(gh)(g'h')^{-1} = (g(g')^{-1})g'(h(h')^{-1})(g')^{-1}$, and the claim follows as W^6 is a normalising neighbourhood for H . Similar arguments (which we omit) show that W/H obeys the identity, inverse, and local associativity axioms.

Next, we give W/H the quotient topology, declaring a set E in W/H open iff its inverse image $\pi^{-1}(E)$ is open in W (or equivalently, in G). One easily verifies that W/H becomes a symmetric local group, and the claim follows. \square

Example 32. — Let G be the additive local group $G := (-2, 2)^2$, and let H be the sub-local group $H := \{0\} \times (-1, 1)$, with normalising neighbourhood $V := (-1, 1)^2$. If we then set $W := (-0.1, 0.1)^2$, then the hypotheses of Lemma B.12 are obeyed, and W/H can be identified with $(-0.1, 0.1)$, with the projection map $\phi : (x, y) \mapsto x$.

Example 33. — Let G be the torus $(\mathbf{R}/\mathbf{Z})^2$, and let H be the sub-local group $H = \{(x, \alpha x) \bmod \mathbf{Z}^2 : x \in (-0.1, 0.1)\}$, where $0 < \alpha < 1$ is an irrational number, with normalising neighbourhood $(-0.1, 0.1)^2 \bmod \mathbf{Z}^2$. Set $W := (-0.01, 0.01)^2 \bmod \mathbf{Z}^2$. Then the hypotheses of Lemma B.12 are again obeyed, and W/H can be identified with the interval $I := (-0.01(1 + \alpha), 0.01(1 + \alpha))$, with the projection map $\phi : (x, y) \bmod \mathbf{Z}^2 \mapsto y - \alpha x$ for $(x, y) \in (-0.01, 0.01)^2$. Note, in contrast, that if one quotiented G by the *global* group $\langle H \rangle = \{(x, \alpha x) \bmod \mathbf{Z}^2 : x \in \mathbf{R}\}$ generated by H , the quotient

would be a non-Hausdorff space (and would also contain a dense set of torsion points, in contrast to the interval I which is “locally torsion free”). It is because of this pathological behaviour of quotienting by global groups that we need to work with local group quotients instead.

Remark B.13. — As we have seen in the above discussion, many familiar concepts in (global) group theory have analogues in the local group setting. We will however mention one important global group-theoretic concept that does not have a convenient local analogue, and that is the notion of the global group $\langle A \rangle$ generated by a set A of generators. The problem is that this global group $\langle A \rangle$ consists of words in A of *arbitrarily* length, whereas in a local group one can typically only multiply together a bounded number of elements of A . However, sets such as A^m or $(A \cup A^{-1} \cup \{\text{id}\})^m$ for various choices of exponent m can sometimes serve as a partial substitute for this concept in local group theory, though one of course has to keep track of the precise value of m throughout the argument.

Locally compact local groups. — Recall that a topological space X is said to be *locally compact* if and every point in X has a compact neighbourhood. In particular, one can speak of a locally compact symmetric local group.

To verify local compactness of a symmetric local group, it suffices to do so at the identity.

Lemma B.14. — *Let G be a symmetric local group. Then G is locally compact if and only if there is a compact symmetric neighbourhood of the identity.*

Proof. — [24, Lemma 2.16] The “only if” part is clear (since id already has a compact neighbourhood). Now we turn to the “if” part. Let K be a compact symmetric neighbourhood of the identity. By continuity, there exists an open neighbourhood V of g such that $g^{-1} \cdot V \cdot V \cdot g^{-1}$ is well-defined and $g^{-1} \cdot V \cdot V \cdot g^{-1} \subset K$. In particular, $h \mapsto g^{-1}h$ is a homeomorphism from $V \cdot V \cdot g^{-1}$ to $g^{-1} \cdot V \cdot V \cdot g^{-1}$ which is inverted by the map $k \mapsto gk$. By Lemma B.9, we conclude that $h \mapsto g^{-1}h$ is also a homeomorphism from \overline{V} to $g^{-1} \cdot \overline{V} = g^{-1} \cdot V$. In particular, since $g^{-1} \cdot V$ is a closed subset of K , it is compact, and so \overline{V} is compact also. Thus g has a precompact neighbourhood as required. \square

Corollary B.15. — *If G is a locally compact symmetric local group, and U is a symmetric open neighbourhood of the identity, then U is also a locally compact local group.*

Proof. — By Lemma B.14, G contains a symmetric precompact open neighbourhood V of the identity. By continuity, one can find a symmetric open neighbourhood W of the identity such that $W \cdot W$ is well-defined in $V \cap U$. By Lemma B.9, we conclude that the closure \overline{W} in U is the same as the closure of \overline{W} in G ; as it is contained in the

precompact set V , it is thus precompact. The claim then follows from another application of Lemma B.14. \square

An important subclass of the locally compact local groups are the *(symmetric) local Lie groups*, defined as those (symmetric) local groups which are also smooth finite-dimensional real manifolds, such that the group operations are smooth on their domain of definition. We have the following basic theorem.

Theorem B.16 (Lie's third theorem). — Every local Lie group is locally isomorphic to a global Lie group. Furthermore, one can take the global Lie group to be both connected and simply connected.

See e.g. [50] for a proof.

We have the following deep structure theorem for locally compact global groups, due to Gleason and Yamabe [61].

Theorem B.17 (Gleason-Yamabe). — Suppose that G is a locally compact global group. Then there is an open subgroup G' of G with the following property: inside any neighbourhood of the identity $U \subseteq G'$, there is a compact normal subgroup H such that G'/H is isomorphic to a connected global Lie group.

The analogous theorem for locally compact local groups was established more recently by Goldbring.

Theorem B.18 (Goldbring). — Suppose that G is a locally compact local group. Then some restriction G' of G to a symmetric neighbourhood of the identity has the following property. Inside any neighbourhood of the identity $U \subseteq G'$, there is a compact normal subgroup H such that G'/H is isomorphic to a local Lie group.

Proof. — The only self-contained proof of Theorem B.18 in the literature is in the thesis [23], where it follows from a combination of Section 4.5 and [23, Proposition 4.7.1]. A more easily accessible account of essentially the same material follows by combining [56, Proposition 4.1] (reduction to the NSS case) with [24, §8] (treatment of the NSS case). Alternatively (though ultimately more circuitously) one may apply the main result of [56], which shows that G has a restriction in common with a *global* locally compact group, followed by Theorem B.17. For our applications, we only need to apply Theorem B.18 when G is metrisable, although the general case can be deduced from the metrisable case without much effort. \square

Appendix C: Nilprogressions and related objects

In this appendix we prove two basic facts about coset nilprogressions in normal form, namely that after shrinking the length parameter slightly they are approximate groups,

and are globalisable: that is to say isomorphic to subsets of a global group. The proofs of these facts are quite short due to the strength of the normal form axioms. One can establish similar assertions without the normal form hypothesis, but the arguments are much more complicated in that they require one to work with an explicit basis for the free nilpotent group. They are not needed in this paper.

Lemma C.1. — *Let $P = P_H(u_1, \dots, u_r; N_1, \dots, N_r)$ be a coset nilprogression in C -normal form. Then for all $\varepsilon > 0$ that are sufficiently small depending on r, C , one has*

$$(C.1) \quad (1 + N_1) \dots (1 + N_r) |H| \ll_{\varepsilon, C, r} |P_H(u_1, \dots, u_r; \varepsilon N_1, \dots, \varepsilon N_r)| \\ \ll_C (1 + N_1) \dots (1 + N_r) |H|$$

and hence, by the volume bounds on P ,

$$|P_H(u_1, \dots, u_r; \varepsilon N_1, \dots, \varepsilon N_r)| \gg_{\varepsilon, C, r} |P|.$$

Furthermore, $P_H(u_1, \dots, u_r; \varepsilon N_1, \dots, \varepsilon N_r)$ is a $O_{\varepsilon, C, r}(1)$ -approximate group.

Proof. — By quotienting out the finite group H , which is normalised by $P_H(u_1, \dots, u_r; \varepsilon N_1, \dots, \varepsilon N_r)$ ⁶ (say) if ε is small enough, we may assume that H is trivial. The upper bound in (C.1) is then immediate from the upper bound in (2.2), while the lower bound follows from the local properness axiom in Definition 2.6.

From (C.1) and the Ruzsa covering lemma we see that for ε small enough,

$$P_H(u_1, \dots, u_r; 2\varepsilon N_1, \dots, 2\varepsilon N_r)$$

is covered by $O_{\varepsilon, C, r}(1)$ translates of $P_H(u_1, \dots, u_r; \varepsilon N_1, \dots, \varepsilon N_r)$, and so the final claim follows from Lemma 5.1. \square

Remark C.2. — It is in fact possible to show that $|P_H(u_1, \dots, u_r; \varepsilon N_1, \dots, \varepsilon N_r)|$ decays at a polynomial rate in ε , and that $P_H(u_1, \dots, u_r; \varepsilon N_1, \dots, \varepsilon N_r)$ is a $O_{C, r}(1)$ -approximate group uniformly in ε , but we will not need these stronger conclusions here.

Lemma C.3. — *Let $P = P_H(u_1, \dots, u_r; N_1, \dots, N_r)$ be a coset nilprogression in C -normal form. Then for all $\varepsilon > 0$ that are sufficiently small depending on r, C , the set $P_H(u_1, \dots, u_r; \varepsilon N_1, \dots, \varepsilon N_r)$ is isomorphic to a subset of a global group G .*

From this lemma (and Lemma C.1) we see that Theorem 2.13 follows immediately from Theorem 2.10.

Proof. — We first establish the claim under the additional hypothesis that the N_1, \dots, N_r are sufficiently large depending on r, C ; we will remove this hypothesis at the end of the argument.

Let v_1, \dots, v_r be lifts of the generators u_1, \dots, u_r of P/H to P . By Definition 2.6 and the normality of H , one has

$$(C.2) \quad [v_i, v_j] \in P \left(v_{j+1}, \dots, v_r; O_C \left(\frac{N_{j+1}}{N_i N_j} \right), \dots, O_C \left(\frac{N_r}{N_i N_j} \right) \right) H$$

for all $1 \leq i < j \leq r$; note that the hypothesis that the N_i are large ensure that the right-hand side is well-defined in P .

Consider a word in $P(v_{j+1}, \dots, v_r; O_C(\frac{N_{j+1}}{N_i N_j}), \dots, O_C(\frac{N_r}{N_i N_j}))$, which therefore contains $O_C(\frac{N_{j+1}}{N_i N_j})$ copies of $v_{j+1}^{\pm 1}$, $O_C(\frac{N_{j+2}}{N_i N_j})$ copies of $v_{j+2}^{\pm 1}$, and so forth. Let us the leftmost copy of $v_{j+1}^{\pm 1}$ and move it all the way to the left. Each time it passes through a $v_k^{\pm 1}$ for some $j+1 < k \leq r$, we use (C.2) and create $O_C(\frac{N_l}{N_{j+1} N_k})$ new copies of $v_l^{\pm 1}$ for each $l > k$, plus an element of H which can be pushed all the way to the right using the normality of H . Thus, if one initially had $a_k \frac{N_k}{N_i N_j}$ copies of $v_k^{\pm 1}$ for each $j+1 < k \leq r$ before one started moving the leftmost $v_{j+1}^{\pm 1}$ to the left, then by the end of the move, one would have

$$(C.3) \quad a_l \frac{N_l}{N_i N_j} + O_C \left(\sum_{j+1 < k < l} \frac{a_k N_k}{N_i N_j} \frac{N_l}{N_{j+1} N_k} \right)$$

copies of $v_l^{\pm 1}$ for each $j+1 < l \leq r$. We may simplify the expression (C.3) as

$$\left(a_l + O_C \left(\sum_{j+1 < k < l} \frac{1}{N_{j+1}} a_k \right) \right) \frac{N_l}{N_i N_j}.$$

Thus we have effectively replaced the sequence $(a_k)_{j+1 < k \leq r}$ by the sequence

$$\left(a_l + O_C \left(\sum_{j+1 < k < l} \frac{1}{N_{j+1}} a_k \right) \right)_{j+1 < l \leq r}.$$

We iterate this process $O_C(\frac{N_{j+1}}{N_i N_j}) = O_C(N_{j+1})$ times, and note that the a_k were initially of size $O_C(1)$, and end up at a sequence, all of whose entries are of size $O_{C,r}(1)$. In other words, after moving all copies of $v_{j+1}^{\pm 1}$ to the left, and all copies of H to the right, we end up with $O_{C,r}(\frac{N_k}{N_i N_j})$ copies of $v_k^{\pm 1}$ in the middle for each $j+1 < k \leq r$. We conclude that

$$[v_i, v_j] \in v_{j+1}^{n_{i,j,j+1}} P \left(v_{j+2}, \dots, v_r; O_{C,r} \left(\frac{N_{j+2}}{N_i N_j} \right), \dots, O_{C,r} \left(\frac{N_r}{N_i N_j} \right) \right) H$$

for some $n_{i,j,j+1} = O_C(\frac{N_{j+1}}{N_i N_j})$; note that as long as the N_i are large enough, all words that appear in this reorganisation will lie inside P and so the algebraic manipulations can be

justified. Iterating this procedure $r - j$ times (which will be justified if the N_i are large enough) we see that

$$(C.4) \quad [v_i, v_j] = v_{j+1}^{n_{i,j,j+1}} \dots v_r^{n_{i,j,r}} h_{i,j}$$

for some $n_{i,j,k} = O_{C,r}(\frac{N_k}{N_i N_j})$ and $h_{i,j} \in H$. Also, one has

$$(C.5) \quad v_i h v_i^{-1} = \phi_i(h)$$

for some (outer) automorphism $\phi_i : H \rightarrow H$ of H .

Now let G be the global group generated by H and formal generators e_1, \dots, e_r , subject to the relations

$$(C.6) \quad [e_i, e_j] = e_{j+1}^{n_{i,j,j+1}} \dots e_r^{n_{i,j,r}} h_{i,j}$$

and

$$(C.7) \quad e_i h e_i^{-1} = \phi_i(h)$$

for $1 \leq i < j \leq r$. We claim that for ε small enough, there is an injective homomorphism from $P_H(u_1, \dots, u_r; \varepsilon N_1, \dots, \varepsilon N_r)$ to G , which will give the claim.

To see this, first observe from the normality of H that

$$P_H(u_1, \dots, u_r; \varepsilon N_1, \dots, \varepsilon N_r) = P(v_1, \dots, v_r; \varepsilon N_1, \dots, \varepsilon N_r)H.$$

Organising the words in $P(v_1, \dots, v_r; \varepsilon N_1, \dots, \varepsilon N_r)$ by moving all occurrences of v_1 to the left (using (C.4)) and all occurrences of H to the right (using the normality of H) we then have

$$(C.8) \quad P_H(u_1, \dots, u_r; \varepsilon N_1, \dots, \varepsilon N_r) \\ \subseteq P(v_1; \varepsilon N_1) P(v_2, \dots, v_r; O_{C,r}(\varepsilon N_2), \dots, O_{C,r}(\varepsilon N_r))H$$

assuming ε is small enough in order to justify all the algebraic manipulations. Iterating this we see that

$$(C.9) \quad P_H(u_1, \dots, u_r; \varepsilon N_1, \dots, \varepsilon N_r) \subseteq P(v_1; O_{C,r}(\varepsilon N_1)) \dots P(v_r; O_{C,r}(\varepsilon N_r))H.$$

Thus it suffices to establish an injective homomorphism ϕ from the set

$$(C.10) \quad \{v_1^{n_1} \dots v_r^{n_r} h : n_i = O_{C,r}(\varepsilon N_i); h \in H\}$$

to G . From the local properness property in Definition 2.6, all the products in (C.10) are distinct if ε is small enough. We may thus define ϕ by the formula

$$\phi(v_1^{n_1} \dots v_r^{n_r} h) := e_1^{n_1} \dots e_r^{n_r} h.$$

Next, we show that ϕ is injective. Indeed, suppose that there exist $n_i, n'_i = O_{C,r}(\varepsilon N_i)$ and $h, h' \in H$ with

$$\phi(v_1^{n_1} \dots v_r^{n_r} h) = \phi(v_1^{n'_1} \dots v_r^{n'_r} h')$$

and thus

$$e_1^{n_1} \dots e_r^{n_r} h = e_1^{n'_1} \dots e_r^{n'_r} h'.$$

By the universal properties of G , there is a homomorphism from G to \mathbf{Z} that maps e_1 to 1 and annihilates the other e_i and H . This implies that $n_1 = n'_1$. We can then eliminate n_1, n'_1 and work with the subgroup G_2 of G generated by e_2, \dots, e_r and H . From abstract nonsense we see that G_2 is universal with respect to the constraints (C.6), (C.7) for $i \geq 2$, and that G is the semidirect product of G_2 with \mathbf{Z} using the conjugation action of e_1 on G_2 defined using (C.6), (C.7) for $i = 1$. In particular, there is a homomorphism from G_2 to \mathbf{Z} that maps e_2 to 1 and annihilates the e_i and H for $i > 2$. This gives $n_2 = n'_2$. Continuing in this fashion we see that $n_i = n'_i$ for all i and hence $h = h'$, which establishes injectivity.

Finally, we need to show that ϕ is a homomorphism. It suffices to show that if $n_i, n'_i, n''_i = O_{C,r}(\varepsilon N_i)$ and $h, h', h'' \in H$ are such that

$$(C.11) \quad v_1^{n_1} \dots v_r^{n_r} h v_1^{n'_1} \dots v_r^{n'_r} h' = v_1^{n''_1} \dots v_r^{n''_r} h''$$

then

$$(C.12) \quad e_1^{n_1} \dots e_r^{n_r} h e_1^{n'_1} \dots e_r^{n'_r} h' = e_1^{n''_1} \dots e_r^{n''_r} h''.$$

To see this, we rearrange the word on the left-hand side of (C.11) by moving all occurrences of v_1 to the left, and all occurrences of elements of H to the right, using (C.4) and (C.5); if ε is small enough, then all manipulations take place inside P and can thus be justified. Iterating this process, we must eventually be able to express this word in the form $v_1^{\tilde{n}_1} \dots v_r^{\tilde{n}_r} \tilde{h}$ for some $\tilde{n}_i = O_{C,r}(\varepsilon N_i)$ and $\tilde{h} \in H$. By injectivity, we then have $\tilde{n}_i = n''_i$ and $\tilde{h} = h''$. But then if one formally replaces all the v_i by e_i and uses (C.6), (C.7) in place of (C.4), (C.5) in the rearrangement procedure just described, we conclude (C.12), and the claim follows.

Now we remove the hypothesis that the N_1, \dots, N_r are sufficiently large depending on r, C . Let $F: \mathbf{R}^+ \rightarrow \mathbf{R}^+$ be a function depending on r, C to be chosen later. By the pigeonhole principle, we can find a threshold $M \geq 1$ with $M = O_F(1)$ such that every length N_i is either less than M , or larger than $F(M)$. If we let $1 \leq i_1 < \dots < i_r \leq r$ be those indices i_j with $N_{i_j} > F(M)$, then we see (if F is sufficiently rapidly growing) that $P_H(u_{i_1}, \dots, u_{i_r}; N_{i_1}, \dots, N_{i_r})$ will be a coset nilprogression in $O_{C,r,M}(1)$ -normal form. For F sufficiently rapidly growing, the preceding argument then applies to conclude that $P_H(u_{i_1}, \dots, u_{i_r}; \varepsilon N_{i_1}, \dots, \varepsilon N_{i_r})$ is isomorphic to a subset of a global group if ε is small enough depending on C, r, M , and the claim follows. \square

Remark C.4. — From (C.9) we see that every element in $P_H(u_1, \dots, u_r; \varepsilon N_1, \dots, \varepsilon N_r)$ takes the form

$$(C.13) \quad v_1^{a_1} \dots v_r^{a_r} h$$

for some integers a_1, \dots, a_r with $a_i = O_{C,r}(\varepsilon N_i)$ and $h \in H$. Conversely, it is clear that if $|a_i| \leq \varepsilon N_i$ then all expressions of the form (C.13) lie in $P_H(u_1, \dots, u_r; \varepsilon N_1, \dots, \varepsilon N_r)$. Informally, we thus see that the nilprogression $P_H(u_1, \dots, u_r; \varepsilon N_1, \dots, \varepsilon N_r)$ is comparable in some sense to the *nilbox*

$$\{v_1^{a_1} \dots v_r^{a_r} h : |a_i| \leq \varepsilon N_i; h \in H\}.$$

We will however not exploit this description of nilprogressions in this paper.

A variant of the above analysis also gives polynomial growth of progressions in C -normal form in the global case.

Proposition C.5 (Polynomial growth). — *Let $P = P_H(u_1, \dots, u_r; N_1, \dots, N_r)$ be a coset nilprogression in C -normal form in a global group. Then for all $m \geq 1$, one has $|P^m| \ll_{C,r} m^{O_{C,r}(1)} |P|$.*

Proof. — We allow all implied constants to depend on C, r . As H is normalised by P , we may quotient out by H and reduce to the case when H is trivial. Then

$$P^m \subseteq P(u_1, \dots, u_r; mN_1, \dots, mN_r)$$

and so it suffices (by the volume bound (2.2)) to show that

$$|P(u_1, \dots, u_r; mN_1, \dots, mN_r)| \ll m^{O(1)} (N_1 + 1) \dots (N_r + 1).$$

By modifying the proof of (C.8), one easily verifies that

$$\begin{aligned} & P(u_1, \dots, u_r; mN_1, \dots, mN_r) \\ & \subseteq P(v_1; mN_1) P(v_2, \dots, v_r; O(m^2 N_2), \dots, O(m^2 N_r)); \end{aligned}$$

iterating this, one sees that

$$P(u_1, \dots, u_r; mN_1, \dots, mN_r) \subseteq P(v_1; O(m^{O(1)} N_1)) \dots P(v_r; O(m^{O(1)} N_r)),$$

and the claim follows. \square

REFERENCES

1. I. BENJAMINI and G. KOZMA, A resistance bound via an isoperimetric inequality, *Combinatorica*, **25** (2005), 645–650.
2. L. BIEBERBACH, *Über einen Satz des Herrn C. Jordan in der Theorie der endlichen Gruppen linearer Substitutionen*, Sitzber. Preuss. Akad. Wiss., Berlin, 1911.
3. Y. BILU, Addition of sets of integers of positive density, *J. Number Theory*, **64** (1997), 233–275.
4. Y. BILU, Structure of sets with small sumset, *Astérisque*, **258** (1999), 77–108. Structure theory of set addition.
5. E. BREUILLARD and B. GREEN, Approximate groups. I: the torsion-free nilpotent case, *J. Inst. Math. Jussieu*, **10** (2011), 37–57.
6. E. BREUILLARD and B. GREEN, Approximate groups. II: the solvable linear case, *Q. J. of Math., Oxf.*, **62** (2011), 513–521.
7. E. BREUILLARD and B. GREEN, Approximate groups. III: the unitary case, *Turk. J. Math.*, **36** (2012), 199–215.
8. E. BREUILLARD, B. GREEN, and T. TAO, Approximate subgroups of linear groups, *Geom. Funct. Anal.*, **21** (2011), 774–819.
9. Y. D. BURAGO and V. A. ZALGALLER, *Geometric Inequalities*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 285, Springer, Berlin, 1988. Translated from the Russian by A. B. Sosinskiĭ, Springer Series in Soviet Mathematics.
10. M.-C. CHANG, A polynomial bound in Freiman’s theorem, *Duke Math. J.*, **113** (2002), 399–419.
11. J. CHEEGER and T. H. COLDING, Lower bounds on Ricci curvature and the almost rigidity of warped products, *Ann. Math.*, **144** (1996), 189–237.
12. L. J. CORWIN and E. GREENLEAF, *Representations of Nilpotent Lie Groups and Their Applications*, CUP, Cambridge, 1990.
13. E. CROOT and O. SISASK, A probabilistic technique for finding almost-periods of convolutions, *Geom. Funct. Anal.*, **20** (2010), 1367–1396.
14. D. FISHER, N. H. KATZ, and I. PENG, Approximate multiplicative groups in nilpotent Lie groups, *Proc. Am. Math. Soc.*, **138** (2010), 1575–1580.
15. G. A. FREIMAN, *Foundations of a Structural Theory of Set Addition*, American Mathematical Society, Providence, 1973. Translated from the Russian, Translations of Mathematical Monographs, vol. 37.
16. K. FUKAYA and T. YAMAGUCHI, The fundamental groups of almost non-negatively curved manifolds, *Ann. Math.*, **136** (1992), 253–333.
17. S. GALLOT, D. HULIN, and J. LAFONTAINE, *Riemannian Geometry*, Universitext, Springer, Berlin, 1987.
18. N. GILL and H. HELFGOTT, *Growth in solvable subgroups of $GL_n(\mathbf{Z}/p\mathbf{Z})$* , preprint (2010), [arXiv:1008.5264](https://arxiv.org/abs/1008.5264).
19. N. GILL and H. HELFGOTT, Growth of small generating sets in $SL_n(\mathbf{Z}/p\mathbf{Z})$, *Int. Math. Res. Not.*, **18** (2011), 4226–4251.
20. A. M. GLEASON, The structure of locally compact groups, *Duke Math. J.*, **18** (1951), 85–104.
21. A. M. GLEASON, Groups without small subgroups, *Ann. Math.*, **56** (1952), 193–212.
22. K. GÖDEL, Consistency of the axiom of choice and of the generalized continuum-hypothesis with the axioms of set theory, *Proc. Natl. Acad. Sci.*, **24** (1938), 556–557.
23. I. GOLDBRING, *Nonstandard Methods in Lie Theory*, Ph.D. Thesis, University of Illinois at Urbana-Champaign, 2009.
24. I. GOLDBRING, Hilbert’s fifth problem for local groups, *Ann. Math.*, **172** (2010), 1269–1314.
25. B. GREEN and I. Z. RUZSA, Freiman’s theorem in an arbitrary abelian group, *J. Lond. Math. Soc.*, **75** (2007), 163–175.
26. B. GREEN and T. TAO, Compressions, convex geometry and the Freiman-Bilu theorem, *Q. J. Math.*, **57** (2006), 495–504.
27. M. GROMOV, Groups of polynomial growth and expanding maps, *Publ. Math. IHÉS*, **53** (1981), 53–73.
28. M. GROMOV, *Metric Structures for Riemannian and Non-Riemannian Spaces*, Modern Birkhäuser Classics, Birkhäuser, Boston, 2007. Based on the 1981 French original, With appendices by M. Katz, P. Pansu and S. Semmes, Translated from the French by Sean Michael Bates.
29. M. HALL JR., *The Theory of Groups*, Chelsea Publishing Co., New York, 1976. Reprinting of the 1968 edition.
30. H. A. HELFGOTT, Growth and generation in $SL_2(\mathbf{Z}/p\mathbf{Z})$, *Ann. Math.*, **167** (2008), 601–623.
31. H. A. HELFGOTT, Growth in $SL_3(\mathbf{Z}/p\mathbf{Z})$, *J. Eur. Math. Soc.*, **13** (2011), 761–851.
32. J. HIRSCHFELD, The nonstandard treatment of Hilbert’s fifth problem, *Trans. Am. Math. Soc.*, **321** (1990), 379–400.
33. E. HRUSHOVSKI, Stable group theory and approximate subgroups, *J. Am. Math. Soc.*, **25** (2012), 189–243.
34. I. KAPLANSKY, *Lie Algebras and Locally Compact Groups*, The University of Chicago Press, Chicago, 1971.
35. V. KAPOVITCH, A. PETRUNIN, and W. TUSCHMANN, Nilpotency, almost nonnegative curvature, and the gradient flow on Alexandrov spaces, *Ann. Math.*, **171** (2010), 343–373.
36. V. KAPOVITCH and B. WILKING, *Structure of fundamental groups of manifolds with Ricci curvature bounded below*, preprint (2011), [arXiv:1105.5955](https://arxiv.org/abs/1105.5955).

37. B. KLEINER, A new proof of Gromov's theorem on groups of polynomial growth, *J. Am. Math. Soc.*, **23** (2010), 815–829.
38. J. LEE and Y. MAKARYCHEV, *Eigenvalue multiplicity and volume growth*, preprint (2008), [arXiv:0806.1745](https://arxiv.org/abs/0806.1745).
39. D. MONTGOMERY and L. ZIPPIN, Small subgroups of finite-dimensional groups, *Ann. Math.*, **56** (1952), 213–241.
40. D. MONTGOMERY and L. ZIPPIN, *Topological Transformation Groups*, Interscience Publishers, New York, 1955.
41. P. J. OLVER, Non-associative local Lie groups, *J. Lie Theory*, **6** (1996), 23–51.
42. C. PITTET and L. SALOFF-COSTE, *A survey on the relationships between volume growth, isoperimetry, and the behavior of simple random walk on Cayley graphs, with examples*, survey, preprint (2000).
43. L. PYBER and E. SZABÓ, Growth in finite simple groups of Lie type of bounded rank, preprint (2010), [arXiv:1005.1858](https://arxiv.org/abs/1005.1858).
44. I. Z. RUZSA, Generalized arithmetical progressions and sumsets, *Acta Math. Hung.*, **65** (1994), 379–388.
45. I. Z. RUZSA, An analog of Freiman's theorem in groups, *Astérisque*, **258** (1999), 323–326.
46. T. SANDERS, *From polynomial growth to metric balls in monomial groups*, preprint (2009), [arXiv:0912.0305](https://arxiv.org/abs/0912.0305).
47. T. SANDERS, On a non-abelian Balog-Szemerédi-type lemma, *J. Aust. Math. Soc.*, **89** (2010), 127–132.
48. T. SANDERS, On the Bogolyubov-Ruzsa lemma. *Anal. Partial Differ. Equ.* (2010), to appear, [arXiv:1011.0107](https://arxiv.org/abs/1011.0107).
49. T. SANDERS, A quantitative version of the non-abelian idempotent theorem, *Geom. Funct. Anal.*, **21** (2011), 141–221.
50. J.-P. SERRE, *Lie Algebras and Lie Groups*, Lecture Notes in Mathematics, vol. 1500, Springer, Berlin, 2006. 1964 lectures given at Harvard University, Corrected fifth printing of the second (1992) edition.
51. Y. SHALOM and T. TAO, A finitary version of Gromov's polynomial growth theorem, *Geom. Funct. Anal.*, **20** (2010), 1502–1547.
52. T. TAO, Product set estimates for non-commutative groups, *Combinatorica*, **28** (2008), 547–594.
53. T. TAO, Freiman's theorem for solvable groups, *Contrib. Discrete Math.*, **5** (2010), 137–184.
54. T. TAO and V. VU, *Additive Combinatorics*, Cambridge Studies in Advanced Mathematics, vol. 105, Cambridge University Press, Cambridge, 2006.
55. W. P. THURSTON, *Three-Dimensional Geometry and Topology, vol. 1*, Princeton Mathematical Series, vol. 35, Princeton University Press, Princeton, 1997. Edited by Silvio Levy.
56. L. van den DRIES and I. GOLDBRING, Globalizing locally compact local groups, *J. Lie Theory*, **20** (2010), 519–524.
57. L. van den DRIES and I. GOLDBRING, *Seminar notes on Hilbert's 5th problem*, preprint (2010).
58. L. van den DRIES and A. J. WILKIE, Gromov's theorem on groups of polynomial growth and elementary logic, *J. Algebra*, **89** (1984), 349–374.
59. N. T. VAROPOULOS, L. SALOFF-COSTE, and T. COULHON, *Analysis and Geometry on Groups*, Cambridge Tracts in Mathematics, vol. 100, Cambridge University Press, Cambridge, 1992.
60. H. YAMABE, A generalization of a theorem of Gleason, *Ann. Math.*, **58** (1953), 351–365.
61. H. YAMABE, On the conjecture of Iwasawa and Gleason, *Ann. Math.*, **58** (1953), 48–54.

E. Breuillard

Laboratoire de Mathématiques, Bâtiment 425,

Université Paris Sud 11,

91405 Orsay, France

emmanuel.breuillard@math.u-psud.fr

B. Green

Centre for Mathematical Sciences,

Wilberforce Road,

Cambridge CB3 0WA, England

b.j.green@dpmms.cam.ac.uk

T. Tao
Department of Mathematics,
UCLA,
405 Hilgard Ave,
Los Angeles, CA 90095, USA
tao@math.ucla.edu

Manuscrit reçu le 7 novembre 2011
Manuscrit accepté le 18 septembre 2012
publié en ligne le 19 octobre 2012.