

QUASIPOLYNOMIAL BOUNDS ON THE INVERSE THEOREM FOR THE GOWERS $U^{s+1}[N]$ -NORM

JAMES LENG, ASHWIN SAH, AND MEHTAAB SAWHNEY

ABSTRACT. We prove quasipolynomial bounds on the inverse theorem for the Gowers $U^{s+1}[N]$ -norm. The proof is modeled after work of Green, Tao, and Ziegler and uses as a crucial input recent work of the first author regarding the equidistribution of nilsequences. In a companion paper, this result will be used to improve the bounds on Szemerédi's theorem.

1. INTRODUCTION

We recall the definition of the Gowers U^s -norm on $\mathbb{Z}/N\mathbb{Z}$ and $[N]$. Throughout we let $[N] = \{1, \dots, N\}$.

Definition 1.1. *Given $f: \mathbb{Z}/N\mathbb{Z} \rightarrow \mathbb{C}$ and $s \geq 1$, we define*

$$\|f\|_{U^s(\mathbb{Z}/N\mathbb{Z})}^{2^s} = \mathbb{E}_{x, h_1, \dots, h_s \in \mathbb{Z}/N\mathbb{Z}} \Delta_{h_1, \dots, h_s} f(x)$$

where $\Delta_h f(x) = f(x)\overline{f(x+h)}$ is the multiplicative discrete derivative (extended to lists by composition). Given a natural number N and a function $f: [N] \rightarrow \mathbb{C}$, we choose a number $\tilde{N} \geq 2^s N$ and define $\tilde{f}: \mathbb{Z}/\tilde{N}\mathbb{Z} \rightarrow \mathbb{C}$ via $\tilde{f}(x) = f(x)$ for $x \in [N]$ and 0 otherwise. Then

$$\|f\|_{U^s[N]} := \|\tilde{f}\|_{U^s(\mathbb{Z}/\tilde{N}\mathbb{Z})} / \|\mathbb{1}_{[N]}\|_{U^s(\mathbb{Z}/\tilde{N}\mathbb{Z})}.$$

Remark. This is known to be well-defined and independent of \tilde{N} , and a norm if $s \geq 2$; see [22, Lemma B.5].

Our main result is quasi-polynomial bounds on the inverse theorem for the Gowers U^{s+1} -norm over the integers.

Theorem 1.2. *Fix $\delta \in (0, 1/2)$. Suppose that $f: [N] \rightarrow \mathbb{C}$ is 1-bounded and*

$$\|f\|_{U^{s+1}[N]} \geq \delta.$$

Then there exists a nilmanifold G/Γ of degree s , complexity at most M , and dimension at most d as well as a function F on G/Γ which is at most K -Lipschitz such that

$$|\mathbb{E}_{n \in [N]} [f(n)\overline{F(g(n)\Gamma)}]| \geq \varepsilon,$$

where we may take

$$d \leq \log(1/\delta)^{O_s(1)} \text{ and } \varepsilon^{-1}, K, M \leq \exp(\log(1/\delta)^{O_s(1)}).$$

Remark. Throughout this paper, we will abusively write \log for $\max(\log(\cdot), e^e)$; this is to avoid issues with small numbers.

We have not formally defined a nilmanifold or notions of complexity; our definition is identical to that in work of Green and Tao [24] and will be recalled precisely in Sections 2 and 3.

In the companion paper to this work [39], we will use Theorem 1.2 in order to improve the long standing bounds of Gowers [11, 13] on Szemerédi's theorem.

Leng was supported by NSF Graduate Research Fellowship Grant No. DGE-2034835. Sah and Sawhney were supported by NSF Graduate Research Fellowship Program DGE-2141064.

Theorem 1.3 (Theorem 1.1 in [39]). *Let $r_k(N)$ denote the size of the largest $S \subseteq [N]$ such that S has no k -term arithmetic progressions. For $k \geq 5$, there is $c_k \in (0, 1)$ such that*

$$r_k(N) \ll N \exp(-(\log \log N)^{c_k}).$$

1.1. History and previous results. A long standing conjecture of Erdős and Turán [8] stated that $r_k(N) = o(N)$. In full generality, this conjecture remained open until a combinatorial tour de force of Szemerédi [47, 48] which established the Erdős and Turán conjecture.

Theorem 1.4. *For $k \geq 3$, we have that*

$$r_k(N) = o_k(N).$$

Due to uses of the van der Waerden theorem and the regularity lemma (which was introduced in this work), Szemerédi’s density saving over the trivial bound was exceedingly small. In particular, Szemerédi’s result provided no improvement on known bounds for van der Waerden’s theorem which was part of Erdős and Turán’s original motivation.

The first result in the effort to prove reasonable bounds for $r_k(N)$, e.g. giving a density saving of at least a finite iterated logarithmic type, came from work of Roth [43] which proved

$$r_3(N) \ll N(\log \log N)^{-1}.$$

Being based on Fourier analysis, the methods used in this paper did not obviously generalize to $k \geq 4$. An estimate for $r_k(N)$ which was “reasonable” would have to wait until pioneering work of Gowers [11, 13].

The starting point of work of Gowers [11, 13] is noting via an iterative application of the Cauchy–Schwarz inequality that if a set A of density δ in $[N]$ has no $(s + 2)$ -term arithmetic progressions then $\|f\|_{U^{s+1}[N]} \geq \delta^{O_k(1)}$ where f is a shifted indicator function of the set. In doing so, Gowers provided the correct notion of “psuedorandomness” generalizing Fourier coefficients which was suitable for understanding arithmetic patterns in subsets of the integers and therefore created “higher order Fourier analysis”. The key technical ingredient in work of Gowers was a certain “local inverse theorem” for the $U^{s+1}[N]$ -norm. Gowers proved that given a 1-bounded function f such that $\|f\|_{U^{s+1}[N]} \geq \delta$, there exists a decomposition of $[N]$ into arithmetic progressions of length roughly N^{c_s} and a 1-bounded function g which is constant along these arithmetic progressions such that

$$\mathbb{E}_{x \in [N]} f(x) \overline{g(x)} \geq \delta^{O_s(1)};$$

i.e., f correlates with g . This result, coupled with the density increment strategy as introduced by Roth [43], provided the bound

$$r_k(N) \ll N(\log \log N)^{-c_k}$$

for Szemerédi’s theorem. These bounds have remained the best known for general k until this work. For the sake of comparison, a long sequence of works have attacked the special case of $k = 3$, culminating in a recent breakthrough work of Kelley and Meka [34] which proved

$$r_3(N) \ll N \exp(-c(\log N)^{1/12});$$

the constant $1/12$ was refined to $1/9$ in work of Bloom and Sisask [4]. The only other improvements to the bound of Gowers were due to works of Green and Tao [20, 25] which ultimately established that

$$r_4(N) \ll N(\log N)^{-c},$$

and very recent work of the authors [38] which handled the case $k = 5$ of Theorem 1.3.

Notice however that the “local inverse theorem” of Gowers only gives correlations on arithmetic progressions of length N^{c_s} and that the converse of this result is not true. In particular, a

function may have small $U^{s+1}[N]$ -norm and still correlate with a function which is constant on progressions of length N^{c_s} . To construct such an example, break $[N]$ into consecutive segments of length \sqrt{N} and include each segment with probability $1/2$; while this set with high probability has large “local correlations” it has polynomially small Gowers norm. To obtain a full inverse result (analogous to the quality of Freiman’s theorem, say), one must carefully pin down the global structure as well. Such a task is not straightforward, since the natural generalization of Fourier characters to exponentials of polynomials does not suffice.

A crucial development in the theory towards the inverse conjecture for the Gowers norm was the discovery of the role of nilpotent Lie groups. In groundbreaking work, Furstenberg [9] gave an alternate proof of Szemerédi based on ergodic theory; this work naturally led to seeking to understand certain nonconventional ergodic averages. In works of Conze and Lesigne [6] and Furstenberg and Weiss [10] regarding nonconventional ergodic averages, *nilmanifolds* G/Γ where G is nilpotent and Γ is a discrete cocompact subgroup were brought to the forefront. Host and Kra [33] and independently Ziegler [53], proved convergence of such nonconventional ergodic averages. Crucial to these works was establishing that such averages are controlled by projections on certain characteristic factors which naturally give rise to nilmanifolds. The role of nilsequences (derived from *polynomial sequences* on nilmanifolds) was further highlighted in work of Bergelson, Host, and Kra [3].

The statement of the inverse conjecture (without the given quantification) we will prove was first formulated in work of Green and Tao [22]. Conditional on this inverse conjecture and that the Möbius function does not correlate with nilsequences, Green and Tao were able to prove asymptotic counts for all linear patterns in the primes of “finite complexity”, vastly generalizing the celebrated Green–Tao theorem [19]. Both of these conjectures were resolved; the second being resolved in work of Green and Tao [23] while the first was resolved in work of Green, Tao, and Ziegler [29]. We remark the cases $s = 2$ and $s = 3$ of the inverse conjecture were proven earlier by Green and Tao [18] and Green, Tao, and Ziegler [27] respectively. A crucial ingredient in the cases $s \geq 3$ was work of Green and Tao [24] on the equidistribution behavior of polynomial orbits on nilmanifolds. Subsequently an alternate proof of the inverse conjecture was given by Szegedy (partly joint with Camarena) [5, 46]; this approach was further developed in work of Gutman, Manners, and Varjú [30–32]. These approaches to the inverse theorem, however, at least formally, gave no bounds on the complexity or dimension of the nilsequences with which the function correlates in the cases $s \geq 4$.

The inverse theorem has had numerous further applications within additive combinatorics; we highlight just two. First, Tao and Ziegler [52] gave an asymptotic for the number of polynomial patterns $x + P_1(y), \dots, x + P_j(y)$ in the primes where $P_1(0) = \dots = P_j(0) = 0$ with top degree terms P_1, \dots, P_j being distinct. Second, works of Green and Tao [21] and Altman [1, 2] used the inverse conjecture in combination with an arithmetic regularity lemma to establish the true complexity conjectures of Gowers and Wolf [16].

Due to its importance in the theory of additive patterns, establishing quantitative bounds on the inverse theorem for the Gowers norm has been seen as a central problem in additive combinatorics, with Green suggesting it as “perhaps the biggest open question in the subject” [17, Problem 56]. For the case of $s = 2$, work of Green and Tao [18] gave quantitative bounds for the inverse theorem over all abelian groups and work of Sanders [44] combined with the strategy in [18] proves Theorem 1.2 for the case of $s = 2$. For general s , until roughly five years ago no quantitative bounds were known for the inverse theorem and this was considered a major open problem. This state of affairs was substantially improved in remarkable work of Manners [40] which proves a version of the inverse theorem where the dimension of the nilmanifold is polynomial in $1/\delta$ and the correlation decays double-exponentially in $1/\delta$ (for

$s \geq 4$). This result was subsequently used as a crucial input in work of Tao and Teräväinen [50] to give an effective result for the counts of linear equations in the primes. We remark that a quantitative version of the inverse conjecture over finite fields of high characteristic was proven in work of Gowers and Milićević [14, 15].

At the highest level, the quantitative proofs of Manners [40] and Gowers and Milićević [14, 15] examine when the iterated derivatives of a function are 0 with positive probability. Deriving useful information from this hypothesis over finite fields and the integers are very different problems but fundamentally one glues information from higher derivatives together into information regarding lower derivatives iteratively.

Our proof instead operates via induction on s and attempts to glue degree $(s-1)$ nilmanifolds into a degree s one exactly as in work of Green, Tao, and Ziegler [29]. Our proof in fact is very closely modeled on their work and borrows large sections of their work essentially verbatim. In fact, we believe that the proof in [29], if appropriately quantified, itself yields a bound involving $O(s^2)$ many iterated exponentials. The primary improvement of our proof over theirs stems from the use of improved quantitative equidistribution results on nilmanifolds [35, 36] rather than the results of [24]. The reason we obtain quasi-polynomial bounds is that our proof, even though it inducts on s , gives quasi-polynomial bounds for each step of the induction. Since an iterated composition of finitely many quasi-polynomial functions is still quasi-polynomial, it follows that our bounds should remain quasi-polynomial. In contrast, we believe that the proof in [29], appropriately quantified, results in adding $O(t)$ iterated exponentials in each step t of the induction, which when iterated totals $O(s^2)$ iterated exponentials. Here the results of [35, 36] play a crucial role in eliminating the logarithms accumulated in the induction step.

1.2. Organization of the paper I. We briefly discuss the next three sections of the paper. In Section 2, we define a number of basic notions regarding nilmanifold and set various conventions which will be used throughout the paper. Our conventions differ in various extremely minor ways from those in the work of Green, Tao, and Ziegler [29] but we record them explicitly to recall a number of definitions which will be used throughout the paper. In Section 3, we set various complexity notions that will be given throughout the paper. In the case of nilmanifolds which are given a degree filtration (as is the case in Theorem 1.2), our conventions match those of Green and Tao [24]. Given these notions in hand, we will be in position to outline the main proof in greater detail in Section 4.

Acknowledgements. The first author thanks Terence Tao for advisement. The authors thank Ben Green and Terence Tao for useful discussions regarding [26, 29]. The authors are grateful to Dan Altman and Ben Green for comments. Finally the authors are especially grateful to Sarah Peluse for exceptionally detailed and useful comments on the manuscript.

2. CONVENTIONS ON NILMANIFOLDS

We will recall a large portion of setup regarding nilsequences. In order to discuss this in a quantitative manner, various complexity notions are required which are formally defined in Section 3. This section contains little more than bare definitions; a number of these concepts are developed and motivated in a beautiful manner in [29, Section 6].

2.1. Basic group theory. We briefly record various basic group theory notations which will be used throughout the paper; our notation is identical to that of [29, Section 3].

Given a group G and a subset A , we define $\langle A \rangle$ to be the subgroup generated by the subset A . Given a collection of subgroups $(H_i)_{i \in I}$ in G , we define $\bigvee_{i \in I} H_i$ to be the smallest subgroup

containing all the H_i . Given $h, k \in G$, we denote the commutator of h and k to be

$$[h, k] = h^{-1}k^{-1}hk.$$

Given a sequence of elements $g_1, \dots, g_r \in G$, we define the set of $(r-1)$ -fold commutators inductively. The 0-fold commutators are g_1, \dots, g_r . For $r > 1$, an $(r-1)$ -fold commutator is $[w, w']$ where w and w' are $(s-1)$ -fold and $(s'-1)$ -fold commutators of g_{i_1}, \dots, g_{i_s} and $g_{i'_1}, \dots, g_{i'_{s'}}$ with $\{i_1, \dots, i_s\} \cup \{i'_1, \dots, i'_{s'}\} = \{1, \dots, r\}$ and $s + s' = r$. For instance, $[[g_3, g_4], [g_1, g_2]]$ and $[g_1, [g_3, [g_2, g_4]]]$ are 3-fold commutators of g_1, g_2, g_3 , and g_4 .

We let $H \leq G$ denote that H is a subgroup of G . Given $H, K \leq G$, we denote the commutator subgroup

$$[H, K] = \langle [h, k] : h \in H, k \in K \rangle.$$

The following pair of elementary lemmas will be used throughout the paper to verify various commutator identities; the first is [29, Lemma 3.1].

Lemma 2.1. *Let $H = \langle A \rangle$ and $K = \langle B \rangle$ be normal subgroups of a nilpotent group G . Then $[H, K]$ is also normal and is generated by the $(i + j - 1)$ -fold iterated commutators of $a_1, \dots, a_i, b_1, \dots, b_j$ over all choices of $a_1, \dots, a_i \in A, b_1, \dots, b_j \in B$ and $i, j \geq 1$.*

This implies (see [29, p. 1242]) that for families $(H_i)_{i \in I}, (K_j)_{j \in J}$ which are normal in a nilpotent group G ,

$$\left[\bigvee_{i \in I} H_i, \bigvee_{j \in J} K_j \right] = \bigvee_{i \in I, j \in J} [H_i, K_j].$$

We next require that normality and various filtration conditions can be checked at the level of generators.

Lemma 2.2. *Suppose $K \leq H$ with $H = \langle A \rangle, K = \langle B \rangle$ where $A = A^{-1}$ and $B = B^{-1}$. Then:*

- *If $[a, b] \in K$ for all $a \in A$ and $b \in B$ then K is normal in H .*
- *Suppose $L \leq K \cap H$ is a normal subgroup with respect to both K and H , and suppose for $a \in A, b \in B$, we have $[a, b] \in L$. Then $[H, K] \leq L$.*

Remark. Suppose we wish to prove that $(G_i)_{i \in I}$ forms an I -filtration (see Definition 2.3). This lemma implies that it suffices to check the commutator filtration conditions simply at the level of generators: if for each $i, j \in I$ we know $[g_i, g_j] \in G_{i+j}$ for all generators g_i for G_i and g_j for G_j , then we can deduce that G_{i+j} is normal in G_i using the first bullet point above, and then deduce that $[G_i, G_j] \leq G_{i+j}$ using the second bullet point above.

Proof. For $a \in A, b \in B$ we have $[a, b] \in K$ hence $a^{-1}b^{-1}a \in K$. Since $B = B^{-1}$ generates K , we find $a^{-1}Ka \leq K$. Since A generates H , we deduce that K is normal in H .

For the second item, note that

$$[xy, z] = y^{-1}[x, z]y \cdot [y, z] \text{ and } [x, zy] = [x, y] \cdot y^{-1}[x, z]y.$$

Repeatedly expanding $[h, k]$ for $h \in H, k \in K$ into generators proves the result. \square

Finally, and most importantly, we will require the following versions of the Baker–Campbell–Hausdorff formula (see [29, (3.2)]). Given g_1, g_2 in a nilpotent group G and $n_1, n_2 \in \mathbb{N}$, we have

$$(2.1) \quad g_1^{n_1} g_2^{n_2} = g_2^{n_2} g_1^{n_1} \prod_a g_a^{P_a(n_1, n_2)}$$

where g_a ranges over all iterated commutators of g_1 and g_2 with at least 1 copy of each and $P_a(n_1, n_2): \mathbb{Z} \times \mathbb{Z} \rightarrow \mathbb{Z}$ is a polynomial in n_1 and n_2 . Furthermore if g_a involves d_1 copies of g_1

and d_2 copies of g_2 we have that P_a has degree at most d_1 in n_1 and degree at most d_2 in n_2 . Here the a have been ordered in some arbitrary manner.

If G is a connected, simply connected nilpotent Lie group, then we denote the Lie algebra of G as $\log G$ and let $\exp: \log G \rightarrow G$ denote the exponential map while $\log: G \rightarrow \log G$ is the inverse (the exponential map being a homeomorphism in this situation). When we refer to nilpotent Lie groups, they will henceforth be connected and simply connected. For $g \in G$ and $t \in \mathbb{R}$, we define

$$g^t = \exp(t \log g).$$

The Baker–Campbell–Hausdorff formula also implies that

$$\exp(t_1 \log g_1 + t_2 \log g_2) = g_1^{t_1} g_2^{t_2} \prod_a g_a^{R_a(t_1, t_2)}$$

where g_a ranges over all iterated commutators of g_1 and g_2 with at least 1 copy of each and R_a is a polynomial with rational coefficients satisfying identical degree constraints to P_a . Finally we require the following, most standard version, of the Baker–Campbell–Hausdorff formula which states that if $X, Y \in \log G$, then

$$\exp(X) \exp(Y) = \exp\left(X + Y + \frac{1}{2}[X, Y] + \dots\right)$$

where the remaining terms in the expansion are iterated commutators in X and Y with all higher terms having at least one “copy” of X and Y within them. In particular, this implies that

$$(2.2) \quad \exp(-X) \exp(-Y) \exp(X) \exp(Y) = \exp([X, Y] + \dots)$$

where all higher order terms have at least one copy of X and Y in them and are r -fold commutators with $r \geq 3$. In all versions of Baker–Campbell–Hausdorff, it is important for us that nilpotency means these expressions are finite.

2.2. Filtrations. We next require the notion of an ordering and an associated filtration (see [29, Definition 6.7]).

Definition 2.3. An ordering $I = (I, \preceq, +, 0)$ is a set I with a distinguished element 0 and binary operation $+: I \times I \rightarrow I$ such that

- $+$ is associative and commutative with 0 acting as an identity element;
- \preceq has 0 as the minimal element;
- For all $i, j, k \in I$, if $i \preceq j$ then $i + k \preceq j + k$;
- The initial segments $\{i \in I: i \preceq d\}$ are finite for all d .

We define the following three orderings, with addition being the standard addition:

- The degree ordering is given by the standard ordering on \mathbb{N} , denoted $I = \mathbb{N}$ for short;
- The degree-rank ordering is given by $\{(d, r) \in \mathbb{N}^2: 0 \leq r \leq d\}$ with the ordering that $(d', r') \preceq (d, r)$ if $d' < d$ or $d' = d$ and $r' \leq r$, denoted $I = \text{DR}$ for short;
- The multidegree ordering is given by \mathbb{N}^k with $(i'_1, \dots, i'_k) \preceq (i_1, \dots, i_k)$ when $i'_j \leq i_j$ for all $1 \leq j \leq k$, denoted $I = \mathbb{N}^k$ for short.

An I -filtration of G is a collection of subgroups $G_I = (G_i)_{i \in I}$ such that $G_0 = G$ and:

- (Nesting) If $i, j \in I$ are such that $i \preceq j$ then $G_i \supseteq G_j$;
- (Commutator) For $i, j \in I$, we have $[G_i, G_j] \subseteq G_{i+j}$.

We say that a filtered group G has degree $\leq d$ (for $d \in I$) if G_i is trivial for $i \not\preceq d$. G has degree $\subseteq J$ for a downset J if G_i is trivial whenever $i \notin J$.

Note that the commutation condition implies nested subgroups are normal within each other. We next define degree, degree-rank, and multidegree filtrations.

Definition 2.4. *Given $d \in \mathbb{N}$, we say a group G is given a degree filtration of degree d if:*

- G is given a \mathbb{N} -filtration $(G_i)_{i \in \mathbb{N}}$ with degree $\leq d$;
- $G_0 = G_1$.

Given $(d, r) \in \mathbb{N}^2$ with $0 \leq r \leq d$, G is given a degree-rank filtration of degree-rank (d, r) if:

- G is given a DR-filtration $(G_i)_{i \in \text{DR}}$ with degree $\leq (d, r)$;
- $G_{(0,0)} = G_{(1,0)}$ and $G_{(i,0)} = G_{(i,1)}$ for $i \geq 1$. (We also let $G_{(i,j)} = G_{(i+1,0)}$ for $j > i$.)

The associated degree filtration with respect to this degree-rank filtration is $(G_{(i,0)})_{i \geq 0}$.

Given $(d_1, \dots, d_k) \in \mathbb{N}^k$, G is given a multidegree filtration of multidegree J (where $J \subseteq \mathbb{N}^k$ is a downset) if:

- G is given a \mathbb{N}^k -filtration $(G_i)_{i \in \mathbb{N}^k}$ with degree $\subseteq J$;
- $G_{\vec{0}} = \bigvee_{i=1}^k G_{\vec{e}_i}$.

The associated degree filtration with respect to the multidegree filtration is $(\bigvee_{|\vec{i}|=i} G_{\vec{i}})_{i \geq 0}$.

Remark. This definition imposes some additional equalities of subgroups in order to say a group is given a degree-rank filtration versus a DR-filtration (for example). In particular, the concept of “degree-rank” filtration and DR-filtration are distinct. The difference is minor, but causes a number of technical checks to be required, most notably in Appendix C. We will almost exclusively operate with these additional conditions; this is so that we can invoke equidistribution theory safely.

We now define polynomial sequences of an I -filtered group. The notion of a polynomial sequence for a group G given a degree-rank filtration will be same as treating this ordering as a DR-filtration; the same applies for degree and multidegree filtrations.

Definition 2.5. *Given $g: H \rightarrow G$ a map between groups (not necessarily a homomorphism) and $h \in H$, we define the derivative $\partial_h g: H \rightarrow G$ via $\partial_h g(n) = g(hn)g(n)^{-1}$ for all $n \in H$. If H, G are I -filtered, we say that this map g is polynomial if for all $m \geq 0$ and $i_1, \dots, i_m \in I$, we have*

$$\partial_{h_1} \cdots \partial_{h_m} g(n) \in G_{i_1 + \cdots + i_m}$$

for all choices of $h_j \in H_{i_j}$ and $n \in H_0$. The space of all polynomial maps with respect to this data is denoted $\text{poly}(H_I \rightarrow G_I)$.

We will require various general properties of polynomial sequences established in [29, Appendix B]. We will only consider $H = \mathbb{Z}^k$ for $k \geq 1$ and the following I -filtrations on H .

Definition 2.6. *We define the following filtrations on $H = \mathbb{Z}^k$:*

- The (domain) degree filtration is with $I = \mathbb{N}$ the degree ordering and $H_0 = H_1 = \mathbb{Z}^k$, and $H_i = \{0\}$ for $i \geq 2$;
- The (domain) multidegree filtration is with $I = \mathbb{N}^k$ the multidegree ordering, $H_{\vec{0}} = \mathbb{Z}^k$, $H_{\vec{e}_i} = \mathbb{Z}\vec{e}_i$ for $i \in [k]$, and $H_{\vec{v}} = \{0\}$ otherwise, where \vec{e}_i forms the standard basis of \mathbb{Z}^k ;
- The (domain) degree-rank filtration is with $I = \text{DR}$ the degree-rank ordering and $H_{(0,0)} = H_{(1,0)} = \mathbb{Z}^k$ and $H_{(d,r)} = \{0\}$ otherwise.

We now define the notion of a nilmanifold, which is essentially a compact quotient of a filtered nilpotent Lie group.

Definition 2.7. *We define an I -filtered nilmanifold G/Γ to be the data of a connected, simply connected nilpotent Lie group G with I -filtration (of Lie subgroups) and discrete cocompact*

subgroup $\Gamma \leq G$ which is rational with respect to G_I (i.e., $\Gamma_i := \Gamma \cap G_i$ is cocompact in G_i for all $i \in I$). We say it has degree $\leq d$ or $\subseteq J$ if G has degree $\leq d$ or $\subseteq J$.

If $I = \mathbb{N}$ and the I -filtration is furthermore a degree filtration with degree $\leq d$, then G/Γ is a degree d nilmanifold. If $I = \text{DR}$ and the I -filtration is furthermore a degree-rank filtration with degree $\leq (d, r)$, then G/Γ is a degree-rank (d, r) nilmanifold. Finally if $I = \mathbb{N}^k$ and the I -filtration is furthermore a multidegree filtration with degree $\subseteq J$, then G/Γ is a multidegree J nilmanifold.

Remark. Note that Γ can naturally be given the structure of an I -filtered group Γ_I .

We finally (very occasionally) will require the lower central series of a group G .

Definition 2.8. Given a nilpotent group G , define the lower central series inductively via $G_{(0)} = G_{(1)} = G$ and $G_{(i+1)} = [G, G_{(i)}]$. The step of G is the minimal j such that $G_{(j+1)} = \text{Id}_G$.

2.3. Horizontal tori and Taylor coefficients. The next notion, that of a horizontal character, plays a vital role when discussing the equidistribution of nilsequences.

Definition 2.9. Given a connected, simply connected nilpotent group G and a discrete, cocompact subgroup Γ , a horizontal character η is a continuous homomorphism $\eta: G \rightarrow \mathbb{R}$ such that $\eta(\Gamma) \subseteq \mathbb{Z}$. We say a horizontal character is nontrivial when η is not identically zero.

Remark. Throughout the literature on nilmanifolds, horizontal characters are continuous homomorphisms $\eta: G \rightarrow \mathbb{R}/\mathbb{Z}$ such that η annihilates Γ . It is straightforward to prove (via using Mal'cev bases) that these two notions are identical up to taking $\text{mod } 1$. The reason we operate with the above definition is that the kernel of η as defined is then a subspace of $G/[G, G] \simeq \mathbb{R}^{\dim(G) - \dim([G, G])}$.

We next require the notion of horizontal tori with respect to a degree-rank filtration. These tori will play a starring role in Sections 8, 9, and 10; our definition is exactly that of [29, Definition 9.6].

Definition 2.10. Let G be a degree-rank filtered nilpotent Lie group with filtration $G_{\text{DR}} = (G_{(d,r)})_{(d,r) \in \text{DR}}$. Given a subgroup Γ of G , we define various horizontal tori for $i \geq 1$ as

$$\begin{aligned} \text{Horiz}_i(G) &:= G_{(i,1)}/G_{(i,2)}, \\ \text{Horiz}_i(\Gamma) &:= (\Gamma \cap G_{(i,1)})/(\Gamma \cap G_{(i,2)}), \\ \text{Horiz}_i(G/\Gamma) &:= \text{Horiz}_i(G)/\text{Horiz}_i(\Gamma). \end{aligned}$$

Given a polynomial sequence $g \in \text{poly}(\mathbb{Z}_{\text{DR}} \rightarrow G_{\text{DR}})$ we define the i -th horizontal Taylor coefficient to be

$$\begin{aligned} \text{Taylor}_i(g) &:= \partial_1 \cdots \partial_1 g(n) \text{ mod } G_{(i,2)} \in \text{Horiz}_i(G), \\ \text{Taylor}_i(g\Gamma) &:= \text{Taylor}_i(g) \text{ mod } \text{Horiz}_i(\Gamma) \in \text{Horiz}_i(G/\Gamma), \end{aligned}$$

where we take i iterated derivatives.

We also require the notion of i -th horizontal characters.

Definition 2.11. Consider a nilmanifold G/Γ with a degree-rank filtration. A continuous homomorphism $\eta: G_{(i,1)} \rightarrow \mathbb{R}$ is an i -th horizontal character if $\eta(G_{(i,2)}) = 0$ and $\eta(G_{(i,1)} \cap \Gamma) \subseteq \mathbb{Z}$.

The name Taylor coefficient is also used in the context of Taylor coefficients of polynomial factorizations. The following elementary lemma relates these two notions; we remark that a very closely related proof appears in [21, Lemma A.8].

Lemma 2.12. *Let G be given a degree-rank filtration of degree-rank (d, r) and consider a sequence $g \in \text{poly}(\mathbb{Z}_{\text{DR}} \rightarrow G_{\text{DR}})$. Then we may write $g(n) = \prod_{i=0}^d g_i^{\binom{n}{i}}$ for elements $g_i \in G_{(i,0)}$ and for $1 \leq i \leq d$ we have*

$$\text{Taylor}_i(g) = g_i \bmod G_{(i,2)}.$$

Proof. The representation of $g(n)$ in the specified product form follows immediately from the existence of Taylor expansion, see [29, Lemma B.9].

We next prove $\text{Taylor}_j(g) = g_j \bmod G_{(j,2)}$ for each $1 \leq j \leq d$ individually. Notice that it suffices to consider $\tilde{g}(n)$ which is $g(n) \bmod G_{(j,2)}$, i.e., we consider the group $G/G_{(j,2)}$ with quotiented filtration. This group is easily seen to be at most j -step nilpotent and furthermore $[G_{(i,1)}/G_{(j,2)}, G_{(j-i,1)}/G_{(j,2)}] = \text{Id}_{G/G_{(j,2)}}$ for $0 \leq i \leq j$ (one should check the cases $j = 1$ and $i \in \{0, j\}$ manually). Let $\tilde{G}_i = G_{(i,1)}/G_{(j,2)}$ for $0 \leq i \leq j$ and note $\tilde{G}_0 = \tilde{G}_1$.

We see that $\tilde{G}_0 \geq \dots \geq \tilde{G}_j$ is an \mathbb{N} -filtration for \tilde{G}_0 with $[\tilde{G}_i, \tilde{G}_{j-i}] = \text{id}_{\tilde{G}_0}$ for all $0 \leq i \leq j$. Note that $\tilde{g}(n) = \prod_{i=0}^j \tilde{g}_i^{\binom{n}{i}}$ where \tilde{g}_i is $g_i \bmod G_{(j,2)}$.

It suffices to prove the claim that $\tilde{g}(n+1)\tilde{g}(n)^{-1} = \prod_{i=0}^{j-1} (\tilde{g}'_i)^{\binom{n}{i}}$ with $\tilde{g}'_i \in \tilde{G}_{i+1}$ and $\tilde{g}'_{j-1} = \tilde{g}_j$. If this is the case, then we may modify the filtration $\tilde{G}_0 \geq \tilde{G}_1 \geq \dots \geq \tilde{G}_j$ by stripping off the top group, which maintains the necessary inductive properties. Iterating this procedure j times we obtain the desired Taylor equality.

This claim is a consequence of the Taylor expansion for general polynomial sequences and the Baker–Campbell–Hausdorff formula and counting the depths of nested commutators. The crucial reason that $\tilde{g}'_{j-1} = \tilde{g}_j$ is that any “higher order” terms which arise in the Baker–Campbell–Hausdorff formula and could contribute are in fact annihilated due to $[\tilde{G}_i, \tilde{G}_{j-i}] = \text{id}_{\tilde{G}_0}$ for $0 \leq i \leq j$. \square

We also have the following linearity of the i -th Taylor coefficients.

Lemma 2.13. *Assume the setup of Lemma 2.12. We have*

$$\text{Taylor}_i(gh) = \text{Taylor}_i(g) \text{Taylor}_i(h)$$

and if

$$g(n) = \exp\left(\sum_{i=0}^d g_i \binom{n}{i}\right)$$

for $g_i \in \log(G_i)$ we have

$$\text{Taylor}_i(g) = \exp(g_i) \bmod G_{(i,2)}.$$

Proof. The first claim follows from Lemma 2.12, the Baker–Campbell–Hausdorff formula, and the commutator relationship that $[G_{(i,0)}, G_{(j-i,0)}] = [G_{(i,1)}, G_{(j-i,1)}] \subseteq G_{(j,2)}$. (Note that this is using $G_{(0,0)} = G_{(0,1)} = G_{(1,0)}$ in the case $i = 0$.)

For the second claim, suppose that

$$g(n) = \exp\left(\sum_{i=0}^d g_i \binom{n}{i}\right) = \prod_{i=0}^s (g'_i)^{\binom{n}{i}}.$$

Via iterated applications of the Baker–Campbell–Hausdorff formula and the commutator relationship that $[G_{(i,0)}, G_{(j-i,0)}] = [G_{(i,1)}, G_{(j-i,1)}] \subseteq G_{(j,2)}$, we see that $g'_j = \exp(g_j) \bmod G_{(j,2)}$ and the result follows. \square

2.4. Vertical tori and nilcharacters. Given a polynomial sequence g on an I -filtered Lie group with $I = \mathbb{N}$, we can define a sequence of vectors by considering a smooth vector-valued function F on G/Γ and looking at $F(g(n)\Gamma)$. However, we will be particularly interested in those which “have a Fourier coefficient” with respect to various subgroups of the center.

Definition 2.14. Consider a nilmanifold G/Γ and a function $F: G/\Gamma \rightarrow \mathbb{C}$. Given a connected, simply connected subgroup T of the center $Z(G)$ which is rational (i.e., $\Gamma \cap T$ is cocompact in T) and a continuous homomorphism $\eta: T \rightarrow \mathbb{R}$ such that $\eta(T \cap \Gamma) \subseteq \mathbb{Z}$, if

$$F(gx) = e(\eta(g))F(x) \text{ for all } g \in T$$

we say that F has a T -vertical character (or T -vertical frequency) η .

Remark. Note that $T/(\Gamma \cap T)$ is isomorphic to a torus and thus one can modify functions under consideration to have vertical characters via appropriate Fourier decomposition.

A particular case which will arise frequently in our applications comes from the fact that given a filtration satisfying the conditions of Definition 2.4, we have that the “bottom group” is contained in the center. For example, a group G given a degree filtration of degree d satisfies $[G, G_d] = [G_1, G_d] = \text{Id}_G$ hence $G_d \leq Z(G)$. One special class of functions with a vertical frequency which will be of particular importance is that of nilcharacters.

Definition 2.15. A nilcharacter of degree d and output dimension D is the following data. Consider an I -filtered nilmanifold G/Γ of degree d such that $[G, G_d] = \text{Id}_G$ and an I -filtered abelian group H . Let $g \in \text{poly}(H_I \rightarrow G_I)$ and consider function $F: G/\Gamma \rightarrow \mathbb{C}^D$ such that:

- $\|F(x)\|_2 = 1$ for all $x \in G/\Gamma$ pointwise;
- $F(g_d x) = e(\eta(g_d))F(x)$ for all $g_d \in G_d$ where η is some continuous homomorphism $G_d \rightarrow \mathbb{R}$ such that $\eta(\Gamma \cap G_d) \subseteq \mathbb{Z}$.

The values of the nilcharacter are given by $\chi: H \rightarrow \mathbb{C}^D$ where $\chi(n) = F(g(n)\Gamma)$ for $n \in H$.

Remark. We work with vector-valued nilcharacters for precisely the same topological reason given in [29, p. 1254].

2.5. Additional miscellaneous conventions. We end this section with a brief discussion of various miscellaneous conventions. Throughout the paper we use $\{\cdot\}$ to denote the map $\mathbb{R} \rightarrow (-1/2, 1/2]$ (or $\mathbb{R}/\mathbb{Z} \rightarrow (-1/2, 1/2]$, abusively) which takes the representative mod 1 closest to 0. Furthermore given $x \in \mathbb{R}/\mathbb{Z}$ and $y \in \mathbb{R}$ we will treat $x - y \in \mathbb{R}/\mathbb{Z}$ in the obvious manner. As used above, we let $e: \mathbb{R}/\mathbb{Z} \rightarrow \mathbb{C}$ denote the exponential function $e(x) = \exp(2\pi i x)$, which is lifted to \mathbb{R} in the obvious manner.

We use standard asymptotic notation. Given functions $f = f(n)$ and $g = g(n)$, we write $f = O(g)$, $f \ll g$, $g = \Omega(f)$, or $g \gg f$ to mean that there is a constant C such that $|f(n)| \leq Cg(n)$ for sufficiently large n . We write $f \asymp g$ or $f = \Theta(g)$ to mean that $f \ll g$ and $g \ll f$, and write $f = o(g)$ or $g = \omega(f)$ to mean $f(n)/g(n) \rightarrow 0$ as $n \rightarrow \infty$. Subscripts indicate dependence on parameters.

Finally in various arguments throughout the paper it will be convenient to denote appropriately bounded functions as $b(n)$ or $b(n_1, \dots, n_k)$, and $B(n), B(n_1, \dots, n_k)$ when vector-valued. When using such notation, the functions b, B may change from line to line and within a line may refer to different functions.

3. VARIOUS COMPLEXITY NOTIONS

3.1. Rationality of bases and Lipschitz norms. We will now discuss the definitions chosen for complexity of nilmanifolds. We start by defining first- and second-kind coordinates given a basis \mathcal{X} for $\log G$.

Definition 3.1. Consider a connected, simply connected nilpotent Lie group G of dimension d . Given a basis $\mathcal{X} = \{X_1, \dots, X_d\}$ of $\log G$ and $g \in G$, there exists $(t_1, \dots, t_d) \in \mathbb{R}^d$ such that

$$g = \exp(t_1 X_1 + t_2 X_2 + \dots + t_d X_d).$$

We define Mal'cev coordinates of first-kind $\psi_{\text{exp}} = \psi_{\text{exp}, \mathcal{X}}: G \rightarrow \mathbb{R}^d$ for g relative to \mathcal{X} by

$$\psi_{\text{exp}}(g) := (t_1, \dots, t_d).$$

Given $g \in G$ there also exists $(u_1, \dots, u_d) \in \mathbb{R}^d$ such that

$$g = \exp(u_1 X_1) \cdots \exp(u_d X_d),$$

and we define the Mal'cev coordinates of second-kind $\psi = \psi_{\mathcal{X}}: G \rightarrow \mathbb{R}^d$ for g relative to \mathcal{X} by

$$\psi(g) := (u_1, \dots, u_d).$$

Note that the above definition does not account for the cocompact subgroup Γ . The next set of definitions account for how “rational” \mathcal{X} is with respect to itself and Γ .

Definition 3.2. The height of a number x is $\max(|a|, |b|)$ if $x = a/b$ with $\gcd(a, b) = 1$ and ∞ if x is irrational.

Definition 3.3. Given a nilmanifold G/Γ of dimension d , consider a basis $\mathcal{X} = \{X_1, \dots, X_d\}$ of $\mathfrak{g} = \log G$. \mathcal{X} is said to be a weak basis of rationality Q with respect to Γ if:

- There exist rationals c_{ijk} of height at most Q such that

$$[X_i, X_j] = \sum_k c_{ijk} X_k;$$

- There exists integer $1 \leq q \leq Q$ such that

$$q \cdot \mathbb{Z}^d \subseteq \psi_{\text{exp}, \mathcal{X}}(\Gamma) \subseteq q^{-1} \cdot \mathbb{Z}^d.$$

\mathcal{X} is a Mal'cev basis of $\log G$ with respect to Γ of rationality Q if:

- There exist rationals c_{ijk} of height at most Q such that

$$[X_i, X_j] = \sum_k c_{ijk} X_k;$$

- $\psi_{\mathcal{X}}(\Gamma) = \mathbb{Z}^d$.

We say that \mathcal{X} has the degree k nesting property if there exist $\ell_1 \leq \dots \leq \ell_k$ such that if $\mathfrak{g}_t = \text{span}_{\mathbb{R}}(X_{\ell_t+1}, \dots, X_m)$ then $[\mathfrak{g}, \mathfrak{g}] \subseteq \mathfrak{g}_1$, $[\mathfrak{g}, \mathfrak{g}_\ell] \subseteq \mathfrak{g}_{\ell+1}$ and $[\mathfrak{g}, \mathfrak{g}_k] = 0$.

Finally we say that a Mal'cev basis is adapted to a sequence of nesting subgroups $G = G_0 \geq G_1 \geq G_2 \geq \dots \geq G_\ell \geq \text{Id}_G$ if

$$\text{span}_{\mathbb{R}}(\{X_j: d - \dim(G_i) < j \leq d\}) = \log G_i$$

for $1 \leq i \leq \ell$.

We next state the definition of the Lipschitz property for a function on G/Γ .

Definition 3.4. We define a metric $d = d_{G, \mathcal{X}}$ on G by

$$d(x, y) := \inf \left\{ \sum_{i=1}^n \min(\|\psi(x_i x_{i+1}^{-1})\|, \|\psi(x_{i+1} x_i^{-1})\|) : n \in \mathbb{N}, x_1, \dots, x_{n+1} \in G, x_1 = x, x_{n+1} = y \right\},$$

where $\|\cdot\|$ denotes the ℓ^∞ -norm on \mathbb{R}^m , and define a metric on G/Γ by

$$d(x\Gamma, y\Gamma) = \inf_{\gamma, \gamma' \in \Gamma} d(x\gamma, y\gamma').$$

Furthermore, for any function $F: G/\Gamma \rightarrow \mathbb{C}$ we define

$$\|F\|_{\text{Lip}} := \|F\|_{\infty} + \sup_{x \neq y \in G/\Gamma} \frac{|F(x) - F(y)|}{d(x, y)}.$$

Given a function $F: G/\Gamma \rightarrow \mathbb{C}^D$ such that $F = (F_1, \dots, F_D)$ we define

$$\|F\|_{\text{Lip}} := \max_{1 \leq i \leq D} \|F_i\|_{\text{Lip}}.$$

Remark. Note that the metric on G is right-invariant. We may omit the subscript \mathcal{X} for the distance function when clear from context.

3.2. Complexity of nilmanifolds. We now define the complexity of a nilmanifold with respect to either a degree or a degree-rank filtration.

Definition 3.5. Let $s \geq 1$ be an integer and let $M \geq 1$. A nilmanifold G/Γ of degree s , dimension d , and complexity at most M consists of a degree s filtration of G along with a Mal'cev basis $\mathcal{X} = \{X_1, \dots, X_d\}$ of $\log G$ which satisfies the following:

- $\{X_1, \dots, X_d\}$ is a Mal'cev basis for $\log G$ with respect to Γ of rationality at most M ;
- \mathcal{X} is adapted to the sequence of subgroups $(G_i)_{i \in \mathbb{N}}$.

Analogously a nilmanifold G/Γ of degree-rank (s, r) , dimension d , and complexity at most M consists of a degree-rank (s, r) filtration of G along with a Mal'cev basis $\mathcal{X} = \{X_1, \dots, X_d\}$ of $\log G$ which satisfies the following:

- $\{X_1, \dots, X_d\}$ is a Mal'cev basis for $\log G$ with respect to Γ of rationality at most M ;
- \mathcal{X} is adapted to the sequence of subgroups $(G_i)_{i \in \text{DR}}$.

Remark. The only difference in complexity for a degree versus degree-rank filtration is that we require the Mal'cev basis to be adapted with respect to the appropriate filtration. This definition unfortunately *does not* extend to the case of multidegree filtrations since the set of subgroups do not nest in a total order. Furthermore note that a degree-rank nilmanifold of complexity M is also a degree nilmanifold of the same complexity by taking the associated degree filtration.

Finally, whenever discussing the complexity of nilmanifolds, this is always with respect to a given Mal'cev basis \mathcal{X} . We will abusively write phrases such as “nilmanifold G/Γ of complexity M ” throughout the paper; such a statement should always be understood with a corresponding implicitly provided adapted Mal'cev basis of the Lie algebra.

Remark. We will also in passing require the notion of a degree 0 nilmanifold. A degree 0 nilmanifold is simply the trivial group id_G . All scalar-valued functions on degree 0 nilmanifolds are constants and the Lipschitz norm is defined to be the absolute value of this constant.

We will next need the notion of a rational subgroup with respect to a Mal'cev basis; this will be crucial when giving the definition of complexity with respect to a multidegree filtration.

Definition 3.6. A closed, connected subgroup $G' \leq G$ is \mathbb{Q} -rational with respect to a basis $\mathcal{X} = \{X_1, \dots, X_m\}$ of $\log G$ if $\log G'$ has a basis $\mathcal{X}' = \{X'_1, \dots, X'_{m'}\}$ where $X'_i = \sum_{j=1}^m c_{ij} X_j$ for $1 \leq i \leq m'$ with $c_{ij} \in \mathbb{Q}$ having heights bounded by Q .

We will repeatedly use the following fact about rational subgroups without further comment.

Fact 3.7. Suppose G is a connected, simply connected nilpotent Lie group of step s and dimension d with a discrete cocompact subgroup Γ . Suppose that G/Γ has a weak basis \mathcal{X} of rationality

at most Q . Let H_1, \dots, H_j be subgroups which are each Q -rational and normal in G . Then

$$H = \bigvee_{i=1}^j H_i$$

is an $O_s(Q^{O_s(d^{O_s(1)})})$ -rational subgroup.

Proof. Let \mathcal{X}^i denote the underlying basis of H_i witnessing low height. By applying Baker–Campbell–Hausdorff, we have that $\log H$ is spanned by taking all $(\leq s)$ -fold commutators of elements in \mathcal{X}^i (possibly for different i). Each such element of the Lie algebra is easily seen to be a $O_s(Q^{O_s(d^{O_s(1)})})$ -rational combination of \mathcal{X} (using the weak basis property of \mathcal{X}). Taking a subset of these commutators which forms a basis of $\log H$ gives the desired result. \square

We are now in position to define the complexity of a multidegree nilsequence. This definition is admittedly rather artificial but is designed to be the most flexible given various lemmas scattered throughout the literature.

Definition 3.8. Consider a downset J with respect to the multidegree ordering on \mathbb{N}^k . Consider a group G with a multidegree filtration of degree $\subseteq J$. Recall the associated degree filtration

$$G_i = \bigvee_{\vec{v}: |\vec{v}|=i} G_{\vec{v}}$$

and define the associated degree to be $\sup_{\vec{v} \in J} |\vec{v}|$. We say a multidegree J nilmanifold G/Γ of dimension d with Mal'cev basis \mathcal{X} has complexity at most M if:

- $\{X_1, \dots, X_d\}$ is a Mal'cev basis for $\log G$ with respect to Γ of rationality at most M ;
- \mathcal{X} is adapted to the sequence of subgroups $(G_i)_{i \in \mathbb{N}}$;
- $G_{\vec{v}}$ is an M -rational subgroup for all $\vec{v} \in \mathbb{N}^k$.

We next note the trivial fact that complexity is bounded appropriately with respect to taking direct products; we implicitly invoke this when handling the complexity of direct products.

Fact 3.9. Consider nilmanifolds G/Γ , H/Γ' given degree s filtrations $(G_i)_{i \geq 0}$, $(H_i)_{i \geq 0}$ and adapted Mal'cev bases $\mathcal{X}, \mathcal{X}'$ each of complexity at most M . Then $(G \times H)/(\Gamma \times \Gamma')$ has complexity at most M with respect to the Mal'cev basis

$$\mathcal{X}^* = \{(X, 0): X \in \mathcal{X}\} \cup \{(0, X'): X' \in \mathcal{X}'\}.$$

\mathcal{X}^* may be adapted to the degree s filtration $G_i \times H_i$ by creating an ordering with suffixes

$$\{(X_j, 0): X_j \in \mathcal{X}, 0 \leq \dim(G) - j < \dim(G_i)\} \cup \{(0, X'_j): X'_j \in \mathcal{X}', 0 \leq \dim(H) - j < \dim(H_i)\}.$$

Furthermore given $F: G/\Gamma \rightarrow \mathbb{C}$ and $F': H/\Gamma' \rightarrow \mathbb{C}$ which are M -Lipschitz,

$$\tilde{F}((g, h)(\Gamma \times \Gamma')) := F(g\Gamma)F'(h\Gamma')$$

is $3M^2$ -Lipschitz on $(G \times H)/(\Gamma \times \Gamma')$. Analogous statements hold for degree-rank filtrations and multidegree filtrations.

We finally end by noting that quotients by normal subgroups of bounded rationality have appropriate complexity.

Lemma 3.10. Consider a nilmanifold G/Γ with G given a degree s filtration (G_i) and of complexity at most M with respect to an adapted Mal'cev basis \mathcal{X} .

Suppose that H is a normal subgroup of G which is M -rational with respect to \mathcal{X} . Then the quotient nilmanifold $(G/H)/(\Gamma/(\Gamma \cap H))$ may be given an adapted Mal'cev basis \mathcal{X}^* , where the degree s filtration is $(G_i/(G_i \cap H))$, which is an $M^{O_s(d^{O_s(1)})}$ -rational combination of

$$\mathcal{X}' = \{X \bmod \log H: X \in \mathcal{X}\}.$$

Analogous statements hold for degree-rank filtrations and multidegree filtrations. Finally if $H \leq Z(G)$ and F is an M -Lipschitz function on G/Γ which is H -invariant then F descends to $(G/H)/(\Gamma/(\Gamma \cap H))$ and is $M^{O_s(d^{O_s(1)})}$ -Lipschitz with respect to \mathcal{X}^* .

Proof. We may find a subset S such that

$$\mathcal{X}' = \{X_i \bmod \log H : X_i \in \mathcal{X}, i \in S\}$$

is a basis for $\log(G/H)$. Using

$$\exp(g) \bmod H = \exp(g \bmod \log H)$$

and [35, Lemma B.2] it follows that \mathcal{X}' is a weak Mal'cev basis for $(G/H)/(\Gamma/(\Gamma \cap H))$ of rationality $M^{O_s(d^{O_s(1)})}$. The desired result then follows from [35, Lemma B.11]. \square

3.3. Size of vertical and horizontal characters. We now define the size of vertical and horizontal characters. We first define the size of a horizontal character.

Definition 3.11. *Given a nilmanifold G/Γ and a Mal'cev basis \mathcal{X} , note that any horizontal character $\eta: G \rightarrow \mathbb{R}$ can be expressed in the form*

$$\eta(g) = k \cdot \psi(g)$$

for some $k \in \mathbb{Z}^{\dim(G)}$. We define the size of the horizontal character as $\|k\|_\infty$.

We next define the size of an i -th horizontal character.

Definition 3.12. *Consider a nilmanifold G/Γ with G given a degree-rank filtration of degree-rank (s, r) and a Mal'cev basis $\mathcal{X} = \{X_1, \dots, X_{\dim(G)}\}$ adapted to the degree-rank filtration. Note that any i -th horizontal character $\eta_i: G_{(i,1)} \rightarrow \mathbb{R}$ can be expressed in the form*

$$\eta_i(g) = k \cdot \psi(g)$$

with $k \in \mathbb{Z}^{\dim(G)}$ for some k which is nonzero only on coordinates between $\dim(G) - \dim(G_{(i,1)}) < j \leq \dim(G) - \dim(G_{(i,2)})$. We define the size of the i -th horizontal character as $\|k\|_\infty$.

We finally define the size of a vertical character.

Definition 3.13. *Consider a nilmanifold G/Γ with G given a degree filtration of degree k and a Mal'cev basis $\mathcal{X} = \{X_1, \dots, X_{\dim(G)}\}$ adapted to the degree filtration. Consider a continuous vertical character $\xi: T \rightarrow \mathbb{R}$ from a rational subgroup $T \leq Z(G)$. We define the height of ξ as*

$$\sup_{x \neq y \in T/(\Gamma \cap T)} \frac{|e(\xi \cdot x) - e(\xi \cdot y)|}{d_G(x\Gamma, y\Gamma)},$$

this will be denoted as $|\xi|$.

Remark. We now justify the terminology ‘‘height’’ given for the complexity of a vertical character. Suppose that G/Γ has complexity M (given \mathcal{X}) with respect to a degree filtration of degree d and T is Q -rational. We have that T has a Mal'cev basis which is a $(QM)^{O_k(d^{O(1)})}$ -rational combination of \mathcal{X} by [35, Lemma B.12]; denote this \mathcal{X}' . By [35, Lemma B.9], we have that for $x, y \in T$,

$$d_{G,\mathcal{X}}(x\Gamma, y\Gamma) \leq (QM)^{O_k(d^{O(1)})} d_{T,\mathcal{X}'}(x(\Gamma \cap T), y(\Gamma \cap T)) \leq (QM)^{O_k(d^{O(1)})} d_{G,\mathcal{X}}(x\Gamma, y\Gamma).$$

With respect to $\mathcal{X}' = \{X'_1, \dots, X'_{\dim(T)}\}$, we have that ξ is an integer vector and the definition of height is equivalent up to a multiplicative factor of $(QM)^{O_k(d^{O(1)})}$ to the height of this vector.

3.4. Correlation. We will also require the notion of a sequence being biased of some order.

Definition 3.14. A function $f: [N] \rightarrow \mathbb{C}^D$ is s -biased of correlation η , complexity M , and dimension d if there exists a nilmanifold G/Γ with a degree s filtration such that G has dimension at most d , G/Γ has complexity at most M , and there exists an M -Lipschitz function F and a polynomial sequence $g \in \text{poly}(\mathbb{Z}_{\mathbb{N}} \rightarrow G_{\mathbb{N}})$ such that

$$\|\mathbb{E}_{n \in [N]}[f(n) \otimes \overline{F(g(n)\Gamma)}]\|_{\infty} \geq \eta.$$

We will denote this as $f \in \text{Corr}(s, \eta, M, d)$.

3.5. Miscellaneous complexity notions. We will also require the following definition regarding smoothness norms of polynomial sequences.

Definition 3.15. Given $\vec{v} \in \mathbb{N}^k$ and $\vec{n} \in \mathbb{N}^k$, we define

$$\binom{\vec{n}}{\vec{v}} = \prod_{i=1}^k \binom{n_i}{v_i}.$$

Any polynomial sequence $g: \mathbb{Z}^k \rightarrow \mathbb{R}$ can be expressed uniquely as

$$g(\vec{n}) = \sum_{\vec{\ell} \in \mathbb{N}^k} \alpha_{\vec{\ell}} \binom{\vec{n}}{\vec{\ell}}$$

with $\alpha_{\vec{\ell}} \in \mathbb{R}$. We define

$$\|g\|_{C^{\infty}[N]} := \max_{\vec{\ell} \neq \vec{0}} N^{|\vec{\ell}|} \cdot \|\alpha_{\vec{\ell}}\|_{\mathbb{R}/\mathbb{Z}}$$

where $|\vec{\ell}| = \sum_{j=1}^k \ell_j$.

Remark. Note that the above definition is only sensitive to the values of $g \bmod 1$.

We now define when a polynomial sequence is rational and smooth.

Definition 3.16. Consider a nilmanifold G/Γ given either a degree, degree-rank, or multidegree filtration with Mal'cev basis \mathcal{X} and g a domain \mathbb{Z}^k polynomial sequence on G with respect to the given filtration. We say that g is (M, N) -smooth if:

- $d_{G, \mathcal{X}}(g(\vec{0}), \text{id}_G) \leq M$;
- $d_{G, \mathcal{X}}(g(\vec{v}), g(\vec{v} + \vec{e}_i)) \leq M \cdot N^{-1}$ for $\vec{v} \in [N]^k$ and $1 \leq i \leq k$.

We say that g is M -rational if for all $\vec{n} \in \mathbb{N}^k$ we have that

$$\psi_{\mathcal{X}}(g(\vec{n})) \in \frac{1}{M} \cdot \mathbb{Z}^{\dim(G)}.$$

4. PROOF OUTLINE

We are now in position to discuss the proof of Theorem 1.2; as our proof is closely modeled on that of Green, Tao, and Ziegler [29], the announcement of [28] may prove a useful starting point for certain readers. For various parts of this outline we will restrict to the case of the U^5 -inverse theorem and discuss the proof as if the analysis were performed with bracket polynomials.

4.1. Induction on degree and additive quadruples. Suppose that $f: [N] \rightarrow \mathbb{C}$ is 1-bounded such that

$$\|f\|_{U^5[N]} \geq \delta.$$

Via the inductive definition of the Gowers norm, we have for $\delta^{O(1)}N$ values of $h \in [N]$ that

$$\|\Delta_h f\|_{U^4[N]} \geq \delta^{O(1)}.$$

Applying Theorem 1.2 inductively (when converted to bracket polynomials; see e.g. [38, Proposition 1.4]) we may choose $d_1, d_2, d_3 \leq \log(1/\delta)^{O(1)}$ and coefficients $a_{i,h}$ etc. such that

$$\left| \mathbb{E}_{n \in [N]} \Delta_h f(n) \cdot e \left(\sum_{i=1}^{d_1} a_{i,h} n [b_{i,h} n] [c_{i,h} n] + \sum_{i=1}^{d_2} d_{i,h} n^2 [e_{i,h} n] + \sum_{i=1}^{d_3} f_{i,h} n [g_{i,h} n] + j_h n^3 + \ell_h n^2 + m_h n \right) \right| \geq \exp(-\log(1/\delta)^{O(1)});$$

we have padded with extra coefficients to make the dimensions d_i not h -dependent. Set

$$\overline{G_h(n)} = e \left(\sum_{i=1}^{d_1} a_{i,h} n [b_{i,h} n] [c_{i,h} n] + \sum_{i=1}^{d_2} d_{i,h} n^2 [e_{i,h} n] + \sum_{i=1}^{d_3} f_{i,h} n [g_{i,h} n] + j_h n^3 + \ell_h n^2 + m_h n \right).$$

For the sake of clarity, we will let $L_h(n)$ denote terms of degree ≤ 2 which are possibly h -dependent. We have

$$\left| \mathbb{E}_{n \in [N]} \Delta_h f(n) \cdot e \left(\sum_{i=1}^{d_1} a_{i,h} n [b_{i,h} n] [c_{i,h} n] + \sum_{i=1}^{d_2} d_{i,h} n^2 [e_{i,h} n] + j_h n^3 + L_h(n) \right) \right| \geq \exp(-\log(1/\delta)^{O(1)}).$$

The first crucial step, via a Cauchy–Schwarz argument due to Gowers [11] (see [27, Proposition 6.1] or Lemma 7.2) is that for many additive quadruples (h_1, h_2, h_3, h_4) , i.e. $h_1 + h_2 = h_3 + h_4$, we have

$$|\mathbb{E}_{n \in [N]} G_{h_1}(n) G_{h_2}(n + h_1 - h_4) \overline{G_{h_3}(n) G_{h_4}(n + h_1 - h_4)}| \geq \exp(-\log(1/\delta)^{O(1)}).$$

4.2. Sunflower and linearization for the top degree-rank. Via bracket polynomial manipulations, we see that the “top degree-rank” term of above expression is

$$\sum_{i=1}^{d_1} (a_{i,h_1} n [b_{i,h_1} n] [c_{i,h_1} n] + a_{i,h_2} n [b_{i,h_2} n] [c_{i,h_2} n] - a_{i,h_3} n [b_{i,h_3} n] [c_{i,h_3} n] - a_{i,h_4} n [b_{i,h_4} n] [c_{i,h_4} n]).$$

The heart of the proof is demonstrating that these “top degree-rank terms line up” in an appropriate sense across a dense set of indices H . Such a conclusion is at least plausible since for generic coefficients the associated bracket polynomial equidistributes mod 1, which would violate the given. One possibility where the top degree-rank term is exactly zero is when we can write $a_{i,h_1} = a_i h_1$, $b_{i,h_1} = b_i^*$, $c_{i,h_1} = c_i^*$. The heart of the matter is that, up to controlled modifications, this is the only way for that to occur in a robust sense.

The first modification is that we can replace in the above example the expression $a_{i,h_1} = a_i h_1$ with $a_{i,h_1} = \Theta_i \{ \Theta'_i h_1 \}$ or more generally a bracket linear form. The second modification is that we may not get a description that respects the presented structure of the sum. Instead the coordinates of the bracket linear form may only appear in these “fixed”, “fixed”, “bracket linear” triples after a linear change of variables. We prove the existence of this structure in two steps, as in [29]. The first step proves that the bracket form is “fixed”, “fixed”, “ h -dependent” and the second step then proves that the “ h -dependent” part in fact has a bracket linear structure. These steps will fall under the names *sunflower* and *linearization* respectively.

4.3. Degree-rank iteration. Once we have learned this refined form for $\sum_{i=1}^{d_1} a_{i,h}n[b_{i,h}n][c_{i,h}n]$, we iterate and then learn the refined form for the next highest degree-rank term $\sum_{i=1}^{d_2} d_{i,h}n^2[e_{i,h}n]$, and then finally we learn the refined form for $j_h n^3$. Given these refined forms, Green, Tao, and Ziegler prove that the top degree terms in fact have the form of a multidegree $(1, 3)$ nilsequence (in variables h and n). Finally given such a correlation, a symmetrization argument as in [29] concludes the proof. We remark here that while terms such as $a_{i,h}$ and $e_{i,h}$ correspond to Taylor coefficients on the top degree horizontal torus, terms such as $d_{i,h}$ belong to the second horizontal torus, and j_h to the third horizontal torus. Furthermore to handle terms of the form $\sum_{i=1}^{d_2} d_{i,h}n^2[e_{i,h}n]$ correctly we must realize such terms via a degree-rank $(3, 2)$ nilmanifold, hence the need for the finer degree-rank notion.

4.4. Nilcharacters and horizontal tori. We now make this description more precise in terms of nilcharacters and horizontal tori. Let $F(g_h(n)\Gamma) = G_h(n)$ be a nilcharacter of degree-rank (s, r) ; here $e(an[bn][cn])$ should be thought of as an “almost” degree-rank $(3, 3)$ nilcharacter and $e(an[bn^2])$ as an “almost” degree-rank $(3, 2)$ nilcharacter. The sunflower step shows proves that the nilsequence $F(g_h(n)\Gamma)$ can be realized as a bracket polynomial whose top degree-rank part is a sum of terms with $(r - 1)$ iterated brackets where each term consists of $(r - 1)$ h -independent phases of g_h , and possibly one h -dependent phase of g_h . Here, “phase” will correspond to components of the Taylor coefficients of g_h , $\text{Taylor}_i(g_h)$. This corresponds to showing that the i -th horizontal torus $G_{(i,1)}/G_{(i,2)}$ contains vector spaces $V_{i,\text{Dep}} \leq V_i$ such that:

- $\text{Taylor}_i(g_h) - \text{Taylor}_i(g_{h'}) \in V_{i,\text{Dep}}$ and $\text{Taylor}_i(g_h) \in V_i$;
- If $i_1 + \dots + i_r = s$, then $[v_{i_1}, v_{i_2}, \dots, v_{i_r}] = 0$ whenever $v_{i_j} \in V_{i_j}$ and there are at least two indices j such that $W_{i_j} = V_{i_j,\text{Dep}}$.

Here we have implicitly descended an iterated commutator to the vector spaces $G_{(i,1)}/G_{(i,2)}$ which corresponds to a multilinear form in this case. Such a result is proven via combining quantitative equidistribution theory of nilsequences [35, 36] with a “Furstenberg–Weiss argument” as in [27, 29, 36]; see [49] for further examples of the Furstenberg–Weiss argument.

The linearization step then proves that the remaining h -dependent phases are “bracket linear” in h . In practice, we require an additional case that the h -dependent phase may be a *petal* phase: a top degree-rank term with the petal phase can be realized as a “lower order term”, or more precisely a bracket phase with at most $(r - 2)$ iterated brackets or of total degree at most $s - 1$. Thus, the statement we ultimately prove is that we may decompose a subspace of the i -th horizontal torus into the sum of three linearly disjoint vector spaces $W_{i,*}$, $W_{i,\text{Lin}}$, and $W_{i,\text{Pet}}$ such that:

$$\begin{aligned} \text{Taylor}_i(g_h) &\in W_{i,*} + W_{i,\text{Lin}} + W_{i,\text{Pet}}, \\ \text{Taylor}_i(g_h) - \text{Taylor}_i(g_{h'}) &\in W_{i,\text{Lin}} + W_{i,\text{Pet}}, \end{aligned}$$

and the projection of $\text{Taylor}_i(g_h)$ onto $W_{i,\text{Lin}}$ is bracket linear. In addition, we require that if $i_1 + \dots + i_r = s$, then $[v_{i_1}, v_{i_2}, \dots, v_{i_r}] = 0$ whenever $v_i \in W_{i,*} + W_{i,\text{Lin}} + W_{i,\text{Pet}}$ and either $v_{i_j} \in W_{i_j,\text{Pet}}$ for at least one index j or $v_{i_j} \in W_{i_j,\text{Lin}}$ for at least two distinct indices j . Thus even though we have not improved our understanding of the Taylor coefficients on $W_{i,\text{Pet}}$ we have improved the vanishing of the top degree-rank commutator bracket on this vector space. The linearization step is proved by a combination of quantitative equidistribution theory of nilmanifolds [35, 36] and inverse sumset theory. We refer the reader to [36] for a simpler case of the argument given here.

4.5. Quantitative bounds. The heart of this paper is performing the sunflower and linearization steps efficiently. Green, Tao, and Ziegler [29] accomplish this (when unwinding the correspondence between nilmanifolds and bracket polynomials) via iteratively learning relations

between the coefficients $a_{i,h}, b_{i,h}, c_{i,h}$ and performing a dimension reduction argument.¹ Furthermore, the underlying equidistribution theorem used in the work of Green, Tao, and Ziegler [29], proven in work of Green and Tao [24], relies on an induction on dimension argument. The use of any induction on dimension argument essentially immediately results in $O(s)$ iterated logarithms and thus must be avoided.

The use of induction on dimension in the equidistribution theorem was avoided in work of the first author [35, 36]. The key point in Sections 8 and 9 therefore is to perform the sunflower and linearization steps without any use of induction on dimension. The precise details, while mainly utilizing elementary linear algebra, require a bit of precision. This argument, extending the case of the U^4 -inverse theorem from [36], demonstrates that a dimension-independent number of applications of equidistribution theory is sufficient to derive the necessary decrease in degree-rank. (Note that the argument in [29] morally uses that one can in fact assume that there are no “short linear relations” between various coefficients, but such a result necessitates exponential in dimension dependencies in the exponent.) Another crucial point in our work is that the length of the associated bracket linear form that is obtained not “very long”. This is, by now, a standard consequence of the quasi-polynomial bounds of Sanders [44] towards the polynomial Bogolyubov conjecture.

We finally remark that the quantitative equidistribution theorem we use is slightly different than the one derived in work of the first author [35, 36]. The work of the first author is most naturally phrased as factoring ill-distributed polynomial sequences into a smooth part, a rational part, and a polynomial sequence which (up to taking a certain quotient) lives in a lower *step* nilmanifold. For our purposes, it is critical to instead lower the *degree* of the nilmanifold. This is most easily seen from the above bracket polynomial example where we are attempting to linearize a function of the form

$$e\left(\sum_{i=1}^{d_3} f_{i,h}n[g_{i,h}n] + j_h n^3 + \ell_h n^2 + m_h n\right).$$

At this step we wish to linearize $j_h n^3$ instead of handling the terms $f_{i,h}n[g_{i,h}n]$; the $j_h n^3$ term, while having the highest degree, does not correspond to the highest step part of the nilmanifold. This phenomenon only occurs when proving the U^{s+1} -inverse theorem for $s \geq 4$. Thus a crucial ingredient in our work is bootstrapping, as a black box, the efficient version of equidistribution with respect to step in order to obtain an efficient version of equidistribution with respect to degree; this is Theorem 5.4.

4.6. Organization of the paper II. In Section 5 we prove the necessary quantitative equidistribution theorem with respect to degree. In Section 6, we perform the setup and give various definitions which will be used to perform the sunflower and linearization steps. In Section 7, we derive that many additive quadruples exhibit a bias. In Section 8 we perform the sunflower step while in Section 9 we perform the linearization step. In Sections 10 and 11 we then convert information regarding the Taylor coefficients into correlation with a multidegree $(1, s - 1)$ nilsequence and a nilsequence of lower degree-rank. Iterating this argument we eventually obtain correlation with a multidegree $(1, s - 1)$ nilsequence. In Section 12, we symmetrize this nilsequence to obtain Theorem 1.2.

Appendix A collects certain standard results regarding approximate homomorphisms (this is ultimately where work of Sanders [45] is invoked). In Appendix B, we collect a number of

¹This is performed in [29, Section 10] via a “rank minimality” argument; this requires passing to an ultralimit. When performed in finitary language this becomes a dimension reduction argument and is also present in the proof of [29, Theorem D.5].

miscellaneous propositions which are deferred throughout the paper. Finally in Appendix C we collect a number of propositions regarding nilcharacters.

5. EFFICIENT EQUIDISTRIBUTION THEORY OF NILSEQUENCES

In order to state the primary equidistribution input of this paper we will need the notion of when an element in $G/[G, G]$ and a horizontal character are orthogonal.

Definition 5.1. *Consider a nilmanifold G/Γ , a horizontal character $\eta: G \rightarrow \mathbb{R}$, and $w \in G/[G, G]$. We say that η and w are orthogonal if $\eta(w) = 0$.*

The primary equidistribution input into our results will be the following result of the first author [35, Theorem 3]. This result is ultimately the driving force of this paper.

Theorem 5.2. *Fix an integer $\ell \geq 1$, $\delta \in (0, 1/10)$, $M, d \geq 1$, and $F: G/\Gamma \rightarrow \mathbb{C}$. Suppose that G is a dimension d , at most s -step nilpotent Lie group with a given degree k filtration, and nilmanifold G/Γ is complexity at most M with respect to this filtration. Let g be a polynomial sequence on G with respect to this filtration.*

Furthermore suppose that $\|F\|_{\text{Lip}} \leq 1$ and F has $G_{(s)}$ -vertical frequency ξ such that the height of ξ is bounded by M/δ . Suppose that $N \geq (M/\delta)^{\Omega_{k,\ell}(d^{\Omega_{k,\ell}(1)})}$ and

$$|\mathbb{E}_{\vec{n} \in [N]^\ell} F(g(\vec{n})\Gamma)| \geq \delta.$$

There exists an integer $0 \leq r \leq \dim(G/[G, G])$ such that:

- *We have horizontal characters $\eta_1, \dots, \eta_r: G \rightarrow \mathbb{R}$ with heights bounded by $(M/\delta)^{O_{k,\ell}(d^{O_{k,\ell}(1)})}$;*
- *For all $1 \leq i \leq r$, we have $\|\eta_i \circ g\|_{C^\infty[N]} \leq (M/\delta)^{O_{k,\ell}(d^{O_{k,\ell}(1)})}$*
- *For any $w_1, \dots, w_s \in G/[G, G]$ such that w_i are orthogonal to all of η_1, \dots, η_r , we have*

$$\xi([[[w_1, w_2], w_3], \dots, w_s]) = 0.$$

Remark 5.3. Note that $G_{(s)}$ (and in fact any group in the lower central series) is seen to be $O_{s,k}(M^{O_{s,k}(1)})$ -rational due to Lemma 2.1. This guarantees that the height definition used in [35] and here are compatible.

Remark. Let $W = \bigcap_{i=1}^r \ker(\eta_i)$. The crucial property of the lemma output is that

$$\tilde{G} := W / \ker(\xi)$$

is trivially seen to be at most $(s-1)$ -step nilpotent. This is due to the fact that defining $W = W_0 = W_1$ and $W_j = [W_1, W_{j-1}]$ for $j \geq 2$ yields $W_{(s)} \leq G_{(s)}$ and $\xi(W_{(s)}) = 0$. Additionally, the statement in [35, Theorem 3] assumes G is exactly s -step nilpotent and ξ is nonzero. In the case when G is strictly less than s -step nilpotent, taking no horizontal characters (i.e., $W = G$) gives the desired statement. Furthermore when ξ is zero we may similarly take no horizontal characters and note that the final statement is vacuous.

The following variant of Theorem 5.2 will essentially be the primary equidistribution tool in our paper. For the sake of argumentation, we first prove the result in the case when the vertical frequency considered lives on a 1-dimensional torus and then bootstrap to the general case.

This theorem and its proof are motivated by [29, Lemma E.11]. The key point is that Theorem 5.2 allows us to give a procedure that relies on an induction on step rather than an induction on dimension. The main technical issue is at each stage we pass to a quotient group given by quotienting the kernel of a certain vertical character and thus we must iteratively “lift” these factorizations.

Theorem 5.4. *Let $\ell \geq 1$ be an integer, $\delta \in (0, 1/10)$, $M \geq 1$, and $F: G/\Gamma \rightarrow \mathbb{C}$. Suppose that G is dimension d , is s -step nilpotent with a given degree k filtration, and nilmanifold G/Γ is complexity at most M with respect to this filtration.*

Suppose that $T \leq Z(G)$ is a 1-dimensional subgroup of the center which is M -rational with respect to G . Further suppose that F has a nonzero T -vertical character ξ with $|\xi| \leq M/\delta$, $\|F\|_{\text{Lip}} \leq M$, $N \geq (M/\delta)^{\Omega_{k,\ell}(d^{\Omega_{k,\ell}(1)})}$, and g is a polynomial sequence with respect to the degree k filtration such that $g(0) = \text{id}_G$. Then if

$$|\mathbb{E}_{\vec{n} \in [N]^\ell} F(g(\vec{n})\Gamma)| \geq \delta$$

there exists a factorization

$$g = \varepsilon g' \gamma$$

such that:

- $\varepsilon(0) = g'(0) = \gamma(0) = \text{id}_G$;
- g' lives in an $(M/\delta)^{O_{k,\ell}(d^{O_{k,\ell}(1)})}$ -rational subgroup H such that $H \cap T = \text{Id}_G$;
- γ is an $(M/\delta)^{O_{k,\ell}(d^{O_{k,\ell}(1)})}$ -rational polynomial sequence;
- ε is an $((M/\delta)^{O_{k,\ell}(d^{O_{k,\ell}(1)})}, N)$ -smooth polynomial sequence.

Proof. The proof proceeds iteratively “simplifying” g to live on successively lower-step nilmanifolds. We treat ℓ as constant and allow implicit constants to depend on ℓ .

Step 1: Iteration setup. We will define a sequence of parameters M_i, δ_i and Q_i, N_i, v_i (where the domain of \vec{n} at stage i will be $v_i + Q_i \cdot [N_i]^\ell$) satisfying:

$$M_{i+1} \leq (M_i/\delta_i)^{O_k(d^{O_k(1)})}, \quad \delta_{i+1} \geq (\delta_i/M_i)^{O_k(d^{O_k(1)})};$$

$$Q_{i+1} \leq Q_i \cdot (M_i/\delta_i)^{O_k(d^{O_k(1)})}, \quad N_{i+1} \geq N_i \cdot (\delta_i/M_i)^{O_k(d^{O_k(1)})}, \quad Q_{i+1} \cdot N_{i+1} + \|v_{i+1}\|_\infty \leq N.$$

During the iteration, we have a sequence of nilpotent Lie groups

$$G^0, G^1, \dots, G^t, \dots$$

such that G^t is at most $(s-t)$ -step nilpotent with associated lattice Γ^t and is complexity at most M_t . This in particular will imply that there are at most s stages in the iteration. We also maintain a sequence of subgroups

$$K^0, \dots, K^t, \dots$$

which are M_t -rational subgroups of G .

We will define homomorphisms $\pi_{t+1}: G^t \rightarrow G^t/\ker(\xi_t) =: \tilde{G}^{t+1}$, where ξ_t is a $G_{(s-t)}^t$ -frequency (recall $H_{(i)}$ denotes the lower central series filtration of a group H). G^{t+1} will be an appropriately rational subgroup of \tilde{G}^{t+1} . We will always maintain the invariant that $\ker(\xi_t) \cap (\pi_t \circ \dots \circ \pi_1(T)) = \text{Id}_{G^t}$. We will furthermore maintain that function F_t has a $\pi_t \circ \dots \circ \pi_1(T)$ -character given by descending ξ on G via $\pi_t \circ \dots \circ \pi_1$.

We inductively maintain the following pair of relations:

- $\pi_t \circ \dots \circ \pi_1(K^t) \leq G^t$;
- $\pi_t \circ \dots \circ \pi_1(g_t) = \tilde{g}_t$;

where g_t and \tilde{g}_t are polynomial sequences living in K^t and G^t respectively.

The iteration terminates when $G^t \cap (\pi_t \circ \dots \circ \pi_1(T)) = \text{Id}_{G^t}$. Before termination note that $G^t \cap (\pi_t \circ \dots \circ \pi_1(T)) = \pi_t \circ \dots \circ \pi_1(T)$ since $\pi_t \circ \dots \circ \pi_1(T)$ is 1-dimensional. Note that this in particular ensures that before the termination of the iteration, $\pi_t \circ \dots \circ \pi_1(T)$ is well-defined even though π_j is not fully defined on the image of π_{j-1} ! Using the invariant that $\ker(\xi_t) \cap (\pi_t \circ \dots \circ \pi_1(T)) = \text{Id}_{G^t}$ we also have that ξ (defined on T) naturally descends to G^t . We define $J^t = \pi_t \circ \dots \circ \pi_1(T)$.

Furthermore at each stage of the iteration we have that

$$g_t = \varepsilon_{t+1} \cdot g_{t+1} \cdot \gamma_{t+1}$$

where:

- ε_{t+1} and γ_{t+1} are polynomial sequences lying in K_t ;
- g_{t+1} is a polynomial sequence lying in K_{t+1} ;
- γ_{t+1} is M_{t+1} -rational;
- ε_{t+1} is (M_{t+1}, N) -smooth.

Finally, in each stage of the iteration we will maintain a function $F_t: G^t/\Gamma^t \rightarrow \mathbb{C}$ such that

$$|\mathbb{E}_{\vec{n} \in v_t + Q_t \cdot [N_t]^\ell} [F_t(\tilde{g}_t(\vec{n})\Gamma^t)]| \geq \delta_t.$$

Throughout the iterations, nilmanifolds at stage i will have complexity bounded by M_i , F_i is M_i -Lipschitz, and various horizontal and vertical characters constructed will have size and height bounded by M_i . The starting conditions are $G^0 = G$, $\Gamma^0 = \Gamma$, $M_0 = M$, $F_0 = F$, $N_0 = N$, $v_0 = 0$, $Q_0 = 1$, $\delta_0 = \delta$, and $K^0 = G$ (and $J^0 = T$), and $g_0 = \tilde{g}_0 = g$.

Step 2: Applying equidistribution. We now run a single step of the iteration. We have

$$|\mathbb{E}_{\vec{n} \in v_t + Q_t \cdot [N_t]^\ell} [F_t(\tilde{g}_t(n)\Gamma^t)]| \geq \delta_t.$$

By definition, we have that F_t has a J^t -frequency (a descent of ξ); this is not sufficient to apply Theorem 5.2. We perform an additional Fourier-analytic step to obtain a $G_{(s-t)}^t$ -vertical frequency. Since F_t is M_t -Lipschitz, via [35, Lemma A.6] we may write

$$F_t(z\Gamma^t) = \sum_{|\xi'| \leq (M_t/\delta_t)^{O_k(d^{O_k(1)})}} F_{\xi',t}(z\Gamma^t) + \tau(z\Gamma^t)$$

such that

- $F_{\xi',t}$ has $G_{(s-t)}^t$ -vertical frequency ξ' ;
- $\|\tau\|_\infty \leq \delta_t/2$;
- $F_{\xi',t}$ is $(M_t/\delta_t)^{O_k(d^{O_k(1)})}$ -Lipschitz on G^t/Γ^t .

Given this representation, recall that F_t has ξ (appropriately descended) as a J^t -vertical frequency. We abusively write this as ξ . Therefore

$$\begin{aligned} F_t(z\Gamma^t) &= \int_{g \in J^t/\Gamma^t} e(-\xi(g)) F_t(zg\Gamma^t) dJ^t(g) \\ &= \sum_{|\xi'| \leq (M_t/\delta_t)^{O_k(d^{O_k(1)})}} \int_{g \in J^t/\Gamma^t} e(-\xi(g)) F_{\xi',t}(zg\Gamma^t) dJ^t(g) + \int_{g \in J^t/\Gamma^t} e(-\xi(g)) \tau(zg\Gamma^t) dJ^t(g) \\ &= \sum_{|\xi'| \leq (M_t/\delta_t)^{O_k(d^{O_k(1)})}} \tilde{F}_{\xi',t}(z\Gamma^t) + \int_{g \in J^t/\Gamma^t} e(-\xi(g)) \tau(zg\Gamma^t) dJ^t(g), \end{aligned}$$

where dJ^t represents the Haar measure on J^t/Γ^t . Thus F_t may be decomposed into a sum of functions with $G_{(s-t)}^t$ -vertical characters up to an L^∞ error of $\delta_t/2$. Furthermore, each vertical character ξ' in question must agree with ξ on $J^t \cap G_{(s-t)}^t$. If not, then the corresponding integral in the second line will average to 0 and we may remove it.

Applying Pigeonhole, there exists $|\xi'| \leq (M_t/\delta_t)^{O_k(d^{O_k(1)})}$ such that

$$(5.1) \quad |\mathbb{E}_{\vec{n} \in v_t + Q_t \cdot [N_t]^\ell} [\tilde{F}_{\xi',t}(\tilde{g}_t(\vec{n})\Gamma^t)]| \geq (\delta_t/M_t)^{O_k(d^{O_k(1)})}.$$

We have the following trichotomy:

- ξ' is nonzero and $J^t \cap G_{(s-t)}^t = J^t$;
- ξ' is nonzero and $J^t \cap G_{(s-t)}^t = \text{Id}_{G^t}$;

- $\xi' = 0$ in $\widehat{G_{(s-t)}^t}$.

We define $\pi_{t+1}: G^t \rightarrow G^t / \ker(\xi') =: \widetilde{G}^{t+1}$ (in particular we let $\xi_{t+1} = \xi'$). Let $\widetilde{\Gamma}^{t+1} = \Gamma^t / (\Gamma^t \cap \ker(\xi'))$. We now apply Theorem 5.2 to (5.1), obtaining horizontal characters $\eta_1, \dots, \eta_r: G^t \rightarrow \mathbb{R}$. Let their common kernel be H^* and let $G^{t+1} = \pi_{t+1}(H^*) \leq \widetilde{G}^{t+1}$. By applying [35, Lemma A.1], we obtain a factorization of $\widetilde{g}_t(Q_t n + v_t)$ into three nilsequences which are “smooth”, supported on a rational subgroup, and “rational”. We may change variables and then apply π_{t+1} to obtain

$$\pi_{t+1}(\widetilde{g}_t) =: \varepsilon_{t+1}^* g_{t+1}^* \gamma_{t+1}^*$$

where:

- $g_{t+1}^* \in G^{t+1}$, and G^{t+1} is at most $(s-t-1)$ -step nilpotent. Furthermore G^{t+1} is trivially seen to be $(M_t/\delta_t)^{O_k(d^{O_k(1)})}$ -rational with respect to \widetilde{G}_{t+1} ;
- γ_{t+1}^* is an $(M_t/\delta_t)^{O_k(d^{O_k(1)})}$ -rational polynomial sequence within \widetilde{G}^{t+1} ;
- ε_{t+1}^* is $((M_t/\delta_t)^{O_k(d^{O_k(1)})}, N)$ -smooth.

We remark that changing variables is easily seen to not affect the smoothness and rationality in a substantial manner due to the bounds on Q_t . We can see that the step of G^{t+1} decreases appropriately.

Step 3: Lifting the factorization data. Note that G^{t+1} can be defined via a set of horizontal characters $\eta'_1, \dots, \eta'_{r'}$ of \widetilde{G}^{t+1} such that

$$G^{t+1} = \{x \in \widetilde{G}^{t+1} : \eta'_i(x) = 0 \text{ for all } 1 \leq i \leq r'\}.$$

(If $s-t \geq 2$ one can directly descend the η_i since $\ker(\xi') \leq G_{(s-t)}^t \leq [G^t, G^t]$ while if $s-t = 1$ then G^t is abelian and the η' obviously exist.)

We define

$$K^{t+1} = \{x \in K^t : \eta'_i(\pi_{t+1} \circ \pi_t \circ \dots \circ \pi_1(x)) = 0 \text{ for all } 1 \leq i \leq r'\}.$$

The trivial (but key) point is that $\pi_{t+1} \circ \pi_t \circ \dots \circ \pi_1(K^{t+1}) \leq G^{t+1}$. The only issue is noting that the map is well-defined; this is because $\pi_t \circ \dots \circ \pi_1(K^t) \leq G^t$ by induction so that we are allowed to apply π_{t+1} to any such values.

Recall by induction that

$$\pi_t \circ \dots \circ \pi_1(g_t) = \widetilde{g}_t$$

and thus

$$\pi_{t+1} \circ \dots \circ \pi_1(g_t) = \varepsilon_{t+1}^* g_{t+1}^* \gamma_{t+1}^*.$$

Applying η'_i , we find that that there exists a nonzero integer $T_i \leq (M_t/\delta_t)^{O_k(d^{O_k(1)})}$ such that

$$(5.2) \quad \|T_i \cdot \eta'_i(\pi_{t+1} \circ \dots \circ \pi_1(g_t))\|_{C^\infty[N]} \leq (M_t/\delta_t)^{O_k(d^{O_k(1)})}.$$

We now claim that $\eta'_i(\pi_{t+1} \circ \dots \circ \pi_1(\cdot))$ is a horizontal character on K^t . It is a homomorphism since the π_i are homomorphisms and it is well-defined by the above. In addition, we may inductively show that $\pi_{t+1} \circ \dots \circ \pi_1(\Gamma \cap K^t) \leq \widetilde{\Gamma}^{t+1}$ and hence $\eta(\pi_{t+1} \circ \dots \circ \pi_1(\Gamma \cap K^t)) \leq \mathbb{Z}$, which verifies the property of being a horizontal character. That the horizontal character has appropriately bounded height is an immediate consequence of induction and the fact that $|\xi'| \leq (M_t/\delta_t)^{O_k(d^{O_k(1)})}$.

Now we use this data to construct the required factorization. By applying [35, Lemma A.1] with the horizontal characters $T_i \cdot \eta'_i(\pi_{t+1} \circ \dots \circ \pi_1)$ defined on K^t with the hypotheses (5.2), we may write

$$g_t = \varepsilon'_{t+1} g'_{t+1} \gamma'_{t+1}$$

where:

- g'_{t+1} takes values in K_{t+1} ;
- ε'_{t+1} and γ'_{t+1} take values in K_t ;
- γ'_{t+1} is an $(M_t/\delta_t)^{O_k(d^{O_k(1)})}$ -rational polynomial sequence;
- ε'_{t+1} is $((M_t/\delta_t)^{O_k(d^{O_k(1)})}, N)$ -smooth.

Then Q' denote the least common multiple of the periods of the ℓ different directions for $\gamma'_{t+1}\Gamma$; note that such periods exist and we have $Q' \leq (M_t/\delta_t)^{O_k(d^{O_k(1)})}$ by [35, Lemma B.14]. Divide $v_t + Q_t \cdot [N_t]^\ell$ into boxes of common difference $Q_t Q'$. By Pigeonhole there exists v' such that

$$|\mathbb{E}_{\vec{n} \in v' + Q' Q_t \cdot [N_t/Q']^\ell} [\tilde{F}_{\xi', t}(\tilde{g}_t(\vec{n})\Gamma^t)]| \geq (\delta_t/M_t)^{O_k(d^{O_k(1)})}.$$

Note that

$$\tilde{g}_t = \pi_t \circ \cdots \circ \pi_1(g_t) = \pi_t \circ \cdots \circ \pi_1(\varepsilon'_{t+1}) \cdot \pi_t \circ \cdots \circ \pi_1(g'_{t+1}) \cdot \pi_t \circ \cdots \circ \pi_1(\gamma'_{t+1}).$$

Since the differences we are considering are divisible by Q' , there is γ_{Rep} such that

$$\gamma_{\text{Rep}}^{-1} \gamma'_{t+1}(v' + Q' Q_t \cdot \vec{n}) \in \Gamma$$

for all $\vec{n} \in \mathbb{Z}^\ell$, where $\gamma_{\text{Rep}} \in K^t$ and $d_G(\gamma_{\text{Rep}}, \text{id}_G) \leq (M_t/\delta_t)^{O_k(d^{O_k(1)})}$. Since $\pi_t \circ \cdots \circ \pi_1(\Gamma \cap K^t) \leq \Gamma^t$ we have that

$$|\mathbb{E}_{\vec{n} \in v' + Q' Q_t \cdot [N_t/Q']^\ell} [\tilde{F}_{\xi', t}(\pi_t \circ \cdots \circ \pi_1(\varepsilon'_{t+1} \gamma_{\text{Rep}}) \cdot \pi_t \circ \cdots \circ \pi_1(\gamma_{\text{Rep}}^{-1} \gamma'_{t+1} \gamma_{\text{Rep}})\Gamma^t)]| \geq (\delta_t/M_t)^{O_k(d^{O_k(1)})}.$$

Step 4: Completing the induction. The first key polynomial sequence we shall define is

$$g_{t+1} = \gamma_{\text{Rep}}^{-1} \cdot g'_{t+1} \cdot \gamma_{\text{Rep}}.$$

Note that K^{t+1} is normal within K^t and since $\gamma_{\text{Rep}} \in K^t$ we have that g_{t+1} takes on values in K^{t+1} as desired. Further let $\varepsilon_{t+1} = \varepsilon'_{t+1} \cdot \gamma_{\text{Rep}}$ and $\gamma_{t+1} = \gamma_{\text{Rep}}^{-1} \cdot \gamma'_{t+1}$; these are trivially seen to lie in K_t and have the necessary rationality and smoothness properties due to the above analysis.

We now break $[N_t/Q']^\ell$ into a collection of boxes of length $N_{t+1} \geq N_t/Q' \cdot (M_t/\delta_t)^{-O_k(d^{O_k(1)})}$. There exists a box such that

$$|\mathbb{E}_{\vec{n} \in v'' + Q' Q_t \cdot [N_{t+1}]^\ell} [\tilde{F}_{\xi', t}(\pi_t \circ \cdots \circ \pi_1(\varepsilon'_{t+1} \cdot \gamma_{\text{Rep}}) \cdot \pi_t \circ \cdots \circ \pi_1(g_{t+1})\Gamma^t)]| \geq (\delta_t/M_t)^{O_k(d^{O_k(1)})}.$$

Taking N_{t+1} sufficiently small, we may replace the initial “smooth” polynomial sequence ε_{t+1}^* by $\varepsilon^* \in K^t$ where $d_G(\varepsilon^*, \text{id}_G) \leq (M_t/\delta_t)^{O_k(d^{O_k(1)})}$ such that

$$|\mathbb{E}_{\vec{n} \in v'' + Q' Q_t \cdot [N_{t+1}]^\ell} [\tilde{F}_{\xi', t}(\pi_t \circ \cdots \circ \pi_1(\varepsilon^*) \cdot \pi_t \circ \cdots \circ \pi_1(g_{t+1})\Gamma^t)]| \geq (\delta_t/M_t)^{O_k(d^{O_k(1)})}.$$

The new function F_{t+1} is given by descending $g \mapsto \tilde{F}_{\xi', t}(\pi_t \circ \cdots \circ \pi_1(\varepsilon^*) \cdot g\Gamma^t)$ from G^t to \tilde{G}^{t+1} (and later we may implicitly restrict to G^{t+1}). Explicitly, for $g \in G^t$ we have

$$\tilde{F}_{\xi', t}(\pi_t \circ \cdots \circ \pi_1(\varepsilon^*)g\Gamma^{t+1}) = F_{t+1}(\pi_{t+1}(g)\tilde{\Gamma}^{t+1})$$

which is possible because $\tilde{F}_{\xi', t}$ has vertical frequency ξ' . Therefore we have

$$(5.3) \quad |\mathbb{E}_{\vec{n} \in v'' + Q' Q_t \cdot [N_{t+1}]^\ell} [F_{t+1}(\pi_{t+1} \circ \pi_t \circ \cdots \circ \pi_1(g_{t+1}(\vec{n}))\tilde{\Gamma}^{t+1})]| \geq (\delta_t/M_t)^{O_k(d^{O_k(1)})}.$$

We let

$$\tilde{g}_{t+1} := \pi_{t+1} \circ \pi_t \circ \cdots \circ \pi_1(g_{t+1})$$

and we may replace $\tilde{\Gamma}^{t+1}$ with $\Gamma^{t+1} = \tilde{\Gamma}^{t+1} \cap G^{t+1}$ in (5.3).

We now check that $\ker(\xi') \cap J^t = \text{id}_{G^t}$, which is one of the invariants we are maintaining (we take $\xi_t = \xi'$). We will have to distinguish between cases:

- If ξ' is nonzero and $J^t \cap G_{(s-t)}^t = J^t$ note that $\ker(\xi') \cap J^t = \text{Id}_{G^t}$. This is due to the fact that ξ' restricted to J^t is (the descended version of) ξ which is nonzero as given.

- If ξ' is nonzero and $J^t \cap G_{(s-t)}^t = \text{id}_G^t$ then note that $\ker(\xi') \cap J^t \leq J^t \cap G_{(s-t)}^t = \text{Id}_{G^t}$.
- If $\xi' = 0$ then note that as ξ (appropriately descended) was nonzero we have that $J^t \cap G_{(s-t)}^t = \text{Id}_G^t$ is forced in this case. The result then follows as in the previous step.

Now, if $G^{t+1} \cap \pi_{t+1} \circ \dots \circ \pi_1(T) = \pi_{t+1} \circ \dots \circ \pi_1(T)$ then we continue with the iteration and do not terminate. If we have reached termination, we therefore have that $G^{t+1} \cap \pi_{t+1} \circ \dots \circ \pi_1(T) = \text{Id}_{G^{t+1}}$. We claim that this implies that $K^{t+1} \cap T = \text{Id}_G$ (and therefore we may take the output group to be $H = K^{t+1}$). For the sake of contradiction, instead suppose $T \leq K^{t+1}$ (since T is 1-dimensional). Applying $\pi_{t+1} \circ \dots \circ \pi_1$ we have that

$$\pi_{t+1} \circ \dots \circ \pi_1(T) \leq \pi_{t+1} \circ \dots \circ \pi_1(K^{t+1}) \leq G^{t+1}$$

which contradicts the termination condition.

Finally, note that if $G^{t+1} \cap \pi_{t+1} \circ \dots \circ \pi_1(T) = \pi_{t+1} \circ \dots \circ \pi_1(T)$ then F_{t+1} when viewed as a function on G^{t+1}/Γ^{t+1} is seen to have a nonzero $\pi_{t+1} \circ \dots \circ \pi_1(T)$ vertical character (which is given by descending ξ on G in through $\pi_{t+1} \circ \dots \circ \pi_1$ in the obvious manner), so one can continue in the iteration in this case.

Step 5: Fixing the value at 0. To see that this completes the proof, if the iteration terminates at some stage t then note that

$$g = \varepsilon_1 \cdots \varepsilon_t \cdot g_t \cdot \gamma_t \cdots \gamma_1.$$

Using that the product of smooth sequences are appropriately smooth and analogously for rational sequences allows us to deduce the necessary outputs. However, we have not guaranteed that the values of the factorization are the id_G at 0. For this, let $g_t(0) = \{g_t(0)\}[g_t(0)]$ with $[g_t(0)] \in K^t \cap \Gamma$ and $d_G(\{g_t(0)\}, \text{id}_G) \leq (M/\varepsilon)^{O_k(d^{O_k(1)})}$. We then have that

$$g = \varepsilon_1 \cdots \varepsilon_t \cdot \{g_t(0)\} \cdot (\{g_t(0)\}^{-1} g_t [g_t(0)]^{-1}) \cdot [g_t(0)] \cdot \gamma_t \cdots \gamma_1.$$

As $g(0) = 0$, we have that $\tau = [g_t(0)] \cdot \gamma_t(0) \cdots \gamma_1(0)$ satisfies $d_G(\tau, \text{id}_G) \leq (M/\varepsilon)^{O_k(d^{O_k(1)})}$ and τ is $(M/\varepsilon)^{O_k(d^{O_k(1)})}$ -rational. Thus

$$g = \varepsilon_1 \cdots \varepsilon_t \cdot \{g_t(0)\} \tau \cdot (\tau^{-1} \{g_t(0)\}^{-1} g_t [g_t(0)]^{-1} \tau) \cdot \tau^{-1} [g_t(0)] \cdot \gamma_t \cdots \gamma_1$$

and note that $(\tau^{-1} \{g_t(0)\}^{-1} g_t [g_t(0)]^{-1} \tau)$ takes value in the conjugated subgroup $\tau^{-1} K^t \tau$ which is $(M/\varepsilon)^{O_k(d^{O_k(1)})}$ -rational by [35, Lemma B.15]. Note however that despite modifying the output group H via conjugation, we have $\tau^{-1} K^t \tau \cap T = \tau^{-1} K^t \tau \cap \tau^{-1} T \tau = \text{Id}_G$ as desired. \square

We now remove the assumption of a 1-dimensional vertical torus via a reduction to this case.

Corollary 5.5. *Let $\ell \geq 1$ be an integer, $\delta \in (0, 1/10)$, $M \geq 1$, and $F: G/\Gamma \rightarrow \mathbb{C}$. Suppose that G is dimension d , is s -step nilpotent with a given degree k filtration, and nilmanifold G/Γ is complexity at most M with respect to this filtration.*

Suppose that $T \leq Z(G)$ is a subgroup of the center which is M -rational. Further suppose that F has a nonzero T -vertical character ξ with $|\xi| \leq M/\delta$, $\|F\|_{\text{Lip}} \leq M$, $N \geq (M/\delta)^{\Omega_{k,\ell}(d^{\Omega_{k,\ell}(1)})}$, and g is a polynomial sequence with respect to the degree k filtration. Then if

$$|\mathbb{E}_{\vec{n} \in [N]^\ell} F(g(\vec{n})\Gamma)| \geq \delta$$

there exists a factorization

$$g = \varepsilon g' \gamma$$

such that:

- g' lives in an $(M/\delta)^{O_{k,\ell}(d^{\Omega_{k,\ell}(1)})}$ -rational subgroup H such that $\xi(H \cap T) = 0$;
- γ is an $(M/\delta)^{O_{k,\ell}(d^{\Omega_{k,\ell}(1)})}$ -rational polynomial sequence;
- ε is an $((M/\delta)^{O_{k,\ell}(d^{\Omega_{k,\ell}(1)})}, N)$ -smooth polynomial sequence.

Furthermore if $g(0) = \text{id}_G$ then we may take $\varepsilon(0) = g'(0) = \gamma(0) = \text{id}_G$.

Proof. We first reduce to the case where $g(0) = \text{id}_G$ as is standard. We factor $g(0) = \{g(0)\}[g(0)]$ such that $[g(0)] \in \Gamma$ and $\psi_G(\{g(0)\}) \in [0, 1]^{\dim(G)}$. Replacing F by $F(\{g(0)\}\cdot)$ and g by $\{g(0)\}^{-1}g\{g(0)\}^{-1}$ we may clearly reduce to the case where $g(0) = \text{id}_G$ at the cost of replacing M by $M^{O_k(d^{O_k(1)})}$ which leaves the conclusion unchanged.

Using Lemma 3.10 to bound the complexity of $G/\ker(\xi)$ and noting that F descends to an $M^{O_k(d^{O_k(1)})}$ -Lipschitz function on $G/\ker(\xi)$, by Theorem 5.4 we have that

$$(g \bmod \ker(\xi)) = \varepsilon g' \gamma$$

where ε, g', γ satisfy:

- $\varepsilon(0) = g'(0) = \gamma(0) = \text{id}_{G/\ker(\xi)}$;
- g' lives in an $(M/\delta)^{O_{k,\ell}(d^{O_{k,\ell}(1)})}$ -rational subgroup H such that $H \cap (T/\ker(\xi)) = \text{id}_{G/\ker(\xi)}$;
- γ is an $(M/\delta)^{O_{k,\ell}(d^{O_{k,\ell}(1)})}$ -rational polynomial sequence;
- ε is an $((M/\delta)^{O_{k,\ell}(d^{O_{k,\ell}(1)})}, N)$ -smooth polynomial sequence.

We now “lift” this factorization. Consider the Mal’cev basis \mathcal{X}' for $G/\ker(\xi)$. For each element $X'_i \in \mathcal{X}'$ we may lift to $Z_i \in \log G$ such that:

- $\exp(X'_i) = \exp(Z_i) \bmod \ker(\xi)$;
- $d_G(\exp(Z_i), \text{id}_G) \leq M^{O_k(d^{O_k(1)})}$;
- Z_i is an $M^{O_k(d^{O_k(1)})}$ -rational combination of the elements of \mathcal{X} .

Writing ε as

$$\varepsilon(\vec{n}) = \exp\left(\sum_{|\vec{i}| \leq k} \varepsilon_{\vec{i}} \begin{pmatrix} \vec{n} \\ \vec{i} \end{pmatrix}\right)$$

where $\varepsilon_{\vec{i}} \in \log(G_{|\vec{i}|}/(\ker(\xi) \cap G_{|\vec{i}|}))$, we lift via the above mapping on \mathcal{X}' to

$$\tilde{\varepsilon}(\vec{n}) = \exp\left(\sum_{|\vec{i}| \leq k} \tilde{\varepsilon}_{\vec{i}} \begin{pmatrix} \vec{n} \\ \vec{i} \end{pmatrix}\right)$$

where $\tilde{\varepsilon}_{\vec{i}} \in \log(G_{|\vec{i}|})$ and analogously for g', γ .

We easily see that $\tilde{\varepsilon}$ is an $((M/\delta)^{O_{k,\ell}(d^{O_{k,\ell}(1)})}, N)$ -smooth polynomial sequence, that $\tilde{\gamma}$ is an $(M/\delta)^{O_{k,\ell}(d^{O_{k,\ell}(1)})}$ -rational polynomial sequence, and that \tilde{g}' takes values in the subgroup $H' = \exp(\log(H) + \log(\ker(\xi)))$. Furthermore H' is seen to be $(M/\delta)^{O_{k,\ell}(d^{O_{k,\ell}(1)})}$ -rational and $\xi(H' \cap T) = \text{Id}_G$. Finally note that $\tilde{\varepsilon} \bmod \ker(\xi) = \varepsilon$ and analogously for $\tilde{g}', \tilde{\gamma}$. Therefore

$$g \cdot (\tilde{\varepsilon} \tilde{g}' \tilde{\gamma})^{-1} \equiv \text{id}_G \bmod \ker(\xi)$$

as polynomial sequences. Thus

$$g = g \cdot (\tilde{\varepsilon} \tilde{g}' \tilde{\gamma})^{-1} \cdot (\tilde{\varepsilon} \tilde{g}' \cdot \tilde{\gamma}) = \tilde{\varepsilon} \cdot ((g \cdot (\tilde{\varepsilon} \tilde{g}' \tilde{\gamma})^{-1}) \cdot \tilde{g}') \cdot \tilde{\gamma}$$

gives the desired factorization noting that $\ker(\xi) \leq H'$ and $\ker(\xi)$ is central and therefore $g \cdot (\tilde{\varepsilon} \tilde{g}' \tilde{\gamma})^{-1}$ may be commuted to the left. \square

6. SETUP FOR SUNFLOWER AND LINEARIZATION ITERATION

We now set up the iteration which will take up the bulk of the following four sections. The idea is to inductively assume the statement of Theorem 1.2 for $s-1$ (i.e., the quantitative inverse theorem for the $U^s[N]$ -norm) and the remaining goal is to prove it for s . The key step is to show that for many $h \in [N]$, $\Delta_h f$ correlates with a multidegree $(1, s-1)$ nilcharacter; this is a quantitative version of [29, Theorem 7.1]. For the remainder of the analysis until Section 12 we

will be concerned with the notion of a correlation structure, which can be thought of as refining the notion in Definition 3.14 with intermediate bracket information.

Definition 6.1. A correlation structure associated to the function $f: [N] \rightarrow \mathbb{C}$ with parameters ρ, M, d , and D and degree-rank $(s-1, r^*)$ is the following data:

- A subset $H \subseteq [N]$ such that $|H| \geq \rho N$;
- A multidegree $(1, s-1)$ nilcharacter $\chi(h, n)$ that lives on a nilmanifold G^*/Γ^* where χ has a $G_{(1, s-1)}^*$ -vertical frequency η^* . Furthermore G^*/Γ^* has dimension bounded by d and complexity bounded by M , the function F^* underlying χ is M -Lipschitz, η^* has height bounded by M , and the output dimension of χ is bounded by D . We let $g(h, n)$ denote the underlying polynomial sequence of χ ;
- A collection of degree-rank $(s-1, r^*)$ nilcharacters $\chi_h(n)$ which live on G/Γ where every χ_h has the same $G_{(s-1, r^*)}$ -vertical frequency η . Furthermore G/Γ has dimension bounded by d and complexity bounded by M (with Mal'cev basis \mathcal{X}), the function underlying χ_h is M -Lipschitz, η has height bounded by M , and χ_h has output dimension bounded by D . Finally, the function underlying χ_h , which we will denote F , is independent of h ;
- The polynomial sequences satisfy $g_h(0) = \text{id}_G$;
- For all $h \in H$ we have

$$\Delta_h f(n) \otimes \overline{\chi(h, n)} \otimes \overline{\chi_h(n)} \in \text{Corr}(s-2, \rho, M, d).$$

If the input function f we are considering for the proof of Theorem 1.2 satisfies

$$\|f\|_{U^{s+1}[N]} \geq \delta,$$

then our proof will always maintain bounds of the form

$$\rho^{-1}, M, D \leq \exp(\log(1/\delta)^{O_s(1)}) \text{ and } d \leq \log(1/\delta)^{O_s(1)}$$

on intermediate correlation structures, although the precise dependence may decay over roughly s stages (wherein we reduce r^* from $s-1$ to 0).

To get started, we first note that given a function f with large U^{s+1} -norm we may associate to it a correlation structure of degree-rank $(s-1, s-1)$; this is little more than chasing definitions and applying induction.

Lemma 6.2. Fix $\delta \in (0, 1/2)$ and $s \geq 2$. Assume Theorem 1.2 for $s-1$. Let $f: [N] \rightarrow \mathbb{C}$ be a 1-bounded function such that

$$\|f\|_{U^{s+1}[N]} \geq \delta.$$

Then there exists a degree-rank $(s-1, s-1)$ correlation structure associated to f with parameters ρ, M, d , and D such that

$$\rho^{-1}, M, D \leq \exp(\log(1/\delta)^{O_s(1)}) \text{ and } d \leq \log(1/\delta)^{O_s(1)}.$$

Proof. Note that $\|f\|_{U^{s+1}[N]} \geq \delta$ implies that

$$\mathbb{E}_{h \in [N]} \|\Delta_h f\|_{U^s[N]}^{2^s} \geq \delta^{O_s(1)};$$

this implicitly uses that $\|\Delta_h f\|_{U^s[N]} = \|\Delta_{-h} f\|_{U^s[N]}$ and that $\Delta_h f$ is identically zero for $|h| > N$.

Therefore there exists $H \subseteq [N]$ with $|H| \geq \delta^{O_s(1)} N$ such that

$$\|\Delta_h f\|_{U^s[N]}^{2^s} \geq \delta^{O_s(1)}$$

for $h \in H$.

By induction on Theorem 1.2, we may assume that for all such $h \in H$ there exists G_h/Γ_h with degree $s-1$ filtration and an associated polynomial sequence $g_h(\cdot)$ such that

$$\mathbb{E}_{h \in [N]} [\Delta_h f(n) \overline{F_h(g_h(n)\Gamma)}] \geq \rho$$

where G_h/Γ_h has complexity bounded by M and dimension bounded by d . We may take

$$M, \rho^{-1} \leq \exp(\log(1/\delta)^{O_s(1)}) \text{ and } d \leq \log(1/\delta)^{O_s(1)}.$$

Note that via writing $g_h(0) = \{g_h(0)\}[g_h(0)]$ where $\psi_{G_h}(\{g_h(0)\}) \in [0, 1]^{\dim(G_h)}$ and $[g_h(0)] \in \Gamma_h$, we have that

$$\begin{aligned} F_h(g_h(n)\Gamma) &= F_h(\{g_h(0)\}\{g_h(0)\}^{-1}g_h(n)[g_h(0)]^{-1} \cdot [g_h(0)]\Gamma) \\ &= F_h(\{g_h(0)\}\{g_h(0)\}^{-1}g_h(n)[g_h(0)]^{-1}\Gamma) \end{aligned}$$

Note that $g'_h(n) = \{g_h(0)\}^{-1}g_h(n)[g_h(0)]^{-1}$ has $g'_h(0) = \text{id}_{G_h}$ and $F'_h = F_h(\{g_h(0)\}\cdot)$ is appropriately Lipschitz (as $\{g_h(0)\}$ has appropriately bounded coordinates by [35, Lemma B.2]). Therefore without loss we may assume that $g_h(0) = \text{id}_G$ for all $h \in H$.

Next note that there are only $O_s(M)^{O_s(d^{O(1)})}$ nilmanifolds with degree $(s-1)$ filtration of complexity bounded by M (up to isomorphism). This follows from Lie's third theorem on the correspondence between Lie algebras and connected, simply connected Lie groups and counting the total possible number of different structure constants and filtration choices for the Lie algebra. Therefore by Pigeonhole we may assume, at the cost of decreasing the size of set H by a multiplicative factor of $O_s(M)^{O_s(-d^{O(1)})}$, that $G_h/\Gamma_h = G/\Gamma$ (and the corresponding filtration) is independent of $h \in H$.

We next remove the dependence on h for the function F_h . Let γ be a parameter to be chosen later; by applying Lemma B.3 we may write

$$F_h(g\Gamma) = \sum_{j \in I} \tau_j(g\Gamma)^2 \cdot F_h(g\Gamma)$$

where $|I| \leq (1/\gamma)^{O_s(d^{O_s(1)})}$, every $g\Gamma$ is supported on at most $2^{O_s(d)}$ many terms, and τ_j are $(M/\gamma)^{O_s(d^{O_s(1)})}$ -Lipschitz. Furthermore each τ_j is supported on a width 2γ cube near the origin (in Mal'cev coordinates); see the third item of Lemma B.3 for a precise description. Since F is an M -Lipschitz function, and choosing γ to be sufficiently small with respect to $(\rho/M)^{O_s(d^{O_s(1)})}$, we find that

$$\sup_{g \in G} |F_h(g\Gamma) - \sum_{j \in I} a_j \tau_j(g\Gamma)^2| \leq \rho/2$$

by taking a_j to be the mean of F_h on the support of τ_j . Note that $|a_j| \leq M$. Pigeonholing over $j \in I$ and decreasing ρ and the size of H by appropriate factors of $O_s(M)^{O_s(-d^{O_s(1)})}$, we may assume that $F_h = F$ for all $h \in H$.

We finally want to replace F by a nilcharacter with a vertical frequency and the claimed output dimension bound. We first give G a degree-rank $(s-1, s-1)$ filtration induced by its degree $s-1$ filtration. This is done via [29, Example 6.11] (i.e., $G_{(d,r)}$ is generated by iterated commutators which either have filtration depths adding to greater than d or adding to exactly d with at least r participating elements). Lemma 2.1 guarantees each subgroup is $M^{O_s(d^{O_s(1)})}$ -rational. Via [35, Lemma B.11], we may give G a Mal'cev basis adapted to this degree-rank $(s-1, s-1)$ filtration with complexity $M^{O_s(d^{O_s(1)})}$.

Via Fourier expansion (see [35, Lemma A.6]) and the triangle inequality we may additionally assume that F has a vertical $G_{(s-1, s-1)}$ -frequency² with height at most $O_s(M/\rho)^{O_s(d^{O_s(1)})} = \exp(\log(1/\delta)^{O_s(1)})$. Given F , there exists a nilcharacter F_η by Lemma B.4 with vertical frequency η , output dimension bounded by $2^{O_s(d)}$, and such that each coordinate is $O_s(M)^{O_s(d^{O_s(1)})}$ -Lipschitz. The function $(F/(2\|F\|_\infty), F_\eta \cdot \sqrt{1 - |F/(2\|F\|_\infty)|^2})$ demonstrates that without loss of generality, we may assume F is a coordinate of a nilcharacter.

²We apply [35, Lemma A.6] to the degree filtration $G_{(0,0)} = G_{(1,0)} \geq G_{(2,0)} \geq \dots \geq G_{(s-1,0)} \geq G_{(s-1,s-1)} \geq \text{Id}_G$.

To complete the deduction, we take G^*/Γ^* to be the trivial nilmanifold and $g(h, n)$ to be a constant sequence. \square

The heart of this paper is the following quantification of [29, Theorem 7.2], the proof of which is the goal of the next few sections culminating in Section 11.2.

Lemma 6.3. *Fix $s \geq 2$ and $1 \leq r^* \leq s - 1$. Suppose $f: [N] \rightarrow \mathbb{C}$ is a 1-bounded function and $N \geq \exp(\Omega_s((d \log(MD/\rho))^{\Omega_s(1)}))$.*

Furthermore suppose that there exists a degree-rank $(s-1, r^)$ correlation structure associated to f with parameters ρ, M, d , and D . Then there exists a degree-rank $(s-1, r^*-1)$ correlation structure associated to f with parameters ρ', M', d' , and D' such that*

$$\rho'^{-1}, M', D' \leq \exp(O_s((d \log(MD/\rho))^{O_s(1)})) \text{ and } d' \leq O_s((d \log(MD/\rho))^{O_s(1)}).$$

Combining Lemma 6.3 along with the observation that degree-rank $(s-1, 0)$ nilmanifolds induce a degree $(s-2)$ filtration (coming from the groups $G_{(i,0)}$), we immediately obtain the following. In particular, these can now be “hidden” inside the nilmanifolds implicit in $\text{Corr}(\cdot, \cdot, \cdot, \cdot)$.

Theorem 6.4. *Fix $\delta \in (0, 1/2)$ and $s \geq 2$. Assume Theorem 1.2 for $s-1$. Let $f: [N] \rightarrow \mathbb{C}$ be a 1-bounded function such that*

$$\|f\|_{U^{s+1}[N]} \geq \delta.$$

Then the following data exists:

- A subset $H \subseteq [N]$ of size at least ρN ;
- A multidegree $(1, s-1)$ nilcharacter $\chi(h, n)$ which lives on a nilmanifold G^*/Γ^* where χ has a $G_{(1, s-1)}^*$ -vertical frequency η^* . Furthermore G^*/Γ^* has dimension bounded by d and complexity bounded by M , the function underlying χ is M -Lipschitz, η^* has height bounded by M , and the output dimension of χ is bounded by D ;
- For all $h \in H$ we have that

$$\Delta_h f(n) \otimes \overline{\chi(h, n)} \in \text{Corr}(s-2, \rho, M, d).$$

Furthermore, we can find such data satisfying

$$\rho^{-1}, M, D \leq \exp(\log(1/\delta)^{O_s(1)}) \text{ and } d \leq \log(1/\delta)^{O_s(1)}.$$

Remark. The case when N is small (i.e., $N \leq \exp(\log(1/\delta)^{O_s(1)})$) is handled via noting that $\|\Delta_h f\|_{L^2[N]} \geq \exp(\log(1/\delta)^{O_s(1)}) \cdot N^{-O(1)}$ for many h and then applying Fourier analysis. Such an analysis always loses factors of N and thus is only useful in this crude edge case. We will not comment further on such issues.

7. ON A CAUCHY-SCHWARZ ARGUMENT OF GOWERS

The proof of Lemma 6.3 is performed in a sequence of stages. We first deduce that the functions correlating with $\Delta_h f$ are not arbitrary. Indeed for many additive quadruples (h_1, h_2, h_3, h_4) we have that the associated tensor product of $\chi_h(n)$ exhibits correlation with a degree $(s-2)$ nilsequence.

We first need the following elementary Fourier-analytic lemma which converts correlation on long progressions to correlation with a major-arc Fourier phase; this is essentially [27, Lemma 3.5(ii)].

Lemma 7.1. *Let $\delta \in (0, 1/2)$. Suppose that $g: [N] \rightarrow \mathbb{C}$ is 1-bounded and there exists an arithmetic progression P of length δN with common difference q within $[N]$ such that*

$$|\mathbb{E}_{n \in P} g(n)| \geq \delta.$$

Then there exists $\Theta \in \mathbb{R}$ such that $\|q\Theta\|_{\mathbb{R}/\mathbb{Z}} \leq \delta^{-O(1)}N^{-1}$ and

$$|\mathbb{E}_{n \in [N]} e(\Theta n) g(n)| \geq \delta^{O(1)}N.$$

Proof. Extend g to be zero beyond the interval $[N]$. Let P' be the arithmetic progression of length $\delta^2 N$ with common difference q centered at 0. We have

$$\left| \sum_{n \in \mathbb{Z}} (\mathbb{1}_P * (|P'|^{-1} \mathbb{1}_{P'}))(n) g(n) \right| \geq \delta^{O(1)}N.$$

Via Fourier inversion, we have

$$\left| \int_{\Theta \in \mathbb{T}} \widehat{g}(\Theta) \overline{\widehat{\mathbb{1}_P(\Theta)} \widehat{\mathbb{1}_{P'}(\Theta)}} d\Theta \right| \geq \delta^{O(1)}N^2.$$

Now via standard bounds on linear exponential sums, we have

$$|\widehat{\mathbb{1}_P}(\Theta)|, |\widehat{\mathbb{1}_{P'}}(\Theta)| \lesssim \min(\|q\Theta\|_{\mathbb{R}/\mathbb{Z}}^{-1}, N).$$

Since $|\widehat{g}(\Theta)| \leq N$, we have that

$$\left| \int_{\|q\Theta\|_{\mathbb{R}/\mathbb{Z}} \geq T/N} \widehat{g}(\Theta) \widehat{\mathbb{1}_P}(\Theta) \overline{\widehat{\mathbb{1}_{P'}(\Theta)}} d\Theta \right| \lesssim N^2/T.$$

Therefore, taking $T = \delta^{-O(1)}$ sufficiently large we have that

$$N^2 \int_{\|q\Theta\|_{\mathbb{R}/\mathbb{Z}} \leq T/N} |\widehat{g}(\Theta)| d\Theta \geq \left| \int_{\|q\Theta\|_{\mathbb{R}/\mathbb{Z}} \leq T/N} \widehat{g}(\Theta) \widehat{\mathbb{1}_P}(\Theta) \overline{\widehat{\mathbb{1}_{P'}(\Theta)}} d\Theta \right| \geq \delta^{O(1)}N^2.$$

Thus

$$\sup_{\|q\Theta\|_{\mathbb{R}/\mathbb{Z}} \leq T/N} |\widehat{g}(\Theta)| \geq \delta^{O(1)}T^{-1}N,$$

which is exactly the desired conclusion (recalling that $T = \delta^{-O(1)}$). \square

The following lemma is due ultimately to Gowers but essentially appears as [27, Proposition 6.1]. We include the proof for the sake of completeness.

Lemma 7.2. *Suppose $\delta \in (0, 1/2)$, $f_1, f_2: [N] \rightarrow \mathbb{C}$ are 1-bounded, and $\chi_h: \mathbb{Z} \rightarrow \mathbb{C}$ are all 1-bounded. Suppose that*

$$\mathbb{E}_{h \in [N]} |\mathbb{E}_{n \in [N]} f_2(n) \Delta_h f_1(n) \overline{\chi_h(n)}| \geq \delta.$$

Then there exists Θ such that $\|\Theta\|_{\mathbb{R}/\mathbb{Z}} \leq \delta^{-O(1)}/N$ and

$$\mathbb{E}_{\substack{h_1+h_2=h_3+h_4 \\ h_i \in [N]}} \left| \mathbb{E}_{n \in [N]} \chi_{h_1}(n) \chi_{h_2}(n+h_1-h_4) \overline{\chi_{h_3}(n) \chi_{h_4}(n+h_1-h_4)} \cdot e(\Theta n) \right| \geq \delta^{O(1)}.$$

Proof. Note that we assume that $\chi_h(n) = 0$ for $h \notin [N]$ and that $\chi_h(n) = 0$ for $n \notin [N]$ via replacing $\chi_h(n)$ with $\chi_h(n) \cdot \mathbb{1}_{n \in [N]}$; we will remove this truncation at the end of the argument. We extend these functions by 0 to $\mathbb{Z}/\tilde{N}\mathbb{Z}$ where \tilde{N} is a prime between $4N$ and $8N$.

By Cauchy–Schwarz, we have

$$\mathbb{E}_{h \in \mathbb{Z}/\tilde{N}\mathbb{Z}} |\mathbb{E}_{n \in \mathbb{Z}/\tilde{N}\mathbb{Z}} f_2(n) \Delta_h f_1(n) \overline{\chi_h(n)}|^2 \gtrsim \delta^2.$$

Expanding, this is equivalent to

$$\mathbb{E}_{h \in \mathbb{Z}/\tilde{N}\mathbb{Z}} \mathbb{E}_{n_1, n_2 \in \mathbb{Z}/\tilde{N}\mathbb{Z}} f_2(n_1) f_1(n_1) \overline{f_2(n_2) f_1(n_2) f_1(n_2+h) \chi_h(n_1) \chi_h(n_2)} \gtrsim \delta^2.$$

We set $n = n_1$, $k = n_2 - n_1$, and $m = n_1 + h$ and find that

$$\mathbb{E}_{m, n \in \mathbb{Z}/\tilde{N}\mathbb{Z}, k \in \mathbb{Z}/\tilde{N}\mathbb{Z}} \Delta_k(f_2 f_1)(n) \Delta_k \overline{f_1(m)} \Delta_k \overline{\chi_{m-n}(n)} \gtrsim \delta^2.$$

This implies that

$$\mathbb{E}_{k \in \mathbb{Z}/\tilde{N}\mathbb{Z}} |\mathbb{E}_{m, n \in \mathbb{Z}/\tilde{N}\mathbb{Z}} \Delta_k(f_2 f_1)(n) \Delta_k \overline{f_1(m)} \Delta_k \overline{\chi_{m-n}(n)}|^4 \gtrsim \delta^8.$$

Recall the box-norm inequality that for a, b, Φ which are 1-bounded, we have

$$\begin{aligned} |\mathbb{E}_{n, m \in \mathbb{Z}/\tilde{N}\mathbb{Z}} a(n) b(m) \Phi(n, m)|^4 &\leq (\mathbb{E}_{n \in \mathbb{Z}/\tilde{N}\mathbb{Z}} |\mathbb{E}_{m \in \mathbb{Z}/\tilde{N}\mathbb{Z}} b(m) \Phi(n, m)|)^4 \\ &\leq (\mathbb{E}_{n \in \mathbb{Z}/\tilde{N}\mathbb{Z}} |\mathbb{E}_{m \in \mathbb{Z}/\tilde{N}\mathbb{Z}} b(m) \Phi(n, m)|^2)^2 \\ &= (\mathbb{E}_{n \in \mathbb{Z}/\tilde{N}\mathbb{Z}} \mathbb{E}_{m, m' \in \mathbb{Z}/\tilde{N}\mathbb{Z}} b(m) \overline{b(m')} \Phi(n, m) \overline{\Phi(n, m')})^2 \\ &= (\mathbb{E}_{m, m' \in \mathbb{Z}/\tilde{N}\mathbb{Z}} |\mathbb{E}_{n \in \mathbb{Z}/\tilde{N}\mathbb{Z}} \Phi(n, m) \overline{\Phi(n, m')}|)^2 \\ &\leq \mathbb{E}_{m, m' \in \mathbb{Z}/\tilde{N}\mathbb{Z}} |\mathbb{E}_{n \in \mathbb{Z}/\tilde{N}\mathbb{Z}} \Phi(n, m) \overline{\Phi(n, m')}|^2 \\ (7.1) \quad &= \mathbb{E}_{n, n', m, m' \in \mathbb{Z}/\tilde{N}\mathbb{Z}} \Phi(n, m) \overline{\Phi(n, m')} \overline{\Phi(n', m)} \Phi(n', m'). \end{aligned}$$

Applying this for each fixed k , we have that

$$\mathbb{E}_{k \in \mathbb{Z}/\tilde{N}\mathbb{Z}} \mathbb{E}_{n, n', m, m' \in \mathbb{Z}/\tilde{N}\mathbb{Z}} \Delta_k \overline{\chi_{m-n}(n)} \Delta_k \overline{\chi_{m'-n'}(n')} \Delta_k \chi_{m-n}(n) \Delta_k \chi_{m'-n'}(n') \gtrsim \delta^8.$$

Take $m' - n = h_1$, $m - n' = h_2$, $m - n = h_3$, $m' - n' = h_4$. Note that $n' - n = h_1 - h_4$ and $h_1 + h_2 = h_3 + h_4$ and noting that n and $n + k$ range over the whole cyclic group, this is exactly

$$\mathbb{E}_{h_1+h_2=h_3+h_4} \left| \mathbb{E}_{n \in \mathbb{Z}/\tilde{N}\mathbb{Z}} \chi_{h_1}(n) \chi_{h_2}(n+h_1-h_4) \overline{\chi_{h_3}(n) \chi_{h_4}(n+h_1-h_4)} \right|^2 \gtrsim \delta^8.$$

Since $\chi_h(n) = 0$ identically for $h \notin [N]$, we in fact have

$$\mathbb{E}_{h_1+h_2=h_3+h_4} \left| \mathbb{E}_{n \in \mathbb{Z}/\tilde{N}\mathbb{Z}} \chi_{h_1}(n) \chi_{h_2}(n+h_1-h_4) \overline{\chi_{h_3}(n) \chi_{h_4}(n+h_1-h_4)} \right|^2 \gtrsim \delta^8.$$

For the inner sum, recall that we “truncated” $\chi_h(n)$ with $\mathbb{1}_{n \in [N]}$. In particular, extracting the truncation term we have that

$$\mathbb{E}_{h_1+h_2=h_3+h_4} \left| \mathbb{E}_{n \in [N]} \mathbb{1}_{1 \leq n+h_1-h_4 \leq N} \chi_{h_1}(n) \chi_{h_2}(n+h_1-h_4) \overline{\chi_{h_3}(n) \chi_{h_4}(n+h_1-h_4)} \right|^2 \gtrsim \delta^8.$$

Via an application of Lemma 7.1, there exist choices of $\Theta_{\vec{h}}$ with $\|\Theta_{\vec{h}}\|_{\mathbb{R}/\mathbb{Z}} \leq \delta^{-O(1)}/N$ such that

$$\mathbb{E}_{h_1+h_2=h_3+h_4} \left| \mathbb{E}_{n \in [N]} \chi_{h_1}(n) \chi_{h_2}(n+h_1-h_4) \overline{\chi_{h_3}(n) \chi_{h_4}(n+h_1-h_4)} e(\Theta_{\vec{h}} n) \right|^2 \gtrsim \delta^{O(1)}.$$

Rounding $\Theta_{\vec{h}}$ to a lattice of spacing $\delta^{O(1)}/N$ and Pigeonholing then gives the desired result. \square

The next proof will require defining the notion when two nilcharacters are “equivalent” (i.e., have the same symbol in a quantified sense of [29, Appendix E]).

Definition 7.3. We say nilcharacters χ, χ' are (M, D, d) -equivalent for multidegree J if χ, χ' have output dimensions bounded by D and all coordinates of

$$\chi \otimes \overline{\chi'}$$

can be represented as sums of at most M nilsequences of multidegree J such that the underlying functions of each nilsequence are M -Lipschitz and the underlying nilmanifolds have complexity bounded by M and dimension bounded by d .

The key reason for the definition of equivalence is the following proposition, which states that given equivalent nilcharacters χ and χ' , correlations with them are equivalent modulo introducing a term of multidegree J . This is a finitary quantification of [29, Lemma E.7].

Lemma 7.4. Given a function $f: \Omega \rightarrow \mathbb{C}^L$ and nilcharacters χ, χ' which are (M, D, d) -equivalent for multidegree J , if

$$\|\mathbb{E}_{\vec{n} \in \Omega} f(\vec{n}) \otimes \chi(\vec{n})\|_\infty \geq \rho$$

then

$$\|\mathbb{E}_{\vec{n} \in \Omega} f(\vec{n}) \otimes \chi'(\vec{n}) \cdot \psi(\vec{n})\|_\infty \geq (\rho/(MD))^{O(1)},$$

where ψ can be taken to be one of nilsequences used as part of a representation of one of the coordinates in $\chi \otimes \overline{\chi'}$. In particular, ψ is a nilsequence of multidegree J such that underlying nilmanifold has complexity bounded by M and dimension bounded by d and the underlying function has Lipschitz constant bounded by M .

Remark. The additional condition that ψ can be taken to be an explicit nilsequence occurring in a witness for the equivalence of χ, χ' is used primarily to allow us to Pigeonhole the choice of ψ in cases where we may need to apply this statement “on average”.

Proof. Notice that since χ' is a nilcharacter, we have that the trace of

$$\chi' \otimes \overline{\chi'}$$

is the constant function 1. Furthermore note that the trace is the sum of at most D coordinates of $\chi' \otimes \overline{\chi'}$ and therefore

$$\|\mathbb{E}_{\vec{n} \in \Omega} f(\vec{n}) \otimes \chi(\vec{n}) \otimes \overline{\chi'(\vec{n})} \otimes \chi'(\vec{n})\|_\infty \geq \rho/D.$$

Consider the coordinate of $f(\vec{n}) \otimes \chi(\vec{n}) \otimes \overline{\chi'(\vec{n})} \otimes \chi'(\vec{n})$ which achieves the L^∞ above, and in particular the associated coordinate of $\chi(\vec{n}) \otimes \overline{\chi'(\vec{n})}$ that contributes. Applying the definition of equivalence and the triangle inequality, there exists $\psi(\vec{n})$ of the desired form such that

$$\|\mathbb{E}_{\vec{n} \in \Omega} f(\vec{n}) \otimes \chi'(\vec{n}) \cdot \psi(\vec{n})\|_\infty \geq (\rho/(MD))^{O(1)}. \quad \square$$

We are now in position to prove the quantification of [29, Proposition 7.3]. We remark that there was an error in the published version of [29, Proposition 8.3] which affected the proof of [29, Proposition 7.3]. We quantify a closely related approach to that given in the erratum [26]. For our proof we require various quantifications of [29, Appendix E]; all of these are completely mechanical.

Lemma 7.5. Fix $s \geq 3$ and $1 \leq r^* \leq s - 1$. Let $f: [N] \rightarrow \mathbb{C}$ is a 1-bounded function. Suppose that f has a correlation structure with parameters ρ, M, d , and D and associated nilcharacters $\chi(h, n)$ and $\chi_h(n)$. Then for at least $(MD/\rho)^{-O_s(d^{O_s(1)})} N^3$ quadruples $h_1, h_2, h_3, h_4 \in H$ with $h_1 + h_2 = h_3 + h_4$ we have

$$\chi_{h_1}(n) \otimes \chi_{h_2}(n + h_1 - h_4) \otimes \overline{\chi_{h_3}(n)} \otimes \overline{\chi_{h_4}(n + h_1 - h_4)} \in \text{Corr}(s - 2, \rho', M', d')$$

with

$$\rho'^{-1}, M' \leq (MD/\rho)^{O_s(d^{O_s(1)})} \text{ and } d' \leq O_s(d^{O_s(1)}).$$

Remark 7.6. For $s = 2$, the same statement holds modulo a correction term of $e(\Theta n)$ where Θ is such that $\|\Theta\|_{\mathbb{R}/\mathbb{Z}} \leq (MD/\rho)^{O_s(d^{O_s(1)})}/N$.

Proof. By definition of correlation structures we have for $h \in H$ that

$$\|\mathbb{E}_{n \in [N]} (\Delta_h f)(n) \otimes \overline{\chi(h, n)} \otimes \overline{\chi_h(n)} \cdot \overline{\psi_h(n)}\|_\infty \geq \rho$$

where ψ_h is a nilsequence of degree $(s - 2)$ whose underlying function is at most M -Lipschitz on a nilmanifold of complexity at most M and dimension at most d . Setting $\chi_h(n)$ to be zero for $h \notin H$ we have

$$\mathbb{E}_{h \in [N]} \|\mathbb{E}_{n \in [N]} f(n) \overline{f(n+h)} \otimes \overline{\chi(h, n)} \otimes \overline{\chi_h(n)} \cdot \overline{\psi_h(n)}\|_\infty \geq \rho^2.$$

Twisting ψ_h by an appropriate h -dependent constant complex phase so as to make the L^∞ values be realized as positive real numbers, we may assume that

$$\|\mathbb{E}_{h \in [N]} \mathbb{E}_{n \in [N]} f(n) \overline{f(n+h)} \otimes \overline{\chi(h, n)} \otimes \overline{\chi_h(n)} \cdot \overline{\psi_h(n)}\|_\infty \geq \rho^2 / D^2.$$

By Lemma C.5, we have that $\chi(h, n)$ is $((MD)^{O_s(d^{O_s(1)})}, (MD)^{O_s(d^{O_s(1)})}, d^{O_s(1)})$ -equivalent for degree $(s-1)$ to some $\tilde{\chi}(h, n, \dots, n)$ which is a multidegree $(1, \dots, 1)$ nilcharacter with output dimension, complexity of underlying nilmanifold, Lipschitz constant of underlying function for each coordinate, and vertical frequency height all bounded by $(MD)^{O_s(d^{O_s(1)})}$. ($\tilde{\chi}$ has s total arguments.) Thus, applying Lemma 7.4, we have that

$$\|\mathbb{E}_{n, h \in [N]} f(n) \overline{f(n+h)} \otimes \overline{\tilde{\chi}(h, n, \dots, n)} \otimes \overline{\chi_h(n)} \cdot \overline{\psi_h(n)} \cdot \tilde{\psi}(h, n)\|_\infty \geq (MD/\rho)^{-O_s(d^{O_s(1)})},$$

where $\tilde{\psi}(h, n)$ is a degree $(s-1)$ nilsequence where the underlying function has Lipschitz norm and complexity of underlying nilmanifold bounded $(MD)^{O_s(d^{O_s(1)})}$ while the dimension of the underlying nilmanifold is bounded by $O_s(d^{O_s(1)})$. The nilsequence $\tilde{\psi}(h, n)$ can also be viewed as a multidegree $(0, s-1) \cup (s-1, s-2)$ nilsequence. (I.e., we take the union of the down-sets generated by these elements.) Furthermore, the underlying function has Lipschitz norm and complexity of underlying nilmanifold bounded $(MD)^{O_s(d^{O_s(1)})}$ while the dimension of the underlying nilmanifold is bounded by $O_s(d^{O_s(1)})$.

Thus, applying Lemma C.6 (splitting) we have

$$\|\mathbb{E}_{n, h \in [N]} f(n) \overline{f(n+h)} \otimes \overline{\tilde{\chi}(h, n, \dots, n)} \otimes \overline{\chi_h(n)} \cdot \overline{\tilde{\psi}_h(n)} \cdot b(n)\|_\infty \geq (MD/\rho)^{-O_s(d^{O_s(1)})}$$

where $\tilde{\psi}_h$ are degree $(s-2)$ nilsequences in n where complexity and Lipschitz constant are bounded by $(MD)^{O_s(d^{O_s(1)})}$ and the dimension of the underlying nilmanifold is bounded by $O_s(d^{O_s(1)})$ while $b(n)$ is $(MD)^{O_s(d^{O_s(1)})}$ -bounded. Therefore, applying Lemma 7.2, we have

$$\begin{aligned} & \mathbb{E}_{\substack{h_1+h_2=h_3+h_4 \\ h_i \in [N]}} \|\mathbb{E}_{n \in [N]} \tilde{\chi}(h_1, n, \dots, n) \otimes \tilde{\chi}(h_2, n+h_1-h_4, \dots, n+h_1-h_4) \otimes \overline{\tilde{\chi}(h_3, n, \dots, n)} \\ & \otimes \overline{\tilde{\chi}(h_4, n+h_1-h_4, \dots, n+h_1-h_4)} \otimes \chi_{h_1}(n) \otimes \chi_{h_2}(n+h_1-h_4) \otimes \overline{\chi_{h_3}(n)} \otimes \overline{\chi_{h_4}(n+h_1-h_4)} \\ & \cdot \overline{\tilde{\psi}_{h_1}(n)} \overline{\tilde{\psi}_{h_2}(n)} \overline{\tilde{\psi}_{h_3}(n)} \overline{\tilde{\psi}_{h_4}(n)} e(\Theta n)\|_\infty \geq (MD/\rho)^{-O_s(d^{O_s(1)})}. \end{aligned}$$

We may combine $\overline{\tilde{\psi}_{h_1}(n)} \overline{\tilde{\psi}_{h_2}(n)} \overline{\tilde{\psi}_{h_3}(n)} \overline{\tilde{\psi}_{h_4}(n)} e(\Theta n)$ to form $\psi_{h_1, h_2, h_3, h_4}^*(n)$ which is degree $(s-2)$ in n and with identical complexity bounds to $\tilde{\psi}_{h_1}$ modulo changing implicit constant. Additionally, we may twist $\psi_{h_1, h_2, h_3, h_4}^*$ by an (h_1, h_2, h_3, h_4) -dependent complex phase to bring the outer expectation inside the norm. Thus we have

$$\begin{aligned} & \|\mathbb{E}_{\substack{h_1+h_2=h_3+h_4 \\ h_i \in [N]}} \mathbb{E}_{n \in [N]} \tilde{\chi}(h_1, n, \dots, n) \otimes \tilde{\chi}(h_2, n+h_1-h_4, \dots, n+h_1-h_4) \otimes \overline{\tilde{\chi}(h_3, n, \dots, n)} \\ & \otimes \overline{\tilde{\chi}(h_4, n+h_1-h_4, \dots, n+h_1-h_4)} \otimes \chi_{h_1}(n) \otimes \chi_{h_2}(n+h_1-h_4) \otimes \overline{\chi_{h_3}(n)} \\ & \otimes \overline{\chi_{h_4}(n+h_1-h_4)} \cdot \psi_{h_1, h_2, h_3, h_4}^*(n)\|_\infty \geq (MD/\rho)^{-O_s(d^{O_s(1)})}. \end{aligned}$$

By Lemma C.5, $\chi(h_2, n+h_1-h_4, \dots, n+h_1-h_4)$ is $((MD)^{O_s(d^{O_s(1)})}, (MD)^{O_s(d^{O_s(1)})}, d^{O_s(1)})$ -equivalent for degree $(s-1)$ to

$$\bigotimes_{k=0}^{s-1} \chi(h_2, n, \dots, n, h_1-h_4, \dots, h_1-h_4)$$

where there are $s-k-1$ copies of n and k copies of h_1-h_4 and we have a similar expansion for $\chi(h_4, n+h_1-h_4, \dots, n+h_1-h_4)$. Note that all terms in this expansion except for $k=0$

may be absorbed into ψ^* . Therefore applying Lemma 7.4, we have that

$$\begin{aligned} & \left\| \mathbb{E}_{\substack{h_1+h_2=h_3+h_4 \\ h_i \in [N]}} \mathbb{E}_{n \in [N]} \widetilde{\chi}(h_1, n, \dots, n) \otimes \widetilde{\chi}(h_2, n, \dots, n) \otimes \overline{\widetilde{\chi}(h_3, n, \dots, n)} \right. \\ & \quad \otimes \overline{\widetilde{\chi}(h_4, n, \dots, n)} \otimes \chi_{h_1}(n) \otimes \chi_{h_2}(n + h_1 - h_4) \otimes \overline{\chi_{h_3}(n)} \\ & \quad \left. \otimes \overline{\chi_{h_4}(n + h_1 - h_4)} \cdot \psi_{h_1, h_2, h_3, h_4}^*(n) \cdot \tau(n, h_1, h_2, h_3, h_4) \right\|_{\infty} \geq (MD/\rho)^{-O_s(d^{O_s(1)})}; \end{aligned}$$

here $\tau(n, h_1, h_2, h_3, h_4)$ is a degree $(s-1)$ nilsequence where the underlying function has Lipschitz norm and complexity of underlying nilmanifold bounded by $(MD)^{O_s(d^{O_s(1)})}$ while the dimension of the underlying nilmanifold is bounded by $O_s(d^{O_s(1)})$ and we have folded certain terms into ψ^* while guaranteeing it is a degree $(s-2)$ nilsequence (and the complexity bounds have not changed modulo implicit constants). Finally via Lemma C.5, we have that

$$\widetilde{\chi}(h_1, n, \dots, n) \otimes \widetilde{\chi}(h_2, n, \dots, n) \otimes \overline{\widetilde{\chi}(h_3, n, \dots, n)} \otimes \overline{\widetilde{\chi}(h_4, n, \dots, n)}$$

and $\widetilde{\chi}(h_1 + h_2 - h_3 - h_4, n, \dots, n)$ are $((MD)^{O_s(d^{O_s(1)})}, (MD)^{O_s(d^{O_s(1)})}, d^{O_s(1)})$ -equivalent for degree $(s-1)$. Thus applying Lemma 7.4, we have

$$\begin{aligned} & \mathbb{E}_{\substack{h_1+h_2=h_3+h_4 \\ h_i \in [N]}} \left\| \mathbb{E}_{n \in [N]} \chi_{h_1}(n) \otimes \chi_{h_2}(n + h_1 - h_4) \otimes \overline{\chi_{h_3}(n)} \otimes \overline{\chi_{h_4}(n + h_1 - h_4)} \right. \\ & \quad \left. \cdot \widetilde{\chi}(0, n, \dots, n) \psi_{h_1, h_2, h_3, h_4}^*(n) \tau(n, h_1, h_2, h_3, h_4) \right\|_{\infty} \geq (MD/\rho)^{-O_s(d^{O_s(1)})}; \end{aligned}$$

here we have folded in various terms into $\tau(n, h_1, \dots, h_4)$ and the complexity bounds have not changed modulo implicit constants. Note that by Lemma C.2, $\widetilde{\chi}(0, n, \dots, n)$ is a degree $(s-1)$ nilsequence in n and thus may abusively also be absorbed into τ . Finally noting that a degree $(s-1)$ nilsequence may also be viewed as a multidegree $(s-1, 0, \dots, 0) \cup (s-2, s-1, \dots, s-1)$ nilsequence and thus applying Lemma C.6 we have

$$\begin{aligned} & \left\| \mathbb{E}_{\substack{h_1+h_2=h_3+h_4 \\ h_i \in [N]}} \mathbb{E}_{n \in [N]} \chi_{h_1}(n) \otimes \chi_{h_2}(n + h_1 - h_4) \otimes \overline{\chi_{h_3}(n)} \otimes \overline{\chi_{h_4}(n + h_1 - h_4)} \cdot \psi_{h_1, h_2, h_3, h_4}^*(n) b(n) \right\|_{\infty} \\ & \quad \geq (MD/\rho)^{-O_s(d^{O_s(1)})}, \end{aligned}$$

where $b(n)$ is an $(MD)^{O_s(d^{O_s(1)})}$ -bounded function and ψ^* has been modified but the underlying complexity bounds have changed by adjusting implicit constants. Note ψ^* is degree $(s-2)$.

We now reparameterize with

$$h_1 = m - n, h_2 = m' - n', h_3 = m' - n, h_4 = m - n'.$$

By approximating with regions where we take m, m', n, n' to live in short intervals, there exist intervals I_1, \dots, I_4 each of density $(MD/\rho)^{-O_s(d^{O_s(1)})}$ in $[\pm 2N]$ such that

$$\begin{aligned} & \left\| \mathbb{E}_{m \in I_1, m' \in I_2, n \in I_3, n' \in I_4} \chi_{m-n}(n) \otimes \chi_{m'-n'}(n') \otimes \overline{\chi_{m'-n}(n)} \otimes \overline{\chi_{m-n'}(n')} \right. \\ & \quad \left. \cdot \psi_{m-n, m'-n', m'-n, m-n'}^*(n) b(n) \right\|_{\infty} \gtrsim (MD/\rho)^{-O_s(d^{O_s(1)})} \end{aligned}$$

where $\psi_{m-n, m'-n', m'-n, m-n'}$ is a degree $(s-2)$ nilsequence. Now by Cauchy-Schwarz, duplicating the variable m and denoting the copies by m, m'' , we obtain

$$\begin{aligned} & \left\| \mathbb{E}_{m, m'' \in I_1, m' \in I_2, n \in I_3, n' \in I_4} \chi_{m-n}(n) \otimes \overline{\chi_{m''-n}(n)} \otimes \overline{\chi_{m'-n'}(n')} \otimes \chi_{m''-n'}(n') \right. \\ & \quad \left. \cdot \psi_{m-n, m'-n', m'-n, m-n', m''-n, m''-n'}^*(n) \right\|_{\infty} \gtrsim (MD/\rho)^{-O_s(d^{O_s(1)})}. \end{aligned}$$

Note that every term not involving m was removed using appropriate boundedness. Now we may Pigeonhole on $m' - n = t$ and deduce

$$\left\| \mathbb{E}_{m,m'' \in I_1, n \in I_3, n' \in I_4} \chi_{m-n}(n) \otimes \overline{\chi_{m''-n}(n)} \otimes \overline{\chi_{m-n'}(n')} \otimes \chi_{m''-n'}(n') \right. \\ \left. \cdot \psi_{m-n, m-n', m''-n, m''-n'}^*(n) \cdot \mathbb{1}[n+t \in I_2] \right\|_{\infty} \gtrsim (MD/\rho)^{-O_s(d^{O_s(1)})}.$$

Let $m'' - n = h_1$, $m - n' = h_2$, $m - n = h_3$, and $m'' - n' = h_4$ (abusively). We have

$$\left\| \mathbb{E}_{n \in [N]} \mathbb{E}_{\substack{h_1+h_2=h_3+h_4 \\ h_i \in [\pm N]}} \chi_{h_1}(n) \otimes \chi_{h_2}(n+h_1-h_4) \otimes \overline{\chi_{h_3}(n)} \otimes \overline{\chi_{h_4}(n+h_1-h_4)} \right. \\ \left. \cdot \chi_{h_1, h_2, h_3, h_4}(n) \cdot \mathbb{1}[n+h_3, n+h_1 \in I_1, n \in I_3, n+h_1-h_4 \in I_4, n+t \in I_2] \right\|_{\infty} \gtrsim (MD/\rho)^{-O_s(d^{O_s(1)})}$$

where $\chi_{h_1, h_2, h_3, h_4}(n)$ is a degree $(s-2)$ nilsequence (for each fixed h_1, h_2, h_3, h_4) where the underlying nilmanifold and Lipschitz constant of underlying function are bounded by $(MD/\rho)^{O_s(d^{O_s(1)})}$ and the dimension is bounded by $O_s(d^{O_s(1)})$.

Therefore, by the triangle inequality we have that

$$\mathbb{E}_{\substack{h_1+h_2=h_3+h_4 \\ h_i \in [\pm N]}} \left\| \mathbb{E}_{n \in [N]} \chi_{h_1}(n) \otimes \chi_{h_2}(n+h_1-h_4) \otimes \overline{\chi_{h_3}(n)} \otimes \overline{\chi_{h_4}(n+h_1-h_4)} \right. \\ \left. \chi_{h_1, h_2, h_3, h_4}(n) \cdot \mathbb{1}[n+h_3, n+h_1 \in I_1, n \in I_3, n+h_1-h_4 \in I_4, n+t \in I_2] \right\|_{\infty} \gtrsim (MD/\rho)^{-O_s(d^{O_s(1)})}.$$

Finally, the last term is the indicator of an \vec{h} -dependent interval. Applying Lemma 7.1 (and noting that $s \geq 3$ allows us to fold in the major arc Fourier term) completes the proof. \square

8. SUNFLOWER STEP

For the next stage of our proof, as outlined in Section 4, we wish to provide more structure on h -dependent nilcharacters χ_h given information about additive quadruples as established in Section 7. As setup we will require the notion of a rational subspace with respect to a specified basis, and establish some basic control over Taylor coefficients of bounded polynomial sequences.

Definition 8.1. *A vector subspace $V' \leq V$ is Q -rational with respect to V given the basis $\mathcal{B} = \{B_1, \dots, B_{\dim(V)}\}$ (of V) if there exists a basis $\mathcal{B}' = \{B'_1, \dots, B'_{\dim(V')}\}$ of V' such that each B'_j is a linear combination of elements of \mathcal{B} with coefficients of height at most Q .*

Lemma 8.2. *Consider a nilmanifold G/Γ given a degree-rank filtration of degree rank (s, r) , dimension d , and complexity at most M . Let \mathcal{X} denote the underlying adapted Mal'cev basis and assign the basis*

$$\mathcal{X}_i = (\mathcal{X} \cap \log(G_{(i,1)})) / \log(G_{(i,2)})$$

for $G_{(i,1)}/G_{(i,2)}$. Suppose ε is a polynomial sequence such that

$$d_{G, \mathcal{X}}(\text{id}_G, \varepsilon(n)) \leq M$$

for $n \in [N]$. Then for $1 \leq i \leq s$, we have

$$d_{G_{(i,1)}/G_{(i,2)}, \mathcal{X}_i}(\text{Taylor}_i(\varepsilon), \text{id}_{G_{(i,1)}/G_{(i,2)}}) \leq M^{O_s(d^{O_s(1)})} N^{-i}.$$

Proof. We may write

$$\varepsilon(n) = \exp \left(\sum_{j=0}^s \varepsilon_j \binom{n}{j} \right)$$

where $\varepsilon_j \in \log(G_{(j,0)})$. By Lemma 2.12, we have

$$\text{Taylor}_i(\varepsilon) = \exp(\varepsilon_i) \bmod G_{(i,2)}.$$

We have that

$$\|\psi_{\text{exp}}(\varepsilon(n))\|_\infty \leq M^{O_s(d^{O_s(1)})}$$

for all $n \in [N]$ by [35, Lemmas B.1, B.3]. This implies that

$$\left\| \sum_{t=0}^j (-1)^t \binom{j}{t} \psi_{\text{exp}}(\varepsilon(t \cdot \lfloor N/(2j) \rfloor + 1)) \right\|_\infty \leq M^{O_s(d^{O_s(1)})}.$$

This is exactly the j -th discrete derivative and thus terms coming from ε_i with $i < j$ vanish. This implies that

$$\|\varepsilon_j N^j \bmod \log(G_{(j,2)})\|_\infty \leq M^{O_s(d^{O_s(1)})},$$

where the basis we assign to $\log(G_{(j,1)}/G_{(j,2)})$ is \mathcal{X}_j . The result follows by dividing by N^{-j} and noting, by say [38, Lemma 2.6], that the distance in first- and second-kind coordinates is comparable. \square

We now come to the first of two crucial arguments in this paper where we “improve” the correlation structure. At the cost of restricting the set H , we force the Taylor coefficients of g_h , the polynomial sequences underlying the χ_h , to live in certain restricted subspaces and their differences to lie in an even finer restriction.

This step is closely related to the “sunflower” arguments of [27, Step 1] and [29, Lemma 11.3]; a quantitative version for the U^4 -inverse theorem due to the first author can be found in [36]. The precise statement of the lemma should also be compared with [29, Theorem 11.1(i)]. We note however that unlike [27, 29], our proof is completely free of any iteration (or equivalently passing to a subgroup where polynomial sequences are “totally equidistributed”, which necessitates too much loss in the relevant parameters).

Thus, the crucial point of the following technical statement is the final condition, which essentially captures that two h -dependent frequencies in the improved correlation structure cannot “simultaneously” affect the bottom degree-rank portion.

Lemma 8.3. *Fix $s \geq 2$ and $1 \leq r^* \leq s - 1$. Let $f: [N] \rightarrow \mathbb{C}$ be a 1-bounded function. Suppose that f has a degree-rank $(s - 1, r^*)$ correlation structure with parameters ρ , M , d , and D and that $N \geq (MD/\rho)^{O_s(d^{O_s(1)})}$ and data labeled as in Definition 6.1. Furthermore let $\mathcal{X}_i = (\mathcal{X} \cap \log(G_{(i,1)}))/\log(G_{(i,2)})$.*

We output a new degree-rank $(s - 1, r^)$ correlation structure for f with parameters*

$$\rho'^{-1} \leq (MD/\rho)^{O_s(d^{O_s(1)})}, \quad M' \leq O(M), \quad D' = D, \quad d' \leq O(d),$$

with set $H' \subseteq H$, with multidegree $(1, s - 1)$ nilcharacter $\chi'(h, n) = F^{'}(g'(h, n)\Gamma^{*'})$ on $(G^*)' = G^* \times \mathbb{R}$, with h -dependent nilcharacters χ'_h having underlying polynomial sequences $g'_h(n) = F'(g'_h(n)\Gamma)$ on $G' = G$. This correlation structure satisfies:*

- $(G^*)'$ is given the multidegree filtration

$$(G^*)'_{(i,j)} = (G^*)_{(i,j)} \times \{0\}$$

if $(i, j) \neq (0, 0)$ or $(0, 1)$. For $(i, j) \in \{(0, 0), (0, 1)\}$ we set

$$(G^*)'_{(i,j)} = (G^*)_{(i,j)} \times \mathbb{R}.$$

We have $F^{'}((x, z)(\Gamma^* \times \mathbb{Z})) = F^*(x\Gamma^*) \cdot e(z)$. We have $g'(h, n) = (g(h, n), \Theta n)$ for some appropriate value of Θ ;*

- There exists a collection of \mathbb{R} -vector spaces $V_{i,\text{Dep}} \leq V_i \leq G_{(i,1)}/G_{(i,2)}$ which are all $(MD/\rho)^{O_s(d^{O_s(1)})}$ -rational with respect to $\exp(\mathcal{X}_i)$ for each i ;
- For $1 \leq i \leq s-1$ and $h, h_1, h_2 \in H^i$ we have

$$\text{Taylor}_i(g'_h) \in V_i, \quad \text{Taylor}_i(g'_{h_1}) - \text{Taylor}_i(g'_{h_2}) \in V_{i,\text{Dep}};$$

- F^i is M^i -Lipschitz and has the same vertical frequency η as F ;
- For integers $i_1 + \dots + i_{r^*} = s-1$, suppose that $v_{i_\ell} \in V_{i_\ell}$ and for at least two distinct indices ℓ_1, ℓ_2 we have $v_{i_{\ell_1}} \in V_{i_{\ell_1},\text{Dep}}$ and $v_{i_{\ell_2}} \in V_{i_{\ell_2},\text{Dep}}$. Then for w which is any (r^*-1) -fold commutator of $v_{i_1}, \dots, v_{i_{r^*}}$, we have

$$\eta(w) = 0.$$

Remark. Consider $g_{i_j} = \exp(X_{i_j})$ with $g_{i_j} \in G_{i_j,0}$ for $1 \leq j \leq r^*$ and $i_1 + \dots + i_{r^*} = s-1$. Fixing any (r^*-1) -fold commutator w of $g_{i_1}, \dots, g_{i_{r^*}}$, repeated application of the commutator version of Baker–Campbell–Hausdorff (e.g. (2.2)) implies that

$$w = \exp([X_{i_1}, \dots, X_{i_j}])$$

where the associated commutator has the same “form” as that defining w . (All higher terms are annihilated since G has degree-rank $(s-1, r^*)$.) Note that this implies that one can define the associated commutator given inputs in $G_{(i_1,1)}/G_{(i_1,2)}, \dots, G_{(i_{r^*},1)}/G_{(i_{r^*},2)}$ and furthermore we see that the associated commutator form on the Lie algebra is a multilinear form of the vector arguments (since $G_{(i,1)}/G_{(i,2)}$ and $G_{(s-1,r^*)}$ are real vector spaces and the commutator bracket on the Lie algebra is multilinear).

Proof. We first note that the statement of the lemma is trivial for $r^* = 1$ since we may take $V_i = V_{i,\text{Dep}} = G'_{(i,1)}/G'_{(i,2)}$; it is impossible to have two distinct indices in the final bullet point. Taking $g'_h(n) = g_h(n)$ and $g'(h, n) = (g(h, n), 0)$ completes the proof in this case. For $s = 2$, the only possible case is $r^* = 1$ and therefore for the remainder of the proof we will consider $s \geq 3$. Similarly, if η is trivial, the result is once again immediate. Thus throughout the remainder of the proof we will assume that $s-1 \geq r^* \geq 2$ and η is nontrivial.

Step 1: Setup for invoking equidistribution theory. By Lemma 7.5, we have

$$\left\| \mathbb{E} \left[\chi_{h_1}(n) \otimes \chi_{h_2}(n+h_1-h_4) \otimes \overline{\chi_{h_3}(n)} \otimes \overline{\chi_{h_4}(n+h_1-h_4)} \cdot \psi_{\vec{h}}(g_{\vec{h}}(n)\Gamma') \right] \right\|_{\infty} \geq (MD/\rho)^{-O_s(d^{O_s(1)})}$$

for at least $(MD/\rho)^{-O_s(d^{O_s(1)})}$ fraction of additive quadruples $h_1 + h_2 = h_3 + h_4$. Furthermore $g_{\vec{h}}(n)$ is a polynomial sequence on a group G_{Error} which has a degree $(s-2)$ filtration and the complexity of $G_{\text{Error}}/\Gamma_{\text{Error}}$ and the Lipschitz constant of the function for $\psi_{\vec{h}}$ are bounded by $(MD/\rho)^{O_s(d^{O_s(1)})}$. Note that a priori $G_{\text{Error}}/\Gamma_{\text{Error}}$ and the associated Mal'cev basis depend on \vec{h} . However, applying Pigeonhole on the choice of the associated structure constants allows us to assume, at the cost of passing to a density $(MD/\rho)^{-O_s(d^{O_s(1)})}$ subset of the additive quadruples, that $G_{\text{Error}}/\Gamma_{\text{Error}}$ is independent of \vec{h} . Finally, we may assume as usual that $g_{\vec{h}}(0) = \text{id}_{G_{\text{Error}}}$ via by-now standard manipulations.

We now consider the group $\tilde{G} = G \times G \times G \times G \times G_{\text{Error}}$. \tilde{G} may naturally be given a degree-rank $(s-1, r^*)$ product filtration (where we use [29, Example 6.11] to assign G_{Error} a degree-rank $(s-2, s-2)$ structure) and Mal'cev basis. Furthermore if $\chi_{h_i}(n) = F(g_{h_i}(n))$ we have that the five-fold function $F(x_1\Gamma) \otimes F(x_2\Gamma) \otimes \overline{F(x_3\Gamma)} \otimes \overline{F(x_4\Gamma)} \cdot \psi_{\vec{h}}(x_5\Gamma_{\text{Error}})$ has a vertical frequency $\eta_{\text{Prod}} = (\eta, \eta, -\eta, -\eta, 0)$. (Note that $(G_{\text{Error}})_{(s-1,i)} = \text{Id}_{G_{\text{Error}}}$ for all $i \geq 0$.)

For the sake of convenience, we set

$$g_{\vec{h}}^*(n) = (g_{h_1}(n), g_{h_2}(n+h_1-h_4), g_{h_3}(n), g_{h_4}(n+h_1-h_4), g_{\vec{h}}(n))$$

and note that the function $F \otimes F \otimes \overline{F} \otimes \overline{F} \cdot \psi_{\vec{h}}$ is seen to be $M^{O_s(d^{O_s(1)})}$ -Lipschitz on \tilde{G} . Note that by the second item of Lemma 10.2, we immediately have that

$$\text{Taylor}_i(g_{h_2}(n + h_1 - h_4)) = \text{Taylor}_i(g_{h_2}(n))$$

for $1 \leq i \leq s - 1$ and analogously for $g_{h_4}(n + h_1 - h_4)$.

Step 2: Invoking equidistribution theory. By applying Corollary 5.5 (since η is nonzero), there exists a $(MD/\rho)^{O_s(d^{O_s(1)})}$ -rational subgroup $J = J_{\vec{h}}$ of \tilde{G} such that $\eta_{\text{Prod}}(J \cap \tilde{G}_{(s-1, r^*)}) = 0$ and such that

$$g_{\vec{h}}^* = \varepsilon_{\vec{h}} \cdot \tilde{g}_{\vec{h}} \cdot \gamma_{\vec{h}}$$

where:

- $\varepsilon_{\vec{h}}(0) = \tilde{g}_{\vec{h}}(0) = \gamma_{\vec{h}}(0) = \text{id}_{\tilde{G}}$;
- $\tilde{g}_{\vec{h}}$ takes values in J ;
- $\gamma_{\vec{h}}$ is $(MD/\rho)^{O_s(d^{O_s(1)})}$ -rational (with respect to the lattice $\Gamma \times \Gamma \times \Gamma \times \Gamma \times \Gamma_{\text{Error}}$);
- $d(\varepsilon(n), \varepsilon(n-1)) \leq (MD/\rho)^{O_s(d^{O_s(1)})} N^{-1}$ for $n \in [N]$.

By passing to a subset of additive quadruples of density $(MD/\rho)^{-O_s(d^{O_s(1)})}$ we may in fact assume that the group J is independent of \vec{h} under consideration.

We define

$$J'_i := (J \cap \tilde{G}_{(i,1)}) / (J \cap \tilde{G}_{(i,2)}), \quad J_i := \tau_i(J'_i)$$

where $\tau_i: \text{Horiz}_i(G)^{\otimes 4} \times \text{Horiz}_i(G_{\text{Error}}) \rightarrow \text{Horiz}_i(G)^{\otimes 4}$ is the natural projection map to the four-fold product. Since $\eta_{\text{Prod}}(J \cap \tilde{G}_{(s-1, r)}) = 0$ (due to the output of Corollary 5.5), we have

$$\eta_{\text{Prod}}([J'_{i_1}, \dots, J'_{i_{r^*}}]) = 0$$

where the commutator bracket is taken with respect to \tilde{G} and $[\cdot, \dots, \cdot]$ denotes any possible $(r^* - 1)$ -fold commutator bracket.

Since G_{Error} has been given a degree-rank $< (s - 1, r^*)$ filtration, we have that in fact

$$\eta_{\text{Prod}}([J_{i_1}, \dots, J_{i_{r^*}}]) = 0$$

where we abusively descend η_{Prod} to $G^{\otimes 4}$. Less formally, we are noting that the final coordinate of elements in \tilde{G} play no role in commutators of the depth being considered.

Step 3: Furstenberg–Weiss commutator argument. We now perform the crucial Furstenberg–Weiss commutator argument. Given $T \subseteq [4]$, we define $\pi_T((v_1, \dots, v_4)) = (v_i)_{i \in T}$ with the coordinates represented in increasing order of index.

We define

$$\begin{aligned} \pi_{123}(J_i)^* &= \pi_{123}(J_i) \cap \{(v, 0, 0) : v \in \text{Horiz}_i(G)\}, \\ \pi_{124}(J_i)^* &= \pi_{124}(J_i) \cap \{(v, 0, 0) : v \in \text{Horiz}_i(G)\}. \end{aligned}$$

Note that $\pi_{123}(J_i)^*$ and $\pi_{124}(J_i)^*$ may (abusively) be viewed as subspaces of $\text{Horiz}_i(G)$. The crucial claim is that

$$\eta([v_{i_1}, \dots, v_{i_{r^*}}]) = 0$$

if $i_1 + \dots + i_{r^*} = s - 1$, each $v_{i_\ell} \in \pi_1(J_{i_\ell})$, and for two distinct indices ℓ_1, ℓ_2 we have that $v_{i_{\ell_1}} \in \pi_{123}(J_{i_{\ell_1}})^*$ and $v_{i_{\ell_2}} \in \pi_{124}(J_{i_{\ell_2}})^*$. Note that η lives on G and the commutator brackets are taken with respect to G , not $G^{\otimes 4}$. The Furstenberg–Weiss commutator argument is required to capture precisely this difference.

Note that an element $v_i \in \pi_1(J_{i_\ell})$ lifts to an element \widetilde{v}_{i_ℓ} of the form $(v_{i_\ell}, \cdot, \cdot, \cdot) \in \text{Horiz}_{i_\ell}(G)^{\otimes 4}$. Furthermore note that $v_{i_\ell} \in \pi_{123}(J_{i_\ell})$ “lifts” to an element \widetilde{v}_{i_ℓ} of the form $(v_{i_\ell}, 0, 0, \cdot) \in \text{Horiz}_{i_\ell}(G)^{\otimes 4}$ while $v_{i_\ell} \in \pi_{124}(J_{i_\ell})$ lifts to an element \widetilde{v}_{i_ℓ} of the form $(v_{i_\ell}, 0, \cdot, 0) \in \text{Horiz}_{i_\ell}(G)^{\otimes 4}$.

Given the above setup, we have

$$[\widetilde{v}_{i_1}, \dots, \widetilde{v}_{i_{r^*}}] = ([v_{i_1}, \dots, v_{i_{r^*}}], \text{id}_G, \text{id}_G, \text{id}_G).$$

To see this note that the iterated commutator of elements in $G \times \text{Id}_G \times \text{Id}_G \times G$ (with any elements in $G^{\otimes 4}$) remains in the subgroup $G \times \text{Id}_G \times \text{Id}_G \times G$; an analogous fact holds true for $G \times \text{Id}_G \times G \times \text{Id}_G$. Since we assumed that our commutator contains elements in both $G \times \text{Id}_G \times \text{Id}_G \times G$ and $G \times \text{Id}_G \times G \times \text{Id}_G$, the commutator must in fact live in $G \times \text{Id}_G \times \text{Id}_G \times \text{Id}_G$, and the first coordinates of the desired commutators is trivially seen to match.

Recalling that we have

$$\eta_{\text{Prod}}([J_{i_1}, \dots, J_{i_{r^*}}]) = 0,$$

and noting that η_{Prod} descends to η on the subgroup $G_{(s-1, r^*)} \times \text{Id}_G^{\otimes 3}$, we have

$$\eta([v_{i_1}, \dots, v_{i_{r^*}}]) = 0$$

as claimed.

Step 4: Finding (h_2, h_3) and (h'_2, h'_4) which extend to many “good” h_1 . Recall that we are looking at the at least $(MD/\rho)^{-O_s(d^{O_s(1)})}$ fraction of additive quadruples $(h_1, h_2, h_3, h_4) \in H^4 \subseteq [N]^4$ which are such that \widetilde{g}_h lives on a specified subgroup J . Call this set of quadruples \mathcal{S} .

So by Markov, there are at least $(MD/\rho)^{-O_s(d^{O_s(1)})}N$ many $h_1 \in [N]$ which extend to at least $(MD/\rho)^{-O_s(d^{O_s(1)})}N^2$ quadruples in \mathcal{S} . Thus there are at least $(MD/\rho)^{-O_s(d^{O_s(1)})}N^5$ pairs of additive tuples of the form

$$(h_1, h_2, h_3, h_1 + h_2 - h_3), (h_1, h'_2, h_1 + h'_2 - h'_4, h'_4) \in \mathcal{S}.$$

By averaging, there exists a pair of pairs (h_2, h_3) and (h'_2, h'_4) such that there are at least $(MD/\rho)^{-O_s(d^{O_s(1)})}N$ many $h_1 \in [N]$ which live in such additive tuples. We fix such a pair of pairs and define \mathcal{T} to denote the set of $h_1 \in [N]$ such that $(h_1, h_2, h_3, h_1 + h_2 - h_3) \in \mathcal{S}$ and $(h_1, h'_2, h_1 + h'_2 - h'_4, h'_4) \in \mathcal{S}$.

Step 5: Extracting coefficient data. Consider $h_1 \in \mathcal{T}$ and define

$$h^{123} = (h_1, h_2, h_3, h_1 + h_2 - h_3), \quad h^{124} = (h_1, h'_2, h_1 + h'_2 - h'_4, h'_4).$$

Recall $\mathcal{X}_i = (\mathcal{X} \cap \log(G_{(i,1)}))/\log(G_{(i,2)})$ and assign the basis $\exp(\mathcal{X}_i)$ to $G_{(i,1)}/G_{(i,2)}$ (viewed as a vector space). Finally we assign the basis $\mathcal{Z}_i = \bigcup_{Y_i \in \exp(\mathcal{X}_i)} \{(Y_i, 0, 0), (0, Y_i, 0), (0, 0, Y_i)\}$ to $(G_{(i,1)}/G_{(i,2)})^{\otimes 3}$.

By Lemma 2.13, we have

$$\text{Taylor}_i(g_{h^{123}}^*) = \text{Taylor}_i(\varepsilon_{h^{123}}) + \text{Taylor}_i(\widetilde{g_{h^{123}}}) + \text{Taylor}_i(\gamma_{h^{123}})$$

$$\text{Taylor}_i(g_{h^{124}}^*) = \text{Taylor}_i(\varepsilon_{h^{124}}) + \text{Taylor}_i(\widetilde{g_{h^{124}}}) + \text{Taylor}_i(\gamma_{h^{124}})$$

Therefore, by Lemma 8.2, for all $h_1 \in \mathcal{T}$ we have

$$\text{dist}(\text{Taylor}_i((g_{h_1}, g_{h_2}, g_{h_3})), \pi_{123}(J_i) + T_{h_1}^{-1} \text{Horiz}_i(\Gamma^{\otimes 3})) \leq (MD/\rho)^{O_s(d^{O_s(1)})} N^{-i},$$

$$\text{dist}(\text{Taylor}_i((g_{h_1}, g_{h'_2}, g_{h'_4})), \pi_{124}(J_i) + T_{h_1}'^{-1} \text{Horiz}_i(\Gamma^{\otimes 3})) \leq (MD/\rho)^{O_s(d^{O_s(1)})} N^{-i},$$

where T_{h_1} and T_{h_1}' are positive integers bounded by $(MD/\rho)^{O_s(d^{O_s(1)})}$. Here we have identified the basis \mathcal{Z}_i (for $(G_{(i,1)}/G_{(i,2)})^{\otimes 3}$) with the standard basis vectors in $\mathbb{R}^{3 \dim(\text{Horiz}_i(G))}$ and taken the L^∞ metric on the latter (for the notion of dist). At the cost of shrinking the set \mathcal{T} by a multiplicative factor of $(MD/\rho)^{-O_s(d^{O_s(1)})}$ we may assume that $T_{h_1} = T$ and $T_{h_1}' = T'$ for all $h_1 \in \mathcal{T}$.

We now consider a basis \mathcal{B}_i for $\pi_{123}(J_i)$ which is in row-echelon form where one orders the coordinates corresponding to second copy of G (in the four-fold $G^{\otimes 4}$) at the front, then the third

copy, and then the first copy. In particular, the “final block” of basis vectors span $\pi_{123}(J_i)^*$. Note that one can take such \mathcal{B}_i such that the coordinates are integers bounded by $(MD/\rho)^{O_s(d^{O_s(1)})}$ due to the rationality of $\pi_{123}(J_i)$.

For $h_1, h'_1 \in \mathcal{T}$, we have

$$(8.1) \quad \text{Taylor}_i((g_{h_1}, g_{h_2}, g_{h_3})) = \sum_{R_j \in \mathcal{B}_i} a_j R_j + T^{-1} \mathbb{Z}^{3 \dim(\text{Horiz}_i(G))} + v_{h_1}$$

$$(8.2) \quad \text{Taylor}_i((g_{h'_1}, g_{h_2}, g_{h_3})) = \sum_{R_j \in \mathcal{B}_i} a'_j R_j + T^{-1} \mathbb{Z}^{3 \dim(\text{Horiz}_i(G))} + v_{h'_1}$$

where $\|v_{h_1}\|_\infty, \|v_{h'_1}\|_\infty \leq (MD/\rho)^{O_s(d^{O_s(1)})} N^{-i}$. For each basis vector $R_j \in \mathcal{B}_i$ where the first nonzero element is either in coordinates corresponding to second or third copy of G , there exists a dual vector which is zero on the coordinates corresponding to the first copy of G and whose inner product with all of \mathcal{B}_i but R_j is zero.

Call this vector v_j and note one may take v_j to have integral coordinates bounded by $(MD/\rho)^{O_s(d^{O_s(1)})}$ and divisible by T . Then from (8.1) and (8.2),

$$0 = v_j \cdot (\text{Taylor}_i((g_{h_1}, g_{h_2}, g_{h_3})) - \text{Taylor}_i((g_{h'_1}, g_{h_2}, g_{h_3}))) = M_j(a_j - a'_j) + \mathbb{Z} \pm (MD/\rho)^{O_s(d^{O_s(1)})} N^{-i},$$

where M_j is a nonzero integer bounded by $(MD/\rho)^{O_s(d^{O_s(1)})}$. (That is, $M_j(a_j - a'_j)$ is within $(MD/\rho)^{O_s(d^{O_s(1)})} N^{-i}$ of an integer.)

We may now use this information about such indices j in conjunction with (8.1) and (8.2). We deduce that for all $h_1, h'_1 \in \mathcal{T}$,

$$\text{dist}(\text{Taylor}_i(g_{h_1}) - \text{Taylor}_i(g_{h'_1}), \pi_{123}(J_i)^* + T_1^{-1} \mathbb{Z}^{\dim(\text{Horiz}_i(G))}) \leq (MD/\rho)^{O_s(d^{O_s(1)})} N^{-i},$$

where T_1 is an integer of size bounded by $(MD/\rho)^{O_s(d^{O_s(1)})}$. Analogously,

$$\text{dist}(\text{Taylor}_i(g_{h_1}) - \text{Taylor}_i(g_{h'_1}), \pi_{124}(J_i)^* + T_1'^{-1} \mathbb{Z}^{\dim(\text{Horiz}_i(G))}) \leq (MD/\rho)^{O_s(d^{O_s(1)})} N^{-i}$$

where T_1' is an integer bounded by $(MD/\rho)^{O_s(d^{O_s(1)})}$. Putting it together, we may deduce that

$$\text{dist}(\text{Taylor}_i(g_{h_1}) - \text{Taylor}_i(g_{h'_1}), \pi_{123}(J_i)^* \cap \pi_{124}(J_i)^* + T_2^{-1} \mathbb{Z}^{\dim(\text{Horiz}_i(G))}) \leq (MD/\rho)^{O_s(d^{O_s(1)})} N^{-i}$$

with T_2 a nonzero integer bounded by $(MD/\rho)^{O_s(d^{O_s(1)})}$. To see this, simply construct a bounded integral basis of the orthogonal complement of $\pi_{123}(J_i)^* \cap \pi_{124}(J_i)^*$ (treated as a subspace of the dual space to $G_{(i,1)}/G_{(i,2)} \simeq \mathbb{R}^{\dim(\text{Horiz}_i(G))}$). Then the two input inequalities imply that any basis vector for the intersection space dual will map $T_1 T_1' (\text{Taylor}_i(g_{h_1}) - \text{Taylor}_i(g_{h'_1}))$ to a near-integral scalar, which gives the claim.

Now by Lemma 2.13 we therefore have

$$(8.3) \quad \text{dist}(\text{Taylor}_i(g_{h_1} g_{h'_1}^{-1}), \pi_{123}(J_i)^* \cap \pi_{124}(J_i)^* + T_2^{-1} \mathbb{Z}^{\dim(\text{Horiz}_i(G))}) \leq (MD/\rho)^{O_s(d^{O_s(1)})} N^{-i}.$$

It is also trivial by restricting the factorization to the first coordinate that

$$(8.4) \quad \text{dist}(\text{Taylor}_i(g_{h_1}), \pi_1(J_i) + T_3^{-1} \mathbb{Z}^{\dim(\text{Horiz}_i(G))}) \leq (MD/\rho)^{O_s(d^{O_s(1)})} N^{-i}$$

with T_3 a nonzero integer bounded by $(MD/\rho)^{O_s(d^{O_s(1)})}$.

Step 6: Extracting initial factorizations. For the remainder of the proof fix $h_1^* \in \mathcal{T}$. Given $h'_1 \in \mathcal{T}$ we have

$$g_{h'_1} = g_{h_1^*}^{-1} \cdot g_{h_1^*}.$$

Note that $\pi_1(J_i)$ and $\pi_{123}(J_i)^* \cap \pi_{124}(J_i)^*$ may each be defined as the kernel of a set of i -th horizontal characters (on G) of height at most $(MD/\rho)^{O_s(d^{O_s(1)})}$. Recall (8.3) and (8.4). Scaling the horizontal characters by at most $(MD/\rho)^{O_s(d^{O_s(1)})}$ and applying Lemma B.2, we may write

$$\begin{aligned} g_{h'_1} g_{h'_1}^{-1} &= \varepsilon_{h'_1} \cdot \widetilde{g}_{h'_1} \cdot \gamma_{h'_1}, \\ g_{h'_1} &= \varepsilon \cdot \widetilde{g}' \cdot \gamma', \end{aligned}$$

where:

- $\varepsilon_{h_1}(0) = \widetilde{g}_{h'_1}(0) = \gamma_{h'_1}(0) = \varepsilon(0) = g'(0) = \gamma'(0) = \text{id}_G$;
- $\text{Taylor}_i(\widetilde{g}_{h'_1}) \in \pi_{123}(J_i)^* \cap \pi_{124}(J_i)^*$ and $\text{Taylor}_i(\widetilde{g}') \in \pi_1(J_i)$;
- $\gamma_{h'_1}, \gamma'$ are $(MD/\rho)^{O_s(d^{O_s(1)})}$ -rational;
- $d_G(\varepsilon_{h_1}(n), \varepsilon_{h_1}(n-1)) + d_G(\varepsilon'(n), \varepsilon'(n-1)) \leq (MD/\rho)^{O_s(d^{O_s(1)})} N^{-1}$ for $n \in [N]$.

Therefore

$$g_{h'_1} = g_{h'_1} g_{h'_1}^{-1} \cdot g_{h'_1} = \varepsilon_{h'_1} \varepsilon \cdot (\varepsilon^{-1} \widetilde{g}_{h'_1} \varepsilon) \cdot (\varepsilon^{-1} \gamma_{h'_1} \varepsilon \gamma_{h'_1}^{-1}) \cdot (\gamma_{h'_1} \cdot \widetilde{g}' \gamma_{h'_1}^{-1}) \cdot \gamma_{h'_1} \gamma'.$$

By Lemma 2.13, we have

$$\begin{aligned} \text{Taylor}_i(\gamma_{h'_1} \widetilde{g}' \gamma_{h'_1}^{-1}) &= \text{Taylor}_i(\widetilde{g}') \in \pi_1(J_i), \\ \text{Taylor}_i((\varepsilon^{-1} \widetilde{g}_{h'_1} \varepsilon) \cdot (\varepsilon^{-1} \gamma_{h'_1} \varepsilon \gamma_{h'_1}^{-1})) &= \text{Taylor}_i(\widetilde{g}_{h'_1}) \in \pi_{123}(J_i)^* \cap \pi_{124}(J_i)^*. \end{aligned}$$

We say that $h'_1, h''_1 \in \mathcal{T}$ have matching rational parts if

$$(\gamma_{h'_1} \gamma')^{-1} \cdot (\gamma_{h''_1} \gamma')$$

is a polynomial sequence valued in Γ . By restricting \mathcal{T} to an appropriate subset of density $(MD/\rho)^{-O_s(d^{O_s(1)})}$, we may assume that all $h'_1 \in \mathcal{T}$ have matching rational parts. (This is most easily seen in first-kind coordinates: if $\gamma_{h'_1} \gamma'$ and $\gamma_{h'_2} \gamma'$ have all coefficients differing by $T_4 \cdot \text{span}(\mathcal{X}, \mathbb{Z})$ where T_4 is an appropriate integer of size bounded by $(MD/\rho)^{O_s(d^{O_s(1)})}$ then two sequences match up to a polynomial sequence in Γ .)

So, ultimately we may assume that for all $h'_1 \in \mathcal{T}$ we have

$$g_{h'_1} = \varepsilon_{h'_1}^* \cdot g_{h'_1}^* \cdot \gamma^* \cdot \widetilde{\gamma}_{h'_1}$$

where:

- $\varepsilon_{h'_1}^*(0) = g_{h'_1}^*(0) = \gamma^*(0) = \widetilde{\gamma}_{h'_1}(0) = \text{id}_G$;
- $\text{Taylor}_i(g_{h'_1}^* \cdot (g_{h'_1}^*)^{-1}) \in \pi_{123}(J_i)^* \cap \pi_{124}(J_i)^*$ and $\text{Taylor}_i(g_{h'_1}^*) \in \pi_1(J_i)$ for all $h''_1 \in \mathcal{T}$;
- γ^* is $(MD/\rho)^{O_s(d^{O_s(1)})}$ -rational
- $\widetilde{\gamma}_{h'_1}$ takes values in Γ ;
- $d_G(\varepsilon_{h'_1}^*(n), \varepsilon_{h'_1}^*(n-1)) \leq (MD/\rho)^{O_s(d^{O_s(1)})} N^{-1}$ for $n \in [N]$.

Step 7: Removing periodic and smooth pieces of factorization. Let Q be the period of $\gamma^* \Gamma$ and define $\delta = (MD/\rho)^{-O_s(d^{O_s(1)})}$ where δ is to be chosen later. We break $[N]$ into a collection of arithmetic progressions with difference Q and length between δN and $2\delta N$; there are at most δ^{-1} such progressions. Call these progressions P_1, \dots, P_ℓ and note that

$$\left\| \mathbb{E}_{n \in [N]} (\Delta_h f)(n) \sum_{i=1}^{\ell} \mathbf{1}_{n \in P_i} \cdot \chi(h, n) \otimes \chi_h(n) \cdot \psi_h(n) \right\|_{\infty} \geq \rho$$

where $\psi_h(n)$ is the degree $(s-2)$ nilsequence coming from the condition $\Delta_h f(n) \otimes \overline{\chi(h, n)} \otimes \overline{\chi_h(n)} \in \text{Corr}(s-1, \rho, M, d)$ from the original correlation structure. For $h \in \mathcal{T}$ we may write

$$\chi_h(n) = F(g_h(n)\Gamma) = F(\varepsilon_h^* g_h^* \gamma^* \Gamma);$$

here we are using that $\widetilde{\gamma}_{h'_1}$ takes values in Γ so may be dropped for the remainder of the analysis.

Since Q is the period of γ^* , we may replace γ^* by a value γ_{P_i} for each progression where $\gamma_{P_i}\gamma^*(n)^{-1} \in \Gamma$ for $n \in P_i$ and $\|\psi(\gamma_{P_i})\|_\infty \leq 1$. Then

$$\left\| \mathbb{E}_{n \in [N]}(\Delta_h f)(n) \cdot \chi(h, n) \otimes \left(\sum_{i=1}^{\ell} \mathbb{1}_{n \in P_i} F(\varepsilon_h^* g_h^* \gamma_{P_i} \Gamma) \right) \cdot \psi_h(n) \right\|_\infty \geq \rho.$$

Furthermore as ε_h^* is sufficiently smooth we may replace ε_h^* with the constant $\varepsilon_{h, P_i} = \varepsilon_h^*(\min(P_i))$ and have

$$\left\| \mathbb{E}_{n \in [N]}(\Delta_h f)(n) \cdot \chi(h, n) \otimes \left(\sum_{i=1}^{\ell} \mathbb{1}_{n \in P_i} F(\varepsilon_{h, P_i} g_h^* \gamma_{P_i} \Gamma) \right) \cdot \psi_h(n) \right\|_\infty \geq \rho/2,$$

as long as δ was chosen sufficiently small.

By the triangle inequality there exists some P_i which is distance at least $\delta^{1/2}N$ from the ends of the interval $[N]$ such that

$$\|\mathbb{E}_{n \in [N]}(\Delta_h f)(n) \cdot \chi(h, n) \otimes \mathbb{1}_{n \in P_i} F(\varepsilon_{h, P_i} g_h^* \gamma_{P_i} \Gamma) \cdot \psi_h(n)\|_\infty \geq \delta^2.$$

By paying a $\delta^{O(1)}$ -fraction in the size of \mathcal{T} we may assume that the choice of index i is independent of h , hence writing $P_i = P$. Furthermore note that there is a $\delta^{O(1)}$ -net of size $\delta^{-O_s(d)}$ for the set of g satisfying $d_G(g, \text{id}_G) \leq (MD/\rho)^{O_s(d^{O_s(1)})}$. If the net size is chosen small enough, we may shift ε_{h, P_i} to a nearby value in the net without much loss. Then we can pay a $\delta^{O_s(d)}$ -fraction in the size of \mathcal{T} to Pigeonhole onto a single point in the net, writing $\varepsilon_{h, P_i} = \varepsilon_P$.

Overall, for all $h \in \mathcal{T}$ we have

$$\|\mathbb{E}_{n \in [N]}(\Delta_h f)(n) \cdot \chi(h, n) \otimes \mathbb{1}_P(n) F(\varepsilon_P g_h^* \gamma_P \Gamma) \cdot \psi_h(n)\|_\infty \geq \delta^3$$

for some P at least $\delta^{1/2}N$ from the endpoints of the interval. Thus by Lemma 7.1, for each $h \in \mathcal{T}$ there exists Θ_h with $\|Q \cdot \Theta_h\|_{\mathbb{R}/\mathbb{Z}} \leq \delta^{-O(1)}N^{-1}$ and

$$(8.5) \quad \|\mathbb{E}_{n \in [N]}(\Delta_h f)(n) \cdot \chi(h, n) \otimes e(\Theta_h n) F(\varepsilon^* g_h^* \gamma_P \Gamma) \cdot \psi_h(n)\|_\infty \geq \delta^{O(1)}.$$

Rounding Θ_h to a net of distance $\delta^{O(1)}N^{-1}$ and paying a $\delta^{O(1)}$ -fraction in the size of \mathcal{T} to Pigeonhole the resulting point, we may write $\Theta_h = \Theta$ for all $h \in \mathcal{T}$. We are now finally in position to define the output data. Define

$$\begin{aligned} g'_h &= \gamma_P^{-1} g_h^* \gamma_P, \\ F' &= F(\varepsilon^* \gamma_P \cdot), \\ \chi'_h(n) &= F'(g'_h(n)), \\ \chi'(h, n) &= \chi(h, n) \cdot e(\Theta n). \end{aligned}$$

Note that $g'(h, n) = (g(h, n), \Theta n)$ is the polynomial sequence underlying χ' , and $\chi'(h, n) = F^{*'}(g'(h, n)\Gamma^{*'})$. It is easy to check the relevant properties of Definition 6.1 to see that we obtain a degree-rank $(s-1, r^*)$ correlation structure with appropriately modified underlying parameters (we set H' to be the final refined version of \mathcal{T}); in particular, (8.5) demonstrates the necessary correlation fact.

Finally, taking

$$V_{i, \text{Dep}} = \pi_{123}(J_i)^* \cap \pi_{124}(J_i)^* \text{ and } V_i = \pi_1(J_i)$$

we finish the proof: in particular, the result from Step 3 demonstrates the final item of the conclusion, and the result from Step 6 demonstrates the third item. \square

9. LINEARIZATION STEP

We now come to the second crucial argument of this paper. Prior this stage we have modified the degree-rank $(s-1, r^*)$ to one in which various Taylor coefficients of g_h for $h \in H$ (upon factoring) differ only on certain special subspaces. In this next stage, we deduce that either these Taylor coefficients differ on a further refined subspace which is seen to be essentially “annihilated” by η or g_h has a certain “bracket linear” form. This step is ultimately where we invoke the results of Sanders [45] on quasi-polynomial bounds for the Bogolyubov lemma.

This step is closely modeled after [27, Step 2] and the closely related proof of [29, Lemma 11.5]; a quantitative version for the U^4 -inverse theorem due to the first author can be found in [36]. The precise statement of the lemma should also be compared with [29, Theorem 11.1(ii)].

Lemma 9.1. *Fix $s \geq 2$ and $1 \leq r^* \leq s-1$. Let $f: [N] \rightarrow \mathbb{C}$ be a 1-bounded function. Suppose that f has a degree rank $(s-1, r^*)$ correlation structure with parameters ρ, M, d , and D and that $N \geq (MD/\rho)^{O_s(d^{O_s(1)})}$. Furthermore let $\mathcal{X}_i = (\mathcal{X} \cap \log(G_{(i,1)}))/\log(G_{(i,2)})$.*

We output a new degree-rank $(s-1, r^)$ correlation structure for f with parameters*

$$\rho'^{-1} \leq \exp(O_s((d \log(MD/\rho))^{O_s(1)})), \quad M' \leq O(M), \quad D' = D, \quad d' \leq O(d),$$

with set $H' \subseteq H$, with multidegree $(1, s-1)$ nilcharacter $\chi'(h, n) = F^{'}(g'(h, n)\Gamma^{*'})$ on $(G^*)' = G^* \times \mathbb{R}$, with h -dependent nilcharacters χ'_h having underlying polynomial sequences $g'_h(n) = F'(g'_h(n)\Gamma)$ on $G' = G$. This correlation structure satisfies:*

- $(G^*)'$ is given the multidegree filtration

$$(G^*)'_{(i,j)} = (G^*)_{(i,j)} \times \{0\}$$

if $(i, j) \neq (0, 0)$ or $(0, 1)$. For $(i, j) \in \{(0, 0), (0, 1)\}$, we set

$$(G^*)'_{(i,j)} = (G^*)_{(i,j)} \times \mathbb{R}.$$

We have $F^{'}((x, z)(\Gamma^* \times \mathbb{Z})) = F^*(x\Gamma^*) \cdot e(z)$. We have $g'(h, n) = (g(h, n), \Theta n)$ for some appropriate value of Θ ;*

- *There is a collection of \mathbb{R} -vector spaces $W_{i,*}, W_{i,\text{Lin}}, W_{i,\text{Pet}} \leq G_{(i,1)}/G_{(i,2)}$ for each i ;*
- *If $W_i := W_{i,*} + W_{i,\text{Lin}} + W_{i,\text{Pet}}$ then $\dim(W_i) = \dim(W_{i,*}) + \dim(W_{i,\text{Lin}}) + \dim(W_{i,\text{Pet}})$, i.e., the three spaces are linearly disjoint;*
- *There exist bases $\mathcal{X}_{i,*}, \mathcal{X}_{i,\text{Lin}}$, and $\mathcal{X}_{i,\text{Pet}}$ of the corresponding spaces which are composed of $(MD/\rho)^{O_s(d^{O_s(1)})}$ -rational combinations of elements of $(\mathcal{X} \cap G_{(i,1)})/G_{(i,2)}$;*
- *For $1 \leq i \leq s-1$ and $h, h_1, h_2 \in H'$ we have*

$$\text{Taylor}_i(g'_h) \in W_{i,*} + W_{i,\text{Lin}} + W_{i,\text{Pet}} = W_i,$$

$$\text{Taylor}_i(g'_{h_1}) - \text{Taylor}_i(g'_{h_2}) \in W_{i,\text{Lin}} + W_{i,\text{Pet}},$$

$$\text{Proj}_{W_{i,\text{Lin}}}(\text{Taylor}_i(g'_h)) = \sum_{Z_{i,j} \in \mathcal{X}_{i,\text{Lin}}} \left(\gamma_{i,j} + \sum_{k=1}^{d^*} \alpha_{i,j,k} \{\beta_k h\} \right) Z_{i,j},$$

with $d^ \leq (d \log(MD/\rho))^{O_s(1)}$ and $\beta_k \in (1/N')\mathbb{Z}$ where N' is a prime in $[100N, 200N]$;*

- *F' is M' -Lipschitz and has the same vertical frequency η as F ;*
- *For any integers $i_1 + \dots + i_{r^*} = s-1$, suppose that $v_{i_j} \in V_{i_j}$ for all j . If for at least one index ℓ we have $v_{i_\ell} \in W_{i_\ell, \text{Pet}}$, then if w is any $(r^* - 1)$ -fold commutator of $v_{i_1}, \dots, v_{i_{r^*}}$ we have*

$$\eta(w) = 0.$$

Furthermore, if instead for at least two indices ℓ_1, ℓ_2 we have $v_{i_{\ell_1}} \in W_{i_{\ell_1}, \text{Lin}}$ and $v_{i_{\ell_2}} \in W_{i_{\ell_2}, \text{Lin}}$, then if w is any $(r^* - 1)$ -fold commutator of $v_{i_1}, \dots, v_{i_{r^*}}$ we have

$$\eta(w) = 0.$$

Remark. The projection map $\text{Proj}_{W_{i, \text{Lin}}} : W_i \rightarrow W_{i, \text{Lin}}$ is well-defined due to the linear disjointness condition. Furthermore we have written Taylor coefficients with additive notation, since $G_{(i,1)}/G_{(i,2)}$ can be identified with $\mathbb{R}^{\dim(\text{Horiz}_i(G))}$.

Proof. For the majority of the proof we will assume $s \geq 3$; we indicate the minor changes required for $s = 2$ for the end of the proof (and the case $s = 2$ is not used in the proof of Theorem 1.2). Note that the case when η is trivial follows via taking $W_{i, \text{Pet}} = G_{(i,1)}/G_{(i,2)}$, $W_{i, \text{Lin}}$ and $W_{i, *}$ to be trivial, $g'_h = g_h$, and $g'(h, n) = (g(h, n), 0)$; therefore we may assume that η is nontrivial for the remainder of the proof.

Step 1: Applying Lemma 8.3 and linear-algebraic setup. We apply Lemma 8.3 and treat the resulting correlation structure as the input to the lemma. Up to changing implicit constants in the output this leaves the lemma unchanged except for noting that

$$\chi(h, n) = e(\Theta n) \cdot F^*(g(h, n)\Gamma^*)$$

which is defined on the group $(G^*)' = G^* \times \mathbb{R}$. In particular, we will abusively overwrite notation and relabel the resulting H' from the application of Lemma 8.3 as H , g'_h as g_h , and $\chi'_h(n)$ as $\chi_h(n)$ and thus assume the output properties without further comment.

It will also be crucial to define certain linear-algebraic operators of V_i . Consider a basis for $V_{i, \text{Dep}}$, an extension to a basis of V_i , and then to $G_{(i,1)}/G_{(i,2)}$ such that all basis elements are $(MD/\rho)^{O_s(d^{O_s(1)})}$ -rational combinations of the basis $\exp(\mathcal{X}_i) \bmod G_{(i,2)}$. In particular, write

$$\begin{aligned} V_{i, \text{Dep}} &= \text{span}_{\mathbb{R}}(w_{i,1}, \dots, w_{i, \dim(V_{i, \text{Dep}})}) \\ V_i &= \text{span}_{\mathbb{R}}(w_{i,1}, \dots, w_{i, \dim(V_{i, \text{Dep}})}, w_{i, \dim(V_{i, \text{Dep}})+1}, \dots, w_{i, \dim(V_i)}) \\ G_{(i,1)}/G_{(i,2)} &= \text{span}_{\mathbb{R}}(w_{i,1}, \dots, w_{i, \dim(\text{Horiz}_i(G))}). \end{aligned}$$

Given $v \in V_i$, there is a unique linear combination

$$v = \sum_{j=1}^{\dim(V_i)} \alpha_j w_{i,j}.$$

We define

$$P_i v = \sum_{j=\dim(V_{i, \text{Dep}})+1}^{\dim(V_i)} \alpha_j w_{i,j}, \quad Q_i v = \sum_{j=1}^{\dim(V_{i, \text{Dep}})} \alpha_j w_{i,j}.$$

By construction $P_i^2 = P_i$, $Q_i^2 = Q_i$, $Q_i(V_i) \cap P_i(V_i) = 0$, and $P_i v + Q_i v = v$ for $v \in V_i$. We also (abusively) extend the operator P_i to $V_i^{\otimes \ell}$ and $(G_{(i,1)}/G_{(i,2)})^{\otimes 4}$ in the obvious manners by acting on each copy of V_i separately (and zeroing out basis elements $w_{i, \dim(V_i)+1}, \dots, w_{i, \dim(\text{Horiz}_i(G))}$).

Step 2: Invoking equidistribution theory. Applying Lemma 7.5 when $s \geq 3$, we have

$$\|\mathbb{E}[\chi_{h_1}(n) \otimes \xi_{h_2}(n + h_1 - h_4) \otimes \overline{\xi_{h_3}(n)} \otimes \overline{\xi_{h_4}(n + h_1 - h_4)} \cdot \psi_{\vec{h}}(g_{\vec{h}}(n)\Gamma')]\|_{\infty} \geq (MD/\rho)^{-O_s(d^{O_s(1)})}$$

for a $(MD/\rho)^{-O_s(d^{O_s(1)})}$ density of additive tuples. We define G_{Error} , \tilde{G} , and η_{Prod} as in the proof of Lemma 8.3 and as before we may assume that $g_{\vec{h}}(0) = \text{id}_{G_{\text{Error}}}$. Define

$$g_{\vec{h}}^*(n) = (g_{h_1}(n), g_{h_2}(n + h_1 - h_4), g_{h_3}(n), g_{h_4}(n + h_1 - h_4), g_{\vec{h}}(n)).$$

By applying Corollary 5.5, we have

$$g_{\vec{h}}^* = \varepsilon_{\vec{h}} \cdot \tilde{g}_{\vec{h}} \cdot \gamma_{\vec{h}}$$

with

- $\varepsilon_{\vec{h}}(0) = \tilde{g}_{\vec{h}}(0) = \gamma_{\vec{h}}(0) = \text{id}_{\tilde{G}}$;
- $\tilde{g}_{\vec{h}}$ takes values in K ;
- $\gamma_{\vec{h}}$ is $(MD/\rho)^{O_s(d^{O_s(1)})}$ -rational;
- $d(\varepsilon(n), \varepsilon(n-1)) \leq (MD/\rho)^{O_s(d^{O_s(1)})} N^{-1}$ for $n \in [N]$.

where $\eta_{\text{Prod}}(K \cap \tilde{G}_{(s-1, r^*)}) = 0$ and K is a $(MD/\rho)^{O_s(d^{O_s(1)})}$ -rational subgroup of \tilde{G} . By passing to a subset of additive quadruples of density $(MD/\rho)^{-O_s(d^{O_s(1)})}$ we may in fact assume that the group K is independent of \vec{h} under consideration.

Step 3: Linear algebra deductions from equidistribution theory. Note that at present the subgroup K does not account for the deductions given in Lemma 8.3; these initial deductions are designed essentially to account for this. Let $\tau_i: \text{Horiz}_i(G)^{\otimes 4} \times \text{Horiz}_i(G_{\text{Error}}) \rightarrow \text{Horiz}_i(G)^{\otimes 4}$ be the natural projection to the four-fold product. We define the following set of vector spaces:

$$\begin{aligned} R_i &:= \{(Q_i v_1, Q_i v_2, Q_i v_3, Q_i v_4) \in V_i^{\otimes 4} : Q_i v_1 + Q_i v_2 - Q_i v_3 - Q_i v_4 = 0\}, \\ K_i &:= \tau_i(K \cap \tilde{G}_{(i,1)} \text{ mod } \tilde{G}_{(i,2)}), \\ S_i &:= \{(v_1, v_2, v_3, v_4) \in V_i^{\otimes 4} : P_i v_1 = P_i v_2 = P_i v_3 = P_i v_4\}, \\ K_{i,1} &:= K_i \cap S_i, \\ \widetilde{K}_i &:= K_{i,1} + R_i, \\ L_i &:= \pi_1(\widetilde{K}_i \cap \{(v_1, v_2, v_3, v_4) \in V_i^{\otimes 4} : Q_i v_2 = Q_i v_3 = Q_i v_4 = 0\}) + Q_i(V_i). \end{aligned}$$

By inspection, we have $R_i \leq S_i$ hence $\widetilde{K}_i \leq S_i$. Note that

$$\eta_{\text{Prod}}([K_{i_1,1}, \dots, K_{i_r^*,1}]) = 0$$

whenever one has that $i_1 + \dots + i_r^* = s - 1$ and $[\cdot, \dots, \cdot]$ denotes any possible $(r^* - 1)$ -fold commutator bracket. This is a consequence of the fact that $\eta_{\text{Prod}}(K \cap \tilde{G}_{(s, r^*)}) = \text{Id}_{\tilde{G}}$ and noting that $K_{i,1} \leq K_i$. Note that we are implicitly using that η_{Prod} is trivial on G_{Error} as well, and we abusively descend η_{Prod} to $G^{\otimes 4}$.

We now claim that

$$\eta_{\text{Prod}}([v_{i_1}, \dots, v_{i_r^*}]) = 0$$

if $v_{i_\ell} \in S_{i_\ell}$ for all ℓ and there is at least one index j such that $v_{i_j} \in R_i$.

To prove this, note by the final bullet point of Lemma 8.3 and multilinearity that

$$\begin{aligned} \eta_{\text{Prod}}([v_{i_1}, \dots, v_{i_r^*}]) &= \eta_{\text{Prod}}([P_{i_1} v_{i_1}, \dots, P_{i_r^*} v_{i_r^*}]) + \sum_{k=1}^r \eta_{\text{Prod}}([P_{i_1} v_{i_1}, \dots, Q_{i_k} v_{i_k}, \dots, P_{i_r^*} v_{i_r^*}]) \\ &= \eta_{\text{Prod}}([P_{i_1} v_{i_1}, \dots, Q_{i_j} v_{i_j}, \dots, P_{i_r^*} v_{i_r^*}]) = 0. \end{aligned}$$

The first equality uses that every bracket with at least two $Q_{i_k} v_{i_k}$ has two $V_{i, \text{Dep}}$ terms so is 0, the second equality uses $P_{i_j} v_{i_j} = 0$, and the third equality follows by noting that

$$P_i v_i \in \{(v_1, v_2, v_3, v_4) \in V_i^{\otimes 4} : P_i v_1 = P_i v_2 = P_i v_3 = P_i v_4, Q_i v_1 = Q_i v_2 = Q_i v_3 = Q_i v_4 = 0\}$$

and $\eta_{\text{Prod}} = (\eta, \eta, -\eta, -\eta)$. Now, we may ultimately deduce

$$\eta_{\text{Prod}}([\widetilde{K}_{i_1}, \dots, \widetilde{K}_{i_r^*}]) = 0$$

because $\widetilde{K}_i = K_{i,1} + R_i$ and $R_i, K_{i,1} \leq S_i$.

Finally, let π_T for $T \subseteq [4]$ is as in the proof of Lemma 8.3 (namely, an appropriate projection map). We have

$$\pi_1(\widetilde{K}_i) \leq L_i.$$

This follows because if

$$((Qv_1, Pv_1), (Qv_2, Pv_2), (Qv_3, Pv_3), (Qv_4, Pv_4)) \in \widetilde{K}_i$$

then

$$((Q(v_1 + v_2 - v_3 - v_4), Pv_1), (0, Pv_2), (0, Pv_3), (0, Pv_4)) \in \widetilde{K}_i.$$

Step 4: Constructing a decomposition of $Q_i(V_i)$. We will now decompose $Q_i(V_i) = V_{i,\text{Dep}}$ into a pair of subspaces. On one of these subspaces we will deduce an improved vanishing for the commutator while on the other subspace we will deduce an approximate linearity for $\text{Taylor}_i(g_h)$. Let

$$L_i^* = \{(v_1, v_2, v_3, v_4) \in S_i : Pv_1 = 0, v_2 = v_3 = v_4 = 0\} \cap \widetilde{K}_i.$$

Note that L_i^* may abusively be viewed as a subspace of V_i (instead of $V_i^{\otimes 4}$) and under this identification $L_i^* \leq Q_i(V_i) = V_{i,\text{Dep}} \leq L_i$.

The key claim in our analysis is if $i_1 + \dots + i_{r^*} = s - 1$, $v_{i_\ell} \in L_{i_\ell}$ for all indices ℓ , and $v_{i_j} \in L_{i_j}^*$ for at least one index j we have

$$\eta([v_{i_1}, \dots, v_{i_\ell}]) = 0.$$

To prove this, note that $Q_{i_j}v_{i_j} = v_{i_j}$ and $P_{i_j}v_{i_j} = 0$ and using the last bullet point of Lemma 8.3, we have

$$\eta([v_{i_1}, \dots, v_{i_\ell}]) = \eta([P_{i_1}v_{i_1}, \dots, Q_{i_j}v_{i_j}, \dots, P_{i_\ell}v_{i_\ell}]),$$

similar to the argument in Step 3.

Next note that $P_i Q_i v = 0$ for all $v \in V_i$ and therefore

$$P_i(L_i) \leq P_i(\pi_1(\widetilde{K}_i \cap \{(v_1, v_2, v_3, v_4) \in V_i^{\otimes 4} : Q_i v_2 = Q_i v_3 = Q_i v_4 = 0\})).$$

Therefore we may lift $P_{i_\ell}v_{i_\ell}$ for $\ell \neq j$ to $\widetilde{v}_{i_\ell} = (P_{i_\ell}v_{i_\ell} + w_{i_\ell}, P_{i_\ell}v_{i_\ell}, P_{i_\ell}v_{i_\ell}, P_{i_\ell}v_{i_\ell}) \in \widetilde{K}_{i_\ell}$ where $w_{i_\ell} \in Q_{i_\ell}(V_{i_\ell})$. We lift v_{i_j} to \widetilde{v}_{i_j} which has the form $(Q_{i_j}v_{i_j}, 0, 0, 0) \in \widetilde{K}_{i_j}$.

Note that we have

$$\begin{aligned} 0 &= \eta_{\text{Prod}}([\widetilde{v}_{i_1}, \dots, \widetilde{v}_{i_{r^*}}]) = \eta([P_{i_1}v_{i_1} + w_{i_1}, \dots, Q_{i_j}v_{i_j}, \dots, P_{i_{r^*}}v_{i_{r^*}} + w_{i_{r^*}}]) \\ &= \eta([P_{i_1}v_{i_1}, \dots, Q_{i_j}v_{i_j}, \dots, P_{i_{r^*}}v_{i_{r^*}}]) \end{aligned}$$

where in the first equality we have used for all ℓ that $\widetilde{v}_{i_\ell} \in \widetilde{K}_{i_\ell}$ and the result from Step 3, in the second equality that \widetilde{v}_{i_j} has the final three coordinates identically zero, and in the final equality that $w_{i_\ell} \in Q_{i_\ell}(V_{i_\ell}) = V_{i_\ell,\text{Dep}}$ and the final item of Lemma 8.3.

The desired decomposition of spaces for the lemma will have

$$W_{i,\text{Pet}} := L_i^*, \quad W_{i,*} := P_i(L_i) \leq L_i \cap P_i(V_i)$$

$W_{i,\text{Lin}}$ will be constructed explicitly in the next step but is chosen so that

$$W_{i,\text{Lin}} \leq Q_i(L_i) = Q_i(V_i) = V_{i,\text{Dep}}$$

and $W_{i,\text{Lin}} + W_{i,\text{Pet}} = V_{i,\text{Dep}} \leq L_i$. Given these properties of $W_{i,\text{Lin}}$, note that the above analysis, along with Lemma 8.3, establishes the final bullet point for our output.

Step 5: Controlling approximate homomorphisms. Recall $\widetilde{K}_i \leq S_i$ and there is a natural isomorphism of groups

$$S_i \simeq \{(v, v_1, v_2, v_3, v_4) : v \in P_i(V_i), v_1, \dots, v_4 \in Q_i(V_i)\}.$$

Using this as an identification, we may write

$$\widetilde{K}_i = \bigcap_{j=1}^{\dim(S_i) - \dim(\widetilde{K}_i)} \ker((\xi_j^{P_i}, \xi_j^{Q_i}, \xi_j^{Q_i}, -\xi_j^{Q_i}, -\xi_j^{Q_i}))$$

where $\xi_j^{P_i} \in P_i(V_i)^\vee$ and $\xi_j^{Q_i} \in Q_i(V_i)^\vee$ (i.e., corresponding dual vector spaces). Note that the annihilators all have the special form of $(\cdot, \xi_j^{Q_i}, \xi_j^{Q_i}, -\xi_j^{Q_i}, -\xi_j^{Q_i})$ since

$$\{(Q_i v_1, Q_i v_2, Q_i v_3, Q_i v_4) \in V_i^{\otimes 4} : Q_i v_1 + Q_i v_2 - Q_i v_3 - Q_i v_4 = 0\} = R_i \leq \widetilde{K}_i.$$

Note that

$$L_i^* = \{v \in V_i : P_i v = 0 \text{ and } \xi_j^{Q_i}(Q_i v) = 0 \text{ for all } j\}$$

since $v \in L_i^*$ is equivalent under this identification to $(0, Q_i v, 0, 0, 0) \in \widetilde{K}_i$. Without loss of generality we may assume that for $1 \leq j \leq \dim(V_{i, \text{Dep}}) - \dim(L_i^*)$, vectors $\xi_j^{Q_i}$ are independent in $Q_i(V_i)^\vee$ (and they must span the orthogonal space to L_i^* within $Q_i(V_i)^\vee$).

By appropriate scaling, we may assume $\xi_j^{Q_i}(w_{i,j})$ is an integer bounded by $(MD/\rho)^{O_s(d^{O_s(1)})}$ for $1 \leq j \leq \dim(V_{i, \text{Dep}})$. We extend each $\xi_j^{Q_i}$ to an operator on $(G_{(i,1)}/G_{(i,2)})^\vee$ by setting $\xi_j^{Q_i}(w_{i,j}) = 0$ for $j > \dim(V_{i, \text{Dep}})$. Possibly at the cost of another $(MD/\rho)^{O_s(d^{O_s(1)})}$ scaling, we may assume that $\xi_j^{Q_i}(\Gamma \cap G_{(i,1)} \bmod G_{(i,2)}) \in \mathbb{Z}$. We extend $\xi_j^{P_i}(\cdot)$ in an analogous manner to $(G_{(i,1)}/G_{(i,2)})^\vee$ by setting $\xi_j^{P_i}(w_{i,j}) = 0$ for $1 \leq j \leq \dim(V_{i, \text{Dep}})$ and $j > \dim(V_i)$. Again, we may scale such that $\xi_j^{P_i}(\Gamma \cap G_{(i,1)} \bmod G_{(i,2)}) \in \mathbb{Z}$. The crucial point here is that now $\xi_j^{P_i}$ and $\xi_j^{Q_i}$ are i -th horizontal characters of height at most $(MD/\rho)^{O_s(d^{O_s(1)})}$.

We have

$$\tau_i(\text{Taylor}_i(\widetilde{g}_h)) \in K_i$$

and thus

$$\text{dist}(\tau_i(\text{Taylor}_i(g_h^*)), S_i + T^{-1} \text{Horiz}_i(\Gamma^{\otimes 4})) \leq (MD/\rho)^{O_s(d^{O_s(1)})} N^{-i}$$

after Pigeonholing \vec{h} appropriately. Here distance is in L^∞ after expressing both of these expressions in the basis $\exp(\mathcal{X}_i)^{\otimes 4}$ and T is an integer bounded by $(MD/\rho)^{O_s(d^{O_s(1)})}$. We have used Lemma 2.13 and the properties of the original factorization; a very similar argument appears in Step 5 of the proof of Lemma 8.3.

Furthermore note that

$$\tau_i(\text{Taylor}_i(g_h^*)) \in S_i$$

by Lemma 8.3. So if we choose a set of horizontal characters of height $(MD/\rho)^{O_s(d^{O_s(1)})}$ relative to $\text{Horiz}_i(\Gamma^{\otimes 4})$ which cut out \widetilde{K}_i as their common kernel, then noting that $K_i \cap S_i \leq \widetilde{K}_i$ and applying Lemma B.2 we may assume that

$$(9.1) \quad \tau_i(\text{Taylor}_i(\widetilde{g}_h)) \in \widetilde{K}_i$$

and $\varepsilon_{\vec{h}}, \gamma_{\vec{h}}$ have identical properties up to changing implicit constants. We will assume this refined property of the factorization for the remainder of our analysis.

Given the factorization of g_h^* , we thus deduce (taking an appropriate least common multiple)

$$(9.2) \quad \left\| T_1 \cdot \left(\xi_j^{P_i}(\text{Taylor}_i(g_{h_2})) + \xi_j^{Q_i}(\text{Taylor}_i(g_{h_1})) + \xi_j^{Q_i}(\text{Taylor}_i(g_{h_2})) \right. \right. \\ \left. \left. - \xi_j^{Q_i}(\text{Taylor}_i(g_{h_3})) - \xi_j^{Q_i}(\text{Taylor}_i(g_{h_4})) \right) \right\|_{\mathbb{R}/\mathbb{Z}} \leq (MD/\rho)^{O_s(d^{O_s(1)})} N^{-i}$$

for all $1 \leq j \leq \dim(V_{i, \text{Dep}}) - \dim(L_i^*)$ where T_1 is an integer bounded by $(MD/\rho)^{O_s(d^{O_s(1)})}$. Here we have used that $\xi_j^{P_i}(\text{Taylor}_i(g_h))$ is equal for all $h \in H$ by Lemma 8.3.

We define functions $f, g: H \rightarrow \mathbb{R}^{\sum_{i=1}^{s-1} \dim(V_{i, \text{Dep}}) - \dim(L_i^*)}$ via

$$f(h) = (T_1 \xi_j^{Q_i}(\text{Taylor}_i(g_h)))_{1 \leq i \leq s-1, 1 \leq j \leq \dim(V_{i, \text{Dep}}) - \dim(L_i^*)}, \\ g(h) = (T_1 \xi_j^{P_i}(\text{Taylor}_i(g_h)) + T_1 \xi_j^{Q_i}(\text{Taylor}_i(g_h)))_{1 \leq i \leq s-1, 1 \leq j \leq \dim(V_{i, \text{Dep}}) - \dim(L_i^*)}.$$

Note that for the additive quadruples on which we have (9.2), we are exactly in the situation necessary to apply results on approximate homomorphisms.

In particular, we may apply Lemma A.1. We see that there exists $H' \subseteq H$ having density at least $\exp(-O_s(d \log(MD/\rho))^{O_s(1)})$ such that for all i, j and $h \in H'$, we have

$$(9.3) \quad \left\| T_1 \xi_j^{Q_i}(\text{Taylor}_i(g_h)) - \left(\gamma_{i,j} + \sum_{k=1}^{d^*} \alpha_{i,j,k} \{ \beta_k h \} \right) \right\|_{\mathbb{R}/\mathbb{Z}} \leq (MD/\rho)^{O_s(d^{O_s(1)})} N^{-i},$$

where:

- $d^* \leq (d \log(MD/\rho))^{O_s(1)}$;
- $\beta_k \in (1/N')\mathbb{Z}$ where N' is a prime between $100N$ and $200N$.

At this point, for each i we find elements $Z_{i,j}$ for $1 \leq j \leq \dim(V_{i,\text{Dep}}) - \dim(L_i^*)$ which are $(MD/\rho)^{O_s(d^{O_s(1)})}$ -rational combinations of $\{w_{i,j} : 1 \leq j \leq \dim(V_{i,\text{Dep}})\}$ such that

$$(9.4) \quad T_1 \xi_j^{Q_i}(Z_{i,j}) = 1 \text{ and } \xi_{j'}^{Q_i}(Z_{i,j}) = 0$$

for $j' \neq j$ such that $1 \leq j, j' \leq \dim(V_{i,\text{Dep}}) - \dim(L_i^*)$. We define

$$W_{i,\text{Lin}} = \text{span}_{\mathbb{R}}((Z_{i,j})_{1 \leq j \leq \dim(V_{i,\text{Dep}}) - \dim(L_i^*)}).$$

We see that there are no nontrivial linear relations between $W_{i,*}$ and $W_{i,\text{Pet}} + W_{i,\text{Lin}}$ since $W_{i,*} \leq P_i(V_i)$ and $W_{i,\text{Pet}} + W_{i,\text{Lin}} \leq Q_i(V_i)$. There are no linear relations between $W_{i,\text{Pet}}$ and $W_{i,\text{Lin}}$ as $W_{i,\text{Pet}}$ lies in the joint kernel of the $\xi_j^{Q_i}$ and therefore using (9.4) one can prove any such relation is trivial. Furthermore, by construction we have $V_{i,\text{Dep}} = Q_i(V_i) = W_{i,\text{Lin}} + W_{i,\text{Pet}}$. Finally $L_i = P_i(L_i) + Q_i(L_i) = W_{i,*} + W_{i,\text{Lin}} + W_{i,\text{Pet}}$; this implicitly uses $Q(L_i) = Q_i(V_i) \leq L_i$.

Step 6: Constructing the desired factorizations and completing the proof. Using the refined factorization (9.1) implies that

$$\pi_1(\tau(\text{Taylor}_i(\tilde{g}_h))) \in L_i$$

since $\pi_1(\tilde{K}_i) \leq L_i$. Applying $g_h^* = \varepsilon_h \cdot \tilde{g}_h \cdot \gamma_h$ in the first coordinate then implies that

$$(9.5) \quad \text{dist}(\text{Taylor}_i(g_{h_1}), L_i + T_2^{-1} \text{Horiz}_i(\Gamma)) \leq (MD/\rho)^{O_s(d^{O_s(1)})} N^{-i}$$

for $h_1 \in H'$ where distance is in L^∞ after expressing values in terms of $\exp(\mathcal{X}_i)$. Here T_2 is an integer bounded by $(MD/\rho)^{O_s(d^{O_s(1)})}$. Furthermore recall from Lemma 8.3 that

$$(9.6) \quad \text{Taylor}_i(g_h) - \text{Taylor}_i(g_{h'}) \in V_{i,\text{Dep}} = Q_i(V_i)$$

for $h, h' \in H' \subseteq H$.

Let $Y_{i,j} \in \text{span}_{\mathbb{R}}(\mathcal{X} \cap \log(G_{(i,1)}) \setminus \mathcal{X} \cap \log(G_{(i,2)}))$ be such that $\exp(Y_{i,j}) \bmod G_{(i,2)} = Z_{i,j}$. Then for $h \in H'$, we define

$$(9.7) \quad \tilde{g}_h(n) = \prod_{i=1}^s \prod_{j=1}^{\dim(W_{i,\text{Lin}})} \exp(Y_{i,j})^{T_1^{-1}(n) \cdot (\gamma_{i,j} + \sum_{k=1}^{d^*} \alpha_{i,j,k} \{ \beta_k h \})}.$$

By construction and Lemma 2.13, for $h, h' \in H'$ we have

$$\begin{aligned} \text{Taylor}_i(\tilde{g}_h^{-1} g_h) - \text{Taylor}_i(\tilde{g}_{h'}^{-1} g_{h'}) &\in Q_i(V_i), \\ \text{dist}(\text{Taylor}_i(\tilde{g}_h^{-1} g_h), L_i + T_2^{-1} \text{Horiz}_i(\Gamma)) &\leq (MD/\rho)^{O_s(d^{O_s(1)})} \cdot N^{-i}, \\ \|T_1 \xi_j^{Q_i}(\tilde{g}_h^{-1} g_h)\|_{\mathbb{R}/\mathbb{Z}} &\leq (MD/\rho)^{O_s(d^{O_s(1)})} \cdot N^{-i}, \end{aligned}$$

where $1 \leq i \leq s-1$ and $1 \leq j \leq \dim(V_{i,\text{Dep}}) - \dim(L_i^*)$. The first line comes from (9.6), the second line from (9.5), and the third from (9.3) and (9.7), in conjunction with (9.7).

We now fix an element $h_2 \in H'$. For each $h_1 \in H'$ we write

$$g'_{h_1} = \tilde{g}_{h_1} \cdot (\tilde{g}_{h_1}^{-1} g'_{h_1}) = \tilde{g}_{h_1} \cdot (\tilde{g}_{h_1}^{-1} g'_{h_1}) \cdot (\tilde{g}_{h_2}^{-1} g'_{h_2})^{-1} \cdot (\tilde{g}_{h_2}^{-1} g'_{h_2}).$$

By applying Lemma B.2, we may write

$$(\tilde{g}_{h_1}^{-1} g'_{h_1}) \cdot (\tilde{g}_{h_2}^{-1} g'_{h_2})^{-1} = \varepsilon_{h_1}^* g_{h_1}^* \gamma_{h_1}^*, \quad (\tilde{g}_{h_2}^{-1} g'_{h_2}) = \varepsilon^* g^* \gamma^*$$

where $\gamma^*, \gamma_{h_1}^*$ are $(MD/\rho)^{O_s(d^{O_s(1)})}$ -rational, $\varepsilon^*, \varepsilon_{h_1}^*$ are $((MD/\rho)^{O_s(d^{O_s(1)})}, N)$ -smooth, and we have $\text{Taylor}_i(g_{h_1}^*) \in L_i^* = W_{i, \text{Pet}}$ using the first and third lines above and $\text{Taylor}_i(g^*) \in L_i^* + P_i(L_i) = W_{i, *}$ using the second and third lines above. (Recall that $L_i^* \leq Q_i(V_i)$ is cut out by the $\xi_j^{Q_i}$.) Additionally, these sequences are the identity at 0.

Therefore, for $h_1 \in H'$ we have

$$\begin{aligned} g'_{h_1} &= \varepsilon_{h_1}^* \varepsilon^* ((\varepsilon_{h_1}^* \varepsilon^*)^{-1} \tilde{g}_{h_1} (\varepsilon_{h_1}^* \varepsilon^*)) ((\varepsilon^*)^{-1} g_{h_1}^* \varepsilon^*) ((\varepsilon^*)^{-1} \gamma_{h_1}^* \varepsilon^* (\gamma_{h_1}^*)^{-1}) (\gamma_{h_1}^* g^* (\gamma_{h_1}^*)^{-1}) (\gamma_{h_1}^* \gamma^*) \\ &=: (\varepsilon_{h_1}^* \varepsilon^*) \cdot g_{h_1}^\Delta \cdot (\gamma_{h_1}^* \gamma^*). \end{aligned}$$

So, for $h_3, h_4 \in H$ we deduce using Lemma 2.13 and the above analysis that

$$\begin{aligned} \text{Taylor}_i(\tilde{g}_{h_3}) &\in W_{i, \text{Lin}}, \\ \text{Taylor}_i(g_{h_3}^\Delta) &\in L_i = W_{i, *} + W_{i, \text{Lin}} + W_{i, \text{Pet}}, \\ \text{Proj}_{W_{i, \text{Lin}}}(\text{Taylor}_i(g_{h_3}^\Delta)) &= \text{Proj}_{W_{i, \text{Lin}}}(\text{Taylor}_i(\tilde{g}_{h_3})), \\ \text{Taylor}_i(\tilde{g}_{h_3}^{-1} g_{h_3}^\Delta) - \text{Taylor}_i(\tilde{g}_{h_4}^{-1} g_{h_4}^\Delta) &\in W_{i, \text{Pet}}. \end{aligned}$$

Furthermore note that $\varepsilon_{h_1}^* \varepsilon^*$ is sufficiently smooth and $\gamma_{h_1}^* \gamma^*$ is appropriately rational. This nearly gives the desired result except we need to remove the rational and smooth parts exactly as in Step 7 of Lemma 8.3; we omit the details, although note that the only difference between g_h^Δ and the output is a conjugation by a fixed element which leaves all properties unchanged and the Fourier phase on the \mathbb{R} part of $(G^*)'$ may be modified. Additionally, the set H' will be made smaller by acceptable factors due to Pigeonhole.

Step 7: Handling the exceptional case $s = 2$. In this exceptional case, we have $r^* = 1$ and $s = 2$, and η is nontrivial. The difference here versus the prior analysis is that the error term $\psi_h(g_{\bar{h}}(n)\Gamma')$ is replaced by $e(\Theta_{\bar{h}}n)$ with $\|\Theta_{\bar{h}}\|_{\mathbb{R}/\mathbb{Z}} \leq (MD/\rho)^{O_s(d^{O_s(1)})} N^{-1}$ by using the Remark 7.6 regarding Lemma 7.5 for $s = 2$.

We take $G^{\text{Error}} = \mathbb{R}$, $\Gamma^{\text{Error}} = \mathbb{Z}$, $g_{\bar{h}}(n) = \Theta_{\bar{h}}n$, and $\psi_{\bar{h}}(z) = e(z)$. \tilde{G} is defined as before. Taking $\eta^* = (\eta, \eta, -\eta, -\eta, 1)$, by Corollary 5.5 we may factor

$$g_{\bar{h}}^* = \varepsilon_{\bar{h}} \cdot \tilde{g}_{\bar{h}} \cdot \gamma_{\bar{h}}$$

where $\varepsilon_{\bar{h}}$ is $((MD/\rho)^{O_s(d^{O_s(1)})}, N)$ -smooth, $\gamma_{\bar{h}}$ is $(MD/\rho)^{O_s(d^{O_s(1)})}$ -rational, and $\tilde{g}_{\bar{h}}$ lies in a $(MD/\rho)^{O_s(d^{O_s(1)})}$ -rational subgroup K such that $\eta^*(K \cap \tilde{G}_{(1,1)}) = 0$. Note however that

$$g_{\bar{h}}^* = (\text{id}_G, \text{id}_G, \text{id}_G, \text{id}_G, \Theta_{\bar{h}}n) \cdot (\tau(g_{\bar{h}}^*), 0)$$

where $\tau: \tilde{G} \rightarrow G^{\otimes 4}$ is the natural projection. Let $K^* = K \cap (G^{\otimes 4} \times \{0\})$ and note that K^* can be defined as the joint kernel of certain horizontal characters of height $(MD/\rho)^{O_s(d^{O_s(1)})}$ (namely, ones defining K along with one of the form $(0, 0, 0, 0, 1)$). Since $\Theta_{\bar{h}}$ is small and is the only part in the fifth coordinate, arguments similar to before allow us to refine the first factorization (up to changing implicit constants) and instead assume that $\tilde{g}_{\bar{h}}$ lies in K^* .

Furthermore note that if $\eta_{\text{Prod}} = (\eta, \eta, -\eta, -\eta, 0)$ we have that $\eta_{\text{Prod}}(K^*) = 0$ as η_{Prod} and η^* agree on the initial four groups. At this point we are exactly in the situation of the earlier analysis and we may complete the proof.³ \square

We remark that modulo minor annoyances, the strategy of using Lemma 7.5, deducing an approximate homomorphism, and then applying results coming from the Bogolyubov lemma was introduced by Gowers [11] in his seminal work on four-term arithmetic progressions. It was similarly applied in work of Green and Tao [18] on the U^3 -inverse theorem. In certain sense, the previous two sections can be thought of as showing that, given an appropriate equidistribution theorem and defining a number of notions for nilmanifolds, this analysis can be modified to make sense in the greater generality of nilmanifolds where the group is not abelian.

10. SETUP FOR EXTRACTING A $(1, s - 1)$ -NILSEQUENCE

Before diving into the formal proof, we motivate how we extract the “top degree-rank” part and why lifting to the universal nilmanifold plays a role in our argument at this stage. We remark that Green, Tao, and Ziegler [29] work with the universal nilmanifold throughout their argument (in the form of a representation of a degree-rank nilcharacter; see [29, Definition 9.11]).

Recall the bracket polynomial U^5 -inverse sketch discussed in Section 4; we started with functions

$$e\left(\sum_{i=1}^{d_1} a_{i,h}n[b_{i,h}n][c_{i,h}n] + \sum_{i=1}^{d_2} d_{i,h}n^2[e_{i,h}n] + \sum_{i=1}^{d_3} f_{i,h}n[g_{i,h}n] + j_hn^3 + \ell_hn^2 + m_hn\right)$$

which correlate with $\Delta_h f$. At this point, we have proven that

$$\sum_{i=1}^{d_1} a_{i,h}n[b_{i,h}n][c_{i,h}n]$$

is equivalent to a bracket polynomial up to lower order terms of degree-rank of the form

$$\sum_{i=1}^{d'_1} \delta_i \{\varepsilon_i h\} n[\beta_{i,*}n][\gamma_{i,*}n].$$

Our goal at this stage is to isolate

$$e\left(\sum_{i=1}^{d'_1} \delta_i \{\varepsilon_i h\} n[\beta_{i,*}n][\gamma_{i,*}n]\right);$$

in the next section we will then convert this “top degree-rank” bracket phase into a $(1, s - 1)$ -nilsequence.

The reason lifting to a universal nilmanifold proves so technically useful is that it enables us to isolate various components of the horizontal tori as “separate subgroups”. For the sake of simplicity, consider a 2-step group G in the U^4 -inverse case given the degree-rank filtration $G_{(0,0)} = G_{(1,0)} = G_{(1,1)} = G$, $G_{(2,0)} = G_{(2,1)} = G_{(2,2)} = [G, G]$, where the remaining groups are trivial. In this case, the output of Lemma 9.1 gives the linearly disjoint subspaces W_* , W_{Lin} , W_{Pet} of $V = G/[G, G]$ such that the commutator of any two elements in $W_{\text{Lin}} + W_{\text{Pet}}$ vanishes and the commutator of any element of W_* and W_{Pet} vanishes.

Let \mathcal{Z}_* denote the rational basis of $\log(W_*) \bmod [G, G]$ and \mathcal{Z}_{Lin} and \mathcal{Z}_{Pet} be analogous. We also have a decomposition of our polynomial

$$g_h = g_{h,*} + g_{h,\text{Lin}} + g_{h,\text{Pet}} \bmod [G, G]$$

³Various simplifications are possible in the case since the underlying groups are all abelian here; in particular, invoking Corollary 5.5 reduces to summing a geometric series.

where

$$g_{h,\text{Lin}}(n) = \prod_{i=1}^{\dim(W_{\text{Lin}})} \exp(\delta_i n \{\varepsilon_i h\} Z_i^{\text{Lin}}), \quad g_{h,*}(n) = \prod_{i=1}^{\dim(W_*)} \exp(\beta_i n Z_i^*),$$

$$g_{h,\text{Pet}}(n) = \prod_{i=1}^{\dim(W_{\text{Pet}})} \exp(\gamma_i^h n Z_i^{\text{Pet}}).$$

Therefore we may write

$$g_h(n) = g_{h,*}(n) g_{h,\text{Lin}}(n) g_{h,\text{Pet}}(n) g_{h,\text{Rem}}(n)$$

with $g_{h,\text{Rem}}(n) \in [G, G]$ pointwise. The top order term which we seek to isolate is heuristically similar to

$$e \left(\sum_{i=1}^{\dim(W_{\text{Lin}})} \sum_{j=1}^{\dim(W_*)} \delta_i n \{\varepsilon_i h\} [\beta_j n] [Z_i^{\text{Lin}}, Z_j^*] \right).$$

Note that given the factorization of g_h , we have established no control over $g_{h,\text{Pet}}$ and $g_{h,\text{Rem}}$. This may suggest that we wish to quotient out by the subgroup $W_{\text{Pet}}[G, G]$ in order to kill these terms; note however that $G/(W_{\text{Pet}}[G, G])$ is now abelian and such a projection “kills” the higher order degree-rank term calculated above. This suggests that the group $W_{\text{Pet}}[G, G]$ is “too large” a quotient. The solution is to “enlarge” the group G so that the subgroup $[W_{\text{Lin}}, W_*]$ and the subgroup G' that corresponds to the remaining phases $\gamma_h n^2 + \delta_h n$ are disjoint. We can then quotient by $W_{\text{Pet}}G'$. This disjointness is accomplished by lifting to the universal nilmanifold of degree-rank $(2, 2)$.

10.1. Unwinding the output of Lemma 9.1. We first require the following elementary lemma regarding lattice elements when presented in first-kind coordinates.

Lemma 10.1. *Fix an integer $k \geq 1$. Consider a nilmanifold G/Γ of dimension d with a Mal'cev basis $\mathcal{X} = \{X_1, \dots, X_d\}$ of $\log G$ which is Q -rational and such that \mathcal{X} has the degree k nesting property. Then there exists a positive integer $Q' \leq O_k(Q^{O_k(d^{O_k(1)})})$ such that if $z_j \in Q' \cdot \mathbb{Z}$ then*

$$\exp \left(\sum_{j=1}^d z_j X_j \right) \in \Gamma.$$

Proof. Note that $\Gamma = \psi_{\mathcal{X}}(\mathbb{Z}^d)$. By [35, Lemma B.1], $\psi_{\mathcal{X}} \circ \psi_{\text{exp}, \mathcal{X}}^{-1}$ is a degree $O_k(1)$ polynomial with coefficients of height at most $Q^{O_k(d^{O_k(1)})}$. The desired result then follows by taking Q' to the least common multiple of all denominators of all coefficients present in this polynomial (since there are only $O_k(d^{O_k(1)})$ total coefficients). Note that the polynomial corresponding to $\psi_{\mathcal{X}} \circ \psi_{\text{exp}, \mathcal{X}}^{-1}$ has no constant term by observing the image of id_G . \square

We next require the following additional elementary lemma which gives a Taylor series expansion which is “graded by the Mal'cev basis”.

Lemma 10.2. *Consider a nilmanifold G/Γ of degree k with an adapted Mal'cev basis $\mathcal{X} = \{X_1, \dots, X_{\dim(G)}\}$ and a polynomial sequence $g(n)$. There exists a representation*

$$g(n) = \prod_{i=0}^k \prod_{j=\dim(G)-\dim(G_k)+1}^{\dim(G)} \exp(X_j)^{\alpha_{i,j} \cdot \frac{n^i}{i!}}$$

where $\alpha_{i,j} \in \mathbb{R}$.

Proof. Note via Baker–Campbell–Hausdorff and existence of Taylor expansions, we may write

$$g(n) = \exp\left(\sum_{i=0}^s g_i \cdot \frac{n^i}{i!}\right)$$

with $g_i \in \log(G_i)$. Let $g_0(n) = g(n)$ and $g_{0,i} = g_i$. Then iteratively define $g_{\ell+1}(n)$ by the following process: write $\sum_{j=\dim(G)-\dim(G_{(\ell,0)})+1}^{\dim(G)} \alpha_{\ell,j} X_j = g_{\ell,\ell}$. Then let

$$g_{\ell+1}(n) := \left(\prod_{j=\dim(G)-\dim(G_\ell)+1}^{\dim(G)} \exp(X_j)^{\alpha_{\ell,j} \cdot \frac{n^\ell}{\ell!}} \right)^{-1} g_\ell(n)$$

and write

$$g_{\ell+1}(n) = \exp\left(\sum_{i=\ell+1}^s g_{\ell+1,i} \cdot \frac{n^i}{i!}\right)$$

in order to define $g_{\ell+1,i}$. There exists a valid choice of $\alpha_{\ell,j}$ at each step since \mathcal{X} is a filtered Mal'cev basis and there exists a valid choice of $g_{\ell+1,i}$ for $i \geq \ell+1$ by Baker–Campbell–Hausdorff. This process terminates with the identity sequence, and unraveling gives the desired. \square

Remark. Note that in the above proof, the reason we do not use the basis $\binom{n}{i} \binom{n}{j}$ is that $\binom{n}{i} \binom{n}{j}$ is *not* a linear combination of polynomials of the form $\binom{n}{t}$ for $t \geq \max(i, j) + 1$ and hence the Baker–Campbell–Hausdorff to construct $g_{\ell+1,i}$ fails (one needs lower-degree terms with $i \leq \ell$).

We now explicitly unwind, for the sake of clarity, the conclusion of Lemma 9.1. We will use the notation and conclusions here throughout the Sections 10 and 11. Suppose we have a 1-bounded function $f: [N] \rightarrow \mathbb{C}$ with a degree-rank $(s-1, r^*)$ correlation structure with parameters ρ, M, D, d .⁴ Then by Lemma 9.1 and some relabeling there exists a degree-rank $(s-1, r^*)$ correlation structure with parameters

$$\rho'^{-1} \leq \exp(O_s((d \log(MD/\rho))^{O_s(1)})), \quad M' \leq O(M), \quad D' = D, \quad d' \leq O(d)$$

and

- A subset $H \subseteq [N]$ with $|H| \geq \rho'N$;
- A multidegree $(1, s-1)$ nilcharacter $\chi(h, n)$ with a frequency η^* with height at most M . Furthermore χ lives on a nilmanifold $(G^* \times \mathbb{R})/(\Gamma^* \times \mathbb{Z})$ with dimension bounded by d' , output dimension bounded by D' , complexity bounded by M' , and the function underlying χ is M' -Lipschitz. We let $g(h, n)$ denote the underlying polynomial sequence;
- A collection of degree-rank $(s-1, r^*)$ nilcharacters $\chi_h(n)$ with a frequency η with of height at most M . Furthermore χ_h lives on a nilmanifold G/Γ with dimension bounded by d , output dimension bounded by D , G/Γ has complexity bounded by M and the function underlying χ_h (which is independent of h) is M' -Lipschitz. We let g_h denote the underlying polynomial sequence and we have $g_h(0) = \text{id}_G$;
- For all $h \in H$, we have

$$\Delta_h f(n) \otimes \chi(h, n) \otimes \chi_h(n) \in \text{Corr}(s-2, \rho', M', d');$$

- Then there exists a collection of subspaces $W_{i,*}, W_{i,\text{Lin}}, W_{i,\text{Pet}} \leq G_{(i,1)}/G_{(i,2)}$ for $1 \leq i \leq s-1$ which are $(MD/\rho)^{O_s(d^{O_s(1)})}$ -rational with respect to $\exp(\mathcal{X}) \cap G_{(i,1)} \bmod G_{(i,2)}$;
- If $W_i = W_{i,*} + W_{i,\text{Lin}} + W_{i,\text{Pet}}$ then $\dim(W_i) = \dim(W_{i,*}) + \dim(W_{i,\text{Lin}}) + \dim(W_{i,\text{Pet}})$;

⁴We apologize to the reader; there is a rather incredible amount of data which is floating around at this point. The crucial details to track are data regarding Taylor coefficient and the associated decompositions of the vector spaces corresponding to horizontal tori.

- Let $Z_{i,1}^*, \dots, Z_{i,\dim(W_{i,*})}^*$ a sequence of integral linear combinations of $\mathcal{X} \cap G_{(i,1)} \setminus \mathcal{X} \cap G_{(i,2)}$ such that $\text{span}_{\mathbb{R}}(\exp(Z_{i,1}^* \bmod G_{(i,2)}), \dots, \exp(Z_{i,\dim(W_{i,*})}^* \bmod G_{(i,2)})) = W_{i,*}$. We may let the coefficients of $Z_{i,j}^*$ be $(MD/\rho)^{O_s(d^{O_s(1)})}$ -bounded and $\exp(Z_{i,j}^*) \in \Gamma$.
- Let $Z_{i,1}^{\text{Lin}}, \dots, Z_{i,\dim(W_{i,\text{Lin}})}^{\text{Lin}}$ be a sequence of integral linear combinations of $\mathcal{X} \cap G_{(i,1)} \setminus \mathcal{X} \cap G_{(i,2)}$ such that $\text{span}_{\mathbb{R}}(\exp(Z_{i,1}^{\text{Lin}} \bmod G_{(i,2)}), \dots, \exp(Z_{i,\dim(W_{i,\text{Lin}})}^{\text{Lin}} \bmod G_{(i,2)})) = W_{i,\text{Lin}}$. We may let the coefficients of $Z_{i,j}^{\text{Lin}}$ are $(MD/\rho)^{O_s(d^{O_s(1)})}$ -bounded and $\exp(Z_{i,j}^{\text{Lin}}) \in \Gamma$.
- Let $Z_{i,1}^{\text{Pet}}, \dots, Z_{i,\dim(W_{i,\text{Pet}})}^{\text{Pet}}$ be a sequence of integral linear combinations of $\mathcal{X} \cap G_{(i,1)} \setminus \mathcal{X} \cap G_{(i,2)}$ such that $\text{span}_{\mathbb{R}}(\exp(Z_{i,1}^{\text{Pet}} \bmod G_{(i,2)}), \dots, \exp(Z_{i,\dim(W_{i,\text{Pet}})}^{\text{Pet}} \bmod G_{(i,2)})) = W_{i,\text{Pet}}$. We may let the coefficients of $Z_{i,j}^{\text{Pet}}$ be $(MD/\rho)^{O_s(d^{O_s(1)})}$ -bounded and let $\exp(Z_{i,j}^{\text{Pet}}) \in \Gamma$.
- For $1 \leq i \leq s-1$ and $h \in H$, we have

$$\begin{aligned} \text{Taylor}_i(g_h) &= \prod_{j=1}^{\dim(W_{i,*})} \exp(Z_{i,j}^*)^{z_{i,j}^*} \cdot \prod_{j=1}^{\dim(W_{i,\text{Pet}})} \exp(Z_{i,j}^{\text{Pet}})^{z_{i,j}^{h,\text{Pet}}} \\ &\quad \cdot \prod_{j=1}^{\dim(W_{i,\text{Lin}})} \exp(Z_{i,j}^{\text{Lin}})^{z_{i,j}^{h,\text{Lin}}} \bmod G_{(i,2)} \end{aligned}$$

where

$$z_{i,j}^{h,\text{Lin}} = \gamma_{i,j} + \sum_{k=1}^{d^*} \alpha_{i,j,k} \{\beta_k h\}$$

where $d^* \leq (d \log(MD/\rho))^{O_s(1)}$ and $\beta_k \in (1/N')\mathbb{Z}$ where N' is a prime in $[100N, 200N]$.

- For any integers $i_1 + \dots + i_{r^*} = s-1$, suppose that $v_{i_j} \in W_{i_j}$ for all j . If for at least one index ℓ we have $v_{i_\ell} \in W_{i_\ell, \text{Pet}}$, then if w is any $(r^* - 1)$ -fold commutator of $v_{i_1}, \dots, v_{i_{r^*}}$, we have

$$\eta(w) = 0.$$

Furthermore, if instead for at least two indices ℓ_1, ℓ_2 we have $v_{i_{\ell_1}} \in W_{i_{\ell_1}, \text{Lin}}$ and $v_{i_{\ell_2}} \in W_{i_{\ell_2}, \text{Lin}}$, then if w which is any $(r^* - 1)$ -fold commutator of $v_{i_1}, \dots, v_{i_{r^*}}$ we have

$$\eta(w) = 0.$$

We have relabeled as H' by H , g'_h by g_h , $g'(h, n)$ by $g(h, n)$, χ' by χ , and χ'_h by χ_h . We have applied Lemma 10.1 and scaling to guarantee that $\exp(Z_{i,j}^*), \exp(Z_{i,j}^{\text{Lin}}), \exp(Z_{i,j}^{\text{Pet}}) \in \Gamma$.

Let $\mathcal{X} = \{X_1, \dots, X_{\dim(G)}\}$ denote the filtered Mal'cev basis given for G/Γ . Via Lemma 10.2, for $h \in H$ we may define

$$\begin{aligned} g_h^* &= \prod_{i=1}^{s-1} \prod_{j=1}^{\dim(W_{i,*})} \exp(Z_{i,j}^*)^{z_{i,j}^* \cdot \frac{n^i}{i!}} \cdot \prod_{i=1}^{s-1} \prod_{j=1}^{\dim(W_{i,\text{Lin}})} \exp(Z_{i,j}^{\text{Lin}})^{\gamma_{i,j} \cdot \frac{n^i}{i!}}, \\ g_h^{\text{Lin}} &= \prod_{i=1}^{s-1} \prod_{j=1}^{\dim(W_{i,\text{Lin}})} \exp(Z_{i,j}^{\text{Lin}})^{(z_{i,j}^{h,\text{Lin}} - \gamma_{i,j}) \cdot \frac{n^i}{i!}}, \\ g_h^{\text{Pet}} &= \prod_{i=1}^{s-1} \prod_{j=1}^{\dim(W_{i,\text{Pet}})} \exp(Z_{i,j}^{\text{Pet}})^{z_{i,j}^{h,\text{Pet}} \cdot \frac{n^i}{i!}}, \end{aligned}$$

and define g_h^{Rem} via

$$g_h = g_h^* \cdot g_h^{\text{Lin}} \cdot g_h^{\text{Pet}} \cdot g_h^{\text{Rem}}.$$

Using Lemma 10.2 again, we may write

$$g_h^{\text{Rem}} = \prod_{i=1}^{s-1} \prod_{j=\dim(G)-\dim(G_{i,1})+1}^{\dim(G)} \exp(X_j)^{\kappa_{i,j}^h \cdot \frac{n^i}{i!}}.$$

The fact that when applying Lemma 10.2 for g_h^{Rem} we observe no coefficients for $\frac{n^i}{i!}$ corresponding to basis elements in $\mathcal{X} \cap \log(G_{(i,1)}) \setminus \mathcal{X} \cap \log(G_{(i,2)})$ follows from the fact that g_h and $g_h^* \cdot g_h^{\text{Lin}} \cdot g_h^{\text{Pet}}$ have Taylor coefficients which match exactly for $1 \leq i \leq s-1$.

We now reach the first stage of “rewriting” where we realize the nilsequence $\chi_h(n) = F(g_h(n)\Gamma)$ on a universal nilmanifold.

10.2. Rewriting degree-rank nilsequences on the universal nilmanifold. We recall the universal nilmanifold of a given degree-rank (see [29, Definition 9.1]).

Definition 10.3. *The universal nilmanifold of degree-rank $(s-1, r^*)$ and the associated discrete cocompact subgroup are defined as follows. We write $G_{\text{Univ}} = G_{\text{Univ}}^{\vec{D}}$ where $\vec{D} = \vec{D}^* + \vec{D}^{\text{Lin}} + \vec{D}^{\text{Pet}}$ with $\vec{D}^*, \vec{D}^{\text{Lin}}, \vec{D}^{\text{Pet}} \in (\mathbb{Z}_{\geq 0})^{s-1}$. We specify $G_{\text{Univ}}^{\vec{D}}$ by formal generators of the Lie algebra $e_{i,j}$ for $1 \leq i \leq s-1$ and $1 \leq j \leq D_i$ where $D_i = D_i^* + D_i^{\text{Lin}} + D_i^{\text{Pet}}$ with the relations:*

- Any $(r-1)$ -fold commutator of $e_{i_1, j_1}, \dots, e_{i_r, j_r}$ with $i_1 + \dots + i_r > (s-1)$ vanishes;
- Any $(r-1)$ -fold commutator of $e_{i_1, j_1}, \dots, e_{i_r, j_r}$ with $i_1 + \dots + i_r = (s-1)$ and $r > r^*$ vanishes.

The associated discrete group which we will be concerned with is Γ_{Univ} which is the discrete group generated by $\exp(e_{i,j})$ for $1 \leq i \leq s-1$ and $1 \leq j \leq D_i$.

Remark. Note that in this definition, $G_{\text{Univ}}^{\vec{D}}$ depends only on \vec{D} ; however, the quotient we will consider later depends on $\vec{D}^*, \vec{D}^{\text{Lin}}, \vec{D}^{\text{Pet}}$. Furthermore, we have presented G_{Univ} as a Lie algebra and not as a Lie group; via the general theory of nilpotent Lie algebras this is sufficient. Note that the Lie algebra defined is trivially seen to be nilpotent. By the Birkhoff Embedding Theorem (see remark following [7, Theorem 1.1.11]), we may realize any real nilpotent Lie algebra \mathfrak{g} as a Lie subalgebra of the $n \times n$ real strictly upper triangular matrices. The proof of [7, Theorem 1.2.1] then realizes the $n \times n$ real strictly upper triangular matrices as a logarithm of a connected, simply connected Lie group N_n where the exponential map is bijective. The Baker–Campbell–Hausdorff formula then demonstrates \mathfrak{g} is the logarithm of a connected, simply connected subgroup $G \leq N_n$ (and by construction the logarithm is a bijection between G and \mathfrak{g}). The group G constructed is unique up to isomorphism by Lie’s third theorem.

We first prove the fact that G_{Univ} may be given a degree-rank filtration and that $G_{\text{Univ}}/\Gamma_{\text{Univ}}$ has reasonable complexity.

Lemma 10.4. *Let $G_{\text{Univ}} = G_{\text{Univ}}^{\vec{D}}$ and define $(G_{\text{Univ}})_{(d,r)}$ by taking the group generated by all $(r'-1)$ -fold iterated commutators of $\exp(t_{i_1, j_1} e_{i_1, j_1}), \dots, \exp(t_{i_{r'}, j_{r'}} e_{i_{r'}, j_{r'}})$ with $t_{i_k, j_k} \in \mathbb{R}$, and either $i_1 + \dots + i_{r'} > d$ or $i_1 + \dots + i_{r'} = d$ and $r' \geq r$.*

Then $(G_{\text{Univ}})_{(d,r)}$ forms a valid degree-rank $(s-1, r^)$ filtration of G_{Univ} . Furthermore the dimension of G_{Univ} is bounded by $O_s(\|D\|_{\infty}^{O_s(1)})$ and one may find an adapted Mal’cev basis $\mathcal{X}_{\text{Univ}}$ such that the complexity of $G_{\text{Univ}}/\Gamma_{\text{Univ}}$ is at most $\exp(\|D\|_{\infty}^{O_s(1)})$.*

Proof. We will be brief with details; that the associated filtration is valid follows via a straightforward computation with Lemma 2.2. Note $(G_{\text{Univ}})_{(i,0)} = (G_{\text{Univ}})_{(i,1)}$ as $r' \geq 1$ in the set of generators always. Also, $(G_{\text{Univ}})_{(0,0)} = (G_{\text{Univ}})_{(0,1)}$ since for all generators $e_{i,j}$ we have $i \geq 1$.

To establish the complexity bounds, the key point is noting that taking all $(r'-1)$ -fold iterated commutators of $e_{i_1, j_1}, \dots, e_{i_{r'}, j_{r'}}$ with $i_1 + \dots + i_{r'} \leq s-2$ or $i_1 + \dots + i_{r'} = s-1$ and $r' \leq r^*$

gives a spanning set for $\log(G_{\text{Univ}})$. This immediately gives the specified dimension bound. These generators are not linearly independent; however, all relations are generated by either antisymmetry ($[x, y] + [y, x] = 0$) or the Jacobi identity ($[x, [y, z]] + [y, [z, x]] + [z, [x, y]] = 0$) applied to the set of generators specified.

To simplify matters, note that all linear relations can be reduced to those between these generators with the “same type” (i.e., relations between the set of $(r' - 1)$ -fold commutators of a given set of generators $e_{i_1, j_1}, \dots, e_{i_{r'}, j_{r'}}$). These can be collected into disconnected non-interacting “components” which are $O_s(1)$ in size. We may take a linearly spanning set within each group; each generator not in the spanning set may be written as a linear combination of height $O_s(1)$. Define \mathcal{X} to be the union of all these spanning elements in $\log(G_{\text{Univ}})$. This gives us a basis. Note the subspaces $\log((G_{\text{Univ}})_{(d,r)})$ are clearly compatible with natural subsets of these “components” and their associated spanning sets, demonstrating that the basis is appropriate adapted to these vector spaces $\log((G_{\text{Univ}})_{(d,r)})$.

The last matter to check is that there exists $C_s \geq 1$ such that $C_s \mathbb{Z}^{\dim(G_{\text{Univ}})} \subseteq \psi_{\exp, \mathcal{X}}(\Gamma_{\text{Univ}}) \subseteq C_s^{-1} \mathbb{Z}^{\dim(G_{\text{Univ}})}$. This follows by noting that each element $\gamma \in \Gamma$ may be written as

$$\gamma = \prod_{k=1}^t \exp(e_{i_k, j_k})^{s_k}$$

with $s_k \in \mathbb{Z}$. We prove the first implication first; we prove that $\log(\gamma)$ may be written as a linear combination of iterated commutators where $(r' - 1)$ -fold commutators have denominator bounded by $C_s^{r'}$. This is trivial to prove inductively via Baker–Campbell–Hausdorff and noting that all s -fold commutators vanish.

For the reverse direction, consider expressions of the form

$$\gamma' = \exp\left(\sum_{\alpha} c_{\alpha} e_{\alpha}\right)$$

where e_{α} ranges over all possible iterated commutators (here e.g. $e_{[(1,2),(1,3)]} := [e_{(1,2)}, e_{(1,3)}]$) where c_{α} are sufficiently divisible integers. Let f_{α} be defined as the commutator of the exponential of associated elements; e.g. $f_{[(1,2),(1,3)]} = [\exp(e_{(1,2)}), \exp(e_{(1,3)})]$. Choose a generator α' with the fewest number of commutators in γ' such that $c_{\alpha'} \neq 0$. It is straightforward to see via Baker–Campbell–Hausdorff that there is an integer M_s such that if c_{α} are all divisible by M_s then

$$f_{\alpha'}^{-c_{\alpha'}} \gamma' = \exp\left(\sum_{\alpha} c_{\alpha}^* e_{\alpha}\right)$$

has each c_{α}^* still divisible by M_s and $c_{\alpha'}^* = 0$ (without introducing backwards corrections).

The desired result then follows from [35, Lemma B.11], noting that $(G_{\text{Univ}})_{(d,r)}$ is the degree-rank ordering forming a nested sequence of subgroups. \square

We now represent the nilsequences $\chi_h(n) = F(g_h(n))$ on the universal nilmanifold. We define

$$D_i^* = \dim(W_{i,*}) + \dim(W_{i,\text{Lin}}), \quad D_i^{\text{Pet}} = \dim(W_{i,\text{Pet}}) + \dim(G_{(i,2)}), \quad D_i^{\text{Lin}} = d^* \dim(W_{i,\text{Lin}}).$$

Note that $D_i^{\text{Lin}} \leq (d \log(MD/\rho))^{O_s(1)}$ and trivially $D_i^*, D_i^{\text{Lin}} \leq d$.

Recall that $\mathcal{X} = \{X_1, \dots, X_{\dim(G)}\}$ is the filtered Mal'cev basis and $Z_{i,j}^*$, $Z_{i,j}^{\text{Pet}}$, $Z_{i,j}^{\text{Lin}}$ are representative of $\log(W_{i,*})$, $\log(W_{i,\text{Lin}})$, and $\log(W_{i,\text{Pet}})$ respectively.

We define a homomorphism $\phi: G_{\text{Univ}} \rightarrow G$ by defining the map on generators. Define

$$\phi(\exp(e_{i,j})) = \begin{cases} \exp(Z_{i,j}^*) & \text{if } 1 \leq j \leq \dim(W_{i,*}), \\ \exp(Z_{i,j-\dim(W_{i,*})}^{\text{Lin}}) & \text{if } \dim(W_{i,*}) + 1 \leq j \leq \dim(W_{i,*}) + \dim(W_{i,\text{Lin}}) = D_i^*, \\ \exp(Z_{i,\ell}^{\text{Lin}}) & \text{if } 1 + (\ell - 1)d^* \leq j - D_i^* \leq \ell d^* \text{ for } 1 \leq \ell \leq \dim(W_{i,\text{Lin}}), \\ \exp(Z_{i,j-D_i^*-D_i^{\text{Lin}}}^{\text{Pet}}) & \text{if } D_i^* + D_i^{\text{Lin}} + 1 \leq j \leq D_i^* + D_i^{\text{Lin}} + \dim(W_{i,\text{Pet}}), \\ \exp(X_{j-D_i+\dim(G)}) & \text{if } D_i - \dim(G_{(i,2)}) + 1 \leq j \leq D_i. \end{cases}$$

That this is a homomorphism is an immediate consequence of the fact that the only relations on the universal nilmanifold are forced on the group G since it has degree-rank $(s-1, r^*)$.

The function with which will be concerned is

$$\tilde{F}(g\Gamma_{\text{Univ}}) := F(\phi(g)\Gamma).$$

This is well-defined since $\phi(\Gamma_{\text{Univ}}) \leq \Gamma$; it suffices to check that the generators $\exp(e_{i,j})$ map to within Γ but this is trivial by construction. (This is precisely why we scaled $Z_{i,\cdot}^*$, $Z_{i,\cdot}^{\text{Lin}}$, and $Z_{i,\cdot}^{\text{Pet}}$ so that when exponentiated they live within Γ .)

We now note a series of basic properties of \tilde{F} and the homomorphism ϕ .

Lemma 10.5. *Given the above setup we have:*

- $\|\tilde{F}\|_2 = 1$ for all $g \in G_{\text{Univ}}$;
- F has a vertical frequency η_{Univ} with height at most $(MD/\rho)^{O_s(\dim(G_{\text{Univ}})^{O_s(1)})}$;
- \tilde{F} is $(MD/\rho)^{O_s(\dim(G_{\text{Univ}})^{O_s(1)})}$ -Lipschitz
- Consider $e_{i_1,j_1}, \dots, e_{i_{r^*},j_{r^*}}$ with $j_1 + \dots + j_{r^*} = s-1$. If for at least one index ℓ we have $j_\ell > D_{i_\ell}^* + D_{i_\ell}^{\text{Lin}}$, then

$$\eta_{\text{Univ}}([\exp(e_{i_1,j_1}), \dots, \exp(e_{i_{r^*},j_{r^*}})]) = 0.$$

Furthermore, if instead for two indices ℓ_1, ℓ_2 we have $j_{\ell_1} > D_{i_{\ell_1}}^*$ and $j_{\ell_2} > D_{i_{\ell_2}}^*$ then

$$\eta_{\text{Univ}}([\exp(e_{i_1,j_1}), \dots, \exp(e_{i_{r^*},j_{r^*}})]) = 0.$$

Proof. The first property is trivial. For the second property, note that ϕ is an I -filtered homomorphism (e.g. $\phi((G_{\text{Univ}})_{(s,r^*-1)}) \leq G_{(s,r^*-1)}$). Thus given $g \in G_{\text{Univ}}$, $g' \in (G_{\text{Univ}})_{(s,r^*-1)}$ we have

$$\tilde{F}(gg'\Gamma_{\text{Univ}}) = F(\phi(g)\phi(g')\Gamma) = e(\eta(\phi(g')))F(\phi(g)\Gamma)$$

and therefore we may set $\eta_{\text{Univ}} = \eta \circ \phi$. To check the complexity of η_{Univ} it suffices to check the magnitude of η_{Univ} on $[\exp(e_{i_1,j_1}), \dots, \exp(e_{i_{r^*},j_{r^*}})]$ where we use Remark 5.3 to convert between this notion and the notion of height defined. The resulting magnitude is bounded because $Z_{i,j}^*$, $Z_{i,j}^{\text{Lin}}$, $Z_{i,j}^{\text{Pet}}$ are appropriately bounded integral combinations of elements in \mathcal{X} which itself has bounded complexity.

We omit a careful justification that \tilde{F} has an appropriately bounded Lipschitz constant. The crucial point is that the Mal'cev basis constructed in Lemma 10.4 is made up of appropriately bounded linear combinations of commutators of $e_{i_1,j_1}, \dots, e_{i_{r^*},j_{r^*}}$ and each such commutator is seen to map to a bounded element of G since ϕ maps each generator to a bounded element.

The final property is an immediate consequence of the properties of $W_{i,\text{Lin}}$, $W_{i,\text{Pet}}$, and $W_{i,*}$ established in Lemma 9.1 and recorded above. The additional generators which are lifted to the ‘‘petal’’ position on the i -th level come from $G_{(i,2)}$ and otherwise we have only artificially placed certain elements in the ‘‘linear’’ class upward to the ‘‘*’’ class. (These will correspond to the constant terms in the linear part of the nilsequences.) \square

We now lift the polynomial sequences in question to the universal nilmanifold. We define:

$$\begin{aligned}
g_h^{*,\text{Univ}}(n) &= \prod_{i=1}^{s-1} \prod_{j=1}^{\dim(W_{i,*})} \exp(e_{i,j}) z_{i,j}^* \cdot \frac{n^i}{i!} \prod_{i=1}^{s-1} \prod_{j=1}^{\dim(W_{i,\text{Lin}})} \exp(e_{i,j+\dim(W_{i,*})}) \gamma_{i,j} \cdot \frac{n^i}{i!}, \\
g_h^{\text{Lin},\text{Univ}}(n) &= \prod_{i=1}^{s-1} \prod_{j=1}^{\dim(W_{i,\text{Lin}})} \prod_{k=1}^{d^*} \exp(e_{i,D_i^*+(j-1)d^*+k}) \alpha_{i,j,k} \{\beta_k h\} \cdot \frac{n^i}{i!}, \\
g_h^{\text{Pet},\text{Univ}}(n) &= \prod_{i=1}^{s-1} \prod_{j=1}^{\dim(W_{i,\text{Pet}})} \exp(e_{i,j+D_i^*+D_i^{\text{Lin}}}) z_{i,j}^{h,\text{Pet}} \cdot \frac{n^i}{i!}, \\
g_h^{\text{Rem},\text{Univ}}(n) &= \prod_{i=1}^{s-1} \prod_{j=1}^{\dim(G_{(i,2)})} \exp(e_{i,j+D_i-\dim(G_{(i,2)})}) \kappa_{i,j}^h \cdot \frac{n^i}{i!}.
\end{aligned}$$

We define

$$g_h^{\text{Univ}} := g_h^{*,\text{Univ}} \cdot g_h^{\text{Lin},\text{Univ}} \cdot g_h^{\text{Pet},\text{Univ}} \cdot g_h^{\text{Rem},\text{Univ}}.$$

The key claim, which is trivial by construction, is the following equality.

Claim 10.6. *Given the above setup, we have*

$$\tilde{F}(g_h^{\text{Univ}}(n)\Gamma_{\text{Univ}}) = F(g_h(n)\Gamma) = \chi_h(n).$$

Proof. The final equality is by definition of $\chi_h(n)$. The first equality follows by checking that $\phi(g_h^{*,\text{Univ}}) = g_h^*$, $\phi(g_h^{\text{Lin},\text{Univ}}) = g_h^{\text{Lin}}$, $\phi(g_h^{\text{Pet},\text{Univ}}) = g_h^{\text{Pet}}$, and $\phi(g_h^{\text{Rem},\text{Univ}}) = g_h^{\text{Rem}}$ by construction. Therefore since ϕ is a homomorphism we conclude that $\phi(g_h^{\text{Univ}}) = g_h$. \square

Note that at this stage we have simply replace the group G in our correlation structure with G^{Univ} as the cost of replacing d by $\dim(G_{\text{Univ}}) = d^{O_s(1)} \log(MD\rho^{-1})^{O_s(1)}$ and M by $\exp(d^{O_s(1)} \log(MD\rho^{-1})^{O_s(1)})$.

This may seem as if we have gone backwards, the key point is that in Lemma 10.5 we have encoded various “vanishing conditions” on the commutator brackets at the level of the generators of the group. This will allow us to translate the “vanishing conditions” obtained in Lemma 9.1 into realizing we can, up to a degree-rank $(s-1, r^*-1)$ -error term.

10.3. Passing to a quotient nilmanifold. We now construct two additional nilmanifolds; there are essentially G^* and \tilde{G} certain quotients constructed in [29, Section 12].⁵

Definition 10.7. *We define $G_{\text{Rel}} = G_{\text{Rel}}^{\vec{D}^*, \vec{D}^{\text{Lin}}, \vec{D}^{\text{Pet}}}$ as the Lie subgroup of G_{Univ} where $\log(G_{\text{Rel}})$ is spanned by:*

- Any $(r-1)$ -fold commutator of $e_{i_1, j_1}, \dots, e_{i_r, j_r}$ with at least one index ℓ such that $j_\ell > D_{i_\ell}^* + D_{i_\ell}^{\text{Lin}}$;
- Any $(r-1)$ -fold commutator of $e_{i_1, j_1}, \dots, e_{i_r, j_r}$ with $j_\ell > D_{i_\ell}^*$ for at least two distinct indices ℓ .

We then define G_{Quot} as $G_{\text{Quot}} := G_{\text{Univ}}/G_{\text{Rel}}$ and $\Gamma_{\text{Quot}} = \Gamma_{\text{Univ}}/(\Gamma_{\text{Univ}} \cap G_{\text{Rel}})$.

Remark 10.8. Note that we may set $r = 1$ in the definition of G_{Rel} ; in particular $\exp(e_{i,j}) \in G_{\text{Rel}}$ for $j > D_i^* + D_i^{\text{Lin}}$. Additionally, $\log(G_{\text{Quot}})$ may be realized as the following. Consider formal generators of a Lie algebra, $\tilde{e}_{i,j}$ for $1 \leq j \leq D_i^* + D_i^{\text{Lin}}$, with the property that:

- Any $(r-1)$ -fold commutator of $\tilde{e}_{i_1, j_1}, \dots, \tilde{e}_{i_r, j_r}$ with either $i_1 + \dots + i_r > (s-1)$ or $i_1 + \dots + i_r = (s-1)$ and $r > r^*$ vanishes;

⁵There is a minor issue in [29, p. 1309] when defining G^* ; we follow the definitions given in the erratum [26].

- Any $(r-1)$ -fold commutator of $\tilde{e}_{i_1, j_1}, \dots, \tilde{e}_{i_r^*, j_r^*}$ with $j_\ell > D_{i_\ell}^*$ for at least two distinct indices ℓ vanishes.

This realization is given by taking $\tilde{e}_{i,j} := \log(\exp(e_{i,j}) \bmod G_{\text{Rel}})$.

We first check that G_{Quot} is well-defined.

Claim 10.9. For $\vec{D}^*, \vec{D}^{\text{Lin}}, \vec{D}^{\text{Pet}} \in (\mathbb{Z}_{\geq 0})^{s-1}$, G_{Rel} is a well-defined normal subgroup of G_{Univ} .

Proof. It is clear from definition that $\log(G_{\text{Rel}})$ is closed under brackets, so forms a Lie subalgebra within $\log(G_{\text{Univ}})$. Thus G_{Rel} is indeed a Lie subgroup. To prove that G_{Rel} is normal it suffices to prove that it is furthermore a Lie algebra ideal, i.e., $[\log(G_{\text{Univ}}), \log(G_{\text{Rel}})] \leq \log(G_{\text{Rel}})$.

Recall that $\log(G_{\text{Univ}})$ is spanned by all the $(r-1)$ -fold commutators $e_{i_1, j_1}, \dots, e_{i_r, j_r}$ (although as discussed in Lemma 10.4 this is not a basis). It suffices to check the containment at the level of generators of the respective Lie algebras. The result then follows since taking a commutator does not decrease the number of ‘‘petal’’ or ‘‘linear’’ generators. \square

We also have the following complexity bound on G_{Quot} . This may be done via the Lie algebra presentation given in Remark 10.8 and repeating the proof in Lemma 10.4, or via noting that G_{Rel} is a sufficient rational subgroup of G_{Univ} and applying Lemma 3.10. We omit the details.

Lemma 10.10. Given the above setup, let $G_{\text{Quot}} = G_{\text{Quot}}^{\vec{D}}$ and note that G_{Quot} has a degree-rank $(s-1, r^*)$ filtration given by

$$(G_{\text{Quot}})_{(d,r)} = (G_{\text{Univ}})_{(d,r)} / ((G_{\text{Univ}})_{(d,r)} \cap G_{\text{Rel}}).$$

Furthermore the dimension of G_{Univ} is bounded by $O_s(\|D\|_\infty^{O_s(1)})$ and one may find an adapted Mal'cev basis $\mathcal{X}_{\text{Quot}}$ such that the complexity of $G_{\text{Quot}}/\Gamma_{\text{Quot}}$ is $\exp(\|D\|_\infty^{O_s(1)})$.

A key point in this analysis is that this quotient is compatible with η_{Univ} .

Lemma 10.11. Given the above setup, define $\eta_{\text{Quot}} : (G_{\text{Quot}})_{(s-1, r^*)} \rightarrow \mathbb{R}$ via

$$\eta_{\text{Quot}}(g \bmod G_{\text{Rel}}) := \eta_{\text{Univ}}(g)$$

for all $g \in (G_{\text{Univ}})_{(s-1, r^*)}$. The map η_{Quot} is well-defined and in fact is a vertical character of G_{Quot} of height at most $(MD/\rho)^{O_s(\dim(G_{\text{Univ}})^{O_s(1)})}$.

Proof. To be well-defined as a map, it suffices to show that $G_{\text{Rel}} \cap (G_{\text{Univ}})_{(s-1, r^*)} \leq \ker(\eta_{\text{Univ}})$. This comes exactly from the final item of Lemma 10.5. That η is a vertical character then follows as $\Gamma_{\text{Quot}} = \Gamma_{\text{Univ}} / (\Gamma_{\text{Univ}} \cap G_{\text{Rel}})$.

To bound the height of η_{Quot} note that taking a quotient by G_{Rel} maps $\exp(e_{i,j})$ to $\exp(\tilde{e}_{i,j})$ in the sense of Remark 10.8. Furthermore the construction of $\mathcal{X}_{\text{Quot}}$ has the property that $\mathcal{X}_{\text{Quot}} \cap \log((G_{\text{Quot}})_{(s-1, r^*)})$ are sufficiently rational combinations of (r^*-1) -fold commutators of $\tilde{e}_{i_1, j_1}, \dots, \tilde{e}_{i_r^*, j_r^*}$ with $i_1 + \dots + i_r^* = s-1$. By Baker–Campbell–Hausdorff, we have that the (r^*-1) -fold commutator of $\exp(\tilde{e}_{i_1, j_1}), \dots, \exp(\tilde{e}_{i_r^*, j_r^*})$ is the same mod G_{Rel} as the corresponding one for $\exp(e_{i_1, j_1}), \dots, \exp(e_{i_r^*, j_r^*})$. However, η_{Univ} maps the latter commutator to a sufficiently bounded integer by the complexity bound on η_{Univ} and the result follows. \square

We will require

$$(10.1) \quad \begin{aligned} g_h^{*, \text{Quot}}(n) &= \prod_{i=1}^{s-1} \prod_{j=1}^{\dim(W_{i,*})} \exp(\tilde{e}_{i,j})^{z_{i,j}^* \cdot \frac{n^i}{i!}} \prod_{i=1}^{s-1} \prod_{j=1}^{\dim(W_{i,\text{Lin}})} \exp(\tilde{e}_{i,j+\dim(W_{i,*})})^{\gamma_{i,j} \cdot \frac{n^i}{i!}} \\ g_h^{\text{Lin}, \text{Quot}}(n) &= \prod_{i=1}^{s-1} \prod_{j=1}^{\dim(W_{i,\text{Lin}})} \prod_{k=1}^{d^*} \exp(\tilde{e}_{i, D_i^* + (j-1)d^* + k})^{\alpha_{i,j,k} \{\beta_k h\} \cdot \frac{n^i}{i!}}; \end{aligned}$$

note that

$$g_h^{*,\text{Quot}} = g_h^{*,\text{Univ}} \bmod G_{\text{Rel}}, \quad g_h^{\text{Lin},\text{Quot}} = g_h^{\text{Lin},\text{Univ}} \bmod G_{\text{Rel}}.$$

Furthermore we have

$$g_h^{\text{Pet},\text{Univ}} \bmod G_{\text{Rel}} = g_h^{\text{Rem},\text{Univ}} \bmod G_{\text{Rel}} = \text{id}_{G_{\text{Quot}}}$$

pointwise. Finally we define

$$g_h^{\text{Quot}} := g_h^{*,\text{Quot}} \cdot g_h^{\text{Lin},\text{Quot}}.$$

For the remainder of this section and Section 11, fix a nilcharacter F^* on G_{Quot} with a $G_{(s-1,r^*)}$ -vertical frequency η_{Quot} . Furthermore by Lemma B.4⁶ we may take F^* which is $(MD/\rho)^{O_s(\dim(G_{\text{Univ}})^{O_s(1)})}$ -Lipschitz with output dimension bounded by $2^{O_s(\dim(G_{\text{Univ}}))}$.

The reason it will be sufficient to study $F^*(g_h^{\text{Quot}}\Gamma_{\text{Quot}})$ will be the following lemma which proves that it is equal to $\tilde{F}(g_h^{\text{Univ}}\Gamma_{\text{Univ}})$ up a term which is lower-order in degree-rank.

Lemma 10.12. *Given the above setup, let*

$$G_{\text{Univ}}^{\Delta} := \{(g, g \bmod G_{\text{Rel}}) \in G_{\text{Univ}} \times G_{\text{Quot}} : g \in G_{\text{Univ}}\}$$

which is given the degree-rank filtration

$$(G_{\text{Univ}}^{\Delta})_{(d,r)} := \{(g, g \bmod G_{\text{Rel}}) \in (G_{\text{Univ}})_{(d,r)} \times (G_{\text{Quot}})_{(d,r)} : g \in (G_{\text{Univ}})_{(d,r)}\}.$$

Define $\Gamma_{\text{Univ}}^{\Delta} = G_{\text{Univ}}^{\Delta} \cap (\Gamma_{\text{Univ}} \times \Gamma_{\text{Quot}})$. We have:

- $(g_h^{\text{Univ}}, g_h^{\text{Quot}})$ is a polynomial sequence on G_{Univ}^{Δ} with respect to the given degree-rank filtration;
- The function

$$(g, g') \mapsto \tilde{F}(g\Gamma_{\text{Univ}}) \otimes \overline{F^*}(g'\Gamma_{\text{Quot}})$$

for $(g, g') \in G_{\text{Univ}}^{\Delta}$ is $(G_{\text{Univ}}^{\Delta})_{(s-1,r^*)}$ -invariant;

- G_{Univ}^{Δ} has complexity bounded by $(MD/\rho)^{O_s(\dim(G_{\text{Univ}})^{O_s(1)})}$;
- Each coordinate of $\tilde{F}(g\Gamma_{\text{Univ}}) \otimes \overline{F^*}(g'\Gamma_{\text{Quot}})$ is $(MD/\rho)^{O_s(\dim(G_{\text{Univ}})^{O_s(1)})}$ -Lipschitz.

Remark 10.13. The second item implies that $\tilde{F}(g\Gamma_{\text{Univ}}) \otimes \overline{F^*}(g'\Gamma_{\text{Quot}})$ is $(G_{\text{Univ}}^{\Delta})_{(s-1,r^*)}$ -invariant and thus can be realized on a degree-rank $(s-1, r^*-1)$ nilmanifold $G_{\text{Univ}}^{\Delta}/(G_{\text{Univ}}^{\Delta})_{(s-1,r^*)}$ with $\Gamma_{\text{Univ}}^{\Delta}/(\Gamma_{\text{Univ}}^{\Delta} \cap (G_{\text{Univ}}^{\Delta})_{(s-1,r^*)})$ being the lattice.

Proof. It is trivial to verify that the degree-rank filtration on G_{Univ}^{Δ} is valid. Noting that

$$\{(X_i, X_i \bmod \log(G_{\text{Rel}})) : X_i \in \mathcal{X}_{\text{Univ}}\}$$

is a valid Mal'cev basis for G_{Univ}^{Δ} bounds the complexity of G_{Univ}^{Δ} . The complexity bounds on $\tilde{F}(g\Gamma_{\text{Univ}}) \otimes \overline{F^*}(g'\Gamma_{\text{Quot}})$ follow by noting that \tilde{F} is appropriately Lipschitz on $G_{\text{Univ}}/\Gamma_{\text{Univ}}$ and similar for $\overline{F^*}$. For F^* , we note that each coordinate of $\{X_i \bmod \log(G_{\text{Rel}})\}$ is appropriately rational with respect to the Mal'cev basis for $\mathcal{X}_{\text{Quot}}$, by construction.

Furthermore for $(h, h \bmod G_{\text{Rel}}) \in (G_{\text{Univ}}^{\Delta})_{(s-1,r^*)}$ we have

$$\begin{aligned} & \tilde{F}(gh\Gamma_{\text{Univ}}) \otimes \overline{F^*}(g'(h \bmod G_{\text{Rel}})\Gamma_{\text{Quot}}) \\ &= \tilde{F}(g\Gamma_{\text{Univ}}) \otimes \overline{F^*}(g'\Gamma_{\text{Quot}}) \cdot e(\eta_{\text{Univ}}(h))\overline{e(\eta_{\text{Quot}}(h \bmod G_{\text{Rel}}))} \\ &= \tilde{F}(g\Gamma_{\text{Univ}}) \otimes \overline{F^*}(g'\Gamma_{\text{Quot}}) \end{aligned}$$

⁶The lemma is stated for degree filtrations. However, one can give G_{Quot} the degree filtration

$$(G_{\text{Quot}})_{(0,0)} = (G_{\text{Quot}})_{(1,0)} \geq (G_{\text{Quot}})_{(2,0)} \geq \cdots \geq (G_{\text{Quot}})_{(s-1,0)} \geq (G_{\text{Quot}})_{(s-1,r^*)} \geq \text{Id}_{G_{\text{Quot}}};$$

a vertical nilcharacter with respect to this filtration is a vertical nilcharacter with respect to the original degree-rank filtration. $\mathcal{X}_{\text{Quot}}$ is adapted to this degree-filtration (as it is adapted to the original degree-rank filtration).

where in the final line we have used the definition of η_{Quot} .

Finally to verify that $(g_h^{\text{Univ}}, g_h^{\text{Quot}})$ is a polynomial sequence with respect to this degree-rank filtration, note via Taylor expansion (e.g. [29, Lemma B.9]) that all polynomial sequences h with respect to G_{Univ}^{Δ} of the form $(h', h' \bmod G_{\text{Rel}})$ where h' is a polynomial sequence with respect to G_{Univ} (and its specified degree-rank filtration). The result then follows due to the property

$$g_h^{\text{Univ}} \bmod G_{\text{Rel}} = g_h^{\text{Quot}}$$

noted above, which was by construction. \square

11. EXTRACTING A $(1, s-1)$ -NILSEQUENCE

The goal of this section is to realize

$$F^*(g_h^{\text{Quot}}(n)\Gamma_{\text{Quot}})$$

as a multidegree $(1, s-1)$ nilsequence in (h, n) . We accomplish this via a construction of Green, Tao, and Ziegler [29, Section 12] and then use this construction in order to complete the proof of Lemma 6.3. After this, the main business of the paper is essentially done and all that remains to prove Theorem 1.2 is the symmetrization argument which will be carried out in the next section.

11.1. Constructing the $(1, s-1)$ -nilsequence. Our analysis at this point is essentially verbatim that of [29, pp. 1313-1315]. We reproduce the details here (and discuss various complexity issues which are completely routine in the appendix). For the sake of simplicity, we may clean up notation from (10.1) and write

$$g_h^{\text{Quot}}(n) = \prod_{i=1}^{s-1} \prod_{j=1}^{D_i^*} \exp(\tilde{e}_{i,j})^{\gamma_{i,j} \cdot \frac{n^i}{i!}} \cdot \prod_{i=1}^{s-1} \prod_{j=D_i^*+1}^{D_i^*+D_i^{\text{Lin}}} \exp(\tilde{e}_{i,j})^{\alpha_{i,j} \{\beta_{i,j} h\} \cdot \frac{n^i}{i!}},$$

where we have abusively reindexed various coefficients γ, α, β but nothing else.

We now define G_{Lin} to be the Lie subgroup of G_{Quot} such that $\log(G_{\text{Lin}})$ is the subspace generated by all $(r-1)$ -fold iterated commutators (with $r \geq 1$) of $\tilde{e}_{i_1, j_1}, \dots, \tilde{e}_{i_r, j_r}$ with $j_\ell > D_{i_\ell}^*$ for exactly one index ℓ . We have the following pair of basic observations.

Claim 11.1. *We have that G_{Lin} is well-defined, abelian, and normal with respect to G_{Quot} .*

Proof. Similar to the proof of Claim 10.9, G_{Lin} is well-defined and normal. The only modification to the proof is noting that a commutator of \tilde{e}_{i_k, j_k} with at least two indices ℓ with $j_\ell > D_{i_\ell}^*$ vanishes by the definition of G_{Quot} .

To see that G_{Lin} is abelian, it suffices to prove that the commutator of any pair of generators is the identity. This immediately follows from the fact that commutators with at least two generators of the form $\tilde{e}_{i_\ell, j_\ell}$ with $j_\ell > D_{i_\ell}^*$ vanish. \square

Due to normality, G_{Quot} acts on G_{Lin} via conjugation. In particular, we define $G_{\text{Quot}} \ltimes G_{\text{Lin}}$ with the group law given by

$$(g, g_1)(g', g'_1) := (gg', g'_1 g'_1) = (gg', ((g')^{-1} g_1 g') g'_1).$$

We now introduce a manner in which the additive group $R = \mathbb{R}^{\sum_{i=1}^{s-1} D_i^{\text{Lin}}}$, with elements denoted

$$t = (t_{i,j})_{1 \leq i \leq s-1, D_{i,*} < j \leq D_i + D_i^{\text{Lin}}},$$

acts on $G_{\text{Quot}} \ltimes G_{\text{Lin}}$. Specifically, we will define an action $\rho(t)$ on this group for all $t \in R$ and use this to construct

$$G_{\text{Multi}} = R \ltimes_{\rho} (G_{\text{Quot}} \ltimes G_{\text{Lin}}).$$

This action will allow us to simultaneously “raise” parts of G_{Lin} to various different fractional powers of h , allowing us to incorporate our “ h -linear” family of nilsequences into a multidegree $(1, s - 1)$ nilsequence (in variables (h, n)).

For each $t \in R$, we define the homomorphism $g \mapsto g^t$ from G_{Quot} to itself on generators. We map $\exp(\tilde{e}_{i,j}) \rightarrow \exp(\tilde{e}_{i,j})^{t_{i,j}}$ for $1 \leq i \leq s - 1$ and $D_i^* < j \leq D_i^* + D_i^{\text{Lin}}$ while $\exp(\tilde{e}_{i,j})$ is fixed for $1 \leq i \leq s - 1$ and $1 \leq j \leq D_i^*$. The defining relations of G_{Quot} are preserved by this transformation, so this is easily seen to be a well-defined homomorphism. At the Lie algebra this transformation is essentially replacing appropriate $\tilde{e}_{i,j}$ by $t_{i,j}\tilde{e}_{i,j}$.

For $g \in G_{\text{Quot}}$ and $t, t' \in R$ we have

$$(g^t)^{t'} = g^{tt'},$$

and for $g, g' \in G_{\text{Lin}}$ we have

$$g^t g'^t = g^{t+t'} \text{ and } g^t g'^t = (gg')^t.$$

This are trivial since G_{Lin} is abelian.

We next claim that if $g \in G_{\text{Quot}}$ and $g' \in G_{\text{Lin}}$ then

$$(11.1) \quad (gg'g^{-1})^t = gg'^t g^{-1}.$$

To prove this note that it suffices to prove the claim for powers of generators of the groups G_{Quot} and G_{Lin} (since conjugation and $g \mapsto g^t$ are homomorphisms). If $g \in G_{\text{Lin}}$ the result is trivial due to the abelian property, and if $g \notin G_{\text{Lin}}$ (and is the power of a generator) then $g^t = g$ by definition so $(gg'g^{-1})^t = g^t g'^t (g^{-1})^t = gg'^t g^{-1}$ as desired.

We now define $\rho: R \rightarrow \text{Aut}(G_{\text{Quot}} \times G_{\text{Lin}})$ by

$$\rho(t)(g, g_1) := (g \cdot g_1^t, g_1).$$

The map $\rho(t)$ is clearly bijective and we have

$$\rho(s)(\rho(t)(g, g_1)) = \rho(s)((g \cdot g_1^t, g_1)) = (g \cdot g_1^{t+s}, g_1) = \rho(t+s)(g, g_1),$$

so to check this is a group action it suffices to show $\rho(t)$ gives a valid homomorphism of $G_{\text{Quot}} \times G_{\text{Lin}}$. This follows because

$$\begin{aligned} \rho(t)((g, g_1) \cdot (g', g'_1)) &= \rho(t)(gg', (g')^{-1}g_1g'g'_1) \\ &= (gg'(g')^{-1}g_1^t g'^t (g'_1)^t, (g')^{-1}g_1g'g'_1) = (gg_1^t g'^t (g'_1)^t, (g')^{-1}g_1g'g'_1), \end{aligned}$$

by (11.1), while

$$\begin{aligned} \rho(t)(g, g_1)\rho(t)(g', g'_1) &= (gg_1^t, g_1) \cdot (g'(g'_1)^t, g'_1) = (gg_1^t g'^t (g'_1)^t, (g')^{-1}g_1g'g'_1) \\ &= (gg_1^t g'^t (g'_1)^t, (g'_1)^{-t}((g')^{-1}g_1g')(g'_1)^t g'_1) = (gg_1^t g'^t (g'_1)^t, (g'_1)^{-t}(g'_1)^t((g')^{-1}g_1g')g'_1) \\ &= (gg_1^t g'^t (g'_1)^t, (g')^{-1}g_1g'g'_1), \end{aligned}$$

where we have used that G_{Lin} is abelian and normal.

We are now in position to define the group of interest which will support the multidegree $(1, s - 1)$ nilsequence. Let

$$G_{\text{Multi}} = R \times_{\rho} (G_{\text{Quot}} \times G_{\text{Lin}})$$

where multiplication is given by

$$(t, (g, g_1))(t', (g', g'_1)) = (t + t', (\rho(t')(g, g_1)) \cdot (g', g'_1)).$$

This is seen to be a connected, simply connected Lie group. We give it a multidegree filtration $(G_{\text{Multi}})_{(d_1, d_2)}$ defined by:

- If $d_1 > 1$ then $(G_{\text{Multi}})_{(d_1, d_2)} = \text{Id}_{G_{\text{Multi}}}$;
- If $d_2 > 0$ then $(G_{\text{Multi}})_{(1, d_2)} = \{(0, (g, \text{id}_{G_{\text{Lin}}})) : g \in (G_{\text{Quot}})_{(d_2, 0)} \cap G_{\text{Lin}}\}$;

- $(G_{\text{Multi}})_{(1,0)} = \{(t, (g, \text{id}_{G_{\text{Lin}}})) : t \in R, g \in (G_{\text{Quot}})_{(0,0)} \cap G_{\text{Lin}}\}$ or equivalently just $\{(t, (g, \text{id}_{G_{\text{Lin}}})) : t \in R, g \in G_{\text{Lin}}\}$;
- If $d_2 > 0$ then $(G_{\text{Multi}})_{(0,d_2)} = \{(0, (g, g_1)) : g \in (G_{\text{Quot}})_{(d_2,0)}, g_1 \in (G_{\text{Quot}})_{(d_2,0)} \cap G_{\text{Lin}}\}$;
- $(G_{\text{Multi}})_{(0,0)} = G_{\text{Multi}}$.

Claim 11.2. $(G_{\text{Multi}})_{(d_1,d_2)}$ is a valid multidegree filtration on G_{Multi} .

Proof. Note that

$$(t, (g, g_1)) = (t, (\text{id}_{G_{\text{Quot}}}, \text{id}_{G_{\text{Lin}}})) \cdot (0, (g, g_1))$$

and therefore $(G_{\text{Multi}})_{(0,0)} = (G_{\text{Multi}})_{(1,0)} \vee (G_{\text{Multi}})_{(0,1)}$. We next check various commutator relations. First note that

$$[(G_{\text{Multi}})_{(1,0)}, (G_{\text{Multi}})_{(1,0)}] = \text{Id}_{G_{\text{Multi}}}.$$

This follows because if $g, h \in G_{\text{Lin}}$ we have $gh = hg$ hence

$$(t, (g, \text{id}_{G_{\text{Lin}}})) \cdot (t', (h, \text{id}_{G_{\text{Lin}}})) = (t + t', (gh, \text{id}_{G_{\text{Lin}}})) = (t', (h, \text{id}_{G_{\text{Lin}}})) \cdot (t, (g, \text{id}_{G_{\text{Lin}}})).$$

Therefore it suffices to verify that

$$\begin{aligned} [(G_{\text{Multi}})_{(0,a)}, (G_{\text{Multi}})_{(0,b)}] &\leq (G_{\text{Multi}})_{(0,a+b)}, \\ [(G_{\text{Multi}})_{(1,a)}, (G_{\text{Multi}})_{(0,b)}] &\leq (G_{\text{Multi}})_{(1,a+b)}. \end{aligned}$$

We first tackle the first claim, in which we may reduce to the case $a, b > 0$. We wish to show

$$[(g, g_1), (g', g'_1)] \in \{(h, h_1) : h \in (G_{\text{Quot}})_{(a+b,0)}, h_1 \in (G_{\text{Quot}})_{(a+b,0)} \cap G_{\text{Lin}}\}$$

if $g, g_1 \in (G_{\text{Quot}})_{(a,0)}$, $g', g'_1 \in (G_{\text{Quot}})_{(b,0)}$, and $g_1, g'_1 \in G_{\text{Lin}}$. Via Lemma 2.2, it suffices to prove $(G_{\text{Quot}})_{(a+b,0)}$ is normal in $(G_{\text{Quot}})_{(a,0)}$ and $(G_{\text{Quot}})_{(b,0)}$ and then check at the level of generators.

To check normality, we have

$$\begin{aligned} (g, g_1)(g', g'_1)(g, g_1)^{-1} &= (g, g_1)(g', g'_1)(g^{-1}, gg_1^{-1}g^{-1}) \\ &= (gg', (g')^{-1}g_1g' \cdot g'_1)(g^{-1}, gg_1^{-1}g^{-1}) \\ &= (gg'g^{-1}, (g(g')^{-1})g_1(g'g^{-1}) \cdot gg'_1g^{-1} \cdot gg_1^{-1}g^{-1}) \end{aligned}$$

and the result follows noting that $G_{\text{Lin}}, (G_{\text{Quot}})_{(j,0)}$ are normal in G_{Quot} for all $j \geq 0$.

Since

$$(g, g_1) = (g, \text{id}_{G_{\text{Lin}}}) \cdot (\text{id}_{G_{\text{Quot}}}, g_1)$$

and it suffices to check the claim on generators, we may reduce to the case where exactly one of g, g_1 and exactly one of g_1, g'_1 are the identity. The result is clear when g, g' are trivial, and the case when g_1, g'_1 are trivial follows from the fact that we have a valid filtration on G_{Quot} . In the remaining cases we may assume by symmetry that $g_1 = \text{id}_{G_{\text{Lin}}}$ and $g' = \text{id}_{G_{\text{Quot}}}$. We have

$$(g^{-1}, \text{id}_{G_{\text{Lin}}})(\text{id}_{G_{\text{Quot}}}, (g'_1)^{-1})(g, \text{id}_{G_{\text{Lin}}})(\text{id}_{G_{\text{Quot}}}, g'_1) = (\text{id}_{G_{\text{Quot}}}, g^{-1}(g'_1)^{-1}gg'_1)$$

and we see that the final coordinate satisfies $[g, g'_1] \in (G_{\text{Quot}})_{(a+b,0)} \cap G_{\text{Lin}}$. We have finished verifying the first claim.

Now note that $\{(h, \text{id}_{G_{\text{Lin}}}) : h \in G_{\text{Lin}}\}$ is a normal subgroup of $G_{\text{Quot}} \times G_{\text{Lin}}$, since G_{Lin} is abelian. Thus combining with the first claim gives the second claim, namely

$$[(G_{\text{Multi}})_{(1,a)}, (G_{\text{Multi}})_{(0,b)}] \leq (G_{\text{Multi}})_{(1,a+b)},$$

for $a > 0$.

The only nontrivial case left is $a = 0$ and $b > 0$ for the second claim. Furthermore, combining what we know it suffices to check the case when $(t, (\text{id}_{G_{\text{Quot}}}, \text{id}_{G_{\text{Lin}}}))$ is the element from $(G_{\text{Multi}})_{(1,0)}$. Note however that

$$\begin{aligned} & (t, (\text{id}_{G_{\text{Quot}}}, \text{id}_{G_{\text{Lin}}})) \cdot (0, (g, g_1)) \cdot (-t, (\text{id}_{G_{\text{Quot}}}, \text{id}_{G_{\text{Lin}}})) \cdot (0, (g, g_1))^{-1} \\ &= (t, (g, g_1)) \cdot (-t, (\text{id}_{G_{\text{Quot}}}, \text{id}_{G_{\text{Lin}}})) \cdot (0, (g^{-1}, gg_1^{-1}g^{-1})) = (0, (gg_1^{-t}, g_1)) \cdot (0, (g^{-1}, gg_1^{-1}g^{-1})) \\ &= (0, (gg_1^{-t}g^{-1}, \text{id}_{G_{\text{Lin}}})) \end{aligned}$$

and the fact that if $g, g_1 \in (G_{\text{Quot}})_{(b,0)}$ and $g_1 \in G_{\text{Lin}}$ then $gg_1^{-t}g^{-1} \in (G_{\text{Quot}})_{(b,0)}$. This follows because if $g_1 \in (G_{\text{Quot}})_{(b,0)} \cap G_{\text{Lin}}$ then g_1^t is in the same group. \square

Writing $t = (t_{i,j})_{1 \leq i \leq s-1, D_i^* < j \leq D_i^* + D_i^{\text{Lin}}}$, we define

$$\Gamma_{\text{Multi}} = \{(t, (g, g_1)) : t_{i,j} \in \mathbb{Z}, g \in \Gamma_{\text{Quot}}, g_1 \in \Gamma_{\text{Quot}} \cap G_{\text{Lin}}\}.$$

To see this is a group, observe that for $g_1 \in \Gamma_{\text{Quot}}$ we have $g^t \in \Gamma_{\text{Quot}}$ if all coordinates of t are integral. This is clear for the generators of Γ_{Quot} and the rest follows from recalling that “taking t -th powers” is a homomorphism on G_{Quot} .

We now define the relevant functions which will be used to represent $F^*(g_h^{\text{Quot}}(n)\Gamma^{\text{Quot}})$. Let $\delta = \exp(-O_s((d \log(MD/\rho))^{O_s(1)}))$, where the implicit constants are chosen sufficiently large.

Let $\phi: \mathbb{R} \rightarrow \mathbb{R}$ be a 1-bounded, 1-periodic function such that:

- $\phi(x) = 1$ if $|\{x\}| \leq 1/2 - 2\delta$;
- $\phi(x) = 0$ if $|\{x\}| \geq 1/2 - \delta$;
- ϕ is $O(1/\delta)$ -Lipschitz.

Define $H^* \subseteq H$ such that for all $1 \leq i \leq s-1$ and $D_i^* < j \leq D_i^* + D_i^{\text{Lin}}$ we have $|\{\beta_{i,j}h\}| \geq 1/2 - \delta$. Using that $\beta_{i,j} \in (1/N')\mathbb{Z}$ where N' is a prime between $100N$ and $200N$, we see that there are at most $O(\delta \cdot N \cdot \sum_{i=1}^{s-1} D_i^{\text{Lin}})$ indices which do not satisfy the criterion and choosing δ sufficiently small, we may assume that H^* is at least half the size of H .

Given $(t, (g, g_1)) \in G_{\text{Multi}}$, we may find $(t', (g', g'_1)) \in (t, (g, g_1))\Gamma_{\text{Multi}}$ such that $(t')_{i,j} \in (-1/2, 1/2]$ for all i, j . Define

$$F_{\text{Multi}}((t, (g, g_1))\Gamma_{\text{Multi}}) = F^*(g'\Gamma_{\text{Quot}}) \cdot \prod_{\substack{1 \leq i \leq s-1 \\ D_i^* < j \leq D_i^* + D_i^{\text{Lin}}}} \phi(t'_{i,j});$$

we check that this in fact gives a well-defined function on $G_{\text{Multi}}/\Gamma_{\text{Multi}}$. Note that if $(t', (g', g'_1)) \in (t, (g, g_1))\Gamma_{\text{Multi}}$ and $t'_{i,j} \in (-1/2, 1/2]$ then $t'_{i,j} = \{t_{i,j}\}$ and hence t' is unique. Furthermore note that

$$(t', (g', g'_1)) \cdot (0, (\gamma', \gamma'_1)) = (t', (g', g'_1)) \cdot (\gamma', \gamma'_1) = (t', (g'\gamma', (\gamma')^{-1}g'_1\gamma'_1))$$

and trivially

$$F^*(g'\Gamma_{\text{Quot}}) = F^*(g'\gamma'\Gamma_{\text{Quot}})$$

if $\gamma' \in \Gamma_{\text{Quot}}$. Now recall that

$$g_h^{\text{Quot}}(n) = \prod_{i=1}^{s-1} \prod_{j=1}^{D_i^*} \exp(\tilde{e}_{i,j} \gamma_{i,j} \cdot \frac{n^i}{i!}) \cdot \prod_{i=1}^{s-1} \prod_{j=D_i^*+1}^{D_i^*+D_i^{\text{Lin}}} \exp(\tilde{e}_{i,j})^{\alpha_{i,j} \{\beta_{i,j}h\} \cdot \frac{n^i}{i!}}.$$

We set

$$g_0(n) = \prod_{i=1}^{s-1} \prod_{j=1}^{D_i^*} \exp(\tilde{e}_{i,j} \gamma_{i,j} \cdot \frac{n^i}{i!}), \quad g_1(n) = \prod_{i=1}^{s-1} \prod_{j=D_i^*+1}^{D_i^*+D_i^{\text{Lin}}} \exp(\tilde{e}_{i,j})^{\alpha_{i,j} \cdot \frac{n^i}{i!}}$$

and define

$$\begin{aligned} g_{\text{Final}}(h, n) &= (0, (g_0(n), g_1(n))) \cdot ((\beta_{i,j}h)_{\substack{1 \leq i \leq s-1 \\ D_i^* < j \leq D_i^* + D_i^{\text{Lin}}}}, (\text{id}_{G_{\text{Quot}}}, \text{id}_{G_{\text{Lin}}})) \\ &= (0, (g_0(n), \text{id}_{G_{\text{Lin}}})) \cdot (0, (\text{id}_{G_{\text{Quot}}}, g_1(n))) \cdot ((\beta_{i,j}h)_{\substack{1 \leq i \leq s-1 \\ D_i^* < j \leq D_i^* + D_i^{\text{Lin}}}}, (\text{id}_{G_{\text{Quot}}}, \text{id}_{G_{\text{Lin}}})) \end{aligned}$$

$g_{\text{Final}}(h, n)$ is seen to be a polynomial sequence with respect to the filtration given to G_{Multi} as each piece is trivially a polynomial sequence and the polynomial sequences form a group under pointwise multiplication (see [29, Corollary B.4]).

Note that for all $h \in H$ we have

$$\begin{aligned} g_{\text{Final}}(h, n)\Gamma_{\text{Multi}} &= (0, (g_0(n), g_1(n))) \cdot ((\beta_{i,j}h)_{\substack{1 \leq i \leq s-1 \\ D_i^* < j \leq D_i^* + D_i^{\text{Lin}}}}, (\text{id}_{G_{\text{Quot}}}, \text{id}_{G_{\text{Lin}}}))\Gamma_{\text{Multi}} \\ &= ((\beta_{i,j}h)_{\substack{1 \leq i \leq s-1 \\ D_i^* < j \leq D_i^* + D_i^{\text{Lin}}}}, (\text{id}_{G_{\text{Quot}}}, \text{id}_{G_{\text{Lin}}})) \cdot (0, (g_{0,h}^*(n), g_1(n)))\Gamma_{\text{Multi}}, \end{aligned}$$

writing

$$g_{0,h}^*(n) = g_0(n)(g_1(n))^{t(h)}$$

where $t(h) = (\{\beta_{i,j}h\})_{1 \leq i \leq s-1, D_i^* < j \leq D_i^* + D_i^{\text{Lin}}} \in R$. This is precisely the desired sense, discussed earlier, in which we have used the group action to “raise” parts of G_{Lin} to h -fractional powers.

Therefore, for all $h \in H^*$ we have

$$(11.2) \quad F_{\text{Multi}}(g_{\text{Final}}(h, n)\Gamma_{\text{Multi}}) = F^*(g_h^{\text{Quot}}(n)\Gamma_{\text{Quot}}).$$

We now state various complexity claims regarding $G_{\text{Multi}}/\Gamma_{\text{Multi}}$ and the Lipschitz nature of the function F_{Multi} . We defer the rather uninspiring task of checking these bounds to the end of Appendix B.

Lemma 11.3. *Given the above setup, we have that $G_{\text{Multi}}/\Gamma_{\text{Multi}}$ has the structure of a multi-degree $(1, s-1)$ nilmanifold and it may be given a basis $\mathcal{X}_{\text{Multi}}$ of complexity bounded by $\exp(O_s((d \log(MD/\rho))^{O_s(1)}))$. Furthermore F_{Multi} is $\exp(O_s((d \log(MD/\rho))^{O_s(1)}))$ -Lipschitz under this metric.*

11.2. Extracting correlation. We now complete the proof of Lemma 6.3. The proof is little more than stitching results proven in this and the previous section and noting that if two nilcharacters “differ by a lower degree-rank term” then one may pass from the other at the cost of introducing a lower order term. (This is essentially [29, Lemma E.7].)

Proof of Lemma 6.3. We return to the correlation structure discussed in Section 10 (that is output by Lemma 9.1). Again, we will abuse notation slightly as discussed. So, for all $h \in H$ (where $|H| \geq \rho'N$) we have

$$\|\mathbb{E}_{n \in [N]}(\Delta_h f)(n) \otimes \overline{\chi(h, n)} \otimes \overline{\chi_h(n)} \cdot \overline{\psi_h(n)}\|_{\infty} \geq \exp(-O_s((d \log(MD/\rho))^{O_s(1)}))$$

where ψ_h is a complexity M' nilsequence of degree $(s-2)$ and dimension at most d' . We adopt the notation developed in Sections 10 and 11. Applying Claim 10.6, we have

$$\|\mathbb{E}_{n \in [N]}(\Delta_h f)(n) \otimes \overline{\chi(h, n)} \otimes \overline{\widetilde{F}(g_h^{\text{Univ}}(n)\Gamma_{\text{Univ}})} \cdot \overline{\psi_h(n)}\|_{\infty} \geq \exp(-O_s((d \log(MD/\rho))^{O_s(1)})).$$

Next note that

$$F^*(g'\Gamma_{\text{Quot}}) \otimes \overline{F^*(g'\Gamma_{\text{Quot}})}$$

has trace equal to 1 as F^* is a nilcharacter. Since the output dimension of F^* is bounded by $\exp((d \log(MD/\rho))^{O_s(1)})$, we have for all $h \in H$ that

$$\begin{aligned} \|\mathbb{E}_{n \in [N]}(\Delta_h f)(n) \otimes \overline{\chi(h, n)} \otimes \overline{\widetilde{F}(g_h^{\text{Univ}}(n)\Gamma_{\text{Univ}})} \otimes F^*(g_h^{\text{Quot}}(n)\Gamma_{\text{Quot}}) \\ \otimes \overline{F^*(g_h^{\text{Quot}}(n)\Gamma_{\text{Quot}})} \cdot \overline{\psi_h(n)}\|_{\infty} \geq \exp(-O_s((d \log(MD/\rho))^{O_s(1)})). \end{aligned}$$

Using (11.2), we in fact may write for $h \in H^*$ that

$$\begin{aligned} & \|\mathbb{E}_{n \in [N]}(\Delta_h f)(n) \otimes \overline{\chi(h, n)} \otimes \overline{\widetilde{F}(g_h^{\text{Univ}}(n) \Gamma_{\text{Univ}})} \otimes F^*(g_h^{\text{Quot}}(n) \Gamma_{\text{Quot}}) \otimes \overline{F_{\text{Multi}}(g_{\text{Final}}(h, n))} \cdot \overline{\psi_h(n)}\|_{\infty} \\ & \geq \exp(-d^{O_s(1)} \log(MD\rho^{-1})^{O_s(1)}). \end{aligned}$$

Now we may pay a cost of $\exp(-\overline{(d \log(MD/\rho))^{O_s(1)}}$ in the size of H^* by Pigeonhole to choose a single coordinate function of $\widetilde{F}(g_h^{\text{Univ}}(n) \Gamma_{\text{Univ}}) \otimes F^*(g' \Gamma_{\text{Quot}})$, call it $\psi_h^*(n)$, such that

$$\begin{aligned} & \|\mathbb{E}_{n \in [N]}(\Delta_h f)(n) \otimes \overline{\chi(h, n)} \otimes \psi_h^*(n) \otimes \overline{F_{\text{Multi}}(g_{\text{Final}}(h, n))} \cdot \overline{\psi_h(n)}\|_{\infty} \\ & \geq \exp(-O_s((d \log(MD/\rho))^{O_s(1)})). \end{aligned}$$

By Lemma 10.12 and using Remark 10.13, $\psi_h^*(n)$ can be realized on a nilmanifold with a degree-rank $(s-1, r^*-1)$ filtration. Furthermore the function underlying $\psi_h^*(n)$ has Lipschitz constant bounded by $\exp((d \log(MD/\rho))^{O_s(1)})$ and the nilmanifold it lives on has dimension at most $(d \log(MD/\rho))^{O_s(1)}$ and complexity bounded by $\exp((d \log(MD/\rho))^{O_s(1)})$ due to Lemma 10.12.

By applying [35, Lemma A.6] with subgroup corresponding to the $(s-1, r^*-1)$ degree-rank and Pigeonholing in the associated vertical frequency, we may assume that $\psi_h^*(n)$ has a vertical frequency with height bounded by $\exp((d \log(MD/\rho))^{O_s(1)})$; this may reduce the subset of H^* under consideration by a further admissible fraction. We then extend $\psi_h^*(n)$ to a nilcharacter by using Lemma B.4⁷; we refer to this nilcharacter as ψ_h^{Output} and note it is a degree-rank $(s-1, r^*-1)$ nilcharacter with appropriate complexity. We thus have

$$\begin{aligned} & \|\mathbb{E}_{n \in [N]}(\Delta_h f)(n) \otimes \overline{\chi(h, n)} \otimes \psi_h^{\text{Output}}(n) \otimes \overline{F_{\text{Multi}}(g_{\text{Final}}(h, n))} \cdot \overline{\psi_h(n)}\|_{\infty} \\ & \geq \exp(-O_s((d \log(MD/\rho))^{O_s(1)})). \end{aligned}$$

By Pigeonholing in h once again we may pass to F_{Multi}^* , which is a fixed coordinate of F_{Multi} ,

$$\begin{aligned} & \|\mathbb{E}_{n \in [N]}(\Delta_h f)(n) \otimes \overline{\chi(h, n)} \otimes \psi_h^{\text{Output}}(n) \cdot \overline{F_{\text{Multi}}^*(g_{\text{Final}}(h, n))} \cdot \overline{\psi_h(n)}\|_{\infty} \\ & \geq \exp(-d^{O_s(1)} \log(MD\rho^{-1})^{O_s(1)}). \end{aligned}$$

on a $\exp(-(d \log(MD/\rho))^{O_s(1)})$ fraction of indices. F_{Multi}^* lives on the group G_{Multi} and via [35, Lemma A.6], Pigeonholing in h so that we have the same frequency, and embedding in a nilcharacter via Lemma B.4 similar to the above argument, we have for all $h \in H^*$ that

$$\begin{aligned} & \|\mathbb{E}_{n \in [N]}(\Delta_h f)(n) \otimes \overline{\chi(h, n)} \otimes \psi_h^{\text{Output}}(n) \otimes \overline{F_{\text{Multi}}^{\text{Output}}(g_{\text{Final}}(h, n))} \cdot \overline{\psi_h(n)}\|_{\infty} \\ & \geq \exp(-O_s((d \log(MD/\rho))^{O_s(1)})) \end{aligned}$$

where $F_{\text{Multi}}^{\text{Output}}$ is a multidegree $(1, s-1)$ nilcharacter on G_{Multi} with vertical frequency height, output dimension, and Lipschitz constant of each coordinate bounded by $\exp((d \log(MD/\rho))^{O_s(1)})$ while the dimension of the underlying nilmanifold is bounded by $(d \log(MD/\rho))^{O_s(1)}$.

This completes the proof with $\chi(h, n) \otimes F_{\text{Multi}}^{\text{Output}}(g_{\text{Final}}(h, n))$ being the new multidegree $(1, s-1)$ nilcharacter, $\overline{\psi_h(n)^{\text{Output}}}$ being the degree-rank $(s-1, r^*-1)$ nilcharacter and noting that the density of indices h which remain is at least $\exp(-O_s((d \log(MD/\rho))^{O_s(1)}))$. \square

⁷We have that ψ_h^* lives on the group $G_{\text{Univ}}^{\Delta}/(G_{\text{Univ}}^{\Delta})_{(s-1, r^*)}$. We may give it the degree filtration

$$\begin{aligned} G_{\text{Univ}}^{\Delta}/(G_{\text{Univ}}^{\Delta})_{(s-1, r^*)} &= (G_{\text{Univ}}^{\Delta})_{(1,0)}/(G_{\text{Univ}}^{\Delta})_{(s-1, r^*)} \geq (G_{\text{Univ}}^{\Delta})_{(2,0)}/(G_{\text{Univ}}^{\Delta})_{(s-1, r^*)} \\ &\geq \cdots \geq (G_{\text{Univ}}^{\Delta})_{(s-1,0)}/(G_{\text{Univ}}^{\Delta})_{(s-1, r^*)} \geq (G_{\text{Univ}}^{\Delta})_{(s-1, r^*-1)}/(G_{\text{Univ}}^{\Delta})_{(s-1, r^*)} \geq \text{Id}_{G_{\text{Univ}}^{\Delta}/(G_{\text{Univ}}^{\Delta})_{(s-1, r^*)}} \end{aligned}$$

and we apply Lemma B.4 to this filtration to get a nilcharacter H . We then embed ψ_h^* by taking the underlying function, call it Q , and taking the nilcharacter $(Q/(2 \cdot \|Q\|_{\infty}), \sqrt{1 - |Q/(2 \cdot \|Q\|_{\infty})|^2} \cdot H)$.

12. SYMMETRIZATION ARGUMENT

We now perform the necessary symmetrization argument. In particular, at this stage in the argument due to Theorem 6.4 we have shown that for many h , $\Delta_h f$ correlates with $\chi(h, n)$ which is a multidegree $(1, s - 1)$ nilcharacter. We now demonstrate that $\chi(h, n)$ is “symmetric up to lower order terms” in h and n (after multilinearizing the n variable) via an argument of Green, Tao, and Ziegler [29], which in turn is closely related to an earlier argument of Green and Tao [18] which proved such a result for the U^3 -norm. Our treatment is slightly simpler than in [29]. Importantly, this argument is fundamentally based on a finite number of applications of Cauchy–Schwarz and a single call to equidistribution theory and therefore naturally comes with good bounds.

All references to Appendix C are simply quantified versions of lemmas which appear in the work of Green, Tao, and Ziegler [29, Appendix E] and a discussion of the correspondence is given more carefully in Appendix C. The reader may benefit from glancing at the statements in Appendix C or those in [29, Appendix E].

For the remainder of this section and Appendix C, to lighten statements, we say a nilsequence χ has complexity (M, d) if the underlying nilmanifold G/Γ has complexity M , the underlying function is M -Lipschitz, and the dimension of G is bounded by d . We will say a nilcharacter χ has complexity (M, d) if the underlying nilmanifold G/Γ has complexity M , the output dimension of χ is bounded by M , the underlying function has all coordinates being M -Lipschitz, the vertical character underlying χ has height bounded by M , and the dimension of G is bounded by d . In this section, M will always be of the form $M(\delta) := \exp(\log(1/\delta)^{O_s(1)})$ while the underlying d will be of the form $d(\delta) := \log(1/\delta)^{O_s(1)}$ in our analysis, where the implicit constants may, by abuse of notation, vary from line to line.

We now recall the output of Theorem 6.4. We have

$$\mathbb{E}_{h \in [N]} \|\mathbb{E}_{n \in [N]} \Delta_h f(n) \otimes \chi(h, n) \psi_h(n)\|_\infty \geq M(\delta)^{-1}.$$

Here $\psi_h(n)$ is a degree $(s - 2)$ nilsequence and $\chi(h, n) = F(g(h, n)\Gamma)$ is a multidegree $(1, s - 1)$ nilcharacter. Furthermore χ has complexity $(M(\delta), d(\delta))$ while $\psi_h(n)$ has complexity $(M(\delta), d(\delta))$.

Our first step is to multilinearize χ in the n variable, replacing it by a multidegree $(1, 1, \dots, 1)$ nilcharacter which is symmetric in the final $(s - 1)$ variables.

Lemma 12.1. *Fix $s \geq 2$. Suppose that*

$$\mathbb{E}_{h \in [N]} \|\mathbb{E}_{n \in [N]} \Delta_h f(n) \otimes \chi(h, n) \cdot \psi_h(n)\|_\infty \geq 1/M(\delta)$$

with $\chi(h, n)$ being a periodic multidegree $(1, s - 1)$ -nilcharacter and $\psi_h(n)$ are degree $(s - 2)$ nilsequences each of complexity $(M(\delta), d(\delta))$.

There exists $\tilde{\chi}$ a multidegree $(1, \dots, 1)$ nilcharacter (with s ones), $\tilde{\psi}$ a degree $(s - 1)$ nilsequence, and there exist $\tilde{\psi}_h(n)$ which are degree $(s - 2)$ nilcharacters all having complexity complexity $(M(\delta), d(\delta))$ such that

$$\mathbb{E}_{h \in [N]} \|\mathbb{E}_{n \in [N]} \Delta_h f(n) \otimes \tilde{\chi}(h, n, \dots, n) \cdot \tilde{\psi}(n) \otimes \tilde{\psi}_h(n)\|_\infty \geq 1/M(\delta).$$

Furthermore $\tilde{\chi}$ is symmetric in the final $(s - 1)$ coordinates, i.e., for any $\sigma \in \mathfrak{S}_{s-1}$ we have

$$\tilde{\chi}(h, n_1, \dots, n_{s-1}) = \tilde{\chi}(h, n_{\sigma(1)}, \dots, n_{\sigma(s-1)}).$$

Proof. This is essentially an immediate consequence of multilinearization (see e.g. [29, Theorem E.10]). By applying Lemma C.5, there is multidegree $(1, \dots, 1)$ nilcharacter $\tilde{\chi}$ of complexity $(M(\delta), d(\delta))$ such that $\chi(h, n)$ and $\tilde{\chi}(h, n, \dots, n)$ are $(M(\delta), M(\delta), d(\delta))$ -equivalent for degree $(s - 1)$. Furthermore $\tilde{\chi}$ is symmetric in the final $(s - 1)$ coordinates.

Thus applying Lemma 7.4 (and the remark following), there exists a nilsequence $\psi^*(h, n)$ of degree $\leq (s - 1)$ and complexity $(M(\delta), d(\delta))$ such that

$$\mathbb{E}_{h \in [N]} \|\mathbb{E}_{n \in [N]} \Delta_h f(n) \otimes \tilde{\chi}(h, n, \dots, n) \otimes \psi^*(h, n) \cdot \psi_h(n)\|_\infty \geq 1/M(\delta).$$

Note that a degree $(s - 1)$ nilsequence of complexity $(M(\delta), d(\delta))$ in two variables (h, n) is also a multidegree $(0, s - 1) \cup (s - 1, s - 2)$ nilsequence of complexity $(M(\delta), d(\delta))$ via taking the filtration $G_i^* := G_{|i|}^*$. Therefore by Lemma C.6 and the first item of Lemma C.2, there exist nilsequences $\tilde{\psi}(n)$ and $\psi_h^*(n)$ of degree $(s - 1)$ and $(s - 2)$ respectively and complexity $(M(\delta), d(\delta))$ such that

$$\mathbb{E}_{h \in [N]} \|\mathbb{E}_{n \in [N]} \Delta_h f(n) \otimes \tilde{\chi}(h, n, \dots, n) \otimes \tilde{\psi}(n) \cdot \psi_h^*(n) \cdot \psi_h(n)\|_\infty \geq 1/M(\delta).$$

Now, $\psi_h^*(n) \cdot \psi_h(n)$ is a degree $(s - 2)$ nilsequence of complexity $(M(\delta), d(\delta))$. Applying [35, Lemma A.6], we may replace this product by $\psi'_h(n)$ which is a degree $(s - 2)$ nilsequence of complexity $(M(\delta), d(\delta))$ with a vertical frequency of height $\exp(\log(1/\delta)^{O_s(1)})$. Finally apply Lemma B.4 and embed $\psi'_h(n)$ as a coordinate of a nilcharacter $\tilde{\psi}_h(n)$ of complexity $(M(\delta), d(\delta))$, similar to in the proof of Lemma 6.3. We thus have

$$\mathbb{E}_{h \in [N]} \|\mathbb{E}_{n \in [N]} \Delta_h f(n) \otimes \tilde{\chi}(h, n, \dots, n) \cdot \tilde{\psi}(n) \otimes \tilde{\psi}_h(n)\|_\infty \geq 1/M(\delta)$$

where $\tilde{\psi}$ and $\tilde{\psi}_h$ have the appropriate properties. \square

We are now in position to complete the proof of Theorem 1.2 via a symmetrization argument. Our argument is analogous to that of Green, Tao, and Ziegler [29, Section 13] modulo certain minor simplifications to the underlying Cauchy–Schwarz arguments.

Proof of Theorem 1.2. We may assume that $s \geq 3$. The case $s = 0$ is trivial, $s = 1$ is standard Fourier analysis, and the case $s = 2$ follows from work of Sanders [45] (see [37, Theorem 8]). Furthermore, throughout the analysis we will assume implicitly that $N \geq \exp(\log(1/\delta)^{\Omega_s(1)})$; in the case when N is small one may deduce the statement via Fourier analysis. We proceed by induction, assuming that the inverse theorem is known for smaller s .

By Theorem 6.4 and then Lemma 12.1 we may assume that

$$(12.1) \quad \mathbb{E}_{h \in [N]} \|\mathbb{E}_{n \in [N]} \Delta_h f(n) \otimes \chi(h, n, \dots, n) \cdot \psi(n) \otimes \psi_h(n)\|_\infty \geq 1/M(\delta).$$

Here χ is a multidegree $(1, \dots, 1)$ nilcharacter which is symmetric in the final $(s - 1)$ variables, ψ is a degree $(s - 1)$ nilsequence, and ψ_h are degree $(s - 2)$ nilcharacters with complexities bounded by complexity $(M(\delta), d(\delta))$. For $h \notin [N]$, we take $\psi_h(n)$ to be the constant function 1 (which is a degree 0 nilcharacter) throughout the argument. Additionally, we may use differently indexed versions of functions ψ that are defined at intermediate stages of the argument; although an abuse of notation, it will always be clear from context.

Step 1: Initial setup for Cauchy–Schwarz argument. For the sake of shorthand, we will denote $\tilde{\chi}(h, n) = \chi(h, n, \dots, n)$ where there are $(s - 1)$ copies of the variable n . By Lemma 7.2 (taking $f_1 = f$ and $f_2 = \psi(n)$), we have

$$\begin{aligned} \mathbb{E}_{\substack{h_1 + h_2 = h_3 + h_4 \\ h_i \in [N]}} \|\mathbb{E}_{n \in [N]} \tilde{\chi}(h_1, n) \otimes \tilde{\chi}(h_2, n + h_1 - h_4) \otimes \overline{\tilde{\chi}(h_3, n)} \otimes \overline{\tilde{\chi}(h_4, n + h_1 - h_4)} \\ \otimes \psi_{h_1}(n) \otimes \psi_{h_2}(n + h_1 - h_4) \otimes \overline{\psi_{h_3}(n)} \otimes \overline{\psi_{h_4}(n + h_1 - h_4)} \cdot e(\Theta n)\|_\infty \geq 1/M(\delta) \end{aligned}$$

for some $\|\Theta\|_{\mathbb{R}/\mathbb{Z}} \leq M(\delta)/N$. Note that Lemma 7.2 is stated for scalar function; here we are using that we may Pigeonhole on coordinates of the vector $\chi(h, n, \dots, n) \cdot \psi(n) \otimes \psi_h(n)$ before using Lemma 7.2.

We next change variables with $h_1 = h + x$, $h_2 = h + y$, $h_3 = h + x + y$, and $h_4 = h$. The above then implies that

$$\mathbb{E}_{h \in [N], x, y \in [\pm N]} \|\mathbb{E}_{n \in [N]} \tilde{\chi}(h + x, n) \otimes \tilde{\chi}(h + y, n + x) \otimes \overline{\tilde{\chi}(h + x + y, n)} \otimes \overline{\tilde{\chi}(h, n + x)} \\ \otimes \psi_{h+x}(n) \otimes \psi_{h+y}(n + x) \otimes \overline{\psi_{h+x+y}(n)} \otimes \overline{\psi_h(n + x)} e(\Theta n)\|_\infty \geq 1/M(\delta).$$

By the first item of Lemma C.3, $\psi_{h+y}(n + x)$ and $\psi_{h+y}(n)$ are $(M(\delta), M(\delta), d(\delta))$ -equivalent for degree $(s - 3)$. We use that $s \geq 3$ precisely here so that this is a well-defined term.

Therefore by Lemma 7.4, there exists a collection $\psi_{h,x,y}(n)$ of degree $(s - 3)$ nilsequences each of complexity $(M(\delta), d(\delta))$ such that

$$\mathbb{E}_{h \in [N], x, y \in [\pm N]} \|\mathbb{E}_{n \in [N]} \tilde{\chi}(h + x, n) \otimes \tilde{\chi}(h + y, n + x) \otimes \overline{\tilde{\chi}(h + x + y, n)} \otimes \overline{\tilde{\chi}(h, n + x)} \\ \otimes \psi_{h+x}(n) \otimes \psi_{h+y}(n) \otimes \overline{\psi_{h+x+y}(n)} \otimes \overline{\psi_h(n + x)} \psi_{h,x,y}(n) \cdot e(\Theta n)\|_\infty \geq 1/M(\delta).$$

We will use B to denote vector-valued functions (which may vary term to term) with coordinates which are 1-bounded such that the dimension is bounded by $M(\delta)$. The key point is that nearly all terms may be folded into 1-bounded terms. In particular, we have

$$\mathbb{E}_{h \in [N], x, y \in [\pm N]} \|\mathbb{E}_{n \in [N]} \tilde{\chi}(h + y, n + x) \cdot \psi_{h,x,y}(n) \otimes B(h, x, n) \otimes B(h, y, n) \otimes B(h, x + y, n)\|_\infty \geq 1/M(\delta).$$

Noting that $\psi_{h,x,y}(n)$ may be twisted by an appropriate complex phase depending on h , we may in fact assume that

$$\|\mathbb{E}_{h, n \in [N], x, y \in [\pm N]} \tilde{\chi}(h + y, n + x) \cdot \psi_{h,x,y}(n) \otimes B(h, x, n) \otimes B(h, y, n) \otimes B(h, x + y, n)\|_\infty \geq 1/M(\delta).$$

By applying Pigeonhole in h , we may fix h^* such that

$$\|\mathbb{E}_{n \in [N], x, y \in [\pm N]} \tilde{\chi}(h^* + y, n + x) \cdot \psi_{h^*,x,y}(n) \otimes B(x, n) \otimes B(y, n) \otimes B(x + y, n)\|_\infty \geq 1/M(\delta).$$

Taking the coordinate which achieves the infinity norm, we may assume that $B(\cdot, \cdot)$ are in fact all scalar and thus

$$\|\mathbb{E}_{x, y \in [\pm N]} \mathbb{E}_{n \in [N]} \tilde{\chi}(h^* + y, n + x) \cdot \psi_{x,y}(n) \cdot b(x, n) \cdot b(y, n) \cdot b(x + y, n)\|_\infty \geq 1/M(\delta);$$

we have dropped h^* in one subscript here.

By applying the second item of Lemma C.3 and the second item of Lemma C.2, we have that $\tilde{\chi}(h^* + y, n + x)$ and $\tilde{\chi}(y, n + x)$ are $(M(\delta), M(\delta), d(\delta))$ -equivalent for degree $(s - 1)$. Thus by Lemma 7.4 there exists a nilsequence $\psi^*(x, y, n)$ of degree $(s - 1)$ and complexity $(M(\delta), d(\delta))$ such that

$$\|\mathbb{E}_{x, y \in [\pm N]} \mathbb{E}_{n \in [N]} \tilde{\chi}(y, n + x) \cdot \psi^*(x, y, n) \cdot \psi_{x,y}(n) \cdot b(x, n) b(y, n) b(x + y, n)\|_\infty \geq 1/M(\delta).$$

Note that a degree $(s - 1)$ nilsequence in variables x, y, n is a multidegree $(s - 1, s - 1, s - 3) \cup (1, 0, s - 2) \cup (0, 1, s - 2) \cup (0, 0, s - 1)$ -nilsequence. Therefore applying Lemma C.6 and applying Pigeonhole, we may adjust $\psi_{x,y}$ and the 1-bounded functions and remove ψ^* and thus we may assume that

$$\|\mathbb{E}_{x, y \in [\pm N]} \mathbb{E}_{n \in [N]} \tilde{\chi}(y, n + x) \cdot \psi_{x,y}(n) \cdot b(x, n) b(y, n) b(x + y, n)\|_\infty \geq 1/M(\delta);$$

note that $\psi_{x,y}$ and B have all been modified but we have abusively maintained the same notation. In particular, $\psi_{x,y}(n)$ is degree $(s - 3)$.

By Lemma C.4, the second item of Lemma C.2, and Lemma C.1 (and symmetry of χ in the final $(s - 1)$ coordinates), we have that $\tilde{\chi}(y, n + x)$ and

$$\bigotimes_{k=0}^{s-1} \chi(y, n, \dots, n, x, \dots, x)^{\otimes \binom{s-1}{k}}$$

are $(M(\delta), M(\delta), d(\delta))$ -equivalent for degree $(s-1)$. In this notation there are k copies of n and $s-1-k$ copies of x . Now by Lemma 7.4, we have

$$\left\| \mathbb{E}_{x,y \in [\pm N]} \mathbb{E}_{n \in [N]} \psi^*(x, y, n) \cdot \bigotimes_{k=0}^{s-1} \chi(y, n, \dots, n, x, \dots, x)^{\otimes \binom{s-1}{k}} \cdot \psi_{x,y}(n) \cdot b(x, n) b(y, n) b(x+y, n) \right\|_{\infty} \geq 1/M(\delta)$$

where $\psi^*(x, y, n)$ is a new degree $(s-1)$ nilsequence of complexity $(M(\delta), d(\delta))$. Applying Lemma C.6 as before, we may adjust $\psi_{x,y}(n)$ and the 1-bounded functions and remove this term to have that

$$\left\| \mathbb{E}_{x,y \in [\pm N]} \mathbb{E}_{n \in [N]} \bigotimes_{k=0}^{s-1} \chi(y, n, \dots, n, x, \dots, x)^{\otimes \binom{s-1}{k}} \cdot \psi_{x,y}(n) b(x, n) b(y, n) b(x+y, n) \right\|_{\infty} \geq 1/M(\delta).$$

Note that the only terms of $\bigotimes_{k=0}^{s-1} \chi(y, n, \dots, n, x, \dots, x)^{\otimes \binom{s-1}{k}}$ which involve all of x, y, n with n appearing at least $s-2$ times have exactly one copy of x , one copy of y and n exactly $(s-2)$ times. Therefore taking the coordinate of

$$\bigotimes_{\substack{0 \leq k \leq s-1 \\ k \neq s-2}} \chi(y, n, \dots, n, x, \dots, x)^{\otimes \binom{s-1}{k}}$$

which achieves the infinity norm and adjusting $\psi_{x,y}$, b , and adding a term $b(x, y)$ we have

$$\left\| \mathbb{E}_{x,y \in [\pm N]} \mathbb{E}_{n \in [N]} \chi(y, x, n, \dots, n)^{\otimes (s-1)} \cdot \psi_{x,y}(n) b(x, n) b(y, n) b(x+y, n) b(x, y) \right\|_{\infty} \geq 1/M(\delta).$$

Step 2: Cauchy–Schwarz to remove 1-bounded functions. Applying Cauchy–Schwarz to each coordinate of the associated vector, duplicating the variable y , and using that $b(x, n)$ is 1-bounded, we find that

$$\left\| \mathbb{E}_{n \in [N], x \in [\pm N]} \mathbb{E}_{y, y' \in [\pm N]} \chi(y, x, n, \dots, n)^{\otimes (s-1)} \otimes \overline{\chi(y', x, n, \dots, n)^{\otimes (s-1)}} \cdot \psi_{x,y}(n) \overline{\psi_{x,y'}(n)} \cdot b(y, n) \overline{b(y', n)} \cdot b(x+y, n) \overline{b(x+y', n)} \cdot b(x, y) \overline{b(x, y')} \right\|_{\infty} \geq 1/M(\delta).$$

By Lemma C.4, Lemma C.2, and Lemma C.1, we have that

$$\chi(y, x, n, \dots, n)^{\otimes (s-1)} \otimes \overline{\chi(y', x, n, \dots, n)^{\otimes (s-1)}} \text{ and } \chi(y-y', x, n, \dots, n)^{\otimes (s-1)}$$

are $(M(\delta), M(\delta), d(\delta))$ -equivalent for degree $(s-1)$. Therefore by Lemma 7.4, there exists $\psi^*(x, y, y', n)$ a degree $(s-1)$ nilsequence of complexity $(M(\delta), d(\delta))$ such that

$$\left\| \mathbb{E}_{n \in [N], x, y, y' \in [\pm N]} \chi(y-y', x, n, \dots, n)^{\otimes (s-1)} \cdot \psi^*(x, y, y', n) \cdot \psi_{x,y}(n) \overline{\psi_{x,y'}(n)} \cdot b(y, n) \overline{b(y', n)} \cdot b(x+y, n) \overline{b(x+y', n)} \cdot b(x, y) \overline{b(x, y')} \right\|_{\infty} \geq 1/M(\delta).$$

Note that $z = x + y + y'$ ranges in the set $[-3N, 3N]$. Take $\rho = \exp(-\log(1/\delta)^{O_s(1)})$ sufficiently small. Then there exists z^* such that $z^* \in [-(3-\rho)N, (3-\rho)N]$ such that

$$\left\| \mathbb{E}_{n \in [N]} \mathbb{E}_{\substack{x, y, y' \in [\pm N] \\ x+y+y'=z^*}} \chi(y-y', x, n, \dots, n)^{\otimes (s-1)} \cdot \psi^*(x, y, y', n) \cdot \psi_{x,y}(n) \overline{\psi_{x,y'}(n)} \cdot b(y, n) \overline{b(y', n)} \cdot b(x+y, n) \overline{b(x+y', n)} \cdot b(x, y) \overline{b(x, y')} \right\|_{\infty} \geq 1/M(\delta).$$

This implies that

$$\left\| \mathbb{E}_{n \in [N]} \mathbb{E}_{\substack{x, y, y' \in [\pm N] \\ x+y+y'=z^*}} \chi(y-y', z^*-y-y', n, \dots, n)^{\otimes (s-1)} \cdot \psi^*(z^*-y-y', y, y', n) \cdot \psi_{x,y}(n) \overline{\psi_{x,y'}(n)} \cdot b(y, n) \overline{b(y', n)} \cdot b(z^*-y', n) \overline{b(z^*-y, n)} \cdot b(z^*-y-y', y) \overline{b(z^*-y-y', y')} \right\|_{\infty} \geq 1/M(\delta).$$

By applying the first item of Lemma C.3, Lemma C.4, Lemma C.2, and Lemma C.1 we have that

$$\chi(y - y', z^* - y - y', n, \dots, n)^{\otimes(s-1)}$$

and

$$\chi(y', y', n, \dots, n)^{\otimes(s-1)} \chi(y', y, n, \dots, n)^{\otimes(s-1)} \overline{\chi(y, y, n, \dots, n)^{\otimes(s-1)}} \overline{\chi(y, y', n, \dots, n)^{\otimes(s-1)}}$$

are $(M(\delta), M(\delta), d(\delta))$ equivalent for degree $(s-1)$. Thus by Lemma 7.4 and letting $\tilde{\psi}$ denote a degree $(s-1)$ nilsequence in y, y', n of complexity $(M(\delta), d(\delta))$ we have that

$$\begin{aligned} & \|\mathbb{E}_{n \in [N]} \mathbb{E}_{\substack{y, y' \in [\pm N] \\ |z^* - y - y'| \leq N}} \chi(y', y', n, \dots, n)^{\otimes(s-1)} \chi(y', y, n, \dots, n)^{\otimes(s-1)} \\ & \overline{\chi(y, y, n, \dots, n)^{\otimes(s-1)}} \overline{\chi(y, y', n, \dots, n)^{\otimes(s-1)}} \cdot \tilde{\psi}(y, y', n) \cdot \psi_{z^* - y - y', y}(n) \overline{\psi_{z^* - y - y', y'}(n)} \\ & \cdot b(y, n) \overline{b(y', n)} \cdot b(z^* - y', n) \overline{b(z^* - y, n)} \cdot b(z^* - y - y', y) \cdot \overline{b(z^* - y - y', y')} \|_{\infty} \geq 1/M(\delta). \end{aligned}$$

Here we have “folded” in $\psi^*(z^* - y - y', y, y', n)$ via Lemma C.2 in $\tilde{\psi}$. We may collapse various 1-bounded functions (and pass to the coordinates of $\chi(y', y', n, \dots, n)^{\otimes(s-1)}$ and $\overline{\chi(y, y, n, \dots, n)^{\otimes(s-1)}}$ which achieve the L^∞ norm) and obtain

$$\begin{aligned} & \|\mathbb{E}_{n \in [N]} \mathbb{E}_{\substack{y, y' \in [\pm N] \\ |z^* - y - y'| \leq N}} \chi(y', y, n, \dots, n)^{\otimes(s-1)} \overline{\chi(y, y', n, \dots, n)^{\otimes(s-1)}} \\ & \cdot \tilde{\psi}(y, y', n) \cdot \psi_{y, y'}(n) \cdot b(y, n) b(y', n) b(y, y') \|_{\infty} \geq 1/M(\delta); \end{aligned}$$

here the $\tilde{\psi}_{y, y'}(n)$ are degree $(s-3)$ nilsequences of complexity $(M(\delta), d(\delta))$. Furthermore as $\tilde{\psi}(y, y', n)$ is a degree $(s-1)$ nilsequence, we have that it is a multidegree $(s-1, s-1, s-3) \cup (1, 0, s-1) \cup (0, 1, s-1)$ nilsequence. Therefore by Lemma C.6 and Lemma C.2, we may remove $\tilde{\psi}$ at the cost of adjusting b and $\psi_{y, y'}$ to obtain

$$\begin{aligned} & \|\mathbb{E}_{n \in [N]} \mathbb{E}_{\substack{y, y' \in [\pm N] \\ |z^* - y - y'| \leq N}} \chi(y', y, n, \dots, n)^{\otimes(s-1)} \overline{\chi(y, y', n, \dots, n)^{\otimes(s-1)}} \\ & \cdot \psi_{y, y'}(n) \cdot b(y, n) b(y', n) b(y, y') \|_{\infty} \geq 1/M(\delta). \end{aligned}$$

This implies that

$$\begin{aligned} & \|\mathbb{E}_{n \in [N]} \mathbb{E}_{y, y' \in [\pm N]} \chi(y', y, n, \dots, n)^{\otimes(s-1)} \overline{\chi(y, y', n, \dots, n)^{\otimes(s-1)}} \\ & \cdot \psi_{y, y'}(n) \cdot b(y, n) b(y', n) b(y, y') \cdot \mathbb{1}_{|z^* - y - y'| \leq N} \|_{\infty} \geq 1/M(\delta). \end{aligned}$$

as $\mathbb{P}[\mathbb{1}_{|z^* - y - y'| \leq N}] \gtrsim \rho^2$. Note that the final indicator may be absorbed into $b(y, y')$ to obtain

$$\begin{aligned} & \|\mathbb{E}_{n \in [N]} \mathbb{E}_{y, y' \in [\pm N]} \chi(y', y, n, \dots, n)^{\otimes(s-1)} \overline{\chi(y, y', n, \dots, n)^{\otimes(s-1)}} \\ & \cdot \psi_{y, y'}(n) \cdot b(y, n) b(y', n) b(y, y') \|_{\infty} \geq 1/M(\delta). \end{aligned}$$

Define $G(y, y', n) = \chi(y, y', n, \dots, n)^{\otimes(s-1)} \otimes \overline{\chi(y', y, n, \dots, n)^{\otimes(s-1)}}$ and we have

$$\|\mathbb{E}_{n \in [N], y, y' \in [\pm N]} G(y, y', n) \cdot \overline{\psi_{y, y'}(n)} \cdot b(y, n) b(y', n) b(y, y') \|_{\infty} \geq 1/M(\delta).$$

Applying Cauchy–Schwarz in n , then y , and then y' (analogously to as in Lemma 7.2) we may remove the bounded functions b and we have

$$\begin{aligned} & \|\mathbb{E}_{n_1, n_2 \in [N], y_1, y_2, y'_1, y'_2 \in [\pm N]} \bigotimes_{\varepsilon \in \{1, 2\}^3} \mathcal{C}^{|\varepsilon| - 1} G(y_{\varepsilon_1}, y'_{\varepsilon_2}, n_{\varepsilon_3}) \cdot \psi_{y_1, y_2, y'_1, y'_2}(n_1) \cdot \overline{\psi_{y_1, y_2, y'_1, y'_2}(n_2)} \|_{\infty} \\ & \geq 1/M(\delta), \end{aligned}$$

where \mathcal{C} denotes conjugation and the $\psi_{y_1, y_2, y'_1, y'_2}(n_i)$ are degree at most $(s-3)$ nilsequences of complexity $(M(\delta), d(\delta))$. Applying Pigeonhole in n_2, y_2, y'_2 and applying Lemma C.2 to specialize variables, reindexing n_1, y_1, y'_1 to n, y, y' , and taking the maximal coordinate we have

$$\|\mathbb{E}_{n \in [N], y, y' \in [\pm N]} G(y, y', n) \cdot \psi_1(y, n) \cdot \psi_2(y', n) \cdot \psi_{y, y'}(n)\|_\infty \geq 1/M(\delta).$$

Here $\psi_{y, y'}$ is degree at most $(s-3)$ in n while $\psi_1(y, n)$ and $\psi_2(y, n)$ are multidegree $(1, s-2)$ and all have complexity $(M(\delta), d(\delta))$. Finally by the triangle inequality we have

$$\mathbb{E}_{y, y' \in [\pm N]} \|\mathbb{E}_{n \in [N]} G(y, y', n) \cdot \psi_1(y, n) \cdot \psi_2(y', n) \cdot \psi_{y, y'}(n)\|_\infty \geq 1/M(\delta).$$

Step 3: Converse of the inverse theorem and polarization. By the converse of the inverse theorem, see Lemma B.5, we have that

$$\mathbb{E}_{y, y' \in [N]} \|G(y, y', \cdot) \psi_1(y, \cdot) \psi_2(y', \cdot)\|_{U^{s-2}[N]}^{2^{s-2}} \geq 1/M(\delta).$$

Expanding out the definition of the U^{s-2} -norm, we find that

$$\left\| \mathbb{E}_{y, y' \in [\pm N]} \mathbb{E}_{n \in [N], h_1, \dots, h_{s-2} \in [\pm N]} \bigotimes_{\varepsilon \in \{0,1\}^{s-2}} \left(\mathcal{C}^{|\varepsilon|+s} (G(y, y', n + \varepsilon \cdot \vec{h}) \cdot \psi_1(y, n + \varepsilon \cdot \vec{h}) \cdot \psi_2(y', n + \varepsilon \cdot \vec{h})) \cdot \mathbb{1}_{n + \varepsilon \cdot \vec{h} \in [N]} \right) \right\|_\infty \geq 1/M(\delta).$$

The crucial point is that by repeatedly applying Lemma C.4, Lemma C.2, and Lemma C.1 we have that

$$\bigotimes_{\varepsilon \in \{0,1\}^{s-2}} \mathcal{C}^{|\varepsilon|+s} (G(y, y', n + \varepsilon \cdot \vec{h}))$$

and

$$\chi(y, y', h_1, \dots, h_{s-2})^{\otimes(s-1)!} \cdot \overline{\chi(y', y, h_1, \dots, h_{s-2})}^{\otimes(s-1)!}$$

are $(M(\delta), M(\delta), d(\delta))$ -equivalent for degree $(s-1)$. Therefore by Lemma 7.4, there exists a nilsequence $\tilde{\psi}$ of degree $(s-1)$ and complexity $(M(\delta), d(\delta))$ such that

$$\left\| \mathbb{E}_{y, y' \in [\pm N]} \mathbb{E}_{n \in [N], h_1, \dots, h_{s-2} \in [\pm N]} \chi(y, y', h_1, \dots, h_{s-2})^{\otimes(s-1)!} \cdot \overline{\chi(y', y, h_1, \dots, h_{s-2})}^{\otimes(s-1)!} \cdot \tilde{\psi}(n, y, y', h_1, \dots, h_{s-2}) \cdot \mathbb{1}_{n + \varepsilon \cdot \vec{h} \in [N]} \right\|_\infty \geq 1/M(\delta).$$

Via Fourier expansion (a multidimensional version of the argument in Lemma 7.1), we may fold in $\mathbb{1}_{n + \varepsilon \cdot \vec{h} \in [N]}$ into $\tilde{\psi}(n, y, y', h_1, \dots, h_{s-2})$.⁸ We reduce to

$$\left\| \mathbb{E}_{y, y' \in [\pm N]} \mathbb{E}_{n \in [N], h_1, \dots, h_{s-2} \in [\pm N]} \chi(y, y', h_1, \dots, h_{s-2})^{\otimes(s-1)!} \cdot \overline{\chi(y', y, h_1, \dots, h_{s-2})}^{\otimes(s-1)!} \cdot \tilde{\psi}(n, y, y', h_1, \dots, h_{s-2}) \right\|_\infty \geq 1/M(\delta).$$

Applying Pigeonhole in n and applying the first item of Lemma C.2, we reduce to

$$(12.2) \quad \left\| \mathbb{E}_{y, y' \in [\pm N]} \mathbb{E}_{h_1, \dots, h_{s-2} \in [\pm N]} \chi(y, y', h_1, \dots, h_{s-2})^{\otimes(s-1)!} \cdot \overline{\chi(y', y, h_1, \dots, h_{s-2})}^{\otimes(s-1)!} \otimes \tilde{\psi}(y, y', h_1, \dots, h_{s-2}) \right\|_\infty \geq 1/M(\delta);$$

once again we have abusively updated $\tilde{\psi}$, which has degree $(s-1)$.

Step 4: Invoking equidistribution theory. This is the unique moment we have the ability to apply equidistribution theory; up to this point we have been applying “elementary” facts regarding nilsequences. Let

$$\chi(y, y', h_1, \dots, h_{s-2}) = F(g(y, y', h_1, \dots, h_{s-2})\Gamma)$$

⁸To be precise, we convolve $\mathbb{1}_{n + \varepsilon \cdot \vec{h} \in [N]}$ with $\mathbb{1}_{|n| \leq \rho N} \cdot \prod_{i=1}^{s-2} \mathbb{1}_{|h_i| \leq \rho N}$ where $\rho = 1/M(\delta)$ is sufficiently small. This function has the necessary Fourier decay to apply the analysis in Lemma 7.1

and let ξ denote the vertical $G_{(1,\dots,1)}$ frequency of F on the multidegree $(1, \dots, 1)$ nilmanifold G/Γ . We write

$$\tilde{\psi}(y, y', h_1, \dots, h_{s-2}) = \tilde{F}(g^*(y, y', h_1, \dots, h_{s-2})\Gamma')$$

on the multidegree $(s-1)$ nilmanifold G'/Γ' . Note that

$$(g(y, y', h_1, \dots, h_{s-2}), g(y', y, h_1, \dots, h_{s-2}), g^*(y, y_1, h_1, \dots, h_{s-2}))$$

may be viewed as a polynomial sequence on $G \times G \times G'$ where G' is given a degree $(s-1)$ filtration. $G \times G \times G'$ is given a degree s filtration where the t -th group is

$$(G \times G \times G')_t = \bigvee_{|\vec{i}|=t} G_{\vec{i}} \times \bigvee_{|\vec{i}|=t} G_{\vec{i}} \times (G')_t.$$

Note that $F \otimes \overline{F} \otimes \tilde{F}$ has $(G \times G \times G')_s$ -vertical frequency $\xi' = (\xi, -\xi, 0)$, noting that $(G')_s = \text{Id}_{G'}$. By applying Corollary 5.5 with (12.2) to

$$F(g(y, y', h_1, \dots, h_{s-2})\Gamma)^{\otimes(s-1)!} \otimes \overline{F(g(y', y, h_1, \dots, h_{s-2})\Gamma)}^{\otimes(s-1)!} \cdot \tilde{F}(g^*(y, y', h_1, \dots, h_{s-2})\Gamma'),$$

and restricting the factorization to $G \times G$, we have

$$\begin{aligned} & (g(y, y', h_1, \dots, h_{s-2}), g(y', y, h_1, \dots, h_{s-2})) \\ &= \varepsilon(y, y_1, h_1, \dots, h_{s-2}) \cdot g^{\text{Output}}(y, y_1, h_1, \dots, h_{s-2}) \cdot \gamma(y, y_1, h_1, \dots, h_{s-2}), \end{aligned}$$

where

- g^{Output} lives in an $M(\delta)$ -rational subgroup H such that $\xi'(H \cap (G \times G)_s) = 0$;
- γ is an $M(\delta)$ -rational polynomial sequence;
- ε is $(M(\delta), N)$ -smooth.

Note that when apply Corollary 5.5 the vertical frequency of the function we have is $(s-1)! \cdot \xi'$ and we obtain $(s-1)! \xi'(H \cap (G \times G)_s) = 0$; we may divide by $(s-1)!$ to obtain the above. Additionally, we have implicitly used that ξ' is trivial in the G' part and abuse notation to descend ξ' to $G \times G$.

Let $F^* = F \otimes \overline{F}$ and note that therefore

$$\chi(h, n, \dots, n) \otimes \overline{\chi(n, h, n, \dots, n)} = F^*(\varepsilon(h, n, \dots, n)g^{\text{Output}}(h, n, \dots, n) \cdot \gamma(h, n, \dots, n)(\Gamma \times \Gamma)).$$

Step 5: The finishing touch. We now recall from (12.1) that

$$\mathbb{E}_{h \in [N]} \|\mathbb{E}_{n \in [N]} \Delta_h f(n) \otimes \chi(h, n, \dots, n) \cdot \psi(n) \cdot \psi_h(n)\|_{\infty} \geq 1/M(\delta);$$

here we have restricted to a coordinate of ψ_h and we treat it as a degree $(s-2)$ nilsequence (rather than using the nilcharacter). By applying Pigeonhole there exist $q, q' \in [s]$ such that

$$\mathbb{E}_{h \in [N/s]} \|\mathbb{E}_{n \in [N/s]} \Delta_{sh+q'} f(sn+q) \otimes \chi(sh+q', sn+q, \dots, sn+q) \cdot \psi(sn+q) \cdot \psi_h(sn+q)\|_{\infty} \geq 1/M(\delta).$$

By Lemma C.3, we have that

$$\chi(sh+q', sn+q, \dots, sn+q) \text{ and } \chi(sh, sn, \dots, sn)$$

are $(M(\delta), M(\delta), d(\delta))$ -equivalent for degree $(s-1)$. Applying Lemma C.6 (splitting) and adjusting ψ, ψ_h , we may instead assume that

$$\mathbb{E}_{h \in [N/s]} \|\mathbb{E}_{n \in [N/s]} \Delta_{sh+q'} f(sn+q) \otimes \chi(sh, sn, \dots, sn) \cdot \psi(n) \cdot \psi_h(n)\|_{\infty} \geq 1/M(\delta)$$

for ψ of degree $(s-1)$ and ψ_h of degree $(s-2)$. Now define

$$T(h, n) := \chi(n+h, \dots, n+h) \otimes \overline{\chi(h, n, \dots, n)} \otimes \overline{\chi(n, h, \dots, n)}^{\otimes(s-1)} \otimes \overline{\chi(n, n, \dots, n)}.$$

Since this is a nilcharacter, we automatically know

$$\mathbb{E}_{h \in [N/s]} \left\| \mathbb{E}_{n \in [N/s]} \Delta_{sh+q'} f(sn+q) \otimes \chi(sh, sn, \dots, sn) \cdot \psi(n) \cdot \psi_h(n) \right. \\ \left. \otimes T(h, n)^{\otimes s^{s-1}} \otimes \overline{T(h, n)^{\otimes s^{s-1}}} \right\|_{\infty} \geq 1/M(\delta).$$

We define

$$\tilde{f}_1(n) = f(sn+q) \cdot \overline{\chi(n, \dots, n)^{\otimes s^{s-1}}}, \\ \tilde{f}_2(n+h) = \overline{f(s(n+h)+q+q')} \cdot \chi(n+h, \dots, n+h)^{\otimes s^{s-1}},$$

which yields

$$\mathbb{E}_{h \in [N/s]} \left\| \mathbb{E}_{n \in [N/s]} \tilde{f}_1(n) \otimes \tilde{f}_2(n+h) \otimes \chi(sh, sn, \dots, sn) \cdot \psi(n) \cdot \psi_h(n) \right. \\ \left. \otimes \overline{(\chi(h, n, \dots, n) \otimes \chi(n, h, n, \dots, n)^{\otimes (s-1)})^{\otimes s^{s-1}}} \otimes \overline{T(h, n)^{\otimes s^{s-1}}} \right\|_{\infty} \geq 1/M(\delta).$$

By applications of Lemma C.4, Lemma C.2, and Lemma C.1 we have that $T(h, n)$ and

$$\bigotimes_{k=1}^{s-1} \chi(h, h, \dots, h, n, \dots, n)^{\binom{s-1}{k}} \otimes \bigotimes_{k=2}^{s-1} \chi(n, h, \dots, h, n, \dots, n)^{\binom{s-1}{k}}$$

are $(M(\delta), M(\delta), d(\delta))$ -equivalent for degree $(s-1)$. (There are $k+1$ many h 's in the first term and k many h 's in the second term.) Applying Lemma C.6, we may approximate each coordinate as a sum of products of multidegree $(s-1, s-2)$ and $(0, s-1)$ nilsequences in variables (h, n) . Furthermore, by the second item of Lemma C.2 this new nilsequence is of similar type. So, folding everything into $\psi(n)$ of degree $(s-1)$ and the $\psi_h(n)$ of degree $(s-2)$, we find

$$\mathbb{E}_{h \in [N/s]} \left\| \mathbb{E}_{n \in [N/s]} \tilde{f}_1(n) \otimes \tilde{f}_2(n+h) \otimes \chi(sh, sn, \dots, sn) \cdot \psi(n) \cdot \psi_h(n) \right. \\ \left. \otimes \overline{(\chi(h, n, \dots, n) \otimes \chi(n, h, \dots, n)^{\otimes (s-1)})^{\otimes s^{s-1}}} \right\|_{\infty} \geq 1/M(\delta).$$

Furthermore note by Lemma C.3 that

$$\chi(sh, sn, \dots, sn) \text{ and } \chi(h, n, \dots, n)^{\otimes s^s}$$

are $(M(\delta), M(\delta), d(\delta))$ -equivalent for degree $(s-1)$. Applying Lemma 7.4 and Lemma C.6 and adjusting ψ and ψ_h yet again we have

$$\mathbb{E}_{h \in [N/s]} \left\| \mathbb{E}_{n \in [N/s]} \tilde{f}_1(n) \otimes \tilde{f}_2(n+h) \otimes \chi(h, n, \dots, n)^{\otimes s^s} \cdot \psi(n) \cdot \psi_h(n) \right. \\ \left. \otimes \overline{(\chi(h, n, \dots, n) \otimes \chi(n, h, \dots, n)^{\otimes (s-1)})^{\otimes s^{s-1}}} \right\|_{\infty} \geq 1/M(\delta).$$

Now by Lemma C.3 we have that

$$\chi(h, n, \dots, n)^{\otimes s^s} \otimes \overline{\chi(h, n, \dots, n)^{\otimes s^{s-1}}} \text{ and } \chi(h, n, \dots, n)^{\otimes (s-1) \cdot s^{s-1}}$$

are $(M(\delta), M(\delta), d(\delta))$ -equivalent for degree $(s-1)$. Thus applying Lemma 7.4 and Lemma C.6 and adjusting ψ and ψ_h once again we have

$$\mathbb{E}_{h \in [N/s]} \left\| \mathbb{E}_{n \in [N/s]} \tilde{f}_1(n) \otimes \tilde{f}_2(n+h) \cdot \psi(n) \cdot \psi_h(n) \otimes (\chi(h, n, \dots, n) \otimes \overline{\chi(n, h, \dots, n)^{\otimes (s-1) \cdot s^{s-1}}}) \right\|_{\infty} \\ \geq 1/M(\delta).$$

This is finally where we may apply our earlier factorization for $\chi(h, n, \dots, n) \otimes \overline{\chi(n, h, \dots, n)}$. Recall that

$$\chi(h, n, \dots, n) \otimes \overline{\chi(n, h, n, \dots, n)} = F^*(\varepsilon(h, n, \dots, n)g^{\text{Output}}(h, n, \dots, n) \cdot \gamma(h, n, \dots, n)(\Gamma \times \Gamma))$$

where γ is $M(\delta)$ -periodic and ε is $(M(\delta), N)$ -smooth. Let Q denote the period of γ (i.e., changing any argument by a multiple of Q keeps its $\Gamma \times \Gamma$ coset the same) and take $\rho = \exp(-\log(1/\delta)^{O_s(1)})$ where the implicit constant is sufficiently large. Break $[N/s]$ into arithmetic progressions of length roughly ρN and common difference Q ; call these $\mathcal{P}_1, \dots, \mathcal{P}_\ell$. There exist $\varepsilon_{\mathcal{P}_{i,h}}$ and $\gamma_{\mathcal{P}_{i,h}}$ such that

$$\begin{aligned} \mathbb{E}_{h \in [N/s]} \left\| \mathbb{E}_{n \in [N/s]} \sum_{i=1}^{\ell} \mathbb{1}_{n \in \mathcal{P}_i} \tilde{f}_1(n) \otimes \tilde{f}_2(n+h) \cdot \psi(n) \cdot \psi_h(n) \right. \\ \left. \otimes (F^*(\varepsilon_{\mathcal{P}_{i,h}} \gamma_{\mathcal{P}_{i,h}} (\gamma_{\mathcal{P}_{i,h}}^{-1} g^{\text{Output}}(h, n, \dots, n) \gamma_{\mathcal{P}_{i,h}}))(\Gamma \times \Gamma)))^{\otimes (s-1)s^{s-1}} \right\|_{\infty} \geq 1/M(\delta) \end{aligned}$$

where $d_{G \times G}(\varepsilon_{\mathcal{P}_{i,h}}, \text{id}_{G \times G}) + d_{G \times G}(\gamma_{\mathcal{P}_{i,h}}, \text{id}_{G \times G}) \leq \exp(\log(1/\delta)^{O_s(1)})$ and $\gamma_{\mathcal{P}_{i,h}}$ is $\exp(\log(1/\delta)^{O_s(1)})$ -rational.

By Pigeonhole, there exists an index i such that

$$\begin{aligned} \mathbb{E}_{h \in [N/s]} \left\| \mathbb{E}_{n \in [N/s]} \mathbb{1}_{n \in \mathcal{P}_i} \cdot \tilde{f}_1(n) \otimes \tilde{f}_2(n+h) \cdot \psi(n) \cdot \psi_h(n) \right. \\ \left. \otimes (F^*(\varepsilon_{\mathcal{P}_{i,h}} \gamma_{\mathcal{P}_{i,h}} (\gamma_{\mathcal{P}_{i,h}}^{-1} g^{\text{Output}}(h, n, \dots, n) \cdot \gamma_{\mathcal{P}_{i,h}}))(\Gamma \times \Gamma)))^{\otimes (s-1)s^{s-1}} \right\|_{\infty} \geq 1/M(\delta). \end{aligned}$$

As $\gamma_{\mathcal{P}_{i,h}}$ is $\exp(\log(1/\delta)^{O_s(1)})$ -rational and bounded, it takes on only $\exp(\log(1/\delta)^{O_s(1)})$ possible values. Thus by Pigeonhole, there is $\gamma \in \Gamma \times \Gamma$ such that $d_{G \times G}(\gamma, \text{id}_{G \times G}) \leq \exp(\log(1/\delta)^{O_s(1)})$ and γ is $\exp(\log(1/\delta)^{O_s(1)})$ -rational such that

$$\begin{aligned} \mathbb{E}_{h \in [N/s]} \left\| \mathbb{E}_{n \in [N/s]} \mathbb{1}_{n \in \mathcal{P}_i} \tilde{f}_1(n) \otimes \tilde{f}_2(n+h) \cdot \psi(n) \cdot \psi_h(n) \right. \\ \left. \otimes (F^*(\varepsilon_{\mathcal{P}_{i,h}} \gamma g^{\text{Conj}}(h, n, \dots, n)(\Gamma \times \Gamma)))^{\otimes (s-1)s^{s-1}} \right\|_{\infty} \geq 1/M(\delta), \end{aligned}$$

where $g^{\text{Conj}} = \gamma^{-1} g^{\text{Output}} \gamma$. Finally, rounding $\varepsilon_{\mathcal{P}_{i,h}} \gamma$ to a $\exp(-\log(1/\delta)^{O_s(1)})$ -net and noting it is $\exp(\log(1/\delta)^{O_s(1)})$ -bounded, there exists ε such that

$$\begin{aligned} \mathbb{E}_{h \in [N/s]} \left\| \mathbb{E}_{n \in [N/s]} \mathbb{1}_{n \in \mathcal{P}_i} \cdot \tilde{f}_1(n) \otimes \tilde{f}_2(n+h) \cdot \psi(n) \cdot \psi_h(n) \right. \\ \left. \otimes (F^*(\varepsilon g^{\text{Conj}}(h, n, \dots, n)(\Gamma \times \Gamma)))^{\otimes (s-1)s^{s-1}} \right\|_{\infty} \geq 1/M(\delta) \end{aligned}$$

and $d_{G \times G}(\varepsilon, \text{id}_{G \times G}) \leq \exp(\log(1/\delta)^{O_s(1)})$, as long as ρ was chosen small enough.

By Lemma 7.1, there exists Θ_h such that

$$\begin{aligned} \mathbb{E}_{h \in [N/s]} \left\| \mathbb{E}_{n \in [N/s]} e(\Theta_h n) \cdot \tilde{f}_1(n) \otimes \tilde{f}_2(n+h) \cdot \psi(n) \cdot \psi_h(n) \right. \\ \left. \otimes (F^*(\varepsilon g^{\text{Conj}}(h, n, \dots, n)(\Gamma \times \Gamma)))^{\otimes (s-1)s^{s-1}} \right\|_{\infty} \geq 1/M(\delta). \end{aligned}$$

As $(s - 2) \geq 1$, we may absorb $e(\Theta_h n)$ into $\psi_h(n)$ and obtain

$$\mathbb{E}_{h \in [N/s]} \left\| \mathbb{E}_{n \in [N/s]} \tilde{f}_1(n) \otimes \tilde{f}_2(n+h) \cdot \psi(n) \cdot \psi_h(n) \right. \\ \left. \otimes (F^*(\varepsilon g^{\text{Conj}}(h, n, \dots, n)(\Gamma \times \Gamma)))^{\otimes (s-1)s^{s-1}} \right\|_{\infty} \geq 1/M(\delta).$$

Replacing F^* with $F^{\text{Final}}(g) = F^*(\varepsilon g \Gamma)$ and writing $g^{\text{Final}}(h, n) = g^{\text{Conj}}(h, n, \dots, n)$, we have

$$\mathbb{E}_{h \in [N/s]} \left\| \mathbb{E}_{n \in [N/s]} \tilde{f}_1(n) \otimes \tilde{f}_2(n+h) \cdot \psi(n) \cdot \psi_h(n) \right. \\ \left. \otimes (F^{\text{Final}}(g^{\text{Final}}(h, n)(\Gamma \times \Gamma)))^{\otimes (s-1)s^{s-1}} \right\|_{\infty} \geq 1/M(\delta).$$

Now $g^{\text{Final}}(h, n)$ takes values in $\gamma^{-1}H\gamma$ such that $\xi'(\gamma^{-1}H\gamma \cap (G \times G)_s) = 0$. The key point is to note that F^{Final} is right-invariant under $(\gamma^{-1}H\gamma) \cap (G \times G)_s$ since it has $(G \times G)_s$ -vertical frequency ξ' . Note that $\gamma^{-1}H\gamma$ has complexity bounded by $M(\delta)$ due to [35, Lemma B.15]. Furthermore F^{Final} is $M(\delta)$ -Lipschitz on $\gamma^{-1}H\gamma$ by [35, Lemma B.9, B.15]. Taking the quotient by $(\gamma^{-1}H\gamma) \cap (G \times G)_s$ gives that each coordinate of $(F^{\text{Final}}(g^{\text{Final}}(h, n)(\Gamma \times \Gamma)))^{\otimes s^{s-1}}$ may be realized a complexity $(M(\delta), d(\delta))$ nilsequence of degree $(s - 1)$.

Applying Pigeonhole in the coordinates of $(F^{\text{Final}})^{\otimes s^{s-1}}$ and then Lemma C.6 to approximate as a sum of products of multidegree $(s - 1, s - 2)$ and $(0, s - 1)$ nilsequences in variables (h, n) . So again folding everything into $\psi(n)$ of degree $(s - 1)$ and the $\psi_h(n)$ of degree $(s - 2)$, we find

$$\mathbb{E}_{h \in [N/s]} \left\| \mathbb{E}_{n \in [N/s]} \tilde{f}_1(n) \otimes \tilde{f}_2(n+h) \cdot \psi(n) \cdot \psi_h(n) \right\|_{\infty} \geq 1/M(\delta)$$

The functions \tilde{f}_1 and \tilde{f}_2 are vector-valued, but by Pigeonhole there exist coordinates j_1, j_2 are coordinates of the vectors \tilde{f}_1 and \tilde{f}_2 such that

$$\mathbb{E}_{h \in [N/s]} \left| \mathbb{E}_{n \in [N/s]} \tilde{f}_{1,j_1}(n) \tilde{f}_{2,j_2}(n+h) \cdot \psi(n) \cdot \psi_h(n) \right| \geq 1/M(\delta).$$

Since $\psi_h(n)$ is a nilsequence of degree $(s - 2)$ and complexity $(M(\delta), d(\delta))$, by the converse of the inverse theorem (see Lemma B.5) we have that

$$\mathbb{E}_{h \in [N/s]} \left\| \tilde{f}_{1,j_1}(\cdot) \tilde{f}_{2,j_2}(\cdot + h) \psi(\cdot) \right\|_{U^{s-1}[N/s]}^{2s-1} \geq 1/M(\delta).$$

By the Gowers–Cauchy–Schwarz inequality (e.g. [12, Lemma 3.8]), we have that

$$\left\| \mathbb{E}_{n \in [N/s]} \tilde{f}_{1,j_1}(n) \psi(n) \right\|_{U^s[N/s]} \geq 1/M(\delta).$$

By induction, there is a nilsequence $\Theta(n)$ of degree $(s - 1)$ and complexity $(M(\delta), d(\delta))$ such that

$$\left| \mathbb{E}_{n \in [N/s]} \tilde{f}_{1,j_1}(n) \psi(n) \Theta(n) \right| \geq 1/M(\delta).$$

Now recall that

$$\tilde{f}_1(n) = f(sn + q) \cdot \overline{\chi(n, \dots, n)}^{\otimes s^{s-1}}.$$

Each coordinate of $\overline{\chi(n, \dots, n)}^{\otimes s^{s-1}}$ is a degree s nilsequence of complexity $(M(\delta), d(\delta))$; say j_1 -th coordinate is $\Theta'(n)$ and thus we have

$$\left| \mathbb{E}_{n \in [N/s]} f(sn + q) \Theta'(n) \psi(n) \Theta(n) \right| \geq 1/M(\delta).$$

This is equivalent to

$$\left| \mathbb{E}_{n \in [N]} \mathbb{1}[n \equiv q \pmod{s}] f(n) \Theta'((n - q)/s) \psi((n - q)/s) \Theta((n - q)/s) \right| \geq 1/M(\delta).$$

Note the condition

$$\mathbb{1}[n \equiv q \pmod{s}] = s^{-1} \sum_{j=0}^{s-1} e\left(\frac{j \cdot (n - q)}{s}\right)$$

and thus there j such that

$$|\mathbb{E}_{n \in [N]} f(n) \Theta'((n-q)/s) \psi((n-q)/s) \Theta((n-q)/s) e(jn/s)| \geq 1/M(\delta).$$

The desired nilsequence is then

$$\overline{\Theta'((n-q)/s) \psi((n-q)/s) \Theta((n-q)/s) e(jn/s)}$$

which is seen to have degree s and complexity $(M(\delta), d(\delta))$. We have finally won. \square

APPENDIX A. ON APPROXIMATE HOMOMORPHISMS

In this section, we give a number of basic results regarding approximate homomorphisms. The results in this section are, by now, well known consequences of work of Sanders [45]. The proof we give is essentially that in [36], modulo being forced to deal with slight error terms and operating over \mathbb{Z} . We dispose of these error terms via a rounding trick of Green, Tao, and Ziegler [27, Appendix C].

Lemma A.1. *Fix $\delta \in (0, 1/2)$, let $H_1, H_2, H_3, H_4 \subseteq [N]$ and let functions $f_i: H_i \rightarrow \mathbb{R}^d$ be such that there are at least δN^3 additive tuples $h_1 + h_2 = h_3 + h_4$ with*

$$\|(f_1(h_1) + f_2(h_2) - f_3(h_3) - f_4(h_4))_j\|_{\mathbb{R}/\mathbb{Z}} \leq \varepsilon_j$$

for all $1 \leq j \leq d$. Then there exists $H'_1 \subseteq H_1$ with $|H'_1| \geq \exp(-(d \log(1/\delta))^{O(1)})N$ such that

$$\left\| \left(f_1(h) - \sum_{i=1}^{d'} a_i \{ \alpha_i h \} - b \right)_j \right\|_{\mathbb{R}/\mathbb{Z}} \leq \varepsilon_j$$

for all $h \in H'_1$, for appropriate choices of $d' \leq (d \log(1/\delta))^{O(1)}$, $a_i, b \in \mathbb{R}^d$, and $\alpha_i \in (1/N')\mathbb{Z}$ where N' is a prime between $100N$ and $200N$.

We deduce the result from the following variant which is the same statement modulo not having an error term.

Lemma A.2. *Fix $\delta \in (0, 1/2)$. Let $H_1, H_2, H_3, H_4 \subseteq [N]$ and $f_i: H_i \rightarrow \mathbb{R}^d$ be such that there are at least δN^3 additive tuples $h_1 + h_2 = h_3 + h_4$ with*

$$f_1(h_1) + f_2(h_2) - f_3(h_3) - f_4(h_4) \in \mathbb{Z}^d.$$

Then there exists $H'_1 \subseteq H_1$ with $|H'_1| \geq \exp(-\log(1/\delta)^{O(1)})N$ such that

$$f_1(h) - \sum_{i=1}^{d'} a_i \{ \alpha_i h \} - b \in \mathbb{Z}^d$$

for all $h \in H'_1$, for appropriate choices of $d' \leq \log(1/\delta)^{O(1)}$, $a_i, b \in \mathbb{R}^d$, and $\alpha_i \in (1/N')\mathbb{Z}$ where N' is a prime between $100N$ and $200N$.

We briefly give the deduction, and then in the sequel focus on Lemma A.2.

Proof of Lemma A.1 given Lemma A.2. Round each value of f_i to the nearest point in the lattice $(\varepsilon_1 \mathbb{Z}, \dots, \varepsilon_d \mathbb{Z})$ to form \tilde{f}_i (breaking ties arbitrarily). We have that

$$\|(\tilde{f}_1(h_1) + \tilde{f}_2(h_2) - \tilde{f}_3(h_3) - \tilde{f}_4(h_4))_j\|_{\mathbb{R}/\mathbb{Z}} \leq 5\varepsilon_j$$

for at least δN^4 additive tuples.

Note however that

$$\tilde{f}_1(h_1) + \tilde{f}_2(h_2) - \tilde{f}_3(h_3) - \tilde{f}_4(h_4) \in (\varepsilon_1 \mathbb{Z}, \dots, \varepsilon_d \mathbb{Z})$$

and that there are at most 11^d lattice points in $(\varepsilon_1\mathbb{Z}, \dots, \varepsilon_d\mathbb{Z})$ which are at most $5\varepsilon_j$ in the j -th direction from the origin in all d directions. Thus there is a vector $w \in (\varepsilon_1\mathbb{Z}, \dots, \varepsilon_d\mathbb{Z})$ such that

$$\tilde{f}_1(h_1) + \tilde{f}_2(h_2) - \tilde{f}_3(h_3) - \tilde{f}_4(h_4) + w \in \mathbb{Z}^d$$

for at least $11^{-d}\delta N^4$ additive tuples. Applying Lemma A.2 with $\tilde{f}_1, \tilde{f}_2, \tilde{f}_3$, and $\tilde{f}_4 - w$ immediately gives the desired result. \square

We now require the notion of a Bohr set in an abelian group.

Definition A.3. *Given an abelian group G and a set $S \subseteq \widehat{G}$, we define the Bohr set of radius ρ to be*

$$B(S, \rho) := \{x \in G : \|s \cdot x\|_{\mathbb{R}/\mathbb{Z}} \leq \rho \text{ for all } s \in S\}.$$

We first require the fact that the four-fold sumset of a set with small doubling contains a Bohr set of small dimension and large radius. This is an immediate consequence of work of Sanders [45, Theorem 1.1] which produces a large symmetric coset progression and a proposition of Milićević [41, Proposition 27] which produces a Bohr set inside a large symmetric coset progression. This is explicitly [41, Corollary 28].

Lemma A.4. *Let $A \subseteq \mathbb{Z}/N\mathbb{Z}$ be such that $|A| \geq N/K$. Then there exists $S \subseteq \widehat{\mathbb{Z}/N\mathbb{Z}}$ with $|S| \leq \log(2K)^{O(1)}$ and $1/\rho \leq \log(2K)^{O(1)}$ such that $B(S, \rho) \subseteq 2A - 2A$.*

We next require the notion of a Freiman homomorphism.

Definition A.5. *A function $f: A \rightarrow B$ (with A and B being subsets of possibly different abelian groups) is a k -Freiman homomorphism if for all $a_i, a'_i \in A$ satisfying*

$$a_1 + \dots + a_k = a'_1 + \dots + a'_k$$

we have

$$f(a_1) + \dots + f(a_k) = f(a'_1) + \dots + f(a'_k).$$

When k is not specified, we will implicitly have $k = 2$.

We will also require the follow basic lemma which converts the Freiman homomorphism on a Bohr set into a “bracket” linear function on a slightly smaller Bohr set; the proof is a simplification of [18, Proposition 10.8].

Lemma A.6. *Consider $S \subseteq \widehat{\mathbb{Z}/N\mathbb{Z}}$ and $\rho \in (0, 1/4)$ with Freiman homomorphism $f: B(S, \rho) \rightarrow \mathbb{R}/\mathbb{Z}$. Taking $\rho' = \rho \cdot |S|^{-2|S|}$, we have for all $n \in B(S, \rho')$ that*

$$f(n) - \left(\sum_{\alpha_i \in S} a_i \{\alpha_i n\} + \gamma \right) \in \mathbb{Z},$$

for appropriate choices of $a_i, \gamma \in \mathbb{R}$.

Proof. By [18, Proposition 10.5], we have that

$$B(S, \rho \cdot |S|^{-2|S|}) \subseteq P \subseteq B(S, \rho)$$

where P is a proper generalized arithmetic progression $\{\sum_{i=1}^d \ell_i n_i : n_i \in [\pm N_i]\}$ of rank $d \leq |S|$. Furthermore $(\{\alpha \cdot \ell_i\})_{\alpha \in S}$ for $1 \leq i \leq d$ are linearly independent as vectors in \mathbb{R}^S .

Note that for $|n_i| \leq N_i$, we have

$$(A.1) \quad f\left(\sum_{i=1}^d \ell_i n_i\right) - f(0) = \sum_{i=1}^d n_i (f(\ell_i) - f(0)).$$

Furthermore letting $\Phi: B(S, \rho) \rightarrow \mathbb{R}^S$ denote $\Phi(x) = (\{\alpha \cdot x\})_{\alpha \in S}$ we have that

$$\Phi(x) + \Phi(y) = \Phi(x + y);$$

we have used crucially that $\rho < 1/4$ here. Therefore, by a simple inductive argument we see

$$\Phi\left(\sum_{i=1}^d \ell_i n_i\right) = \sum_{i=1}^d n_i \Phi(\ell_i)$$

if $n_i \in [\pm N_i]$ for all $1 \leq i \leq d$.

By the above linear independence, there exists $u_i \in \mathbb{R}^S$ such that $u_i \cdot \Phi(\ell_i) = 1$ and $u_i \cdot \Phi(\ell_j) = 0$ for $j \neq i$. Therefore if $n \in P$ is such that $n = \sum_{i=1}^d \ell_i n_i$, we have that

$$n_i = u_i \cdot \sum_{i=1}^d n_i \Phi(\ell_i) = u_i \cdot \Phi(n) = \sum_{\alpha \in S} (u_i)_\alpha \cdot \{\alpha n\}.$$

The lemma then follows by plugging into (A.1). \square

We now recall the definition of additive energy.

Definition A.7. *Given (finite) subsets A_1, A_2, A_3, A_4 of an abelian group G , define the additive energy $E(A_1, A_2, A_3, A_4)$ to be*

$$E(A_1, A_2, A_3, A_4) = \sum_{x_i \in A_i} \mathbb{1}[x_1 + x_2 = x_3 + x_4]$$

and let $E(A) = E(A, A, A, A)$.

Note that one has the trivial bound $E(A) \leq |A|^3$. Furthermore via a standard Cauchy–Schwarz argument (similar to e.g. [51, Corollary 2.10]) we have

$$E(A_1, A_2, A_3, A_4) \leq \prod_{i=1}^4 E(A_i)^{1/4}.$$

Proof of Lemma A.2. Let $\Gamma_i = \{(h_i, f_i(h_i) \bmod \mathbb{Z}^d) : h_i \in H_i\} \subseteq \mathbb{Z} \times (\mathbb{R}/\mathbb{Z})^d$, which is a graph (i.e., for every $x \in \mathbb{Z}$ there is at most one $y \in (\mathbb{R}/\mathbb{Z})^d$ with $(x, y) \in \Gamma_i$). By assumption we have

$$E(\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4) \geq \delta N^3.$$

We have

$$E(\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4) \leq \prod_{i=1}^4 E(\Gamma_i)^{1/4} \leq E(\Gamma_1)^{1/4} N^{9/4}$$

and therefore $E(\Gamma_1) \geq \delta^4 N^3$. By Balog–Szemerédi–Gowers (see [18, Theorem 5.2]), there is $\Gamma' \subseteq \Gamma_1$ such that $|\Gamma'| \geq \delta^{O(1)} N$ while $|\Gamma' - \Gamma'| \leq \delta^{-O(1)} N$.

Let $A = (8\Gamma' - 8\Gamma') \cap (\{0\} \times (\mathbb{R}/\mathbb{Z})^d)$. Since Γ' is a graph, we have that $|\Gamma' + A| = |\Gamma'| |A|$. However $|\Gamma' + A| \leq |9\Gamma' - 8\Gamma'| \leq \delta^{-O(1)} N$ by the Plünnecke–Ruzsa inequality (e.g. [18, Theorem 5.3]) and thus $|A| \leq \delta^{-O(1)}$.

Now, by abuse of notation we may view A as a subset of $(\mathbb{R}/\mathbb{Z})^d$. We claim there exists $T \subseteq \mathbb{Z}^d$ with $|T| \leq O(\log(1/\delta))$ such that $A \cap B(T, 1/4) = \{0\}$; we give a proof which is essentially identical to that in [18, Lemma 8.3]. Note that given any $w \in (\mathbb{R}/\mathbb{Z})^d \setminus \{0\}$ we have

$$\limsup_{M \rightarrow \infty} \mathbb{P}_{v \in \{-M, \dots, M\}^d} [\|v \cdot w\|_{\mathbb{R}/\mathbb{Z}} < 1/4] \leq 3/4.$$

This follows immediately from noting that if w has an irrational coordinate the probability tends to $1/2$ by Weyl’s equidistribution criterion while if w is rational the limiting probability is at most say $2/3$. Choosing an integer vector v which kills at least $1/4$ of the set iteratively then immediately gives the desired lemma.

Let $\psi: (\mathbb{R}/\mathbb{Z})^d \rightarrow (\mathbb{R}/\mathbb{Z})^T$ be defined as $\psi(\xi) = (t(\xi))_{t \in T}$. Now let $\tau = 2^{-7}$. By averaging there exists a cube $Q = \vec{x} + [0, \tau]^T$ such that

$$\tilde{\Gamma} := \{(h, f_1(h)) \in \Gamma' : \psi(f_1(h)) \in Q\}$$

with $|\tilde{\Gamma}| \geq \tau^{|T|} |\Gamma'|$, so $|\tilde{\Gamma}| \geq \delta^{O(1)} N$. Fix such a cube Q .

We claim that $4\tilde{\Gamma} - 4\tilde{\Gamma}$ is a graph. For the sake of contradiction suppose not. Then there exist h_1, \dots, h_8 and h'_1, \dots, h'_8 such that

$$h_1 + \dots + h_4 - h_5 - \dots - h_8 = h'_1 + \dots + h'_4 - h'_5 - \dots - h'_8,$$

$$f_1(h_1) + \dots + f_1(h_4) - f_1(h_5) - \dots - f_1(h_8) \not\equiv f_1(h'_1) + \dots + f_1(h'_4) - f_1(h'_5) - \dots - f_1(h'_8) \pmod{1}.$$

However,

$$\left\| \psi((f_1(h_1) + \dots + f_1(h_4) - f_1(h_5) - \dots - f_1(h_8)) - (f_1(h'_1) + \dots + f_1(h'_4) - f_1(h'_5) - \dots - f_1(h'_8))) \right\|_{\infty} \leq 16 \cdot \tau < 1/4$$

by definition of $\tilde{\Gamma}$. Since $A \cap B(T, 1/4) = \{0\}$, it follows that

$$(f_1(h_1) + \dots + f_1(h_4) - f_1(h_5) - \dots - f_1(h_8)) - (f_1(h'_1) + \dots + f_1(h'_4) - f_1(h'_5) - \dots - f_1(h'_8)) \in \mathbb{Z}^d$$

as desired.

Let H^* denote the projection of $\tilde{\Gamma}$ onto the first coordinate. Since f_1 is an 8-Freiman homomorphism on H^* (because $4\tilde{\Gamma} - 4\tilde{\Gamma}$ is a graph), we have that f_1 is a Freiman homomorphism on $2H^* - 2H^*$ (where f_1 is extended via linearity). We now view H^* (which is a subset of integers) as a subset of $\mathbb{Z}/N'\mathbb{Z}$ where N' is a prime in $[100N, 200N]$. Note here that $H^* \subseteq [-4N, 4N]$ and thus $4\tilde{\Gamma} - 4\tilde{\Gamma}$ when viewed as a subset of $(\mathbb{Z}/N'\mathbb{Z}) \times (\mathbb{R}/\mathbb{Z})^d$ is still a graph. Note that $|H^*| \geq \delta^{O(1)} N$.

By Lemma A.4, we have that $2H^* - 2H^*$ contains a Bohr set $B(S, \rho)$ with $|S|, \rho^{-1} \leq (\log(1/\delta))^{O(1)}$. Then by applying Lemma A.6 to each coordinate of f_1 on $B(S, \rho') \subseteq 2H^* - 2H^*$ with $\rho'^{-1} \leq \exp(\log(1/\delta))^{O(1)}$, we have that

$$(A.2) \quad f_1(h_1) = \sum_{\alpha_i \in S} a_i \{\alpha_i h_1\} + \gamma \pmod{1}$$

for all $h_1 \in B(S, \rho')$, for appropriate choices of $a_i, \gamma \in \mathbb{R}^d$. Here $\alpha_i \in (1/N')\mathbb{Z}$.

We now undo this transformation and we abusively view $B(S, \rho') \subseteq 2H^* - 2H^*$ as a subset of integers in $[-4N, 4N]$ instead of $\mathbb{Z}/N'\mathbb{Z}$, noting that the fractional part remains identical in both cases. As a slight technical annoyance, $B(S, \rho')$ might not intersect H^* . But, by Pigeonhole there exists $x^* \in [-5N, 5N]$ such that $|(x^* + B(S, \rho'/2)) \cap H^*| \geq \exp(-(\log(1/\delta))^{O(1)}) N$. (This requires a lower bound on the size of a Bohr set, see [51, Lemma 4.20].)

Fix $h^* \in B(S, \rho'/2)$ such that $x^* + h^* \in H^*$ and consider any $h_1 \in B(S, \rho'/2)$ such that $h_1 + x^* \in H^*$ we have that

$$f_1(h_1 - h^*) + f_1(x^* + h^*) = f_1(h_1 + x^*) + f(0) \pmod{1}$$

since $4\tilde{\Gamma} - 4\tilde{\Gamma}$ is a graph (note that $h_1 - h^* \in B(S, \rho') \subseteq 2H^* - 2H^*$). Thus we have

$$\begin{aligned} f_1(h_1 + x^*) &= f_1(h_1 - h^*) + f_1(x^* + h^*) - f(0) \pmod{1} \\ &= \sum_{\alpha_i \in S} a_i \{\alpha_i((h_1 + x^*) - (x^* + h^*))\} + \gamma' \pmod{1} \end{aligned}$$

The second line holds since x^*, h^* are viewed as fixed and $h_1 - h^* \in B(S, \rho')$ hence we may apply (A.2).

So, letting H' be the set of values $h_1 + x^* \in H_1$ where $h_1 \in B(S, \rho'/2)$, this nearly gives the desired result. The only issue is that there are shifts inside the brackets. Note that

$$\{z_1 + z_2\} = \begin{cases} \{z_1\} + \{z_2\} - 1 & \text{if } \{z_1\} + \{z_2\} > 1/2, \\ \{z_1\} + \{z_2\} + 1 & \text{if } \{z_1\} + \{z_2\} \leq -1/2, \\ \{z_1\} + \{z_2\} & \text{otherwise.} \end{cases}$$

Given this, we may Pigeonhole possible values $h_1 + x^*$ into one of $3^{|S|}$ cases based on the corresponding shift for each $\alpha_i \in S$. Applying the above relation with $z_1 = \alpha_i(h_1 + x^*)$ and $z_2 = -\alpha_i(x^* + h^*)$ and taking the most common case then gives the desired result. \square

APPENDIX B. MISCELLANEOUS DEFERRED RESULTS

We first require the following elementary lemma which will be used in the following deduction.

Lemma B.1. *Fix an integer $H \geq 2$. Consider vectors $v_1, \dots, v_\ell \in \mathbb{Z}^d$ with integer coordinates bounded by H and $w \in \mathbb{R}^d$ such that $\text{dist}(v_i \cdot w, \mathbb{Z}) \leq \delta$ for $1 \leq i \leq \ell$. We may write $w = w_{\text{small}} + w_{\text{rat}} + (w - w_{\text{small}} - w_{\text{rat}})$ where w_{rat} has coordinates which are rationals with denominators bounded by $H^{O(d^{O(1)})}$, $\|w_{\text{small}}\|_\infty \leq \delta \cdot H^{O(d^{O(1)})}$, and $v_i \cdot (w - w_{\text{small}} - w_{\text{rat}}) = 0$ for $1 \leq i \leq \ell$.*

Proof. Note that by passing to a subset we may assume that $v_1, \dots, v_\ell \in \mathbb{Z}^d$ are linearly independent. By Cramer's rule, there exist $w_1, \dots, w_\ell \in \mathbb{R}^d$ which have coordinates which are height $H^{O(d^{O(1)})}$ rationals such that $w_j \cdot v_k = \mathbb{1}_{j=k}$. Taking $w_{\text{rat}} = \sum_{j=1}^{\ell} (v_j \cdot w - \{v_j \cdot w\}) \cdot w_j$ and $w_{\text{small}} = \sum_{j=1}^{\ell} \{v_j \cdot w\} \cdot w_j$ we immediately have the desired result. Recall that we have chosen the fractional part $\{\cdot\}$ to live within $(-1/2, 1/2]$. \square

We now prove the following elementary lemma which takes a set of horizontal characters (at potentially different levels) and produces a factorization.

Lemma B.2. *Consider a nilmanifold G/Γ of degree-rank (s, r) of dimension d and complexity M . Consider a polynomial sequence g such that $g(0) = \text{id}_G$ and consider a set of horizontal characters $\psi_{i,j}$ for $1 \leq j \leq \ell_i$ and where $\psi_{i,\cdot}$ is an i -th horizontal character of height at most H . Furthermore suppose that for all i, j ,*

$$\text{dist}(\psi_{i,j}(\text{Taylor}_i(g)), \mathbb{Z}) \leq H \cdot N^{-i}.$$

Then one may factor

$$g = \varepsilon \cdot g' \cdot \gamma$$

where:

- $\varepsilon(0) = g'(0) = \gamma(0) = \text{id}_G$;
- $\psi_{i,j}(\text{Taylor}_i(g')) = 0$;
- γ is $(MH)^{O_s(d^{O_s(1)})}$ -rational;
- $d_G(\varepsilon(n), \varepsilon(n-1)) \leq (MH)^{O_s(d^{O_s(1)})} \cdot N^{-1}$ for $n \in [N]$.

Proof. By the classification of polynomial sequences in terms of coordinates of the second-kind, we have that

$$g(n) = \exp\left(\sum_{k=1}^s \binom{n}{k} g_k\right)$$

for some $g_k \in \log(G_{(k,0)}) = \log(G_{(k,1)})$. Note that

$$\text{Taylor}_i(g) = \exp(g_k) \text{ mod } G_{(i,2)}$$

and note that each $\psi_{i,j}$ can be descended to a linear map on $\log(G_{(i,1)})$ with the property that $\psi_{i,j}(\log(\Gamma \cap G_{(i,1)})) \in \mathbb{Z}$ and $\psi_{i,j}(\log(G_{(i,2)})) = 0$. That $\psi_{i,j}$ descends uses the fact that

$\log(x) + \log(y) \equiv \log(xy) \pmod{\log(G_{(i,2)})}$ for $x, y \in G_{(i,1)}$, which follows from Baker–Campbell–Hausdorff.

We now apply Lemma B.1. As $\text{dist}(\psi_{i,j}(\text{Taylor}_i(g)), \mathbb{Z}) \leq H \cdot N^{-i}$ by assumption, we may write $g_i = g_{i,\text{small}} + g_{i,\text{rat}} + (g_i - g_{i,\text{small}} - g_{i,\text{rat}})$ such that $g_{i,\text{rat}}$ is an $H^{O_s(d^{O_s(1)})}$ -rational combination of elements in $\mathcal{X} \cap \log(G_{(i,1)})$, such that $\|g_{i,\text{small}}\|_\infty \leq (MH)^{O_s(d^{O_s(1)})} \cdot N^{-i}$, and such that $\psi_{i,j}(g_i - g_{i,\text{small}} - g_{i,\text{rat}}) = 0$. Defining

$$\gamma := \exp\left(\sum_{k=1}^s \binom{n}{k} g_{k,\text{rat}}\right), \quad \varepsilon := \exp\left(\sum_{k=1}^s \binom{n}{k} g_{k,\text{small}}\right),$$

and $g' := \varepsilon^{-1} g \gamma^{-1}$, we immediately have that $\gamma \Gamma$ is $(MH)^{O_s(d^{O_s(1)})}$ -periodic by [35, Lemma B.14]. That ε is sufficiently smooth is an immediate consequence of [35, Lemmas B.1, B.3]. \square

We next require the following result regarding the existence of a nilmanifold partition of unity. As a remark, a similar statement (e.g. with $\sum_j \tau_j = 1$) appears as [38, Lemma 2.4]. The proof there, strangely, does not adapt in a straightforward manner to here.

Lemma B.3. *Fix $\varepsilon \in (0, 1/2)$ and a nilmanifold G/Γ of degree s , dimension d , and complexity M . There exists an index set I and a collection of nonnegative smooth functions $\tau_j: G/\Gamma \rightarrow \mathbb{R}^{\geq 0}$ for $j \in I$ such that:*

- For all $g \in G$, we have $\sum_{j \in I} \tau_j(g\Gamma)^2 = 1$;
- $|I| \leq (1/\varepsilon)^{O_s(d^{O_s(1)})}$;
- For each $j \in I$, there exists $\beta \in [-2, 2]^d$ so that for any $g\Gamma \in \text{supp}(\tau_j)$ there exists $g' \in g\Gamma$ such that $\psi_G(g') \in \prod_{i=1}^d [\beta_i - \varepsilon, \beta_i + \varepsilon]$;
- τ_j are $(M/\varepsilon)^{O_s(d^{O_s(1)})}$ -Lipschitz on G/Γ ;
- For any $g \in G$, $g\Gamma$ is contained in the support of at most $2^{O_s(d)}$ terms.

Proof. We will prove the statement inductively based on the degree of the nilmanifold. For degree 1 nilmanifolds G , note that $G \simeq \mathbb{T}^d$. There exists a set of function $\rho_1, \dots, \rho_{2k}: \mathbb{T} \rightarrow \mathbb{R}^{\geq 0}$ such that:

- $\text{supp}(\rho_j) \subseteq [j/(2k), j/(2k) + 1/k] \pmod{1}$;
- $\sum_{j=1}^{2k} \rho_j^2 = 1$;
- ρ_j are $O(1/k)$ -Lipschitz.

Taking $k = O(\varepsilon^{-1})$, we have that

$$1 = \sum_{(j_1, \dots, j_d) \in [2k]^d} \prod_{\ell=1}^d \rho_{j_\ell}((\psi_G(g))_\ell)^2$$

where $(\psi_G)_\ell$ denotes the ℓ -th coordinate of ψ_G . For $\vec{j} \in [2k]^d$ we take

$$\tau_{\vec{j}}(g) = \prod_{\ell=1}^d \rho_{j_\ell}((\psi_G(g))_\ell)$$

and note that this function is Γ -invariant since multiplying by an element in Γ shifts all coordinates by an integer. Furthermore, by [35, Lemma B.3] we have that the standard ℓ^∞ -metric on G/Γ is equivalent to $d_{G/\Gamma}$ up to a factor of $O(M)^{O(d^{O(1)})}$. This completes the proof in this case.

When considering the case of a degree $s \geq 2$ filtration on G , suppose that $G_0 = G_1 \geq G_2 \geq \dots \geq G_s \geq \text{Id}_G$ is the given filtration. Note that if $\mathcal{X} = \{X_1, \dots, X_d\}$ is the adapted Mal'cev basis for G/Γ then

$$\tilde{\mathcal{X}} := \{X_1, \dots, X_{\dim(G) - \dim(G_s)}\} \pmod{\log(G_s)}$$

is a valid Mal'cev basis for $\tilde{G} := G/G_s$. Furthermore define $\tilde{\Gamma} := \Gamma/(\Gamma \cap G_s)$. The complexity of $\tilde{\mathcal{X}}$ is always bounded by M by definition. The filtration on \tilde{G} is lower degree.

By induction, we have functions $(\tau_j)_{j \in I}$ with $|I| \leq (M/\varepsilon)^{O_s(d^{O_s(1)})}$ such that

$$1 = \sum_{j \in I} \tilde{\tau}_j(\tilde{g}\tilde{\Gamma})^2$$

and satisfying various other appropriate properties. We may lift these functions to G/Γ via

$$\tau_j(g\Gamma) = \tilde{\tau}_j((g \bmod G_s)\tilde{\Gamma}).$$

Note that this is well-defined since $g\Gamma \bmod G_s = (g \bmod G_s) \cdot (\Gamma \bmod G_s) = (g \bmod G_s)\tilde{\Gamma}$.

We view each τ_j as a function on $\prod_{i=1}^{\dim(\tilde{G})} (\beta_i - 1/2, \beta_i + 1/2] \times \mathbb{T}^{\dim(G_s)}$ which only depends on the first $\dim(\tilde{G})$ coordinates and such that the support is only within some $\prod_{i=1}^{\dim(\tilde{G})} [\beta_i - \varepsilon, \beta_i + \varepsilon] \times \mathbb{T}^{\dim(G_s)}$. This is via identifying the fundamental domain of G/Γ via Mal'cev coordinates of the second-kind (see the proof of [35, Lemma B.6]). We let $\psi_\beta: G/\Gamma \rightarrow \prod_{i=1}^{\dim(\tilde{G})} (\beta_i - 1/2, \beta_i + 1/2] \times \mathbb{T}^{\dim(G_s)}$ denote this identification. (Note that the choice of β depends on $j \in I$, which we will fix through the remainder of the proof.)

We now have

$$\tau_j(g\Gamma)^2 = \tilde{\tau}_j((g \bmod G_s)\tilde{\Gamma})^2 \cdot \sum_{(t_1, \dots, t_{\dim(G_s)}) \in [2k]^{\dim(G_s)}} \prod_{\ell=1}^{\dim(G_s)} \rho_{t_\ell}((\psi_\beta(g\Gamma))_{\ell + \dim(\tilde{G})})^2$$

where $k = O(1/\varepsilon)$ and ρ are defined as above.

The fact that each piece

$$\tau_{j,\vec{t}}(g\Gamma)^2 := \tau_j(g\Gamma)^2 \cdot \prod_{\ell=1}^{\dim(G_s)} \rho_{t_\ell}((\psi_\beta(g\Gamma))_{\ell + \dim(\tilde{G})})^2$$

is Γ -invariant on the right is trivial by construction, and the sum of squares property is trivial.

Identifying ρ_j with a function $\mathbb{R} \rightarrow \mathbb{R}^{\geq 0}$ where $\text{supp}(\rho_j) \subseteq [j/(2k), j/(2k) + 1/k]$, we may identify $\tau_{j,\vec{t}}$ with a function on the fundamental domain (with respect to second-kind coordinates) of the form

$$\prod_{i=1}^{\dim(\tilde{G})} (\beta_i - 1/2, \beta_i + 1/2] \times \prod_{\ell=1}^{\dim(G_s)} ((t_\ell + 1)/(2k) - 1/2, (t_\ell + 1)/(2k) + 1/2].$$

To check that this function is sufficiently Lipschitz, we note that each element $g\Gamma$ has a unique representative in this domain.

Consider $\tau_{j,\vec{t}}(x\Gamma)$ and $\tau_{j,\vec{t}}(y\Gamma)$; by multiplying by the lattice we may assume that $\psi(x), \psi(y)$ are in the specified fundamental domain. Furthermore if $d_{G/\Gamma}(x\Gamma, y\Gamma) \geq \varepsilon' = M^{-O_s(d_s^{O(1)})}$ we immediately win as $\tau_{j,\vec{t}}$ is 1-bounded. We claim that if $d_{G/\Gamma}(x\Gamma, y\Gamma) \leq \varepsilon'$ then $d_{G/\Gamma}(x\Gamma, y\Gamma) = d_G(x, y)$. In particular, note that

$$d_{G/\Gamma}(x\Gamma, y\Gamma) = \min_{\gamma \in \Gamma} d_G(x\gamma, y)$$

and that

$$\min_{\gamma \in \Gamma \setminus \{\text{id}_G\}} d_G(x\gamma, y) \geq M^{-O_s(d_s^{O(1)})} \cdot \min_{\gamma \in \Gamma \setminus \{\text{id}_G\}} d_G(\gamma, x^{-1}y) \geq M^{-O_s(d_s^{O(1)})}$$

which gives the desired contradiction assuming that various implicit constants defining ε' are chosen appropriately.

Now we may assume that x, y are such that

$$\psi(x), \psi(y) \in \prod_{i=1}^{\dim(\tilde{G})} [\beta_i - 2\varepsilon, \beta_i + 2\varepsilon] \times \prod_{\ell=1}^{\dim(G_s)} [t_\ell/(2k) - \varepsilon, t_\ell/(2k) + 1/k + \varepsilon],$$

else both function values vanish (again supposing ε' is sufficiently small). This is because $d_G(x, y)$ is equivalent to $\|\psi(x) - \psi(y)\|_\infty$ (up to a factor of $M^{O_s(d^{O_s(1)})}$) for bounded elements by [35, Lemma B.3]), and due to the condition on the support of ρ_{t_ℓ} .

In particular, $\psi(x), \psi(y)$ are seen to lie in the interior of the domain. The result then follows immediately noting that τ_j is appropriately Lipschitz and ρ_{t_ℓ} is an appropriately Lipschitz function on \mathbb{R} . The claim that $g\Gamma$ is contained in the support of at most $2^{O_s(d)}$ terms follows trivially by construction. \square

Given this we are now in position to show the existence of nilcharacters on G/Γ .

Lemma B.4. *Fix $\varepsilon \in (0, 1/2)$ and a nilmanifold G/Γ of degree s , dimension d , and complexity M . Fix η a vertical G_s -frequency with height bounded by M . There exists a nilcharacter F with frequency η such that the output dimension is bounded by $2^{O_s(d^{O_s(1)})}$ and each coordinate is $O_s(M)^{O_s(d^{O_s(1)})}$ -Lipschitz.*

Proof. Let $\tilde{G} = G/G_s$ and $\tilde{\Gamma} = \Gamma/(\Gamma \cap G_s)$. Apply Lemma B.3 on $\tilde{G}/\tilde{\Gamma}$ with $\varepsilon = 1/4$ to obtain $\tilde{\tau}_j$ for $j \in I$. For $\eta = 0$, we may take the coordinates of F to be

$$\tau_j(g\Gamma) = \tilde{\tau}_j((g \bmod G_d)\tilde{\Gamma}).$$

In general, for appropriate β depending on j , we have that $g\Gamma$ is naturally identified with a unique point inside $\prod_{i=1}^{\dim(\tilde{G})} (\beta_i - 1/2, \beta_i + 1/2] \times \mathbb{T}^{\dim(G_s)}$ as in the proof of Lemma B.3 and we let $\psi_\beta(g\Gamma)$ denote this map. The key point is to write

$$\tau_j(g\Gamma) = \tilde{\tau}_j((g \bmod G_s)\tilde{\Gamma}) \cdot \exp(\eta \cdot \psi_\beta(g\Gamma))$$

and note that $\sum_{j \in I} |\tau_j(g\Gamma)|^2 = 1$ as before. Here we have identified η with an integer vector using the last $\dim(G_s)$ elements of the Mal'cev basis and extending by 0. Note that this is trivially a function on G/Γ and by construction it has the G_s -vertical frequency η . The only technical point is verifying that this function is indeed Lipschitz, which we check for each coordinate τ_j .

Consider $x\Gamma$ and $y\Gamma$. If $\tau_j(x\Gamma) = \tau_j(y\Gamma) = 0$ the Lipschitz condition is obviously satisfied. Thus at least one value is nonzero, and without loss of generality we may assume $\tau_j(x\Gamma) \neq 0$. Furthermore, noting that τ_j is 1-bounded, we may assume that $d_{G/\Gamma}(x\Gamma, y\Gamma) \leq M^{-O_s(d^{O_s(1)})}$. As $\tau_j(x\Gamma) \neq 0$, possibly shifting x on the right by an element in the lattice allows us to assume

$$\psi(x) \in \prod_{i=1}^{\dim(\tilde{G})} (\beta_i - 1/4, \beta_i + 1/4] \times (0, 1]^{\dim(G_s)}.$$

Via an argument analogous to that in the proof of Lemma B.3, there exists y' such that $y'\Gamma = y\Gamma$,

$$\psi(y') \in \prod_{i=1}^{\dim(\tilde{G})} (\beta_i - 1/3, \beta_i + 1/3] \times (-1/2, 3/2]^{\dim(G_s)},$$

and $\|\psi(x) - \psi(y')\|_\infty \leq M^{O_s(d^{O_s(1)})} d_{G/\Gamma}(x\Gamma, y\Gamma)$. Since $\vec{z} \mapsto \exp(\eta \cdot \vec{z})$ is an appropriately Lipschitz function on the torus if $\eta \in \mathbb{Z}^{\dim(G_s)}$, the desired result follows immediately. \square

We will also require the following converse of the U^{s+1} -inverse theorem; this is verbatim in [27, Appendix G] modulo various complexity details being omitted.

Lemma B.5. Fix $\varepsilon \in (0, 1/2)$ and let G/Γ be a degree s nilmanifold of dimension d and complexity M , and let $g(n)$ be a polynomial sequence with respect to this filtration. Furthermore let $F: G/\Gamma \rightarrow \mathbf{C}$ satisfy $\|F\|_{\text{Lip}} \leq M$. If $f: [N] \rightarrow \mathbf{C}$ is a 1-bounded function such that

$$|\mathbb{E}_{n \in [N]} f(n) \overline{F(g(n)\Gamma)}| \geq \varepsilon,$$

then

$$\|f\|_{U^{s+1}[N]} \geq (\varepsilon/M)^{O_s(d^{O_s(1)})}.$$

Proof. In the degenerate case when $s = 0$, we take a degree s nilsequence of complexity M to be a constant function ψ bounded by M . This implies that

$$|\mathbb{E}_{n \in [N]} f(n)| \geq \varepsilon/M$$

and by Cauchy–Schwarz we have

$$\mathbb{E}_{n, n' \in [N]} f(n) \overline{f(n')} \geq (\varepsilon/M)^2.$$

By unwinding definitions this implies the case $s = 0$.

For larger s , by applying [35, Lemma A.6] we may assume that

$$|\mathbb{E}_{n \in [N]} f(n) \overline{F_\xi(g(n)\Gamma)}| \geq (\varepsilon/M)^{O_s(d^{O_s(1)})}$$

where F_ξ is a $(M/\varepsilon)^{O_s(d^{O_s(1)})}$ -Lipschitz function with G_s -vertical frequency ξ bounded in height by $(M/\varepsilon)^{O_s(d^{O_s(1)})}$, after Pigeonhole. Cauchy–Schwarz implies that

$$\mathbb{E}_{n, n' \in [N]} f(n) \overline{f(n')} F_\xi(g(n')\Gamma) \overline{F_\xi(g(n)\Gamma)} \geq (\varepsilon/M)^{O_s(d^{O_s(1)})}.$$

Note that we may rewrite this as

$$\mathbb{E}_{n \in [N], h \in [\pm N]} f(n) \overline{f(n+h)} F_\xi(g(n+h)\Gamma) \overline{F_\xi(g(n)\Gamma)} \geq (\varepsilon/M)^{O_s(d^{O_s(1)})},$$

where we extend f by 0 in the usual manner. We define

$$G^\square = \{(g, g') : g, g' \in G, g^{-1}g' \in G_2\}$$

and note that this has a filtration $(G^\square)_i = \{(g, g') : g, g' \in G_i, g^{-1}g' \in G_{i+1}\}$ by [35, Lemma A.3] (with $G^\square = (G^\square)_1$). Let $\Gamma^\square = (\Gamma \times \Gamma) \cap G^\square$ and note that

$$\tilde{F}_\xi((x, y)(\Gamma \times \Gamma)) := F_\xi(x\Gamma) \overline{F_\xi(y\Gamma)}$$

is invariant under G_s^\square . Note that \tilde{F}_ξ is $(M/\varepsilon)^{O_s(d^{O_s(1)})}$ -Lipschitz on $G \times G$ and on G^\square , and G^\square/Γ^\square is a nilmanifold of appropriate complexity by [35, Lemma A.3].

Let

$$(g(0), g(h)) = \{(g(0), g(h))\} \cdot [(g(0), g(h))]$$

with $d_{G \times G}(\{(g(0), g(h))\}) \leq M^{O_s(d^{O_s(1)})}$ and $[(g(0), g(h))] \in \Gamma \times \Gamma$. Define

$$g'_h(n) = \{(g(0), g(h))\}^{-1} (g(n), g(n+h)) [(g(0), g(h))]^{-1};$$

this is easily seen to be a polynomial sequence with respect to G^\square . Thus

$$\mathbb{E}_{n \in [N], h \in [\pm N]} f(n) \overline{f(n+h)} \tilde{F}_\xi(\{(g(0), g(h))\} g'_h(n)(\Gamma \times \Gamma)) \geq (\varepsilon/M)^{O_s(d^{O_s(1)})}.$$

Define $\tilde{F}_{\xi, h}(x, y) := \tilde{F}_\xi(\{(g(0), g(h))\}(x, y)(\Gamma \times \Gamma))$ and note that it is $(M/\varepsilon)^{O_s(d^{O_s(1)})}$ -Lipschitz on $G \times G$ and on G^\square by [35, Lemma B.4]. Applying the triangle inequality and restricting to G^\square we have

$$\mathbb{E}_{h \in [\pm N]} \left| \mathbb{E}_{n \in [N]} \Delta_h f(n) \cdot \overline{\tilde{F}_{\chi, h}(g'_h(n)\Gamma^\square)} \right| \geq (\varepsilon/M)^{O_s(d^{O_s(1)})}.$$

Since \tilde{F}_ξ is invariant under $(G^\square)_s$, passing to $G^\square/(G^\square)_s$ gives a nilmanifold of degree $(s-1)$ and complexity $M^{O_s(d^{O_s(1)})}$. Thus we may apply by induction, and deduce that

$$\mathbb{E}_{h \in [\pm N]} \|\Delta_h f\|_{U^s[N]} \geq (\varepsilon/M)^{O_s(d^{O_s(1)})}.$$

Since

$$\mathbb{E}_{h \in [\pm N]} \|\Delta_h f\|_{U^s[N]}^{2^s} \lesssim_s \|f\|_{U^{s+1}[N]}^{2^{s+1}},$$

the desired result follows. \square

We now check the deferred Lemma 11.3.

Proof of Lemma 11.3. We first construct a weak basis for $G_{\text{Quot}} \times G_{\text{Lin}}$. Note that each element in $(g, g') \in G_{\text{Quot}} \times G_{\text{Lin}}$ may be written as

$$(g, g') = (g, \text{id}_{G_{\text{Lin}}}) \cdot (\text{id}_{G_{\text{Quot}}}, g').$$

Consider $\tilde{e}_{i,j}$ and consider $(r-1)$ -fold commutators of $\tilde{e}_{i_1, j_1}, \dots, \tilde{e}_{i_r, j_r}$ with $i_1 + \dots + i_r \leq s-2$ or $i_1 + \dots + i_r = s-1$, $r \leq r^*$ and at most one generator has $i_\ell > D_{i_\ell}^*$. We define the type of the commutator to be given by the multiset $\{\tilde{e}_{i_1, j_1}, \dots, \tilde{e}_{i_r, j_r}\}$ and we say that said type is linear if $i_\ell > D_{i_\ell}^*$ for exactly one index ℓ . We define the degree of a commutator to be $i_1 + \dots + i_r$. As discussed in Lemmas 10.4 and 10.10, commutators of all types span $\log(G_{\text{Quot}})$ and commutators of linear type span $\log(G_{\text{Lin}})$, and all relations between these elements are spanned by relations between commutators of the same type of height $O_s(1)$.

Given this, for each collection of commutators of a given type choose a subset which ‘‘spans the type’’ (similar to in the proof of Lemma 10.4). Let \mathcal{X}_1 denote the set of selected commutators and \mathcal{X}_2 denote the selected commutators which are of linear type. Our weak basis for $G_{\text{Quot}} \times G_{\text{Lin}}$ will be

$$\mathcal{X} = \{(X, 0) : X \in \mathcal{X}_1\} \cup \{(0, X) : X \in \mathcal{X}_2\};$$

this is seen to be a basis for the Lie algebra of $G_{\text{Quot}} \times G_{\text{Lin}}$. That it spans is trivial, and if there were a relation note that there could be no elements of the form $(X, 0)$ in the relation since projecting onto the first coordinate we recover multiplication in G_{Quot} . Given that there are no elements of the form $(X, 0)$, within this relation multiplication then acts exactly as in G_{Lin} and the result claimed independence follows.

We give $G_{\text{Quot}} \times G_{\text{Lin}}$ a multidegree filtration by taking the multidegree filtration of G_{Multi} and intersecting with the subgroup of elements of the form $(0, (g, g_1))$. We see that all the subgroups of the filtration are in fact spanned subsets by subsets of \mathcal{X} . This is simply by taking the generators in \mathcal{X} of the appropriate degree-rank; for instance

$$(G_{\text{Quot}} \times G_{\text{Lin}})_{(0,d)} = \{(g, g_1) : g \in (G_{\text{Quot}})_{(d,0)}, g_1 \in (G_{\text{Quot}})_{(d,0)} \cap G_{\text{Lin}}\}$$

and we take the subsets of $\{(X, 0) : X \in \mathcal{X}_1\}$ and $\{(0, X) : X \in \mathcal{X}_2\}$ where X has degree at least d . This is similarly true for $\bigvee_{|\vec{i}|=k} (G_{\text{Quot}} \times G_{\text{Lin}})_{(i_1, i_2)}$ which will ultimately form the underlying degree filtration for $G_{\text{Quot}} \times G_{\text{Lin}}$. Furthermore ordering the basis according to whether they lie in the degree ordering associated to $G_{\text{Quot}} \times G_{\text{Lin}}$ proves that the basis has the degree $O_s(1)$ nesting property. Thus it suffices to check the complexity of various commutators.

Note the identity

$$[V, W] = \frac{d}{ds} \frac{d}{dt} \exp(sV) \exp(tW) \exp(-sV) \exp(-tW) \Big|_{s,t=0}$$

which holds for any Lie group and the associated Lie bracket. It is therefore immediate that

$$[(X, 0), (X', 0)] = ([X, X'], 0) \text{ and } [(0, X), (0, X')] = (0, [X, X']),$$

and we have

$$\begin{aligned}
& [(X, 0), (0, X')] \\
&= \frac{d}{ds} \frac{d}{dt} (\exp(sX), \text{id}_{G_{\text{Lin}}}) \cdot (\text{id}_{G_{\text{Quot}}}, \exp(tX')) \cdot (\exp(-sX), \text{id}_{G_{\text{Lin}}}) \cdot (\text{id}_{G_{\text{Quot}}}, \exp(-tX')) \Big|_{s,t=0} \\
&= \frac{d}{ds} \frac{d}{dt} (\text{id}_{G_{\text{Quot}}}, \exp(sX) \exp(tX') \exp(-sX) \exp(-tX')) \Big|_{s,t=0} \\
&= (0, [X, X']).
\end{aligned}$$

This immediately implies that the structure constants associated to the weak basis \mathcal{X} are of height $O_s(1)$.

When including the semi-direct action, we will use the weak basis given by taking elements $\log((\vec{e}_{ij}, (\text{id}_{G_{\text{Quot}}}, \text{id}_{G_{\text{Lin}}}))$) where $\vec{e}_{i,j}$ denotes the elementary basis vector in the corresponding direction in R , placed at the start of \mathcal{X} . This is easily seen to preserve the nesting property.

To compute the associated structure constants, first note that

$$[(\vec{e}_{ij}, (\text{id}_{G_{\text{Quot}}}, \text{id}_{G_{\text{Lin}}})) , (0, (g, \text{id}_{G_{\text{Lin}}}))] = \text{id}_{G_{\text{Multi}}}$$

and thus the all Lie bracket structure constants of the corresponding form vanish. Furthermore note that

$$[\log((\vec{e}_{ij}, (\text{id}_{G_{\text{Quot}}}, \text{id}_{G_{\text{Lin}}})) , (0, (0, X')))] = \frac{d}{ds} \frac{d}{dt} (0, (\exp(tX')^{s \cdot \vec{e}_{i,j}}, \text{id}_{G_{\text{Lin}}})) \Big|_{s,t=0}.$$

We have that if the type of X' does not contain $\vec{e}_{i,j}$ then $\exp(tX')^{s \cdot \vec{e}_{i,j}} = \text{id}_{G_{\text{Quot}}}$ and otherwise $\exp(tX')^{s \cdot \vec{e}_{i,j}} = \exp(stX')$ (recall the definition of exponentiation by elements of R given in Section 11.1). In either case the structure constant is appropriately rational. Therefore we may construct a Mal'cev basis adapted to G_{Multi} with the appropriate complexity by applying [35, Lemma B.11] to \mathcal{X} to construct a Mal'cev basis for $G_{\text{Quot}} \times G_{\text{Lin}}$ and adding the semi-direct Mal'cev basis elements described above to the front of the list. We define this basis to be $\mathcal{X}_{\text{Multi}}$ and define initial segment corresponding to the semi-direct Mal'cev basis elements to be $\mathcal{X}_{\text{Multi},R}$ and the remaining elements to be $\mathcal{X}_{\text{Multi},G_{\text{Quot}} \times G_{\text{Lin}}}$.

We finally check that F_{Multi} is an appropriately Lipschitz function. Let δ be defined as in Section 11.1. Fix a pair $x, y \in G_{\text{Multi}}$. Note that if

$$d_{G_{\text{Multi}}/\Gamma_{\text{Multi}}}(x\Gamma_{\text{Multi}}, y\Gamma_{\text{Multi}}) \geq \delta^{O_s(d^{O_s(1)})}$$

we have that

$$\frac{F_{\text{Multi}}(x\Gamma_{\text{Multi}}) - F_{\text{Multi}}(y\Gamma_{\text{Multi}})}{d_{G_{\text{Multi}}/\Gamma_{\text{Multi}}}(x\Gamma_{\text{Multi}}, y\Gamma_{\text{Multi}})} \leq \delta^{-O_s(d^{O_s(1)})} \cdot 2 \|F_{\text{Multi}}\|_{\infty}$$

which is sufficiently bounded. Therefore to check the Lipschitz constant it suffices to consider x, y such that $d_{G_{\text{Multi}}/\Gamma_{\text{Multi}}}(x\Gamma_{\text{Multi}}, y\Gamma_{\text{Multi}}) \leq \delta^{O_s(d^{O_s(1)})}$ (where the implicit constants are chosen sufficiently large for the remainder of the argument). By multiplying by elements in the lattice, we may assume that $d_{G_{\text{Multi}}/\Gamma_{\text{Multi}}}(x\Gamma_{\text{Multi}}, y\Gamma_{\text{Multi}}) = d_{G_{\text{Multi}}}(x, y)$, that $F_{\text{Multi}}(x\Gamma_{\text{Multi}}) \neq 0$, and

$$\psi_{\mathcal{X}_{\text{Multi}}}(x) \in [-1/2, 1/2]^{\dim(G_{\text{Multi}})}.$$

Note that to assume that $F_{\text{Multi}}(x\Gamma_{\text{Multi}}) \neq 0$ we may need to swap x and y (if both are zero there is nothing to check with respect to the Lipschitz constant).

Since $F_{\text{Multi}}(x\Gamma_{\text{Multi}}) \neq 0$ we in fact have that the first $\sum_{i=1}^{s-1} D_i^{\text{Lin}}$ coordinates of $\psi_{\mathcal{X}_{\text{Multi}}}(x)$ are in $[-1/2 + \delta, 1/2 - \delta]$. This implies, due to the distance bound between x and y and by [35,

Lemma B.3], that the first $\sum_{i=1}^{s-1} D_i^{\text{Lin}}$ coordinates of $\psi_{\mathcal{X}_{\text{Multi}}}(y)$ are in $[-1/2 + \delta/2, 1/2 - \delta/2]$. Therefore if $x = (t_1, (g_1, g'_1))$ and $y = (t_2, (g_2, g'_2))$ then

$$F_{\text{Multi}}(x\Gamma_{\text{Multi}}) = F^*(g'_1\Gamma_{\text{Quot}}) \cdot \prod_{\substack{1 \leq i \leq s-1 \\ D_i^* < j \leq D_i^* + D_i^{\text{Lin}}}} \phi((t_1)_{i,j}),$$

$$F_{\text{Multi}}(y\Gamma_{\text{Multi}}) = F^*(g'_2\Gamma_{\text{Quot}}) \cdot \prod_{\substack{1 \leq i \leq s-1 \\ D_i^* < j \leq D_i^* + D_i^{\text{Lin}}}} \phi((t_2)_{i,j}).$$

Note that

$$|F_{\text{Multi}}(x\Gamma_{\text{Multi}}) - F_{\text{Multi}}(y\Gamma_{\text{Multi}})| \leq \|F^*\|_{\infty} \cdot \left| \prod_{\substack{1 \leq i \leq s-1 \\ D_i^* < j \leq D_i^* + D_i^{\text{Lin}}}} \phi((t_1)_{i,j}) - \prod_{\substack{1 \leq i \leq s-1 \\ D_i^* < j \leq D_i^* + D_i^{\text{Lin}}}} \phi((t_2)_{i,j}) \right|$$

$$+ |F^*(g'_1\Gamma_{\text{Quot}}) - F^*(g'_2\Gamma_{\text{Quot}})|,$$

where we have used that ϕ is 1-bounded. Next note that distance in $\psi_{\mathcal{X}, \text{exp}}$ controls the distance in $\psi_{\mathcal{X}}$ for bounded elements by [35, Lemma B.1] and distance in $\psi_{\mathcal{X}}$ controls distance in $d_{G_{\text{Multi}}}$ by [35, Lemma B.3]. The first term is therefore sufficiently bounded as ϕ is $O(1/\delta)$ -Lipschitz.

Finally note that

$$(g_1, g'_1) = \prod_{(X_i, X'_i) \in \mathcal{X}_{\text{Multi}, G_{\text{Quot}} \times G_{\text{Lin}}}} \exp((X_i, X'_i))^{x_i},$$

$$(g_2, g'_2) = \prod_{(X_i, X'_i) \in \mathcal{X}_{\text{Multi}, G_{\text{Quot}} \times G_{\text{Lin}}}} \exp((X_i, X'_i))^{y_i},$$

where x_i and y_i are the coordinates of x and y in $\psi_{\mathcal{X}_{\text{Multi}}}$ in the coordinates corresponding to $\mathcal{X}_{\text{Multi}, G_{\text{Quot}} \times G_{\text{Lin}}}$. This is using that

$$(t, (\text{id}_{G_{\text{Quot}}}, \text{id}_{G_{\text{Lin}}})) \cdot (0, (g, g')) = (t, (g, g')).$$

Therefore we have that

$$g_1 = \prod_{(X_i, X'_i) \in \mathcal{X}_{\text{Multi}, G_{\text{Quot}} \times G_{\text{Lin}}}} \exp(X_i)^{x_i}$$

$$g_2 = \prod_{(X_i, X'_i) \in \mathcal{X}_{\text{Multi}, G_{\text{Quot}} \times G_{\text{Lin}}}} \exp(X_i)^{y_i};$$

and note that $\exp(X_i)$ are appropriately bounded elements in G_{Quot} since $\mathcal{X}_{\text{Multi}, G_{\text{Quot}} \times G_{\text{Lin}}}$ are low height combinations of elements in \mathcal{X} . Via telescoping, and using that the metric $d_{G_{\text{Quot}}}$ is right-invariant and essentially left-invariant under multiplication by bounded elements (e.g. [35, Lemma B.4]), we have that

$$d_{G_{\text{Quot}}}(g_1, g_2) \leq \delta^{-O_s(d^{O_s(1)})} \cdot \sum_{(X_i, X'_i) \in \mathcal{X}_{\text{Multi}, G_{\text{Quot}}}} d_{G_{\text{Quot}}}(\exp(X_i)^{x_i - y_i}, \text{id}_{G_{\text{Quot}}})$$

$$\leq \delta^{-O_s(d^{O_s(1)})} \cdot \sum_{(X_i, X'_i) \in \mathcal{X}_{\text{Multi}, G_{\text{Quot}}}} |x_i - y_i|$$

$$\leq \delta^{-O_s(d^{O_s(1)})} \cdot \|\psi_{\mathcal{X}_{\text{Multi}}}(x) - \psi_{\mathcal{X}_{\text{Multi}}}(y)\| \leq \delta^{-O_s(d^{O_s(1)})} \cdot d_{G_{\text{Multi}}}(x, y).$$

This completes the proof upon noting that F^* is appropriately Lipschitz on G_{Quot} . \square

APPENDIX C. NILCHARACTERS

This section is essentially a straightforward quantification of various statements regarding nilcharacters proven in [29, Appendix E].

We first require that two nilcharacters being equivalent is a transitive relationship; this is a quantified version of [29, Lemma E.7]. Recall the notion of complexity (M, d) that we carry over from Section 12.

Lemma C.1. *Consider three nilcharacters χ_1, χ_2, χ_3 each of complexity (M, d) and such that the pair χ_1 and χ_2 and the pair χ_2 and χ_3 are (M, D, d) -equivalent for multidegree J . Then χ_1 and χ_3 are $((MD)^{O_{|J|}(1)}, (MD)^{O_{|J|}(1)}, O(d))$ -equivalent for multidegree J .*

Proof. Notice that each coordinate of $\chi_1 \otimes \overline{\chi_3}$ may be expressed as the sum of at most D coordinates of the nilcharacter

$$\chi_1 \otimes (\overline{\chi_2} \otimes \chi_2) \otimes \overline{\chi_3};$$

this follows since the trace of $\overline{\chi_2} \otimes \chi_2$ is 1. The result then follows by rewriting

$$\chi_1 \otimes (\overline{\chi_2} \otimes \chi_2) \otimes \overline{\chi_3} = (\chi_1 \otimes \overline{\chi_2}) \otimes (\chi_2 \otimes \overline{\chi_3})$$

and applying the assumption. □

We will generally require the following specialization lemmas; these are rather straightforward consequences of the definitions modulo the need to handle slight filtration issues.

Lemma C.2. *We have the following:*

- *Consider a nilsequence $\chi(h_1, \dots, h_k)$ of multidegree (s_1, \dots, s_k) and complexity (M, d) . Given $h^* \in \mathbb{Z}$, the function $\chi(h^*, h_2, \dots, h_k)$, treating h^* as fixed, is a multidegree (s_2, \dots, s_k) nilsequence of complexity $(M^{O_{|s|}(d^{O_{|s|}(1)}), d)$.*
- *Consider homomorphisms $L_i: \mathbb{Z}^\ell \rightarrow \mathbb{Z}$. If $\chi(h_1, \dots, h_k)$ is a nilsequence of degree s of complexity (M, d) then $\chi(L_1(t_1, \dots, t_\ell), \dots, L_k(t_1, \dots, t_\ell))$ is a degree s nilsequence in variables t_1, \dots, t_ℓ of complexity (M, d) .*
- *If $\chi(h_1, \dots, h_k)$ is a nilsequence of multidegree (s_1, \dots, s_k) of complexity (M, d) then it is also a nilsequence of degree $s_1 + \dots + s_k$ of complexity (M, d) .*

Remark. This result allows us to interpret expressions such as $\chi(h_1 + h'_1, h_2, \dots, h_k)$ as an appropriate degree nilcharacter in $k + 1$ variables, if χ is a nilcharacter in k variables with multidegree (s_1, \dots, s_k) .

Proof. We handle these items in reverse order (as this is also the difficulty of these claims). Let

$$\chi(h_1, \dots, h_k) = F(g(h_1, \dots, h_k)\Gamma)$$

with the underlying nilmanifold being G/Γ and the specified Mal'cev basis being \mathcal{X} .

For the last claim, note that \mathcal{X} (by the definition of complexity for multidegree nilmanifolds) is adapted to the degree filtration $G_t = \bigvee_{|\vec{i}|=t} G_{\vec{i}}$. Furthermore by the inclusion given on [29, p. 1264] or direct inspection given the Taylor expansion in [29, Lemma B.9], we have that $g(h_1, \dots, h_k)$ is a polynomial sequence with respect to the degree filtration $G_0 = G_1 \geq G_2 \geq \dots \geq G_{s_1 + \dots + s_k} \geq \text{Id}_G$. The desired result follows immediately.

For the second item, notice that if a polynomial $P(x_1, \dots, x_k)$ has total degree s , then for any linear maps $L_i: \mathbb{R}^\ell \rightarrow \mathbb{R}$ we have that $P(L_1(y_1, \dots, y_\ell), \dots, L_k(y_1, \dots, y_\ell))$ has total degree s . This coupled with Taylor expansion [29, Lemma B.9] and the fact that the set of polynomial sequences with respect to a given I -filtration is a group (by [29, Corollary B.4]) implies the result.

We now handle the first item; this is the only nontrivial part. Write $g(h^*, 0, \dots, 0) = \{g(h^*, 0, \dots, 0)\}[g(h^*, 0, \dots, 0)]$ with $\psi_{G, \mathcal{X}}(\{g(h^*, 0, \dots, 0)\}) \in [0, 1]^{\dim(G)}$ and $[g(h^*, 0, \dots, 0)] \in \Gamma$. We replace the polynomial sequence g by $g' = \{g(h^*, 0, \dots, 0)\}^{-1}g[h(h^*, 0, \dots, 0)]^{-1}$ and F by the function $F'(\cdot) = F(\{g((h^*, 0, \dots, 0))\} \cdot)$. We may thus assume, at the cost of replacing M by $M^{O_{|\vec{s}|}(d^{O_{|\vec{s}|}(1)})}$, that $g(h^*, 0, \dots, 0) = \text{id}_G$.

We now apply [29, Lemma B.9] to see

$$g(h_1, \dots, h_k) = \prod_{i_1, \dots, i_k} g_{(i_1, \dots, i_k)}^{\binom{h_1}{i_1} \dots \binom{h_k}{i_k}}$$

where we order (i_1, \dots, i_k) lexicographically with indices considered in reverse order in the product (in particular, the first few terms are $(0, \dots, 0)$, $(1, \dots, 0)$, $(2, \dots, 0)$, and so on) and $g_{(i_1, \dots, i_k)} \in G_{(i_1, \dots, i_k)}$. As $g(h^*, 0, \dots, 0) = \text{id}_G$, we have that

$$g(h^*, h_2, \dots, h_k) = \prod_{\substack{i_2, \dots, i_k \\ i_2 + \dots + i_k > 0}} g_{(i_1, \dots, i_k)}^{\binom{h^*}{i_1} \binom{h_2}{i_2} \dots \binom{h_k}{i_k}}.$$

It then follows that $g(h^*, h_2, \dots, h_k)$ is a polynomial with respect to

$$G^* = \bigvee_{\ell=2}^k G_{\vec{e}_\ell}$$

which we give a multidegree filtration $G^*_{(i_2, \dots, i_k)} = G_{(0, i_2, \dots, i_k)}$ for $i_2 + \dots + i_k > 0$ and $G^*_{(0, \dots, 0)} = G^*$. Note that all subgroups in this filtration are M -rational with respect to \mathcal{X} and that $(G^*)_t = \bigvee_{|\vec{i}|=t} G^*_{\vec{i}}$ is a degree $|\vec{s}| - s_1$ filtration. Therefore applying [35, Lemma B.11] guarantees that we may find a Mal'cev basis \mathcal{X}^* for G^* (which is $M^{O_{|\vec{s}|}(d^{O_{|\vec{s}|}(1)})}$ -rational with respect to \mathcal{X}). Descending F to G^* gives the desired result with the necessary Lipschitz bound following from [35, Lemma B.9]. \square

We now state a quantified version of [29, Lemma E.8]. Recall the notion of equivalence (Definition 7.3).

Lemma C.3. *Consider a nilcharacter χ with complexity (M, d) of multidegree $\vec{s} = (s_1, \dots, s_k)$ with $|\vec{s}| = s_1 + \dots + s_k$. We have that:*

- *The nilcharacters*

$$\chi(\cdot) \text{ and } \chi(\cdot)$$

*are $(M^{O_{|\vec{s}|}(1)}, M^{O_{|\vec{s}|}(1)}, O(d))$ -equivalent for multidegree $< (s_1, \dots, s_k)$.*⁹

- *Fix $h^* \in \mathbb{Z}$. The nilcharacters*

$$\chi(\cdot + h^* \vec{e}_j) \text{ and } \chi(\cdot)$$

are $(M^{O_{|\vec{s}|}(d^{O_{|\vec{s}|}(1)})}, M^{O_{|\vec{s}|}(d^{O_{|\vec{s}|}(1)})}, O(d))$ -equivalent for multidegree $< (s_1, \dots, s_k)$.

- *Fix $q \in \mathbb{Z}$. Then*

$$\chi^{\otimes q^{|\vec{s}|}}(\cdot) \text{ and } \chi(q \cdot)$$

are $(M^{O_{|\vec{s}|, q}(d^{O_{|\vec{s}|, q}(1)})}, M^{O_{|\vec{s}|, q}(d^{O_{|\vec{s}|, q}(1)})}, d^{O_q(1)})$ -equivalent for multidegree $< (s_1, \dots, s_k)$.

- *Fix $q \in \mathbb{Z}^{>0}$. There exists a nilcharacter $\tilde{\chi}$ of complexity $(M^{O_{|\vec{s}|, q}(d^{O_{|\vec{s}|, q}(1)})}, d^{O_q(1)})$ such that*

$$\chi(\cdot) \text{ and } \tilde{\chi}^{\otimes q}(\cdot)$$

are $(M^{O_{|\vec{s}|, q}(d^{O_{|\vec{s}|, q}(1)})}, M^{O_{|\vec{s}|, q}(d^{O_{|\vec{s}|, q}(1)})}, d^{O_q(1)})$ -equivalent for multidegree $< (s_1, \dots, s_k)$.

⁹This means we take the down-set generated by (s_1, \dots, s_k) and then remove (s_1, \dots, s_k) .

Remark. $\chi^{-\otimes q}$ for $q \in \mathbb{Z}^{>0}$ is interpreted as $\overline{\chi}^{\otimes q}$.

Proof. Throughout the proof, we let

$$\chi(\vec{n}) = F(g(\vec{n})\Gamma)$$

where the underlying nilmanifold is G/Γ and the underlying Mal'cev basis is \mathcal{X} . When going from item to item, we may reuse variables (e.g., G' will be defined in multiple different manners throughout the proof). Additionally, the following analysis implicitly uses that $|\vec{s}| \geq 1$; in the remaining case $\vec{s} = 0$ all nilsequences become fixed constants and the result is obvious.

For the first item, note that coordinates of $\chi \otimes \overline{\chi}$ are multidegree (s_1, \dots, s_k) polynomial sequences with respect to group $G' = \{(g, g) : g \in G\}$ given the filtration

$$G'_i = \{(g, g) : g \in G'_i\}.$$

As all coordinates of χ have the same vertical frequency, we have that the coordinates of $\chi \otimes \overline{\chi}$ are invariant under $G'_{(s_1, \dots, s_k)}$. This immediately gives the desired result upon taking a quotient and using Lemma 3.10.

For the second item, note that $G^{+\vec{e}_j} = (G'_{i+\vec{e}_j})_{i \in I}$ is a shifted filtration. Note that this is an I -filtration with respect to the multidegree ordering. We define the group

$$G' = \bigvee_{\ell=1}^k (G'_{\vec{e}_\ell} \times \text{Id}_G) \vee \{(g, g) : g \in G\},$$

let $\Gamma' = G' \cap (\Gamma \times \Gamma)$, and define the following I -filtration with respect to the multidegree ordering:

$$G'_i = \bigvee_{\ell=1}^k (G'_{i+\vec{e}_\ell} \times \text{Id}_G) \vee \{(g, g) : g \in G'_i\}.$$

By using Lemma 2.2 we may see that this is a valid. We define the cocompact groups similarly. Now the proof of [29, Lemma E.8] shows that

$$(g(\vec{n} + h\vec{e}_j), g(\vec{n}))$$

is a polynomial sequence with respect to this filtration and that

$$\widetilde{F}((x, y)(\Gamma \times \Gamma)) = F(x\Gamma) \otimes \overline{F(y\Gamma)}$$

is invariant under the action of $G'_{\vec{s}} = \{(g, g) : g \in G_{(s_1, \dots, s_k)}\}$.

We first construct a Mal'cev basis \mathcal{X}' on G' . Define $G_t^* = \bigvee_{|\vec{i}|=t} G'_{i+\vec{e}_j}$ and note that $G^* = G_0^*$ has a degree filtration

$$G_0^* = G_0^* \geq G_1^* \geq G_2^* \geq \dots$$

and all these subgroups are $M^{O_{|\vec{s}|}(1)}$ -rational with respect to \mathcal{X} . Therefore G^* has a Mal'cev basis \mathcal{X}^* which is adapted to this filtration and all elements are height at most $M^{O_{|\vec{s}|}(d^{O_{|\vec{s}|}(1)})}$ combinations of elements in \mathcal{X} by [35, Lemma B.11]. Note that

$$\{(X, X) : X \in \mathcal{X}\} \cup \{(X^*, 0) : X^* \in \mathcal{X}^*\}$$

is easily shown to be a weak basis of rationality $M^{O_{|\vec{s}|}(d^{O_{|\vec{s}|}(1)})}$ for G'/Γ' and has the degree $O_{|\vec{s}|}(1)$ nesting property. Letting $G'_t = \bigvee_{|\vec{i}|=t} G'_i$ we see that

$$G' = G' \geq G'_1 \geq G'_2 \geq \dots$$

form a sequence of subgroups such that $[G', G'_i] \leq G'_{i+1}$ for $i \geq 0$ (with $G'_0 = G'$). Thus by [35, Lemma B.11] we can find a Mal'cev basis \mathcal{X}' adapted to this sequence such that each element is a height $M^{O_{|\vec{s}|}(d^{O_{|\vec{s}|}(1)})}$ linear combination of

$$\{(X, X) : X \in \mathcal{X}\} \cup \{(X^*, 0) : X^* \in \mathcal{X}^*\}.$$

At present, however, we see that G' has not been given a multidegree filtration (only an I -filtration with respect to the multidegree ordering; recall Definition 2.4). We replace G' by

$$\tilde{G} = \bigvee_{\ell=1}^k G'_{\vec{e}_\ell} = G'_1$$

and note that \tilde{G} is appropriately rational with respect to \mathcal{X}' and $G'_i \leq \tilde{G}$ for $\vec{i} \neq 0$. Note that \tilde{G} is easily seen to have a multidegree (s_1, \dots, s_k) filtration. Furthermore, removing the initial $\dim(G') - \dim(\tilde{G})$ elements, we see that the truncation of \mathcal{X}' is valid Mal'cev basis for \tilde{G} of complexity $M^{O_{|\vec{s}|}(d^{O_{|\vec{s}|}(1)})}$ and all subgroups in the multidegree filtration are $M^{O_{|\vec{s}|}(d^{O_{|\vec{s}|}(1)})}$ -rational.

We now write $(g(h^* \vec{e}_j), g(0)) = \{(g(h^* \vec{e}_j), g(0))\}[(g(h^* \vec{e}_j), g(0))]$ where

$$\|\psi_{\mathcal{X}'}(\{(g(h^* \vec{e}_j), g(0))\})\|_\infty \leq M^{O_{|\vec{s}|}(d^{O_{|\vec{s}|}(1)})}$$

and $[(g(h^* \vec{e}_j), g(0))] \in \Gamma'$. We consider the modified polynomial sequence

$$g'(\vec{n}) = \{(g(h^* \vec{e}_j), g(0))\}^{-1} (g(\vec{n} + h^* \vec{e}_j), g(\vec{n})) [(g(h^* \vec{e}_j), g(0))]^{-1};$$

evaluating at $\vec{n} = 0$ this is now seen to be a polynomial sequence in \tilde{G} . Defining

$$F'((x, y)(\Gamma \times \Gamma)) = \tilde{F}(\{(g(h^* \vec{e}_j), g(0))\}(x, y)(\Gamma \times \Gamma)),$$

we have that F' is invariant under $(\tilde{G})_{(s_1, \dots, s_k)}$ and

$$F'(g'(\vec{n})) = F(g(\vec{n} + h^* \vec{e}_j)) \overline{F}(g(\vec{n})).$$

We may pass to the quotient group $\tilde{G}/\tilde{G}_{(s_1, \dots, s_k)}$ and the desired result is essentially an immediate consequence of Lemma 3.10.

We now come to the third item; we only maintain the notation from the first sentence of the proof. Note that via writing $g(0) = \{g(0)\}[g(0)]$ with $[g(0)] \in \Gamma$ and $\psi_{\mathcal{X}}(\{g(0)\}) \in [0, 1]^{\dim(G)}$, replacing $g(\vec{n})$ by $\{g(0)\}^{-1} g(\vec{n}) [g(0)]^{-1}$ and F by $F(\{g(0)\} \cdot)$, up to replacing M by $M^{O_{|\vec{s}|}(1)}$ we may assume that $g(0) = \text{id}_G$.

Define

$$G'_i = \bigvee_{\vec{j} > \vec{i}} (G_{\vec{j}} \times G_{\vec{j}}) \vee \bigvee \{(g^{q^{|\vec{i}|}}, g) : g \in G_{\vec{j}}\};$$

here $\vec{j} > \vec{i}$ means \vec{j} is coordinate-wise at least as large as \vec{i} and not identical. Furthermore $\Gamma' = G' \cap (\Gamma \times \Gamma)$; note that G' is isomorphic to $G \times G$ however we have given the group an alternate filtration. This is verified to be an I -filtration with respect to the multidegree ordering in [29, p. 1356]. Furthermore the proof of [29, Lemma E.8] shows that

$$(g(q\vec{n}), g(\vec{n}))$$

is a polynomial sequence with respect to this filtration and that

$$\tilde{F}((x, y)(\Gamma \times \Gamma)) = F(x\Gamma) \otimes \overline{F(y\Gamma)}^{\otimes q^{|\vec{s}|}}$$

is invariant under the action of $G_{\vec{s}} = \{(g^{q^{|\vec{s}|}}, g) : g \in G_{(s_1, \dots, s_k)}\}$. The primary technical issue, as before, is that while this is an I -filtration with respect to the multidegree ordering this is not a multidegree filtration (Definition 2.4).

We first give G' a Mal'cev basis. Note that \mathcal{X} is adapted to the degree filtration on G given by $G_t = \bigvee_{|\vec{i}|=t} G'_i$ (Definition 3.8). It is immediate to see that

$$\mathcal{X}^* = \{(X, 0) : X \in \mathcal{X} \cap G_1\} \cup \{(0, X) : X \in \mathcal{X} \cap G_1\}$$

is a Mal'cev basis for the product filtration on G . Then using [35, Lemma B.11] on

$$G'_0 = G'_0 \geq G'_1 \geq \dots$$

where $G'_t = \bigvee_{|\vec{i}|=t} G'_{\vec{i}}$, which is seen to satisfy $[G'_0, G'_t] \leq G'_{t+1}$ for $t \geq 0$, we easily construct a Mal'cev basis \mathcal{X}' for G'/Γ' coming from combinations of \mathcal{X}^* . (We implicitly use that $G = \bigvee_j G_{\vec{e}_j}$.) As G has complexity M , it is trivial to see that all subgroups in the filtration of G' are $M^{O_{|\vec{s}|,q}(d^{O_{|\vec{s}|,q}(1)})}$ -rational. The Mal'cev basis \mathcal{X}' clearly has the nesting property of order $|\vec{s}|$ since \mathcal{X} does.

We define \tilde{G} as

$$\tilde{G} = \bigvee_{\ell=1}^k G'_{\vec{e}_\ell}$$

and this group is seen to be appropriately rational with respect to \mathcal{X}' and is given the multi-degree filtration $\tilde{G}_{\vec{i}} = G'_{\vec{i}}$ for $\vec{i} \neq 0$. Noting that the constant term of the Taylor expansion of $(g(q\vec{n}), g(\vec{n}))$ is $(\text{id}_G, \text{id}_G)$, we have that this is in fact a polynomial sequence with respect to the multidegree filtration given to \tilde{G} . (This is where we use that we reduced to $g(0) = \text{id}_G$.) Furthermore, letting $\tilde{G}_t = \bigvee_{|\vec{i}|=t} \tilde{G}_{\vec{i}}$, we see that a truncation of \mathcal{X}' is an adapted Mal'cev basis to $\tilde{G}_0 = \tilde{G}_1 \geq \tilde{G}_2 \geq \dots$ where each element is an $M^{O_{|\vec{s}|,q}(d^{O_{|\vec{s}|,q}(1)})}$ -rational combination of \mathcal{X}^* . As \tilde{F} is invariant under $\tilde{G}_{(s_1, \dots, s_k)}$, by passing to the quotient $\tilde{G}/\tilde{G}_{(s_1, \dots, s_k)}$ and applying Lemma 3.10 we immediately finish the proof.

We finally deduce the fourth item from the third item. Let $g'(n) = g(n/q)$; note that g may be extended to take on rational input via using Mal'cev coordinates and we may treat g' as a valid polynomial sequence. By applying the third item, we have that

$$F(g(n)) \text{ and } F(g'(n))^{\otimes q^{|\vec{s}|}}$$

are $(M^{O_{|\vec{s}|,q}(d^{O_{|\vec{s}|,q}(1)})}, M^{O_{|\vec{s}|,q}(d^{O_{|\vec{s}|,q}(1)})}, d^{O_q(1)})$ -equivalent for multidegree $< (s_1, \dots, s_k)$. Outputting $F(g'(n))^{\otimes q^{|\vec{s}|-1}}$ then gives the desired result. \square

The next lemma is a quantified version of [29, Lemma 13.2]. The proof is once again essentially identical modulo noting slight changes in the filtration notions.

Lemma C.4. *Consider $\chi: \mathbb{Z}^k \rightarrow \mathbb{C}$ which is a multidegree $(1, \dots, 1)$ nilcharacter of complexity (M, d) . Then*

$$\chi(h_1 + h'_1, h_2, \dots, h_k) \text{ and } \chi(h_1, h_2, \dots, h_k) \otimes \chi(h'_1, h_2, \dots, h_k)$$

are $(M^{O_k(d^{O_k(1)})}, d^{O_k(1)})$ -equivalent for degree $(k-1)$.

Proof. Let $\chi(h_1, \dots, h_k) = F(g(h_1, \dots, h_k)\Gamma)$ where the underlying nilmanifold is G/Γ and the underlying Mal'cev basis is \mathcal{X} . Let $g(0, \dots, 0) = \{g(0, \dots, 0)\}[g(0, \dots, 0)]$ with $[g(0, \dots, 0)] \in \Gamma$ and $\psi_{G, \mathcal{X}}(\{g(0, \dots, 0)\}) \in [0, 1]^{\dim(G)}$. We have that

$$F(g(h_1, \dots, h_k)\Gamma) = F(\{g(0, \dots, 0)\} \cdot (\{g(0, \dots, 0)\}^{-1} g(h_1, \dots, h_k) [g(0, \dots, 0)]^{-1} \Gamma)).$$

We let $F'(\cdot) = F(\{g(0, \dots, 0)\} \cdot \cdot)$ and $g'(h_1, \dots, h_k) = (\{g(0, \dots, 0)\}^{-1} g(h_1, \dots, h_k) [g(0, \dots, 0)]^{-1})^{-1}$. Thus may assume replace F by F' and g by g' (at the cost of replacing M by $M^{O_k(d^{O_k(1)})}$) and assume that $g(0) = \text{id}_G$.

Consider, for $t \geq 1$,

$$G'_t = \bigvee_{|\vec{i}|>t} (G_{\vec{i}} \times G_{\vec{i}} \times G_{\vec{i}}) \vee \bigvee_{|\vec{i}|=t} \{(g_1 g_2, g_1, g_2) : g_i \in G_{\vec{i}}\}$$

and take $G' = G'_1$. Via Baker–Campbell–Hausdorff this gives a valid degree k filtration

$$G' = G'_1 \geq G'_2 \geq \cdots \geq G'_k \geq \text{Id}_{G'}.$$

We define $\Gamma' = G' \cap (\Gamma \times \Gamma \times \Gamma)$. We now verify that

$$(g(h_1 + h'_1, h_2, \dots, h_k), g(h_1, h_2, \dots, h_k), g(h'_1, h_2, \dots, h_k))$$

is a polynomial sequence with respect to this degree filtration. This is immediate nothing that by Taylor expansion [29, Lemma B.9] and the condition at 0, we have

$$g(h_1, \dots, h_k) = \prod_{(i_1, \dots, i_k) \in \{0,1\}^k \setminus \{\vec{0}\}} g_{i_1, \dots, i_k}^{(h_1) \dots (h_k)}$$

with $g_{i_1, \dots, i_k} \in G_{(i_1, \dots, i_k)}$. The desired polynomiality of the tripled sequence then follows easily from Baker–Campbell–Hausdorff and the fact that the degree of the exponents in the Taylor expansion for the h_1 term is at most 1.

We now construct a Mal'cev basis for G' . Let $G_t = \bigvee_{|\vec{i}|=t} G'_i$ and note that by definition \mathcal{X} is adapted to G_t . We may prove that

$$\begin{aligned} \mathcal{X}' = & \{(X, 0, 0) : X \in \mathcal{X} \cap \log(G_2)\} \cup \{(0, X, 0) : X \in \mathcal{X} \cap \log(G_2)\} \cup \{(0, 0, X) : X \in \mathcal{X} \cap \log(G_2)\} \\ & \cup \{(X, X, 0) : X \in \mathcal{X} \cap \log(G_1) \setminus \mathcal{X} \cap \log(G_2)\} \cup \{(X, 0, X) : X \in \mathcal{X} \cap \log(G_1) \setminus \mathcal{X} \cap \log(G_2)\} \end{aligned}$$

is a weak basis for G' . Furthermore this basis is easily seen to have the nesting property of order k and that all subgroups G'_t are $M^{O_k(d^{O_k(1)})}$ -rational. Thus applying [35, Lemma B.11] we may find a Mal'cev basis $\tilde{\mathcal{X}}$ for G' adapted to the given filtration of complexity $M^{O_k(d^{O_k(1)})}$ and such that all basis elements are $M^{O_k(d^{O_k(1)})}$ -rational combinations of \mathcal{X}' .

The function \tilde{F} we will consider is

$$\tilde{F}((x, y, z)\Gamma^{\otimes 3}) = F(x\Gamma) \otimes \overline{F(y\Gamma)} \otimes \overline{F(z\Gamma)}.$$

This is easily seen to be Lipschitz on $G \times G \times G$ when given the Mal'cev basis $\mathcal{X}^* = \{(X, 0, 0) : X \in \mathcal{X}\} \cup \{(0, X, 0) : X \in \mathcal{X}\} \cup \{(0, 0, X) : X \in \mathcal{X}\}$. As $\tilde{\mathcal{X}}$ has basis elements which are $M^{O_k(d^{O_k(1)})}$ height rational combination of \mathcal{X}^* , we find that \tilde{F} is $M^{O_k(d^{O_k(1)})}$ -Lipschitz on G'/Γ' .

Note that \tilde{F} is invariant under the group G_k^3 and therefore taking the output quotient group G^3/G_k^3 with lattice $\Gamma^3/(\Gamma^3 \cap G_k^3)$ and using Lemma 3.10 completes the proof. \square

We now come to the most technical of the complexity justifications we will need to perform, multilinearization. We will give a rather barebones analysis (citing much from [29, Proposition E.9, E.10]); the reader may find the discussion in [29, pp. 1360-1363] where an extended example is discussed useful. (We prove a slightly weaker statement which is all that is used in the analysis to ease checking extra complexity details.)

Lemma C.5. *Consider nilcharacter $\chi(h_1, \dots, h_k)$ of multidegree (s_1, \dots, s_k) and complexity (M, d) . There exists a multidegree $(1, \dots, 1)$ nilcharacter*

$$\chi'(h_{1,1}, \dots, h_{1,s_1}, h_{2,1}, \dots, h_{2,s_2}, \dots, h_{k,1}, \dots, h_{k,s_k})$$

of complexity $(M^{O_{|\vec{s}|}(d^{O_{|\vec{s}|}(1)})}, d^{O_{|\vec{s}|}(1)})$ such that

$$\chi(h_1, \dots, h_k) \text{ and } \chi'(h_1, \dots, h_1, h_2, \dots, h_2, \dots, h_k, \dots, h_k)$$

are $(M^{O_{|\vec{s}|}(d^{O_{|\vec{s}|}(1)})}, M^{O_{|\vec{s}|}(d^{O_{|\vec{s}|}(1)})}, d^{O_{|\vec{s}|}(1)})$ -equivalent for degree $|\vec{s}| - 1$ and furthermore, for each $1 \leq i \leq k$, χ' is symmetric in the variables $h_{i,1}, \dots, h_{i,s_i}$.

Remark. We will only require the above lemma for multidegree $(1, s - 1)$ nilsequences.

Proof. By Lemma C.3, there exists χ^* such that

$$\chi(h_1, \dots, h_k) \text{ and } \chi^*(h_1, \dots, h_k)^{\otimes \prod_{i=1}^k s_i!}$$

are $(M^{\mathcal{O}_{|\vec{s}|}}(d^{\mathcal{O}_{|\vec{s}|(1)}}), M^{\mathcal{O}_{|\vec{s}|}}(d^{\mathcal{O}_{|\vec{s}|(1)}}), d^{\mathcal{O}_{|\vec{s}|(1)}})$ -equivalent for degree $|\vec{s}| - 1$. Therefore by Lemma 7.4, it suffices to produce χ' such that

$$\chi^*(h_1, \dots, h_k)^{\otimes \prod_{i=1}^k s_i!} \text{ and } \chi'(h_1, \dots, h_1, h_2, \dots, h_2, \dots, h_k, \dots, h_k)$$

are $(M^{\mathcal{O}_{|\vec{s}|}}(d^{\mathcal{O}_{|\vec{s}|(1)}}), M^{\mathcal{O}_{|\vec{s}|}}(d^{\mathcal{O}_{|\vec{s}|(1)}}), d^{\mathcal{O}_{|\vec{s}|(1)}})$ -equivalent for degree $|\vec{s}| - 1$.

Let

$$\chi'(h_1, \dots, h_k) = F(g(h_1, \dots, h_k)\Gamma)$$

with the underlying nilmanifold being G/Γ and the associated Mal'cev basis being \mathcal{X} . Via a standard manipulation which has been performed several times already, we may assume that $g(0) = \text{id}_G$ (at the cost of an insignificant change in parameters). Furthermore assume that η is the vertical character, so

$$F(g_{(s_1, \dots, s_k)}x\Gamma) = e(\eta(g_{(s_1, \dots, s_k)})) \cdot F(x\Gamma).$$

Given $J \subseteq [|\vec{s}|]$, we denote

$$\|J\| := (|J \cap \{s_1 + \dots + s_{i-1} + 1, \dots, s_1 + \dots + s_{i-1} + s_i\}|)_{1 \leq i \leq k}.$$

The group \tilde{G} we will ultimately use to construct our nilsequence will be given by constructing the associated nilpotent Lie algebra. We take

$$\log(\tilde{G}) = \bigoplus_{\emptyset \neq J \subseteq [|\vec{s}|]} \log(G_{\|J\|})$$

and for each $\emptyset \neq J \subseteq [|\vec{s}|]$ let $\iota_J: \log(G_{\|J\|}) \hookrightarrow \log(\tilde{G})$ denote the embedding into the direct sum. We endow \tilde{G} with a Lie bracket such that if $J \cap K \neq \emptyset$ then

$$[\iota_J(x_J), \iota_K(y_K)] = 0$$

and if $J \cap K = \emptyset$ then

$$[\iota_J(x_J), \iota_K(y_K)] = \iota_{J \cup K}([x_J, y_K]),$$

where the bracket between x_J, x_K is taken in the ambient space $\log(G)$ and is seen to lie in $\log(G_{J \cup K})$ by the commutation property of the original filtration on G .

To verify that this gives a valid Lie algebra it suffices to verify this operation is antisymmetric and satisfies the Jacobi relations. Furthermore to verify it suffices to verify these relations on the generators. For antisymmetry for $\iota_J(x_J), \iota_K(y_K)$, if $J \cap K \neq \emptyset$ it is trivial and otherwise

$$[\iota_J(x_J), \iota_K(y_K)] = \iota_{J \cup K}([x_J, y_K]) = -\iota_{J \cup K}([y_J, x_K]) = -[\iota_K(y_K), \iota_J(x_J)]$$

as desired. For the Jacobi identity, when checked on generators $\iota_J(x_J), \iota_K(y_K), \iota_L(z_L)$, if $(J \cap K) \cup (K \cap L) \cup (L \cap J) \neq \emptyset$ the result is trivial. Otherwise we have

$$\begin{aligned} & [\iota_J(x_J), [\iota_K(y_K), \iota_L(z_L)]] + [\iota_K(y_K), [\iota_L(z_L), \iota_J(x_J)]] + [\iota_L(z_L), [\iota_J(x_J), \iota_K(y_K)]] \\ &= \iota_{J \cup K \cup L}([x_J, [y_K, z_L]] + [y_K, [z_L, x_J]] + [z_L, [x_J, y_K]]) = 0 \end{aligned}$$

as desired.

The associated I -filtration with respect to the multidegree ordering is given as follows. For any $(a_1, \dots, a_{|\vec{s}|}) \in \mathbb{N}^{|\vec{s}|}$, let $\log(\tilde{G}_{(a_1, \dots, a_{|\vec{s}|})})$ be the Lie subalgebra of $\log(\tilde{G})$ generated by $\iota_J(x_J)$ for which $1_J(j) \geq a_j$ for each $j = 1, \dots, |\vec{s}|$, and $x_J \in G_{\|J\|}$. It follows this is an I -filtration with respect to the multidegree ordering because for vectors $a, b \in \{0, 1\}^{|\vec{s}|}$, if $1_J(j) \geq a_j$ and $1_K(j) \geq b_j$ we either have $J \cap K \neq \emptyset$ in which case the commutator is trivial or $1_{K \cup J}(j) \geq a_j + b_j$

in which case the result also follows easily. Noting that by construction $\tilde{G} = \bigvee_{\ell=1}^{|\bar{s}|} \tilde{G}_{\bar{e}_\ell}$, the above immediately implies that we have a multidegree $(1, \dots, 1)$ filtration on \tilde{G} .

We now construct a weak basis for \tilde{G} . Recall we have a Mal'cev basis \mathcal{X} for G . Given $\|J\|$, we define the filtration

$$G_{\|J\|}^t = \bigvee_{|\vec{i}|=t} G_{\|J\|+\vec{i}}.$$

Note that $G_{\|J\|} = G_{\|J\|}^0$ and that $G_{\|J\|}^0 = G_{\|J\|}^0 \geq G_{\|J\|}^1 \geq G_{\|J\|}^2 \geq \dots$ is a valid degree filtration when $\|J\| \neq \vec{0}$. Thus by [29, Lemma B.11], we may find a Mal'cev basis $\mathcal{X}^{\|J\|}$ for each $G_{\|J\|}$ which is an $M^{O_{|\bar{s}|}(d^{O_{|\bar{s}|}(1)})}$ -rational combination of \mathcal{X} .

Define

$$\tilde{\mathcal{X}} = \bigcup_{\emptyset \neq J \subseteq [|\bar{s}|]} \iota_J(\mathcal{X}^{\|J\|}).$$

Furthermore define $\tilde{\Gamma}$ to be the group generated by $\exp(L! \cdot \iota_J(\Gamma \cap G_{\|J\|}))$ where L is a sufficiently large constant depending only on $|\bar{s}|$ (and in particular not on M or d). Direct computation with Baker–Campbell–Hausdorff implies that $\tilde{\Gamma} \cap \tilde{G}_{(1, \dots, 1)}$ is contained in $\iota_{[|\bar{s}|]}(\Gamma \cap G_{(s_1, \dots, s_k)})$. Furthermore we see that $\tilde{G}/\tilde{\Gamma}$ is compact, $\tilde{\mathcal{X}}$ is a weak basis of rationality $M^{O_{|\bar{s}|}(d^{O_{|\bar{s}|}(1)})}$ for \tilde{G} , and $\tilde{\mathcal{X}}$ has the degree $O_{|\bar{s}|}(1)$ nesting property. As all groups within the multidegree filtration are $M^{O_{|\bar{s}|}(d^{O_{|\bar{s}|}(1)})}$ -rational with respect to $\tilde{\mathcal{X}}$, by applying [35, Lemma B.11] we may construct a basis with respect to the canonical associated degree filtration of \tilde{G} which certifies that \tilde{G} with the given multidegree filtration has complexity bounded by $M^{O_{|\bar{s}|}(d^{O_{|\bar{s}|}(1)})}$. Furthermore the adapted Mal'cev basis $\tilde{\mathcal{X}}^*$ is an $M^{O_{|\bar{s}|}(d^{O_{|\bar{s}|}(1)})}$ -rational combination of $\tilde{\mathcal{X}}$ (lifted to $\log(\tilde{G})$ appropriately).

We define the $\tilde{G}_{(1, \dots, 1)}$ -vertical frequency as

$$\tilde{\eta}(\exp(\iota_{(1, \dots, 1)}(\log(g_{(s_1, \dots, s_k)})))) := \eta(g_{(s_1, \dots, s_k)})$$

and it is trivial to use the construction of $\tilde{\mathcal{X}}^*$ to certify that $\tilde{\eta}$ has height bounded by $M^{O_{|\bar{s}|}(d^{O_{|\bar{s}|}(1)})}$. We take \tilde{F} to be a nilcharacter with frequency $\tilde{\eta}$ produced by Lemma B.4 (which is applied to the canonical degree filtration of \tilde{G}) and this construction gives output dimension $M^{O_{|\bar{s}|}(d)}$ and Lipschitz constant $M^{O_{|\bar{s}|}(d^{O_{|\bar{s}|}(1)})}$ with respect to $\tilde{\mathcal{X}}^*$.

We now define \tilde{g} . Note that

$$g(h_1, \dots, h_k) = \prod_{\vec{0} \neq (i_1, \dots, i_k) \leq (s_1, \dots, s_k)} (g_{(i_1, \dots, i_k)}^{\text{Tay}})^{h_1^{i_1} \dots h_k^{i_k}}$$

via [29, Lemma B.9] and the condition at 0 to rule out need a coefficient where $(i_1, \dots, i_k) = \vec{0}$. (We are using monomials instead of binomials, which is a minor but easy alteration.) The product here is taken in increasing lexicographic order. We define

$$\tilde{g}(h_1, \dots, h_{|\bar{s}|}) := \prod_{\vec{0} \neq (i_1, \dots, i_k) \leq (s_1, \dots, s_k)} \exp \left(i_1! \dots i_k! \sum_{\substack{J \subseteq \{1, \dots, |\bar{s}|\} \\ \|J\| = (i_1, \dots, i_k)}} \left(\prod_{i \in J} h_i \right) \cdot \iota_J(\log(g_{(i_1, \dots, i_k)}^{\text{Tay}})) \right).$$

Let G^* denote the subgroup of $G \times \tilde{G}$ generated by

$$G^* := \{(g_{(s_1, \dots, s_k)}, \exp(s_1! \dots s_k! \iota_{(1, \dots, 1)}(\log(g_{(s_1, \dots, s_k)})))) : g_{(s_1, \dots, s_k)} \in G_{(s_1, \dots, s_k)}\}.$$

Note that the function

$$(g, \tilde{g}) \mapsto F(g\Gamma)^{\otimes s_1! \dots s_k!} \otimes \overline{F(\tilde{g}\tilde{\Gamma})}$$

is invariant under the action of G^* .

We will construct G' which is a subgroup of $G \times \tilde{G}$ with a degree $|\vec{s}|$ filtration such that the final group is G^* and such that

$$(g(h_1, \dots, h_k), \tilde{g}(h_1, \dots, h_1, h_2, \dots, h_2, \dots, h_k, \dots, h_k))$$

is a polynomial sequence with respect to this filtration. Let G'_j (for $j \geq 1$) be generated by elements of the form

$$(C.1) \quad \left(g_{(i_1, \dots, i_k)}, \exp \left(i_1! \cdots i_k! \sum_{\substack{J \subseteq \{1, \dots, |\vec{s}|\} \\ \|J\| = (i_1, \dots, i_k)}} \iota_J(\log(g_{(i_1, \dots, i_k)})) \right) \right)$$

where $|\vec{i}| = j$, as well as

$$(g_{(i_1, \dots, i_k)}, \text{id}_{\tilde{G}}), \text{ and } (\text{id}_G, \tilde{G}_J)$$

where $|\vec{i}| \geq j + 1$ in the first case, and $|J| \geq j + 1$ in the second case. Furthermore set $G' = G'_0 := G'_1$; we trivially see that $G_{|\vec{s}|} = G^*$. That this is a filtration follows from liberal application of Baker–Campbell–Hausdorff; we use crucially that the number of ways to break a set of size $(i + j)$ into two labeled sets of size i and j which are disjoint is $(i + j)! / (i! \cdot j!)$, which modifies the factorial prefactors in (C.1) appropriately.

Furthermore it is trivial to see that the G'_j are $M^{O_{|\vec{s}|}(d^{O_{|\vec{s}|}(1)})}$ -rational with respect to the Mal'cev basis for $G \times \tilde{G}$ given by

$$\{(X, 0) : X \in \mathcal{X}\} \cup \{(0, X) : X \in \tilde{\mathcal{X}}^*\},$$

and therefore applying [35, Lemma B.11] we may construct a Mal'cev basis \mathcal{X}' of complexity $M^{O_{|\vec{s}|}(d^{O_{|\vec{s}|}(1)})}$ for $G' / (G' \cap (\Gamma \times \tilde{\Gamma}))$. Furthermore $F^{\otimes s_1! \cdots s_k!} \otimes \tilde{F}$ is appropriately Lipschitz with respect to \mathcal{X}' . Finally, since

$$\left(g_{(i_1, \dots, i_k)}, \exp \left(i_1! \cdots i_k! \sum_{\substack{J \subseteq \{1, \dots, |\vec{s}|\} \\ \|J\| = (i_1, \dots, i_k)}} \iota_J(\log(g_{(i_1, \dots, i_k)})) \right) \right)$$

is in $G'_{i_1 + \dots + i_k}$ by definition, we see that

$$(g(h_1, \dots, h_k), \tilde{g}(h_1, \dots, h_1, h_2, \dots, h_2, \dots, h_k, \dots, h_k))$$

is a polynomial sequence with respect to the filtration. Quotienting out by $G^* = G'_{|\vec{s}|}$ (using that $F^{\otimes s_1! \cdots s_k!} \otimes \tilde{F}$ is invariant under G^*) and using Lemma 3.10, we finally complete the proof. \square

We now reach the final technical lemma of the paper which states that a nilsequence of multidegree $J \cup J'$ can be approximated by a sum of products of nilsequences in J and J' . This “splitting” lemma is a quantified version of [29, Lemma E.4]; the proof here is ever so slightly different as we are forced to not use the Stone–Weierstrass theorem.

Lemma C.6. *Let J and J' be finite downsets in \mathbb{N}^k and fix $\varepsilon \in (0, 1/2)$. Suppose that $\beta(h_1, \dots, h_k)$ is a nilsequence of multidegree $J \cup J'$ with complexity (M, d) . Then there exists $1 \leq L \leq (M/\varepsilon)^{O_{J, J'}(d^{O_{J, J'}(1)})}$ such that*

$$\left\| \beta(h_1, \dots, h_k) - \sum_{j=1}^L \beta_j(h_1, \dots, h_k) \beta'_j(h_1, \dots, h_k) \right\|_{L^\infty(\mathbb{Z}^k)} \leq \varepsilon$$

with the β_j being nilsequences of multidegree J , the β'_j being nilsequences of multidegree J' , and β_j, β'_j having complexity $((M/\varepsilon)^{O_{J, J'}(d^{O_{J, J'}(1)})}, d^{O_{J, J'}(1)})$.

Proof. We let

$$\beta(h_1, \dots, h_k) = F(g(h_1, \dots, h_k)\Gamma)$$

where the underlying nilmanifold is G/Γ . As is standard, we may assume that $g(0, \dots, 0) = \text{id}_G$ up to the insignificant change of adjusting M to $M^{O_{J,J'}(d^{O_{J,J'}(1)})}$. Furthermore let the adapted Mal'cev basis for G be \mathcal{X} .

We have for each $\vec{j} \neq \vec{0}$ that the groups

$$G_t^{\vec{j}} = \bigvee_{|\vec{i}|=t} G_{\vec{j}+\vec{i}}$$

form a degree filtration $G_0^{\vec{j}} = G_0^{\vec{j}} \supseteq G_1^{\vec{j}} \supseteq G_2^{\vec{j}} \supseteq \dots$, where the length of the filtration is $O_{J,J'}(1)$. As these subgroups are all M -rational with respect to \mathcal{X} , there exists a Mal'cev basis $\mathcal{X}^{\vec{j}}$ adapted to this filtration of complexity $M^{O_{J,J'}(d^{O_{J,J'}(1)})}$ where each element is an $M^{O_{J,J'}(d^{O_{J,J'}(1)})}$ -rational combination of elements in \mathcal{X} by [35, Lemma B.11].

Using a variant of Lemma 10.2, adapted to multidegree filtrations, we may write

$$g(h_1, \dots, h_k) = \prod_{\vec{j} \neq \vec{0}} \prod_{X_{\vec{j},i} \in \mathcal{X}^{\vec{j}}} \exp(X_{\vec{j},i})^{\alpha_{\vec{j},i} \prod_{\ell=1}^k (h_\ell^{j_\ell} / j_\ell!)}.$$

The product here is taken in \vec{j} is increasing $|\vec{j}|$ and then lexicographic order and $X_{\vec{j},i}$ taken in increasing order of i . The modified proof of such a representation involves iteratively handling terms in increasing order of $|\vec{j}|$ (and handling these terms in an arbitrary order); we omit a careful proof.

The first key part of the proof is lifting to the universal nilmanifold. We define the universal nilmanifold \tilde{G} to be generated by generators $\exp(e_{\vec{j},i}^{\vec{j}})^{t_{\vec{j},i}}$ for $\vec{j} \neq \vec{0}$, $1 \leq i \leq \dim(G_{\vec{j}})$, and $t_{\vec{j},i} \in \mathbb{R}$. The only relations these generators satisfy is that any $(r-1)$ -fold commutator for $r \geq 1$ between $\exp(e_{\vec{j}_1, i_1}^{\vec{j}_1}), \dots, \exp(e_{\vec{j}_r, i_r}^{\vec{j}_r})$ vanishes if $\vec{j}_1 + \dots + \vec{j}_r$ is not in $J \cup J'$. We give \tilde{G} the structure of a multidegree $J \cup J'$ nilmanifold by letting $(\tilde{G})_{\vec{j}^*}$ be generated by the set of $(r-1)$ -fold commutators (for any $r \geq 1$) of $\exp(e_{\vec{j}_1, i_1}^{\vec{j}_1}), \dots, \exp(e_{\vec{j}_r, i_r}^{\vec{j}_r})$ where $\vec{j}_1 + \dots + \vec{j}_r \geq \vec{j}^*$ (here \geq means that each coordinate is larger). This is easily proven to be an I -filtration with respect to the multidegree ordering and note that since we have no generators with $\vec{j} = \vec{0}$, this is in fact a multidegree filtration. Finally we let $\tilde{\Gamma}$ be the lattice generated by $\exp(e_{\vec{j},i}^{\vec{j}})$.

The analysis in Lemma 10.4 can easily be extended to prove that \tilde{G} has a filtered Mal'cev basis $\tilde{\mathcal{X}}$ of complexity $M^{O_{J,J'}(d^{O_{J,J'}(1)})}$ where the basis elements are height $M^{O_{J,J'}(d^{O_{J,J'}(1)})}$ linear combinations of $(r-1)$ -fold commutators of $e_{\vec{j}_1, i_1}^{\vec{j}_1}, \dots, e_{\vec{j}_r, i_r}^{\vec{j}_r}$. Furthermore note that the dimension of \tilde{G} is $d^{O_{J,J'}(1)}$.

We now lift β to $\tilde{G}/\tilde{\Gamma}$. Define the homomorphism $\phi: \tilde{G} \rightarrow G$ via

$$\phi(\exp(e_{\vec{j},i}^{\vec{j}})) = \exp(X_{\vec{j},i});$$

here we are writing $\mathcal{X}^{\vec{j}} = \{X_{\vec{j},1}, \dots, X_{\vec{j}, \dim(G_{\vec{j}})}\}$. That this is a homomorphism follows from noting that all relations in \tilde{G} are present in G because G has multidegree $J \cup J'$. We next lift the polynomial sequence g to

$$\tilde{g}(h_1, \dots, h_k) = \prod_{\vec{j} \neq \vec{0}} \prod_{i=1}^{\dim(G_{\vec{j}})} \exp(e_{\vec{j},i}^{\vec{j}})^{\alpha_{i,\vec{j}} \prod_{\ell=1}^k (h_\ell^{j_\ell} / j_\ell!)}$$

and F to \tilde{F} via

$$\tilde{F}(\tilde{g}\tilde{\Gamma}) = F(\phi(\tilde{g})\Gamma).$$

Note that since $\phi(\tilde{\Gamma}) \leq \Gamma$, this is a well-defined function on $\tilde{G}/\tilde{\Gamma}$. Furthermore, noting various properties of $\tilde{\mathcal{X}}$ and that elements of $\mathcal{X}^{\vec{j}}$ are appropriately bounded and rational linear combinations of $\tilde{\mathcal{X}}$, we have that \tilde{F} is $M^{O_{J,J'}(d^{O_{J,J'}(1)})}$ -Lipschitz with respect to the Mal'cev basis specified by $\tilde{\mathcal{X}}$. Therefore for the remainder of the proof we operate with the nilsequence

$$\tilde{F}(\tilde{g}(h_1, \dots, h_k)\tilde{\Gamma}).$$

For the remainder of the analysis we furthermore assume that there exists $g^* \in \tilde{G}$ with $d_{\tilde{G}, \tilde{\mathcal{X}}}(g^*) \leq M^{O_{J,J'}(d^{O_{J,J'}(1)})}$ such that if $\psi_{\exp, \tilde{G}}(g^*) + (-1/2, 1/2]^{\dim(G)}$ is identified with $\tilde{G}/\tilde{\Gamma}$ then $\text{supp}(\tilde{F})$ lies in $\psi_{\exp, \tilde{G}}(g^*) + (-\delta, \delta]^{\dim(G)}$. We will ultimately take $\delta = M^{-O_{J,J'}(d^{O_{J,J'}(1)})}$ sufficiently small. If we prove the proposition with $\varepsilon' = \varepsilon \cdot \delta^{O_{J,J'}(d^{O_{J,J'}(1)})}$ for functions with such restricted support then the result in generality follows by Lemma B.3.

The use of the universal nilmanifold comes precisely when defining the following two nilmanifolds for the split terms. Let $\tilde{G}_{>J}$ be the group generated by $\tilde{G}_{\vec{i}}$ with $\vec{i} \in J' \setminus J$ and $\tilde{G}_{>J'}$ be the group generated by $\tilde{G}_{\vec{i}}$ with $\vec{i} \in J \setminus J'$. It is trivial to see that $\tilde{G}_{>J}, \tilde{G}_{>J'}$ are normal and $M^{O_{J,J'}(d^{O_{J,J'}(1)})}$ -rational with respect to $\tilde{\mathcal{X}}$. Let $\tilde{\mathcal{X}}^J$ and $\tilde{\mathcal{X}}^{J'}$ be bases for the Lie algebras of $\log(\tilde{G}_{>J})$ and $\log(\tilde{G}_{>J'})$ which are $M^{O_{J,J'}(d^{O_{J,J'}(1)})}$ -rational bounded combinations of $\tilde{\mathcal{X}}$.

We consider the nilmanifolds $\tilde{G}/(\tilde{G}_{>J}\tilde{\Gamma})$ and $\tilde{G}/(\tilde{G}_{>J'}\tilde{\Gamma})$. The first is clearly a multidegree J' nilmanifold while the second is a multidegree J nilmanifold, each of complexity $M^{O_{J,J'}(d^{O_{J,J'}(1)})}$. Furthermore we can choose underlying Mal'cev bases which are $M^{O_{J,J'}(d^{O_{J,J'}(1)})}$ -rational combinations of $\tilde{\mathcal{X}} \bmod \tilde{\mathcal{X}}^J$ and $\tilde{\mathcal{X}} \bmod \tilde{\mathcal{X}}^{J'}$, respectively. (See, e.g., the arguments regarding G_{Quot} in Section 10.3.) Note here that $\tilde{\Gamma}_{>J} = \tilde{\Gamma}/(\tilde{\Gamma} \cap \tilde{G}_{>J})$ and analogously for $\tilde{\Gamma}_{>J'}$.

The key point is that by construction, $\tilde{G}_{>J} \cap \tilde{G}_{>J'} = \text{Id}_{\tilde{G}}$. This implies that there exist linear maps A and B such that

$$(C.2) \quad A \circ \psi_{\exp, \tilde{G}/\tilde{G}_{>J}}(z \bmod \tilde{G}_{>J}) + B \circ \psi_{\exp, \tilde{G}/\tilde{G}_{>J'}}(z \bmod \tilde{G}_{>J'}) = \psi_{\exp, \tilde{G}}(z)$$

for all $z \in \tilde{G}$. Furthermore one can take A and B bounded in the sense that

$$d_{\tilde{G}}(A \circ \psi_{\exp, \tilde{G}/\tilde{G}_{>J}}(\exp(\tilde{X}_i) \bmod \tilde{G}_{>J}), \text{id}_{\tilde{G}}) \leq M^{O_{J,J'}(d^{O_{J,J'}(1)})}$$

for all $\tilde{X}_i \in \tilde{\mathcal{X}}$ and analogously for B and J' .

We now identify $\tilde{G}/\tilde{\Gamma}$ via $\psi_{\tilde{G}}$ with the domain $\psi_{\exp, \tilde{G}}(g^*) + (-1/2, 1/2]^{\dim(\tilde{G})}$ and we only have support of \tilde{F} in $\psi_{\exp, \tilde{G}}(g^*) + (-\delta, \delta]^{\dim(\tilde{G})}$. Given $x \in \tilde{G}$ such that $\psi_{\exp, \tilde{G}}(x) \in \psi_{\exp, \tilde{G}}(g^*) + (-\delta, \delta]^{\dim(\tilde{G})}$, we have that

$$\begin{aligned} \psi_{\exp, \tilde{G}/\tilde{G}_{>J}}(x \bmod \tilde{G}_{>J}) &\in \psi_{\exp, \tilde{G}/\tilde{G}_{>J}}(g^* \bmod \tilde{G}_{>J}) + (-\delta, \delta]^{\dim(\tilde{G}/\tilde{G}_{>J})} \cdot M^{O_{J,J'}(d^{O_{J,J'}(1)})}, \\ \psi_{\exp, \tilde{G}/\tilde{G}_{>J'}}(x \bmod \tilde{G}_{>J'}) &\in \psi_{\exp, \tilde{G}/\tilde{G}_{>J'}}(g^* \bmod \tilde{G}_{>J'}) + (-\delta, \delta]^{\dim(\tilde{G}/\tilde{G}_{>J'})} \cdot M^{O_{J,J'}(d^{O_{J,J'}(1)})}. \end{aligned}$$

Given that $\delta = M^{-O_{J,J'}(d^{O_{J,J'}(1)})}$ is sufficiently small, these are contained

$$\begin{aligned} &\psi_{\exp, \tilde{G}/\tilde{G}_{>J}}(g^* \bmod \tilde{G}_{>J}) + (-1/4, 1/4]^{\dim(\tilde{G}/\tilde{G}_{>J})}, \\ &\psi_{\exp, \tilde{G}/\tilde{G}_{>J'}}(g^* \bmod \tilde{G}_{>J'}) + (-1/4, 1/4]^{\dim(\tilde{G}/\tilde{G}_{>J'})}, \end{aligned}$$

respectively.

Identify $\psi_{\exp, \tilde{G}}(g^*) + (-1/2, 1/2]^{\dim(G)}$ with the torus (note that the boundaries are glued differently than in $\tilde{G}/\tilde{\Gamma}$, but we are near the center so it is not an issue). We have that \tilde{F} is an $(M/\delta)^{O_{J,J'}(d^{O_{J,J'}(1)})}$ -Lipschitz function with respect to the standard torus metric (see

e.g. [38, Lemma 2.3] and [35, Lemma B.3]). Thus for $x \in \tilde{G}$ such that $\psi_{\exp, \tilde{G}}(x) \in \psi_{\exp, \tilde{G}}(g^*) + (-1/2, 1/2]^{\dim(G)}$, via standard Fourier approximation (see e.g. [42, Lemma A.8]), for $\xi \in \mathbb{Z}^{\dim(\tilde{G})}$ there exist c_ξ with $|c_\xi| \leq (M/(\delta\varepsilon'))^{O(\dim(\tilde{G}))}$ such that

$$\left\| \tilde{F}(x\tilde{\Gamma}) - \sum_{\|\xi\|_\infty \leq (M/(\delta\varepsilon'))^{O(\dim(\tilde{G}))}} c_\xi e(\xi \cdot \psi_{\exp}(x)) \right\|_\infty \leq \varepsilon'$$

where the sum is over $\xi \in \mathbb{Z}^{\dim(\tilde{G})}$. Using (C.2) we may write this equivalently as

$$\left\| \tilde{F}(x\tilde{\Gamma}) - \sum_{\xi} c_\xi e(\xi \cdot (A \circ \psi_{\exp, \tilde{G}_{>J}}(x \bmod \tilde{G}_{>J})) e(\xi \cdot (B \circ \psi_{\exp, \tilde{G}_{>J'}}(x \bmod \tilde{G}_{>J'}))) \right\|_\infty \leq \varepsilon',$$

where again the sum is over $\|\xi\|_\infty \leq (M/(\delta\varepsilon'))^{O(\dim(\tilde{G}))}$.

For z such that $\psi_{\exp, \tilde{G}/\tilde{G}_{>J}}(z) - \psi_{\exp, \tilde{G}/\tilde{G}_{>J}}(g^* \bmod \tilde{G}_{>J}) \in (-1/2, 1/2]^{\dim(\tilde{G}/\tilde{G}_{>J})}$, we let

$$\tau_{>J, \xi}(z) = \rho(\|\psi_{\exp, \tilde{G}/\tilde{G}_{>J}}(z) - \psi_{\exp, \tilde{G}/\tilde{G}_{>J}}(g^* \bmod \tilde{G}_{>J})\|) \cdot e(\xi \cdot (A \circ \psi_{\exp, \tilde{G}_{>J}}(z)))$$

with $\rho(x) = 1$ for $|x| \leq 1/4$, $\rho(x) = 0$ for $|x| \geq 1/3$, and ρ is $O(1)$ -Lipschitz and extends to $\tilde{G}/(\tilde{G}_{>J}\tilde{\Gamma})$ via periodicity. $\tau_{>J}$ is seen to be an $(M/(\delta\varepsilon'))^{O_{J, J'}(d^{O_{J, J'}(1)})}$ -Lipschitz function on $\tilde{G}/(\tilde{G}_{>J}\tilde{\Gamma})$. This follows via the size of ξ and that distance in $d_{\tilde{G}}$ controls distance in first-kind coordinates (see e.g. [35, Lemmas B.1, B.3]). Define $\tau_{>J', \xi}$ in the same manner. We have that

$$\left\| \tilde{F}(x\tilde{\Gamma}) - \sum_{\|\xi\|_\infty \leq (M/(\delta\varepsilon'))^{O(\dim(\tilde{G}))}} c_\xi \tau_{>J, \xi}((x \bmod \tilde{G}_{>J})\tilde{\Gamma}_{>J}) \tau_{>J', \xi}((x \bmod \tilde{G}_{>J'})\tilde{\Gamma}_{>J'}) \right\|_\infty \leq \varepsilon'.$$

As this holds for all x such that $\psi_{\exp, \tilde{G}}(x) \in \psi_{\exp, \tilde{G}}(g^*) + (-1/2, 1/2]^{\dim(G)}$ and the approximating function is invariant under $\tilde{\Gamma}$, this holds for all $x \in \tilde{G}$. This completes the proof, plugging in $x = \tilde{g}(h_1, \dots, h_k)$ and noting that $\tilde{g} \bmod \tilde{G}_{>J}$ and $\tilde{g} \bmod \tilde{G}_{>J'}$ are multidegree J' and J polynomial sequences on $\tilde{G}/(\tilde{G}_{>J}\tilde{\Gamma})$ and $\tilde{G}/(\tilde{G}_{>J'}\tilde{\Gamma})$ respectively. \square

REFERENCES

- [1] D. Altman, *A non-flag arithmetic regularity lemma and counting lemma*, arXiv:2209.14083.
- [2] D. Altman, *On a conjecture of Gowers and Wolf*, Discrete Anal. (2022), Paper No. 10, 13.
- [3] V. Bergelson, B. Host, and B. Kra, *Multiple recurrence and nilsequences*, Invent. Math. **160** (2005), 261–303, With an appendix by Imre Ruzsa.
- [4] T. F. Bloom and O. Sisask, *An improvement to the Kelley-Meka bounds on three-term arithmetic progressions*, arXiv:2309.02353.
- [5] O. A. Camarena and B. Szegedy, *Nilspaces, nilmanifolds and their morphisms*, arXiv:1009.3825.
- [6] J.-P. Conze and E. Lesigne, *Théorèmes ergodiques pour des mesures diagonales*, Bull. Soc. Math. France **112** (1984), 143–175.
- [7] L. J. Corwin and F. P. Greenleaf, *Representations of nilpotent Lie groups and their applications. Part I*, Cambridge Studies in Advanced Mathematics, vol. 18, Cambridge University Press, Cambridge, 1990, Basic theory and examples.
- [8] P. Erdős and P. Turán, *On Some Sequences of Integers*, J. London Math. Soc. **11** (1936), 261–264.
- [9] H. Furstenberg, *Ergodic behavior of diagonal measures and a theorem of Szemerédi on arithmetic progressions*, J. Analyse Math. **31** (1977), 204–256.
- [10] H. Furstenberg and B. Weiss, *A mean ergodic theorem for $(1/N) \sum_{n=1}^N f(T^n x)g(T^{n^2} x)$* , Convergence in ergodic theory and probability (Columbus, OH, 1993), Ohio State Univ. Math. Res. Inst. Publ., vol. 5, de Gruyter, Berlin, 1996, pp. 193–227.
- [11] W. T. Gowers, *A new proof of Szemerédi’s theorem for arithmetic progressions of length four*, Geom. Funct. Anal. **8** (1998), 529–551.
- [12] W. T. Gowers, *Arithmetic progressions in sparse sets*, Current developments in mathematics, 2000, Int. Press, Somerville, MA, 2001, pp. 149–196.
- [13] W. T. Gowers, *A new proof of Szemerédi’s theorem*, Geom. Funct. Anal. **11** (2001), 465–588.

- [14] W. T. Gowers and L. Milićević, *An inverse theorem for Freiman multi-homomorphisms*, arXiv:2002.11667.
- [15] W. T. Gowers and L. Milićević, *A quantitative inverse theorem for the U^4 norm over finite fields*, arXiv:1712.00241.
- [16] W. T. Gowers and J. Wolf, *The true complexity of a system of linear equations*, Proc. Lond. Math. Soc. (3) **100** (2010), 155–176.
- [17] B. Green, *100 open problems*, manuscript, available on request.
- [18] B. Green and T. Tao, *An inverse theorem for the Gowers $U^3(G)$ norm*, Proc. Edinb. Math. Soc. (2) **51** (2008), 73–153.
- [19] B. Green and T. Tao, *The primes contain arbitrarily long arithmetic progressions*, Ann. of Math. (2) **167** (2008), 481–547.
- [20] B. Green and T. Tao, *New bounds for Szemerédi’s theorem. II. A new bound for $r_4(N)$* , Analytic number theory, Cambridge Univ. Press, Cambridge, 2009, pp. 180–204.
- [21] B. Green and T. Tao, *An arithmetic regularity lemma, an associated counting lemma, and applications*, An irregular mind, Bolyai Soc. Math. Stud., vol. 21, János Bolyai Math. Soc., Budapest, 2010, pp. 261–334.
- [22] B. Green and T. Tao, *Linear equations in primes*, Ann. of Math. (2) **171** (2010), 1753–1850.
- [23] B. Green and T. Tao, *The Möbius function is strongly orthogonal to nilsequences*, Ann. of Math. (2) **175** (2012), 541–566.
- [24] B. Green and T. Tao, *The quantitative behaviour of polynomial orbits on nilmanifolds*, Ann. of Math. (2) **175** (2012), 465–540.
- [25] B. Green and T. Tao, *New bounds for Szemerédi’s theorem, III: a polylogarithmic bound for $r_4(N)$* , Matematika **63** (2017), 944–1040.
- [26] B. Green, T. Tao, and T. Ziegler, *Erratum for “An inverse theorem for the Gowers $U^{s+1}[N]$ -norm”*, manuscript.
- [27] B. Green, T. Tao, and T. Ziegler, *An inverse theorem for the Gowers U^4 -norm*, Glasg. Math. J. **53** (2011), 1–50.
- [28] B. Green, T. Tao, and T. Ziegler, *An inverse theorem for the Gowers $U^{s+1}[N]$ -norm*, Electron. Res. Announc. Math. Sci. **18** (2011), 69–90.
- [29] B. Green, T. Tao, and T. Ziegler, *An inverse theorem for the Gowers $U^{s+1}[N]$ -norm*, Ann. of Math. (2) **176** (2012), 1231–1372.
- [30] Y. Gutman, F. W. R. M. Manners, and P. P. Varjú, *The structure theory of nilspaces II: Representation as nilmanifolds*, Trans. Amer. Math. Soc. **371** (2019), 4951–4992.
- [31] Y. Gutman, F. W. R. M. Manners, and P. P. Varjú, *The structure theory of nilspaces I*, J. Anal. Math. **140** (2020), 299–369.
- [32] Y. Gutman, F. W. R. M. Manners, and P. P. Varjú, *The structure theory of nilspaces III: Inverse limit representations and topological dynamics*, Adv. Math. **365** (2020), 107059, 53.
- [33] B. Host and B. Kra, *Nonconventional ergodic averages and nilmanifolds*, Ann. of Math. (2) **161** (2005), 397–488.
- [34] Z. Kelley and R. Meka, *Strong bounds for 3-progressions*, arXiv:2302.05537.
- [35] J. Leng, *Efficient Equidistribution of Nilsequences*, arXiv:2312.10772.
- [36] J. Leng, *Efficient Equidistribution of Periodic Nilsequences and Applications*, arXiv:2306.13820.
- [37] J. Leng, *Improved Quadratic Gowers Uniformity for the Möbius Function*, arXiv:2212.09635.
- [38] J. Leng, A. Sah, and M. Sawhney, *Improved bounds for five-term arithmetic progressions*, arXiv:2312.10776.
- [39] J. Leng, A. Sah, and M. Sawhney, *Improved Bounds for Szemerédi’s Theorem*, manuscript.
- [40] F. W. R. M. Manners, *Quantitative bounds in the inverse theorem for the Gowers U^{s+1} -norms over cyclic groups*, arXiv:1811.00718.
- [41] L. Milićević, *Bilinear Bogolyubov Argument in Abelian Groups*, arXiv:2109.03093.
- [42] S. Peluse, A. Sah, and M. Sawhney, *Effective bounds for Roth’s theorem with shifted square common difference*, arXiv:2309.08359.
- [43] K. F. Roth, *On certain sets of integers. II*, J. London Math. Soc. **29** (1954), 20–26.
- [44] T. Sanders, *On certain other sets of integers*, J. Anal. Math. **116** (2012), 53–82.
- [45] T. Sanders, *On the Bogolyubov-Ruzsa lemma*, Anal. PDE **5** (2012), 627–655.
- [46] B. Szegedy, *On higher order Fourier analysis*, arXiv:1203.2260.
- [47] E. Szemerédi, *On sets of integers containing no four elements in arithmetic progression*, Number Theory (Colloq., János Bolyai Math. Soc., Debrecen, 1968), Colloq. Math. Soc. János Bolyai, vol. 2, North-Holland, Amsterdam-London, 1970, pp. 197–204.
- [48] E. Szemerédi, *On sets of integers containing no k elements in arithmetic progression*, Acta Arith. **27** (1975), 199–245.

- [49] T. Tao, *Goursat and Furstenberg–Weiss type lemmas*, 2021, blog post. <https://terrytao.wordpress.com/2021/05/07/goursat-and-furstenberg-weiss-type-lemmas/>.
- [50] T. Tao and J. Teräväinen, *Quantitative bounds for Gowers uniformity of the Möbius and von Mangoldt functions*, arXiv:2107.02158.
- [51] T. Tao and V. H. Vu, *Additive combinatorics*, Cambridge Studies in Advanced Mathematics, vol. 105, Cambridge University Press, Cambridge, 2010, Paperback edition [of MR2289012].
- [52] T. Tao and T. Ziegler, *Polynomial patterns in the primes*, Forum Math. Pi **6** (2018), e1, 60.
- [53] T. Ziegler, *Universal characteristic factors and Furstenberg averages*, J. Amer. Math. Soc. **20** (2007), 53–97.

DEPARTMENT OF MATHEMATICS, UCLA, LOS ANGELES, CA 90095, USA
Email address: jamesleng@math.ucla.edu

DEPARTMENT OF MATHEMATICS, MASSACHUSETTS INSTITUTE OF TECHNOLOGY, CAMBRIDGE, MA 02139, USA
Email address: {asah,msawhney}@mit.edu