A Product Theorem in Free Groups

Alexander A. Razborov *

September 29, 2007

Abstract

If A is a finite subset of a free group with at least two non-commuting elements then $|A \cdot A \cdot A| \geq \frac{|A|^2}{(\log |A|)^{O(1)}}$. More generally, the same conclusion holds in an arbitrary virtually free group, unless A generates a virtually cyclic subgroup.

The central part of the proof of this result is carried on by estimating the number of collisions in multiple products $A_1 \cdot \ldots \cdot A_k$. We include a few simple observations showing that in this "statistical" context the analogue of the fundamental Plünnecke-Ruzsa theory looks particularly simple and appealing.

1. Introduction

Let G be a group, and A be its finite subset. Assume that for some fixed $k \geq 2$, $|\underbrace{A \cdot A \cdot \ldots \cdot A}_{k \text{ times}}|$ (where the *product set* $\underbrace{A \cdot A \cdot \ldots \cdot A}_{k \text{ times}}$ is defined as $\{b \in G \mid (\exists a_1, \ldots, a_k \in A)(b = a_1 a_2 \ldots a_k)\}$) is much smaller than $|A|^k$. What can be said about the internal structure of A?

Questions of this (and similar) sort are known in arithmetic combinatorics as *inverse problems* (most of the material briefly surveyed in this section can be found in comprehensive monographs [13, 16]). Originally they were studied for $G = \mathbb{Z}$ (the case $G = \mathbb{Z}^n$ is easily seen to be "essentially equivalent" to this one). And one of the deepest and hardest results in the area is Freiman's

*Institute for Advanced Study, Princeton, US, on leave from Steklov Mathematical Institute, Moscow, Russia, razborov@ias.edu. Supported by the NSF grant ITR-0324906 and by the Russian Foundation for Basic Research.

theorem [21] that provides a complete characterization of sets $A \subseteq \mathbb{Z}$ with $|A + A + \ldots + A| \leq O(|A|)$.

For many applications, however, it is highly desirable to be able to infer at least something intelligent about the structure of A from the weaker assumption

$$|\underbrace{A + A + \ldots + A}_{k \text{ times}}| \leq |A|^{1+o(1)}. \tag{1}$$

And for the case of abelian groups this is a widely open problem (perhaps, *the* central problem in the whole area). This state of the art is particularly embarrassing given the amount of useful information one *can* extract from (1) with the help of powerful *Plünnecke-Ruzsa theory*. As one of the most cited corollaries, let us just mention that the conditions (1) are equivalent for all (fixed) $k \geq 2$, and, moreover, this equivalence still holds if some pluses are replaced by minuses. Further, (1) follows from $|A + B| \leq |A|^{o(1)}|B|$ for an *arbitrary* set B with $|B| \leq |A|$. Unfortunately, these powerful conclusions tell us very little about the *internal* structure of A.

Somewhat surprisingly, inverse problems have turned out to be simpler for more complicated algebraic structures. For example, sum-product estimates in commutative rings by Bourgain, Katz and Tao [3] do give strong inverse results in the range (1) if we append the analogous restriction $|A \cdot A \cdot \ldots \cdot A| \leq |A|^{1+o(1)}$ for product sets.

In this paper we are interested in another class of algebraic structures that has recently sparkled a considerable attention, the class of *non-abelian* groups [7, 15, 4]. One of the reasons for this interest lies in the motivations of the pioneering paper by Helfgott [7] that linked this kind of questions to estimating the diameter of Cayley graphs in certain finite groups, and, via this, to difficult open problems about explicit constructions of expanders. But before reviewing these latest developments, it is worth mentioning that for groups equipped with a length function very similar problems were studied long before, in quite a different context and in a different community. Specifically, *Rapid Decay Property* [6, 8] implies that any set A satisfying (1) (or, in fact, the weaker assumption $|A \cdot A| \leq A^{2-\Omega(1)}$) can not be positioned within a small ball, and must necessarily contain elements of length $|A|^{\Omega(1)}$. Among others, this property is known for free groups [6], groups of polynomial growth and hyperbolic groups [8].

An easy example shows that the Plünnecke-Ruzsa theory does not literally transfer to the non-abelian case: $|A \cdot A|$ can be small, whereas already $|A \cdot A \cdot A|$

is large. Tao [15] and Helfgott [7], however, proved that this theory catches up already at the next level: say, the statements (1) become equivalent for $k = 3, 4, \ldots$ For this reason in the non-abelian case it does make sense to concentrate on the study of sets A with *small tripling* (that is, $k = 3$), as opposed to sets with *small doubling* in the abelian case. And Helfgott [7] indeed proved a strong inverse result for tripling in the range (1) when $G = SL_2(\mathbb{Z}_p)$. Chang [4] proved a similar theorem for $G = SL_2(\mathbb{C})$, and made a very substantial step toward obtaining an analogous result for $G = SL_3(\mathbb{Z})$.

Chang's former result (for $SL_2(\mathbb{C})$) looks in fact rather intriguing since it exhibits the following "threshold behaviour". There exists a *fixed* constant $\delta > 0$ such that the structural conclusion she gets from $|A \cdot A \cdot A| \leq |A|^{1+\delta}$ is *exactly* the same as the conclusion one gets from much stronger bound (1): A generates a virtually abelian sub-group (this reduces the inverse problem for $SL_2(\mathbb{C})$ to the same problem for abelian groups – the best we can hope for without actually solving the latter!) This is very unusual for arithmetic combinatorics where the conclusion usually depends on things like $|A \cdot A|$ or $|A \cdot A \cdot A|$ *numerically* and *smoothly*. Chang also remarked that the same conclusion holds (via any known embedding of F_m into $SL_2(\mathbb{C})$) for free groups F_m,[1] asked for a purely combinatorial proof of this fact and for any estimates of the threshold constant δ.

The main result of our paper provides an answer to her question, and we show that in fact $\delta = 1$ (which is clearly optimal). More precisely, we prove the following:

Main Theorem. *Let A be a finite subset of a free group F_m with at least two non-commuting elements. Then*

$$|A \cdot A \cdot A| \geq \frac{|A|^2}{(\log |A|)^{O(1)}}.$$

More generally, the same conclusion holds for any finite subset A of an arbitrary fixed virtually free group, unless the subgroup generated by A is virtually cyclic. In particular, this is true for the modular group $PSL_2(\mathbb{Z})$, as well as for $SL_2(\mathbb{Z})$ and $GL_2(\mathbb{Z})$, and this makes an improvement over [4, Theorem 5.1] (the latter gave the bound $|A \cdot A \cdot A| \geq |A|^{1+\delta}$ for $SL_2(\mathbb{Z})$ and for an unspecified constant $\delta > 0$.)

[1] Breuillard (personal communication) observed that this can be derived already from the work of Helfgott [7]

Our proof is heavily based on the machinery of combinatorial group theory, and, more specifically, its part known as the theory of (highly) periodic words. It is worth noting that this theory lies in the heart of two of the deepest (and *extremally* involved technically) achievements in that area: the work on Burnside problem [17], and the work on equations in free groups [11, 18, 19, 20] that has recently culminated in independent solutions of Tarski's problem given by Kharlampovich-Miasnikov [9, 10] and Sela [14].

Instead of lower bounds on the cardinalities of sum/product sets, it is often more convenient to go after *upper* bounds on the dual quantities[2] defined like

$$\mathbf{c}(A,B) \stackrel{\text{def}}{=} \left|\left\{ (a,b,a',b') \in (A\times B)^2 \mid ab = a'b' \right\}\right|.$$

These collision numbers are related to the cardinalities of sum/product sets via a simple Cauchy-Schwartz by

$$|A\cdot B| \geq \frac{|A|^2|B|^2}{\mathbf{c}(A,B)},$$

but display much more analytical (and in many cases more convenient) behaviour than $|A\cdot B|$. The Balog-Szemerédi-Gowers theorem shows how to go in the opposite direction (from large $\mathbf{c}(A,B)$ to large subsets $A_0 \subseteq A,\ B_0 \subseteq B$ with small $|A_0 \cdot B_0|$) without losing too much. But we would also like to note that one of the most striking recent applications of arithmetic combinatorics [1, 2] actually *needs* upper bounds on collision numbers/probabilities rather than lower bounds on the size of sum/product sets.

The most crucial part of our argument (contained in Section 5) also works entirely in this framework (that we, following [1] once more, will call *statistical*) and essentially utilizes all its versatility. This has motivated us to wonder how far we can get in the world in which all quantities like $|A_1\cdot A_2\cdot \ldots \cdot A_k|$ are *systematically* replaced by their statistical counterparts $\mathbf{c}(A_1,\ldots,A_k)$. We contribute to this a few simple remarks showing that the statistical version of Plünnecke-Ruzsa theory looks particularly simple and appealing, without ever mentioning cardinalities $|A_1 \cdot A_2 \cdot \ldots \cdot A_k|$, Menger's theorem or Ruzsa's covering lemma inherent to its "classical" versions.

[2]Accordingly, they appeared in the literature under many different names, e.g. *quadruples* [5] or *additive energy* [15, 16]. In order to stress our purely combinatorial treatment, we prefer to follow the lead of [1] and call them *collision numbers* or, after appropriate normalization, *collision probabilities*.

These remarks are given in the concluding Section 6, and all the preceding part of the paper is entirely devoted to the proof of Main Theorem. In Section 2 we give the necessary background, mostly from combinatorial group theory. In Section 3 we get rid of cancellations, and also show that when lower bounding $|A \cdot B \cdot C|$ in a free semi-group, we can assume w.l.o.g. that A is a prefix chain, and C is a suffix chain. In Section 4 we further reduce our problem to the case when the triple (A, B, C) has "enough aperiodicity" in it. And then in Section 5 comes the central part of our proof: we upper bound the collision numbers $\mathbf{c}(A, B, C)$, ruling out the only unpleasant case with the help of "aperiodicity constraints" enforced in the previous Section 4.

2. Background

All the material in this section related to the combinatorial group theory can be found e.g. in [12, 17].

We let $[n] \stackrel{\text{def}}{=} \{1, 2, \ldots, n\}$.

Let F_m be the free group with the basis $\{x_1, \ldots, x_m\}$. A word w in the alphabet $\{x_1, x_1^{-1}, \ldots, x_m, x_m^{-1}\}$ is *reduced* if for any $i \in [m]$, x_i and x_i^{-1} never appear in w as adjacent letters. The elements of F_m are in one-to-one correspondence with the set of reduced words, and we will always represent them in this form. The unit element is the empty word, denoted by Λ. $|w|$ is the length of the word w.

The notation $\equiv$ stands for *graphical* (or *letter-for-letter*) equality: for $u_1, \ldots, u_r, v_1, \ldots, v_s \in F_m$, $u_1 u_2 \ldots u_r \equiv v_1 v_2 \ldots v_s$ by definition means that $u_1 u_2 \ldots u_r = v_1 v_2 \ldots v_s$ in F_m *and* both words $u_1 u_2 \ldots u_r$, $v_1 v_2 \ldots v_s$ are reduced. In the opposite direction, $u = vw$ in F_m if and only if there exist (uniquely defined) $v', c, w' \in F_m$ such that $v \equiv v'c$, $w \equiv c^{-1} w'$ and $u \equiv v'w'$. In this case we say that the word c is the *cancellation* (or *gets canceled*) in the product vw. If $c = v$ [$c = w^{-1}$], then we say that v [w, respectively] *gets completely canceled* in this product. And if $c = \Lambda$, then we say that *there is no cancellation* in vw, or that vw is *reduced*.

A word v is a *subword* of u, denoted $v \subseteq u$, if there exist words L, R such that $u \equiv LvR$. Any such representation is called an *occurrence* of v into u, and L, R are called *wings* of this occurrence. If $L = \Lambda$ [$R = \Lambda$] then we say that u *begins with* v, or that v *is a prefix of* u [u *ends with* v/v *is a suffix of* u, respectively]. A prefix or a suffix v of u is *proper* if $v \neq u$. We let $a \leq b$

denote that a is a prefix of b. This is a partial ordering on the set of all reduced words called the *prefix order*. Let $a \leq^* b$ be the dual *suffix order*.

A reduced word w is *cyclically reduced* if w^2 (and, hence, also all higher powers w^s) is reduced. Two cyclically reduced words u, v are *cyclic shifts* of each other, denoted $u \sim v$, if for some w_1, w_2 we have

$$u \equiv w_1 w_2, \ v \equiv w_2 w_1. \tag{2}$$

$u \sim v$ if and only if cyclically reduced words u, v are conjugated (in the ordinary sense) in F_m, and $\sim$ is an equivalence relation on the set of cyclically reduced words. A *cyclic word* is an equivalence class of this relation. That is, a cyclic word is a cyclically reduced word considered up to cyclic shifts. Cyclic words are in one-to-one correspondence with conjugacy classes of F_m. $u \sim v$ implies $|u| = |v|$, therefore the length of a cyclic word is well-defined.

A cyclically reduced word w is *simple* if it can not be represented in the form $w \equiv v^s, \ s > 1$ (thus, simple words are non-empty). Simple (cyclically reduced) words will be also called *periods*[3] and denoted by capital letters P, Q. If $P \sim u$ and P is a period, then u is a period, too. Different cyclic shifts of a period are also different as words. That is, if in (2) u (and, hence, also v) is a period and both w_1, w_2 are non-empty, then $u \neq v$. Cyclic words consisting of periods will be called *cyclic periods* and denoted by the letters $\mathfrak{p}, \mathfrak{q}$. Thus, cyclic periods are periods considered up to cyclic shifts. It is worth noting that if we further identify $\mathfrak{p}$ with $\mathfrak{p}^{-1}$, then these will be in one-to-one correspondence with maximal cyclic subgroups of F_m considered up to conjugacy.

Let P be a period. A reduced word u is *P-periodic* if $u \subseteq P^s$ for some $s > 0$ and $|u| \geq 2|P|$. We denote by $\mathrm{Per}(P)$ the set of all P-periodic words. u is *periodic* if it is P-periodic for some period P and *aperiodic* otherwise.

Clearly, u is P-periodic if and only if it is representable in the form $Q^s Q'$, where $Q \sim P, \ s \geq 2$ and Q' is a proper prefix of Q (and we will see soon that such a representation is unique). In particular, if $P \sim Q$ and u is P-periodic then it is also Q-periodic. Therefore, for every *cyclic* period $\mathfrak{p}$ we have the well-defined notion $\mathrm{Per}(\mathfrak{p})$ of *$\mathfrak{p}$-periodic words*.

In order to go any further, we need the following simple but very fundamental *Overlapping Lemmas* (see e.g. [17, Section 1.2]).

[3]This is a slight deviation from the notation of [17] where periods are not required to be simple.

Lemma 2.1 (First Overlapping Lemma) *Let P, Q be two periods, and u, v, w be reduced words such that*

$$uv = P'P^s, \ vw = Q^tQ', \tag{3}$$

where $s, t \geq 0$, P' is a proper suffix of P and Q' is a proper prefix of Q. Assume further that

$$|v| \geq |P| + |Q|.$$

Then $P \sim Q$. Moreover, the two representations (3) *are* **compatible in phase** *in the following sense: if $v = P''P^{s'}$, where P'' is a (possibly another) suffix of P, $P = P^{(3)}P''$, then $Q = P''P^{(3)}$.*

Applying Lemma 2.1 in the case when the wings u, w are empty, we find that $\mathrm{Per}(\mathfrak{p}) \cap \mathrm{Per}(\mathfrak{q}) = \emptyset$ for any two different cyclic periods $\mathfrak{p}, \mathfrak{q}$.[4] The *left period* of $u \in \mathrm{Per}(\mathfrak{p})$ is defined as that particular $P \in \mathfrak{p}$ for which $u = P^sP'$ ($s \geq 2$), and *right periods* are defined symmetrically. Then the second part of Lemma 2.1 implies that left and right periods of periodic words are *uniquely defined*. Also, if we know left and right periods of $u \in \mathrm{Per}(\mathfrak{p})$, and also know $|u|$ within additive error $C \cdot |\mathfrak{p}|$, then u itself is completely determined up to $(2C+1)$ possibilities (this simple remark will play a crucial role in Section 5).

Finally, let $u = LvR$ be an occurrence of a $\mathfrak{p}$-periodic word v into a (reduced) word u. Let us choose any *maximal* occurrence $u = L'\hat{v}R'$ of a $\mathfrak{p}$-periodic word subsuming this one (that is, L' is a prefix of L, and R' is a suffix of R). Then Lemma 2.1 again implies that this occurrence is uniquely defined, and we call it *the maximal $\mathfrak{p}$-periodic extension of the occurrence $u = LvR$*. Equivalently, the length of the common part of any two different maximal occurrences of $\mathfrak{p}$-periodic words into the same word is less than $2|\mathfrak{p}|$ (otherwise they could have been combined into one larger occurrence of a $\mathfrak{p}$-periodic word by Lemma 2.1).

First Overlapping Lemma basically says that occurrences of sufficiently periodic words can not overlap "accidentally", and this is what one needs for the problems where the periodical structure is given to us a priori (which is the case e.g. for the Burnside problem). On the contrary, the Second Overlapping Lemma tells us how to *extract* such structure from any two occurrences of an *arbitrary* word, provided they are close enough. This lemma

[4]Note that $|v| \geq 2|P|$ and $|v| \geq 2|Q|$ imply $|v| \geq |P| + |Q|$.

lies in the heart of the research on equations in free groups cited in Introduction.

Lemma 2.2 (Second Overlapping Lemma) *Let $u \equiv LvR$, $u \equiv L'vR'$ be two different occurrences of the same word v into u. Assume that*

$$||L'| - |L|| \leq \frac{1}{3}|v|.$$

Then $v \in \mathrm{Per}(\mathfrak{p})$ for some cyclic period $\mathfrak{p}$ and, moreover, these two occurrences of v into u have the same maximal $\mathfrak{p}$-periodic extension.

If G is a group and $A_1, \ldots, A_k \subseteq G$ then

$$A_1 \cdot \ldots \cdot A_k \stackrel{\text{def}}{=} \{b \in G \mid (\exists (a_1, \ldots, a_k) \in A_1 \times \ldots \times A_k)(b = a_1 a_2 \ldots a_k)\}.$$

Throughout the paper we use the asymptotic notation $O, \Omega, \widetilde{O}, \widetilde{\Omega}$ quite customary in Combinatorics and Theoretical Computer Science. Thus[5], $f \leq O(g)$ $[f \geq \Omega(g)]$ means "there exists an absolute constant $C > 0$ $[\epsilon > 0]$ such that $f \leq Cg$ $[f \geq C\epsilon$, respectively] for all possible values of parameters assumed in f, g explicitly or implicitly". Its "soft" version $f \leq \widetilde{O}(g)$ and $f \geq \widetilde{\Omega}(g)$ can be used when all parameters $n_1, \ldots, n_t$ to f, g are integer and given explicitly (or, at least, are clear from the context). $f(n_1, \ldots, n_t) \leq \widetilde{O}(g(n_1, \ldots, n_t))$ $[f(n_1, \ldots, n_t) \geq \widetilde{\Omega}(g(n_1, \ldots, n_t))]$ means that there exist absolute constants $C, k > 0$ $[\epsilon, k > 0]$ such that $\forall n_1, \ldots, n_t (f(n_1, \ldots, n_t) \leq C \cdot \log^k(n_1 + \ldots + n_t) g(n_1, \ldots, n_t))$ $[\forall n_1, \ldots, n_t (f(n_1, \ldots, n_t) \geq \epsilon g(n_1, \ldots, n_t) / \log^k(n_1 + \ldots + n_t))$, respectively].

Thus, in this notation our main result looks as follows.

Theorem 2.3 *Let $A \subseteq F_m$ be a finite subset of the free group F_m with at least two non-commuting elements. Then $|A \cdot A \cdot A| \geq \widetilde{\Omega}(|A|^2)$.*

Remark 1 In one place of our proof (namely, Lemma 3.5) constants assumed in the asymptotic notation do become dependent on the number of generators m. But this dependence can be eliminated by considering any fixed embedding $\phi : F_m \longrightarrow F_2$, and applying Theorem 2.3 to $\phi(A)$ (instead of applying it to the original $A \subseteq F_m$).

[5]Most people would have used here the equality sign, but we find the combination of this notation with $\leq, \geq$ particularly expressive and instructive.

In fact, our main Lemma 3.2 readily implies a more general result. Recall that a group G is *virtually free* [*virtually cyclic*] if it contains a free [cyclic, respectively] subgroup of finite index.

Theorem 2.4 *Let G be any fixed virtually free group and $A \subseteq G$ be its finite subset such that the subgroup generated by A is* **not** *virtually cyclic. Then $|A \cdot A \cdot A| \geq \tilde{\Omega}(|A|^2)$.*

In particular, it is well-known that the modular group $PSL_2(\mathbb{Z}) \approx \mathbb{Z}_2 * \mathbb{Z}_3$ is virtually free (e.g. because its commutant is torsion-free, therefore it is a free subgroup (of index 6) by the Kurosh subgroup theorem [12, Theorem IV.1.10]). The same is true for $SL_2(\mathbb{Z})$ (every *free* subgroup of $PSL_2(\mathbb{Z})$ can be lifted to $SL_2(\mathbb{Z})$), as well as for $GL_2(\mathbb{Z})$. Therefore, Theorem 2.4 improves upon [4, Theorem 5.1] (that, under the same assumptions, stated the bound $|A \cdot A \cdot A| \geq |A|^{1+\delta}$ for $SL_2(\mathbb{Z})$ and for an unspecified constant $\delta > 0$).

3. Reduction: combinatorial part

This and the next two sections are entirely devoted to the proof of Theorems 2.3, 2.4. Our overall strategy is to analyze a potential counterexample by exhibiting in it "sufficiently large" subsets with "sufficiently rich" structure (accordingly, most of the proof is written in the distinct "top-down" style). And, as stated, Theorem 2.3 turns out to be very inconvenient for this purpose. Our first task is to replace it with a stronger (and much clumsier) statement specifically designed with several types of reduction in mind.

Definition 3.1 For a finite subset $A \subseteq F_m$, $\Delta(A)$ is the maximal possible size of the intersection $A \cap C$, where C runs over all cosets of maximal cyclic subgroups[6] in F_m.

Note that Δ is monotone ($\Delta(A) \leq \Delta(B)$ if $A \subseteq B$) and invariant under left and right shifts ($\Delta(A) = \Delta(uA) = \Delta(Au)$).

Lemma 3.2 (Main Lemma) *Let $A, B, C \subseteq F_m$ be finite subsets, and assume that*

$$|A|, |C| \leq O(|B|).$$

[6]Since this class of subgroups is invariant under conjugacy, it does not matter whether we consider left or right cosets in this definition.

Then one of the following two is true.

a) $|A \cdot B \cdot C| \geq \tilde{\Omega}(|A| \cdot |C|)$;

b) $\Delta(B) \geq \Omega(|B|)$.

For the benefit of the reader who may feel uncomfortable with that much of asymptotic notation, we provide a translation of this statement to the ϵ/δ-language (and in all analogous places below the translation is quite similar).

Lemma 3.3 (Main Lemma, ϵ/δ-version) *For every $D > 0$ there exist $\epsilon, K > 0$ such that the following is true. For all finite $A, B, C \subseteq F_m$ with $|A|, |C| \leq D \cdot |B|$, either $|A \cdot B \cdot C| \geq \epsilon \cdot \frac{|A| \cdot |C|}{\log^K(|A|+|B|+|C|)}$ or $\Delta(B) \geq \epsilon \cdot |B|$ holds.*

Proof of Theorems 2.3, 2.4 from Lemma 3.2. Since every virtually cyclic subgroup of a free group is cyclic, Theorem 2.4 implies Theorem 2.3, and we only have to prove the former.

Let G be a virtually free group, and $F \leq G$ be a free subgroup of finite index; w.l.o.g. we can assume that F is normal. Let $A \subseteq G$ be finite; represent it as $A = \dot{\bigcup}_{u \in U} (uA_u)$, where U is an arbitrary set of representatives for cosets of F and $A_u \subseteq F$. Choose that $u \in U$ for which $|A_u|$ is maximal (thus, $|A_u| \geq \Omega(|A|)$), and note that $(uA_u)(uA_u)(uA_u) = u^2(u^{-1}A_u u)A_u(uA_u u^{-1})u$. We apply Lemma 3.2 with $A := u^{-1}A_u u$, $B := A_u$, $C := uA_u u^{-1}$. If the conclusion a) holds, we are done. If $\Delta(A_u) \geq \Omega(|A_u|) \geq \Omega(|A|)$, there exists a maximal cyclic subgroup $C \leq F$ and $v \in F$ such that $|A_u \cap (vC)| \geq \Omega(|A|)$. Denoting $w = uv$, we conclude that $|A \cap (wC)| \geq \Omega(|A|)$. Let $N \leq G$ be the normalizer of C.

If $w \notin N$, we are done: since C and (wCw^{-1}) are *different* maximal cyclic subgroups in F, they have empty intersection. Therefore, all products c_1c_2 $(c_1, c_2 \in (wC))$ are pairwise distinct and $|A \cdot A \cdot A| \geq |A \cdot A| \geq |A \cap (wC)|^2 \geq \Omega(|A|^2)$.

Assume $w \in N$. Since $N \cap F = C$, C has a finite index in N and, therefore, N is virtually cyclic. Since A does not generate a virtually cyclic subgroup, $A \not\subseteq N$; fix arbitrarily $a \in A \setminus N$. And now we are done by the same argument as above, applied to the product $(wC)a(wC)$.■

Remark 2 The statement of Lemma 3.2 allows the following three types of reductions that we are going to use.

- Let $u, v \in F_m$, $A_0 \stackrel{\text{def}}{=} Au^{-1}$, $B_0 \stackrel{\text{def}}{=} uBv$ and $C_0 \stackrel{\text{def}}{=} v^{-1}C$. Then the validity of Lemma 3.2 for the triple (A_0, B_0, C_0) implies its validity for the original (A, B, C).

- The same conclusion holds if $A_0 \subseteq A$, $B_0 \subseteq B, C_0 \subseteq C$ are arbitrary subsets with the only restriction $|A_0| \geq \tilde{\Omega}(|A|)$, $|B_0| \geq \Omega(|B|)$, $|C_0| \geq \tilde{\Omega}(|C|)$.

- Assume that $A = A_1 \dot{\cup} \ldots \dot{\cup} A_{\ell_A}$ and $C = C_1 \dot{\cup} \ldots \dot{\cup} C_{\ell_C}$ are decompositions of A and C into disjoint unions of subsets, and further assume that all $\ell_A \ell_C$ sets $A_i B C_j$ ($i \in [\ell_A]$, $j \in [\ell_C]$) are pairwise disjoint. Then the validity of Lemma 3.2 for all triples (A_i, B, C_j) implies its validity for (A, B, C).

In the reduction of the last type we of course require the *uniform* dependence of assumed constants (that is, ϵ, K on D in the notation of Lemma 3.3).

After this preparatory work, we begin the real proof with getting rid of cancellations.

Lemma 3.4 *For any finite $A \subseteq F_m$, there exists $u \in F_m$ such that for any letter $y \in \{x_1, x_1^{-1}, \ldots, x_m, x_m^{-1}\}$ at least $\frac{1}{4m}|A|$ words in Au^{-1} do* **not** *end with y.*

Proof. Let us call $u \in F_m$ *populated* if it is a suffix of at least $\frac{1}{4|m|}|A|$ words in A. Λ is populated whereas sufficiently long words are not. Choose the longest populated word u; we claim that it has the required property.

Indeed, every one of the words yu ($y \in \{x_1, x_1^{-1}, \ldots, x_m, x_m^{-1}\}$, u does not begin with y^{-1}) is not populated and therefore may appear as a suffix in $\leq \frac{1}{4m}|A|$ words from A. Hence u is a suffix of at most $\frac{1}{2}|A|$ words in A (and on the other hand, it is a suffix of at least $\frac{1}{4m}|A|$ words since u itself is populated). It only remains to note that if u is not a suffix of $a \in A$, then au^{-1} ends with the same letter as u^{-1}, and if it is its suffix, then au^{-1} ends with a different letter, unless it is empty.■

Lemma 3.5 *For any finite $A, B, C \subseteq F_m$ with $|B| \geq 2$ there exist $u, v \in F_m$ and $A_0 \subseteq Au^{-1}$, $B_0 \subseteq uBv$, $C_0 \subseteq v^{-1}C$ such that $|A_0| \geq \Omega(|A|)$, $|B_0| \geq \Omega(|B|)$, $|C_0| \geq \Omega(|C|)$ and all products abc ($a \in A_0$, $b \in B_0$, $c \in C_0$) are reduced.*

Proof. Apply Lemma 3.4 to A, and apply its dual version to C; let u, v be the resulting elements. Removing from uBv the empty word (if it is there), we find a subset $B_0 \subseteq uBv$ with $|B_0| \geq \frac{1}{4m^2}(|B| - 1)$ such that all words in B_0 begin with the same letter y and end with the same letter z. Finally, let $A_0 \subseteq Au^{-1}$ consist of all those words that do not end with y^{-1}, and similarly for $C_0 \subseteq v^{-1}C$. $|A_0| \geq \Omega(|A|)$ and $|C_0| \geq \Omega(|C|)$ hold by Lemma 3.4.■

From this point on, cancellations will never appear again, and the reader may freely assume that we are working in a free *semi-group*. Note that if $abc = a'b'c'$ is a collision in the product $A \cdot B \cdot C$, then a, a' are comparable in the prefix order and c, c' are comparable in the suffix order. This suggests that the most difficult case should be when the elements of A form a *prefix chain* (defined as a set of words mutually comparable in the prefix order), and C forms a *suffix chain*. The following lemma makes this intuition precise.

Definition 3.6 Two prefix [suffix] chains A_1, A_2 are *incomparable* if any two $a_1 \in A_1,\ a_2 \in A_2$ are incomparable in the prefix [suffix, respectively] order.

In particular, incomparable prefix/suffix chains are necessarily disjoint. Also, two prefix chains A_1, A_2 are incomparable if and only if their *minimal* elements are incomparable.

Lemma 3.7 *Every finite set of words A contains a collection $A_1, \ldots, A_\ell \subseteq A$ of mutually incomparable prefix chains such that*

$$|A_1 \cup \ldots \cup A_\ell| = \sum_{i=1}^{\ell} |A_i| \geq \widetilde{\Omega}(|A|), \tag{4}$$

and a similar statement holds for suffix chains.

Proof. Consider the restriction of the prefix order $\leq$ onto A. For $a \in A$, let $h(a)$ be its *height* defined as the maximal possible length of a prefix chain having a as its minimal element (and entirely contained in A). All elements of the same height h are mutually incomparable; let ℓ_h be their number. Then

$$|A| = \sum_{h=1}^{|A|} \ell_h,$$

and also for every h there exist ℓ_h mutually incomparable prefix chains of length h each (for every element a of height h include an arbitrarily chosen prefix chain of height h with the minimal element a).

Thus, if t is the maximal possible value of $|A_1 \cup \ldots \cup A_\ell|$ in (4), then $t \geq h\ell_h$ for each h, which implies

$$|A| \leq t \cdot \sum_{h=1}^{|A|} \frac{1}{h} \leq O(t \log |A|)$$

and, therefore, $t \geq \widetilde{\Omega}(|A|)$.■

Now, by Lemma 3.5 we may assume in Lemma 3.2 that all products abc $(a \in A, b \in B, c \in C)$ are reduced. By Lemma 3.7 we may also assume that A $[C]$ can be decomposed as a union of mutually incomparable prefix [suffix, respectively] chains; say, $A = A_1 \dot\cup \ldots \dot\cup A_{\ell_A}$, $C = C_1 \dot\cup \ldots \dot\cup C_{\ell_C}$. But if $i \neq i' \in [\ell_A]$ then $A_i BC$ and $A_{i'} BC$ are disjoint (since A_i and $A_{i'}$ are incomparable in the prefix order), and similarly for $j \neq j' \in [\ell_C]$. Which means that we can apply the reduction of the third type from Remark 2.

Summarizing what we have achieved so far, *in Lemma 3.2 we can assume w.l.o.g. that all products abc $(a \in A, b \in B, c \in C)$ are reduced, and that, moreover, A is a prefix chain, and C is a suffix chain.*

4. Reduction: finding aperiodicity

At this point we bring into the analysis periodic words, and the rest of the proof is split into two almost independent parts. Namely (thinking in terms of a hypothetical counterexample to Lemma 3.2), we want to show that:

- if $|A \cdot B \cdot C|$ is small, there is enough "periodical structure" in A, B, C;
- if $\Delta(B)$ is small then some large subsets A_0, B_0, C_0 display enough "aperiodicity" in them,

and these two conclusions will contradict each other. Of these two, the first task is much more difficult, interesting and natural to start with. But for technical reasons we have to begin with the second.

Definition 4.1 Let $a, b \in F_m$, and assume that the product ab is reduced. We say that ab is *left regular* if b is periodic, and a ends with P^2, where P is the left period of b (equivalently, $b \in \mathrm{Per}(\mathfrak{p})$ for some cyclic period $\mathfrak{p}$, and its maximal $\mathfrak{p}$-periodic extension in ab has length $\geq |b| + 2|\mathfrak{p}|$). ab is *left singular*

in all other cases. *Right regular* and *right singular* products bc are defined by symmetry.

Definition 4.2 Let P be a period, and $A \subseteq F_m$ be a finite set. We define $\Delta_{\ell,P}$ as the maximal possible size of the intersection $A \cap C$, where C runs over all sets of the form

$$\left\{ LP^t \mid t \geq 0 \right\} \ (L \in F_m,\ LP \text{ reduced}).$$

$\Delta_{r,P}(A)$ is defined by symmetry.

Clearly, $\Delta_{\ell,P}(A), \Delta_{r,P}(A) \leq \Delta(A)$.

Lemma 4.3 *Let $A, B, C \subseteq F_m$ be finite sets, and assume that all products abc $(a \in A,\ b \in B,\ c \in C)$ are reduced. Then either*

$$\Delta(B) \geq \Omega(|B|), \tag{5}$$

or there exist $A_0 \subseteq A$, $B_0 \subseteq B$, $C_0 \subseteq C$ with $|A_0| \geq \Omega(|A|)$, $|B_0| \geq \Omega(|B|)$, $|C_0| \geq \Omega(|C|)$ such that at least one of the following three is true.

- **a)** *At least $\frac{1}{2}|A_0||B_0|$ products ab $(a \in A_0,\ b \in B_0)$ are left singular.*
- **b)** *At least $\frac{1}{2}|B_0||C_0|$ products bc $(b \in B_0,\ c \in C_0)$ are right singular.*
- **c)** *For every period P which is the left period of at least one periodic word in B_0, $\Delta_{\ell,P}(A_0) \leq O(1)$, and the dual conclusion holds for right periods.*

Proof. Either at least half of all words in B are aperiodic, or at least half of them is periodic. In the first case both a), b) hold trivially. Removing from B all aperiodic words in the second case, we may assume w.l.o.g. that all words in B are periodic.

Consider now any individual cyclic period $\mathfrak{p}$ for which $B_\mathfrak{p} \stackrel{\text{def}}{=} B \cap \mathrm{Per}(\mathfrak{p})$ is non-empty. If there exists $P \in \mathfrak{p}$ that appears as either the left period in at least half of all words from $B_\mathfrak{p}$ or the right period in at least half of them, remove from $B_\mathfrak{p}$ all words violating this. Repeating this procedure once more if necessary, we will find $B'_\mathfrak{p} \subseteq B_\mathfrak{p}$ with $|B'_\mathfrak{p}| \geq \Omega(B_\mathfrak{p})$ and such that one of the following is true.

a) Every period $P \in \mathfrak{p}$ appears as the left period in $\leq \frac{1}{2}|B'_{\mathfrak{p}}|$ words from $B'_{\mathfrak{p}}$.

b) Every period $P \in \mathfrak{p}$ appears as the right period in $\leq \frac{1}{2}|B'_{\mathfrak{p}}|$ words from $B'_{\mathfrak{p}}$.

c) All words in $B'_{\mathfrak{p}}$ have the same left and right periods.

Let $B' \stackrel{\text{def}}{=} \bigcup_{\mathfrak{p}} B'_{\mathfrak{p}}$. At the expense of decreasing $|B'|$ by at most a factor of three, we may assume that one and the same of these three alternatives holds for *every* cyclic period $\mathfrak{p}$ for which $B'_{\mathfrak{p}}$ is non-empty.

Alternatives a) and b) (along with Proposition 2.1) immediately apply the corresponding conclusions in the statement of Lemma 4.3 (with $A_0 := A,\ B_0 := B',\ C_0 := C$) since then in (say) case a), for every $a \in A$ and every cyclic period $\mathfrak{p}$ there would be at most $\leq \frac{1}{2}|B'_{\mathfrak{p}}|$ words $b \in B'_{\mathfrak{p}}$ for which ab is left regular. So, we are left with the case when for every $\mathfrak{p}$, all words in $B'_{\mathfrak{p}}$ have the same left and right periods. Note that in this case $B'_{\mathfrak{p}}$ is a subset of the coset $\{P'P^tP'' \mid t \in \mathbb{Z}\}$ of a cyclic subgroup and, therefore,

$$|B'_{\mathfrak{p}}| \leq |\Delta(B)|. \tag{6}$$

Let $\{\mathfrak{p}_1, \ldots, \mathfrak{p}_d\}$ be the enumeration of all cyclic periods $\mathfrak{p}$ for which $B'_{\mathfrak{p}} \neq \emptyset$ in the order of non-decreasing length:

$$|\mathfrak{p}_1| \leq |\mathfrak{p}_2| \leq \ldots \leq |\mathfrak{p}_d|. \tag{7}$$

Choose the minimal ℓ for which $\sum_{i=1}^{\ell} |B'_{\mathfrak{p}_i}| \geq \frac{1}{3}|B'|$ (thus, $\sum_{i=1}^{\ell-1} |B'_{\mathfrak{p}_i}| < \frac{1}{3}|B'|$). If $\sum_{i=1}^{\ell} |B'_{\mathfrak{p}_i}| \geq \frac{2}{3}|B'|$ then $|B'_{\mathfrak{p}_\ell}| \geq \frac{1}{3}|B'|$, and hence (6) implies (5).

Otherwise, $\sum_{i=\ell+1}^{d} |B'_{\mathfrak{p}_i}| \geq \frac{1}{3}|B'|$, and we first try out the set $\bigcup_{i=\ell+1}^{d} B'_{\mathfrak{p}_i}$ as B_0. If at least $\frac{1}{2}|A||B_0|$ products ab ($a \in A,\ b \in B_0$) are left singular, or at least $\frac{1}{2}|B_0||C|$ products bc ($b \in B_0,\ c \in C$) are right singular, we are done.

Otherwise, there exist *fixed* $b_\ell, b_r \in \bigcup_{i=\ell+1}^{d} B'_{\mathfrak{p}_i}$ such that for at least half of all $a \in A$ the product ab_ℓ is left regular, and for at least half of all $c \in C$, $b_r c$ is right regular. We remove from A and C all elements violating these properties, and let A_0, C_0 be the result of this removal. Set also

$$B_0 \stackrel{\text{def}}{=} \bigcup_{i=1}^{\ell} B_{\mathfrak{p}_i}.$$

We finally claim that A_0, B_0, C_0 satisfy the alternative c) in Lemma 4.3, and, by symmetry, it is sufficient to check this only on the left side.

Indeed, all words in A_0 end with Q^2, where Q is the left period of b_ℓ (and hence $Q \in \mathfrak{p}_j$ for some $j \geq \ell + 1$). If a period P appears as the left period of some word in B_0 then $P \in \mathfrak{p}_i$ for some $i \leq \ell$. In particular, $P \not\sim Q$ and, by (7),

$$|P| \leq |Q|. \tag{8}$$

According to Definition 4.2, consider any fixed wing L such that LP is reduced. If $LP^t \in A_0$ then LP^t ends with Q^2. The word Q^2, however, is not $\mathfrak{p}_i$-periodic, therefore, due to (8), it can not be a subword of P^s for any s. Which means that P^t is a suffix of Q^2. Moreover, if $t \geq 2$ then the maximal $\mathfrak{p}_i$-periodic extension of P^t in LP^t is a proper suffix of Q^2 and, therefore, has the same length as its maximal $\mathfrak{p}_i$-periodic extension in Q^2. In particular, *this extension does not depend on* t. This implies that there can be at most one value $t \geq 2$ for which LP^t ends with Q^2. Which shows that $\Delta_{\ell,P}(A_0) \leq 3$ and completes the proof of Lemma 4.3.■

To summarize, so far we have reduced Lemma 3.2 to its partial case described as follows (the alternative b) in the statement of Lemma 3.2 has already been used up in (5), and we do not need to carry it any longer).

Lemma 4.4 *Let $A, B, C \subseteq F_m$ be finite sets such that*

$$|A|, |C| \leq O(|B|).$$

Assume that all products abc ($a \in A,\ b \in B,\ c \in C$) are reduced, that A is a prefix chain, and that C is a suffix chain. Moreover, assume that one of the following three is true.

- **a)** *At least $\frac{1}{2}|A||B|$ products ab ($a \in A,\ b \in B$) are left singular.*
- **b)** *At least $\frac{1}{2}|B||C|$ products bc ($b \in B,\ c \in C$) are right singular.*
- **c)** *For every period P which is the left period of at least one periodic word in B, $\Delta_{\ell,P}(A) \leq O(1)$, and the symmetric conclusion holds for the right periods.*

Then

$$|A \cdot B \cdot C| \geq \tilde{\Omega}(|A| \cdot |C|).$$

5. Finding periodicity with collision numbers

In this section we prove Lemma 4.4, thereby completing the proof of our main result.

Fix $A, B, C \subseteq F_m$ satisfying all the premises of Lemma 4.4. Define $T \subseteq A \times B \times C$ as follows. If one of the alternatives a), b) holds, T consists of those triplets (a, b, c) for which *either* ab is left singular *or* bc is right singular. In the remaining case c), we simply let $T := A \times B \times C$. Note that in any case

$$|T| \geq \Omega(|A| \cdot |B| \cdot |C|). \tag{9}$$

We define the *collision number* $\mathbf{c}_T(A, B, C)$ as

$$\mathbf{c}_T(A, B, C) \stackrel{\text{def}}{=} \left| \left\{ ((a, b, c), (a', b', c')) \in T^2 \mid abc = a'b'c' \right\} \right|.$$

For $u \in F_m$, let

$$n(u) \stackrel{\text{def}}{=} |\{ (a, b, c) \in T \mid abc = u \}|.$$

Then by Cauchy-Schwartz and (9),

$$\left.\begin{aligned} \mathbf{c}_T(A, B, C) &= \sum_{u \in A \cdot B \cdot C} n(u)^2 \geq \frac{1}{|A \cdot B \cdot C|} \left(\sum_{u \in A \cdot B \cdot C} n(u) \right)^2 \\ &= \frac{|T|^2}{|A \cdot B \cdot C|} \geq \Omega\left(\frac{|A|^2 |B|^2 |C|^2}{|A \cdot B \cdot C|} \right). \end{aligned}\right\} \tag{10}$$

Thus, in order to complete our proof, we only have to show that

$$\mathbf{c}_T(A, B, C) \leq \tilde{O}(|A||B|^2|C|). \tag{11}$$

Our next task is to set stage for the Second Overlapping Lemma 2.2, and for this we need one more reduction (this time in terms of collision numbers). But now the reduction is slightly more subtle than those based on Remark 2 seen in previous sections. For this reason we prefer to change the gears, and we first *formulate* the statement we are reducing to.

Lemma 5.1 *Let $A, B, C \subseteq F_m$ be finite sets such that*

$$|A|, |C| \leq O(|B|). \tag{12}$$

Assume that all products abc ($a \in A,\ b \in B,\ c \in C$) are reduced, that A is a prefix chain of even length, and that C is a suffix chain of even length: $A = \{a_1, \ldots, a_{2n_A}\},\ C = \{c_1, \ldots, c_{2n_C}\}$, where $a_1 < a_2 < \ldots < a_{2n_A}$ and $c_1 <^* c_2 <^* \ldots <^* c_{2n_C}$. Let $T \subseteq A \times B \times C$ be such that either (A, B, C) satisfies property c) *in the statement of Lemma 4.4, or for every $(a, b, c) \in T$ either ab is left singular or bc is right singular. Then*

$$\left.\begin{aligned} &|\{((a_i, b, c_j), (a_{i'}, b', c_{j'})) \in T^2 \mid a_i b c_j = a_{i'} b c_{j'}, \\ &\qquad |\{i, i'\} \cap \{1, 2, \ldots, n_A\}| = 1, \\ &\qquad |\{j, j'\}| \cap \{1, 2, \ldots, n_C\}| = 1\} \\ &\quad \le O(|A||B|^2|C|). \end{aligned}\right\} \quad (13)$$

Thus, the only difference in the conclusion from (11) is that we additionally require that the "middle" prefix a_{n_A} of a_{2n_A} separates i from i', and the same holds for j, j'.

Proof of (11) **from Lemma 5.1.** Let (A, B, C) satisfy the assumptions of Lemma 4.4, and let T be defined as in the beginning of this section. Assume for simplicity that $|A|$ and $|C|$ are powers of 2, and represent A and C similarly to the statement of Lemma 5.1: $A = \{a_1, \ldots, a_{n_A}\},\ C = \{c_1, \ldots, c_{n_C}\}$, where $a_1 < a_2 < \ldots < a_{n_A}$ and $c_1 <^* c_2 <^* \ldots <^* c_{n_C}$. For $d \le \log_2 n_A,\ d^* \le \log_2 n_C$ and integers α, γ, let

$$\begin{aligned} A^d_\alpha &\stackrel{\text{def}}{=} \left\{ a_i \in A \,\middle|\, \lfloor i/2^d \rfloor = \alpha \right\}; \\ C^{d^*}_\gamma &\stackrel{\text{def}}{=} \left\{ c_j \in C \,\middle|\, \lfloor j/2^{d^*} \rfloor = \gamma \right\}. \end{aligned}$$

For any fixed values of d, d^*, α, γ we can apply Lemma 5.1 to the triple $(A^d_\alpha, B, C^{d^*}_\gamma)$ letting $T := T \cap (A^d_\alpha \times B \times C^{d^*}_\gamma)$. Summing up the right-hand sides of the resulting estimates (13), we get

$$O\left(\sum_{d,d^*}\sum_{\alpha,\gamma} |A^d_\alpha||B|^2|C^{d^*}_\gamma|\right) = O\left(\sum_{d,d^*} |A||B|^2|C|\right) \le \widetilde{O}(|A||B|^2|C|),$$

as d, d^* take on only logarithmically many values.

On the other hand, the sets in the left-hand sides of (13) give a partition of all those tuples $((a_i, b, c_j), (a_{i'}, b', c_{j'})) \in T^2$ for which $a_i b c_j = a_{i'} b' c_{j'}$ *and*

$i \neq i',\ j \neq j'$. Namely, such a tuple is counted in that $(A^d_\alpha, B, C^{d^*}_\gamma)$ where d is the most significant bit in which binary representations of i and i' differ, d^* is defined in the same way from j, j' and $\alpha = \lfloor i/2^d \rfloor (= \lfloor i'/2^d \rfloor),\ \gamma = \lfloor j/2^{d^*} \rfloor$.

And there are at most $2|A||B|^2|C|$ tuples $((a_i, b, c_j), (a_{i'}, b', c_{j'}))$ with $a_i b c_j = a_{i'} b' c_{j'}$ for which either $i = i'$ or $j = j'$.■

Now we prove Lemma 5.1, and at this point we have to break the symmetry by assuming (w.l.o.g) that

$$|C| \leq |A|. \tag{14}$$

For two words a, a' comparable in the prefix order, we let $\delta(a, a')$ denote their difference (that is, $a = a'\delta(a, a')$ or $a' = a\delta(a, a')$, depending on which of the two is longer). Let $\mathcal{P}$ be the set of all those cyclic periods $\mathfrak{p}$ for which there exists an occurrence

$$a_{2n_A} = L_\mathfrak{p} u_\mathfrak{p} R_\mathfrak{p} \tag{15}$$

of a $\mathfrak{p}$-periodical word $u_\mathfrak{p}$ in a_{2n_A} which is "non-trivially cut" by a_{n_A} in the following sense:

$$a_{n_A} = L_\mathfrak{p} v,\ v \text{ is a prefix of } u_\mathfrak{p} \text{ with } |v| \geq 2|\mathfrak{p}| \text{ and } |\delta(v, u_\mathfrak{p})| \geq 2|\mathfrak{p}|. \tag{16}$$

It follows that for any fixed $\mathfrak{p}$, maximal $\mathfrak{p}$-periodic extensions of all such occurrences coincide, and we choose (15) to be this maximal (and uniquely defined) occurrence.

Next, let $A_\mathfrak{p}$ be the set of all $a_i \in A$ for which we, like in (16), still have $a_i = L_\mathfrak{p} v$, where v is a prefix of $u_\mathfrak{p}$ with $|v| \geq 2|\mathfrak{p}|$ and $|\delta(v, u_\mathfrak{p})| \geq 2|\mathfrak{p}|$, but now *we also additionally require that* $|\delta(a_i, a_{n_A})| \geq 2|\mathfrak{p}|$. This new condition implies in particular that $a_{n_A} \notin A_\mathfrak{p}$. In fact, it implies that for $a_i \in A_\mathfrak{p}$ the word $\delta(a_i, a_{n_A})$ is $\mathfrak{p}$-periodic; therefore, $A_\mathfrak{p} \cap A_\mathfrak{q}$ *for every two different cyclic periods* $\mathfrak{p}, \mathfrak{q}$.

After this set-up, we begin proving the bound (13). First we drop from circulation the condition $|\{j, j'\} \cap \{1, 2, \ldots, n_C\}| = 1$, and simplify the dual one by insisting that $i \leq n_A < i'$. That is, we will prove (13) in the form

$$\left.\begin{aligned} &|\{((a_i, b, c_j), (a_{i'}, b', c_{j'})) \in T^2 \mid a_i b c_j = a_{i'} b' c_{j'}, \\ &\qquad i \leq n_A,\ i' \geq n_A + 1\}| \\ &\quad \leq O(|A||B|^2|C|). \end{aligned}\right\} \tag{17}$$

We do it by case analysis according to the structural properties of a tuple $((a_i, b, c_j), (a_{i'}, b', c_{j'}))$ contributing to the left-hand side. In every of the four cases our strategy will be the same: we will show that four out of six elements of the tuple $((a_i, b, c_j), (a_{i'}, b', c_{j'}))$ already determine it up to $O(1)$ possibilities. But the exact choice of these four entries will depend on the case.

Case 1. There is no cyclic period $\mathfrak{p}$ such that $\{a_i, a_{i'}\} \subseteq A_{\mathfrak{p}}$. Let us call such pairs $(a_i, a_{i'})$ *singular*. First we claim that every fixed $d \in F_m$ can be realized in the form $\delta(a_i, a_{i'})$ for at most 12 singular pairs $(a_i, a_{i'})$.

Indeed, any such realization $a_{i'} = a_i d$ defines the occurrence $a_{2n_A} = a_i d \delta(a_{i'}, a_{2n_A})$ of d into a_{2n_A}, and, moreover, $|a_{n_A}| - |d| \le |a_i| \le |a_{n_A}|$. Suppose for the sake of contradiction that d possesses ≥ 13 realizations. Then, by the pigeon-hole principle, we could choose five of them $d = \delta(a_{i_1}, a_{i'_1}) = \ldots = \delta(a_{i_5}, a_{i'_5})$ $(i_1 \le \ldots \le i_5)$ such that $||a_{i_\alpha}| - |a_{i_\beta}|| \le |d|/3$ for all $\alpha, \beta \in [5]$. Therefore, we could apply Lemma 2.2 and conclude that $d \in \mathrm{Per}(\mathfrak{p})$ for some cyclic period $\mathfrak{p}$ and, moreover, all five selected occurrences of d into a_{2n_A} would be contained in the same maximal occurrence of a $\mathfrak{p}$-periodic word in a_{2n_A}. Further, they would be compatible in phase (in the sense of Lemma 2.1), that is all $||a_{i_\alpha}| - |a_{i_\beta}||$ would be multiples of $|\mathfrak{p}|$. Which would readily imply that this maximal occurrence would necessarily be the occurrence (15), and that $\{a_{i_3}, a_{i'_3}\} \subseteq A_{\mathfrak{p}}$, a contradiction.

Now we only have to observe that $a_i b c_j = a_{i'} b' c_{j'}$ implies $\delta(a_i, a_{i'}) = \delta(bc_j, b'c_{j'})$, that is $b, c_j, b', c_{j'}$ determine $\delta(a_i, a_{i'})$. Therefore, they also determine $a_i, a_{i'}$ up to ≤ 12 possibilities, and hence the contribution of Case 1 to (17) is estimated as $O(|B|^2|C|^2)$ which is $O(|A||B|^2|C|)$ by (14).

Case 2. $\{a_i, a_{i'}\} \subseteq A_{\mathfrak{p}}$ for some cyclic period $\mathfrak{p}$ and $|b| \le 2|\mathfrak{p}|$. In this case we claim that the tuple can be retrieved (again, up to $O(1)$ possibilities) from $a_i, c_j, a_{i'}, b'$. Indeed, since $a_i, a_{i'} \in A_{\mathfrak{p}}$, we have $|\delta(a_i, a_{i'})| = |\delta(a_i, a_{n_A})| + |\delta(a_{i'}, a_{n_A})| \ge 4|\mathfrak{p}|$. This implies that $|a_{i'}| - |a_i b| \ge 2|\mathfrak{p}|$ and hence $\delta(a_i b, a_{i'})$ is $\mathfrak{p}$-periodic. Its left period is completely determined by c_j (as $\delta(a_i b, a_{i'}) \le c_j$), and its right period is determined by $a_{i'}$ (as $\delta(a_i b, a_{i'}) \le^* a_{i'}$). Finally, since $|b| \le 2|\mathfrak{p}|$, we can estimate its length as $|a_{i'}| - |a_i| - 2|\mathfrak{p}| \le |\delta(a_i b, a_{i'})| \le |a_{i'}| - |a_i|$. Thus, given $a_i, c_j, a_{i'}$, there are at most 3 possibilities for $\delta(a_i b, a_{i'})$, and once we know it, we also know b and then $c'_j = \delta(a_i b c_j, a_{i'} b')$.

Thus, Case 2 contributes at most $O(|A|^2|B||C|)$ which is $O(|A||B|^2|C|)$ by (12).

Case 3. $\{a_i, a_{i'}\} \subseteq A_{\mathfrak{p}}$ **for some cyclic period** $\mathfrak{p}$, $|b| \geq 2|\mathfrak{p}|$ **but either** $b \notin \mathbf{Per}(\mathfrak{p})$ **or** $b \in \mathbf{Per}(\mathfrak{p})$ **and the product** bc_j **is right singular.**
This time the tuple is determined by $b, a_{i'}, b', c_{j'}$ (as always, up to $O(1)$ possibilities). Indeed, from these four entries we know $u = a'_i b' c_{j'} \Rightarrow a_i b c_j$, as well as the occurrence

$$u \Rightarrow a_{n_A} \delta(a_{n_A}, a_{i'})(b' c_{j'}) \tag{18}$$

of the $\mathfrak{p}$-periodic word $\delta(a_{n_A}, a_{i'})$ into it. The prefix v of b of length $2|\mathfrak{p}|$ is a prefix of $\delta(a_i, a_{i'})$ and thus $\mathfrak{p}$-periodic; let $b \Rightarrow vw$ and $R \stackrel{\text{def}}{=} wc_j$. Now consider its (yet unknown!) occurrence

$$u \Rightarrow a_i v R \tag{19}$$

into u. These two occurrences of $\mathfrak{p}$-periodic words into u possess a common (also unknown) $\mathfrak{p}$-periodic extension $u \Rightarrow a_i \delta(a_i, a_{i'})(b' c_{j'})$. Therefore, the *maximal* $\mathfrak{p}$-periodic extension $u \Rightarrow \tilde{a}_i \hat{v} R'$ of (19) is the same as the maximal $\mathfrak{p}$-periodic extension of the known occurrence (18), and hence *is also determined by* $(a_{i'}, b', c_{j'})$. Further, if $\hat{v}_1$ is the maximal $\mathfrak{p}$-periodic extension of the prefix v in the word bc_j, then it should have the same right wing R': $bc_j \Rightarrow \hat{v}_1 R'$. And the assumptions of Case 3 imply that $|\hat{v}_1|$ (and hence also $|c_j|$ since b and R' are already known) is determined within accuracy $2|\mathfrak{p}|$ *by the word* b *only*. Namely, it can not exceed by more than $2|\mathfrak{p}|$ the length of the maximal $\mathfrak{p}$-periodic extension of v in b. Therefore, $|a_i|$ and then $\delta(a_i, a_{n_A})$ are also determined within that accuracy. But the left and right periods of the latter words are known (it is a suffix of a_{n_A}, and has v as its prefix), hence this word (and then a_i) is determined up to $O(1)$ possibilities.

Case 4. $\{a_i, a_{i'}\} \subseteq A_{\mathfrak{p}}$ **for some cyclic period** $\mathfrak{p}$, $b \in \mathbf{Per}(\mathfrak{p})$ **and the product** bc_j **is right regular.**
In this final case we also claim that the information can be retrieved from $b, a_{i'}, b', c_{j'}$ (but for entirely different reasons). Namely, recalling the definition (15), the word $\delta(L_{\mathfrak{p}}, a_i)b$ is $\mathfrak{p}$-periodic and $|\delta(L_{\mathfrak{p}}, a_i)| \geq 2|\mathfrak{p}|$. Hence, the product $a_i b$ is left regular. Since $(a_i, b, c_j) \in T$, this implies (recall the statement of Lemma 5.1) that (A, B, C) must necessarily satisfy property c) in the statement of Lemma 4.4. In particular, $\Delta_{\ell,P}(A) \leq O(1)$, where P is the left period of b. Let $L'_{\mathfrak{p}}$ be the prefix of $L_{\mathfrak{p}} u_{\mathfrak{p}}$ in (15) with $|L_{\mathfrak{p}}| \leq |L'_{\mathfrak{p}}| \leq |L_{\mathfrak{p}}| + |\mathfrak{p}|$ and such that the left period of $\delta(L'_{\mathfrak{p}}, L_{\mathfrak{p}} u_{\mathfrak{p}})$ is equal to P. Then a_i must necessarily have the form $L'_{\mathfrak{p}} P^t$ for some integer t. And now the condition $\Delta_{\ell,P}(A) \leq O(1)$ again pinpoints it down to $O(1)$ possibilities.

We have shown that every one of four logically possible cases contributes at most $O(|A||B|^2|C|)$ to the left-hand side of (17). This completes the proof of Lemma 5.1, (11), Lemmas 4.4, 3.2 and Theorems 2.3, 2.4.

6. Statistical version of Plünnecke-Ruzsa inequalities

In this section G will be an abelian group. For its finite subsets $A_1, \dots, A_k$, define the *collision number* $\mathbf{c}(A_1, \dots, A_k)$ as

$$\mathbf{c}(A_1, \dots, A_k) \stackrel{\text{def}}{=} |\{((a_1, \dots, a_k), (a'_1, \dots, a'_k)) \in (A_1 \times \dots \times A_k)^2 | \\ a_1 + \dots + a_k = a'_1 + \dots + a'_k\}|.$$

These qualities were extensively used in additive combinatorics, mostly for the case $k = 2$. In the previous section we saw their application (in the non-abelian case) for $k = 3$. And here we observe how extremally natural and appealing the Plünnecke-Ruzsa theory looks in this setting.

By "the setting" we mean the following. By Cauchy-Schwartz (cf. (10)),

$$\mathbf{c}(A_1, \dots, A_k) \ge \frac{|A_1|^2 \dots |A_k|^2}{|A_1 \cdot \dots \cdot A_k|},$$

so we have the lower bound

$$|A_1 \cdot \dots \cdot A_k| \ge \frac{|A_1|^2 \dots |A_k|^2}{\mathbf{c}(A_1, \dots, A_k)}. \qquad (20)$$

And assuming we are willing to accept the right-hand side as a "good enough" substitute for $|A_1 \cdot \dots \cdot A_k|$, we can infer Plünnecke-Ruzsa inequalities as follows.

Lemma 6.1

$$\mathbf{c}(B_1, \dots, B_k, A, A) \ge \frac{\mathbf{c}(B_1, \dots, B_k, A)^2}{|B_1| \cdot (|B_2| \cdot \dots \cdot |B_k|)^2}.$$

Proof. For $\vec{b} = (b_1, \dots, b_k) \in B_1 \times \dots \times B_k$, let $n(\vec{b})$ be the number of tuples $(\vec{b'}, a, a')$ such that $b_1 + \dots + b_k + a = b'_1 + \dots + b'_k + a'$; thus,

$$\mathbf{c}(B_1, \dots, B_k, A) = \sum_{\vec{b}} n(\vec{b}).$$

On the other hand, for any fixed $\vec{b}$, every couple of tuples $(\vec{b}^{(1)}, a_1, a'_1), (\vec{b}^{(2)}, a_2, a'_2)$ contributing to $n(\vec{b})$ as

$$\begin{aligned} b_1 + \ldots + b_k + a_1 &= b_1^{(1)} + \ldots + b_k^{(1)} + a'_1 \\ b_1 + \ldots + b_k + a_2 &= b_1^{(2)} + \ldots + b_k^{(2)} + a'_2 \end{aligned}$$

also contributes to $\mathbf{c}(B_1, \ldots, B_k, A, A)$ as

$$b_1^{(1)} + \ldots + b_k^{(1)} + a'_1 + a_2 = b_1^{(2)} + \ldots + b_k^{(2)} + a'_2 + a_1.$$

And every such contribution is counted at most $|B_2| \cdot \ldots \cdot |B_k|$ times (as this is an upper bound on the number of tuples $\vec{b}$ for which $b_1 + \ldots + b_k$ takes on the prescribed value $b_1^{(1)} + \ldots + b_k^{(1)} + a'_1 - a_1$). Which implies

$$\mathbf{c}(B_1, \ldots, B_k, A, A) \geq \frac{1}{|B_2| \cdot \ldots \cdot |B_k|} \cdot \sum_{\vec{b}} n(\vec{b})^2,$$

and make our lemma the result of yet another application of Cauchy-Schwartz.■

Lemma 6.2

$$\mathbf{c}(A_1, \ldots, A_k) \geq \frac{\mathbf{c}(B, A_1, \ldots, A_k)}{|B|^2}.$$

Proof. Applying the union bound to all possible choices of b, b',

$$\mathbf{c}(B, A_1, \ldots, A_k) \leq |B|^2 \cdot \max_{d \in G} \mathbf{c}_d(A_1, \ldots, A_k),$$

where $\mathbf{c}_d(A_1, \ldots, A_k)$ is the "shifted" version of $\mathbf{c}(A_1, \ldots, A_k)$:

$$\begin{aligned} \mathbf{c}_d(A_1, \ldots, A_k) \stackrel{\text{def}}{=} |\{((a_1, \ldots, a_k), (a'_1, \ldots, a'_k)) \in (A_1 \times \cdots \times A_k)^2 | \\ a_1 + \ldots + a_k + d = a'_1 + \ldots + a'_k\}|. \end{aligned}$$

But

$$\mathbf{c}(A_1, \ldots, A_k) \geq \mathbf{c}_d(A_1, \ldots, A_k) \tag{21}$$

is easy (and well-known). Namely, if $n(e)$ is the number of representations of $e \in G$ in the form $a_1 + \ldots + a_k$, then

$$\begin{aligned} \mathbf{c}(A_1, \ldots, A_k) &= \sum_e n(e)^2 \\ \mathbf{c}_d(A_1, \ldots, A_k) &= \sum_e n(e) n(e+d), \end{aligned}$$

and since the vectors $(n(e)\,|\,e \in G)\,, (n(e+d)\,|\,e \in G)$ have the same ℓ_2 norm, (21) follows by Cauchy-Schwartz.■

Theorem 6.3

$$\mathbf{c}(\underbrace{\pm A, \pm A, \ldots, \pm A}_{k \text{ times}}) \geq \frac{\mathbf{c}(B,A)^{2^{k-1}}}{|B|^{(2^{k-1}+1)}|A|^{2^k-2k}}.$$

Proof. $\mathbf{c}(A_1, \ldots, A_k)$ is clearly invariant under negating components, so we may assume that all signs are actually plus signs. Applying Lemma 6.1 to $B_1 := B,\ B_2 := \ldots := B_k := A$, we find

$$\mathbf{c}(B, \underbrace{A, \ldots, A}_{k \text{ times}}) \geq \frac{1}{|B| \cdot |A|^{2(k-2)}} \cdot \mathbf{c}(B, \underbrace{A, \ldots, A}_{k-1}) \ (k \geq 2).$$

By induction on k,

$$\mathbf{c}(B, \underbrace{A, \ldots, A}_{k \text{ times}}) \geq \frac{\mathbf{c}(B,A)^{2^{k-1}}}{|B|^{(2^{k-1}-1)}|A|^{2^k-2k}}.$$

Applying Lemma 6.2 finishes the proof.■

In order to interpret this result, recall that the standard doubling constant $K_{A,B}$ given by

$$|A \cdot B| = K_{A,B}|B|$$

in our framework corresponds, via (20), to

$$\mathbf{c}(A,B) = \epsilon_{A,B}|A|^2|B| \ (\epsilon_{A,B} = K_{A,B}^{-1}).$$

In this notation Theorem 6.3 can be re-written as

$$\mathbf{c}(\underbrace{\pm A, \ldots, \pm A}_{k \text{ times}}) \geq \epsilon_{A,B}^{2^{k-1}} \cdot \frac{|A|^{2k}}{|B|},$$

which (again, via (20)) corresponds exactly to the "classical" conclusion $|\pm A \pm A \pm \ldots \pm A| \leq K_{A,B}^{O(1)}|B|$.

The material in this section can be readily generalized to convolutions of discrete probability measures (replacing uniform distributions on $A_1, \ldots, A_k$).

Namely, the *collision probability* $\mathbf{cp}(\mu)$ of a discrete probability measure μ is defined as

$$\mathbf{cp}(\mu) \stackrel{\text{def}}{=} \mathbf{P}[\boldsymbol{a} = \boldsymbol{a'}],$$

where $\boldsymbol{a}, \boldsymbol{a'}$ are two random variables picked independently at random according to μ. We also let

$$\ell_\infty(\mu) \stackrel{\text{def}}{=} \max_{a \in \mathrm{Sup}(\mu)} \mu(\{a\})$$

(thus, the *min-entropy* $H^\infty(\mu)$ is equal to $-\log_2 \ell_\infty(\mu)$). If A is the support of μ then clearly

$$\ell_\infty(\mu) \geq \mathbf{cp}(\mu) \geq \frac{1}{|A|}.$$

For probability measures $\mu_1, \ldots, \mu_k$ on an abelian group G, we denote by $\mu_1 + \ldots + \mu_k$ their *convolution*, that is the measure corresponding to the random variable $\boldsymbol{a_1} + \ldots + \boldsymbol{a_k}$, where $\boldsymbol{a_1}, \ldots, \boldsymbol{a_k}$ are picked uniformly at random according to the measures $\mu_1, \ldots, \mu_k$.

And in this notation the proof of Theorem 6.3 can be easily generalized to give the inequality

$$\frac{1}{\ell_\infty(\eta)} \cdot \mathbf{cp}(\underbrace{\pm\mu \pm \mu \pm \ldots \pm \mu}_{k \text{ times}}) \geq \left(\frac{1}{\ell_\infty(\eta)} \cdot \mathbf{cp}(\mu + \eta) \right)^{2^{k-1}}$$

for any two discrete probability measures μ, η on G.

A further generalization is apparently possible in the continuous setting of Tao [15]. It is not clear, however, whether any interesting analogue of this exists in the non-abelian case.

Acknowledgment

I am greatly indebted to Jean Bourgain for posing this problem, and for his constant encouragement afterward. I am also grateful to Avi Wigderson for very useful conversations and for finding a gap in the first version of the argument. My thanks are due to Akshay Venkatesh for pointing out the references [6, 8], and to Lisa Carbone for her question that inspired the extension of the main result to virtually free groups.

References

[1] B. Barak, R. Impagliazzo, and A. Wigderson. Extracting randomness using few independent sources. *SIAM Journal on Computing*, 36(4):1095–1118, 2006.

[2] B. Barak, A. Rao, R. Shaltiel, and A. Wigderson. 2-source dispersers for sub-polynomial entropy and Ramsey graphs beating the Frankl-Wilson construction. In *Proceedings of the* 38*th ACM STOC*, pages 671–680, 2006.

[3] J. Bourgain, N. Katz, and T. Tao. A sum-product estimate in finite fields, and applications. *Geometric and Functional Analysis*, 14(1):27–57, 2004.

[4] M.-C. Chang. Product theorems in SL_2 and SL_3. To appear in *Journal of the Institute of Mathematics of Jussieu*, 2006.

[5] T. Gowers. A new proof of Szemeredi's theorem for arithmetic progressions of length four. *Geometric and Functional Analysis*, 8:529–551, 1998.

[6] U. Haagerup. An example of nonnuclear C^*-algebra which has the metric approximation property. *Inv. Math.*, 50:279–293, 1979.

[7] H. A. Helfgott. Growth and generation in $SL_2(\mathbb{Z}/p\mathbb{Z})$. Technical Report math/0509024, arXiv e-print, 2005.

[8] P. Jollisaint. Rapidly decreasing functions in reduced C^*-algebras of groups. *Trans. Amer. Math. Soc.*, 317:167–196, 1990.

[9] O. Kharlampovich and A. Myasnikov. Implicit function theorem over free groups. *Journal of Algebra*, 290:1–203, 2005.

[10] O. Kharlampovich and A. Myasnikov. Elementary theory of free non-abelian group. *Journal of Algebra*, 302:451–552, 2006.

[11] R. C. Lyndon. Equations in free groups. *Trans. Amer. Math. Soc.*, 96:445–457, 1960.

[12] R. C. Lyndon and P. E. Shupp. *Combinatorial Group Theory*. Springer-Verlag, New York/Berlin, 1977.

[13] M. B. Nathanson. *Additive number theory: inverse problems and the geometry of sumsets.* Graduate texts in mathematics 165. Springer, 1996.

[14] Z. Sela. Diophantine geometry over groups VI: The elementary theory of a free group. *Geometric and Functional Analysis*, 16:707–730, 2006.

[15] T. Tao. Product set estimates in noncommutative groups. Technical Report math/0601431, arXiv e-print, 2006.

[16] T. Tao and V. Vu. *Additive combinatorics.* Cambridge University Press, 2006.

[17] С. И. Адян. *Проблема Бернсайда и тождества в группах.* Наука, Москва, 1975. Engl. transl.: S. I. Adian, The Burnside Problem and Identities in Groups, Springer-Verlag, 1979.

[18] В. К. Булитко. Об уравнениях и неравенствах в свободной группе и свободной полугруппе. *Учёные записки математической кафедры ТГПИ*, 2:242–253, 1970. V. K. Bulitko, On equations and inequalities in a free group and in a free semigroup, *Proceedings of the math. dep. Tula State Institute for Elementary School*, 2(1970), 242-253.

[19] Г. С. Маканин. Уравнения в свободной группе. *Известия АН СССР, сер. матем.*, 46(6):1199–1273, 1982. G. S. Makanin, Equations in a free group, *Math. USSR Izvestiya*, 21(1983), 483-546.

[20] А. А. Разборов. О системах уравнений в свободной группе. *Известия АН СССР, сер. матем.*, 48(4):779–832, 1984. A. A. Razborov, On systems of equations in a free group, *Math. USSR Izvestiya*, 25(1):115-162, 1985.

[21] Г. А. Фрейман. *Начала структурной теории сложения множеств.* КГПИ, 1966. Engl. transl.: G. A. Freiman, Foundations of a structural theory of set addition, American Math. Soc., 1973.