ON GRAHAM'S REARRANGEMENT CONJECTURE OVER $\mathbb{F}_2^n$

BENJAMIN BEDERT[1], MATIJA BUCIĆ[2], NOAH KRAVITZ[3], RICHARD MONTGOMERY[4], AND ALP MÜYESSER[5]

Abstract. A sequence $s_1, s_2, \dots, s_k$ of elements of a group G is called a valid ordering if the partial products $s_1, s_1s_2, \dots, s_1 \cdots s_k$ are all distinct. A long-standing problem in combinatorial group theory asks whether, for a given group G, every subset $S \subseteq G \setminus \{\mathrm{id}\}$ admits a valid ordering; the instance of the additive group $\mathbb{F}_p$ is the content of a well-known 1971 conjecture of Graham. Most partial progress to date has concerned the edge cases where either S or $G \setminus S$ is quite small. Our main result is an essentially complete resolution of the problem for $G = \mathbb{F}_2^n$: we show that there is an absolute constant $C > 0$ such that every subset $S \subseteq \mathbb{F}_2^n \setminus \{0\}$ of size at least C admits a valid ordering. Our proof combines techniques from additive and probabilistic combinatorics, including the Freiman–Ruzsa theorem and the absorption method.

Along the way, we also solve the general problem for moderately large subsets: there is a constant $c > 0$ such that for every group G (not necessarily abelian), every subset $S \subseteq G \setminus \{\mathrm{id}\}$ of size at least $|G|^{1-c}$ admits a valid ordering. Previous work in this direction concerned only sets of size at least $(1 - o(1))|G|$. A main ingredient in our proof is a structural result, similar in spirit to the Arithmetic Regularity Lemma, showing that every Cayley graph can be efficiently decomposed into mildly quasirandom components.

1. Introduction

1.1. **The main problem.** A sequence $g_1, g_2, \dots, g_n$ of elements of a (multiplicative) group G is a *valid ordering* if the partial products

$$g_1, \quad g_1g_2, \quad g_1g_2g_3, \quad \dots, \quad g_1 \cdots g_n$$

are all distinct. Which subsets of groups admit valid orderings? Variants of this natural problem have been studied in many different cases over the years.

The first question in this direction appeared in 1961, when Gordon [17], motivated by constructions of complete Latin squares, asked for which finite groups the entire group has a valid ordering. Gordon gave a complete characterization in the abelian case: A finite (additive), nontrivial abelian group G admits a valid ordering if and only if $\sum_{g \in G} g \neq 0$, this being the obvious necessary condition for the existence of such an ordering. In 1974, Ringel [40] posed the closely related problem of characterising the groups G whose elements can be ordered as $g_1, \dots, g_n$ in such a way that $g_1 = g_1g_2 \cdots g_n = \mathrm{id}$ but otherwise all partial products are distinct. The motivation for this question came from Ringel's solution [41] of the Heawood map colouring conjecture.

The nonabelian case of Gordon's problem is more subtle, since there are some small nonabelian groups (such as S_3) that for no apparent reason fail to have valid orderings. In 1981, Keedwell [29] posed the bold conjecture that every sufficiently large nonabelian group has a valid ordering. Müyesser and Pokrovskiy [36] recently proved Keedwell's conjecture as a consequence of their more general probabilistic analogue of the Hall–Paige Conjecture [11, 25] concerning the existence of transversals in multiplication tables. This work also shows that large groups have an ordering, in the sense that Ringel asked for, if and only if the product of all group elements (in any order) is an element of the commutator subgroup[1].

In this paper we will be concerned not only with the case when an entire group G admits a valid ordering but with the more general question of when an *arbitrary* subset S of a given group G admits a valid ordering. Notice that when S contains the identity element, every possible valid ordering of S must start with the identity, since otherwise two consecutive partial products would be equal. Thus, if G is abelian and $\sum_{g \in S} g = 0$, then there

(1) University of Oxford, UK. Email: benjamin.bedert@maths.ox.ac.uk.

(2) Department of Mathematics, Princeton University, Princeton, USA. Research supported in part by NSF Award DMS-2349013. Email: mb5225@princeton.edu.

(3) Department of Mathematics, Princeton University, Princeton, USA. Email: nkravitz@princeton.edu.

(4) University of Warwick, UK. Email: richard.montgomery@warwick.ac.uk. Supported by the European Research Council (ERC) under the European Union Horizon 2020 research and innovation programme (grant agreement No. 947978).

(5) University of Oxford, UK. Email: alp.muyesser@new.ox.ac.uk.

[1]This condition is equivalent to existence of an ordering $g_1, \dots, g_n$ such that $g_1g_2 \cdots g_n = \mathrm{id}$.

cannot be a valid ordering of S. In order to avoid this obstruction, we restrict our attention to subsets S not containing the identity, and the following is our central question.

Question 1.1. *For which groups G does every subset $S \subseteq G \setminus \{\mathrm{id}\}$ admit a valid ordering?*

It seems feasible that the answer to this question is affirmative for every finite group G. At a first glance, finding valid orderings for smaller subsets S might seem like an easier task, since there is more space to place the partial products without creating collisions. However, the potential obstructions for small S are at least as rich as for Gordon's setting $S = G \setminus \{\mathrm{id}\}$, since a small set S may itself be a subgroup of G, or could be a complicated conglomeration of approximate subgroups and random-like sets. In the graph-theoretic formulation of these problems, which we will describe below, Gordon's setting corresponds to the complete graph case (in particular, a directed variant of a well-known conjecture of Andersen [2]), whereas Question 1.1 corresponds to a sparse analogue. Such sparse analogues in extremal graph theory tend to be harder and less well understood than their dense counterparts.

The simplest instance of Question 1.1 is when $G = \mathbb{F}_p$, for a prime p. This problem was first posed by Graham [19] in 1971 and later reiterated in an open problems book of Erdős and Graham [12].

Conjecture 1.2 (Graham)**.** *Let p be prime. Then every subset of $\mathbb{F}_p \setminus \{0\}$ admits a valid ordering.*

Most previous work towards Conjecture 1.2 has concerned the edge cases where either S or $\mathbb{F}_p \setminus S$ is very large. The best result for small sets S is due to Bedert and Kravitz [4], who showed that every set $S \subseteq \mathbb{F}_p \setminus \{0\}$ of size at most $e^{\log^{1/4} p}$ has a valid ordering. For very large sets S, the aforementioned result of Müyesser and Pokrovskiy [36] establishes Conjecture 1.2 for all sets $S \subseteq \mathbb{F}_p \setminus \{0\}$ of size at least $(1-o(1))p$ (and indeed proves an analogous result for all finite groups; see Theorem 7.1). The intermediate regime remains open.

Various groups of authors (see, e.g., [1, 10, 26]) have considered instances of Question 1.1 other than $G = \mathbb{F}_p$. In particular, Alspach [8] conjectured an affirmative answer to Question 1.1 for all finite abelian groups G, and Alspach and Liversidge [1] confirmed this for subsets of size up to 11. For extensions of this problem to a nonabelian setting, see [9, 37] and the dynamic survey of Ollis [38].

In a different direction, Bucić, Frederickson, Müyesser, Pokrovskiy, and Yepremyan [5] have recently provided an affirmative answer to an "approximate" relaxation of Question 1.1. They showed that every finite subset S of any group G has an ordering in which all but $o(|S|)$ partial products are distinct.

1.2. **Main results.** Despite the partial progress discussed above, there is no infinite class of groups G for which we have a complete understanding of Question 1.1. Our main result remedies this situation for the family of groups $\mathbb{F}_2^n$.

Theorem 1.3. *There is an absolute constant C such that for all $n \in \mathbb{N}$, every set $S \subseteq \mathbb{F}_2^n \setminus \{0\}$ of size at least C has a valid ordering.*

We remark that our methods allow us to obtain the same result for the class of finite abelian groups of exponent at most K for any K. For example, with the same method, for p fixed and $n \to \infty$, we can conclude that any subset of $\mathbb{F}_p^n \setminus \{0\}$ of size at least $C = C(p)$ has a valid ordering. For clarity of exposition, we describe only the 2-torsion case in this paper.

One can view Theorem 1.3 as resolving the "finite-field model" version of Conjecture 1.2. The study of additive combinatorial problems over finite-field models is a well-established topic in its own right; see the decennial surveys by Green [20], Wolf [45], and Peluse [39]. One of the key structural advantages of high-dimensional vector-spaces over finite fields is their rich subgroup structure. Perhaps more unexpectedly, another key advantage —crucial for our purposes— is that any moderately dense $S \subset \mathbb{F}_2^n$ contains an abundance of small subsets whose elements sum to 0. This is surprising given that 0-sum subsets are precisely what we need to avoid in valid orderings. We refer the interested reader to Section 2 for a high-level overview of our proof strategy.

Although $\mathbb{F}_2^n$ has its advantages, the simplest setting for Question 1.1 turns out to be $\mathbb{Z}$, where a simple inductive argument produces a valid ordering of any finite subset of $\mathbb{Z} \setminus \{0\}$ (see [30]). This fact plays a key role in the work of Bedert and Kravitz [4], who resolve Conjecture 1.2 for subsets S of quasipolynomial size by leveraging the fact that $\mathbb{F}_p$ looks locally like $\mathbb{Z}$. Unfortunately, there is no such "lifting" trick in the finite-field model.

Our proof of Theorem 1.3 treats the "sparse S" and "dense S" regimes separately. Our argument for the sparse case makes use of the specific structure of $\mathbb{F}_2^n$, but our argument for the dense case applies to general

(even nonabelian) groups. In particular, we are able to provide an affirmative answer to Question 1.1 if one restricts attention to subsets S of size at least $|G|^{1-c}$; this significantly improves on the result of Müyesser and Pokrovskiy [36], which treats only subsets S of size $(1-o(1))|G|$.

Theorem 1.4. *There is an absolute constant $c > 0$ such that for any finite (possibly nonabelian) group G, every subset $S \subseteq G \setminus \{\mathrm{id}\}$ of size at least $|G|^{1-c}$ admits a valid ordering.*

1.3. **Connections to designs.** Let us say a few words about the relation between Question 1.1 and the theory of combinatorial designs. Gordon was initially interested in groups with valid orderings because their multiplication tables can be used to construct complete Latin squares. A *Latin square*, also called a *quasigroup*, is a group without the axiom of associativity; equivalently, a Latin square is an n by n grid filled with the symbols $\{1, 2, \dots, n\}$ in such a way that each symbol appears exactly once in each row and in each column. A Latin square is called *complete* if for each pair of distinct symbols (i, j), the symbol j appears immediately after the symbol i in exactly one row and in exactly one column. The additional degree of symmetry in complete Latin squares gives them practical uses in the design of experiments (see, e.g., [3]), and they have applications to the study of graph decompositions (see [38]). We point an interested reader to a wonderful book [28] on the topic with a plethora of further connections and applications.

1.4. **A weak nonabelian arithmetic regularity lemma.** The proofs of Theorems 1.3 and 1.4 use a combination of the absorption method and various tools from additive combinatorics. We will give a more detailed overview in the following section, but for now we will highlight one key intermediate result which may be of independent interest. Recall that for a subset X of a group G, the *right Cayley graph* of G with respect to X, denoted $\mathrm{Cay}_G(X)$, is the directed graph with vertex set G where there is a directed edge from g to gx for each $g \in G$ and $x \in X$. The *adjacency matrix* of a directed graph $\Gamma = (V, E)$ is the $|V| \times |V|$ matrix M_Γ with rows and columns indexed by V, where the (u, v)-entry equals 1 if (u, v) is a directed edge and equals 0 otherwise. Note that M_Γ is not necessarily symmetric, so it may have complex eigenvalues. When every vertex of Γ has out-degree d, the adjacency matrix M_Γ always has d as a *trivial eigenvalue* (and in fact d is the largest eigenvalue in absolute value).

Theorem 1.5. *Let $\sigma \in (0, 1]$ and $\varepsilon \in (0, 1/2)$. Let G be a finite (not necessarily abelian) group, and let $S \subseteq G$ be a subset with density $\sigma = |S|/|G|$. Then there is a subgroup H of G such that:*

(1) $|S \cap H| \geq (1-\varepsilon)|S|$;
(2) all non-trivial eigenvalues of the adjacency matrix of $\mathrm{Cay}_H(S \cap H)$ have real part at most $(1-\eta)|S \cap H|$, where $\eta := \varepsilon\sigma^2/1000$.

Condition (2) asserts that $\mathrm{Cay}_H(S \cap H)$ has a positive spectral gap, which turns out to be a natural mild expansion condition for our purposes. In particular, this spectral condition allows us to lower bound the number of edges across any cut of $\mathrm{Cay}_H(S \cap H)$. We say that an *η-sparse cut* in a finite directed graph Γ is a partition $X_1 \sqcup X_2$ of the vertex set of Γ such that there are fewer than $\eta|X_1| \cdot |X_2|$ (directed) edges from X_1 to X_2. We will see below (Lemma 4.5) that (2) implies the purely combinatorial condition that $\mathrm{Cay}_H(S \cap H)$ has no $\eta\sigma$-sparse cut.

In a sense, Theorem 1.5 is analogous to the more familiar Arithmetic Regularity Lemma (ARL) of Green [21] (see also [23]). Roughly speaking, the ARL offers a more refined decomposition where (2) is strengthened by replacing $(1-\eta)|S|$ with $\eta|S|$. This stronger condition allows one to count occurrences of additive patterns such as 3-term arithmetic progressions. Theorem 1.5 is unfortunately unable to count such delicate "local" substructures, but in the context of Question 1.1 the mild quasirandomness condition (2) already provides sufficiently strong information, and we shall see that it has several further redeeming qualities.

One advantage of our weak ARL is that it can handle nonabelian groups. Although there has been some prior interest in nonabelian analogues of the ARL (e.g., model-theoretic approaches [6] can be used to give structure theorems for sets with bounded VC-dimension), our weak ARL is the first such result that applies to arbitrary subsets S. We further note that the decomposition of $\mathrm{Cay}_H(S \cap H)$ provided by Theorem 1.5 is particularly simple, in that it allows us to partition the vertex set G into the cosets of a subgroup H so that each of the induced graphs $\mathrm{Cay}_{xH}(S \cap H)$ is isomorphic to the mildly quasirandom Cayley graph $\mathrm{Cay}_H(S \cap H)$. The fact that these structured components are cosets makes the application in the context of Question 1.1 very convenient. It is known on the other hand that if one wants to find such a decomposition where each component

is "strongly quasirandom" as in Green's ARL, then already in the abelian setting one has to work with more complicated components than subgroups, such as Bohr sets.

Another advantage of our Theorem 1.5 lies in the quantitative aspect. The polynomial dependence of η on σ is ultimately the source of the polynomial improvement in Theorem 1.4. By contrast, it is well-established that tower-type dependences are essential to the usual versions of regularity lemmas [21, 27]. Even for the weak graph regularity lemma of Frieze and Kannan [16], exponential dependencies are required [7]. Therefore, usual versions of regularity lemmas give useful information only about dense subsets, even in the simplest case of Cayley graphs over $\mathbb{F}_2^n$. In contrast, our Theorem 1.5 gives information about polynomially sparse sets S. Note also that the index of H in G is $O(1/\sigma)$ by condition (1), so our decomposition of G into H-cosets (with each $\mathrm{Cay}_{xH}(S \cap H)$ mildly quasirandom) uses only $O(1/\sigma)$ pieces.

We also mention that Theorem 1.5 is closely related to a purely graph-theoretical result of Kühn, Lo, Osthus, and Staden [32] (see also [24, 33]) that provides a similar structural decomposition for *dense* d-regular graphs. More precisely, these authors show that any regular graph of density σ can be decomposed into clusters in such a way that there are very few edges between different clusters, and there are no $f(\sigma)$-sparse cuts within any single cluster; we refer the reader to [24] for further details. Our Theorem 1.5 is a more specialised result because it pertains only to Cayley graphs, but it has the dual advantages of giving group-theoretic information about the clusters, and of enjoying polynomial bounds (as contrasted with the exponential bounds in [32]).

We anticipate that Theorem 1.5 will find further applications in the study of Cayley graphs. For example, in upcoming work, Bedert, Draganić, Müyesser, and Pavez-Signé apply Theorem 1.5 to the well-known conjecture of Lovász asserting that every (connected) Cayley graph is Hamiltonian.

1.5. **Organization of the paper.** In Section 2 we give a high-level overview of our main ideas. The results in this section are only for expository purposes and are not used in the remainder of the paper. Section 3 contains notation and other preliminaries. We then turn to our weak nonabelian regularity lemma in Section 4, which is split into one subsection for the special case of $\mathbb{F}_2^n$ and one subsection for the case of general finite groups. In Section 5 we prove a very flexible asymptotic result for the dense setting under the assumption of a certain expansion condition (as guaranteed by the natural output of Section 4). In Section 6 we prove our absorption lemmas. This section is divided into Section 6.1, where we show how to build our absorbing structure, and Section 6.2, where we show how this structure lets us absorb a small set of leftover colours. In Section 7 we establish our main result over $\mathbb{F}_2^n$ (Theorem 1.3) in the dense case. In Section 8 we complete the proof of Theorem 1.3 by analysing the sparse case. This section is split into Section 8.1, where we deal with the "structured" case, and Section 8.2, where we deal with the "random-like" case. In Section 9 we prove Theorem 1.4, which provides an affirmative answer to Question 1.1 for polynomial-density subsets of general groups. Finally, we make some concluding remarks in Section 10.

We remark that the arguments about $\mathbb{F}_2^n$ in Sections 4.1 and 7 are not strictly speaking necessary since they are subsumed by the more general results in Sections 4.2 and 9. We include the analysis of these special cases separately because several of the arguments simplify, leading to a more direct and streamlined proof of Theorem 1.3. This case also provides an opportunity to build intuition for the more technical general results that follow.

Acknowledgements. We are grateful to Mira Tartarotti and Julia Wolf for remarks concerning arithmetic regularity lemmas, and we thank Zach Hunter for helpful comments on a draft of this paper. The first author gratefully acknowledges financial support from the EPSRC. The second and fifth authors were supported by the National Science Foundation under Grant No. DMS-1928930 during their Spring 2025 residence at the Simons Laufer Mathematical Sciences Institute in Berkeley, California. The third author was supported in part by the NSF Graduate Research Fellowship Program under grant DGE–203965.

2. Overview

Our arguments combine several ideas from different parts of combinatorics, including *inverse problems* and *Fourier analysis* from additive combinatorics, *absorption* from probabilistic combinatorics, and *robust expansion* from extremal combinatorics. In the interest of making our proofs accessible to a wide audience, we will first give a high-level overview of the main ideas in a simplified context. This purely expository section is not logically necessary for the rest of the paper.

It is useful to recast the main problem in the language of finding rainbow paths in Cayley graphs. In general, a *rainbow* subgraph of an edge-coloured graph is a subgraph all of whose edges have different colours; see [5, 34, 43] for more context on the rich study of rainbow subgraphs from a graph-theoretic perspective. We can view the Cayley graph $\mathrm{Cay}_G(S)$ (recall the definition from above) as an edge-coloured digraph with colour set S, where the directed edge from g to gx has the colour x for each $g \in G$ and $x \in S$.

Observation 2.1. *Let S be a finite subset of a group G. Then, S has a valid ordering if and only if $\mathrm{Cay}_G(S)$ has a directed rainbow path with $|S|-1$ edges.*

Proof. If $s_1, \dots, s_{|S|}$ is a valid ordering of S, then

$$s_1 \to s_1 s_2 \to \cdots \to s_1 s_2 \cdots s_{|S|}$$

is a directed rainbow path in $\mathrm{Cay}_G(S)$ with $|S|-1$ edges. Conversely, any directed rainbow path in $\mathrm{Cay}_G(S)$ with $|S|-1$ edges is of the form

$$g s_{\sigma(1)} \to g s_{\sigma(1)} s_{\sigma(2)} \to \cdots \to g s_{\sigma(1)} s_{\sigma(2)} \cdots s_{\sigma(|S|)}$$

for some permutation σ of $[|S|]$ and $g \in G$, and then $s_{\sigma(1)}, s_{\sigma(2)}, \dots, s_{\sigma(|S|)}$ is a valid ordering of S. □

Therefore, our goal is to find a rainbow path of length $|S|-1$ in $\mathrm{Cay}_G(S)$. We use a "99% → 100% framework", more commonly known in the world of probabilistic combinatorics as the "absorption method" since its codification by Rödl, Ruciński, and Szemerédi [42] in 2008 (though its origins can be traced back farther to [13]). The rough idea is that we first find a rainbow path of length $0.99|S|$ and then upgrade this partial rainbow path to a rainbow path of length $|S|-1$.[2] We carry out this upgrade using a certain "absorbing structure" that we set aside before finding the 99% rainbow path. We treat these two steps in the following two subsections.

2.1. **99%-results.** In this subsection we will describe how to find a rainbow path of length $0.99|S|$ in $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$. Such an approximate result was already established recently in [5, Theorem 1.5], but this result is not robust enough for our framework to be able to convert it into a 100% result. The approach we use in the present paper for the 99% part is significantly different and in particular more robust in several ways. A key advantage of our new methods is that we can establish the existence of rainbow paths of length $0.99|S|$ in random subgraphs of $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$, and this flexibility is crucial for the second step of our 99% → 100% framework.

A central idea is the *dichotomy between structure and randomness* from additive combinatorics. We will decompose our given subset $S \subseteq \mathbb{F}_2^n$ into a "structured" part and a "random-like" part. We measure structure/randomness according to the *doubling constant* $|S+S|/|S|$, where we have written $S+S := \{x+y \colon x, y \in S\}$. Small doubling corresponds to structure; and its opposite is "everywhere-expansion", in the following sense.

Definition 2.2 (Everywhere-expansion)**.** Let $\gamma, K > 0$. A subset $E \subseteq \mathbb{F}_2^n$ is (γ, K)-*everywhere-expanding* if every subset $E' \subseteq E$ of size $\gamma|E|$ satisfies $|E'+E'| \geq K|E'|$.

To obtain our decomposition of S, we iteratively remove subsets of size at least $\gamma|S|$ and doubling at most K as long as such subsets exist; the remainder is then guaranteed to be (γ, K)-everywhere-expanding. The following lemma codifies the outcome of this procedure.

Proposition 2.3. *Let $K \gg \alpha \gg \gamma > 0$. We can decompose any subset $S \subseteq \mathbb{F}_2^n$ as $S = S_1 \cup S_2 \cup \cdots \cup S_t \cup E$, where*

(1) $|S_i| \geq \gamma|S|$ and $|S_i + S_i| \leq K|S_i|$ for all i;
(2) E is $(\gamma/\alpha, K/\alpha)$-everywhere-expanding set of size $\alpha|S|$.

Here, one should think of the S_i's as the structured pieces of S and of E as the random-like piece. Two extreme possible outcomes of the above lemma are $E = \emptyset$ and $E = S$. In the former case S completely decomposes into structured pieces, while in the latter case all of S is random-like; these two cases naturally require different treatments. Our analysis of the general case splits into two cases depending on the size of E.

We start by illustrating how to solve the 99% problem when the random-like part E is all of S. For this we will need the following standard additive-combinatorial tool (see [44, Lemma 2.6]).

Lemma 2.4 (Ruzsa triangle inequality)**.** *For subsets V, S of an abelian group, we have $|V+S|^2 \geq |V| \cdot |S+S|$.*

[2] Of course, the constants 0.01 and 0.99 serve schematic purposes and should not be taken too literally.

We can now establish a 99%-result for the model case of an everywhere-expanding set S. This case per se does not figure in our main argument, but it serves as an excellent illustration of the ideas involved. The strategy is that we will build a long rainbow path two vertices at a time, and at each step we will make sure that we have enough options to continue extending the path at the subsequent step. Extending two vertices at a time instead of one vertex at a time is what allows us to make use of the everywhere-expanding hypothesis (which guarantees that sumsets of large subsets of S grow).

Proposition 2.5. *Let $0 < \gamma < 1/10$ and $K > 0$ satisfy $K > 10/\gamma^4$. Suppose that $S \subseteq \mathbb{F}_2^n \setminus \{0\}$ is a (γ, K)-everywhere-expanding set of size $|S| \geq 2/\gamma$. Then, $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$ has a rainbow path of length $(1-2\gamma)|S|$.*

Proof. For each $t = 0, 1, 2, \ldots, (1/2-\gamma)|S|$, we will build a rainbow path

$$P_t = (v_0 \to v_1 \to \cdots \to v_{2t})$$

in $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$ such that v_{2t} has at most $\gamma|S|$ neighbours in P_t, i.e.,

$$|(v_{2t} + S) \cap \{v_0, \ldots, v_{2t}\}| \leq \gamma|S|.$$

For $t = 0$, we can take v_0 to be any element of $\mathbb{F}_2^n$. Suppose that we have already constructed P_t and we want to extend it to P_{t+1}. Since v_{2t} has at most $\gamma|S|$ neighbours in P_t, among the $|S| - 2t > 2\gamma|S|$ colours not appearing in P_t, there is a set $S' \subseteq S$ consisting of $2\gamma|S| - \gamma|S| = \gamma|S|$ colours such that

$$(v_{2t} + S') \cap \{v_0, \ldots, v_{2t}\} = \emptyset; \tag{1}$$

let S'' consist of some $\gamma|S|$ of the remaining colours not appearing in P_t. The Ruzsa triangle inequality and the (γ, K)-everywhere-expanding hypothesis give

$$|v_{2t} + S' + S''| \geq \sqrt{|S'| \cdot |S'' + S''|} \geq \sqrt{\gamma|S| \cdot K\gamma|S|} = \sqrt{K} \cdot \gamma|S|. \tag{2}$$

We will obtain the path P_{t+1} by setting

$$v_{2t+1} := v_{2t} + s', \quad v_{2t+2} := v_{2t} + s' + s''$$

for suitable $s' \in S'$, $s'' \in S''$. Our definitions of the sets S', S'' guarantee that P_{t+1} is a rainbow walk; we show that we can choose s', s'' so that this walk is in fact a path. Note that v_{2t+1} is disjoint from P_t by (1) for all choices of $s' \in S'$. We must check that v_{2t+2} does not lie on P_t and that v_{2t+2} has at most $\gamma|S|$ neighbours in $P_t \cup \{v_{2t+1}\}$.

Say that a vertex $v \in \mathbb{F}_2^n$ is *bad* if it either lies on P_t or has at least $(\gamma/2)|S|$ neighbours in P_t. Since there are at most $|S|$ vertices on P_t and each is incident to $|S|$ edges, the number of bad vertices is at most

$$(2t+1) + \frac{|S| \cdot |S|}{(\gamma/2)|S|} \leq |S| + (2/\gamma)|S| < \sqrt{K} \cdot \gamma|S|.$$

So by (2), we can choose $s' \in S'$, $s'' \in S''$ so that v_{2t+2} is not bad. It follows that v_{2t+2} has at most $(\gamma/2)|S| + 1 \leq \gamma|S|$ neighbours in $P_t \cup \{v_{2t+1}, v_{2t+2}\}$, as desired. □

This proof has a fair bit of flexibility. For example, we had plenty of viable choices, say, at least $\frac{1}{2}\sqrt{K}\gamma|S|$ choices, for v_{2t+2} at each step. Now, if P' is a fixed rainbow path of length 1000 (say) with colours not appearing in P_t, then we can append a translate of P' to one of our viable choices for v_{2t+2} in such a way that we still get a path, and that the final vertex of the resulting long rainbow path has few neighbours on the new path itself. In other words, at the cost of using two colours from the given everywhere-expanding set, we can incorporate 1000 *arbitrary* colours into our rainbow path. A careful implementation of this idea leads to the following proposition ensuring a 99% rainbow path in $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$ whenever the unstructured piece of S has size at least $0.01|S|$ (see Theorem 8.10 for more details).

Proposition 2.6. *Let $S \subseteq \mathbb{F}_2^n$, and suppose that there is a $(0.001, 10^{20})$-everywhere-expanding subset $E \subseteq S$ of size at least $0.01|S|$. Then $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$ has a rainbow path of length $0.99|S|$.*

In order to show that $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$ has a rainbow path of length $0.99|S|$ for all choices of S, it remains only to handle the case where at least 99% of S is structured, in the sense of Proposition 2.3. To this end, suppose that at least 99% of S can be expressed as the union of sets $S_1, \ldots, S_t$ each with size at least $\gamma|S|$ and doubling at most K. Notice that $t \leq 1/\gamma$ is bounded. Provided that we can (somewhat flexibly) find a 99% rainbow path in each $\mathrm{Cay}_{\mathbb{F}_2^n}(S_i)$ individually, we will be able to concatenate translates of these paths using ideas similar to those sketched above (see Lemma 8.6 for more details).

With this in mind, let us turn our attention to the 99% problem for a single structured piece. Our analysis of this case starts with the celebrated Freiman–Ruzsa Theorem, which provides a description of sets of small doubling. Green and Tao [22] proved a strong result of this type in $\mathbb{F}_2^n$, and we will use the following slight improvement later formulated in [14].

Theorem 2.7. *Let $K \geq 1$. If $S \subseteq \mathbb{F}_2^n$ satisfies $|S+S| \leq K|S|$, then there is a subspace H of $\mathbb{F}_2^n$ such that $S \subseteq H$ and $|H| \leq 2^{2K}|S|$.*

The recently proven Polynomial Freiman–Ruzsa Conjecture over $\mathbb{F}_2^n$ [18] provides the additional information that any subset $S \subseteq \mathbb{F}_2^n$ of doubling at most K can be covered by $K^{O(1)}$ translates of a "small" subspace of $\mathbb{F}_2^n$. Using this result in place of Theorem 2.7 would improve the quantitative dependencies among the various parameters in our proof, but such an improvement would be inconsequential for the final result Theorem 1.3. Hence, we prefer to work with the conceptually simpler Theorem 2.7 despite its quantitative inefficiency.

Theorem 2.7 effectively reduces the structured case to the case of dense subsets of subspaces of $\mathbb{F}_2^n$, which, of course, are isomorphic to $\mathbb{F}_2^m$ for $m \leq n$. Such a reduction is useful because it gives us access to so-called "robust expansion" tools, as in the work of Lo, Kühn, Osthus, and Staden [32] mentioned above, which generally apply only in the setting of *dense* graphs. We will return to this theme in Section 5; in the meantime we refer the reader to [5, Sections 4 and 5] and [24, 32] for more context.

Once we reduce to the dense case, we can apply a result from [5] (based on robust expansion tools) to obtain a 99% path in each $\mathrm{Cay}_{\mathbb{F}_2^n}(S_i)$. This is not sufficient, however: For other parts of our argument (concatenating the paths for different S_i's and carrying out the later absorption step), we need additional flexibility in prescribing *where* within $\mathrm{Cay}_{\mathbb{F}_2^n}(S_i)$ the 99% path lives. It is here that Theorem 1.5 comes to the rescue by allowing us to pass from the Cayley graph of a dense set to a robust expander whose vertex set corresponds to a subgroup of $\mathbb{F}_2^n$.

We will prove Theorem 1.5 in full generality in Section 4. The proof of Theorem 1.3 requires only the special case of Cayley graphs on $\mathbb{F}_2^n$, where the following slightly stronger result holds.

Lemma 2.8. *Let $\varepsilon \in (0, 1/2)$ and write $N = 2^n$. Let $S \subseteq \mathbb{F}_2^n$ have size $|S| \geq \sigma N$. Then, there is a subspace H of $\mathbb{F}_2^n$ satisfying*

(1) $|S \cap H| \geq (1-\varepsilon)|S|$;
(2) $\mathrm{Cay}_H(S \cap H)$ has no $\varepsilon\sigma/2$-sparse cuts.

The proof of Theorem 1.5 simplifies considerably in the special setting of $\mathbb{F}_2^n$, and we include a separate proof of this case in Section 4 since it is all that is needed for the proof of Theorem 1.3.

In particular, the reader who wishes only to see a proof of Theorem 1.3 need not bother with our nonabelian Fourier-analytic arguments for general groups.

Lemma 2.8 tells us that by sacrificing a tiny proportion of the structured set S, we may assume that S generates a Cayley graph with good expansion properties within the subspace generated by S. This subspace property will later prove useful since we will be able to "jump" among cosets when linking up translates of various paths; see Lemma 5.7 below.

2.2. **99% to 100%-results.** In this subsection we will discuss how to upgrade a 99% result to a 100% result. The main framework has three steps:

Step 1. Build a flexible "absorbing" structure within $\mathrm{Cay}_G(S)$.
Step 2. Run the 99% strategy to obtain a rainbow path using 99% of the colours in S.
Step 3. Use the absorbing structure to integrate the remaining 1% of colours of S into the rainbow path.

Let us break this down step by step.

Step 1. The main idea for building our flexible structure is exploiting *popular sums*. For simplicity, consider the case where the group G is abelian. Suppose $S \subseteq G$ contains elements a, b, c summing to 0, and let d be some other element of S. Then, for any $v \in G$ we can build a path from v to $v+d$ either directly as $v \to v+d$ (using only the colour d) or as

$$v \;\to\; v+a \;\to\; v+a+d \;\to\; v+a+d+b \;\to\; v+a+d+b+c = v+d$$

(using the colours a, b, c, d). See Figure 1. We note that for the latter case, some mild conditions on a, b, c, d are required in order for this to be an actual path rather than a walk. Thus, if we have a rainbow path containing an edge of colour d, and the above alternative route does not intersect the path elsewhere, then we may choose whether or not to add the colours a, b, c in addition to d.

We will see in Section 6 that with some minor caveats (including using 6-tuples instead of triples), we can find not only a single quadruple (a, b, c, d) as above but rather many disjoint such quadruples (a_i, b_i, c_i, d_i) for $1 \le i \le |S|/10$ (say) with $a_i + b_i + c_i = 0$. This is possible in $\mathbb{F}_2^n$ because 0 is a "popular sum" for any sufficiently large subset $S \subset \mathbb{F}_2^n$. In Lemma 6.6, we will see how to string together the gadgets from the previous paragraph to obtain a long rainbow absorbing path in which for each i, there is a shortcut that avoids precisely a_i, b_i, c_i. (When we refer to an absorbing path, we mean the path that takes the long route through each gadget.) Thus we may choose, independently for each i, whether or not to take the colours a_i, b_i, c_i out of our rainbow path. See Figure 2. The benefit of this manoeuvre is that we may later flexibly use the freed-up triples a_i, b_i, c_i elsewhere, and below in Step 3 we will see how this flexibility will turn out to be very useful.

For nonabelian groups G, our absorbing structure is more delicate because we cannot rely on an abundance of small subsets of S with the same product. We will instead use a variant of the so-called "distributive absorption" strategy, first introduced in [35]. We defer further explanation to Section 9.

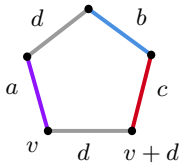

FIGURE 1. Two paths from v to $v+d$, one using only the colour d, and the other using the colours a, b, c, d.

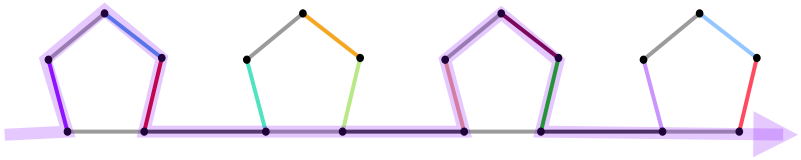

FIGURE 2. An absorbing path of gadgets. The path indicated in purple shows a subpath that uses some triples of colours a_i, b_i, c_i, but not others.

Step 2. We take the last vertex of the absorbing path from Step 1 and use it as the first vertex for a 99% rainbow path as described in the previous subsection. See Figure 3. (More precisely, the 99% path will use 99% of the colours not already used in the absorbing path.) We need to ensure that the absorbing path is vertex-disjoint from the 99% path. In the everywhere-expanding case from the previous subsection, this is not too difficult since we always have enough choices to avoid an absorbing path fixed from the outset. In the structured case, however, we do not have such freedom, so instead we will build the absorbing path and the 99% path in disjoint random subsets of G; this introduces several technical difficulties that we will gloss over for now.

FIGURE 3. An absorbing path connected to a 99% path (drawn dashed).

Step 3. We have now built a rainbow path P that uses 99% of the colours of S and contains a long absorbing path. The heart of the matter is using the flexibility of our absorbing path to integrate the remaining 1% of the colours. Let $\mathcal{L}$ denote the set of "leftover" colours not yet used. The key insight is that we can iteratively reduce the size of $\mathcal{L}$ by "activating" an absorbing gadget a_i, b_i, c_i, d_i and using the freed-up colours a_i, b_i, c_i elsewhere.

As long as $|\mathcal{L}| \ge 3$, choose some three elements $\ell_1, \ell_2, \ell_3 \in \mathcal{L}$. Consider all of the 4-edge extensions of P using the colours $a_i, \ell_1, \ell_2, \ell_3$ in this order, for i ranging over the indices of the absorbing gadgets that have not yet been activated. See Figure 4. This figure is a bit misleading since the 4-edge paths may intersect one another or earlier parts of P, but let us suppose for the moment that we can find some 4-edge path, corresponding to the index i_0, which does not intersect P. Then, we modify P as follows: we "activate" the gadget indexed by i_0 and free up the colours $a_{i_0}, b_{i_0}, c_{i_0}$ by taking the shortcut along the colour d_{i_0}; and we extend P by adding the length-4 path with index i_0. We then update the leftover set $\mathcal{L}$ by removing ℓ_1, ℓ_2, ℓ_3 and adding b_{i_0}, c_{i_0}. In total, we have succeeded in reducing the size of $\mathcal{L}$ by 1.

We have omitted two important technical points from this discussion. The first point concerns ensuring that we can always find a length-4 path that does not intersect other parts of our structure. How we ensure this

FIGURE 4. Various options for extending our long rainbow path by a 4-edge path. To append one of the 4-edge paths, we activate the corresponding gadget.

depends on whether or not S has an everywhere-expanding part. If S does have an everywhere-expanding part, then we can use this expansion to make our candidate length-4 extensions "spread out". We will show that in the remaining case of structured S, we can carry out Steps 1, 2, and 3 in disjoint random vertex subsets of $\mathrm{Cay}_G(S)$, effectively avoiding this issue altogether. The second point is that our iterative procedure terminates once $|\mathcal{L}|$ drops below 3. Saturating the remaining 2 colours is a delicate matter that we will discuss later in the paper.

3. Notation and Prerequisites

Given a set X, a *p-random subset* of X is obtained by sampling each element of X with probability p, independently of all other elements.

We will need the following basic concentration bound.

Lemma 3.1 (Chernoff's inequality)**.** *Let X be a sum of independent Bernoulli random variables with $\mathbb{E}(X) = \mu$. Then, for every $t > 0$,*

- $\mathbb{P}(X \leq \mu - t) \leq \exp(-t^2/(2\mu))$;
- $\mathbb{P}(X \geq \mu + t) \leq \exp(-t^2/(2\mu + t))$.

The digraphs we consider are loopless, and for each pair (u, v) of distinct vertices, we allow at most one edge from u to v, which we denote by (u, v). We do, however, allow both edges (u, v) and (v, u) to appear in the same digraph. If G is a (possibly edge-coloured) digraph, then for $U, V \subseteq V(G)$, we write $e_G(U, V)$ to denote the number of edges (u, v) with $u \in U$ and $v \in V$. As special cases, for a vertex $v \in V(G)$, we denote the *out-degree* of v by $\deg^+_G(v) := e_G(\{v\}, V(G))$ and the *in-degree* of v by $\deg^-_G(v) := e_G(V(G), \{v\})$. We denote the minimum out-degree and in-degree of G by $\delta^+(G)$ and $\delta^-(G)$, respectively, and we write $\delta^\pm(G) := \min\{\delta^+(G), \delta^-(G)\}$ for the *minimum semi-degree* of G.

Nonabelian Fourier analysis. We shall make use of some nonabelian Fourier analysis for finite groups in order to prove the regularity result in Theorem 1.5; we record all the basic properties that we need here. Again, we mention that to prove Lemma 2.8, it suffices to use Fourier analysis over $\mathbb{F}_2^n$. The reader who wishes to focus on this result may skip ahead to the end of this section, where we separately state the basic properties of Fourier analysis over $\mathbb{F}_2^n$.

Let G be a finite (possibly nonabelian) group. We use the following standard notation:

- $|G|$: the order of G,
- $\widehat{G}$: the set of irreducible complex representations of G,
- $\rho \in \widehat{G}$: a representation $\rho : G \to \mathrm{GL}(V_\rho)$, which means that ρ is a group homomorphism from G to $\mathrm{GL}(V_\rho)$,
- $d_\rho = \dim V_\rho$ is the *degree* of the representation ρ, and note that $\sum_{\rho \in \widehat{G}} d_\rho^2 = |G|$.

We will write triv for the trivial irreducible representation. For a vector $v \in V_\rho$, we will write $\|v\|^2_{V_\rho} = \langle v, v\rangle_{V_\rho}$ where we take $\langle \cdot, \cdot \rangle_{V_\rho}$ to be a Hermitian inner product in each of the vector spaces V_ρ, for each irreducible representation ρ. By Weyl's unitary trick, we can and will always assume that each of the representations ρ is unitary with respect to $\langle \cdot, \cdot \rangle_{V_\rho}$, meaning that all the matrices $\rho(g), g \in G$ are unitary so that $\langle \rho(g)v, \rho(g)v\rangle_{V_\rho} = \langle v, v\rangle_{V_\rho}$ for all $v \in V_\rho$ and $g \in G$. Recall also that a matrix A is said to be *unitary* if it satisfies $\overline{A}^T = A^{-1}$.

For a function $f : G \to \mathbb{C}$, the Fourier transform of f at $\rho \in \widehat{G}$ is given by

$$\widehat{f}(\rho) = \sum_{x\in G} f(x)\,\rho(x) \in \mathbb{C}^{d_\rho \times d_\rho}.$$

We shall exclusively consider Fourier transforms $\hat{1}_T(\rho), \rho \in \hat{G}$ of indicator functions of sets $T \subset G$ in this paper. We note that the eigenvalues of the Fourier coefficient matrices $\hat{1}_T(\rho), \rho \in \hat{G}$ are precisely the eigenvalues of the adjacency matrix of the directed graph $\mathrm{Cay}_G(T)$. In fact, each of the d_ρ eigenvalues of $\hat{1}_T(\rho)$ appears with multiplicity d_ρ as an eigenvalue of the adjacency matrix of $\mathrm{Cay}_G(T)$ (which has $|G| = \sum_{\rho\in\hat{G}} d_\rho^2$ eigenvalues in total). The value of the function f at $y \in G$ can be recovered from its Fourier transform via the inversion formula

$$f(y) = \frac{1}{|G|} \sum_{\rho\in\widehat{G}} d_\rho \cdot \mathrm{Tr}\left(\widehat{f}(\rho)\cdot\rho(y)^*\right)$$

where $\rho(x)^* = \overline{\rho(x)}^T$ denotes the conjugate transpose of $\rho(x)$, and where Tr is the trace. Parseval's identity states that for $f : G \to \mathbb{C}$ we have

$$\sum_{x\in G} |f(x)|^2 = \frac{1}{|G|} \sum_{\rho\in\widehat{G}} d_\rho \cdot \|\widehat{f}(\rho)\|_F^2$$

where $\|\cdot\|_F$ is the Frobenius norm $\|A\|_F^2 = \sum_{i,j} |A_{ij}|^2 = \mathrm{Tr}(AA^*)$.

We shall also make use of the following important property for the Fourier transform of the convolution of two functions $f, g : G \to \mathbb{C}$, which is defined as $(f * g)(x) = \sum_{y\in G} f(y)g(y^{-1}x)$. The Fourier transform of the convolution satisfies $\widehat{f * g}(\rho) = \widehat{f}(\rho)\cdot\widehat{g}(\rho)$.

An important observation is that, for sets $X, Y, T \subset G$, the number of solutions $(x, y, t) \in X \times Y \times T$ to $xyt = \mathrm{id}$ can be written using convolutions as $1_X * 1_Y * 1_T(\mathrm{id})$, and hence the formula for the Fourier transform of a convolution and Fourier inversion allow us to express this count as $\frac{1}{|G|}\sum_{\rho\in\widehat{G}} d_\rho \cdot \mathrm{Tr}\left(\widehat{1}_X(\rho)\hat{1}_Y(\rho)\hat{1}_T(\rho)\right)$. Let us state one more elementary fact, namely that the Fourier transform of the indicator function of the whole group G is given by

$$\hat{1}_G(\rho) = \begin{cases} |G|, & \text{if } \rho = \mathrm{triv} \\ (0)_{d_\rho\times d_\rho}, & \text{if } \rho \in \hat{G} \text{ is non-trivial,} \end{cases}$$

where $(0)_{d_\rho\times d_\rho}$ denotes the d_ρ by d_ρ zero matrix. These are all the properties that we require in this paper, for a more extensive overview of the basics of nonabelian Fourier analysis, we refer the reader to [15].

Fourier analysis over $\mathbb{F}_2^n$. For the convenience of the reader who wants a streamlined proof of Theorem 1.3, we briefly discuss the results above in the specialised setting where $G = \mathbb{F}_2^n$. The dual group $\hat{G} = \hat{\mathbb{F}}_2^n$ of characters on $\mathbb{F}_2^n$ is isomorphic to $\mathbb{F}_2^n$, with each character $\gamma \in \hat{G}$ being of the form

$$\gamma_\xi : \mathbb{F}_2^n \to \mathbb{R} : x \mapsto (-1)^{\langle\xi,x\rangle}$$

for a $\xi \in \mathbb{F}_2^n$, where $\langle\xi, x\rangle = \sum_{i=1}^n \xi_i x_i$ is the standard dot product over $\mathbb{F}_2^n$. For a function $f : \mathbb{F}_2^n \to \mathbb{C}$, its Fourier transform is

$$\widehat{f}(\gamma) = \sum_{x\in\mathbb{F}_2^n} f(x)\gamma(x).$$

We have the inversion formula $f(x) = \frac{1}{|G|}\sum_{\gamma\in\hat{G}} \widehat{f}(\gamma)\gamma(x)$, and Parseval's formula states that

$$\sum_{x\in\mathbb{F}_2^n} |f(x)|^2 = \frac{1}{|G|}\sum_{\gamma\in\hat{G}} |\widehat{f}(\gamma)|^2.$$

For $f, g : \mathbb{F}_2^n \to \mathbb{C}$, their convolution $f * g$ is defined as $(f * g)(x) = \sum_{y\in\mathbb{F}_2^n} f(y)g(x + y)$, and the Fourier transform of a convolution satisfies $\widehat{f * g}(\gamma) = \widehat{f}(\gamma)\widehat{g}(\gamma)$. Finally, we note that the Fourier transform of the indicator function of the whole group $G = \mathbb{F}_2^n$ is given by

$$\hat{1}_G(\gamma) = \begin{cases} |G|, & \text{if } \gamma = 0 \text{ is the trivial character} \\ 0, & \text{if } \gamma \in \hat{G}\setminus\{0\}. \end{cases}$$

4. A weak nonabelian regularity lemma for finding expander Cayley subgraphs

The goal of this section is to prove the nonabelian regularity lemma Theorem 1.5, which says that given a subset S of a finite group G, we can find a subgroup H of G such that H contains most of S and $\mathrm{Cay}_H(H \cap S)$ is mildly quasirandom. We restate Theorem 1.5 now for the reader's convenience.

Theorem 1.5. *Let $\sigma \in (0,1]$ and $\varepsilon \in (0,1/2)$. Let G be a finite (not necessarily abelian) group, and let $S \subseteq G$ be a subset with density $\sigma = |S|/|G|$. Then there is a subgroup H of G such that:*

(1) $|S \cap H| \geq (1-\varepsilon)|S|$;
(2) all non-trivial eigenvalues of the adjacency matrix of $\mathrm{Cay}_H(S \cap H)$ have real part at most $(1-\eta)|S \cap H|$, where $\eta := \varepsilon\sigma^2/1000$.

Let us digress and make a few remarks. First, the quadratic dependence of η on σ is optimal, as illustrated by the case where S is an arithmetic progression in $\mathbb{F}_p$. Second, our arguments can be modified to produce a subgroup H and an element $x \in G$ such that $|S \cap (x^{-1}H)| \geq (1-\varepsilon)|S|$, and all non-trivial eigenvalues of the adjacency matrix of $\mathrm{Cay}_H(H \cap (xS))$ have *absolute value* (rather than real part) at most $(1-\eta)|(xS) \cap H|$. Third, another minor variation of the proof shows that there is a subgroup H such that $|S \cap H| \geq (1-\varepsilon)|S|$, and all non-trivial eigenvalues of the adjacency matrix of $\mathrm{Cay}_H(S \cap H)$ have real part at most $(1-\eta)|S \cap H|$, where $\eta := \frac{\varepsilon\delta^2}{1000|\log\sigma|}$ and $\delta := |S \cap H|/|H|$ is the density of S within H (rather than the density σ of S within G).

Recall from Section 3 that the spectrum of the adjacency matrix of $\mathrm{Cay}_G(S)$ is precisely the union of the spectra of the Fourier coefficient matrices $\hat{1}_S(\rho)$ for $\rho \in \hat{G}$. Thus, the second condition in Theorem 1.5 can be formulated in terms of $S \cap H \subseteq H$ having a spectral gap bounded away from 0, in the sense of Definition 4.3 below. The connection between spectral gaps, quasirandomness, and edge-expansion is by now a standard theme in spectral graph theory; see, e.g., the survey [31]. Our formulation of this connection, encapsulated in Lemma 4.5 below, relies on the notion of an η-sparse cut (as defined in the introduction following the statement of Theorem 1.5) and leads to the following corollary of Theorem 1.5 that will be convenient for our later applications.

Corollary 4.1. *Let $\sigma \in (0,1]$ and $\varepsilon \in (0,1/2)$. Let G be a finite group (not necessarily abelian), and let $S \subseteq G$ be a subset with density $\sigma = |S|/|G|$. Then there is a subgroup H of G such that*

(1) $|S \cap H| \geq (1-\varepsilon)|S|$;
(2) $\mathrm{Cay}_H(S \cap H)$ has no $\varepsilon\sigma^3/1000$-sparse cuts.

We remind the reader that $\mathrm{Cay}_H(S \cap H)$ has no η-sparse cuts if for every partition $H = X_1 \sqcup X_2$ we have

$$\#\{(x_1,x_2) \in X_1 \times X_2 : x_1^{-1}x_2 \in S\} \geq \eta|X_1||X_2|.$$

Our proof of Theorem 1.5 uses nonabelian Fourier analysis. As a warm-up, we will start by proving Theorem 1.5 for the group $\mathbb{F}_2^n$, where the argument simplifies considerably due to the nature of the Fourier transform on $\mathbb{F}_2^n$. This simplified argument also yields a somewhat better (in fact, optimal) quantitative dependence of the sparse-cut parameter on the density σ. We remark that only this special case is necessary for the proof of Theorem 1.3, so the reader who is interested only in that result may safely skip our treatment of the general case of Theorem 1.5.

4.1. The $\mathbb{F}_2^n$ case.

Here is Lemma 2.8 restated for the reader's convenience.

Lemma 2.8. *Let $\varepsilon \in (0,1/2)$ and write $N = 2^n$. Let $S \subseteq \mathbb{F}_2^n$ have size $|S| \geq \sigma N$. Then, there is a subspace H of $\mathbb{F}_2^n$ satisfying*

(1) $|S \cap H| \geq (1-\varepsilon)|S|$;
(2) $\mathrm{Cay}_H(S \cap H)$ has no $\varepsilon\sigma/2$-sparse cuts.

Before proving this lemma, we will need the following auxiliary result which states that the Cayley graph of a subset $T \subset \mathbb{F}_2^n$ has no sparse cuts provided that it has a spectral gap. As we discussed at the start of this section, results of this flavour are well-known.

Lemma 4.2. *Let $T \subset H = \mathbb{F}_2^n$ have a spectral gap*

$$\max_{\gamma \in \hat{H}:\gamma \neq 0} \hat{1}_T(\gamma) \leq (1-\beta)|T|,$$

then $\mathrm{Cay}_H(T)$ has no $\beta\tau$-sparse cuts, where $\tau = |T|/|H|$.

Proof of Lemma 4.2. Let $H = X_1 \sqcup X_2$ be a partition of $H = \mathbb{F}_2^n$, so that our goal is to show that $\#\{(x_1,x_2) \in X_1 \times X_2 : x_1 + x_2 \in T\} \geqslant \beta\tau|X_1||X_2|$. By the formula for the Fourier transform of a convolution and Fourier inversion, we can write

$$\#\{(x_1,x_2) \in X_1 \times X_2 : x_1 + x_2 \in T\} = 1_{X_1} * 1_{X_2} * 1_T(0) = \frac{1}{|H|}\sum_{\gamma \in \hat{H}} \hat{1}_{X_1}(\gamma)\hat{1}_{X_2}(\gamma)\hat{1}_T(\gamma).$$

Since the contribution from the trivial character $\gamma = 0$ to the right hand side is $|X_1||X_2||T|/|H| = \tau|X_1||X_2|$, it suffices to show that

$$\frac{-1}{|H|}\sum_{\gamma \in \hat{H}:\gamma \neq 0} \hat{1}_{X_1}(\gamma)\hat{1}_{X_2}(\gamma)\hat{1}_T(\gamma) \leq (1-\beta)\tau|X_1||X_2| \tag{3}$$

as this would show precisely that $\#\{(x_1,x_2) \in X_1 \times X_2 : x_1 + x_2 \in T\} \geq \beta\tau|X_1||X_2|$. Note that as X_1, X_2 partition H, we have that

$$\hat{1}_{X_1} + \hat{1}_{X_2} = \hat{1}_H = \begin{cases} |H|, & \text{if } \gamma = 0 \\ 0, & \text{otherwise.} \end{cases}$$

So $\hat{1}_{X_2}(\gamma) = -\hat{1}_{X_1}(\gamma)$ at all non-zero characters γ. Hence,

$$\frac{-1}{|H|}\sum_{\gamma \in \hat{H}:\gamma \neq 0} \hat{1}_{X_1}(\gamma)\hat{1}_{X_2}(\gamma)\hat{1}_T(\gamma) = \frac{1}{|H|}\sum_{\gamma \in \hat{H}:\gamma \neq 0} |\hat{1}_{X_1}(\gamma)|^2\hat{1}_T(\gamma).$$

Now we simply invoke the spectral gap assumption to bound this by

$$\begin{aligned}\frac{(1-\beta)|T|}{|H|}\sum_{\gamma \neq 0}|\hat{1}_{X_1}(\gamma)|^2 &= (1-\beta)|T|\left(|X_1| - \frac{|X_1|^2}{|H|}\right) \\ &= (1-\beta)\tau|X_1||X_2|\end{aligned}$$

where we used Parseval to calculate $\frac{1}{|H|}\sum_{\gamma \neq 0}|\hat{1}_{X_1}(\gamma)|^2 = |X_1| - \hat{1}_{X_1}(0)^2/|H| = |X_1| - |X_1|^2/|H| = |X_1||X_2|/|H|$ as $|X_2| = |H| - |X_1|$ since X_1, X_2 partition H. This establishes (3) and hence completes the proof. □

It is now a simple matter to prove our result over $\mathbb{F}_2^n$ due to the nature of the Fourier transform in this group. Namely, the characters $\gamma \in \hat{\mathbb{F}}_2^n$ are precisely those functions of the form $\gamma_\xi : x \in \mathbb{F}_2^n \mapsto (-1)^{\langle x,\xi\rangle}$ for some vector $\xi \in \mathbb{F}_2^n$, where $\langle x,\xi\rangle = \sum_{j=1}^n x_j\xi_j$ is the standard dot product in $\mathbb{F}_2^n$. Hence, if we write

$$\langle\gamma\rangle^\perp = \{x \in \mathbb{F}_2^n : \gamma(x) = 1\} = \{x \in \mathbb{F}_2^n : \langle x,\xi\rangle = 0\}$$

for the codimension one subspace defined by γ, then for any subset $T \subset \mathbb{F}_2^n$, the Fourier transform of T at γ is simply given by

$$\hat{1}_T(\gamma) = \sum_{x \in T}\gamma(x) = \sum_{x \in T}(-1)^{\langle x,\xi\rangle} = |T \cap \langle\gamma\rangle^\perp| - |T \cap (x_0 + \langle\gamma\rangle^\perp)|,$$

where $x_0 + \langle\gamma\rangle^\perp$ is the non-trivial coset of $\langle\gamma\rangle^\perp$ in $\mathbb{F}_2^n$. In particular, if T has **no** spectral gap (in the sense of Lemma 4.2), one immediately sees that most of T is contained in the proper subspace $\langle\gamma\rangle^\perp$ of $\mathbb{F}_2^n$.

Proof of Lemma 2.8. Let $\varepsilon \in (0,1/2)$ be given. Let $S \subset \mathbb{F}_2^n$ and we define $\sigma = |S|/N$ and $\delta = \varepsilon\sigma/2$. We proceed by a basic density increment argument, starting with $S_0 = S$ and $H_0 = \mathbb{F}_2^n$. We will iteratively construct subgroups $H_j < H_{j-1}$ and sets $S_j := S_{j-1} \cap H_j$ satisfying for all j that:

$$|S_{j+1}| \geq (1 - \frac{\varepsilon}{2^{j+1}})|S_j|. \tag{4}$$

Suppose now that we have constructed $H_j < H_{j-1} < \cdots < H_0$ and $S_i = S \cap H_i$, for $i \leq j$, satisfying (4). Then we certainly have

$$|S_j| \geq |S|\prod_{i=0}^{\infty}(1 - \frac{\varepsilon}{2^{i+1}}) \geq (1-\varepsilon)|S|. \tag{5}$$

So item (1) from the conclusion of Lemma 2.8 is satisfied for all S_j. If there is no δ-sparse cut in $\mathrm{Cay}_{H_j}(S_j)$, then item (2) is also satisfied and we are done. Else, by Lemma 4.2 there is a non-trivial character $\gamma \in \hat{H}_j$

satisfying

$$\hat{1}_{S_j}(\gamma) \geq \left(1 - \delta\frac{|H_j|}{|S_j|}\right)|S_j| \geq \left(1 - \frac{\delta}{2^{j-1}\sigma}\right)|S_j|$$

since we noted that S_j has size at least $(1-\varepsilon)|S| \geq |S|/2$ as $\varepsilon < 1/2$, and hence S_j has density at least $2^{j-1}\sigma$ in H_j because H_j is a subgroup of H_0 of codimension j (note that at each stage i we find H_i which is a proper subgroup of H_{i-1}). Now define $H_{j+1} = \langle\gamma\rangle^\perp$ which is a subspace of H_j of codimension 1. As $\hat{1}_{S_j}(\gamma) = |S_j \cap H_{j+1}| - |S_j \cap (\text{non-trivial coset of } H_{j+1})|$, we get from the bound on the Fourier coefficient at γ from above that $S_{j+1} = S_j \cap H_{j+1}$ has size

$$|S_{j+1}| \geq \left(1 - \frac{\delta}{2^j\sigma}\right)|S_j| \geq \left(1 - \frac{\varepsilon}{2^{j+1}}\right)|S_j|,$$

as $\delta = \varepsilon\sigma/2$. Hence, we have shown that if $\mathrm{Cay}_{H_j}(S_j)$ has a δ-sparse cut, then we can continue and find a proper subgroup $H_{j+1} < H_j$ such that the set $S_{j+1} = S \cap H_{j+1}$ still satisfies (4).

Observe that the process must trivially halt after a finite number of steps, since there is no infinite chain of subgroups H_j in the finite group $\mathbb{F}_2^n$. In fact, since each S_j has density at least $2^{j-1}\sigma$ in H_j, the process terminates after $O(\log 1/\sigma)$ steps. The final set S_j in this process then has the property that $\mathrm{Cay}_{H_j}(S_j)$ has no δ-sparse cuts, and moreover it satisfies (5), so H_j and $S_j = S \cap H_j$ are the desired sets from the conclusion of the lemma. □

4.2. **The general case.** Next, we prove the result in Theorem 1.5 in full generality, i.e. for a general finite group. We emphasise again that this general result is needed to establish Theorem 1.4, but that it is not required for our resolution of the rearrangement problem over $\mathbb{F}_2^n$ (Theorem 1.3). We begin by defining the correct notion, at least for our application, of a spectral gap for subsets of a general (not necessarily abelian) group. Note that for a subset T of a nonabelian group G, its Fourier coefficients $\hat{1}_T(\rho)$ are matrices, rather than scalars as is the case when G is abelian (such as $G = \mathbb{F}_2^n$ in the previous subsection).

Definition 4.3 (Spectral gap)**.** Let H be a finite possibly nonabelian group, and let $T \subset H$. We say that T has a *β-spectral gap* if for every non-trivial irreducible representation $\rho \in \hat{H}$ and every unit vector $v \in V_\rho$ the following holds:

$$\Re\langle\hat{1}_T(\rho)v, v\rangle_{V_\rho} \leq (1-\beta)|T|.$$

Remark 4.4. We note that this definition of the spectral gap is equivalent to the condition that all eigenvalues of the matrices $\hat{1}_T(\rho)$ have real part at most $(1-\beta)|T|$, for all non-trivial irreducible representations ρ. In particular, recalling from Section 3 that the adjacency matrix of the directed graph $\mathrm{Cay}_H(T)$ has precisely the same eigenvalues as the matrices $\hat{1}_T(\rho)$ as ρ ranges over $\hat{H}$, we further note T has a β-spectral gap if and only if all non-trivial eigenvalues of the adjacency matrix of $\mathrm{Cay}_H(T)$ have real part at most $(1-\beta)|T|$.

Intuitively speaking, T has no (or only a small) spectral gap if there is some non-trivial irreducible representation ρ and some vector v such that $\hat{1}_T(\rho)v \approx |T|v$. We also remark that $\hat{1}_T(\rho) = \sum_{t\in T}\rho(t)$ is a sum of $|T|$ unitary matrices, and hence we always have the trivial bound $\Re\langle\hat{1}_T(\rho)v, v\rangle_{V_\rho} \leq \|\hat{1}_T(\rho)v\|_{V_\rho} \leq |T|$ for unit vectors v. The following lemma generalises Lemma 4.2, showing that if a subset T of a finite group H has a spectral gap, then $\mathrm{Cay}_H(T)$ has no sparse cut. In particular, it immediately shows that Corollary 4.1 follows from Theorem 1.5.

Lemma 4.5. *Let H be a finite group, and suppose that $T \subset H$ has a β-spectral gap in the sense of Definition 4.3:*

$$\sup_{\rho\in\hat{H}:\rho\neq\mathrm{triv}} \sup_{\substack{v\in V_\rho \\ \|v\|_{V_\rho}=1}} \Re\langle\hat{1}_T(\rho)v, v\rangle_{V_\rho} \leq (1-\beta)|T|.$$

Then $\mathrm{Cay}_H(T)$ has no $\beta\tau$-sparse cuts, where $\tau = |T|/|H|$.

Proof of Lemma 4.5. Let $H = X_1 \sqcup X_2$ be a partition of H, so that our goal is to show that $\#\{(x_1,x_2) \in X_1 \times X_2 : x_1^{-1}x_2 \in T\} \geqslant \beta\tau|X_1||X_2|$. By the formula for the Fourier transform of a convolution and Fourier inversion, we can write

$$\#\{(x_1,x_2) \in X_1 \times X_2 : x_1^{-1}x_2 \in T\} = 1_{X_1} * 1_T * 1_{X_2^{-1}}(\mathrm{id}_H) = \frac{1}{|H|}\sum_{\rho\in\hat{H}} d_\rho \,\mathrm{Tr}\left(\hat{1}_{X_1}(\rho)\hat{1}_T(\rho)\overline{\hat{1}_{X_2}(\rho)}^T\right),$$

where we used that all irreducible representations are unitary to note that $\hat{1}_{X_2^{-1}}(\rho) = \sum_{x_2 \in X_2} \rho(x_2)^{-1} = \sum_{x_2 \in X_2} \overline{\rho(x_2)}^T = \overline{\hat{1}_{X_2}(\rho)}^T$. Since the contribution from the trivial representation to the right hand side is $|X_1||X_2||T|/|H| = \tau|X_1||X_2|$, and by using that the trace is invariant under cyclic shifts, it suffices to show that

$$\Re \frac{-1}{|H|} \sum_{\rho \in \hat{H}: \rho \neq \text{triv}} d_\rho \operatorname{Tr}(\overline{\hat{1}_{X_2}(\rho)}^T \hat{1}_{X_1}(\rho) \hat{1}_T(\rho)) \leq (1-\beta)\tau|X_1||X_2|, \tag{6}$$

as plugging this into the first equation would show precisely that $\#\{(x_1, x_2) \in X_1 \times X_2 : x_1^{-1}x_2 \in T\} \geq \beta\tau|X_1||X_2|$. Note that as X_1, X_2 partition H, we have that

$$\hat{1}_{X_1} + \hat{1}_{X_2} = \hat{1}_H = \begin{cases} |H|, & \text{if } \rho = \text{triv} \\ 0, & \text{otherwise.} \end{cases}$$

So $\hat{1}_{X_2}(\rho) = -\hat{1}_{X_1}(\rho)$ at all non-trivial representations ρ. Hence,

$$\Re \frac{-1}{|H|} \sum_{\rho \in \hat{H}: \rho \neq \text{triv}} d_\rho \operatorname{Tr}(\overline{\hat{1}_{X_2}(\rho)}^T \hat{1}_{X_1}(\rho) \hat{1}_T(\rho)) = \frac{1}{|H|} \sum_{\rho \in \hat{H}: \rho \neq \text{triv}} d_\rho \Re \operatorname{Tr}(\overline{\hat{1}_{X_1}(\rho)}^T \hat{1}_{X_1}(\rho) \hat{1}_T(\rho)). \tag{7}$$

The matrix $\overline{\hat{1}_{X_1}(\rho)}^T \hat{1}_{X_1}(\rho)$ is conjugate symmetric (so self-adjoint with respect to $\langle \cdot, \cdot \rangle_{V_\rho}$) and positive semi-definite, so there is an orthonormal basis of vectors $v_1, v_2, \ldots, v_{d_\rho}$ of V_ρ which are eigenvectors, and with real non-negative eigenvalues $\lambda_1, \lambda_2, \ldots, \lambda_{d_\rho} \geqslant 0$ whose sum is equal to $\operatorname{Tr}(\hat{1}_{X_1}(\rho)\overline{\hat{1}_{X_1}(\rho)}^T)$. By noting that for any linear map $A : V_\rho \to V_\rho$ and any orthonormal basis w_j we have that $\operatorname{Tr}(A) = \sum_j \langle Aw_j, w_j \rangle_{V_\rho}$, we get for any non-trivial irreducible representation ρ that

$$\begin{aligned} \Re \operatorname{Tr}(\overline{\hat{1}_{X_1}(\rho)}^T \hat{1}_{X_1}(\rho) \hat{1}_T(\rho)) &= \Re \sum_{j=1}^{d_\rho} \langle \overline{\hat{1}_{X_1}(\rho)}^T \hat{1}_{X_1}(\rho) \hat{1}_T(\rho) v_j, v_j \rangle_{V_\rho} \\ &= \Re \sum_{j=1}^{d_\rho} \langle \hat{1}_T(\rho) v_j, \overline{\hat{1}_{X_1}(\rho)}^T \hat{1}_{X_1}(\rho) v_j \rangle_{V_\rho} \\ &= \sum_{j=1}^{d_\rho} \lambda_j \Re \langle \hat{1}_T(\rho) v_j, v_j \rangle_{V_\rho}, \end{aligned}$$

where we used that $\overline{\hat{1}_{X_1}(\rho)}^T \hat{1}_{X_1}(\rho)$ is self-adjoint in the second line. The spectral gap assumption states that for any unit vector v we have an upper bound $\Re \langle \hat{1}_T(\rho) v, v \rangle_{V_\rho} \leqslant (1-\beta)|T|$. Using this spectral gap bound in the equation above, we get the following upper bound for every non-trivial irreducible representation ρ:

$$\begin{aligned} \Re \operatorname{Tr}(\overline{1_{X_1}(\rho)}^T \hat{1}_{X_1}(\rho) \hat{1}_T(\rho)) &\leqslant (1-\beta)|T| \sum_{j=1}^{d_\rho} \lambda_j \\ &= (1-\beta)|T| \operatorname{Tr}(\overline{\hat{1}_{X_1}(\rho)}^T \hat{1}_{X_1}(\rho)), \end{aligned}$$

as $\operatorname{Tr}(\overline{\hat{1}_{X_1}(\rho)}^T \hat{1}_{X_1}(\rho)) = \sum_{j=1}^{d_\rho} \lambda_j$. Finally, we can plug these trace bounds in the right hand side of (7) and bound this by

$$\begin{aligned} \frac{(1-\beta)|T|}{|H|} \sum_{\rho \neq 0} d_\rho \operatorname{Tr}(\overline{\hat{1}_{X_1}(\rho)}^T \hat{1}_{X_1}(\rho)) &= (1-\beta)|T| \left(|X_1| - \frac{|X_1|^2}{|H|} \right) \\ &= (1-\beta)\tau|X_1||X_2| \end{aligned}$$

where we used Parseval to calculate $\frac{1}{|H|} \sum_{\rho \neq \text{triv}} d_\rho \operatorname{Tr}(\overline{\hat{1}_{X_1}(\gamma)}^T \hat{1}_{X_1}(\rho)) = |X_1| - \hat{1}_{X_1}(\text{triv})^2/|H| = |X_1| - |X_1|^2/|H| = |X_1||X_2|/|H|$ as $|X_2| = |H| - |X_1|$ since X_1, X_2 partition H. This establishes (6) and hence completes the proof of the lemma. □

Recall that over $\mathbb{F}_2^n$, a lemma of the type that we just proved could immediately be combined with a density increment argument to conclude Lemma 2.8, basically because a subset T of $\mathbb{F}_2^n$ having no spectral gap is trivially equivalent to most of T being contained in a proper subgroup. Such a statement is more delicate in general groups, and in fact only true in a weaker sense. The next auxiliary lemma is a result of this type that is true in general groups. It states that if a set $T \subset G$ has no β-spectral gap for some β which is sufficiently small in terms of the density of T in G, then one can again conclude that most of T lies in a proper subgroup.

Lemma 4.6. *Let H be a finite group, and suppose that $T \subset H$ is a subset for which there exists a non-trivial irreducible representation ρ and a unit vector $v \in V_\rho$ such that*

$$\Re \langle \hat{1}_T(\rho)v, v\rangle \geqslant (1-\beta)|T|.$$

Let $\tau = |T|/|H|$ and assume that $\beta \leq \tau^2/1000$ (say). Then there is a proper subgroup H' of H which contains at least $|T \cap H'| \geq (1-50\beta/\tau^2)|T|$ of the elements of T.

Proof. Suppose that there exists a non-trivial irreducible representation ρ and a unit vector $v \in V_\rho$ such that $\Re \langle \hat{1}_T(\rho)v, v\rangle_{V_\rho} \geqslant (1-\beta)|T|$. Throughout this proof, ρ will be fixed and hence we will simply write $\|\cdot\|$ and $\langle \cdot, \cdot\rangle$ for $\|\cdot\|_{V_\rho}, \langle \cdot, \cdot\rangle_{V_\rho}$. We will consider the Bohr sets

$$B(\eta) := \{h \in H : \|\rho(h)v - v\| \leq \eta\}$$

for $\eta \in [0,2]$ (these sets also depend on v, but we consider v to be fixed in this proof). Recall that $\Re \langle \rho(t)v, v\rangle \leqslant 1$ for any $t \in T$, as $\rho(t)$ is unitary. From the assumption that $\Re \langle \hat{1}_T(\rho)v, v\rangle \geq (1-\beta)|T|$ we thus deduce that the set

$$T_m := \{t \in T : \Re \langle \rho(t)v, v\rangle \geq 1 - m\beta\}$$

satisfies

$$(1-\beta)|T| \leq \Re \langle \hat{1}_T(\rho)v, v\rangle = \sum_{t\in T} \Re \langle \rho(t)v, v\rangle \leq |T_m| + (1-m\beta)|T \setminus T_m| = |T| - m\beta|T \setminus T_m|$$

and hence has size $|T_m| \geq (1-1/m)|T|$. T_m is not (necessarily) a Bohr set, but we show that it is efficiently contained in a Bohr set. Since v and hence $\rho(t)v$ are unit vectors, we have for $t \in T_m$, i.e. for t satisfying $\Re\langle \rho(t)v, v\rangle \geqslant (1-m\beta)$, that $\|\rho(t)v - v\|^2 = 2 - 2\Re\langle \rho(t)v, v\rangle \leqslant 2m\beta$. We conclude that the set

$$T^{(m)} := T \cap B(\sqrt{2m\beta}) = \{t \in T : \|\rho(t)v - v\| \leqslant \sqrt{2m\beta}\}$$

contains T_m and thus has size $|T^{(m)}| \geq (1-1/m)|T|$. In other words, we have shown that $T^{(m)}$ contains 'most' of T and is contained in a Bohr set $B(\sqrt{2m\beta})$ of rather short width.

We make the following basic observation about the Bohr sets $B(\eta)$: if $x \in B(\eta)$ then $xB(\eta') \subset B(\eta + \eta')$. The proof of this is simply observing that if $\|\rho(x)v - v\| \leqslant \eta$ and $\|\rho(y)v - v\| \leqslant \eta'$, then

$$\begin{aligned}\|\rho(xy)v - v\| &= \|\rho(x)\rho(y)v - v\| \\ &= \|\rho(x)\rho(y)v - \rho(x)v + \rho(x)v - v\| \\ &\leqslant \|\rho(y)v - v\| + \|\rho(x)v - v\| \leqslant \eta + \eta',\end{aligned}$$

by the triangle inequality and as $\rho(x)$ is unitary.

We now choose $m = \tau^2/(50\beta)$ and we will write $\eta := \tau/5$, so we have shown that $T^{(m)} := T \cap B(\sqrt{2m\beta}) = T \cap B(\eta)$ contains at least $(1-1/m)|T| \geqslant (1-50\beta/\tau^2)|T|$ elements of T. In particular, as we are assuming that $\beta \leq \tau^2/1000$, we certainly have that $|T^{(m)}| \geqslant 0.9|T|$. Note also that $T^{(m)}$ satisfies the size requirement from the conclusion of the lemma, so to complete the proof of this lemma, it only remains to show that $T^{(m)}$ is contained in a proper subgroup of H. It thus suffices to show that $B(\eta)$ is contained in a proper subgroup, and to do this we will establish the following claim.

Claim 4.7. *There exists some integer $k < 4/\tau$ such that $B(k\eta) \setminus B((k-1)\eta) = \emptyset$, where $\eta = \tau/5$.*

First let us see how, assuming this claim, we can easily deduce the desired conclusion that $B(\tau/5) = B(\eta)$ is contained in a proper subgroup of H. Indeed, we have the basic fact that $xB((k-1)\eta) \subset B(k\eta)$ for any $x \in B(\eta)$, and hence the claim that $B(k\eta) \setminus B((k-1)\eta) = \emptyset$ implies that $B(\eta) \cdot X \subset X$ where $X = B((k-1)\eta)$. Iterating this, we see that $B(\eta)^j \cdot X \subset X$ for all j and as $X = B((k-1)\eta)$ clearly contains the identity element, X must therefore contain the subgroup generated by $B(\eta)$. Finally, to see that this subgroup is proper we may simply note that X is not the whole of H since

$$X = B((k-1)\eta) \subset B(k\eta) \subset B((4/\tau)\tau/5) \subset B(4/5) = \{h \in H : \|\rho(h)v - v\| \leqslant 4/5\}$$

cannot contain the whole of H: recall the orthogonality relation $\sum_{h\in H} \rho(h) = 0$ which holds as ρ is a non-trivial irreducible representation, so $\sum_{h\in H} \rho(h)v = 0$ and hence

$$|G| = \left\| |G|v - \sum_{g\in G} \rho(g)v \right\| \leqslant \frac{4}{5}|B(4/5)| + 2|G \setminus B(4/5)|$$

which shows that $|G \setminus B(4/5)| \geqslant |G|/10$.

It only remains to prove the claim. Suppose for a contradiction that it is not true, then for each integer $j \leqslant 4/\tau$ we can find an element $h_j \in B(j\eta) \setminus B((j-1)\eta)$. Consider the elements $h_1, h_4, \ldots, h_{3r+1}$ with indices which are 1 (mod 3) up to $4/\tau$. We claim that the sets $h_1B(\eta), h_4B(\eta), \ldots, h_{3r+1}B(\eta)$ are pairwise disjoint subsets of H. Indeed, pick $x \in h_{3i+1}B(\eta)$ and $y \in h_{3j+1}B(\eta)$ for some $i > j$. So $x = h_{3i+1}x_0$ for some $x_0 \in B(\eta)$. Then we calculate

$$\begin{aligned}\|\rho(x)v - \rho(y)v\| &\geqslant \|\rho(x)v - v\| - \|\rho(y)v - v\| \\ &\geq \|(\rho(h_{3i+1})v - v) + \rho(h_{3i+1})(\rho(x_0)v - v)\| - (3j+2)\eta,\end{aligned}$$

where we used that $y \in h_{3j+1}B(\eta) \subset B((3j+2)\eta)$. Hence, using that $\rho(h_{3i+1})$ is unitary (so distance preserving) and that $x_0 \in B(\eta)$, we get

$$\begin{aligned}\|\rho(x)v - \rho(y)v\| &\geqslant \|\rho(h_{3i+1})v - v\| - \|\rho(x_0)v - v\| - (3j+2)\eta \\ &> 3i\eta - (3j+3)\eta\end{aligned}$$

where the strictness in the final inequality holds as $h_{3i+1} \in B((3i+1)\eta) \setminus B(3i\eta)$. As $i > j$ this implies that $x \neq y$ so that indeed the sets $h_{3i+1}B(\eta), h_{3j+1}B(\eta)$ are disjoint as we claimed. Finally, we note that this gives us the required contradiction since we showed above that $T^{(m)} \subset B(\tau/5) = B(\eta)$ and that $T^{(m)}$ contains at least $(1 - 50\beta/\tau^2)|T| \geq 0.9|T| = 0.9\tau|H|$ elements, by the assumption of the lemma that $\beta \leq \tau^2/1000$. Hence the sets $h_{3j+1}B(\tau/5)$ for $1 \leqslant 3j+1 \leqslant 4/\tau$ would give us $4/(3\tau)$ disjoints sets of size at least $0.9\tau|H|$ inside H. This is of course absurd, and we deduce that there must be some $j \leqslant 4/\tau$ for which $B(j\eta) \setminus B((j-1)\eta) = \emptyset$, proving the claim. □

We can now prove Theorem 1.5 by repeatedly applying Lemma 4.6.

Proof. Let $\varepsilon \in (0, 1/2)$ be given. Let $S \subset G$ and we define $\sigma = |S|/|G|$ and $\eta = \varepsilon\sigma^2/1000$. We proceed by a density increment argument, starting with $S_0 = S$ and $H_0 = G$. We will iteratively construct subgroups $H_j < H_{j-1}$ and sets $S_j := S_{j-1} \cap H_j$ satisfying for all j that:

$$|S_{j+1}| \geq \left(1 - \frac{\varepsilon}{2^{j+1}}\right)|S_j|. \tag{8}$$

Suppose now that we have constructed $H_j < H_{j-1} < \cdots < H_0$ and $S_i = S \cap H_i$, for $i \leq j$, satisfying (8). Then we certainly have

$$|S_j| \geq |S| \prod_{i=0}^{\infty} \left(1 - \frac{\varepsilon}{2^{i+1}}\right) \geq (1-\varepsilon)|S|. \tag{9}$$

So item (1) is satisfied for all S_j. Hence, either item (2) is also satisfied in which case the desired conclusion from Theorem 1.5 holds, or the adjacency matrix of $\mathrm{Cay}_{H_j}(S_j)$ has a non-trivial eigenvalue with real part at least $(1-\eta)|S_j|$, where $\eta = \varepsilon\sigma^2/1000$. Following the remark after Definition 4.3, this is equivalent to $S_j \subset H_j$ having no η-spectral gap meaning that there are a non-trivial irreducible representation $\rho \in \hat{H}_j$ and a unit vector $v \in V_\rho$ satisfying

$$\Re \langle \hat{1}_{S_j}(\rho)v, v\rangle_{V_\rho} \geqslant (1-\eta)\,|S_j|.$$

By (9), we have that S_j has size at least $(1-\varepsilon)|S| \geq |S|/2$ as $\varepsilon < 1/2$, and hence S_j has density at least $2^{j-1}\sigma$ in $|H_j|$ because H_j is a subgroup of H_0 of index at least 2^j (note that at each stage i we find H_i which is a proper subgroup of H_{i-1}). In particular, as $\eta = \varepsilon\sigma^2/1000$, we see that $\eta \leq (|S_j|/|H_j|)^2/1000$ so that the assumption of Lemma 4.6 is satisfied. This lemma concludes that there is a proper subgroup $H_{j+1} < H_j$ such that $S_{j+1} = S \cap H_{j+1}$ has size at least

$$|S_{j+1}| \geq \left(1 - \eta \cdot \frac{50}{(|S_j|/|H_j|)^2}\right)|S_j| \geq \left(1 - \eta \cdot \frac{50}{4^{j-1}\sigma^2}\right)|S_j| \geq \left(1 - \frac{\varepsilon}{2^{j+1}}\right)|S_j|,$$

as $\eta = \varepsilon\sigma^2/1000$. Hence, we have shown that if $\mathrm{Cay}_{H_j}(S_j)$ does not satisfy condition (2), then we can continue and find a proper subgroup $H_{j+1} < H_j$ such that $S_{j+1} = S \cap H_{j+1}$ still satisfies (8).

Observe that the process must trivially halt after a finite number of steps, since there is no infinite chain of subgroups H_j in the finite group G. The final set S_j in this process then has the property that $\mathrm{Cay}_{H_j}(S_j)$ satisfies condition (2), and moreover it satisfies (9), so H_j and $S_j = S \cap H_j$ are the desired sets from the conclusion of Theorem 1.5. □

5. A flexible 99% result

The main result of this section is Lemma 5.7, a flexible asymptotic statement about finding rainbow paths in dense robust expander digraphs. This Lemma 5.7 will play a crucial role at several stages in the remainder of the paper.

As we mentioned in the introduction, it is shown in [5] that if $\mathrm{Cay}_G(S)$ is a robust expander, then it contains a rainbow path of length $(1-o(1))|S|$. This result is insufficient for our applications because we will need to obtain the same conclusion even if we restrict to random vertex subsets of $\mathrm{Cay}_G(S)$ and forbid a small number of colours from S. Due to this additional flexibility requirement, our proof of Lemma 5.7 diverges significantly from the approach in [5].

5.1. Tools.

In this section we introduce some notation and previous results.

We start with the following lemma ([36, Lemma 3.8]), which combines Thomason's jumbledness criterion with the Rödl nibble. The power of the lemma is that the set C' can be chosen completely arbitrarily *after* the random sets A', B' are revealed. We say that a tripartite 3-uniform hypergraph is (γ, p, n)-*typical* if each partite set has $(1\pm\gamma)n$ vertices; each vertex has degree $(1\pm\gamma)pn$; and for each pair of vertices u, v in the same partite set, there are $(1\pm\gamma)p^2n$ vertices w in each other partite set such that $(u, w, x), (v, w, y)$ are edges for some x, y in the third partite set. A hypergraph is linear if through every two vertices there exists at most one edge.

Lemma 5.1 ([36], Lemma 3.8)**.** *Let $H = (A, B, C)$ be a $(0, 1, n)$-typical tripartite linear hypergraph, and let $p \geq n^{-1/600}$. Let $A' \subseteq A$ and $B' \subseteq B$ be (not necessarily independent) p-random subsets. Then with probability at least $1-n^{-2}$, the following holds: For any $C' \subseteq C$ of size $(1\pm n^{-0.2})pn$, there is a matching covering all but $2n^{1-1/500}$ vertices in $A' \cup B' \cup C'$.*

We will always use the above lemma in the form of the following corollary, which picks out the special case of the multiplication hypergraph of a group G (which is always $(0, 1, n)$-typical).

Corollary 5.2. *Let G be a group on n elements, and let $p \geq n^{-1/600}$. Let $A, B \subseteq G$ be disjoint p-random subsets. Then with probability at least $1-n^{-2}$, the following holds: For any $C \subseteq G$ of size $(1\pm n^{-0.2})pn$, there is a rainbow matching in $\mathrm{Cay}_G(C)$ from A to B covering all but $2n^{1-1/500}$ vertices in $A \cup B$ and using all but at most $2n^{1-1/500}$ colours from C.*

We next introduce the notion of robust expansion (following [33]). As we will see shortly, robust expansion is implied[3] by the absence of sparse cuts. We will use robust expansion only through our invocation of Lemma 5.5 below (from [5]); the notion will not otherwise figure in the paper.

Definition 5.3 (Robust expansion)**.** Let G be a directed graph on n vertices. For $U \subseteq V(G)$ and $\nu > 0$, the *ν-robust out-neighbourhood* of U in G is the set

$$RN^+_{\nu,G}(U) := \{v \in V(G) : |N^-(v) \cap U| \geq \nu n\}.$$

We say that G is a *robust (ν, τ)-out-expander* if every $U \subseteq V(G)$ with $\tau n \leq |U| \leq (1-\tau)n$ satisfies

$$|RN^+_{\nu,G}(U) \setminus U| \geq \nu n.$$

Similarly, the *ν-robust in-neighbourhood* of U in G is the set

$$RN^-_{\nu,G}(U) := \{v \in V(G) : |N^+(v) \cap U| \geq \nu n\},$$

and we say that G is a *robust (ν, τ)-in-expander* if every $U \subseteq V(G)$ with $\tau n \leq |U| \leq (1-\tau)n$ satisfies

$$|RN^-_{\nu,G}(U) \setminus U| \geq \nu n.$$

We say that an undirected graph G is a *robust (ν, τ)-expander* if the directed graph obtained by replacing each edge with two directed edges (one in each direction) is a robust (ν, τ)-out-expander (or, equivalently, a robust (ν, τ)-in-expander).

The following elementary proposition shows that a graph with no sparse cuts, as in the definition following Theorem 1.5, is a robust expander. After quoting Lemma 5.5 from [5], we will work with only the no-sparse-cuts property in the rest of this paper.

[3] The two notions are in fact equivalent up to a constant factor loss in parameters.

Proposition 5.4. *Let $0 \le \tau \le 1/2$ and $0 \le \zeta \le 1$. Then any digraph H with no ζ-sparse cuts is a $(\zeta\tau/8, \tau)$-robust-out-expander.*

Proof. Set $n := |H|$. Let U be any subset of $V(H)$ of size $\tau n \le |U| \le (1-\tau)n$ Since $U \sqcup (H \setminus U)$ is not a ζ-sparse cut, there must be at least $\zeta\tau(1-\tau)n^2$ edges from U to $H \setminus U$. The $\zeta\tau/8$-non-robust out-neighbourhood of U can pick up at most $(\zeta\tau/8)n^2$ of these edges, so the $\zeta\tau/8$-robust out-neighbourhood of U has size at least $(\zeta\tau(1-\tau) - \zeta\tau/8)n \ge (\zeta\tau/8)n$. □

In a sufficiently dense robust expander, one can find short paths connecting any two given vertices, and one can moreover guarantee that all vertices and edge-colours in the connecting path come from specified random subsets. The following lemma makes this precise (here we we state only one of the several properties in the lemma from [5]). The proof consists of elementary applications of Chernoff's bound and an application of the definition of robust expansion.

Lemma 5.5 ([5], Lemma 4.3)**.** *Let $\nu, \tau, p \le 1$ be positive constants. Let G be a properly edge-coloured directed graph on n vertices, where $p^3\nu^2 n \ge 144 \log n$. Suppose that G is a robust (ν, τ)-out-expander, with $\delta^{\pm}(G) \ge (\nu+\tau)n$. Let $V_0 \subseteq V(G), C_0 \subseteq C(G)$ be independent p-random subsets. Then with probability at least $1 - 5/n$, the following holds:*

For any distinct vertices $u, v \in V(G)$, and for any vertex subset $V_1 \subseteq V_0$ and colour subset $C_1 \subseteq C_0$ with $|V_1|, |C_1| \le (p^3\nu/100)n$, there exists a rainbow directed path of length at most $\nu^{-1} + 1$ from u to v in G whose internal vertices lie in $V_0 \setminus V_1$ and whose colours lie in $C_0 \setminus C_1$.

Iterative applications of this lemma yield the following corollary.

Corollary 5.6. *Let $0 < \nu, \tau, p \le 1$. Let G be a properly edge-coloured directed graph on n vertices, where $p^3\nu^2 n \ge 144 \log n$. Suppose that G is a robust (ν, τ)-out-expander with $\delta^{\pm}(G) \ge (\nu+\tau)n$. Let $V_0 \subseteq V, C_0 \subseteq C(G)$ be independent p-random subsets. Then with probability at least $1 - 5/n$, the following holds:*

For any collection $(v_i, w_i)_{i \in [k]}$ of $k \le \frac{p^3\nu^2}{300} n$ disjoint pairs of vertices, we can find a rainbow collection of vertex-disjoint paths $P_1, \dots, P_k$ (meaning that the union of the P_i's is rainbow), where each P_i goes from v_i to w_i, and the vertices of the P_i's lie in V_0 and use colours from C_0.

Proof. With probability $1 - 5/n$, the conclusion of Lemma 5.5 holds for V_0, C_0. We construct the paths P_i one at a time. Suppose we have already constructed $P_1, \dots, P_\ell$ for some $\ell < k$. Let V_1 denote the union of the internal vertices in $P_1, \dots, P_\ell$, and let C_1 denote the set of colours in $P_1, \dots, P_\ell$. Notice that

$$|C_1|, |V_1| \le (\nu^{-1} + 2)\ell \le (\nu^{-1} + 2)k \le (p^3\nu/100)n.$$

Then Lemma 5.5 with this choice of V_1, C_1 produces the desired path $P_{\ell+1}$ from $v_{\ell+1}$ to $w_{\ell+1}$. □

5.2. **The 99% lemma.** We have nearly arrived at the main lemma, which establishes a very flexible asymptotic result in the dense setting. This lemma allows us to find a rainbow path of length $(1 - o(1))|S|$ inside a (large) random vertex subset of $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$ with high probability. We can in fact guarantee a bit more: For S_F contained in a random $S' \subseteq S$, we want to find a rainbow path in $\mathrm{Cay}_{\mathbb{F}_2^n}(S \setminus S_F)$ of length $(1 - o(1))|S \setminus S_F|$; our lemma guarantees that with high probability, the restriction of $\mathrm{Cay}_{\mathbb{F}_2^n}(S \setminus S_F)$ to our random vertex set contains such a path for *all* eligible choices of S_F simultaneously. This flexibility will be useful later in the argument, for instance when we want our 99% path to avoid the absorbing structure that we set aside initially.

The statement of our lemma involves many different parameters, objects, and quantifiers. To help the reader get their bearing, we gloss over some of the characters involved. The main thrust of the lemma is that a nicely expanding Cayley graph with a generating set S of size at least $8n^{1-1/9500}$ has a rainbow path which uses all but a few colours from S. For our later applications, we will need to be able to impose further restrictions on this long rainbow path:

- If M is a randomly sampled vertex subset, then with high probability for any two vertices u, v we can require the long rainbow path to start at u, end at v, and have all of its internal vertices lying in M.
- We require the path to avoid a small (adversarially-chosen) deterministic vertex set J.
- We also require the colour set of the path to avoid an adversarially-chosen subset S_F of a randomly sampled subsets $S' \subseteq S$.
- Our path should use all but a small fraction of the colours in $S \setminus S_F$.

We now give the formal statement of our flexible 99% lemma.

Lemma 5.7. *Let G be an N-element group. Let $8N^{-1/9500} \le \zeta, \mu, \varepsilon, q \le 1$.*

- *Let $S \subseteq G$ have $|S| \ge \varepsilon N$, and suppose that $\mathrm{Cay}_G(S)$ has no ζ-sparse cuts.*
- *Let $J \subseteq G$ have $|J| \le 2^{-28} q^3 \mu^3 \varepsilon^2 \zeta^2 N$.*
- *Let $M \subseteq G$ be a q-random subset of G, with $q \ge (1+\mu)|S|/N$.*
- *Let $S' \subseteq S$ be a q'-random subset of S, with $q' \le 1 - \mu q/4$.*

Then with probability at least $1 - 7/N$, the following holds for every choice of $S_F \subseteq S'$ and every pair of distinct vertices $u, v \in G$: There exists a rainbow path from u to v in $\mathrm{Cay}_G(S \setminus S_F)$, using all but at most $\mu q N$ colours of $S \setminus S_F$, such that all of the internal vertices of the path lie in $M \setminus J$.

Proof. Let $H := \mathrm{Cay}_G(S)$ and set $\tau := \frac{3}{4}\varepsilon$. Due to Proposition 5.4, the no ζ-sparse cuts hypothesis implies that H is a (ν, τ)-robust out-expander for $\nu := \zeta\tau/8$.

Since $|J| \le \zeta\varepsilon/32 \cdot N \le \zeta\tau/16 \cdot N$, the graph $H \setminus J$ is still a $(\nu/2, \tau)$-robust out-expander with minimum degree at least $|S| - |J| \ge \frac{7}{8}\varepsilon N$. Note that $\frac{7}{8}\varepsilon \ge \zeta\tau/16 + \tau$, so $H \setminus J$ satisfies the minimum-degree requirement of Corollary 5.6.

Let $t := \frac{2^{28}}{q^2\mu^3\zeta^2\varepsilon^2}$. We now randomly partition G (the vertex set of H) into sets $R, M_1, \dots, M_t$ and a junk set by placing each vertex into R with probability $\tilde{p} := \mu q/4$, into each M_i with probability $p := (q - \tilde{p})/t$, and into the junk set otherwise (independently for each vertex). Hence $R \subseteq G$ is a $\tilde{p}$-random subset, each $M_i \subseteq G$ is a p-random subset, and

$$M := R \cup M_1 \cup \ldots \cup M_t \subseteq G$$

is a q-random subset of G; of course the sets $R, M_1, \dots, M_t$ are all disjoint. (We discard the junk set.) Our choice of t guarantees that

$$p \ge \frac{q}{2t} \ge q^3\mu^3\zeta^2\varepsilon^2/2^{29} \ge N^{-1/600}.$$

We can now describe the plan for the proof, depicted schematically in Figure 5. We will use Corollary 5.2 to obtain an almost-complete rainbow matching between M_i and M_{i+1} for each i (using a fresh set of colours for each new pair); this produces a large rainbow path forest with few components. We will then use Corollary 5.6 to find rainbow paths in R (depicted in gray in Figure 5) linking together the components of the path forest; this step will use colours from a reserved random set C'_R.

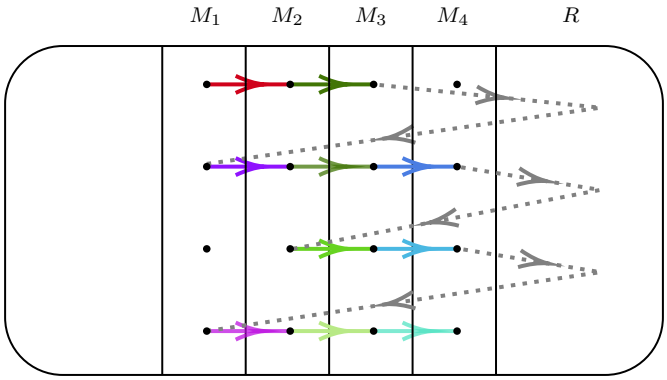

FIGURE 5. An illustration of the argument in Lemma 5.7.

In order to carry out this strategy, we need to upper-bound the probability of failure in our applications of Corollaries 5.2 and 5.6 to various random sets. Let us start with the latter. Let $C_R \subseteq S$ be a $\frac{\tilde{p}}{1-q'}$-random subset, and define $C'_R := C_R \setminus S'$, which is a $\tilde{p}$-random subset of S. Now Corollary 5.6 applied to $H \setminus J$ tells us that with probability at least $1 - 5/(N - |J|)$, the following property holds: We can link any collection of up to

$$\frac{\tilde{p}^3(\nu/2)^2}{300} \cdot (N - |J|) \ge \frac{\tilde{p}^3\zeta^2\varepsilon^2}{2^{19}} \cdot N \tag{10}$$

disjoint pairs of vertices with a rainbow path forest such that the paths use colours only from C'_R and all of their internal vertices lie in $R \setminus J$. Note that the final hypothesis of Corollary 5.6 is satisfied since the right-hand side of (10) is much larger than $\log N$.

The second desirable property is that for each $1 \le i \le t-1$ and *every* subset $C \subseteq G$ of pN colours, the graph $\mathrm{Cay}_G(C)$ contains a rainbow matching between M_i and M_{i+1} covering all but at most $2N^{1-1/500}$ vertices of $M_i \cup M_{i+1}$ and using all but at most $2N^{1-1/500}$ colours of C. This happens for each fixed i with probability at least $1 - 1/N^2$ by Corollary 5.2 applied to the whole of $\mathrm{Cay}_G(G)$ (which is $(0,1,n)$-typical) since M_i, M_{i+1} are both p-random subsets with $p \ge N^{-1/600}$. Notice that independence of M_i, M_{i+1} is *not* required for the application of Corollary 5.2.

By a union bound we can ensure that with probability at least $1 - 7/N$, the properties from the previous two paragraphs simultaneously hold, and we have $|C'_R| \le 2\tilde{p}N$ and $|M_i| \le 2pN$ for all i (using Chernoff bounds). We will now establish the conclusion of the lemma under the assumption that this is the case.

Using the second property, we can find a rainbow matching between M_1 and M_2 using at least $pN - 2N^{1-1/500}$ colours from $S \setminus (S_F \cup C'_R)$. We then remove these newly-used colours from consideration and use the second property to obtain a rainbow matching between M_2 and M_3 using at least $pN - 2N^{1-1/500}$ colours, and so on. We continue until there are fewer than pN unused colours of $S \setminus (S_F \cup C'_R)$ remaining; this happens after at most $t-1$ steps because otherwise we would have used up

$$(t-1)(pN - 2N^{1-1/500}) \ge tpN - pN - 2tN^{1-1/500} = N(q - \tilde{p} - p - 2tN^{-1/500}) > (1-\mu/2)qN \ge |S|$$

colours, which is impossible. (The last inequality uses the hypothesis on the size of q.)

Consider the union of the matchings constructed in the previous paragraph, and throw out all edges incident to $J \cup \{u, v\}$ The remainder is a rainbow (directed) path forest using all but at most

$$pN + |C'_R| + |J| + 2 \le pN + 2\tilde{p}N + \mu qN/4 + 2 \le \mu qN$$

colours of $S \setminus S_F$. Since each matching left at most $2N^{1-1/500}$ uncovered vertices on each side of (M_i, M_{i+1}), the total number of degree-1 vertices in this path forest is at most

$$4pN + t \cdot 2N^{1-1/500} + 2|J| + 4 \le \frac{4q}{t} \cdot N + \frac{2^{29}N}{q^2\mu^3\zeta^2\varepsilon^2 \cdot N^{1/500}} + \frac{q^3\mu^3\zeta^2\varepsilon^2}{2^{27}} \cdot N + 4 \le \frac{q^3\mu^3\zeta^2\varepsilon^2}{2^{25}} \cdot N = \frac{\tilde{p}^3\zeta^2\varepsilon^2}{2^{19}} \cdot N.$$

Fix an ordering $P_1, \dots, P_m$ of the paths in our path forest, where $m \le \frac{\tilde{p}^3\zeta^2\varepsilon^2}{2^{19}} \cdot N$. By the linking-up property guaranteed above, we can find vertex-disjoint paths in $R \setminus J$ using colours in C'_R that connect u to the initial vertex of P_1, connect the final vertex of P_i to the initial vertex of P_{i+1} for each $1 \le i \le m-1$, and connect the final vertex of P_m to v. Putting everything together produces the desired long rainbow path. □

6. The absorption (99% → 100%) lemmas

In this section we prove several lemmas which will allow us to run the absorption argument. We start with the simplest one, in part to illustrate an argument which, in a somewhat more complicated form, will appear in several later lemmas.

Lemma 6.1. *Let $0 < p \le 1$, and let G be a finite group. Suppose $J \subseteq E \subseteq G \setminus \{\mathrm{id}\}$ satisfy*

$$|E|p^2 \ge \max(40|J|, 96 \log |G|).$$

Let A be a p-random subset of G. Then with high probability, we can find, for each vertex $u \in G$, a rainbow path in $\mathrm{Cay}_G(E)$ that starts at u, has all other vertices in A, and contains all of the colours in J.

Proof. Set $N := |G|$. For each vertex $v \in G$ and colour $j \in J$, let $E_{v,j}$ be the event that there are at least $5|J|$ paths of the form

$$v, vg, vgj$$

with $g \in E \setminus \{j\}$ and $vg, vgj \in A$, and these paths are vertex-disjoint except on v.

We will show that these events are very likely. Fix some $v \in G, j \in J$. There are at least $|E| - 2$ candidate paths v, vg, vgj in $\mathrm{Cay}_G(E)$ (since we may have to exclude $g = j^{-1}$ to guarantee $vgj \ne v$), and each such path intersects at most two other paths (since $vg = vg'j$ implies that $g = g'j$). Thus we can greedily find a collection at least $(|E|-2)/3$ disjoint such paths. Each path in this collection survives in A with probability p^2, and these events are independent. Hence the number of surviving paths dominates $\mathrm{Bin}(|E|/4, p^2)$, and a Chernoff bound tells us that at least $|E|p^2/8 > 5|J|$ of them survive with probability at least $1 - \exp(-|E|p^2/32) \ge 1 - 1/N^3$. Thus $\mathbb{P}[E_{v,j}] \ge 1 - 1/N^3$. By a union bound, we conclude that with probability at least $1 - 1/N$ all of the events $E_{v,j}$ simultaneously occur.

Suppose we are in such an outcome. We can find our desired path by starting at u and repeatedly adding a length-2 path containing an arbitrary hitherto-unused element of J. Indeed, since we have at least $5|J|$ candidate extensions at each step, we can ensure that the colour g is hitherto unused (there are at most $2|J|-2$ colours already used) and that the two new vertices do not intersect the part of the path (of length at most $2|J|-2$) that we have already built. □

In the remainder of this section, we shall work specifically over $\mathbb{F}_2^n$ since our absorbing structures for general groups have a very different form.

6.1. **Building an absorbing path.** In this section we describe our absorbing path and show how to find it robustly. By an *ordered subset* of $\mathbb{F}_2^n$ we mean a subset $F \subseteq \mathbb{F}_2^n$ together with an ordering on its elements. We write f_i for the i-th element of F, and we write $\langle F \rangle$ for the subspace generated by F.

Definition 6.2 (Gadget)**.** Let $S \subseteq \mathbb{F}_2^n$. An ordered subset $F \subseteq S$ is a *gadget* in S if $|F| \leq 6$, the elements of F sum to 0, and no proper subset of F is 0-sum. A family $\mathcal{F}$ of gadgets in S is *flexible* if the following all hold:

- **F1** The elements of $\mathcal{F}$ are pairwise disjoint.
- **F2** The sets of partial sums $\{f_1, f_1+f_2, \ldots, f_1+\ldots+f_{|F|-1}\}$ for $F \in \mathcal{F}$ are all disjoint.
- **F3** For any distinct $F_1, F_2 \in \mathcal{F}$, we have $|\langle F_1 \rangle \cap \langle F_2 \rangle| \leq 2$.

Equivalently, F is a gadget if and only $|F| \leq 6$ and starting at any vertex v and following the edges of the colours of F (in order) produces a rainbow cycle. Removing the last edge of such a rainbow cycle produces a rainbow path starting from v associated to the gadget F. If $\mathcal{F}$ is a flexible family of gadgets, then for each vertex v, the rainbow paths from v associated to the gadgets in $\mathcal{F}$ are vertex-disjoint except for v. (This fact uses only **F1** and **F2**. The role of **F3** will become clear later; at a high level, it ensures that different gadgets do not interact too much.) The union of these paths is a rainbow tree which we will refer to as an *out-spider* of $\mathcal{F}$. An *in-spider* of $\mathcal{F}$ is an out-spider of the gadget obtained from $\mathcal{F}$ by reversing the ordering of each gadget.

For example, if v is a vertex and $F = \{f_1, f_2, f_3, f_4\}$ is a gadget in a flexible family $\mathcal{F}$, then the path $v, v+f_1, v+f_1+f_2, v+f_1+f_2+f_3$ forms a leg of the out-spider of $\mathcal{F}$ and $v, v+f_4, v+f_4+f_3, v+f_4+f_3+f_2$ forms a leg of an in-spider of $\mathcal{F}$. Notice that an out-spider and an in-spider with the same starting vertex v have the same vertex set, since for any gadget F we have $\{f_1, f_1+f_2, \ldots, f_1+\ldots+f_{|F|-1}\} = \{f_2+\ldots+f_{|F|}, f_3+f_4+\ldots+f_{|F|}, \ldots, f_{|F|}\}$ due to the 0-sum assumption. See Figure 6 for an illustration.

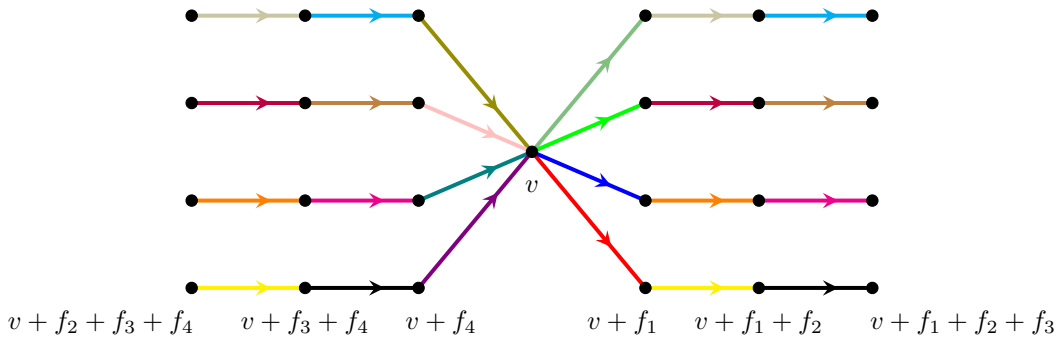

FIGURE 6. An out-spider and an in-spider of a flexible family $\mathcal{F}$. The figure is misleading in representing the out-spider and in-spider on different vertex sets. The bottom two legs correspond to the gadget $F = \{f_1, f_2, f_3, f_4\}$.

Our absorbing structure will allow us to choose, for each gadget $F \in \mathcal{F}$, either to leave all of the colours of F in the absorbing structure or to free them all up for use embedding other colours elsewhere. In an idealised scenario (which provides good intuition), each F would consist of a single colour, and then $\bigcup \mathcal{F}$ would represent a set of flexible colours which we may absorb into our absorbing structure at the very end of the argument if they ended up being unneeded elsewhere. Since of course there are no non-trivial 0-sum single elements, we must package our flexible colours in short tuples (of size at most six), as encoded by our gadgets.

We can find a large flexible family in any reasonably large subset of $\mathbb{F}_2^n$, essentially by the pigeonhole principle.

Lemma 6.3. *Let $0 < \varepsilon \le 1$. If $E \subseteq \mathbb{F}_2^n$ has size $|E| \ge \varepsilon N^{1/2}$, where $N := 2^n$, then E contains a flexible family with at least $\lfloor \varepsilon^2|E|/2^{49} \rfloor$ gadgets.*

Proof. We may assume that $|E| \ge 2^{49}/\varepsilon^2$, as otherwise the statement is trivial. Let us take $\mathcal{F}$ to be a maximal flexible family of gadgets in E. Towards a contradiction, let us assume that $|\mathcal{F}| < \varepsilon^2|E|/2^{49}$. Define the set $B := \bigcup_{F\in\mathcal{F}}\langle F\rangle$ of *blocked* vertices. Note that each $F \in \mathcal{F}$ is a 0-sum set of size at most 6, so $|\langle F\rangle| \le 32$ and hence $|B| \le 32|\mathcal{F}| \le \varepsilon^2|E|/32$.

Now consider triples $\{e_1, e_2, e_3\} \subseteq E$ of linearly independent elements such that $\langle e_1, e_2, e_3\rangle \cap B = \{0\}$. We have at least $(1-\varepsilon^2/32)|E|-1 \ge 2^{-1/3}|E|$ such choices for $e_1 \ne 0$ (ensuring $e_1 \notin B$), then $(1-\varepsilon^2/16)|E|-2 \ge 2^{-1/3}|E|$ choices for $e_2 \notin \langle e_1\rangle$ (ensuring $e_2, e_1+e_2 \notin B$) and $(1-\varepsilon^2/8)|E|-4 \ge 2^{-1/3}|E|$ choices for $e_3 \notin \langle e_1, e_2\rangle$ (ensuring the remaining four subsums are not in B). Since we counted each triple 6 times, there are at least $|E|^3/12$ many such triples. Fix an ordering of the elements of $\mathbb{F}_2^n$, and let s_i denote the number of triples summing to the i-th element of $\mathbb{F}_2^n$. Then $s_1+\ldots+s_N \ge |E|^3/12$, and (by convexity) there are at least $\binom{s_1}{2}+\ldots+\binom{s_N}{2} \ge |E|^6/(512N)$ 6-tuples $(e_1, e_2, e_3, e_4, e_5, e_6)$ such that $e_1 + e_2 + e_3 = e_4 + e_5 + e_6$ and $\langle e_1, e_2, e_3\rangle, \langle e_4, e_5, e_6\rangle$ are disjoint from $B \setminus \{0\}$; let us call such 6-tuples *good*.

Since $\dim(\langle e_1, e_2, e_3, e_4, e_5, e_6\rangle \le 5$, there are at most 32^6 good 6-tuples with a given span. Thus we can find a subcollection of at least $|E|^6/(2^{39}N)$ good 6-tuples spanning pairwise distinct subspaces. We will be done if we can show that some such good 6-tuple $F' = (e_1, e_2, e_3, e_4, e_5, e_6)$ satisfies $|\langle F'\rangle \cap \langle F\rangle| \le 2$ for all $F \in \mathcal{F}$, since then we can add a suitable 0-sum subset of F' to $\mathcal{F}$, contradicting the maximality of $\mathcal{F}$.

There are at most $32^2 \cdot |\mathcal{F}|$ pairs of distinct nonzero elements (a, b) such that $a, b \in \langle F\rangle$ for some $F \in \mathcal{F}$. Each such pair (a, b) is contained in at most $1 + |E| + \binom{|E|}{2} + \binom{|E|}{3} \le |E|^3$ subspaces of the form $\langle F'\rangle$ as F' ranges over our subcollection of good 6-tuples (since any such subspace can be obtained as the span of a, b and at most 3 elements of E). When we range over the pairs (a, b), there are at most

$$32^2 \cdot |\mathcal{F}| \cdot |E|^3 < \varepsilon^2|E|^4/2^{39} \le |E|^6/(2^{39}N)$$

such subspaces in total. In particular, we can choose a good tuple $F' = (e_1, e_2, e_3, e_4, e_5, e_6)$ for which there are no such pairs (a, b); this means that $|\langle e_1, e_2, e_3, e_4, e_5, e_6\rangle \cap \langle F\rangle| \le 2$ for every $F \in \mathcal{F}$. Now let F'' be a minimal 0-sum subset of $\{e_1, e_2, e_3, e_4, e_5, e_6\}$ and fix an ordering of F'' which first traverses the elements from $\{e_1, e_2, e_3\}$ and only afterwards traverses the elements from $\{e_4, e_5, e_6\}$; then F'' is a new gadget which can be added to $\mathcal{F}$, giving the desired contradiction.

Let us check more explicitly that $\mathcal{F} \cup \{F''\}$ is a flexible family. **F1** holds as we chose each $e_i \notin \bigcup_{F\in\mathcal{F}}\langle F\rangle$. We chose F' to satisfy $|\langle F'\rangle \cap \langle F\rangle| \le 2$ for all $F \in \mathcal{F}$; a fortiori the same holds with F' replaced by F'', so **F3** holds. It remains to verify **F2**. Write $F'' = \{e_{i_1}, \ldots, e_{i_t}\}$. Each partial sum $e_{i_1} + \cdots + e_{i_r} = e_{i_{r+1}} + \cdots + e_{i_t}$ is in either $\langle e_1, e_2, e_3\rangle$ or $\langle e_4, e_5, e_6\rangle$ according to whether or not $i_r \le 3$; either way, the sum is by construction not in B. □

To gain intuition on a first read-through, the reader may wish to think of the properties **F1**–**F3** in the definition of a flexible family as saying that $\langle F\rangle \cap \langle F'\rangle = \{0\}$ for distinct $F, F' \in \mathcal{F}$. This stronger property implies all of **F1**–**F3**. There is, however, one instance where we wish to find such a family but we will not be able to ensure this stronger zero-intersection property.

The following easy proposition allows us to obtain a short rainbow path from a gadget and an arbitrary element not in the gadget. This will come in handy in several places.

Proposition 6.4. *Let F be a gadget, and let $x \notin F$ be any nonzero element. Then we can order the elements of F in such a way that x is not equal to any contiguous subsum of F. In particular, inserting x into this ordering of F in any position except for the first or the last produces a valid ordering of $F \cup \{x\}$.*

Proof. If $x \notin \langle F\rangle$, then any ordering of F will do, so suppose that $x \in \langle F\rangle$. Since F is a gadget, it has no nontrivial zero subsums. Thus there is a nonempty subset $T \subsetneq F$, unique up to complementation, such that

$$x = \sum_{f\in T} f = \sum_{f\in F\setminus T} f.$$

Our task is to show that the elements of F can be ordered in such a way that neither the elements of T, nor the elements of $F \setminus T$ appear as a contiguous subsequence. Since $x \notin F$, we know that $|T|, |F \setminus T| \ge 2$. We can build our desired ordering by taking all but one of the elements of T, then one element of $F \setminus T$, then the last element of T, then the remaining elements of $F \setminus T$ (in any way). □

The next step is incorporating a flexible family of gadgets into an absorbing path in $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$.

Definition 6.5 (Absorbing path)**.** We say a rainbow path P in $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$ is *$\mathcal{F}$-absorbing* for a flexible family $\mathcal{F}$ of gadgets in S if there exists an injective function $c : \mathcal{F} \to S \setminus \bigcup \mathcal{F}$ such that for each $F \in \mathcal{F}$ we can find a subpath of P using precisely the colours in $F \cup c(F)$. We say the colours of P not in $\bigcup \mathcal{F}$ are the *fixed colours* of P.

In an $\mathcal{F}$-absorbing path P, for each $F \in \mathcal{F}$ we can delete the subpath of P consisting of the edges with colours $F \cup c(F)$. Doing so leaves two subpaths of P, which we can join using a single edge of colour $c(F)$ (since F is zero-sum). We will denote the resulting subpath by $P - F$; see Figure 2 for an illustration.

The following lemma will let us find an absorbing path inside a random vertex subset while avoiding a small set of forbidden vertices.

Lemma 6.6. *Let $p \in (0,1]$, let $\mathcal{F}$ be a flexible family of gadgets in $E \subseteq \mathbb{F}_2^n$, and let $U \subseteq \mathbb{F}_2^n$ be a subset of size $|U| \le |\mathcal{F}|$. Suppose $p^8|E| \ge 2^{12}|\mathcal{F}| \ge 2^{13}n$. Let R be a p-random subset of $\mathbb{F}_2^n$. Then with high probability, we can find, for each $u \in \mathbb{F}_2^n$, an $\mathcal{F}$-absorbing rainbow path in $\mathrm{Cay}_{\mathbb{F}_2^n}(E)$ of length at most $8|\mathcal{F}|$ that starts at $u \in \mathbb{F}_2^n$ and has all other vertices in $R \setminus U$.*

Proof. Fix $N = 2^n$. First we add a fixed, unique colour $c_F \in E \setminus \bigcup \mathcal{F}$ to each gadget $F \in \mathcal{F}$ and construct a rainbow path P_F that starts at 0 and uses the colours $\{c_F\} \cup F$ (which is possible by Proposition 6.4). Write $P_{F,y}$ for the translate of P_F starting at the vertex $y \in G$. Let

$$X := E \setminus \bigcup_{F \in \mathcal{F}} (\{c_F\} \cup F)$$

be the set of unused colours from E, and notice that that $|X| \ge |E| - 7|\mathcal{F}| \ge |E|/2$. For each vertex $v \in \mathbb{F}_2^n$ and gadget $F \in \mathcal{F}$, we define $E_{v,F}$ to be the event that we can find a collection of at least $10|\mathcal{F}|$ elements $x \in X$ whose corresponding paths $P_{F,v+x}$ are all vertex-disjoint and contained in R.

We will show that these events are (very) likely. Fix some $v \in \mathbb{F}_2^n$, $F \in \mathcal{F}$. We will find many paths $P_{F,v+x}$ which are disjoint and do not contain v. There are $|X|$ paths in total. Of these, at most $|P_F| + 1$ contain v, since the position of v in a translate of P_F determines the translate. Each path $P_{F,v+x}$ can intersect at most $(|P_F|+1)^2$ other such paths, since again the translate of the other path is determined by the relative positions of the intersection point in the two paths. Thus there is a family of $|X|/100$ vertex-disjoint paths $P_{F,v+x}$ avoiding v. Each such path is contained in R with probability $p^{|P_F|+1}$, and these events are independent. Hence the number of surviving paths dominates $\mathrm{Bin}(|X|/100, p^{|P_F|+1})$, and by a Chernoff bound at least

$$|X|p^{|P_F|+1}/200 \ge |E|p^8/400 \ge 10|\mathcal{F}|$$

survive with probability at least $1 - \exp(-|X|p^{|P_F|+1}/800) \ge 1 - 1/N^3$ (using $|E|p^8 \ge 2^{12}|\mathcal{F}| \ge 2^{13}n$). Thus $\mathbb{P}[E_{v,F}] \ge 1 - 1/N^3$, and by a union bound we conclude that with probability at least $1 - 1/N$ all of the events $E_{v,F}$ occur.

Suppose we are in such an outcome. We find our $\mathcal{F}$-absorbing path by incorporating gadgets F one at a time, as in the proof of Lemma 6.1. We start our path P at the vertex u and iteratively add on paths of the form $P_{F,x+v}$, where v is the current endpoint of P. At each step, we identify a hitherto-unincorporated gadget F and consider the $10|\mathcal{F}|$ paths $P_{F,x+v}$ identified in the previous paragraph. Of these, at least $9|\mathcal{F}|$ correspond to values of x that have not yet been used. Since $|P| < 8|\mathcal{F}|$, there are more than $|\mathcal{F}|$ paths $P_{F,x+v}$ that remain disjoint from P. Finally, since $|U| \le |\mathcal{F}|$, we can choose a path $P_{F,x+v}$ that is also disjoint from U (notice that each element of U eliminates at most one choice of x since the paths $P_{F,x+v}$ are vertex disjoint); we choose one such path and concatenate P with it. □

6.2. **The absorbing lemma.** In this subsection we establish a lemma which will eventually allow us to "absorb" any small subset of colours using the flexibility provided by an absorbing path. We also need the ability to work within a random vertex subset and guarantee that we avoid a given small subset of forbidden vertices.

Lemma 6.7. *Let $p \in (0,1]$, let $\mathcal{F}$ be a flexible family of at least $2^{12}p^{-7}n$ gadgets in $S \subseteq \mathbb{F}_2^n \setminus \{0\}$, let $U \subseteq \mathbb{F}_2^n$ be a set of size $|U| \le |\mathcal{F}|/128$. Let $T \subseteq \mathbb{F}_2^n$ be a p-random set. Then with high probability, the following holds for* every *$L \subseteq S$ of size $|L| \le |\mathcal{F}|p^7/2^{12}$ and every vertex $v \in \mathbb{F}_2^n$: There exist a subfamily of gadgets $\mathcal{F}' \subseteq \mathcal{F}$ and a rainbow path in $\mathrm{Cay}_{\mathbb{F}_2^n}(L \cup \bigcup_{F \in \mathcal{F}'} F)$ that starts at v, is otherwise contained in $T \setminus U$, and uses all except possibly one colour from $L \cup \bigcup_{F \in \mathcal{F}'} F$.*

Proof. Consider a pair of distinct colours $a, b \in S$. Our first goal is to construct a subfamily $\mathcal{F}_{a,b} \subseteq \mathcal{F}$ consisting of at least $|\mathcal{F}|/64$ gadgets $F \in \mathcal{F}$ (possibly not inheriting the original orderings $\{f_1, f_2, \dots, f_{|F|}\}$) such that extending each leg of the $\mathcal{F}_{a,b}$-out-spider starting at 0 by the edge of colour a and then the edge of colour b produces a family of vertex-disjoint paths (except for the shared initial vertex 0).

For each $F \in \mathcal{F}$, consider the walk P_F that starts at 0 and then follows the edges of colours $f_1, \dots, f_{|F|-1}, a, b$ (recall that $F = \{f_1, \dots, f_{|F|}\}$). Note that P_F is a bona fide path as long as $a, a+b \notin \langle F \rangle$. We claim that each path P_F can intersect at most 11 other paths $P_{F'}$ at vertices other than 0. Indeed, P_F can intersect $P_{F'}$ only if $\{f'_1, f'_1 + f'_2, \dots, f'_1 + \cdots + f'_{|F'|-1}\}$ intersects the set

$$\{f_1 + \cdots + f_i, f_1 + \cdots + f_i + a, f_1 + \cdots + f_i + a + b : 1 \le i \le |F| - 1\} \cup \{f_1 + \cdots + f_{|F|-1} + b\}.$$

Property **F2** in the definition of flexibility ensures that $f_1 + \cdots + f_i$ can never appear in $\{f'_1, f'_1 + f'_2, \dots, f'_1 + \cdots + f'_{|F'|-1}\}$. This leaves us with at most $2(|F| - 1) + 1 \le 2(5) + 1 = 11$ possible collisions.

It follows that if we can find a collection of $|\mathcal{F}|/3$ gadgets $F \in \mathcal{F}$ for which P_F is a path (as opposed to just a walk), then we can find the desired subfamily $\mathcal{F}_{a,b}$ consisting of at least $|\mathcal{F}|/36$ gadgets F whose corresponding paths P_F are vertex-disjoint (except for 0).

Suppose instead that for some $x \in \{a, a+b\}$ there are at least $|\mathcal{F}|/3$ gadgets $F \in \mathcal{F}$ with $x \in \langle F \rangle$. As the sets $F \in \mathcal{F}$ are disjoint by property **F1** of flexibility, there is at most one such F which contains x; let us remove it from consideration (if it exists). For each remaining F we have $x \in \langle F \rangle \setminus F$; Proposition 6.4 provides an ordering of the elements of F such that x is not equal to any contiguous subsum of F.

The walk P_F with respect to this ordering of F is a bona fide path. The spans of any two such F's intersect precisely in $\langle x \rangle$ by **F3**, so there are no collisions among the sets $\{f_1, f_1 + f_2, \dots, f_1 + \cdots + f_{|F|-1}\}$. We can thus repeat the above argument from the second paragraph of the proof in order to find the desired family $\mathcal{F}_{a,b}$.

Let $\ell := |\mathcal{F}|p^7/2^{12}$. For each vertex $u \in \mathbb{F}_2^n \setminus U$ and two colours $a, b \in S$, let $E_{u,a,b}$ be the event that we can find a collection of at least 10ℓ gadgets $F \in \mathcal{F}_{a,b}$ for which the translate of P_F starting at u is contained in T (except possibly u) and does not intersect U. By the above considerations, there are at least $|\mathcal{F}|/128$ such paths which avoid U. The number of surviving such paths in R dominates $\mathrm{Bin}(|\mathcal{F}|/128, p^7)$. By Chernoff's bound, at least $|\mathcal{F}|p^7/2^8 > 10\ell$ of these paths survive (i.e., $E_{u,a,b}$ occurs) with probability at least $1 - \exp(-|\mathcal{F}|p^7/2^{10}) \ge 1 - 1/N^4$. A union bound over u, a, b ensures that with probability at least $1 - 1/N$ all of the events $E_{u,a,b}$ occur.

Suppose we are in such an outcome. We will construct a sequence of subsets $L = L_0, L_1, \dots, L_m$ and a sequence of directed rainbow paths $P_0 \subset P_1 \subset \cdots \subset P_m$ starting at v such that for each $0 \le i \le m \le |L| - 1$, we have $|L_i| \le |L_{i-1}| - 1$, and the path P_i is contained in $T \setminus U$, has size $|P_i| \le |P_{i-1}| + 7$, and contains $L \setminus L_i$. The path P_m will satisfy the conclusion of the lemma. Suppose we have already done this for some $i < |L| - 1$. Then $|L_i| \ge 2$. Pick some distinct $a, b \in L_i$. By the above considerations, we can find $10\ell - 2 \ge 10|L| - 2$ vertex-disjoint rainbow paths, each of which uses edges with colours a, b and all but one of the elements of some $F \in \mathcal{F}$, and starts at the endpoint of P_i and has all other vertices in $T \setminus U$. One of these paths uses a new F, does not use any of the already used colours and is vertex-disjoint from P_i since $2i + |P_i| \le 9i + 1 \le 9|L|$. Now append this path to P_i to obtain P_{i+1}. To obtain L_{i+1} from L_i, remove $a, b, F \cap L_i$ and add the unused element of F. □

7. Proof of the dense case for $\mathbb{F}_2^n$

In this section we prove Graham's conjecture over $\mathbb{F}_2^n$ in the dense case, namely, the case where the size of the set S is linear in $N := 2^n$. The results of Section 9 show that any subset $S \subset G \setminus \{0\}$ of size $|S| \ge |G|^{1-c}$, in any finite group G, admits a valid ordering, so those results subsume the results in this section. We include a short proof of the weaker result here to demonstrate the implementation of the tools from the previous two sections, which we will also need for the sparse case of $\mathbb{F}_2^n$. We also note that the simpler results in this section already suffice for proving Theorem 1.3 using only the basic absorption argument (similar in spirit to one used in [13]), rather than the distributive absorption tools that we will need for the general dense case in Section 9.

The following theorem handles the extremely dense case. It is convenient to isolate this regime since in the dense-but-not-extremely-dense case our absorption arguments will make essential use of the resulting extra vertex space. The result that we need is contained in [36]; see Appendix A for more details.

Theorem 7.1 ([36]). *Let $\gamma > 0$. Then for all sufficiently large N the following holds: For every group G of order N, every subset $S \subseteq G \setminus \{\mathrm{id}\}$ with $|S| \geq N - N^{1-\gamma}$ has a valid ordering.*

For the rest of this section, assume that $S \subset \mathbb{F}_2^n \setminus \{0\}$ has linear size in $N = 2^n$. The case $|S| \geq \frac{3}{4}N$, which requires a separate treatment due to tighter space constraints, will serve as a simple illustration of our strategy. The main idea is that we first set aside an absorbing path, then find a rainbow path using nearly all of the remaining colours of S, and finally use some of the gadgets from the absorbing path to integrate the remaining few colours of S. To prevent unwanted collisions among the rainbow paths produced in these three steps, we carry out each step in its own random vertex subset.

Theorem 7.2. *Let n be sufficiently large, and set $N := 2^n$. If $S \subseteq \mathbb{F}_2^n \setminus \{0\}$ is a subset of size $|S| \geq \frac{3}{4}N$, then $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$ has a rainbow path of length $|S| - 1$.*

Proof. Set $\gamma := 2^{-14}$. If $|S| \geq N - N^{1-\gamma/32}$, then we are done by Theorem 7.1. It remains to consider the case $|S| \leq N - N^{1-\gamma/32}$. Set $p := N^{-\gamma/16}/2$.

Let E be a $\frac{1}{4}$-random subset of S. Partition $\mathbb{F}_2^n$ into three sets $R \sqcup M \sqcup T$ by independently assigning each vertex to R, M, T with probabilities $p, 1-2p, p$, respectively.

We apply Lemma 5.7 with $S = S$, $J = \emptyset$, $M = M$, $S' = E$ and parameters

$$\varepsilon = 3/4, \quad q = 1 - 2p, \quad q' = 1/4, \quad \zeta = 1/8, \quad \mu = N^{-\gamma}.$$

The graph $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$ has no $\frac{1}{4}$-sparse cuts since $|S| \geq \frac{3}{4}N$, and the other hypotheses of the lemma are easy to check ($|J| = 0$, $1 - 2p = q \geq (1+\mu)(1 - N^{-\gamma/32})$, and $q' \leq 1 - \mu q/4$). Thus with high probability we have:

P1 For any $S_F \subseteq E$ and any two vertices $u, v \in \mathbb{F}_2^n$, we can find a rainbow path from u to v in $\mathrm{Cay}_{\mathbb{F}_2^n}(S \setminus S_F)$, using all but at most μq colours from $S \setminus S_F$, such that all of the internal vertices of the path lie in M.

By a Chernoff bound we have $|E| \geq N/8$ with high probability. In any such outcome, Lemma 6.3 (with $\varepsilon = 1$) lets us find a flexible family $\mathcal{F}$ of gadgets in E of size $Np^8/2^{15}$ (since this is at most $|E|/2^{50}$); fix such a flexible family for each outcome. Lemma 6.6 with $\mathcal{F} = \mathcal{F}, E = E, U = \emptyset, R = R$ (note that $|\mathcal{F}| \leq |E|p^8/2^{12}$) guarantees that with high probability we have:

P2 There is an $\mathcal{F}$-absorbing rainbow path in $\mathrm{Cay}_{\mathbb{F}_2^n}(E)$ starting from any given vertex and otherwise contained in R.

Lemma 6.7 with $\mathcal{F} = \mathcal{F}, S = S, U = \emptyset, T = T$ (note that $|\mathcal{F}| \geq 2^{12}p^{-7}\log N$) shows that with high probability we have:

P3 For any $L \subseteq S$ of size $|L| \leq |\mathcal{F}|p^7/2^{12}$ and any vertex $v \in \mathbb{F}_2^n$, there is some $\mathcal{F}' \subseteq \mathcal{F}$ such that $\mathrm{Cay}_{\mathbb{F}_2^n}(L \cup \bigcup_{F \in \mathcal{F}'} F)$ has a rainbow path that starts at v, is otherwise contained in T, and uses all except possibly one the colours from $L \cup \bigcup_{F \in \mathcal{F}'} F$.

From now on consider an outcome for E, M, R, T where conclusions **P1**-**P3** hold.

Fix some distinct vertices $u, v \in M$. **P2** provides an $\mathcal{F}$-absorbing rainbow path P_A starting at u, otherwise contained in R, and using only colours from E; write S_F for the the set of colours from E appearing in P_A. Now **P1** allows us to find a rainbow path P_M from u to v which is contained in M and saturates all but some set L of up to $\mu q N$ colours from $S \setminus S_F$. Note that $|L| \leq \mu q N \leq N^{1-\gamma} \leq Np^{15}/2^{27} = |\mathcal{F}|p^7/2^{12}$. Now $P_A \cup P_M$ is a rainbow path using precisely the colours in $S \setminus L$. Next, as $|L| \leq |\mathcal{F}|p^7/2^{12}$, by **P3** we can find a subfamily of gadgets $\mathcal{F}' \subseteq \mathcal{F}$ and a rainbow path P_T starting at v and otherwise contained in T which uses all except possibly one colour of $L \cup \bigcup_{F \in \mathcal{F}'} F$. Now we use the $\mathcal{F}$-absorbing properties of P_A to remove $\bigcup_{F \in \mathcal{F}'} F$ and pass to a shorter path P_A' using only a subset of the vertices of P_A (see the illustrations in Section 2.2). Finally, $P_A' \cup P_M \cup P_T$ is a rainbow path using all but one colour from S, as desired. □

We now turn to the main argument for the dense regime. In the very-dense setting of Theorem 7.2, the Cayley graph $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$ was automatically a robust expander. In the merely-dense regime, we have to use our regularity lemma to locate a robustly expanding part of $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$; this requires setting aside a few colours of S that lie outside of the subspace H from Lemma 2.8, and re-integrating these colours causes some additional technical complications. Also, to avoid the case where $S \cap H$ is too dense in H for Lemma 5.7 to apply, we artificially remove a few of these colours and re-integrate them separately, as with the colours in $S \setminus H$.

Theorem 7.3. *Let $\varepsilon \in (0, 1/16)$, and let n be sufficiently large in terms of ε. Set $N := 2^n$. Then for any $S \subseteq \mathbb{F}_2^n$ of size $|S| \geq \varepsilon N$, the graph $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$ has a rainbow path of length $|S| - 1$.*

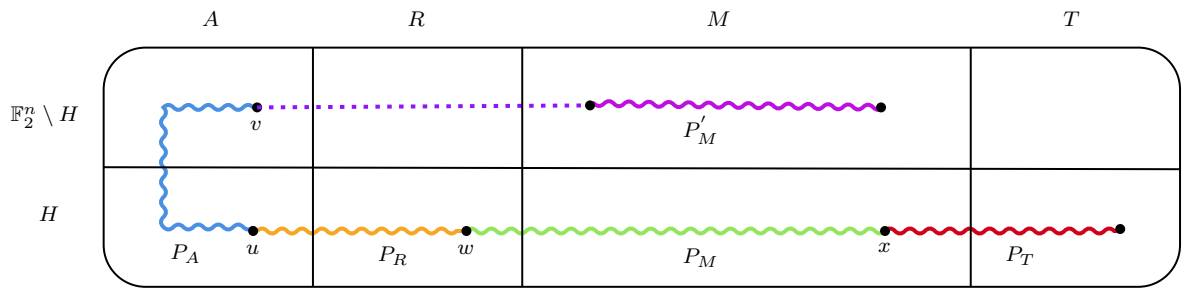

FIGURE 7. Illustration of the rainbow path constructed in the proof of Theorem 7.3. The dashed line indicates that the path P'_M does not intersect R. Colours indicate the different segments of the path, which of course are all rainbow with disjoint colour sets (except for P_T and P_R, whose colours are disjoint only after we activate the appropriate gadgets to replace P_R by P'_R). The picture depicts the case $\sum J \notin H$; the other case looks only marginally different (P_T might "jump" once from $T \cap H$ to $T \setminus H$ at the point where it uses an edge of the colour that we removed from P_A). P_A uses the colours $J \cup S'_F$; P'_M uses the colours $S_1 \cup S''_F$, P_R uses the colours S_F; P_M uses the colours $S_0 \setminus (S_F \cup S'_F \cup S''_F \cup L)$; and P_T uses all but at most one of the colours of $L \cup \bigcup_{F \in \mathcal{F}'} F$.

Proof. We apply our regularity result Lemma 2.8 (with ε replaced by $2\varepsilon^{11}$ and $\sigma = \varepsilon$) to find a subspace H such that $|S \cap H| \geq (1 - 2\varepsilon^{11})|S|$ and $\mathrm{Cay}_H(S \cap H)$ has no ε^{12}-sparse cuts. Identify $H \cong \mathbb{F}_2^m$ and set $J := S \setminus H$; notice that

$$|J| \leq 2\varepsilon^{11}|S|. \tag{11}$$

We now partition S into a large set set S_0 and a small set S_1; how we do so depends on the proportion of H occupied by S. If $|S \cap H| \leq (1 - \varepsilon^3)|H|$, then we set $S_0 := S \cap H$ and $S_1 := \emptyset$. If instead $|S \cap H| > (1 - \varepsilon^3)|H|$, then we let S_0 be an arbitrary subset of $S \cap H$ of size $(1 - \varepsilon^3)|H|$ and set $S_1 := (S \cap H) \setminus S_0$. Notice that either way $|S_1| \leq \varepsilon^3|H|$ and the graph $\mathrm{Cay}_H(S_0)$ has no ε^{12}-sparse cuts[4]. Notice also that when $S_1 \neq \emptyset$, we may assume that $J \neq \emptyset$ since the $J = \emptyset$ case reduces to the situation already handled by Theorem 7.2.

So far, we have a partition $S = S_0 \cup S_1 \cup J$ such that

$$|S_0| \leq (1 - \varepsilon^3)|H|, \quad |S_1| \leq \varepsilon^3|H|, \quad \text{and} \quad \mathrm{Cay}_H(S_0) \text{ has no } \varepsilon^{12}\text{-sparse cuts.} \tag{12}$$

Set $p := \varepsilon^4$. Let S', E_1, E_2 be disjoint $\frac{1}{4}$-random subsets of S_0, and let $A \sqcup R \sqcup M \sqcup T$ be a random partition of $\mathbb{F}_2^n$ where each vertex is (independently) assigned to A, R, M, T with probabilities $p, p, 1 - 3p, p$, respectively. We will activate the lemmas from the previous two sections to show that five desirable properties hold with high probability, and then we will show how to use these properties to find the desired rainbow path in $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$.

We now apply Lemma 5.7 with $S = S_0$, $M = M$, $S' = S' \cup E_1 \cup E_2$, $J = \emptyset$, $\mathbb{F}_2^m \cong H$ and parameters

$$\zeta = \varepsilon^{12}, \quad q = 1 - 3p, \quad \mu = \varepsilon p^{15}/2^{28}, \quad q' = 3/4, \tag{13}$$

and we replace ε in Lemma 5.7 by $\varepsilon/2$. The hypotheses of the lemma are satisfied since

$$|S_0| = |S| - |S_1| - |J| \overset{(12)}{\geq} (\varepsilon - \varepsilon^3 - 2\varepsilon^{11})|H| \geq \frac{\varepsilon}{2}|H|$$

and $q \geq (1 + \mu)|S_0|/|H|$ (which holds since $|S_0| \overset{(12)}{\leq} (1 - \varepsilon^3)|H|$, while $q = 1 - 3p = 1 - 3\varepsilon^4$) and $q' \leq 1 - \mu q/4$. The conclusion of the lemma tells us that with high probability we have:

Q1 For any $S_F \subseteq S' \cup E_1 \cup E_2$ and any distinct vertices $x, w \in H$, we can find a rainbow path in $\mathrm{Cay}_H(S_0 \setminus S_F)$ from x to w, with all other vertices in M, such that the path uses all but at most μq of the colours of $S_0 \setminus S_F$.

[4]This is clear when $S_1 = \emptyset$, and otherwise the large size of S_0 gives the stronger property that $\mathrm{Cay}_H(S_0)$ has no $\frac{1}{4}$-sparse cuts.

The next two properties depend on the random set S'. If $|S'| \leq |S_0|/8$, which by a Chernoff bound occurs with probability $o(1)$, then we declare both properties to fail. So we will work only with outcomes where $|S'| \geq |S_0|/8 \geq |S|/16$. In this case, Lemma 6.3 (with $\varepsilon = 1$) produces a flexible family $\mathcal{F}$ of gadgets in S' of size

$$|\mathcal{F}| = p^8|S|/2^{16}, \tag{14}$$

since this is smaller than $|S|/2^{54} \leq |S'|/2^{50}$.

Lemma 6.6 with $\mathcal{F}, E = S', U = \emptyset$, and R (which is allowed since $|\mathcal{F}| \overset{(14)}{=} p^8|S|/2^{16} \leq |S'|p^8/2^{12}$) tells us that with high probability we have:

Q2 For each vertex $u \in H$, there is an $\mathcal{F}$-absorbing rainbow path in $\mathrm{Cay}_H(S')$ starting from u and otherwise contained in R.

Lemma 6.7 with $\mathcal{F}, U = \emptyset, S$ and T (which is allowed since $|\mathcal{F}| \geq 2^{12}p^{-7}\log N$) tells us that with high probability we have:

Q3 For any $L \subseteq S$ of size $|L| \leq \mu qN + 1$ and any vertex $v \in \mathbb{F}_2^n$, there is some $\mathcal{F}' \subseteq \mathcal{F}$ such that $\mathrm{Cay}_{\mathbb{F}_2^n}(L \cup \bigcup_{F\in\mathcal{F}'} F)$ has a rainbow path that starts at v, is otherwise contained in T, and uses all except possibly one of the colours from $L \cup \bigcup_{F\in\mathcal{F}'} F$.

To check that we may indeed take $|L|$ up to $\mu qN+1$, note that $\mu qN+1 \overset{(13)}{\leq} \varepsilon p^{15}N/2^{28} \leq p^{15}|S|/2^{28} \overset{(14)}{=} |\mathcal{F}|p^7/2^{12}$.

As above, a Chernoff bound tells us that with high probability $|E_1| \geq \frac{1}{8}|S_0| \geq \frac{1}{16}|S| \overset{(11)}{\geq} p^{-2}\cdot\max\{40|J|, 96\log N\}$. In this case we apply Lemma 6.1 with $G = \mathbb{F}_2^n$, $J, E = J \cup E_1$, and our p-random subset A to conclude that with high probability we have:

Q4 For each vertex $u \in \mathbb{F}_2^n$, there is a rainbow path in $\mathrm{Cay}_{\mathbb{F}_2^n}(E_1 \cup J)$, using all of the colours from J, that starts at u and is otherwise contained in A.

Again by Chernoff's bound, we have that with high probability $|E_2|(1-3p)^2 \geq \frac{1}{16}|S| \geq \max\{40|S_1|, 96\log N\}$. In this case, another application of Lemma 6.1, this time with $G = H$, $J = S_1, E = S_1 \cup E_2$, and our $(1-3p)$-random subset M, tells us that with high probability we have:

Q5 For each vertex $u \in \mathbb{F}_2^n$, there is a rainbow path in $\mathrm{Cay}_{\mathbb{F}_2^n}(S_1 \cup E_2)$, using all of the colours from S_1, that starts at u and is otherwise contained in M.

Consider now an outcome for which the properties **Q1**-**Q5** all hold. Fix a vertex $u \in H$.

First, we deal with the junk set J. Using **Q4**, we find a rainbow path P_A, starting at u and otherwise contained in A, such that P_A uses all of the colours from J and some subset S'_F of the colours of E_1. Among all such paths, choose one of minimal length. Let v denote the last vertex in P_A.

Next, we use **Q5** to handle the set S_1. If $S_1 = \emptyset$, then let P'_M be the trivial 1-vertex path at v and set $S''_F := \emptyset$. Now suppose that $S_1 \neq \emptyset$, and recall that in this case we have $J \neq \emptyset$. If $\sum J \notin H$, then we find a rainbow path P'_M, starting at v and otherwise contained in M, such that P'_M uses all of the colours from S_1 and some subset S''_F of the colours of E_2. Notice that P'_M is entirely contained in a single proper H-coset since its starting vertex v is not in H and all of its edges have colours in H. It remains to consider the case where $\sum J \in H$. The minimality of P_A guarantees that the last edge of P_A uses some colour $j^* \in J$. Now truncate P_A by removing v, and let v' denote its new final vertex. Now use **Q5** as above but with v replaced by v' and with J replaced by $J \setminus \{j^*\}$; again note that the resulting P'_M is contained in a proper H-coset. We remark that we will later build another path P_M contained in $M \cap H$, and it will automatically be disjoint from P'_M even though both paths live in the same random subset M; this ability to jump between H-cosets is the key benefit of working in the setting $J \neq \emptyset$.

We now use **Q2** to find an $\mathcal{F}$-absorbing rainbow path P_R, ending[5] at u and otherwise contained in R, whose colour set is some $S_F \subseteq S'$. Let w denote the first vertex of P_R. We have $w \in H$ since $u \in H$ and $S' \subseteq H$.

Now **Q1** gives us a rainbow path P_M in M from x to w which uses all of the colours of $S_0 \setminus (S_F \cup S'_F \cup S''_F)$ except for some set L of size $|L| \leq \mu qN$ (note that **Q1** applies since $S_F \subseteq S'$ and $S'_F \cup S''_F \subseteq E_1 \cup E_2$). The path P_M is fully contained in H since $w \in H$ and $S_0 \subseteq H$; crucially, P_M is disjoint from P'_M. If we removed the last edge

[5]Strictly speaking, we take the reverse of the path produced by **Q2**.

of P_A in our application of **Q5**, then we add j^* to L. So far we have found a rainbow path $P'_M \cup P_A \cup P_R \cup P_M$ which avoids the vertex set T and uses precisely the colours in $S \setminus L$.

Finally, by **Q3** we can find a subfamily of gadgets $\mathcal{F}' \subseteq \mathcal{F}$ and a rainbow path P_T, ending at x and otherwise contained in T, which uses all except possibly one colour of $L \cup \bigcup_{F\in\mathcal{F}'} F$. Since P_R is $\mathcal{F}$-absorbing, we may remove the gadgets in $\mathcal{F}'$ from it to obtain a shorter path P'_R. Now $P'_M \cup P_A \cup P'_R \cup P_M \cup P_T$ is a rainbow path using all but at most one colour from S, as desired. □

8. The sparse case

In this section we treat the sparse case of Theorem 1.3. This takes the following shape.

Theorem 8.1. *There is a constant $\nu > 0$ such that every subset $S \subseteq \mathbb{F}_2^n \setminus \{0\}$ of size $|S| \le \nu \cdot 2^n$ satisfying $S + S = \mathbb{F}_2^n$ has a valid ordering.*

This theorem shows that sparse subsets S of $\mathbb{F}_2^n$ have valid orderings as long as $S + S = \mathbb{F}_2^n$. The following simple lemma shows that this additional assumption is in fact not restrictive.

Lemma 8.2. *Let $S \subseteq \mathbb{F}_2^n$. If $S + S \neq \mathbb{F}_2^n$, then there exists a non-trivial quotient group H of $\mathbb{F}_2^n$ such that the projection map $\pi\colon \mathbb{F}_2^n \to H$ is injective on S. In particular, $\pi(S)$ having a valid ordering in H implies that S has a valid ordering in $\mathbb{F}_2^n$.*

Proof. Let $v \in \mathbb{F}_2^n \setminus (S + S)$, and set $H := \mathbb{F}_2^n/\langle v\rangle$. Now π is injective on S because otherwise we would have distinct $s_1, s_2 \in S$ with $\pi(s_1) = \pi(s_2)$, and then we would have $s_1 + s_2 = v$, which is impossible. The second part of the lemma is obvious. □

Before turning to the proof of Theorem 8.1, let us confirm that Theorem 8.1 and Theorem 7.3 indeed combine to establish Theorem 1.3. Lemma 8.2 shows that it suffices to consider sets $S \subseteq \mathbb{F}_2^n$ with $S + S = \mathbb{F}_2^n$. Now Theorem 8.1 handles the regime $|S| \le \nu N$ (with ν as given by Theorem 8.1), and Theorem 7.3 handles the regime $|S| \ge \nu N$.

Our proof of Theorem 8.1 splits into a "structured" (non-expanding) case and "random-like" (expanding) case. Recall that we say a subset $E \subseteq \mathbb{F}_2^n$ is (γ, K)-*everywhere-expanding* if every subset $E' \subseteq E$ of size $\gamma|E|$ satisfies $|E' + E'| \ge K|E'|$.

8.1. The structured case.

We start with the non-expanding case since it will essentially reduce to (several interdependent instances of) the dense case and the argument is similar to what we saw in the previous section.

8.1.1. *Preliminaries.* We always work with a set $S \subset \mathbb{F}_2^n$ with $|S| \le \nu \cdot 2^n$ and $S + S = \mathbb{F}_2^n$. These assumptions already guarantee that S has at least a bit of expansion. The following lemma lets us set aside a small, well-expanding reservoir of colours for later use. As usual, we omit floor and ceiling functions throughout.

Lemma 8.3. *Let $S \subseteq \mathbb{F}_2^n$, and let $2/|S| \le \gamma \le 1$. Then there is a subset $X \subseteq S$ of size $|X| = \gamma|S|$ such that $|X + X| \ge \frac{\gamma^2}{2}|S + S|$.*

Proof. Take a uniformly random subset X of the specified size. Each element of $S + S$ survives in $X + X$ with probability at least $\gamma \cdot \frac{\gamma|S|-1}{|S|-1} \ge \gamma^2/2$, so $X + X$ has size at least $\frac{\gamma^2}{2}|S + S|$ in expectation. □

We will often use Ruzsa's triangle inequality to translate large doubling of $T + T$ into good expansion of $V + T$ for any other (reasonably large) subset V.

Lemma 8.4 (Ruzsa triangle inequality)**.** *For any subsets V, T of an abelian group, we have $|V + T|^2 \ge |V| \cdot |T + T|$.*

We also need a version of the Freiman–Ruzsa Theorem in $\mathbb{F}_2^n$. An asymptotic formulation of the relevant result was first proven by Green and Tao [22], and we will use the following version due to [14].

Theorem 8.5. *Let $K \ge 0$. If $T \subseteq \mathbb{F}_2^n$ satisfies $|T + T| \le K|T|$, then there is a subspace H of $\mathbb{F}_2^n$ such that $T \subseteq H$ and $|H| \le 2^{2K}|T|$.*

If a set S lacks everywhere-expanding subsets, then we can (almost) partition it into a small number of subsets with small doubling, and each such subset is dense in a (smaller) subspace. We will analyse most of the subsets within their respective subspaces. The following lemma will allow us to link up the resulting pieces. Here and in the subsequent theorem, one should think of ν, γ, K as constants where $K > 0$ is sufficiently large in terms of γ and $\nu > 0$ is sufficiently small in terms of K. We work with concrete dependences among these constants to make the calculations easier to verify.

Lemma 8.6. *Let $0 < \nu, \gamma \le 1 \le K$ satisfy $2\nu^{1/48} \le 2^{1-K} \le \gamma$. Suppose $s, n \in \mathbb{N}$ are such that $8\gamma^{-2} \le s \le \nu N$, where $N := 2^n$. Let $t \le 2/\gamma$, and let $\{H_i\}_{i\in[t]}$ be a sequence of (not necessarily distinct) subspaces of $\mathbb{F}_2^n$, each of size between γs and $2^{2K}s$. If $X \subseteq \mathbb{F}_2^n \setminus \{0\}$ satisfies $|X| \le \gamma s$ and $|X + X| \ge (\gamma^2/5)N$, then there exist $w_1, \dots, w_t \in \mathbb{F}_2^n$ such that the following holds with $W_i := w_i + H_i$:*

(1) For each $i \in [t]$, we have $|W_i \cap \bigcup_{\ell\in[t]\setminus\{i\}} W_\ell| \le \nu^{1/4}s$.
(2) There is a sequence of distinct elements $x_1, y_1, \dots, x_{t-1}, y_{t-1} \in \mathbb{F}_2^n$ such that $x_i \in W_i, y_i \in W_{i+1}$, and $x_i + y_i \in X$ (i.e., (x_i, y_i) is an edge in $\mathrm{Cay}_{\mathbb{F}_2^n}(X)$) for each $i \in [t-1]$, and the $x_i + y_i$'s are distinct.

Proof. We find suitable elements w_{i+1}, x_i, y_i one value of i at a time. Start with $w_1 := 0$, so that $W_1 = H_1$. Suppose we have already found $w_1, \dots, w_m, x_1, y_1, \dots, x_{m-1}, y_{m-1}$ such that

$$\left| W_i \cap \bigcup_{\ell\in[m]\setminus\{i\}} W_\ell \right| \le (m + 1 - i)\nu^{1/4}\gamma s/2$$

for each $i \in [m]$ and the conditions in part (2) of the lemma statement are so far satisfied.

As long as $m < t$, we will find w_{m+1}, x_m, y_m preserving these conditions (with m replaced by $m + 1$).

Set $X_{\text{free}} := X \setminus \{x_1 + y_1, \dots, x_{m-1} + y_{m-1}\}$. Then

$$|X_{\text{free}} + X_{\text{free}}| \ge |X + X| - t|X| \ge (\gamma^2/5)N - t\gamma s \ge (\gamma^2/5)N - 2\nu N \ge (\gamma^2/6)N$$

since $\nu \le \gamma^2/60$ (with room to spare). Likewise, set $W_{\text{free}} := W_m \setminus \{x_1, y_1, \dots, x_{m-1}, y_{m-1}\}$, so that

$$|W_{\text{free}}| \ge |W_m| - 2t \ge \gamma s - 4/\gamma \ge \gamma s/2.$$

Now the Ruzsa triangle inequality (Lemma 8.4) gives

$$|W_{\text{free}} + X_{\text{free}}| \ge \sqrt{|W_{\text{free}}|}\sqrt{|X_{\text{free}} + X_{\text{free}}|} \ge \sqrt{\gamma s/2}\sqrt{(\gamma^2/6)N} \ge \sqrt{\frac{\gamma^3}{12\nu}}s \ge \nu^{-2/5}s,$$

where we used $N \ge s/\nu$ and $\nu \le \gamma^{15}/12^5$ (say).

Thus there are at least $\frac{\nu^{-2/5}s}{2^{2K}s} \ge \nu^{-1/3}$ cosets of H_{m+1} which intersect $W_{\text{free}} + X_{\text{free}}$. At most $2m - 2$ of these cosets intersect $\{x_1, y_1, \dots, x_{m-1}, y_{m-1}\}$, and at most $\frac{2^{3K+1}s}{\nu^{1/4}\gamma s} \le \nu^{-1/3} - 2m$ of them contain at least $\nu^{1/4}\gamma s/2$ elements of $W_1 \cup \cdots \cup W_m$ (the last inequality uses $m \le t \le 2/\gamma$ and $\nu \le 2^{-36K-24}\gamma^{12}$). Hence there is a coset $W_{m+1} = w_{m+1} + H_{m+1}$ which intersects $W_{\text{free}} + X_{\text{free}}$ in some element $y_m \notin \{x_1, y_1, \dots, x_{m-1}, y_{m-1}\}$ and contains at most $\nu^{1/4}\gamma s/2$ elements of $W_1 \cup \cdots \cup W_m$. In particular, there is some $x_m \in W_{\text{free}}$ such that $x_m + y_m \in X_{\text{free}}$ (notice that $x_m \ne y_m$ since $0 \notin X_{\text{free}}$); this choice of x_m, y_m works for our induction.

Once we reach $m = t$, we have $|W_i \cap \bigcup_{\ell\in[t]\setminus\{i\}} W_\ell| \le t\nu^{1/4}\gamma s/2 \le \nu^{1/4}s$ for every $i \in [t]$, as desired. □

8.1.2. *The main argument.* We are now ready to handle the fully-structured case. The main idea is that if S is "fully-structured", then we can decompose it into sets $S_1, \dots, S_t$ of small doubling, each of which is fairly dense in some subspace. We then obtain cosets $W_1, \dots, W_t$ of these subspaces which are nearly disjoint, and in each W_i we build a rainbow path that uses most of the colours of S_i; at the end we join up these short rainbow paths and absorb the "junk set" of hitherto-unused colours. See Figure 8.

In the following theorem, we write $1/s, \nu \ll 1/K \ll \gamma \ll \alpha \ll 1$ to mean that $\alpha \in (0, 1)$ is a sufficiently small constant; γ is sufficiently small in terms of α; K is sufficiently large in terms of γ; and $\nu, 1/s$ are sufficiently small in terms of K.

Theorem 8.7. *Suppose $1/s, \nu \ll 1/K \ll \gamma \ll \alpha \ll 1$. Let $S \subseteq \mathbb{F}_2^n \setminus \{0\}$ be a set of size $s := |S| \le \nu N$ (where $N := 2^n$ as usual), and suppose that $S + S = \mathbb{F}_2^n$. If S has no $(\gamma/\alpha, K/\gamma)$-everywhere-expanding subset E of size $|E| = \alpha s$, then $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$ has a rainbow path of length $|S| - 1$.*

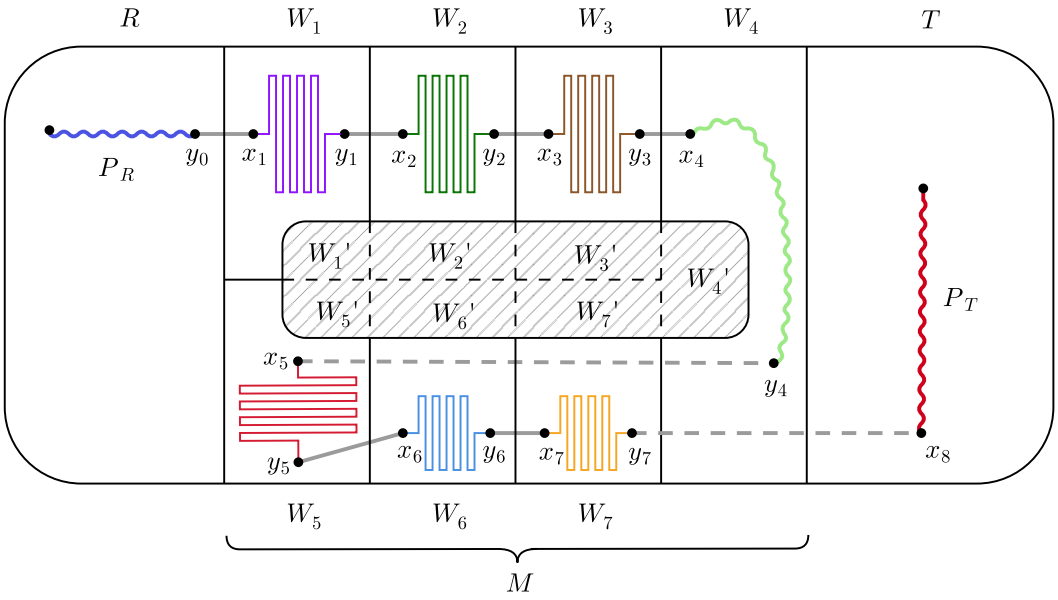

FIGURE 8. Illustration of the argument in Theorem 8.7 with $t = 4, m = 3$. W_1, W_5; W_2, W_6; W_3, W_7 are pairs of cosets of the same subspace, and the W_i's are chosen to have very small intersections W_i' with the other W_j's. We also have fixed "connection points" $y_0, x_1, y_1, \ldots, x_7, y_7, x_8$. P_R is an absorbing path built inside a random subset R. Each $P_{M,i}$ is a rainbow path from x_i to y_i built in the intersection of W_i with a random set M, while avoiding W_i' and all of the other connection points. The path P_T uses up all of the unused colours, with the help of some gadgets activated in P_R.

Proof. By assumption, every subset of S of size at least αs contains a subset of size γs with doubling constant at most K. We first extract $X \subseteq S$ of size γs having $|X + X| \geq \gamma^2 |S + S|/2 = \gamma^2 N/2$ by Lemma 8.3. We can iteratively remove disjoint subsets $S_1, \ldots, S_t \subseteq S \setminus X$, each of size γs (so $t \leq 1/\gamma$), such that each $|S_i + S_i| \leq K|S_i|$ and $S_1 \cup \ldots \cup S_t$ covers all of $S \setminus X$ except for a set J_0 of size at most αs. By Theorem 8.5, this implies that there exist subspaces $H_i \supseteq S_i$ such that $|H_i| \leq 2^{2K}|S_i|$. Next, we can apply the regularity-type result Lemma 2.8 to each S_i (which has density at least 2^{-2K} inside of its H_i) to find a subset $S_i' \subseteq S_i$ of size at least $(1-\alpha)|S_i|$ and a subspace H_i' of H_i containing S_i' such that $\mathrm{Cay}_{H_i'}(S_i')$ has no $\alpha 2^{-2K-1}$-sparse cuts. We add the leftover elements $\bigcup S_i \setminus S_i'$ (there are at most $\sum_i \alpha |S_i| \leq \alpha s$ such elements in total) to the set J_0 to obtain the set J_1, which has size $|J_1| \leq 2\alpha s$.

Let us re-index so that the quantities $|S_i'|/|H_i'|$ are non-increasing with i. Let m be the largest index for which $|S_m'|/|H_m'| \geq 3/4$. Next, we invoke Lemma 8.6 with the sequence of subspaces $H_1', \ldots, H_t', H_1', \ldots, H_m'$ (we list all of the subspaces in order and then list the first m again) and the set X. This gives us cosets $W_1, \ldots, W_{t+m}$ and distinct elements $x_1, y_1, \ldots, x_{t+m-1}, y_{t+m-1} \in \mathbb{F}_2^n$ such that each W_i is a coset of $H_{i \bmod t}'$, each W_i intersects the union of other W_j's in at most $\nu^{1/4} s$ elements, each $x_i \in W_i, y_i \in W_{i+1}, x_i + y_i \in X$, and the quantities $x_i + y_i$ are distinct. Our rainbow path will end up including the $t + m - 1$ edges (x_i, y_i) (with colours $x_i + y_i$). As such, we mark the set of vertices $U := \{x_1, y_1, \ldots, x_{t+m-1}, y_{t+m-1}\}$ as "already used". We also add the unused colours from X, namely, $X \setminus \{x_1 + y_1, \ldots, x_{t+m-1} + y_{t+m-1}\}$, to J_1 to obtain J_2; thus J_2 represents a "junk set" of colours that we will need to absorb into our rainbow path later. Write $W_i' := W_i \cap \bigcup_{j \neq i} W_j$ for each i, and note that $|W_i'| \leq \nu^{1/4} s$.

Next, for each $i \in [t]$, let E_i be a $\frac{1}{4}$-random subset of S_i'. Set $E := \bigcup E_i$.

For each $i \leq m$, take random subsets $S_{i,1}, S_{i,2}$ of S_i' by assigning each element of S_i' independently to $S_{i,1}$ with probability $\frac{1}{4}$, to $S_{i,2}$ with probability $\frac{1}{4}$, and to both $S_{i,1}, S_{1,2}$ with probability $\frac{1}{2}$. So $S_{i,1}, S_{i,2}$ are $\frac{3}{4}$-random subsets of S_i', and $S_{i,1} \cup S_{i,2} = S_i'$. Although $S_{i,1}, S_{i,2}$ are not independent, it is true that after we reveal $S_{i,1}$ (respectively, $S_{i,2}$), the intersection $S_{i,1} \cap S_{i,2}$ is a $\frac{2}{3}$-random subset of $S_{i,1}$ (respectively, $S_{i,2}$).

Set $p := 1/16$. Let $R \sqcup M \sqcup T$ be random partition of $\mathbb{F}_2^n$, where each vertex is assigned to R, M, T with probabilities $p, 1-2p, p$, respectively.

For each i, we want to apply Lemma 5.7 in $W_i \cong \mathbb{F}_2^{n_i}$ with the sets

$$S = \begin{cases} S_{i,1} & \text{if } i \le m, \\ S'_i & \text{if } m < i \le t, \\ S_{i,2} & \text{if } i > t, \end{cases} \quad J = W'_i \cup U, \quad M = M \cap W_i, \quad S' = \begin{cases} E_i \cup (S_{i,1} \cap S_{i,2}) & \text{if } i \le m, \\ E_i & \text{if } m < i \le t, \\ E_i \cup (S_{i,1} \cap S_{i,2}) & \text{if } i > t. \end{cases}$$

Here we choose the parameters[6]

$$q = 1 - 2p, \quad \varepsilon = 2^{-2K-1}, \quad \zeta = \alpha 2^{-2K-1}, \quad \mu = \gamma\alpha 2^{-2K}, \quad q' = \begin{cases} 1/4 & \text{if } m < i \le t, \\ 3/4, & \text{otherwise,} \end{cases} \quad j = 2\nu^{1/4}.$$

Let us say a brief word about the order in which the random sets are revealed. For the case $i \le m$, we first reveal $S_{i,1}$; at this point, the set $S_{i,1} \cap S_{i,2}$ is a $\frac{2}{3}$-random subset of $S_{i,1}$, and we wish to apply Lemma 5.7 with this *deterministic* choice of $S_{i,1}$ (as long as this set has the appropriate size, which is the case with high probability) and the random subset $S_{i,1} \cap S_{i,2}$. The case $i > t$ goes the same way. Let us check that the remaining hypotheses of the lemma hold with high probability (since there are only $2/\gamma = O(1)$ indices i, it suffices to check that the hypotheses hold with high probability for each index individually):

- We have $|S_{i,1}|, |S_{i,2}|, |S'_i| \ge 2^{-2K-1}|W_i|$, since $|S'_i| \ge (1-\alpha)|S_i| \ge (1-\alpha)2^{-2K}|W_i| \ge \frac{3}{4} \cdot 2^{-2K}|W_i|$ (deterministically) and $|S_{i,1}|, |S_{i,2}| \ge \frac{2}{3}|S'_i|$ (with high probability by Chernoff's bound).
- The graphs $\mathrm{Cay}_{W_i}(S_{i,1}), \mathrm{Cay}_{W_i}(S_{i,2}), \mathrm{Cay}_{W_i}(S'_i)$ have no ζ-cuts. This holds for $m < i \le t$ since Lemma 2.8 guarantees that $\mathrm{Cay}_{W_i}(S'_i)$ has no $\alpha 2^{-2K-1}$ cuts. For $i \le m$ and $i > t$, a Chernoff bound guarantees that with high probability $|S_{i,1}|, |S_{i,2}| \ge \frac{17}{32}|W_i|$, since $S_{i,1}, S_{i,2}$ are $\frac{3}{4}$-random subsets of S'_i, which itself has size at least $\frac{3}{4}|H'_i|$; thus $\mathrm{Cay}_{W_i}(S \setminus S')$ lacks $\frac{1}{32}$-sparse cuts just by density considerations.
- We have $q|W_i|/(1+\mu) \ge |S'_i|$ if $m < i \le t$ (since $|S'_i|/|W_i| \le 3/4$ for such i), and $q|W_i|/(1+\mu) \ge |S_{i,1}|, |S_{i,2}|$ if $i \le m$ or $i > t$ (since $|S_{i,j}|/|W_i| \le 5/6$ with high probability, again by Chernoff's bound).
- We have $q' \le 1 - \mu q/4$ with room to spare.

Now Lemma 5.7 tells us that with high probability, for all i we have:

Z1 For every subset S_F of E_i (if $m < i \le t$) or of $E_i \cup (S_{i,1} \cap S_{i,2})$ (if $i \le m$, or $i > t$), we can find a rainbow path in $\mathrm{Cay}_{\mathbb{F}_2^n}(S'_i)$, from x_i to y_i, that has all of its internal vertices in $M \setminus (W'_i \cup U)$ and uses all but at most $\mu q|W_i|$ of the colours of $S'_i \setminus S_F$ (if $m < i \le t$), of $S_{i,1} \setminus S_F$ (if $i \le m$), or of $S_{i,2} \setminus S_F$ (if $i > t$).

We now reveal the sets $E_1, \dots, E_t$. Chernoff's bound implies that with high probability $|E_i| \ge \frac{1}{8}|S'_i|$ for each i; suppose we are in such an outcome. Then $|E| \ge s/16 \ge N^{1/2}/16$ (since $s = |S| \ge N^{1/2}$ due to $S + S = \mathbb{F}_2^n$). Lemma 6.3 (with $\varepsilon = 1/16$) provides a flexible family $\mathcal{F}$ of $s/2^{62}$ gadgets in E; note that the assumption of this lemma is satisfied since $s/2^{62} \le |E|/2^{58}$.

We now apply Lemma 6.6 with the random set R, the set U, the set E, and this flexible family $\mathcal{F}$. The hypotheses of the lemma are satisfied since $|\mathcal{F}| \ge 2/\gamma \ge |U|$ and $2^{13} \log N \le 2^{12}|\mathcal{F}| \le |E|p^8$. The lemma guarantees that with high probability we have:

Z2 For any vertex y_0, there is a rainbow $\mathcal{F}$-absorbing path in $\mathrm{Cay}_{\mathbb{F}_2^n}(E)$ that ends at y_0 and is otherwise contained in $R \setminus U$.

Finally, apply Lemma 6.7 with the random set T, the set U of already-used vertices, and our flexible family $\mathcal{F}$. To verify the hypotheses of the lemma, note that $|\mathcal{F}| \ge \max\{2^{12}p^{-7}\log N, 128|U|\}$ (by a huge margin). The lemma says that with high probability we have:

Z3 For any set $L \subseteq S$ of at most $(t+m)\mu q 2^{2K} s + 2\alpha s + \gamma s \le |\mathcal{F}|p^7/2^{12}$ colours and any vertex x_{t+m}, there are a sub-family $\mathcal{F}' \subseteq \mathcal{F}$ and a rainbow path in $\mathrm{Cay}_{\mathbb{F}_2^n}(L \cup \bigcup_{F \in \mathcal{F}'} F)$ that starts at x_{t+m}, is otherwise contained in $T \setminus U$, and uses all except possibly one colour from $L \cup \bigcup_{F \in \mathcal{F}'} F$.

Consider an outcome where **Z1**-**Z3** hold. We will deduce the existence of the desired rainbow path. Using **Z2**, we find an $\mathcal{F}$-absorbing rainbow path P_R that ends at some vertex y_0, is otherwise contained in $R \setminus U$, and uses the colours of some subset $S_F \subseteq E$. Next using **Z1** we find rainbow paths $P_{M,1}, \dots, P_{M,t+m}$ with internal vertices in $M \setminus U$, where each P_i goes from y_{i-1} to x_i (with x_{t+m} chosen arbitrarily) and uses all but some set L_i of at most $\mu q|W_i|$ colours from $S_{i,1} \setminus S_F$ (for $i \le m$), from $S'_i \setminus S_F$ (for $m < i \le t$), and from $S_{i,2} \setminus (S_F \cup c(P_{i-t}))$ (for $i > t$). Let $L := J_2 \cup \bigcup_{i=1}^{t+m} L_i$.

[6]It is possible to make "tighter" choices of parameters in some of the cases, but the choices listed here work for all cases and are tight enough for later applications.

We can use the edges (x_i, y_i) to join the paths $P_R, P_{M,1}, \dots, P_{M,t+m}$. The result is a rainbow path that avoids $T \setminus U$ and uses precisely the colours in $S \setminus L$. Note also that

$$|L| \le |J_2| + (t+m)\mu q \cdot \max_i |W_i| \le |J_1| + |X| + (t+m)\mu q 2^{2K} s \le 2\alpha s + \gamma s + (t+m)\mu q 2^{2K} s,$$

so the hypothesis of **Z3** is satisfied. Now **Z3** provides a subfamily $\mathcal{F}' \subseteq \mathcal{F}$ and a rainbow path P_T that starts at x_{t+m}, is otherwise contained in $T \setminus U$, and uses all except possibly one colour from $L \cup \bigcup_{F \in \mathcal{F}'} F$. We use the $\mathcal{F}$-absorbing properties of P_R to remove $\bigcup_{F \in \mathcal{F}'} F$ and pass to a shorter path P_R' using only a subset of vertices of P_R. Now the concatenation of $P_R', P_{M,1}, \dots, P_{M,t+m}, P_T$ gives the desired a rainbow path using all but at most one colour from S. □

8.2. **Expanding case.** In this section we will show how to find a rainbow path of length $|S|-1$ in $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$ in the case where S has a subset E with suitable everywhere-expanding properties. In order to incorporate the final few colours at the end of our argument, we will need to replace absorbing paths with slightly larger structures. Recall the definition of an in-spider from the discussion following Definition 6.2.

Definition 8.8 (Absorbing fork)**.** Let $\mathcal{F}$ be a flexible family of gadgets in $S \subseteq \mathbb{F}_2^n$. A a corresponding *absorbing fork* (P, Q) consists of a (directed) $\mathcal{F}$-absorbing path P and an in-spider Q of $\mathcal{F}$ such that Q is based at the initial endpoint of P and otherwise P, Q are disjoint. We refer to the final endpoint of P as the *final* vertex of the entire absorbing fork.

We now show how to robustly embed absorbing forks in $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$ as long as S is large enough.

Lemma 8.9. *Let $N = 2^n$, and let $E \subseteq S \subseteq \mathbb{F}_2^n$ be such that $|S| \ge 2^{11}|E|$. Let $\mathcal{F}$ be a flexible family of subsets of E. Then $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$ contains an $\mathcal{F}$-absorbing fork (P, Q) with $|P| \le 8|\mathcal{F}| + 1$.*

Proof. Choose an element $c(F) \in S \setminus \bigcup \mathcal{F}$ for each $F \in \mathcal{F}$; since $|S| - |\bigcup \mathcal{F}| \ge |\mathcal{F}|$, we can ensure that these elements $c(F)$ are all distinct. Proposition 6.4 tells us that for each $F \in \mathcal{F}$, there is a rainbow path P_F using precisely the colours in $F \cup c(F)$. It remains to concatenate these paths and join them to an in-spider of $\mathcal{F}$.

To start, let Q be the in-spider of $\mathcal{F}$ based at the vertex 0. We start our path P at the vertex 0 and iteratively extend the path by adding both a single edge (with a colour from S) and one of the rainbow paths P_F from the previous paragraph. At each step, we identify a hitherto-unincorporated gadget F and consider extensions of our current path by a single edge with some colour $x \in S \setminus (\bigcup_{F \in \mathcal{F}} (F \cup \{c(F)\}))$ followed by P_F. Note that $|S \setminus (\bigcup_{F \in \mathcal{F}} (F \cup \{c(F)\}))| \ge |S| - 2|E|$. The structure that we have already built has (crudely) at most $|Q| + 2|E| < 4|E|$ vertices. The new path that we wish to append has at most 8 vertices, so there are fewer than $32|E| < |S| - 2|E|$ choices of x (i.e., translates of this new path) that cause collision, so we choose some $x \in S \setminus E$ that does not cause any collisions. □

We will build our long rainbow path in stages. We will start with an absorbing fork, as guaranteed by Lemma 8.9, which embeds a large family of gadgets. In each subsequent step we will append a new short path to the final vertex of the fork at the cost of "activating" up to two gadgets. In order to iterate this multi-stage procedure, we need to make sure that the final vertex of the fork is not "blocked" by other vertices of the fork (see Proposition 2.5 from our proof overview for a model version of this argument). With this in mind, we say a final vertex v of an $\mathcal{F}$-absorbing fork (P, Q) is *t-extendable* if at least t of the paths of the out-spider of $\mathcal{F}$ based at v are disjoint from $P \cup Q$ (except at v, of course). See Figure 9 for a picture of the proof strategy. It will be notationally convenient to phrase the iteration argument in terms of analysing the properties of a suitable maximal substructure.

We the symbol $\ll$ in the following theorem as we used it in Theorem 8.7, viz., $x \ll y$ means that x is sufficiently small with respect to y.

Theorem 8.10. *Suppose that $1/s \ll 1/K \ll \gamma \ll \alpha \ll 1$. Let $S \subseteq \mathbb{F}_2^n \setminus \{0\}$ be a subset of size $s := |S| \ge \sqrt{N}/2$ (where as usual $N := 2^n$). If S has a $(\gamma/\alpha, K/\gamma)$-everywhere-expanding subset E of size αs, then $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$ has a rainbow path of length $|S|-1$.*

Proof. Lemma 6.3 applied to E provides a flexible family $\mathcal{F}_0$ of gadgets in E with $|\mathcal{F}_0| = \alpha^2|E|/2^{52} = \alpha^3 s/2^{52}$. Let $\beta := \alpha^3/2^{52}$, so $|\mathcal{F}_0| = \beta s$. Now Lemma 8.9 tells us that $\mathrm{Cay}_{\mathbb{F}_2^n}(S)$ contains an $\mathcal{F}_0$-absorbing fork (P_0, Q_0) with $|P_0| \le 9|\mathcal{F}_0| = 9\beta s$; let v_0 denote its final vertex.

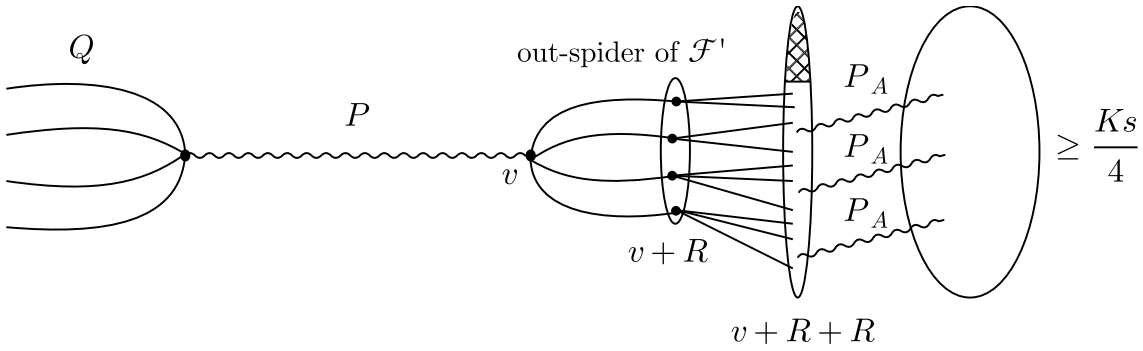

FIGURE 9. Illustration of the proof of Theorem 8.10. We take a maximal $\mathcal{F}$-absorbing fork (P,Q) whose final vertex is suitably extendable. The everywhere-expansion of the set E provides many potential ways to extend P by following a leg of the out-spider, then taking one more step using a colour from a different gadget, and finally appending a translate of a short rainbow path P_A. The maximality hypothesis ensures that in fact P already used nearly all of the colours of S, and we can integrate the remaining colours using one of the legs of Q.

There are at least $s-9\beta s$ colours not appearing in P_0. Let N_0 denote the set of vertices in $\mathbb{F}_2^n \setminus (P_0 \cup Q_0)$ that are reachable from v_0 via single edge with one of these colours. We have $|N_0| \geq s - 24\beta s \geq s/2$. Consider an auxiliary bipartite graph whose left side is the vertex set of $P_0 \cup Q_0$ and whose right side is N_0, with an edge between $a \in P_0 \cup Q_0$ and $b \in N_0$ if the out-spider of $\mathcal{F}_0$ based at b contains a. Equivalently, a is adjacent to b if the in-spider of $\mathcal{F}_0$ based at a contains b, so the degree of each vertex on the left is at most $5\beta s$. Recall that the left side has at most $15\beta s$ vertices, so the graph has at most $75\beta^2 s^2$ edges in total. At the same time, $|N_0| \geq s/2$, so there is some $u \in N_0$ with degree at most $150\beta^2 s$ in the auxiliary bipartite graph. Extending P_0 to such a vertex gives a new absorbing fork whose final vertex u is γs-extendable (since $|\mathcal{F}_0| - 150\beta^2 s \geq \gamma s$).

Now, let (P,Q) be a maximal-size $\mathcal{F}$-absorbing fork such that its final vertex v is γs-extendable, $\mathcal{F} \subseteq \mathcal{F}_0$, and $\frac{1}{2}|\mathcal{F}_0 \setminus \mathcal{F}| \leq |P| - (1-\gamma)\min\{|P|, s - 8/\gamma\}$. Note that (P_0, Q_0) is such a fork for $\mathcal{F}_0$, so this is well-defined. Let $A \subseteq S$ denote the subset of *available* colours not used in P. We claim that $\mathrm{Cay}_{\mathbb{F}_2^n}(A)$ has a rainbow path P_A using precisely $4/\gamma$ colours if $|A| \geq 8/\gamma$, and using $\min\{|A|, 7\}$ colours from A otherwise. If $|A| \geq 14$ this follows from the standard greedy argument which always allows us to find a path of length at least $|A|/2$ (see [5, Observation 2.2]). If $|A| < 14$, then we may simply use the fact from [1] that all sets of size at most 7 have valid orderings.

Since v is γs-extendable, there is some $\mathcal{F}' \subseteq \mathcal{F}$ of size γs such that the out-spider of $\mathcal{F}'$ based at v is disjoint from the vertices of $P \cup Q$ (other than v). Let R denote the set of last elements of gadgets in $\mathcal{F}'$ (i.e., if we write each gadget as $F = \{f_1, \ldots, f_{|F|}\}$, then $R = \{f_{|F|} : F \in \mathcal{F}'\}$). Note that since each gadget is 0-sum, the set of leaves of this spider is precisely $v + R$. The subset $R \subseteq E$ and has size γs, so our everywhere-expanding assumption on E ensures that $|R+R| \geq Ks$. This means that if we take a second step with an edge with colour in R, then we can reach at least Ks vertices in $v + R + R$. There are at least $Ks - 1 - s - 5\beta s - 5\gamma s \geq Ks/2$ such vertices which are not in $P \cup Q$ (which contains at most $1 + s + 4\beta s$ vertices) or in the out-spider of $\mathcal{F}'$ based at v (which contains at most $5\beta s$ vertices). Further, as P_A is a path using at most $8/\gamma$ colours, at most $(8/\gamma) \cdot (1 + 5\beta + 5\gamma)s \leq Ks/4$ translations of the path P_A starting at the vertices in $v + R + R$ can intersect $P \cup Q$ or the out-spider of $\mathcal{F}'$ based at v; here we used $K \gg \alpha, \gamma$. So at least $Ks/4$ such translates of P_A will be disjoint from both $P \cup Q$ and the out-spider of $\mathcal{F}'$ based at v.

Let $z = v + f_{|F_1|} + f_{|F_2|} \in v + R + R$, with $F_1, F_2 \in \mathcal{F}'$, be such a *good* vertex, in the sense that if we start at z and follow the colours of P_A in order, then we obtain a rainbow path whose vertices are disjoint from $P \cup Q$ and the out-spider of $\mathcal{F}'$ based at v. Now we can extend the path P in our current $\mathcal{F}$-absorbing fork (P,Q) adding the edges with the colours from $F_1 \setminus \{f_{|F_1|}\}$, then the edge with the colour $f_{|F_2|}$ (to reach z), and finally the translation of P_A. This produces a genuine path since $F_1 \in \mathcal{F}'$ and the out-spider of $\mathcal{F}'$ based at v is disjoint from $P \cup Q$, while the choice of z guarantees that z and the translation of P_A do not cause collisions. This procedure has produced a new $\mathcal{F} \setminus \{F_1, F_2\}$-absorbing fork by activating the gadgets F_1, F_2 (we replace P by $P - F_1 - F_2$ to maintain rainbowness). Note that we reintegrated all except the last colour of F_1 into our extended rainbow path, but we reintegrated only the first colour of F_2. Hence, as our gadgets have size at most

6, our new absorbing path P' has size

$$|P'| \geq |P| - 6 + 4/\gamma \geq \frac{1}{2\gamma}|\mathcal{F}_0 \setminus \mathcal{F}| + \frac{2}{\gamma} \geq \frac{1}{2\gamma}|\mathcal{F}_0 \setminus (\mathcal{F} \setminus \{F_1, F_2\})|$$

if $|A| = s - |P| \geq 8/\gamma$. It still has size

$$|P'| \geq |P| + 1 \geq \frac{1}{2}|\mathcal{F}_0 \setminus \mathcal{F}| + (1-\gamma)(s - 8/\gamma) + 1 = \frac{1}{2}|\mathcal{F}_0 \setminus (\mathcal{F} \setminus \{F_1, F_2\})| + (1-\gamma)(s-8/\gamma)$$

if $7 \leq |A| = s - |P| < 8/\gamma$. The only remaining case is where P' is missing exactly 6 colours, including the last 5 colours of F_2.

In the final case, we can just append the leg of the in-spider Q corresponding to the 5 leftover colours from F_2 at the beginning of the path P'; this produces the desired rainbow path of length $|S| - 1$. We will be done if we show that the first two cases are actually impossible by the maximality assumption on (P, Q). For this, we need to verify that at least one of the $Ks/4$ possible new final vertices is γs-extendable in its new fork. We can repeat the argument from above with the auxiliary bipartite graph; this time there are at most $s + 5\beta s + 5\gamma s \leq 2s$ vertices on the left side (consisting of $P \cup Q$ together with the vertices of the out-spider of $\mathcal{F}'$ based at v), each sending at most αs edges[7] to the right side (consisting of the final vertices of our potential extended forks, namely, the translations of good vertices by the path P_A), which has size at least $Ks/4$. So some vertex v' on the right has degree at most $\frac{8\alpha}{K}s \leq \gamma s$. At most γs of the leaves of the out-spider of $\mathcal{F}$ based at v' are blocked by vertices in $P \cup Q$ or the out-spider of $\mathcal{F}'$ based at v. Besides these, at most $8/\gamma$ additional leaves of the out-spider of $\mathcal{F}$ based at v' are blocked by the translate of P_A that we used to reach v'. In total, our new endpoint v' is at least $|\mathcal{F}| - \gamma s - 8/\gamma \geq \gamma s$-extendable, since

$$|\mathcal{F}| = |\mathcal{F}_0| - |\mathcal{F}_0 \setminus \mathcal{F}| \geq \beta s - 2\max\{\gamma s, s - (1-\gamma)(s-8/\gamma)\} \geq \alpha^3 s/2^{52} - 2\gamma s \geq 2\gamma s + 8/\gamma$$

(here using $\beta = \alpha^3/2^{52}$ and the assumption that $\gamma \ll \alpha$). □

9. General dense case

In this section we prove Theorem 1.4, which we restate for convenience.

Theorem 1.4. *There is an absolute constant $c > 0$ such that for any finite (possibly nonabelian) group G, every subset $S \subseteq G \setminus \{\mathrm{id}\}$ of size at least $|G|^{1-c}$ admits a valid ordering.*

Due to a lack of 0-sum subsets, we will need to work with a different type of gadget.

Definition 9.1 (g-pair)**.** Let G be a group, and let $g \in G$. A *g-pair* is a pair of distinct elements $a, b \in G$ such that $ab = g$. A *family of g-pairs* in a subset $S \subseteq G$ is a collection of disjoint g-pairs contained in S; we say that the number of such pairs is the *size* of the family.

The following easy lemma lets us find a large family of g-pairs in any large subset of a finite group.

Lemma 9.2. *Let G be a group, and let $S \subseteq G$. If $|S| \geq |G|^{1-\varepsilon} \geq 2$, then for some $g \in G$ there exists a family of g-pairs in S of size at least $|G|^{1-2\varepsilon}/6$.*

Proof. There are $|S|(|S|-1)$ ordered pairs of distinct elements $a, b \in S$. There are at most $|G|$ possible values for the product ab, so by the pigeonhole principle there is some $g \in G$ such that S contains a collection of at least $|S|(|S|-1)/|G| \geq |G|^{1-2\varepsilon}/2$ g-pairs. Each such g-pair (a, b) can have elements in common with at most two other g-pairs (a', b'). So we can find a sub-collection consisting of at least $|G|^{1-2\varepsilon}/6$ disjoint g-pairs. □

Consider a vertex v of $\mathrm{Cay}_G(S)$ and a family of g-pairs of S. For each g-pair (a, b), the two-edge path $v \to va \to vab = vg$ using the colours a, b terminates at the vertex vg. The union of these paths over all of the g-pairs forms a *theta-graph*. We will later stitch together several such theta-graphs along a path-like structure which we call a *waveform*; see Figure 10 for an illustration (formal definitions forthcoming).

We have great flexibility in taking a path through a waveform, since for each theta-graph we may (independently) choose which length-two path to include. The downside here, compared to the zero-sum gadget approach from the previous sections, is that the remaining pairs in each theta-graph remain unused. As each colour pair must eventually appear somewhere, we need to carefully specify which colour pairs occur in which theta-graphs. This procedure will be simplest to describe using the language of a certain auxiliary bipartite graph on parts (X, Z),

[7] We note that the edges in our auxiliary bipartite graph are defined by the *full* $\mathcal{F}$ out-spiders and not just the $\mathcal{F}'$ ones.

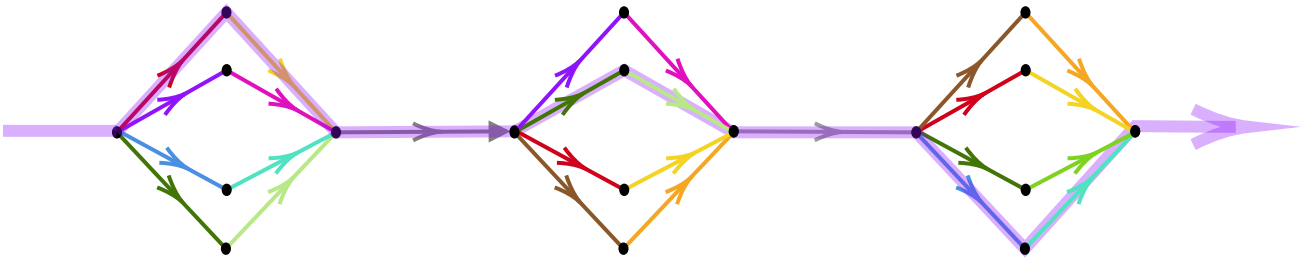

FIGURE 10. A waveform with three theta-graphs. Notice that some colour pairs repeat across the theta-graphs. A rainbow path omitting the brown-orange colour pair is highlighted.

where vertices of Z correspond to theta-graphs and vertices of X represent g-pairs as given by Lemma 9.2. We wish to build a waveform where the theta-graph corresponding to the vertex $z \in Z$ contains precisely the g-pairs in its neighbourhood $N(z) \subseteq X$. No theta-graph should be too big (as then finding absorbing structures in random subgraphs would be problematic), so our bipartite graph should have small maximum degree. Furthermore, after we later (during an absorption step) saturate some fraction of the g-pairs elsewhere, we will wish to integrate exactly the remaining (unused) g-pairs into the path coming from our waveform. In the language of the bipartite graph, we wish to find perfect matchings between $X \setminus X'$ and Z for a wide variety of choices of subsets X'.

The following proposition from [35] shows that bipartite graphs with such properties exist.

Proposition 9.3 ([35], Lemma 10.7)**.** *Let $\ell \leq k$ be positive integers. There exists a bipartite graph B with bipartition $(X \cup Y, Z)$, where X and Y are disjoint, $|X| = k + \ell, |Y| = 2k$ and $|Z| = 3k$, such that:*

(1) The maximum degree is at most 40.
(2) For any subset $X' \subset X$ with $|X'| = k$, there is a perfect matching between $X' \cup Y$ and Z.

We remark that [35, Lemma 10.7] states this result only with $\ell = k$, but one can obtain the more general version above by simply deleting an appropriate number of vertices from X. The construction from [35] is essentially a union of 40 perfect matchings, sampled uniformly at random, and it is not hard to show that the desired properties hold with positive probability.

We can now precisely describe the type of absorbing structure we will utilise.

Definition 9.4 (Absorbing family)**.** Let G be a group, $g \in G$, and $S \subseteq G$. Let $\mathcal{P}_{\text{flex}} \subseteq \mathcal{P}$ be families of g-pairs in S. A *$(\mathcal{P}_{\text{flex}}, \ell)$-absorbing family in $\mathcal{P}$* consists of a sequence $\mathcal{P}_1, \dots, \mathcal{P}_{|\mathcal{P}|-\ell} \subseteq \mathcal{P}$ of sub-families such that each $|\mathcal{P}_i| \leq 40$ and for every $\mathcal{P}' \subseteq \mathcal{P}_{\text{flex}}$ with $|\mathcal{P}'| = \ell$ there exists a system of distinct representatives $p_i \in \mathcal{P}_i \setminus \mathcal{P}'$.

The following is essentially an immediate consequence of Proposition 9.3.

Corollary 9.5. *Let G be a group. Let $g \in G$, and let $S \subseteq G$, and let $\mathcal{P}$ be a family of g-pairs in S of size $3t + \ell$, with $\ell \leq t$. Then for some sub-family $\mathcal{P}_{\text{flex}} \subseteq \mathcal{P}$ of size $t + \ell$ there is a $(\mathcal{P}_{\text{flex}}, \ell)$-absorbing family in $\mathcal{P}$.*

Proof. Consider the bipartite graph B from Proposition 9.3. We identify $X \cup Y$ with $\mathcal{P}$, set $\mathcal{P}_{\text{flex}} := X$, and define $\mathcal{P}_i$ to be the set of neighbours of the i-th vertex in Z. Note that each $|\mathcal{P}_i| \leq 40$ since B has maximum degree at most 40. The matching property of B precisely translates to $\mathcal{P}_1, \dots, \mathcal{P}_{3t}, \mathcal{P}_{\text{flex}}$ having the desired property about distinct representatives. □

As we alluded to above, the building blocks for our absorbing structures are theta-graphs.

Definition 9.6 (Theta-graph)**.** Let $\mathcal{P}$ be a family of g-pairs in a group G. The *theta-graph* of $\mathcal{P}$ anchored at v, denoted $T(v, \mathcal{P})$, is the union of the $|\mathcal{P}|$ length-two paths obtained by starting at v and following the edges of colours a, b for each g-pair $(a, b) \in \mathcal{P}$.

Notice that each of the paths in $T(v, \mathcal{P})$ terminates at the vertex $vab = vg$.

Definition 9.7 (Waveform)**.** Let G be a group, $g \in G$, $S \subseteq G$, and let $\mathcal{P}_1, \dots, \mathcal{P}_t$ be families of g-pairs in S. A corresponding *waveform* starting at a vertex $u \in G$ is a subgraph of $\text{Cay}_G(S)$ consisting of u and $T(v_1, \mathcal{P}_1), \dots, T(v_t, \mathcal{P}_t)$, together with the edges $(u, v_1), (v_1 g, v_2), (v_2 g, v_3), \dots, (v_{t-1} g, v_t)$, such that:

- the theta-graphs $T(v_i, \mathcal{P}_i)$ are disjoint from one another and from u;

- the edges $(u, v_1), (v_1g, v_2), (v_2g, v_3), \dots, (v_{t-1}g, v_t)$ are all edges of $\mathrm{Cay}_G(S)$ of distinct colours, and none of these colours appear in the $T(v_i, \mathcal{P}_i)$'s.

A waveform of a $(\mathcal{P}_{\text{flex}}, \ell)$-absorbing family in $\mathcal{P}$ has the key property that by choosing a single length-two path from each theta-graph, we can construct a rainbow path which uses all the pairs of colours in $\mathcal{P}$ apart from any desired subcollection of ℓ pairs of $\mathcal{P}_{\text{flex}}$; this allows us to "absorb" such subfamilies. We will refer to the specification of such a subpath of the waveform as *collapsing* the waveform. We now present a waveform analogue of Lemma 6.6, with a similar proof.

Lemma 9.8. *Let G be a group, $g \in G$, and $E \subseteq G$. Let $\mathcal{P}_1, \dots, \mathcal{P}_t$ be families of g-pairs in E with each $|\mathcal{P}_i| \le 40$. Let $0 < p \le 1$ be such that $|E|p^{42}/2^{18} \ge \max\{t, \log|G|\}$. Let R be a p-random subset of G. Then with high probability, for every vertex $u \in G$ we can find a waveform in $\mathrm{Cay}_G(E)$ corresponding to $\mathcal{P}_1, \dots, \mathcal{P}_t$ which starts at the vertex u and is otherwise contained in R.*

Proof. Let $N := |G|$. For each vertex $v \in G$ and family $\mathcal{P}_i$, let $E_{v,i}$ be the event that we can find a collection of at least $50t$ elements $e \in E \setminus \bigcup_{i=1}^t \mathcal{P}_i$ whose corresponding theta-graphs $T(ve, \mathcal{P}_i)$ are all vertex-disjoint and contained in R. Notice that each theta-graph $T(ve, \mathcal{P}_i)$ intersects at most 42^2 other such theta-graphs, since each theta-graph has at most 42 vertices and the translate of the other theta-graph is determined by the relative positions of the intersection point on the two paths. Thus we can find a collection of at least $(|E| - 80t)/(42^2 + 1) \ge |E|/2^{11}$ vertex-disjoint $T(ve, \mathcal{P}_i)$'s. Each survives in R with probability at least p^{42}, and these events are independent. Hence the number of surviving theta-graphs stochastically dominates $\mathrm{Bin}(|E|/2^{11}, p^{42})$, and by a Chernoff bound at least

$$|E|p^{42}/2^{12} \ge 50t$$

survive with probability at least $1 - \exp(-|E|p^{42}/2^{14}) \ge 1 - 1/N^3$. Thus $\mathbb{P}[E_{v,i}] \ge 1 - 1/N^3$, and by a union bound we conclude that with probability at least $1 - 1/N$ all of the events $E_{v,i}$ occur.

Suppose we are in such an outcome. We find our waveform by incorporating theta-graphs one at a time, as in the proof above. We start our waveform W at the vertex u and iteratively add on graphs of the form $T(ve, \mathcal{P}_i)$, where v is the current endpoint of W. At each step, we identify a hitherto-unincorporated $\mathcal{P}_i$ and consider the $50t$ theta-graphs $T(ve, \mathcal{P}_i)$ identified in the previous paragraph. Of these, at least $49t$ correspond to colours e that have not yet been used. Since $|W| < 42t$, there are at least $7t$ theta-graphs $T(ve, \mathcal{P}_i)$ that remain disjoint from W; we choose one such theta-graph and add it to the end of our waveform W. □

We will also need a slightly tweaked version of the absorbing lemma (Lemma 6.7).

Lemma 9.9. *Let $p \in (0, 1]$, let $\mathcal{P}_{\text{flex}}$ be a family of g-pairs in a group G with $|\mathcal{P}_{\text{flex}}| \ge t + \ell \ge 2^9 p^{-3} \log|G|$, and let T be a p-random subset of G. Then with high probability, the following holds for every $L \subseteq G \setminus \bigcup \mathcal{P}_{\text{flex}}$ of size $|L| \le \ell < tp^3/80$ and every vertex $v \in G$: There exist a subfamily $\mathcal{P}' \subseteq \mathcal{P}_{\text{flex}}$ of size ℓ and a rainbow path in $\mathrm{Cay}_{\mathbb{F}_2^n}(L \cup \bigcup \mathcal{P}')$ that starts at v, is otherwise contained in T, and uses all except possibly one colour from $L \cup \bigcup \mathcal{P}'$.*

Proof. Consider a pair of distinct colours $a, b \in G$. For each $(a_i, b_i) \in \mathcal{P}_{\text{flex}}$, consider the length-three path that starts at the vertex id and then traverses the edges of colours a_i, a, b. Our first goal is to construct a family $\mathcal{P}_{a,b}$ of at least $t/10$ such paths that are vertex-disjoint (except at the shared vertex id). Note that each such path can intersect at most nine other paths, so we can find a vertex-disjoint collection of $|\mathcal{P}_{\text{flex}}|/10 \ge t/10$ of them.

For each vertex $u \in G$ pair of distinct colours $a, b \in S$, let $E_{u,a,b}$ be the event that we can find a collection of more than 4ℓ g-pairs $P \in \mathcal{P}_{a,b}$ such that the (left-)translates by u of the corresponding length-three paths are all contained in T (except for possibly u). The number of surviving paths stochastically dominates $\mathrm{Bin}(t/10, p^3)$, so by a Chernoff bound at least $tp^3/20 > 4\ell$ survive with probability at least $1 - \exp(tp^3/80) \ge 1 - 1/N^4$. Thus $\mathbb{P}[E_{u,a,b}] \ge 1 - 1/N^4$. A union bound over all u, a, b ensures that with probability at least $1 - 1/N$ all of the events $E_{u,a,b}$ occur. Suppose we are in such an outcome.

We will construct a sequence of sets $L = L_0, L_1, \dots, L_{|L|-1}$ of sizes $|L_i| = |L| - i$ and a sequence of directed rainbow paths $v = P_0 \subset P_1 \subset \cdots \subset P_{|L|-1}$ of sizes $|P_i| = 3i+1$, as follows. Suppose we have already constructed L_i, P_i, and suppose that $|L_i| \ge 2$. Pick some distinct $a, b \in L_i$. There is a collection of more than $4\ell \ge 4|L|$ vertex-disjoint rainbow paths in T, where each starts at the endpoint of P_i and then traverses the edges with colours a_j, a, b for some $(a_j, b_j) \in \mathcal{P}_{\text{flex}}$. Some such path uses a colour a_j from a new pair and is vertex-disjoint from P_i; we append it to P_i to obtain P_{i+1}. To obtain L_{i+1} from L_i, we remove a, b and add b_j (so indeed

$|L_{i+1}| = |L_i| - 1$). Thus the process can indeed run $|L| - 1$ steps, provided we can always find a suitable short path to extend by. Note also that by construction P_i uses colours from at most i pairs in $\mathcal{P}_{\text{flex}}$ (besides the colours from L) and has length $3i$. So when constructing P_{i+1} at most i pairs are already used, and at most $|P_i| - 1 = 3i$ of our short paths can intersect P_i, so we indeed can always choose a short path to extend P_i by into P_{i+1}.

At the end of this process we used at most $|L| - 1 \leq \ell$ pairs from $\mathcal{P}_{\text{flex}}$ and embedded all the colours from these pairs as well as from L except possibly one. To ensure we use up exactly ℓ pairs, we continue the process to find $L = L_{|L|}, \dots, L_\ell$ each of size one and $P_\ell \supset \dots \supset P_{|L|} \supset P_{|L|-1}$ such that for each $i \geq |L|$ we have $|P_i| = |P_{i-1}| + 2$, and $P_i \setminus P_{i-1}$ uses the colour in L_{i-1} together with a new colour from $\bigcup \mathcal{P}_{\text{flex}}$. To see that we can do this, suppose that we are at stage $i - 1 \in [|L| - 1, \ell)$ and that the current path P_{i-1} ends at the vertex v. Then we pick a to be the (unique) colour in L_{i-1} and $b \neq a$ to be an arbitrary other colour. As the event $E_{v,a,b}$ holds, there is some new $\mathcal{P}_i = (a_i, b_i)$ (not yet used on the path) such that we may extend our current path by appending the edges going from v to va_i to va_ia (and we simply do not use the colour b edge in the path guaranteed by the above process), in order to construct P_i (and replace a with b_i in L_i). Since we have more than 4ℓ choices, we can continue until we have used up exactly ℓ pairs from $\mathcal{P}_{\text{flex}}$, as desired. □

As in Section 7 we start by establishing the main result in the regime $|S| \geq \frac{3}{4}N$, which is slightly different due to tighter space constraints.

Theorem 9.10. *Let G be a group of order N, where N is sufficiently large. If $S \subseteq G$ is a subset of size $|S| \geq \frac{3}{4}N$, then $\mathrm{Cay}_G(S)$ has a rainbow path of length $|S| - 1$.*

Proof. Set $\gamma := 2^{-20}$. If $|S| \geq N - N^{1-\gamma}$, then we are done by Theorem 7.1 so let us assume $|S| \leq N - N^{1-\gamma}$. Let us also set $p = N^{-2\gamma}/2$.

Let E be a $\frac{1}{8}$-random subset of S. Let us also partition $\mathbb{F}_2^n$ into three sets $R \sqcup M \sqcup T$ by independently assigning each vertex to R, M, T with probabilities $p, 1 - 2p, p$, respectively.

We apply Lemma 5.7 with $S = S, J = \emptyset, M = M, S' = E$ and the parameters

$$\varepsilon = \frac{3}{4}, \quad \zeta = \frac{1}{4}, \quad q = 1 - 2p, \quad q' = \frac{1}{8}, \quad \mu = N^{-90\gamma}/2^{76}.$$

Since $1 - 2p = q \geq (1 + \mu)(1 - N^{-\gamma})$, and $q' \leq 1 - \mu q/4$, we may indeed apply the lemma. Thus with high probability we have:

D1 For any $S_F \subseteq E$ and any two vertices $v, u \in \mathbb{F}_2^n$, we can find a rainbow path from v to u in $\mathrm{Cay}_{\mathbb{F}_2^n}(S \setminus S_F)$, using all but at most μq colours from $S \setminus S_F$, such that all of the internal vertices of the path lie in M.

Let us now reveal the random subset E. Chernoff's bound guarantees that with high probability $|E| \geq N/16$ (as before, if this is not the case then we declare failure and do not apply the following lemmas). Let

$$t := N^{1-84\gamma}/2^{66}, \quad \ell := tp^3/2^7 = N^{1-90\gamma}/2^{76}.$$

Using Lemma 9.2 we can find some $g \in G$ and a family $\mathcal{P}$ of g-pairs in E of size $3t + \ell \leq N/2^{11}$. Using Corollary 9.5 we can find a $(\mathcal{P}_{\text{flex}}, \ell)$-absorbing family $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_{3t}$ in $\mathcal{P}$ for some $\mathcal{P}_{\text{flex}} \subseteq \mathcal{P}$ of size $t + \ell$.

We apply Lemma 9.8 to $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_{3t}$ with the random set R; note that this lemma applies since $|E|p^{42}/2^{18} \geq 3t \geq \log N$. Thus with high probability we have:

D2 For any vertex v, there is a waveform corresponding to $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_{3t}$ that ends at v and is otherwise contained in R.

Finally, we apply Lemma 9.9 to $\mathcal{P}_{\text{flex}}$ with random set T; we can do so since $t + \ell \geq 2^9 p^{-3} \log N$ and $\ell < tp^3/80$. Thus with high probability we have:

D3 For any $L \subseteq S$ of size $|L| \leq \ell$ and any vertex u, there is a sub-family $\mathcal{P}' \subseteq \mathcal{P}_{\text{flex}}$ of size ℓ such that $\mathrm{Cay}_G(L \cup \bigcup \mathcal{P}')$ contains a rainbow path that starts at u, is otherwise contained in T, and uses all but possibly one colour from $L \cup \bigcup \mathcal{P}'$.

Let us fix an outcome in which all three of the above properties hold. Fix some distinct vertices $v \in M$, $u \in T$. **D2** gives us a waveform corresponding to $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_{3t}$ that is completely contained in R and ends at v. This waveform uses a total of $6t + \ell$ colours: It uses $3t + \ell$ colours within the theta-graphs (we have $\mathcal{P}_1 \cup \dots \cup \mathcal{P}_{3t} = \mathcal{P}$ by the properties of an absorbing family) and an additional $3t$ colours connecting the theta-graphs to one

another (and to the final vertex v). Let $S_F \subseteq E$ denote this set of $6t+\ell$ colours. Now **D1** produces a rainbow path P_M, starting at v and ending u and otherwise contained in M, which saturates all but some set L of at most $\mu q N \leq \ell$ colours from $S \setminus S_F$. Finally, by **D3** we can find $\mathcal{P}' \subseteq \mathcal{P}_{\text{flex}}$ of size precisely ℓ and a rainbow path P_T contained in T which uses all except possibly one colour from $L \cup \bigcup_{(a_i,b_i)\in\mathcal{P}'}\{a_i,b_i\}$. By the absorbing property, there is a system of distinct representatives for $\{\mathcal{P}_i \setminus \mathcal{P}' : i \in [3t]\}$. Now we can follow the length-two paths in our waveform corresponding to these pairs in order to use up the colours in $S_F \setminus \bigcup_{(a_i,b_i)\in\mathcal{P}'}\{a_i,b_i\}$. The concatenation of these three paths uses all except possibly one of the colours of S, as desired. □

We are now ready to prove the main theorem of the section, Theorem 1.4, in the following more precise form.

Theorem 9.11. *Set $\gamma := 2^{-20}$. Let G be a group of order N, where N is sufficiently large. Then for any subset $S \subseteq G$ of size $|S| \geq N^{1-\gamma}$, the graph $\mathrm{Cay}_G(S)$ has a rainbow path of length $|S|-1$.*

Proof. Set $\varepsilon := N^{-\gamma}$, so that $|S| \geq \varepsilon N$. We may assume that $\varepsilon \leq 2^{-20}$. We apply our regularity result Corollary 4.1 (with $\sigma = |S|/|G| \geq \varepsilon$) to find a subgroup H of G such that $|S\cap H| \geq (1-\varepsilon)|S|$ and $\mathrm{Cay}_H(S\cap H)$ has no $\varepsilon^4/1000$-sparse cuts. Let $S_0 := S \cap H$ and $J := S \setminus H$. We now define two subsets S_1, S_2 of S_0; how we do so depends on the proportion of H occupied by S_0.

Case 1. $|S_0| \leq \frac{3}{4}|H|$.
In this case we set $S_1 := S_0$ and $S_2 := \emptyset$.

Case 2. $|S_0| \geq \frac{3}{4}|H|$.
Note that if $S \setminus H = \emptyset$, then we are done by Theorem 9.10, so we may assume that $J = S \setminus H \neq \emptyset$. Take random subsets S_1, S_2 of S_0 by assigning each element of S_0 independently to S_1 with probability $\frac{1}{4}$, to S_2 with probability $\frac{1}{4}$, and to both S_1, S_2 with probability $\frac{1}{2}$. So S_1, S_2 are $\frac{3}{4}$-random subsets of S_0, and $S_1 \cup S_2 = S_0$. Although S_1, S_2 are not independent, it is true that after we reveal S_1 (respectively, S_2), the intersection $S_1 \cap S_2$ is a $\frac{2}{3}$-random subset of S_1 (respectively, S_2).

Let S' and E be disjoint $\frac{1}{4}$-random subsets of S_0. Set $p := 1/32$, and let $A \sqcup R \sqcup M \sqcup T$ be a random partition of G where each vertex is (independently) assigned to A, R, M, T with probabilities $p, p, 1-3p, p$, respectively.

We now reveal S_1[8]. We have

$$\frac{5}{8}|S_0| \leq |S_1| \leq \frac{5}{6}|H| \tag{15}$$

deterministically in Case 1, and with high probability by Chernoff in Case 2; suppose that this inequality holds. Fix a coset sH of H and apply Lemma 5.7 with $G = sH, S = S_1, J = \emptyset, M = M, S' = S' \cup E \cup (S_1 \cap S_2)$ and the parameters

$$\zeta = \varepsilon^4/1000, \quad q = 1-3p, \quad \mu = N^{-2\gamma}/2^{38}, \quad q' = 5/6,$$

and we replace ε in Lemma 5.7 with $\varepsilon/2$. Let us check that the assumptions of the lemma are satisfied. By (15) we have $|S_1| \geq \frac{5}{8}|S_0| \geq \frac{5}{8}(1-\varepsilon)|S| \geq \frac{\varepsilon}{2}N \geq \frac{\varepsilon}{2}|H|$. We also need $\mathrm{Cay}_H(S_1)$ to have no ζ-sparse cuts: This holds in Case 1, since $S_1 = S_0$ has no ζ-sparse cuts by construction; and it holds in Case 2 by (15) which implies $\mathrm{Cay}_H(S_1)$ has no $\frac{1}{8}$-sparse cuts. Further, we have $q \geq (1+\mu)|S_1|/|H|$, since $|S_1| \leq \frac{5}{6}|H|$, by (15). Finally, we have $q' \leq 1 - \mu q/4$ with plenty of room to spare. Hence, Lemma 5.7 tells us that with probability at least $1 - 7/|H|$ we have:

U1 For any $S_F \subseteq S' \cup E \cup (S_1 \cap S_2)$ and any vertices $w \in sH$, there is a rainbow path in $\mathrm{Cay}_{sH}(S_1)$ starting at w, with all other vertices in M, such that the path uses all but at most $\mu q|H|$ of the colours of $S_1 \setminus S_F$.[9]

Moreover, since there are only $\frac{|G|}{|H|} \leq \frac{|G|}{(1-\varepsilon)|S|} \leq 2N^{\gamma} = o(|H|)$ left-cosets of H, with high probability the conclusion **U1** holds simultaneously for all cosets sH.

If we are in Case 2, we will need a second application of Lemma 5.7, this time with $S = S_2$ and with all of the other sets and parameters the same as in our first application (this time revealing S_2 but leaving S_1 as unrevealed so that $S_1 \cap S_2$ is a genuinely $\frac{2}{3}$-random subset of S_2). By the same reasoning, we conclude that with high probability for all cosets sH we have:

U2 For any $S_F \subseteq S' \cup E \cup (S_1 \cap S_2)$ and any vertex $z \in sH$, there is a rainbow path in $\mathrm{Cay}_{sH}(S_2)$ starting at z, with all other vertices in M, such that the path uses all but at most $\mu q|H|$ of the colours of $S_2 \setminus S_F$.

[8]Note that $S_1 \cap S_2$ remains a $\frac{2}{3}$-random subset of S_1.
[9]Lemma 5.7 also gives us the freedom to specify the other endpoint of the path, but we will not need to do so in this proof.

We know that $|E| \geq |S_0|/8 \geq |S|/16 \geq p^{-2}\max\{40|J|, 96\log|G|\}$ with high probability, so Lemma 6.1 (applied with $G = G$, $E = E \cup J$ and $J = J$) tells us that with high probability we have:

U3 For any vertex $v \in G$, there is a rainbow path in $\mathrm{Cay}_G(E \cup J)$, starting from u and otherwise contained in A, that uses all of the colours of J.

We now reveal S'. By a Chernoff bound, with high probability we have $|S'| \geq |S_0|/8 \geq |S|/16$; we henceforth assume that we are in such an outcome. The next two lemmas use only the randomness in the subsets R, T (respectively); the key point is that these are independent of S'. Set

$$t := N^{1-2\gamma}/2^{14}, \quad \ell := tp^3/81.$$

Using Lemma 9.2, we can find an element $g \in H$ and a family $\mathcal{P}$ of g-pairs in S' of size $3t + \ell \leq N^{1-2\gamma}/2^{12}$. Corollary 9.5 provides a $(\mathcal{P}_{\text{flex}}, \ell)$-absorbing family $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_{3t}$ in $\mathcal{P}$ with some $\mathcal{P}_{\text{flex}} \subseteq \mathcal{P}$ of size $t + \ell$.

We next apply Lemma 9.8 with $G = G, g = g, E = S'$, the p-random set R, and the families of g-pairs $\mathcal{P}_1, \dots, \mathcal{P}_{3t}$. The hypotheses of the lemma are satisfied since $|S'|p^{42}/2^{18} \geq \max\{3t, \log|G|\}$. Thus with high probability we have:

U4 For any vertex $u \in G$, there is a waveform in $\mathrm{Cay}_G(S')$ corresponding to $\mathcal{P}_1, \dots, \mathcal{P}_{3t}$ that starts at u and is otherwise contained in R.

Finally, we apply Lemma 9.9 with $G = G, g = g$, a p-random subset T and the family $\mathcal{P}_{\text{flex}}$. The hypotheses of the lemma are satisfied since $|\mathcal{P}_{\text{flex}}| = t + \ell \geq 2^9 p^{-3}\log N$. Thus with high probability we have:

U5 For any L of size $|L| \leq \ell < tp^3/80$ and any vertex $y \in G$, there is a subfamily $\mathcal{P}' \subseteq \mathcal{P}_{\text{flex}}$ with exactly ℓ pairs such that $\mathrm{Cay}_G(L \cup \bigcup \mathcal{P}')$ contains a rainbow path that starts at y, is otherwise contained in T, and uses all except possibly one of the colours from $L \cup \bigcup \mathcal{P}'$.

Consider an outcome where all of **U1**–**U5** occur. Fix any vertex $z \in M \cap H$.

If we are in Case 1, then let $P_{M,1}$ be the empty path. If we are in Case 2, then use **U2** to find a rainbow path $P_{M,1}$, starting at z, that uses all of the colours from $S_2 \setminus (S' \cup E)$ except for some some subset L_2 of size at most $\mu q|H|$. Note that since $S_2 \subseteq S_0 \subseteq H$ and $z \in H$, the path $P_{M,1}$ is completely contained in H. Let $v \in H$ be its other endpoint, and let S''_F be the set of colours appearing in $P_{M,1}$.

Next, let P_A be a minimal-length rainbow path that starts at v (if we are in Case 1, we instead let it start from an arbitrary vertex $v \in A \cap H$), is otherwise contained in A, and uses all of the colours from J and some subset of the colours from E. Such a path P_A exists by **U3**. Let u be the endpoint of P_A. If $J \neq \emptyset$, then the minimality of P_A guarantees that the last edge of P_A uses some colour $j^* \in J = S \setminus H$. If in addition $u \in H$, then we delete this last edge from P_A so that its endpoint u now lies in a proper coset sH. Let S'_F denote the set of colours from E appearing in P_A. Notice that P_A is the empty path (so $u = v$) if $J = \emptyset$.

We now use **U4** to find a waveform W_R corresponding to $\mathcal{P}_1, \dots, \mathcal{P}_{3t}$ that starts at u, is otherwise contained in R, and uses the colours of some subset $S_F \subseteq S'$. Let w denote the other endpoint of W_R. Note that $w \in sH$ since $S' \subseteq H$.

Using **U1**, we find a rainbow path $P_{M,2}$, starting at w and otherwise contained in M, which uses all of the colours from $S_1 \setminus (S_F \cup S'_F \cup S''_F)$ except for some subset L_1 of size at most $\mu q|H|$. If we are in Case 2, since $w \in sH$ and $S_1 \subseteq H$, the path $P_{M,2}$ is completely contained in sH; in particular it is vertex-disjoint from $P_{M,1}$ (which is contained in H). (Notice that there is no $P_{M,1}$ to avoid when $J = \emptyset$ in Case 1.) Let y denote the final vertex of $P_{M,2}$.

Let $L := L_1 \cup L_2$ be the set of colours that we have yet to integrate in our rainbow path; we also include the element j^* if we deleted it from P_A previously. Notice that $|L| \leq 2\mu q|H| + 1 \leq \ell$.

For any family $\mathcal{P}'$ of precisely ℓ pairs from $\mathcal{P}_{\text{flex}}$, we can collapse the waveform W_R into a path P_R by following the system of distinct representatives for $\mathcal{P}_1 \setminus \mathcal{P}', \dots, \mathcal{P}_{3t} \setminus \mathcal{P}'$ guaranteed by the absorbing property (so in total we use the pairs in $\mathcal{P} \setminus \mathcal{P}'$). Then, $P_{M,1} \cup P_A \cup P_R \cup P_{M,2}$ is a rainbow path using precisely the colours in $S \setminus \left(L \cup \bigcup_{(a_i,b_i)\in\mathcal{P}'}\{a_i, b_i\}\right)$. This path avoids the vertex set T because $P_{M,1} \subseteq H \cap M$, $P_A \setminus \{v\} \subseteq A$, $P_R \setminus \{u\} \subseteq R$, and $P_{M,2} \setminus \{w\} \subseteq sH \cap M$.

Finally, by **U5** we can find a subfamily $\mathcal{P}' \subseteq \mathcal{P}_{\text{flex}}$ of size precisely ℓ and a rainbow path P_T, starting at y and otherwise contained in T, which uses all except possibly one of the colours from $L \cup \bigcup_{(a_i,b_i)\in\mathcal{P}'}\{a_i, b_i\}$. Take

the corresponding path P_R from the previous paragraph. Then $P_{M,1} \cup P_A \cup P_R \cup P_{M,2} \cup P_T$ is a rainbow path using all but one colour from S, as desired. □

10. Concluding remarks

As we have seen in Section 9, our methods in the case of dense subsets $S \subset G$ apply to Problem 1.1 just as well over arbitrary groups as in the specialised setting of $\mathbb{F}_2^n$. The basic randomness vs. structure dichotomy (see Section 2) also translates well to general groups. However, a key complication for general groups is that the structure of subsets with bounded doubling is more complicated; already for $\mathbb{F}_p$ one has to work with generalised arithmetic progressions in place of proper subgroups. In particular, over $\mathbb{F}_p$, we have no means of passing to a robust expander of size $O(|S|)$ and finishing most of the job there. There are also further complications over $\mathbb{F}_p$ for the absorption part of the argument, not least because we no longer have access to popular sums as we did over $\mathbb{F}_2^n$, or as we did in the dense case. Novel ideas are required to settle both of these issues in order to use our framework to settle Graham's conjecture for large p.

References

[1] B. Alspach and G. Liversidge, *On strongly sequenceable abelian groups*, Art Discrete Appl. Math. (2020). 2, 33

[2] L. D. Andersen, *Hamilton circuits with many colours in properly edge-coloured complete graphs*, Math. Scand. **64** (1989), no. 1, 5–14. 2

[3] S. T. Bate and B. Jones, *A review of uniform cross-over designs*, J. Statist. Plann. Inference **138** (2008), no. 2, 336–351. 3

[4] B. Bedert and N. Kravitz, *Graham's rearrangement conjecture beyond the rectification barrier*, Isr. J. Math. (to appear), arXiv preprint 2409.07403. 2

[5] M. Bucić, B. Frederickson, A. Müyesser, A. Pokrovskiy, and L. Yepremyan, *Towards Graham's rearrangement conjecture via rainbow paths*, arXiv preprint 2503.01825 (2025). 2, 5, 7, 17, 18, 33

[6] G. Conant, A. Pillay, and C. Terry, *Structure and regularity for subsets of groups with finite VC-dimension*, J. Eur. Math. Soc. **24** (2021), no. 2, 583–621. 3

[7] D. Conlon and J. Fox, *Bounds for graph regularity and removal lemmas*, Geometric and Functional Analysis **22** (2012), no. 5, 1191–1256. 4

[8] S. Costa and M. A. Pellegrini, *Some new results about a conjecture by Brian Alspach*, Arch. Math. **115** (2020), no. 5, 479–488. 2

[9] S. Costa, S. Della Fiore, and E. Engel, *Graham's rearrangement for dihedral groups*, arXiv preprint 2503.18101 (2025). 2

[10] S. Costa, F. Morini, A. Pasotti, and M. A. Pellegrini, *A problem on partial sums in abelian groups*, Discrete Math. **341** (2018), no. 3, 705–712. 2

[11] S. Eberhard, F. Manners, and R. Mrazović, *An asymptotic for the Hall–Paige conjecture*, Adv. Math. **404** (2022), 108423. 1

[12] P. Erdős and R. L. Graham, *Old and new problems and results in combinatorial number theory*, Université de Genève, L'Enseignement Mathématique, Geneva, 1980. 2

[13] P. Erdős, A. Gyárfás, and L. Pyber, *Vertex coverings by monochromatic cycles and trees*, J. Comb. Theory Ser. B **51** (1991), no. 1, 90–95. 5, 24

[14] C. Even-Zohar, *On sums of generating sets in* $\mathbb{Z}_2^n$, Combin. Probab. Comput. **21** (2012), no. 6, 916–941. 7, 28

[15] G. B. Folland, *A course in abstract harmonic analysis*, 2nd ed., Texts in Pure Mathematics, CRC Press, 2016. 10

[16] A. Frieze and R. Kannan, *Quick approximation to matrices and applications*, Combinatorica **19** (1999), no. 2, 175–220. 4

[17] B. Gordon, *Sequences in groups with distinct partial products*, Pacific J. Math. **11** (1961), 1309–1313. 1

[18] W. T. Gowers, B. Green, F. Manners, and T. Tao, *On a conjecture of Marton*, Annals of Mathematics **201** (2025), no. 2, 515–549. 7

[19] R. L. Graham, *On sums of integers taken from a fixed sequence*, Proceedings of the Washington State University Conference on Number Theory (Washington State Univ., Pullman, Wash., 1971), Washington State University, Department of Mathematics, Pi Mu Epsilon, Pullman, WA, 1971, pp. 22–40. 2

[20] B. Green, *Finite field models in additive combinatorics*, arXiv preprint math/0409420 (2004). 2

[21] B. Green, *A Szemerédi-type regularity lemma in abelian groups, with applications*, Geom. Funct. Anal. **15** (2005), no. 2, 340–376. 3, 4
[22] B. Green and T. Tao, *Freiman's theorem in finite fields via extremal set theory*, Combin. Probab. Comput. **18** (2009), no. 3, 335–355. 7, 28
[23] B. Green and T. Tao, *An arithmetic regularity lemma, an associated counting lemma, and applications*, An Irregular Mind: Szemerédi is 70, Springer, 2010, pp. 261–334. 3
[24] V. Gruslys and S. Letzter, *Cycle partitions of regular graphs*, Combin. Probab. Comput. **30** (2021), no. 4, 526–549. 4, 7
[25] M. Hall and L. J. Paige, *Complete mappings of finite groups*, Pacific J. Math. **5** (1955), 541–549. 1
[26] J. Hicks, M. A. Ollis, and J. R. Schmitt, *Distinct partial sums in cyclic groups: polynomial method and constructive approaches*, J. Combin. Des. **27** (2019), no. 6, 369–385. 2
[27] K. Hosseini, S. Lovett, G. Moshkovitz, and A. Shapira, *An improved lower bound for arithmetic regularity*, Mathematical Proceedings of the Cambridge Philosophical Society, vol. 161, Cambridge University Press, 2016, pp. 193–197. 4
[28] A. D. Keedwell and J. Dénes, *Latin squares and their applications*, second ed., Elsevier/North-Holland, Amsterdam, 2015. 3
[29] A. Keedwell, P. Cameron, J. Hirschfeld, and D. Hughes, *Sequenceable groups: a survey*, LMS Lecture Notes **49** (1981), 205–215. 1
[30] N. Kravitz, *Rearranging small sets for distinct partial sums*, arXiv preprint 2407.01835 (2024). 2
[31] M. Krivelevich and B. Sudakov, *Pseudo-random graphs*, More Sets, Graphs and Numbers (G. Erdős, D. Miklós, and V. T. Sós, eds.), Bolyai Society Mathematical Studies, vol. 15, Springer, 2006, pp. 199–262. 11
[32] D. Kühn, A. Lo, D. Osthus, and K. Staden, *The robust component structure of dense regular graphs and applications*, Proc. London Math. Soc. **110** (2015), no. 1, 19–56. 4, 7
[33] D. Kühn and D. Osthus, *Hamilton decompositions of regular expanders: applications*, J. Combin. Theory Ser. B **104** (2014), 1–27. 4, 17
[34] R. Montgomery, A. Pokrovskiy, and B. Sudakov, *A proof of Ringel's conjecture*, Geom. Funct. Anal. **31** (2021), no. 3, 663–720. 5
[35] R. Montgomery, *Spanning trees in random graphs*, Adv. Math. **356** (2019). 8, 35
[36] A. Müyesser and A. Pokrovskiy, *A random Hall-Paige conjecture*, Inventiones mathematicae **240** (2025), 779–867. 1, 2, 3, 17, 24, 25, 42
[37] M. A. Ollis, *Sequences in dihedral groups with distinct partial products*, arXiv preprint 1904.07646. 2
[38] M. A. Ollis, *Sequenceable groups and related topics*, Electron. J. Combin. **DS10** (2002), 34. 2, 3
[39] S. Peluse, *Finite field models in arithmetic combinatorics–twenty years on*, arXiv preprint 2312.08100 (2023). 2
[40] G. Ringel, *Cyclic arrangements of the elements of a group.*, Notices Amer. Math. Soc. **21** (1974), no. 1, A–95. 1
[41] G. Ringel, *Map color theorem*, Die Grundlehren der mathematischen Wissenschaften, vol. 209, Springer-Verlag, New York-Heidelberg, 1974. 1
[42] V. Rödl, E. Szemerédi, and A. Ruciński, *An approximate Dirac-type theorem for k-uniform hypergraphs*, Combinatorica **28** (2008), no. 2, 229–260. 5
[43] B. Sudakov, *Restricted subgraphs of edge-colored graphs and applications*, arXiv preprint 2412.13945. 5
[44] T. Tao and V. H. Vu, *Additive combinatorics*, vol. 105, Cambridge University Press, 2006. 5
[45] J. Wolf, *Finite field models in arithmetic combinatorics–ten years on*, Finite Fields and Their Applications **32** (2015), 233–274. 2

Appendix A. Extremely dense case

For brevity, we write $K_G^- := \mathrm{Cay}_G(G \setminus \{\mathrm{id}\})$ in this appendix.

For subsets $R, C \subseteq G$, we write $K_G^-[R; C]$ to denote the subgraph of K_G^- induced on the vertex set R by the edges with colours in C. For disjoint subsets $V_1, V_2 \subseteq G$, we write $K_G^-[V_1, V_2; C]$ to denote the bipartite subgraph of K_G^- obtained by keeping only the directed edges from V_1 to V_2 with colours in C.

The following lemma is part of Lemma 6.22 from [36]. The proof combines the sorting network method and the statement of the random Hall–Paige conjecture. The original statement pertains to both addition and multiplication tables, but to reduce clutter we have included only the part that we will need. In the remainder of this appendix, we will perform many calculations in the abelianisation $G/[G, G]$ of G; since the order of multiplication does not matter in the abelianisation, product notation such as $\prod_{v \in V} v$ is unambiguous.

Lemma A.1. *Let $1/n \ll \gamma, p \leq 1$, and let $(\log n)^7 \leq t \leq (\log n)^8$ be an integer. Set $q := p/(t-1)$. Let G be a group of order n. Let $V_{\mathrm{str}}, V_{\mathrm{mid}}, V_{\mathrm{end}}$ be disjoint random subsets of G with $V_{\mathrm{str}}, V_{\mathrm{end}}$ q-random and V_{mid} p-random. Let C be a $(q+p)$-random subset of G, sampled independently of $V_{\mathrm{str}}, V_{\mathrm{mid}}, V_{\mathrm{end}}$. Then with high probability the following holds for all choices of $C' \subseteq G$ and disjoint subsets $V'_{\mathrm{str}}, V'_{\mathrm{end}}, V'_{\mathrm{mid}} \subseteq G$:*

If $C', V'_{\mathrm{str}}, V'_{\mathrm{end}}, V'_{\mathrm{mid}}$ satisfy

(1) for each $R \in \{V_{\mathrm{str}}, V_{\mathrm{mid}}, V_{\mathrm{end}}, C\}$, we have $|R \Delta R'| \leq n^{1-\gamma}$;
(2) $\prod V'_{\mathrm{end}} \cdot (\prod V'_{\mathrm{str}})^{-1} = \prod C' \pmod{[G, G]}$;
(3) $\mathrm{id} \notin C'$;
(4) $|V'_{\mathrm{str}}| = |V'_{\mathrm{end}}| = |V'_{\mathrm{mid}}|/(t-1) = |C'|/t$,

then for every bijection $f \colon V'_{\mathrm{str}} \to V'_{\mathrm{end}}$, the graph $K_G^-[V'_{\mathrm{str}} \cup V'_{\mathrm{end}} \cup V'_{\mathrm{mid}}; C']$ has a rainbow collection of vertex-disjoint paths $\{P_v : v \in V'_{\mathrm{str}}\}$, where each P_v has length t and starts at v and ends at $f(v)$.

We can now prove the main result of the appendix. Theorem 6.9 of [36] gives a sharper version of this result in the regime $\gamma \geq 1/2$. In fact, the same proof works verbatim for the larger range $1/n \ll \gamma < 1$, but this flexibility is unfortunately not recorded in [36], as the authors did not anticipate that it would have further applications. We follow the proof from [36] quite closely in our discussion here. The main idea is applying Lemma A.1 twice, in a such a way that the starting-vertices of one collection of paths correspond to the ending-vertices of the other collection of paths, and vice versa, so that together all of the paths form a single long path.

Theorem A.2. *Let $1/N \ll \gamma \leq 1$. If G is a group of order N and $S \subseteq G \setminus \{\mathrm{id}\}$ is a subset with $|S| \geq N - N^{1-\gamma}$, then S has a valid ordering, i.e., the Cayley graph $\mathrm{Cay}_G(S)$ has a directed rainbow path with $|S| - 1$ edges.*

Proof. Fix distinct $x, y \in G$ such that that $yx^{-1} = \prod S \pmod{[G, G]}$. We will show that there is a directed rainbow path from x to y with $|S|$ edges. Note that if G is abelian and $\sum S = 0$, then there are no such distinct x and y; in this case we simply delete an element of S so that $\sum S \neq 0$, and applying our argument with this new S still produces the desired rainbow path using all but one of the colours from the original set S.

Set $t := 2\lfloor (\log N)^7 \rfloor$ and $s := |S|$. Set $q := 1/(2t)$ and $p := (t-1)q$. Take a random partition $V_{\mathrm{str}} \sqcup V_{\mathrm{end}} \sqcup V_{\mathrm{mid},1} \sqcup V_{\mathrm{mid},2}$ of G where the former two parts are q-random and the latter two are p-random. Independently, take a random partition of G into 1/2-random sets C_0 and C_1 (note that $1/2 = p + q$ here).

With high probability, Lemma A.1 applies with $V_{\mathrm{str}}, V_{\mathrm{end}}, V_{\mathrm{mid},1}$, C_0 (playing the role of C) and t; and Lemma A.1 applies with V_{str} and V_{end} interchanged, with $V_{\mathrm{mid},2}$ instead of $V_{\mathrm{mid},1}$, and with C_1 playing the role of C. In each of these applications of Lemma A.1, let $\gamma/10$ play the role of γ. Furthermore, we can ensure that with high probability for each $g \in G$, each random set with randomness parameter z contains a disjoint collection of $\Omega(z^3 N)$ triples (a, b, c) of distinct group elements with $abc = g$; call this property $(*)$. This property follows from Chernoff's bound and the fact that for each g we can find a disjoint collection of $\Omega(N)$ distinct triples (a, b, c) with $abc = g$ (and then we union-bound over g). Notice also that with high probability each random set with parameter z has $zN \pm N^{0.6}$ elements, again by Chernoff's bound.

Fix an outcome for the random sets satisfying the properties described in the previous paragraph. We will construct slightly-modified "prime" versions of the sets $V_{\mathrm{str}}, V_{\mathrm{mid},1}, V_{\mathrm{mid},2}, V_{\mathrm{end}}$ so that the hypotheses of Lemma A.1 are satisfied. We will remove small "junk sets" to guarantee the divisibility condition (4), and then we will interchange a few elements among the sets to guarantee the product condition (2).

Let ℓ be the largest integer satisfying $2t\ell - t + 1 \le s + 1$. Note that $\ell = qs + O(1) = qN \pm N^{1-\gamma/2}$. Define $w := (s+1) - (2t\ell - t + 1)$, and note that $w \le (\log n)^{10}$. Observe that we can greedily find an S-rainbow path P_0, from x to some x', which has exactly $w+1$ vertices and does not pass through y; fix such a path.

By modifying at most $n^{1-\gamma/2}$ elements from each of the sets $V_{\mathrm{str}}, V_{\mathrm{mid},1}, V_{\mathrm{mid},2}, V_{\mathrm{end}}$, we can obtain new disjoint sets $V'_{\mathrm{str}}, V'_{\mathrm{mid},1}, V'_{\mathrm{mid},2}, V'_{\mathrm{end}} \subseteq (V \setminus V(P_0)) \cup \{x'\}$ such that $x' \in V'_{\mathrm{str}}$, $y \in V'_{\mathrm{end}}$, and $\ell = |V'_{\mathrm{str}}| = |V'_{\mathrm{end}}| = |V'_{\mathrm{mid},1}|/(t-1) = |V'_{\mathrm{mid},2}|/(t-1)$. Similarly, we partition $S \setminus C(P_0)$ into sets C'_0 of size $t\ell$ and C'_1 of size $t(\ell-1)$ (this is possible due to the divisibility constraint on the size of P_0) satisfying $|C_0 \Delta C'_0|, |C_1 \Delta C'_1| \le n^{1-2\gamma/3}$. Furthermore, we can interchange a few elements, thanks to property $(*)$, to ensure that

$$\prod V'_{\mathrm{str}} \setminus \{x'\} = \prod V'_{\mathrm{end}} \setminus \{y\} = \prod C'_1 = \mathrm{id} \pmod{[G,G]}.$$

This implies that $\prod C'_0 = \prod S \left(\prod C(P_0)\right)^{-1} \pmod{[G,G]}$. Our interchanges maintain the property that $|Z \Delta Z'| \le n^{1-\gamma/2}$ for each set Z.

We now invoke Lemma A.1 for the sets $V'_{\mathrm{str}} \setminus \{x'\}$, $V'_{\mathrm{end}} \setminus \{y\}$, $V'_{\mathrm{mid},2}$, C'_1 with an arbitrary choice of bijection to get a partition into paths of length t with starting points in $V'_{\mathrm{end}} \setminus \{y\}$ and endpoints in $V'_{\mathrm{str}} \setminus \{x'\}$. We wish to now invoke Lemma A.1 in the opposite orientation, with the sets V'_{str}, V'_{end}, $V'_{\mathrm{mid},2}$, C'_1, and a choice of bijection that we will shortly specify (to ensure that everything links up to form a path). First, we check the relevant product condition.

Claim A.3. *We have* $\prod V'_{\mathrm{end}} (\prod V'_{\mathrm{str}})^{-1} = \prod C'_0 \pmod{[G,G]}$.

Proof. We carry out the following calculations modulo $[G,G]$. Note first that $\prod V'_{\mathrm{end}} (\prod V'_{\mathrm{str}})^{-1} = y(x')^{-1}$ as from the previous exchanges we had ensured that $\prod V'_{\mathrm{str}} \setminus \{x'\} = \prod V'_{\mathrm{end}} \setminus \{y\}$. Recall that $yx^{-1} = \prod S$, and that $\prod C(P_0) = x'x^{-1}$, as P_0 is a path from x to x'. Thus $y(x')^{-1} = \prod S (\prod C(P_0))^{-1}$. We also previously saw that $\prod C'_0 = \prod S \left(\prod C(P_0)\right)^{-1}$; this completes the proof of the claim. □

Now, we specify the bijection that will ensure that the concatanation of all paths we have constructed so far yields an S-rainbow path from x to y. Suppose that the previously-constructed collection of paths had endpoints $y_1 \to x_1$, $y_2 \to x_2$, …, $y_{\ell-1} \to x_{\ell-1}$. Then we choose the bijection that maps $x_1 \to y_2$, $x_2 \to y_3$, …, $x_{\ell-2} \to y_{\ell-1}$, $x_{\ell-1} \to y$, $x' \to y_1$. The union of the resulting paths from the two applications of Lemma A.1, together with P_0, yields a rainbow path from x to y whose edges use precisely the colours of S. □