MODEL-THEORETIC ELEKES-SZABÓ FOR STABLE AND O-MINIMAL HYPERGRAPHS

ARTEM CHERNIKOV, YA'ACOV PETERZIL, AND SERGEI STARCHENKO

Abstract. A theorem of Elekes and Szabó recognizes algebraic groups among certain complex algebraic varieties with maximal size intersections with finite grids. We establish a generalization to relations of any arity and dimension, definable in: 1) stable structures with distal expansions (includes algebraically and differentially closed fields of characteristic 0); and 2) o-minimal expansions of groups. Our methods provide explicit bounds on the power saving exponent in the non-group case. Ingredients of the proof include: a higher arity generalization of the abelian group configuration theorem in stable structures, along with a purely combinatorial variant characterizing Latin hypercubes that arise from abelian groups; and Zarankiewicz-style bounds for hypergraphs definable in distal structures.

Contents

Department of Mathematics, University of California, Los Angeles, Los Angeles, CA 90095, USA

Department of Mathematics, University of Haifa, Haifa, Israel

Department of Mathematics, University of Notre Dame, Notre Dame, IN, 46656, USA

E-mail addresses: (Corresonding) chernikov@math.ucla.edu, kobi@math.haifa.ac.il, sstarche@nd.edu.

2010 *Mathematics Subject Classification.* 03C45, 52C10.

1. Introduction

1.1. **History, and a special case of the main theorem.** Erdős and Szemerédi [24] observed the following sum-product phenomenon: there exists $c \in \mathbb{R}_{>0}$ such that for any finite set $A \subseteq \mathbb{R}$,

$$\max \{|A + A|, |A \cdot A|\} \geq |A|^{1+c}.$$

They conjectured that this holds with $c = 1-\varepsilon$ for an arbitrary $\varepsilon \in \mathbb{R}_{>0}$, and by the work of Solymosi [55] and Konyagin and Shkredov [35] it is known to hold with $c = \frac{1}{3} + \varepsilon$ for some sufficiently small ε. Elekes and Rónyai [22] generalized this by showing that for any polynomial $f(x, y) \in \mathbb{R}[x, y]$ there exists $c > 0$ such that for every finite set $A \subseteq \mathbb{R}$ we have

$$|f(A \times A)| \geq |A|^{1+c},$$

unless f is either additive or multiplicative, i.e. of the form $g(h(x)+i(y))$ or $g(h(x) \cdot i(y))$ for some univariate polynomials g, h, i respectively. The bound was improved to $\Omega_{\deg f}\left(|A|^{\frac{11}{6}}\right)$ in [50].

Elekes and Szabó [23] established a conceptual generalization of this result explaining the exceptional role played by the additive and multiplicative forms: for any irreducible polynomial $Q(x, y, z)$ over $\mathbb{C}$ depending on all of its coordinates and such that its set zero set has dimension 2, either there exists some $\varepsilon > 0$ such that F has at most $O(n^{2-\varepsilon})$ zeroes on all finite $n \times n \times n$ grids, or F is in a coordinate-wise finite-to-finite correspondence with the graph of multiplication of an algebraic group (see Theorem (B) below for a more precise statement). In the special Elekes-Rónyai case above, taking Q to be the graph of the polynomial function f, the resulting group is either the additive or the multiplicative group of the field. Several generalizations, refinements and variants of this influential result were obtained recently

[13,30,33,48,49,51,56], in particular for complex algebraic relations of higher dimension and arity by Bays and Breuillard [8].

In this paper we obtain a generalization of the Elekes-Szabó theorem to hypergraphs of any arity and dimension definable in stable structures admitting distal expansions (this class includes algebraically and differentially closed fields of characteristic 0 and compact complex manifolds); as well as for arbitrary *o*-minimal structures. Before explaining our general theorems, we state two very special corollaries.

Theorem (A). *(Corollary 6.21) Assume $s \geq 3$ and $Q \subseteq \mathbb{R}^s$ is semi-algebraic, of description complexity D (i.e. given by at most D polynomial (in-)equalities, with all polynomials of degree at most D, and $s \leq D$), such that the projection of Q to any $s-1$ coordinates is finite-to-one. Then exactly one of the following holds.*

(1) There exists a constant c, depending only on s and D, such that: for any $n \in \mathbb{N}$ and finite $A_i \subseteq \mathbb{R}$ with $|A_i| = n$ for $i \in [s]$ we have

$$|Q \cap (A_1 \times \ldots \times A_s)| \leq cn^{s-1-\gamma},$$

where $\gamma = \frac{1}{3}$ if $s \geq 4$, and $\gamma = \frac{1}{6}$ if $s = 3$.

(2) There exist open sets $U_i \subseteq \mathbb{R}, i \in [s]$, an open set $V \subseteq \mathbb{R}$ containing 0, and analytic bijections with analytic inverses $\pi_i : U_i \to V$ such that

$$\pi_1(x_1) + \cdots + \pi_s(x_s) = 0 \Leftrightarrow Q(x_1, \ldots, x_s)$$

for all $x_i \in U_i, i \in [s]$.

Theorem (B). *(Corollary 5.51) Assume $s \geq 3$, and let $Q \subseteq \mathbb{C}^s$ be an irreducible algebraic variety so that for each $i \in [s]$, the projection of Q to any $s-1$ coordinates is generically finite. Then exactly one of the following holds.*

(1) There exist c depending only on $s, \deg(Q)$ such that: for any $n \in \mathbb{N}$ and $A_i \subseteq \mathbb{C}_i$ with $|A_i| = n$ for $i \in [s]$ we have

$$|Q \cap (A_1 \times \ldots \times A_s)| \leq cn^{s-1-\gamma}$$

where $\gamma = \frac{1}{11}$ if $s \geq 4$, and $\gamma = \frac{1}{22}$ if $s = 3$.

(2) For G one of $(\mathbb{C}, +)$, $(\mathbb{C}, \times)$ or an elliptic curve group, Q is in coordinate-wise correspondence (see Section 5.8) with

$$Q' := \{(x_1, \ldots, x_s) \in G^s : x_1 \cdot \ldots \cdot x_s = 1_G\}.$$

Remark 1.1. Theorem (B) is similar to the codimension 1 case of [8, Theorem 1.4], however our method provides an explicit bound on the exponent in Clause (1).

Remark 1.2. Theorems (A) and (B) correspond to the 1-dimensional case of Corollaries 6.20 and 5.48, respectively, which allow $Q \subseteq \prod_{i\in[s]} X_i$ with $\dim(X_i) = d$ for an arbitrary $d \in \mathbb{N}$.

Remark 1.3. Note the important difference — Theorem (A) is *local*, i.e. we can only obtain a correspondence of Q to a subset of a group after restricting to *some* open subsets U_i. This is unavoidable in an ordered structure since the high count in Theorem (A.2) might be the result of a local phenomenon in Q. E.g. when Q is the union of $Q_1 = \{\bar{x} : x_1 + \cdots + x_s = 0\} \cap (-\varepsilon, \varepsilon)^s$, for some $\varepsilon > 0$, and another set Q_2 for which the count is low.

1.2. **The Elekes-Szabó principle.** We now describe the general setting of our main results. We let $\mathcal{M} = (M, \ldots)$ be an arbitrary first-order structure, in the sense of model theory, i.e. a set M equipped with some distinguished functions and relations. As usual, a subset of M^d is definable if it is the set of tuples satisfying a formula (with parameters). Two key examples to keep in mind are $(\mathbb{C}, +, \times, 0, 1)$ (in which definable sets are exactly the constructible ones, i.e. boolean combinations of the zero-sets of polynomials, by Tarski's quantifier elimination) and $(\mathbb{R}, +, \times, <, 0, 1)$ (in which definable sets are exactly the semialgebraic ones, by Tarski-Seidenberg quantifier elimination). We refer to [40] for an introduction to model theory and the details of the aforementioned quantifier elimination results.

From now on, we fix a structure $\mathcal{M}$, $s \in \mathbb{N}$, definable sets $X_i \subseteq M^{d_i}, i \in [s]$, and a definable relation $Q \subseteq \bar{X} = X_1 \times \ldots \times X_s$. We write $A_i \subseteq_n X_i$ if $A_i \subseteq X_i$ with $|A_i| \leq n$. By a *grid on* $\bar{X}$ we mean a set $\bar{A} \subseteq \bar{X}$ with $\bar{A} = A_1 \times \ldots \times A_s$ and $A_i \subseteq X_i$. By an *n-grid on* $\bar{X}$ we mean a grid $\bar{A} = A_1 \times \ldots \times A_s$ with $A_i \subseteq_n X_i$.

Definition 1.4. For $d \in \mathbb{N}$, we say that a relation $Q \subseteq X_1 \times X_2 \times \ldots \times X_s$ is *fiber-algebraic, of degree* d if for any $i \in [s]$ we have

$$\forall x_1 \in X_1 \ldots \forall x_{i-1} \in X_{i-1} \forall x_{i+1} \in X_{i+1} \ldots \forall x_s \in X_s$$
$$\exists^{\leq d} x_i \in X_i \ (x_1, \ldots, x_s) \in Q.$$

We say that $Q \subseteq X_1 \times X_2 \times \ldots \times X_s$ is *fiber-algebraic* if it is fiber-algebraic of degree d for some $d \in \mathbb{N}$.

In other words, fiber algebraicity means that the projection of Q onto any $s-1$ coordinates is finite-to-one. For example, if $Q \subseteq X_1 \times X_2 \times X_3$ is fiber-algebraic of degree d, then for any $A_i \subseteq_n X_i$ we have $|Q \cap A_1 \times A_2 \times A_3| \leq dn^2$. Conversely, let $Q \subseteq \mathbb{C}^3$ be given by $x_1 + x_2 - x_3 = 0$, and let $A_1 = A_2 = A_3 = \{0, \ldots, n-1\}$. Then $|Q \cap A_1 \times A_2 \times A_3| = \frac{n(n+1)}{2} = \Omega\left(n^2\right)$. This indicates that the upper and lower bounds match for the graph of addition in an abelian group (up to a constant) — and the Elekes-Szabó principle suggests that in many situations this is the only possibility. Before making this precise, we introduce some notation.

1.2.1. *Grids in general position.* From now on we will assume that $\mathcal{M}$ is equipped with some notion of integer-valued dimension on definable sets, to be specified later. A good example to keep in mind is Zariski dimension on constructible subsets of $\mathbb{C}^d$, or the topological dimension on semialgebraic subsets of $\mathbb{R}^d$.

Definition 1.5. (1) Let X be a definable set in $\mathcal{M}$, and let $\mathcal{F}$ be a definable family of subsets of X. For $\nu \in \mathbb{N}$, we say that a set $A \subseteq X$ is *in* $(\mathcal{F}, \nu)$-*general position* if $|A \cap F| \leq \nu$ for every $F \in \mathcal{F}$ with $\dim(F) < \dim(X)$.

(2) Let X_i, $i = 1, \dots, s$, be definable sets in $\mathcal{M}$. Let $\bar{\mathcal{F}} = (\mathcal{F}_1, \dots, \mathcal{F}_s)$, where $\mathcal{F}_i$ is a definable family of subsets of X_i. For $\nu \in \mathbb{N}$ we say that a grid $\bar{A}$ on $\bar{X}$ is in $(\bar{\mathcal{F}}, \nu)$-*general position* if each A_i is in $(\mathcal{F}_i, \nu)$-general position.

For example, when $\mathcal{M}$ is the field $\mathbb{C}$, a subset of $\mathbb{C}^d$ is in a $(\mathcal{F}, \nu)$-general position if any variety of smaller dimension and bounded degree (determined by the formula defining $\mathcal{F}$) can cut out only ν points from it (see the proof of Corollary 5.48). Also, if $\mathcal{F}$ is any definable family of subsets of $\mathbb{C}$, then for any large enough ν, every $A \subseteq X$ is in $(\mathcal{F}, \nu)$-general position. On the other hand, let $X = \mathbb{C}^2$ and let $\mathcal{F}_d$ be the family of algebraic curves of degree less than d. If $\nu \leq d + 1$, then any set $A \subseteq X$ with $|A| \geq \nu$ is not in $(\mathcal{F}_d, \nu - 1)$-general position.

1.2.2. *Generic correspondence with group multiplication.* We assume that $\mathcal{M}$ is a sufficiently saturated structure, and let $Q \subseteq \bar{X}$ be a definable relation and $(G, \cdot, 1_G)$ a connected type-definable group in $\mathcal{M}^{\mathrm{eq}}$. Type-definability means that the underlying set G of the group is given by the intersection of a small (but possibly infinite) collection of definable sets, and the multiplication and inverse operations are relatively definable. Such a group is connected if it contains no proper type-definable subgroup of small index (see e.g. [40, Chapter 7.5]). And $\mathcal{M}^{\mathrm{eq}}$ is the structure obtained from $\mathcal{M}$ by adding sorts for the quotients of definable sets by definable equivalence relations in $\mathcal{M}$ (see e.g. [40, Chapter 1.3]). In the case when $\mathcal{M}$ is the field $\mathbb{C}$, connected type-definable groups are essentially just the complex algebraic groups connected in the sense of Zariski topology (see Section 5.8 for a discussion and further references).

Definition 1.6. We say that Q is in a *generic correspondence with multiplication in* G if there exist a small set $A \subseteq M$ and elements $g_1, \dots, g_s \in G$ such that:

(1) $g_1 \cdot \dots \cdot g_s = 1_G$;
(2) $g_1, \dots, g_{s-1}$ are independent generics in G over A (i.e. each g_i does not belong to any definable set of dimension smaller than G definable over $A \cup \{g_1, \dots, g_{i-1}, g_{i+1}, \dots, g_{s-1}\}$);
(3) For each $i = 1, \dots, s$ there is a generic element $a_i \in X_i$ inter-algebraic with g_i over $\mathcal{A}$ (i.e. $a_i \in \operatorname{acl}(g_i, A)$ and $g_i \in \operatorname{acl}(a_i, A)$, where acl is the model-theoretic algebraic closure), such that $(a_1, \dots, a_s) \in Q$.

Remark 1.7. There are several variants of "generic correspondence with a group" considered in the literature around the Elekes-Szabó theorem. The one that we use arises naturally at the level of generality we work with, and as we discuss in Sections 5.8 and 6.4 it easily specializes to the notions considered previously in several cases of interest (e.g. the algebraic coordinate-wise finite-to-finite correspondence in the case of constructible sets in Theorem (B), or coordinate-wise analytic bijections on a neighborhood in the case of semialgebraic sets in Theorem (A)).

1.2.3. *The Elekes-Szabó principle.* Let $s \geq 3, k \in \mathbb{N}$ and $X_1, \ldots, X_s$ be definable sets in a sufficiently saturated structure $\mathcal{M}$ with $\dim(X_i) = k$.

Definition 1.8. We say that $X_1, \ldots, X_s$ satisfy the *Elekes-Szabó principle* if for any fiber-algebraic definable relation $Q \subseteq \bar{X}$, one of the following holds:

(1) Q *admits power saving*: there exist some $\gamma \in \mathbb{R}_{>0}$ and some definable families $\mathcal{F}_i$ on X_i such that: for any $\nu \in \mathbb{N}$ and any n-grid $\bar{A} \subseteq \bar{X}$ in $(\bar{F}, \nu)$-general position, we have $|Q \cap \bar{A}| = O_\nu\left(n^{(s-1)-\gamma}\right)$;
(2) there exists a type-definable subset of Q of full dimension that is in a generic correspondence with multiplication in some type-definable *abelian* group G of dimension k.

The following are the previously known cases of the Elekes-Szabó principle:

(1) [23] $\mathcal{M} = (\mathbb{C}, +, \times)$, $s = 3$, k arbitrary (no explicit exponent γ in power saving; no abelianity of the algebraic group for $k > 1$);
(2) [48] $\mathcal{M} = (\mathbb{C}, +, \times)$, $s = 3$, $k = 1$ (explicit γ in power saving);
(3) [49] $\mathcal{M} = (\mathbb{C}, +, \times)$, $s = 4$, $k = 1$ (explicit γ in power saving);
(4) [51] $\mathcal{M} = (\mathbb{C}, +, \times)$, $k = 1$, Q is the graph of an s-ary polynomial function for an arbitrary s (i.e. this is a generalization of Elekes-Rónyai to an arbitrary number of variables);
(5) [8] $\mathcal{M} = (\mathbb{C}, +, \times)$, s and k arbitrary, abelianity of the group for $k > 1$ (they work with a more relaxed notion of general position and arbitrary codimension, however no bounds on γ);
(6) [20] $\mathcal{M}$ is any strongly minimal structure interpretable in a distal structure (see Section 2), $s = 3$, $k = 1$.

In the first five cases the dimension is the Zariski dimension, and in the sixth case the Morley rank.

1.3. Main theorem.

We can now state the main result of this paper.

Theorem (C). *The Elekes-Szabó principle holds in the following two cases:*

(1) *(Theorem 5.24) $\mathcal{M}$ is a stable structure interpretable in a distal structure, with respect to $\mathfrak{p}$-dimension (see Section 5.1, and below).*
(2) *(Theorem 6.4) $\mathcal{M}$ is an o-minimal structure expanding a group, with respect to the topological dimension. In this case, on a type-definable generic subset of $\bar{X}$, we get a definable coordinate-wise bijection of Q with the graph of multiplication of an abelian type-definable group G (we stress that this G is a priori unrelated to the underlying group that $\mathcal{M}$ expands).*

Moreover, the power saving bound is explicit in (2) (see the statement of Theorem 6.4), and is explicitly calculated from a given distal cell decomposition for Q in (1) (see Theorem 5.27).

Examples of structures satisfying the assumption of Theorem (C.1) include: algebraically closed fields of characteristic 0, differentially closed fields of characteristic 0 with finitely many commuting derivations, compact complex manifolds. In particular, Theorem (B) follows from Theorem (C.1) with $k = 1$, combined with some basic model theory of algebraically closed fields (see Section 5.8). We refer to [46] for a detailed treatment of stability, and to [57, Chapter 8] for a quick introduction. See Section 2 for a discussion of distality.

Examples of *o*-minimal structures include real closed fields (in particular, Theorem (A) follows from Theorem (C.2) with $k = 1$ combined with some basic *o*-minimality, see Section 6.4), $\mathbb{R}_{\text{exp}} = (\mathbb{R}, +, \times, e^x)$, $\mathbb{R}_{\text{an}} = \left(\mathbb{R}, +, \times, f \restriction_{[0,1]^k}\right)$ for $k \in \mathbb{N}$ and f ranging over all functions real-analytic on some neighborhood of $[0, 1]^k$, or the combination of both $\mathbb{R}_{\text{an,exp}}$. We refer to [58] for a detailed treatment of *o*-minimality, or to [52, Section 3] and reference there for a quick introduction.

Remark 1.9. The assumption that $\mathcal{M}$ is an *o*-minimal expansion of a group in Theorem (C.2) can be relaxed to the more general assumption that $\mathcal{M}$ is an *o*-minimal structure with *definable Skolem functions* (see e.g. [21] for a detailed discussion of Skolem functions and related notions), but possibly with a weaker bound on the power saving exponent than the one stated in Theorem 6.4. Indeed, the γ in the γ-power saving stated in Theorem 6.4 depends on γ in the γ-ST property, and hence on $t = 2d_1 - 2$, in Fact 2.15(2) — the proof of which uses that $\mathcal{M}$ is an *o*-minimal expansion of a group. However, Fact 2.15(2) is known to hold in an arbitrary *o*-minimal structure with (at least) the weaker bound $t = 2d_1 - 1$ (see [1, Theorem 4.1]). To carry out the rest of the arguments in the proof of Theorem 6.4 in Section 6 we only use the existence of definable Skolem functions. Thus any *o*-minimal structure with definable Skolem functions satisfies the conclusion of Theorem 6.4 with $\gamma = \frac{1}{8m-3}$ if $s \geq 4$ and $\gamma = \frac{1}{16m-6}$ if $s = 3$.

1.4. **Outline of the paper.** In this section we outline the structure of the paper, and highlight some of the key ingredients of the proof of the main theorem. The proofs of (1) and (2) in Theorem (C) have similar strategy at the general level, however there are considerable technical differences. In each of the cases, the proof consists of the following key ingredients.

(1) Zarankiewicz-type bounds for distal relations (Section 2, used for both Theorem (C.1) and (C.2)).
(2) A higher arity generalization of the abelian group configuration theorem (Section 3 for the *o*-minimal case Theorem (C.2), and Section 4 for the stable case Theorem (C.1)).

(3) The dichotomy between an incidence configuration, in which case the bounds from (1) give power saving, and existence of a family of functions (or finite-to-finite correspondences) associated to Q closed under generic composition, in which case a correspondence of Q to an abelian group is obtained using (2). This is Section 5 for the stable case (C.1) and Section 6 for the *o*-minimal case (C.2).

We provide some further details for each of these ingredients, and discuss some auxiliary results of independent interest.

1.4.1. *Zarankiewicz-type bounds for distal relations (Section 2).* Distal structures constitute a subclass of purely unstable NIP structures [54] that contains all *o*-minimal structures, various expansions of the field $\mathbb{Q}_p$, and many other valued fields and related structures [2] (we refer to the introduction of [19] for a general discussion of distality in connection to combinatorics and references). Distality of a graph can be viewed as a strengthening of finiteness of its VC-dimension retaining stronger combinatorial properties of semialgebraic graphs. In particular, it is demonstrated in [15, 18, 19] that many of the results in semialgebraic incidence combinatorics generalize to relations definable in distal structures. In Section 2 we discuss distality, in particular proving the following generalized "Szemerédi-Trotter" theorem:

Theorem (D)**.** *(Theorem 2.8) For every $d \in \mathbb{N}, t \in \mathbb{N}_{\geq 2}$ and $c \in \mathbb{R}$ there exists some $C = C(d,t,c) \in \mathbb{R}$ satisfying the following.*

Assume that $E \subseteq X \times Y$ admits a distal cell decomposition $\mathcal{T}$ such that $|\mathcal{T}(B)| \leq c|B|^t$ for all finite $B \subseteq Y$. Then, taking $\gamma_1 := \frac{(t-1)d}{td-1}, \gamma_2 := \frac{td-t}{td-1}$ we have: for all $\nu \in \mathbb{N}_{\geq 2}$ and $A \subseteq_m X, B \subseteq_n Y$ such that $E \cap (A \times B)$ is $K_{d,\nu}$-free,

$$|E \cap (A \times B)| \leq C\nu \left(m^{\gamma_1} n^{\gamma_2} + m + n\right).$$

In particular, if $E \subseteq U \times V$ is a binary relation definable in a distal structure and E is $K_{s,2}$-free for some $s \in \mathbb{N}$, then there is some $\gamma > 0$ such that: for all $A \subseteq_n U, B \subseteq_n V$ we have $|E \cap A \times B| = O(n^{\frac{3}{2}-\gamma})$. The exponent strictly less that $\frac{3}{2}$ requires distality, and is strictly better than e.g. the optimal bound $\Omega(n^{\frac{3}{2}})$ for the point-line incidence relation on the affine plane over a field of positive characteristic. In the proof of Theorem (C), we will see how this γ translates to the power saving exponent in the non-group case. More precisely, for our analysis of the higher arity relation Q, we introduce the so-called γ-*Szemerédi-Trotter property*, or γ-*ST property* (Definition 2.12), capturing an iterated variant of Theorem (D), and show in Proposition 2.14 that Theorem (D) implies that every binary relation definable in a distal structure satisfies the γ-ST property for some $\gamma > 0$ calculated in terms of its distal cell decomposition. We conclude Section 2 with a discussion of the explicit bounds on γ for the γ-ST property in several particular structures of interest needed to deduce the explicit bounds on the power saving in Theorems (A) and (B).

1.4.2. *Reconstructing an abelian group from a family of bijections (Section 3).* Assume that $(G, +, 0)$ is an abelian group, and consider the s-ary relation $Q \subseteq \prod_{i \in [s]} G$ given by $x_1 + \ldots + x_s = 0$. Then Q is easily seen to satisfy the following two properties, for any permutation of the variables of Q:

$$\forall x_1, \ldots, \forall x_{s-1} \exists ! x_s Q(x_1, \ldots, x_s), \tag{P1}$$

$$\forall x_1, x_2 \forall y_3, \ldots y_s \forall y_3', \ldots, y_s' \Big(Q(\bar{x}, \bar{y}) \wedge Q(\bar{x}, \bar{y}') \rightarrow \tag{P2}$$
$$\big(\forall x_1', x_2' Q(\bar{x}', \bar{y}) \leftrightarrow Q(\bar{x}', \bar{y}') \big) \Big).$$

In Section 3 we show a converse, assuming $s \geq 4$:

Theorem (E)**.** *(Theorem 3.21) Assume $s \in \mathbb{N}_{\geq 4}$, $X_1, \ldots, X_s$ and $Q \subseteq \prod_{i \in [s]} X_i$ are sets, so that Q satisfies (P1) and (P2) for any permutation of the variables. Then there exists an abelian group $(G, +, 0_G)$ and bijections $\pi_i : X_i \rightarrow G$ such that for every $(a_1, \ldots, a_s) \in \prod_{i \in [s]} X_i$ we have*

$$Q(a_1, \ldots, a_s) \iff \pi_1(a_1) + \ldots + \pi_s(a_s) = 0_G.$$

Moreover, if Q is definable and X_i are type-definable in a sufficiently saturated structure $\mathcal{M}$, then we can take G to be type-definable and the bijections π_i relatively definable in $\mathcal{M}$.

On the one hand, this can be viewed as a purely combinatorial higher arity variant of the Abelian Group Configuration theorem (see below) in the case when the definable closure in $\mathcal{M}$ is equal to the algebraic closure (e.g. when $\mathcal{M}$ is *o*-minimal). On the other hand, if $X_1 = \ldots = X_s$, property (P1) is equivalent to saying that the relation Q is an $(s-1)$-*dimensional permutation* on the set X_1, or a *Latin* $(s-1)$-*hypercube*, as studied by Linial and Luria in [37, 38] (where Latin 2-hypercube is just a Latin square). Thus the condition (P2) in Theorem (E) characterizes, for $s \geq 3$, those Latin s-hypercubes that are given by the relation "$x_1 + \ldots + x_{s-1} = x_s$" in an abelian group. We remark that for $s = 2$ there is a known "quadrangle condition" due to Brandt characterizing those Latin squares that represent the multiplication table of a group, see e.g. [28, Proposition 1.4].

1.4.3. *Reconstructing a group from an abelian s-gon in stable structures (Section 4).* Here we consider a generalization of the group reconstruction method from a fiber-algebraic Q of degree 1 to a fiber-algebraic Q of arbitrary degree, which moreover only satisfies (P2) *generically*, and restricting to Q definable in a stable structure.

Working in a stable theory, it is convenient to formulate this in the language of generic points. By an *s-gon* over a set of parameters A we mean a tuple $a_1, \ldots, a_s$ such that any $s-1$ of its elements are (forking-) independent over A, and any element in it is in the algebraic closure of the other ones and A. We say that an s-gon is *abelian* if, after any permutation of its elements, we have

$$a_1a_2 \underset{\mathrm{acl}_A(a_1a_2)\cap\mathrm{acl}_A(a_3\ldots a_m)}{\downarrow} a_3\ldots a_m.$$

Note that this condition corresponds to the definition of a 1-*based* stable theory, but localized to the elements of the s-gon.

If $(G,+)$ is a type-definable abelian group, $g_1,\ldots,g_{s-1}$ are independent generics in G and $g_s := g_1+\ldots+g_{s-1}$, then $g_1,\ldots,g_s$ is an abelian s-gon (associated to G). In Section 4 we prove a converse:

Theorem (F)**.** *(Theorem 4.6) Let $a_1,\ldots,a_s$ be an abelian s-gon, $s\geq 4$, in a sufficiently saturated stable structure $\mathcal{M}$. Then there is a type-definable (in $\mathcal{M}^{eq}$) connected abelian group $(G,+)$ and an abelian s-gon $g_1,\ldots,g_s$ associated to G, such that after a base change each g_i is inter-algebraic with a_i.*

It is not hard to see that a 4-gon is essentially equivalent to the usual abelian group configuration, so Theorem (F) is a higher arity generalization. After this work was completed, we have learned that independently Hrushovski obtained a similar (but incomparable) unpublished result [31, 32].

1.4.4. *Elekes-Szabó principle in stable structures with distal expansions — proof of Theorem (C.1) (Section 5).* We introduce and study the notion of $\mathfrak{p}$-*dimension* in Section 5.1, imitating the topological definition of dimension in o-minimal structures, but localized at a given tuple of commuting definable global types. Assume we are given $\mathfrak{p}$-*pairs* $(X_i,\mathfrak{p}_i)$ for $1\leq i\leq s$, i.e. X_i is an $\mathcal{M}$-definable set and $\mathfrak{p}_i\in S(\mathcal{M})$ is a complete stationary type on X_i for each $1\leq i\leq s$ (see Definition 5.2). We say that a definable set $Y\subseteq X_1\times\ldots\times X_s$ is $\mathfrak{p}$-*generic*, where $\mathfrak{p}$ refers to the tuple $(\mathfrak{p}_1,\ldots,\mathfrak{p}_s)$, if $Y\in(\mathfrak{p}_1\otimes\ldots\otimes\mathfrak{p}_s)|_{\mathbb{M}}$. Finally, we define the $\mathfrak{p}$-dimension via $\dim_{\mathfrak{p}}(Y)\geq k$ if for some projection π of $\bar{X}$ onto k components, $\pi(Y)$ is $\mathfrak{p}$-generic. We show that $\mathfrak{p}$-dimension enjoys definability and additivity properties crucial for our arguments that may fail for Morley rank in general ω-stable theories such as DCF_0. However, if X is a definable subset of finite Morley rank k and degree one, taking $\mathfrak{p}_X$ to be the unique type on X of Morley rank k, we have that $k\cdot\dim_{\mathfrak{p}}=\mathrm{MR}$ (this is used to deduce Theorem (B) from Theorem (C.1)).

In Section 5.2 we consider the notion of irreducibility and show that every fiber-algebraic relation is a union of finitely many absolutely $\mathfrak{p}$-irreducible sets. In Section 5.3 we consider finite grids in general position with respect to $\mathfrak{p}$-dimension and prove some preliminary power-saving bounds. In Section 5.4 we state a more informative version of Theorem (C.1) (Theorem 5.24 + Theorem 5.27 concerning the bound γ in power saving) and make some preliminary reductions. In particular, we may assume $\dim(Q)=s-1$, and let $\bar{a}=(a_1,\ldots,a_s)$ be a generic tuple in Q. As Q is fiber-algebraic, $\bar{a}$ is an s-gon. We then establish the following key structural dichotomy.

Theorem (G)**.** *(Theorem 5.35 and its proof) Assuming $s \geq 3$, one of the following is true:*

(1) For $u = (a_1, a_2)$ and $v = (a_3, \ldots, a_s)$ we have $u \mathop{\smile\!\!\!\!\!\!\;|}_{\mathrm{acl}(u)\cap\mathrm{acl}(v)} v$.

(2) Q, as a binary relation on $U \times V$, for $U = X_1 \times X_2$ and $V = X_3 \times \ldots \times X_s$, is a "pseudo-plane". By which we mean here that, ignoring a smaller dimensional ($\dim_{\mathfrak{p}} < s-2$) set of $v \in V$, every fiber $Q_v \subseteq U$ has a zero-dimensional intersection $Q_v \cap Q_{v'}$ for all $v' \in V$ outside a smaller dimensional set (more precisely, the $\mathfrak{p}$-dimension of the set Z defined in terms of Q in Section 5.5 is $< s-2$).

This notion of a "pseudo-plane" generalizes the usual definition requiring that any two "lines" in V intersect on finitely many "points" in U, viewing Q as the incidence relation.

In case (2) the relation Q satisfies the assumption on the intersection of its fibers in Definition 2.12, hence the incidence bound from Theorem (D) can be applied inductively to obtain power saving for Q (see Section 5.5). Thus we may assume that for any permutation of $\{1, \ldots, s\}$ we have

$$a_1 a_2 \mathop{\smile\!\!\!\!\!\!\;|}_{\mathrm{acl}(a_1 a_2)\cap\mathrm{acl}(a_3\ldots a_s)} a_3 \ldots a_s,$$

i.e. the s-gon $\bar{a}$ is abelian. Assuming that $s \geq 4$, Theorem (F) can be applied to establish a generic correspondence with a type-definable abelian group (Section 5.6). The case $s = 3$ of Theorem (C.1) is treated separately in Section 5.7 by reducing it to the case $s = 4$ (similar to the approach in [48]).

In Section 5.8 we combine Theorem (C.1) with some standard model theory of algebraically closed fields to deduce Theorem (B) and its higher dimensional version.

1.4.5. *Elekes-Szabó principle in o-minimal structures — proof of Theorem (C.2) (Section 6).* Our proof of the o-minimal case is overall similar to the stable case, but is independent from it. In Section 6.1 we formulate a more informative version of Theorem (C.2) with explicit bounds on power saving (Theorem 6.4) and reduce it to Theorem 6.9 — which is an appropriate analog of Theorem (G): (1) either Q is a "pseudo-plane", or (2) it contains a subset Q^* of full dimension so that the property (P2) from Theorem (E) holds in a neighborhood of every point of Q^*. In Case (1), considered in Section 6.2, we show that Q satisfies the required power saving using Theorem (D) (or rather, its refinement for o-minimal structures from Fact 2.15). In Case (2), we show in Section 6.3 that one can choose a generic tuple $(a_1, \ldots, a_s)$ in Q and (type-definable) infinitesimal neighborhoods μ_i of a_i so that the relation $Q \cap (\mu_1 \times \ldots \times \mu_s)$ satisfies (P1) and (P2) from Theorem (E) — applying it we obtain a generic correspondence with a type-definable abelian group, concluding the proof of Theorem (C.2) for $s \geq 4$. The case $s = 3$ is reduced to $s = 4$ similarly to the stable case.

Finally, in Section 6.4 we obtain a Corollary of Theorem (C.2) that holds in an arbitrary *o*-minimal structure, not necessarily a saturated one - replacing a type-definable group by a definable *local* group (Theorem 6.19). Combining this with the solution of the Hilbert's 5th problem for local groups [27] (in fact, only in the much easier abelian case, see Theorem 8.5 there), we can improve "local group" to a "Lie group" in the case when the underlying set of the *o*-minimal structure $\mathcal{M}$ is $\mathbb{R}$ and deduce Theorem (A) and its higher dimensional analog (Theorem 6.20, see also Remark 6.22). We also observe that for semi-linear relations, in the non-group case we have $(1-\varepsilon)$-power saving for any $\varepsilon > 0$ (Remark 6.24).

1.5. **Acknowledgements.** We are very grateful to the referees for their detailed and insightful reports and many valuable suggestions on improving the paper. We thank Saugata Basu, Martin Bays, Emmanuel Breuillard, Jim Freitag, Rahim Moosa, Tom Scanlon, Pierre Simon, Chieu-Minh Tran and Frank Wagner for some helpful conversations. We thank Institut Henri Poincaré in Paris for its hospitality during the "Model theory, Combinatorics and valued fields" term in the Spring trimester of 2018. Chernikov was partially supported by the NSF CAREER grant DMS-1651321 and by a Simons Fellowship. He thanks Lior Pachter, Michael Kinyon, and Math Twitter for the motivation in the final effort of finishing this paper. Peterzil was partially supported by the Israel Science Foundation grant 290/19. Starchenko was partially supported by the NSF grant DMS-1800806. The results of this paper were announced in [12].

2. Zarankiewicz-type bounds for distal relations

We begin by recalling some of the notions and results about distality and generalized "incidence bounds" for distal relations from [15], and refer to that article for further details. The following definition captures a combinatorial "shadow" of the existence of a nice topological cell decomposition (as e.g. in *o*-minimal theories or in the *p*-adics).

Definition 2.1. [15, Section 2] Let X, Y be infinite sets, and $E \subseteq X \times Y$ a binary relation.

(1) Let $A \subseteq X$. For $b \in Y$, we say that $E_b = \{a \in X : (a,b) \in E\}$ *crosses* A if $E_b \cap A \neq \emptyset$ and $(X \setminus E_b) \cap A \neq \emptyset$.
(2) A set $A \subseteq X$ is *E-complete over* $B \subseteq Y$ if A is not crossed by any E_b with $b \in B$.
(3) A family $\mathcal{F}$ of subsets of X is a *cell decomposition for E over* $B \subseteq Y$ if $X \subseteq \bigcup \mathcal{F}$ and every $A \in \mathcal{F}$ is E-complete over B.
(4) A *cell decomposition for E* is a map $\mathcal{T} : B \mapsto \mathcal{T}(B)$ such that for each finite $B \subseteq Y$, $\mathcal{T}(B)$ is a cell decomposition for E over B.
(5) A cell decomposition $\mathcal{T}$ is *distal* if there exist $k \in \mathbb{N}$ and a relation $D \subseteq X \times Y^k$ such that for all finite $B \subseteq Y$, $\mathcal{T}(B) = \{D_{(b_1,\ldots,b_k)} : b_1, \ldots, b_k \in B \text{ and } D_{(b_1,\ldots,b_k)} \text{ is } E\text{-complete over } B\}$.

(6) For $t \in \mathbb{R}_{>0}$, we say that a cell decomposition $\mathcal{T}$ has *exponent* $\leq t$ if there exists some $c \in \mathbb{R}_{>0}$ such that $|\mathcal{T}(B)| \leq c|B|^t$ for all finite sets $B \subseteq Y$.

Remark 2.2. Note that if $\mathcal{T}$ is a distal cell decomposition, then it has exponent $\leq k$ for k as in Definition 2.1(5).

Remark 2.3. Assume that the binary relation $E \subseteq X \times (Y \times Z)$ admits a distal cell decomposition $\mathcal{T}$ with $|\mathcal{T}(\hat{B})| \leq c|\hat{B}|^t$ for every finite $\hat{B} \subseteq Y \times Z$. Then for every $z \in Z$, the binary relation $E_z \subseteq X \times Y$ admits a distal cell decomposition $\mathcal{T}_z$ with $|\mathcal{T}_z(B)| \leq c|B|^t$ for all finite $B \subseteq Y$.

Proof. Indeed, assume that $D \subseteq X \times (Y \times Z)^k$ is witnessing that $\mathcal{T}$ is distal, i.e. for any finite $\hat{B} \subseteq Y \times Z$ we have

$$\mathcal{T}\left(\hat{B}\right) = \{D_{(b_1,\ldots,b_k)} : b_1, \ldots, b_k \in \hat{B} \text{ and } D_{(b_1,\ldots,b_k)} \text{ is } E\text{-complete over } \hat{B}\}.$$

Fix $z \in Z$, and let

$$D_z := \left\{(x; y_1, \ldots, y_k) \in X \times Y^k : (x; y_1, z, \ldots, y_k, z) \in D\right\} \subseteq X \times Y^k.$$

Given a finite $B \subseteq Y$, we define $\mathcal{T}_z(B)$ as

$$\left\{(D_z)_{(b_1,\ldots,b_k)} : b_1, \ldots, b_k \in B \text{ and } (D_z)_{(b_1,\ldots,b_k)} \text{ is } E_z\text{-complete over } B\right\}.$$

Then $\mathcal{T}_z(B) = \mathcal{T}(B \times \{z\})$, hence $\mathcal{T}_z$ is a distal cell decomposition for E_z and $|\mathcal{T}_z(B)| = |\mathcal{T}(B \times \{z\})| \leq c|B|^t$. □

Existence of "strong honest definitions" established in [18] shows that every relation definable in a distal structure admits a distal cell decomposition (of some exponent).

Fact 2.4. *(see* [15, Fact 2.9]*) Assume that the relation E is definable in a distal structure $\mathcal{M}$. Then E admits a distal cell decomposition (of some exponent $t \in \mathbb{N}$). Moreover, in this case the relation D in Definition 2.1(5) is also definable in $\mathcal{M}$.*

The following definition abstracts from the notion of cuttings in incidence geometry (see the introduction of [15] for an extended discussion).

Definition 2.5. Let X, Y be infinite sets, $E \subseteq X \times Y$. We say that E *admits cuttings with exponent* $t \in \mathbb{R}$ if there is some constant $c \in \mathbb{R}_{>0}$ satisfying the following. For any $B \subseteq Y$ with $|B| = n$ and any $r \in \mathbb{R}$ with $1 < r < n$ there are some sets $X_1, \ldots, X_s \subseteq X$ covering X with $s \leq cr^t$ and such that for each $i \in [s]$ there are at most $\frac{n}{r}$ elements $b \in B$ so that X_i is crossed by E_b.

In the case $r > n$ in Definition 2.5, an r-cutting is equivalent to a distal cell decomposition (sets in the covering are not crossed at all). And for r varying between 1 and n, r-cutting allows to control the trade-off between the number of cells in a covering and the number of times each cell is allowed to be crossed.

Fact 2.6. *(Distal cutting lemma, [15, Theorem 3.2]) Assume $E \subseteq X \times Y$ admits a distal cell decomposition $\mathcal{T}$ of exponent $\leq t$. Then E admits cuttings with exponent $\leq t$ and with the constant coefficient depending only on t and the constant coefficient of $\mathcal{T}$ (the latter is not stated there explicitly, but follows from the proof). Moreover, every set in this cutting is an intersection of at most two cells in $\mathcal{T}$.*

Remark 2.7. We stress that in the Definition 2.5 of an r-cutting, some of the fibers $E_b, b \in B$ might be equal to each other. This is stated correctly on page 2 of the introduction of [15], but is ambiguous in [15, Definition 3.1] (the family $\mathcal{F}$ there is allowed to have repeated sets, so it is a multi-set of sets) and in the statement of [15, Theorem 3.2] (again, the family $\{\varphi(M; a) : a \in H\}$ there should be viewed as a family of sets with repetitions — this is how it is understood in the proof of Theorem 3.2 there).

The next theorem can be viewed as an abstract variant of the Szemerédi-Trotter theorem. It generalizes (and strengthens) the incidence bound due to Elekes and Szabó [23, Theorem 9] to arbitrary graphs admitting a distal cell decomposition, and is crucial to obtain power saving in the non-group case of our main theorem. Our proof below closely follows the proof of [20, Theorem 2.6] (which in turn is a generalization of [25, Theorem 3.2] and [43, Theorem 4]) making the dependence on s explicit. We note that the fact that the bound in Theorem 2.8 is sub-linear in s was first observed in a special case in [53].

As usual, given $d, \nu \in \mathbb{N}$ we say that a bipartite graph $E \subseteq U \times V$ is *$K_{d,\nu}$-free* if it does not contain a copy of the complete bipartite graph $K_{d,\nu}$ with its parts of size d and ν, respectively.

Theorem 2.8. *For every $d, t \in \mathbb{N}_{\geq 2}$ and $c \in \mathbb{R}_{>0}$ there exists some $C = C(d, t, c) \in \mathbb{R}$ satisfying the following.*

Assume that $E \subseteq X \times Y$ admits a distal cell decomposition $\mathcal{T}$ such that $|\mathcal{T}(B)| \leq c|B|^t$ for all finite $B \subseteq Y$. Then, taking $\gamma_1 := \frac{(t-1)d}{td-1}, \gamma_2 := \frac{td-t}{td-1}$ we have: for all $\nu \in \mathbb{N}_{\geq 2}$ and $A \subseteq_m X, B \subseteq_n Y$ such that $E \cap (A \times B)$ is $K_{d,\nu}$-free,

$$|E \cap (A \times B)| \leq C\nu \left(m^{\gamma_1} n^{\gamma_2} + m + n\right).$$

Before giving its proof we recall a couple of weaker general bounds that will be used. First, a classical fact from [36] giving a bound on the number of edges in $K_{d,\nu}$-free graphs without any additional assumptions (see e.g. [11, Chapter VI.2, Theorem 2.2] for the stated version):

Fact 2.9. *Assume $E \subseteq A \times B$ is $K_{d,\nu}$-free, for some $d, \nu \in \mathbb{N}_{\geq 1}$ and A, B finite. Then $|E \cap A \times B| \leq \nu^{\frac{1}{d}}|A||B|^{1-\frac{1}{d}} + d|B|$.*

Given a set Y and a family $\mathcal{F}$ of subsets of Y, the *shatter function* $\pi_{\mathcal{F}} : \mathbb{N} \to \mathbb{N}$ of $\mathcal{F}$ is defined as

$$\pi_{\mathcal{F}}(z) := \max\{|\mathcal{F} \cap B| : B \subseteq Y, |B| = z\},$$

where $\mathcal{F} \cap B = \{S \cap B : S \in \mathcal{F}\}$.

Second, the following bound for graphs of bounded VC-density is only stated in [25] for $K_{d,\nu}$-free graphs with $d = \nu$ (and with the sides of the bipartite graph exchanged), but the more general statement below, as well as the linear dependence of the bound on ν, follow from its proof.

Fact 2.10. [25, Theorem 2.1] *For every $c \in \mathbb{R}$ and $t, d \in \mathbb{N}$ there is some constant $C = C(c, t, d)$ such that the following holds.*

Let $E \subseteq X \times Y$ be a bipartite graph such that the family $\mathcal{F} = \{E_a : a \in X\}$ of subsets of Y satisfies $\pi_{\mathcal{F}}(z) \leq cz^t$ for all $z \in \mathbb{N}$ (where $E_a = \{b \in Y : (a, b) \in E\}$). Then, for any $A \subseteq_m X, B \subseteq_n Y$ so that $E \cap (A \times B)$ is $K_{d,\nu}$-free, we have

$$|E \cap (A \times B)| \leq C\nu(m^{1-\frac{1}{t}}n + m).$$

Remark 2.11. If $E \subseteq X \times Y$ admits a distal cell decomposition $\mathcal{T}$ with $|\mathcal{T}(B)| \leq c|B|^t$ for all $B \subseteq Y$, then for $\mathcal{F} = \{E_a : a \in X\}$ we have $\pi_{\mathcal{F}}(z) \leq cz^t$ for all $z \in \mathbb{N}$.

Indeed, by Definition 2.1, given any finite $B \subseteq Y$ and $\Delta \in \mathcal{T}(B)$, $B \cap E_a = B \cap E_{a'}$ for any $a, a' \in \Delta$ (and the sets in $\mathcal{T}(B)$ give a covering of X), hence at most $|\mathcal{T}(B)|$ different subsets of B are cut out by the fibers of E.

Proof of Theorem 2.8. Let $A \subseteq_m X, B \subseteq_n Y$ so that $E \cap (A \times B)$ is $K_{d,\nu}$-free be given.

If $n \geq m^d$, then by Fact 2.9 we have

$$|E \cap (A \times B)| \leq \nu^{\frac{1}{d}} m n^{1-\frac{1}{d}} + dn \leq d\nu(n^{\frac{1}{d}} n^{1-\frac{1}{d}} + n) = 2d\nu n. \tag{2.1}$$

Hence we assume $n < m^d$ from now on.

Let $r := \frac{m^{\frac{d}{td-1}}}{n^{\frac{1}{td-1}}}$ (note that $r > 1$ as $m^d > n$), and consider the family $\Sigma = (E_b : b \in B)$ of subsets of X (some of the sets in it might be repeated). By assumption and Fact 2.6, there is a family $\mathcal{C}$ of subsets of X giving a $\frac{1}{r}$-cutting for the family Σ. That is, X is covered by the union of the sets in $\mathcal{C}$, any of the sets $C \in \mathcal{C}$ is crossed by at most $|B|/r$ elements from Σ, and $|\mathcal{C}| \leq \alpha_1 r^t$ for some $\alpha_1 = \alpha_1(c, t)$.

Then there is a set $C \in \mathcal{C}$ containing at least $\frac{m}{\alpha_1 r^t} = \frac{n^{\frac{t}{td-1}}}{\alpha_1 m^{\frac{1}{td-1}}}$ points from A. Let $A' \subseteq A \cap C$ be a subset of size exactly $\left\lceil \frac{n^{\frac{t}{td-1}}}{\alpha_1 m^{\frac{1}{td-1}}} \right\rceil$.

If $|A'| \leq d$, we have $\frac{n^{\frac{t}{td-1}}}{\alpha_1 m^{\frac{1}{td-1}}} \leq |A'| \leq d$, so $n \leq d^{\frac{td-1}{t}} \alpha_1^{\frac{td-1}{t}} m^{\frac{1}{t}}$. By assumption, Remark 2.11 and Fact 2.10, for some $\alpha_2 = \alpha_2(c, t, d)$ we have

$$|E \cap (A \times B)| \leq \alpha_2 \nu (nm^{1-\frac{1}{t}} + m) \leq \alpha_2 \nu (d^{\frac{td-1}{t}} \alpha_1^{\frac{td-1}{t}} m^{\frac{1}{t}} m^{1-\frac{1}{t}} + m),$$

hence

$$|E \cap (A \times B)| \leq \alpha_3 \nu m \text{ for some } \alpha_3 = \alpha_3(c, t, d). \tag{2.2}$$

Hence from now on we assume that $|A'| > d$. Let B' be the set of all points $b \in B$ such that E_b crosses C. We know that

$$|B'| \leq \frac{|B|}{r} \leq \frac{nn^{\frac{1}{td-1}}}{m^{\frac{d}{td-1}}} = \frac{n^{\frac{td}{td-1}}}{m^{\frac{d}{td-1}}} \leq \alpha_1^d |A'|^d.$$

Again by Fact 2.9 we get

$$|E \cap (A' \times B')| \leq d\nu(|A'||B'|^{1-\frac{1}{d}} + |B'|)$$
$$\leq d\nu(|A'|\alpha_1^{d-1}|A'|^{d-1} + \alpha_1^d |A'|^d) \leq \alpha_4 \nu |A'|^d$$

for some $\alpha_4 = \alpha_4(c, t, d)$. Hence there is a point $a \in A'$ such that $|E_a \cap B'| \leq \alpha_4 \nu |A'|^{d-1}$.

Since $E \cap (A \times B)$ is $K_{d,\nu}$-free, there are at most $\nu - 1$ points in $B \setminus B'$ from E_a (otherwise, since none of the sets $E_b, b \in B \setminus B'$ crosses C and C contains A', which is of size $\geq d$, we would have a copy of $K_{d,\nu}$). And we have $|A'| \leq \frac{n^{\frac{t}{td-1}}}{\alpha_1 m^{\frac{1}{td-1}}} + 1 \leq \frac{2}{\alpha_1} \frac{n^{\frac{t}{td-1}}}{m^{\frac{1}{td-1}}}$ as $|A|' > d \geq 1$. Hence

$$|E_a \cap B| \leq |E_a \cap B'| + |E_a \cap (B \setminus B')| \leq \alpha_4 \nu |A'|^{d-1} + (\nu - 1)$$
$$\leq \frac{\alpha_4 2^{d-1}}{\alpha_1^{d-1}} \nu \frac{n^{\frac{t(d-1)}{td-1}}}{m^{\frac{d-1}{td-1}}} + (\nu - 1) \leq \alpha_5 \nu \frac{n^{\frac{t(d-1)}{td-1}}}{m^{\frac{d-1}{td-1}}} + (\nu - 1)$$

for some $\alpha_5 := \alpha_5(c, t, d)$. We remove a and repeat the argument until (2.1) or (2.2) applies. This shows:

$$|E \cap (A \times B)| \leq (2d\nu + \alpha_3 \nu)(n + m) + \sum_{i=n^{\frac{1}{d}}}^{m} \left(\alpha_5 \nu \frac{n^{\frac{t(d-1)}{td-1}}}{i^{\frac{d-1}{td-1}}} + (\nu - 1) \right)$$
$$\leq (2d + \alpha_3)\nu(n + m) + \alpha_5 \nu n^{\frac{t(d-1)}{td-1}} \sum_{i=n^{\frac{1}{d}}}^{m} \frac{1}{i^{\frac{d-1}{td-1}}} + (\nu - 1)m.$$

Note that

$$\sum_{i=n^{\frac{1}{d}}}^{m} \frac{1}{i^{\frac{d-1}{td-1}}} \leq \int_{n^{\frac{1}{d}}-1}^{m} \frac{dx}{x^{\frac{d-1}{td-1}}} = \frac{m^{1-\frac{d-1}{td-1}}}{1 - \frac{d-1}{td-1}} - \frac{\left(n^{\frac{1}{d}} - 1\right)^{1-\frac{d-1}{td-1}}}{1 - \frac{d-1}{td-1}}$$
$$\leq \frac{td-1}{(t-1)d} m^{1-\frac{d-1}{td-1}}$$

using $d, t \geq 2$ and that the second term is non-negative for $n \geq 1$.

Taking $C := 3\max\{2d + \alpha_3, \frac{td-1}{(t-1)d}\alpha_5\}$ — which only depends on c, t, d — we thus have

$$|E \cap (A \times B)| \leq \frac{C}{3}\nu(n + m) + \frac{C}{3}\nu n^{\frac{t(d-1)}{td-1}} m^{1-\frac{d-1}{td-1}} + \frac{C}{3}\nu m$$
$$\leq C\nu(m^{\frac{(t-1)d}{td-1}} n^{\frac{td-t}{td-1}} + m + n)$$

for all m, n. ☐

For our applications to hypergraphs, we will need to consider a certain iterated variant of the bound in Theorem 2.8.

Definition 2.12. Let $\mathcal{E}$ be a family of subsets of $X \times Y$ and $\gamma \in \mathbb{R}$. We say that $\mathcal{E}$ satisfies the *γ-Szemerédi-Trotter property*, or *γ-ST property*, if for any function $C : \mathbb{N} \to \mathbb{N}_{\geq 1}$ there exists a function $C' : \mathbb{N} \to \mathbb{N}_{\geq 1}$ so that: for every $E \in \mathcal{E}$, $s \in \mathbb{N}_{\geq 4}, \nu \in \mathbb{N}_{\geq 2}, n \in \mathbb{N}$ and $A \subseteq X, B \subseteq Y$ with $|A| \leq n^{s-2}, |B| \leq n^2$, if for every $a \in A$ there are at most $C(\nu)n^{s-4}$ elements $a' \in A$ with $|E_a \cap E_{a'} \cap B| \geq \nu$, then $|E \cap (A \times B)| \leq C'(\nu)n^{(s-1)-\gamma}$.

We say that a relation $E \subseteq X \times Y$ satisfies the *γ-ST property* if the family $\mathcal{E} := \{E\}$ does.

Lemma 2.1. Assume that $\mathcal{E}$ is a family of subsets of $X \times Y$ and $\gamma \in \mathbb{R}$.

(1) Assume that X', Y' are some sets and $f : X \to X', g : Y \to Y'$ are bijections. For $E \in \mathcal{E}$, let $E' := \{(x, y) \in X' \times Y' : (f^{-1}(x), g^{-1}(y)) \in E\}$, and let $\mathcal{E}' := \{E' : E \in \mathcal{E}\}$, a family of subsets of $X' \times Y'$. Then $\mathcal{E}$ satisfies the γ-ST property if and only if $\mathcal{E}'$ satisfies the γ-ST property.
(2) Assume that for some $k, \ell \in \mathbb{N}$ we have $X = \bigsqcup_{i \in [k]} X_i, Y = \bigsqcup_{j \in [\ell]} Y_i$, and let $E_{i,j} := E \cap (X_i \times Y_j)$, $\mathcal{E}_{i,j} := \{E_{i,j} : E \in \mathcal{E}\}$. Assume that each $\mathcal{E}_{i,j}$ satisfies the γ-ST property. Then $\mathcal{E}$ also satisfies the γ-ST property.

Proof. (1) is immediate from the definition. In (2), given $C : \mathbb{N} \to \mathbb{N}_{\geq 1}$, assume $C'_{i,j} : \mathbb{N} \to \mathbb{N}_{\geq 1}$ witnesses that $\mathcal{E}_{i,j}$ satisfies the γ-ST property. Then $C' := \sum_{(i,j) \in [k] \times [\ell]} C'_{i,j}$ witnesses that $\mathcal{E}$ satisfies the γ-ST property. ☐

Lemma 2.13. *Assume that $\mathcal{E} \subseteq \mathcal{P}(X \times Y)$, $\gamma_1, \gamma_2 \in \mathbb{R}_{>0}$ with $\gamma_1, \gamma_2 \leq 1, \gamma_1 + \gamma_2 \geq 1$ and $C_0 : \mathbb{N} \to \mathbb{R}$ satisfy:*

(∗) *for every $E \in \mathcal{E}$, $\nu \in \mathbb{N}_{\geq 2}$ and finite $A \subseteq_m X, B \subseteq_n Y$, if $E \cap (A \times B)$ is $K_{2,\nu}$-free, then $|E \cap (A \times B)| \leq C_0(\nu)(m^{\gamma_1} n^{\gamma_2} + m + n)$.*

Then $\mathcal{E}$ satisfies the γ-ST property with $\gamma := 3 - 2(\gamma_1 + \gamma_2) \leq 1$ and $C'(\nu) := 2C_0(\nu)(C(\nu) + 2)$.

Proof. Given $E \in \mathcal{E}$ and finite sets A, B satisfying the assumption of the γ-ST property, we consider the finite graph with the vertex set A and the edge relation R defined by $aRa' \iff |E_a \cap E_{a'} \cap B| \geq \nu$ for all $a, a' \in A$. By the assumption of the γ-ST property, this graph has degree at most $r := C(\nu)n^{s-4}$, so it is $(r+1)$-colorable by a standard fact in graph theory. For each $i \in [r+1]$, let $A_i \subseteq A$ be the set of vertices corresponding to the ith color. Then the sets A_i give a partition of A, and for each $i \in [r+1]$ the restriction of E to $A_i \times B$ is $K_{2,\nu}$-free.

For any fixed i, applying the assumption on E to $A_i \times B$, we have

$$|E \cap (A_i \times B)| \leq C_0(\nu) \left(|A_i|^{\gamma_1} |B|^{\gamma_2} + |A_i| + |B| \right).$$

Then we have

$$
\begin{aligned}
|E \cap (A \times B)| &\leq \sum_{i \in [r+1]} |E \cap (A_i \times B)| \\
&\leq \sum_{i \in [r+1]} C_0(\nu) \left(|A_i|^{\gamma_1} |B|^{\gamma_2} + |A_i| + |B| \right) \\
&\leq C_0(\nu) \left(\sum_{i \in [r+1]} |A_i|^{\gamma_1} |B|^{\gamma_2} + \sum_{i \in [r+1]} |A_i| + \sum_{i \in [r+1]} |B| \right). \qquad (2.3)
\end{aligned}
$$

For the first sum, applying Hölder's inequality with $p = \frac{1}{\gamma_1}$, we have

$$
\begin{aligned}
\sum_{i \in [r+1]} |A_i|^{\gamma_1} |B|^{\gamma_2} &= |B|^{\gamma_2} \sum_{i \in [r+1]} |A_i|^{\gamma_1} \\
&\leq |B|^{\gamma_2} \left(\sum_{i \in [r+1]} |A_i| \right)^{\gamma_1} \left(\sum_{i \in [r+1]} 1 \right)^{1-\gamma_1} \\
= |B|^{\gamma_2} |A|^{\gamma_1} (r+1)^{1-\gamma_1} &\leq n^{2\gamma_2} n^{(s-2)\gamma_1} \left(C(\nu) n^{s-4} + 1 \right)^{1-\gamma_1} \\
&\leq n^{2\gamma_2} n^{(s-2)\gamma_1} \left(C(\nu) + 1 \right)^{1-\gamma_1} n^{(s-4)(1-\gamma_1)} \\
&\leq (C(\nu) + 1) n^{(s-4) + 2(\gamma_1 + \gamma_2)} = (C(\nu) + 1) n^{(s-1) - \gamma}
\end{aligned}
$$

for all n (by definition of γ and as $s \geq 4, C(\nu) \geq 1, 0 < \gamma_1 \leq 1$).

For the second sum, we have

$$
\sum_{i \in [r+1]} |A_i| = |A| \leq n^{s-2}
$$

for all n. For the third sum we have

$$
\sum_{i \in [r+1]} |B| \leq (r+1)|B| \leq (C(\nu) n^{s-4} + 1) n^2 \leq (C(\nu) + 1) n^{s-2}
$$

for all n. Substituting these bounds into (2.3), as $\gamma \leq 1$ we get

$$
|E \cap (A \times B)| \leq 2C_0(\nu)(C(\nu) + 2) n^{(s-1)-\gamma}. \quad \square
$$

We note that the γ-ST property is non-trivial only if $\gamma > 0$. Lemma 2.13 shows that if $\mathcal{E}$ satisfies the condition in Lemma 2.13(∗) with $\gamma_1 + \gamma_2 < \frac{3}{2}$, then $\mathcal{E}$ satisfies the γ-ST property for some $\gamma > 0$. By Theorem 2.8 this condition on $\gamma_1 + \gamma_2$ is satisfied for any relation admitting a distal cell decomposition, leading to the following.

Proposition 2.14. *(1) Assume that $t \in \mathbb{N}_{\geq 2}$ and $E \subseteq X \times Y$ admits a distal cell decomposition $\mathcal{T}$ such that $|\mathcal{T}(B)| \leq c|B|^t$ for all finite $B \subseteq Y$. Then E satisfies the γ-ST property with $\gamma := \frac{1}{2t-1} > 0$ and $C' : \mathbb{N} \to \mathbb{N}_{\geq 1}$ depending only on t, c, C.*

(2) In particular, if the binary relation $E \subseteq X \times (Y \times Z)$ admits a distal cell decomposition $\mathcal{T}$ of exponent t, then the family of fibers

$$\mathcal{E} := \{E_z \subseteq X \times Y : z \in Z\} \subseteq \mathcal{P}(X \times Y)$$

satisfies the γ-ST property $\gamma := \frac{1}{2t-1}$.

Proof. (1) By assumption and Theorem 2.8 with $d := 2$, there exists some $c' = c'(t, c) \in \mathbb{R}$ such that, taking $\gamma_1 := \frac{2t-2}{2t-1}, \gamma_2 := \frac{t}{2t-1}$, for all $\nu \in \mathbb{N}_{\geq 2}, m, n \in \mathbb{N}$ and $A \subseteq_m X, B \subseteq_n Y$ with $E \cap (A \times B)$ is $K_{2,\nu}$-free we have

$$|E \cap (A \times B)| \leq c'\nu \left(m^{\gamma_1} n^{\gamma_2} + m + n\right).$$

Then, by Lemma 2.13, E satisfies the γ-ST property $\gamma := 3 - 2(\gamma_1 + \gamma_2) = 3 - 2\frac{3t-2}{2t-1} = \frac{1}{2t-1} > 0$ and $C'(\nu) := 2c'\nu(C(\nu) + 2)$.

(2) Combining (1) and Remark 2.3. □

The γ in Proposition 2.14 will correspond to the power saving in the main theorem. Stronger upper bounds on γ_1, γ_2 in Lemma 2.13(∗) (than the ones given by Theorem 2.8) are known in some particular distal structures of interest and can be used to improve the bound on γ in Proposition 2.14, and hence in the main theorem. We summarize some of these results relevant for our applications.

Fact 2.15. *Let $\mathcal{M} = (M, <, \ldots)$ be an o-minimal expansion of a group.*

(1) Let $\mathcal{E}$ be a definable family of subsets of $M^2 \times M^{d_2}, d_2 \in \mathbb{N}$, i.e. $\mathcal{E} = \{E_b : b \in Z\}$ for some $d_3 \in \mathbb{N}$ and definable sets $E \subseteq M^2 \times M^{d_2} \times M^{d_3}, Z \subseteq M^{d_3}$. The definable relation E viewed as a binary relation on $M^2 \times M^{d_2+d_3}$ admits a distal cell decomposition with exponent $t = 2$ by [15, Theorem 4.1]*. Then Proposition 2.14(2) implies that $\mathcal{E}$ satisfies the γ-ST property with $\gamma := \frac{1}{3}$. (See also* [6] *for an alternative approach.)*

(2) For general $d_1, d_2 \in \mathbb{N}_{\geq 2}$, every definable relation $E \subseteq M^{d_1} \times M^{d_2+d_3}$ admits a distal cell decomposition with exponent $t = 2d_1 - 2$ by [1] *(this improves on the weaker bound in* [3, Section 4] *and generalizes the semialgebraic case in* [14]*). As in (1), Proposition 2.14(2) implies that any definable family $\mathcal{E}$ of subsets of $M^{d_1} \times M^{d_2}$ satisfies the γ-ST property with $\gamma := \frac{1}{4d_1-5}$.*

In particular this implies the following bounds for semialgebraic and constructible sets of bounded description complexity:

Corollary 2.16. *(1) If $d_1, d_2, D \in \mathbb{N}_{\geq 2}$, and $\mathcal{E}_D$ is the family of semialgebraic subsets of $\mathbb{R}^{d_1} \times \mathbb{R}^{d_2}$ of description complexity D (i.e. every $E \in \mathcal{E}$ is defined by a Boolean combination of at most D polynomial (in-)equalities with real coefficients, with all polynomials of degree at most D), then $\mathcal{E}_D$ satisfies the γ-ST property with $\gamma := \frac{1}{4d_1-5}$ (noting that for a fixed D, the family $\mathcal{E}_D$ is definable in the o-minimal structure $(\mathbb{R}, +, \times, <)$ and using Fact 2.15(2)).*

(2) If $d_1, d_2, D \in \mathbb{N}_{\geq 2}$ and $\mathcal{E}_D$ is the family of constructible subsets of $\mathbb{C}^{d_1} \times \mathbb{C}^{d_2}$ of description complexity D (i.e. every $E \in \mathcal{E}$ is defined by a Boolean combination of at most D polynomial equations with complex coefficients, with all polynomials of degree at most D), then $\mathcal{E}_D$ satisfies the γ-ST property with $\gamma := \frac{1}{8d_1-5}$ (noting that for a fixed D, every $E \in \mathcal{E}_D$ can be viewed as a constructible, and hence semialgebraic, subset of $\mathbb{R}^{2d_1} \times \mathbb{R}^{2d_2}$ of description complexity D, and using (1)).

We note that a stronger bound is known for algebraic sets over $\mathbb{R}$ and $\mathbb{C}$, however in the proof of the main theorem over $\mathbb{C}$ we require a bound for more general families of constructible sets:

Fact 2.17. *(1) ([25, Theorem 1.2], [60, Corollary 1.7]) If $d_1, d_2 \in \mathbb{N}_{\geq 2}$ and $E \subseteq \mathbb{R}^{d_1} \times \mathbb{R}^{d_2}$ is algebraic with each $E_b, b \in \mathbb{R}^{d_2}$ an algebraic variety of degree D in $\mathbb{R}^{d_1}$, then E satisfies the condition in Lemma 2.13$(*)$ with $\gamma_1 = \frac{2(d_1-1)}{2d_1-1}, \gamma_2 = \frac{d_1}{2d_1-1}$ and some function C_0 depending on d_2, D. Hence, by Lemma 2.13, E satisfies the γ-ST property with $\gamma := \frac{1}{2d_1-1}$.*

(2) If $d_1, d_2 \in \mathbb{N}_{\geq 2}$ and $E \subseteq \mathbb{C}^{d_1} \times \mathbb{C}^{d_2}$ is algebraic with each $E_b, b \in \mathbb{C}^{d_2}$ an algebraic variety of degree D, it can be viewed as an algebraic subset of $\mathbb{R}^{2d_1} \times \mathbb{R}^{2d_2}$ with all fibers algebraic varieties of fixed degree, which implies by (1) that E satisfies the γ-ST property with $\gamma := \frac{1}{4d_1-1}$. (This improves the bound in [23, Theorem 9].)

Problem 2.18. *We expect that the same bound on γ as in Fact 2.17(2) should hold for an arbitrary constructible family $\mathcal{E}_D$ over $\mathbb{C}$ in Corollary 2.16(2), and the same bound on γ as in Fact 2.17(1) should hold for an arbitrary definable family $\mathcal{E}$ in an o-minimal structure in Fact 2.15(2). However, the polynomial method used to obtain these stronger bounds for high dimensions in the algebraic case does not immediately generalize to constructible sets, and is not available for general o-minimal structures (see [5]).*

Fact 2.19. *Assume that $d_1, d_2, s \in \mathbb{N}$ and $\mathcal{E}$ is a family of semilinear subsets of $\mathbb{R}^{d_1} \times \mathbb{R}^{d_2}$ so that each $E \in \mathcal{E}$ is defined by a Boolean combination of s linear equalities and inequalities (with real coefficients). Then by [4, Theorem (C)], for every $\varepsilon \in \mathbb{R}_{>0}$ the family $\mathcal{E}$ satisfies the condition in Lemma 2.13$(*)$ with $\gamma_1 + \gamma_2 \leq 1 + \varepsilon$ (and some function C_0 depending on s and ε). It follows that $\mathcal{E}$ satisfies the $(1-\varepsilon)$-ST property for every $\varepsilon > 0$ (which is the best possible bound up to ε).*

Fact 2.20. *It has been shown in [2] that every differentially closed field (with one or several commuting derivations) of characteristic 0 admits a distal expansion. Hence by Fact 2.4, every definable relation admits a distal cell decomposition of some finite exponent t, hence by Proposition 2.14(2) any definable family $\mathcal{E}$ of subsets of $\subseteq M^{d_1} \times M^{d_2}$ in a differentially closed field $\mathcal{M}$ of characteristic 0 satisfies the γ-ST property for some $\gamma > 0$.*

Fact 2.21. *The theory of compact complex manifolds, or CCM, is the theory of the structure containing a separate sort for each compact complex variety,* *with each Zariski closed subset of the cartesian products of the sorts named by a predicate (see [41] for a survey). This is an ω-stable theory of finite Morley rank, and it is interpretable in the o-minimal structure $\mathbb{R}_{an}$. Hence, by Fact 2.4 and Proposition 2.14(2), every definable family $\mathcal{E}$ admits a distal cell decomposition of some finite exponent t, and hence satisfies the γ-ST property for some $\gamma > 0$.*

We remark that in differentially closed fields it is not possible to bound t in terms of d_1 alone. Indeed, the dp-rank of the formula "$x = x$" is $\geq n$ for all $n \in \mathbb{N}$ (since the field of constants is definable, and M is an infinite dimensional vector space over it, see [17, Remark 5.3]). This implies that the VC-density of a definable relation $E \subseteq M \times M^n$ cannot be bounded independent of n (see e.g. [34]), and since t gives an upper bound on the VC-density (see Remark 2.11), it cannot be bounded either.

Problem 2.22. *Obtain explicit bounds on the distal cell decomposition and incidence counting for relations E definable in DCF_0 (e.g., are they bounded in terms of the Morley rank of the relation E?).*

3. Reconstructing an abelian group from a family of bijections

In this and the following sections we provide two higher arity variants of the group configuration theorem of Zilber-Hrushovski (see e.g [46, Chapter 5.4]). From a model-theoretic point of view, our result can be viewed as a construction of a type-definable abelian group in the non-trivial *local* locally modular case, i.e. local modularity is only assumed for the given relation, as opposed to the whole theory, based on a relation of arbitrary arity ≥ 4.

In this section, as a warm-up, we begin with a purely combinatorial abelian group configuration for the case of bijections as opposed to finite-to-finite correspondences. It illustrates some of the main ideas and is sufficient for the application in the *o*-minimal case of the main theorem in Section 6.

In the next Section 4, we generalize the construction to allow finite-to-finite correspondences instead of bijections (model-theoretically, algebraic closure instead of the definable closure) in the stable case, which requires additional forking calculus arguments.

3.1. Q-relations or arity 4.

Throughout this section, we fix some sets A, B, C, D and a quaternary relation $Q \subseteq A \times B \times C \times D$. We assume that Q satisfies the following two properties.

(P1) If we fix any 3 variables, then there is exactly one value for the 4th variable satisfying Q.

(P2) If
$$(\alpha, \beta; \gamma, \delta), (\alpha', \beta'; \gamma, \delta), (\alpha', \beta'; \gamma', \delta') \in Q,$$
then
$$(\alpha, \beta; \gamma', \delta') \in Q;$$
and the same is true under any other partition of the variables into two groups each of size two.

Intuitively, the first condition says that Q induces a family of bijective functions between any two of its coordinates, and the second condition says that this family of bijections satisfies the "abelian group configuration" condition in a strong sense. Our goal is to show that under these assumption there exist an abelian group for which Q is in a coordinate-wise bijective correspondence with the relation defined by $\alpha \cdot \beta = \gamma \cdot \delta$.

First, we can view the relation Q as a 2-parametric family of bijections as follows. Note that for every pair $(c, d) \in C \times D$, the corresponding fiber $\{(a, b) \in A \times B : (a, b, c, d) \in Q\}$ is the graph of a function from A to B by (P1). Let $\mathcal{F}$ be the set of all functions from A to B whose graph is a fiber of Q.

Similarly, let $\mathcal{G}$ be the set of all functions from C to D whose graph is a fiber of Q (for some $(a, b) \in A \times B$). Note that all functions in $\mathcal{F}$ and in $\mathcal{G}$ are bijections, again by (P1).

Claim 3.1. *For every $(a, b) \in A \times B$ there is a unique $f \in \mathcal{F}$ with $f(a) = b$, and similarly for $\mathcal{G}$.*

Proof. We only check this for $\mathcal{F}$, the argument for $\mathcal{G}$ is analogous. Let $(a, b) \in A \times B$ be fixed. Existence: let $c \in C$ be arbitrary, then by (P1) there exists some $d \in D$ with $(a, b, c, d) \in Q$, hence the function corresponding to the fiber of Q at (c, d) satisfies the requirement. Uniqueness follows from (P2) for the appropriate partition of the variables: if $(a, b; c, d), (a, b; c_1, d_1) \in Q$ for some $(c, d), (c_1, d_1) \in C \times D$, then for all $(x, y) \in A \times B$ we have $(x, y, c, d) \in Q \iff (x, y; c_1, d_1) \in Q$. □

Claim 3.2. *For every $f \in \mathcal{F}$ and (x, u) in $A \times C$ there exists a unique $g \in \mathcal{G}$ such that $(x, f(x), u, g(u)) \in Q$ (which then satisfies $(x', f(x'), u', g(u')) \in Q$ for all $(x', u') \in A \times C$).*

And similarly exchanging the roles of $\mathcal{F}$ and $\mathcal{G}$.

Proof. As $x, f(x), u$ are given, by (P1) there is a unique choice for the fourth coordinate of a tuple in Q determining the image of g on u. There is only one such $g \in \mathcal{G}$ by Claim 3.1 with respect to $\mathcal{G}$. □

For $f \in \mathcal{F}$, we will denote by $f^{\perp}$ the unique $g \in \mathcal{G}$ as in Claim 3.2. Similarly, for $g \in \mathcal{G}$, we will denote by $g^{\perp}$ the unique $f \in \mathcal{F}$ as in Claim 3.2.

Remark 3.3. Note that $(f^{\perp})^{\perp} = f$ and $(g^{\perp})^{\perp} = g$ for all $f \in \mathcal{F}, g \in \mathcal{G}$.

Claim 3.4. *Let $f_1, f_2, f_3 \in \mathcal{F}$, and $g_i := f_i^{\perp} \in \mathcal{G}$ for $i \in [3]$. Then $f_3 \circ f_2^{-1} \circ f_1 \in \mathcal{F}$, $g_3 \circ g_2^{-1} \circ g_1 \in \mathcal{G}$ and $(f_3 \circ f_2^{-1} \circ f_1)^{\perp} = g_3 \circ g_2^{-1} \circ g_1$.*

Proof. We first observe the following: given any $a \in A$ and $c \in C$, if we take $b := (f_3 \circ f_2^{-1} \circ f_1)(a) \in B$ and $d := (g_3 \circ g_2^{-1} \circ g_1)(c) \in D$, then $(a, b, c, d) \in Q$. Indeed, let $b_1 := f_1(a)$, $a_2 := f_2^{-1}(b_1)$, then $b = f_3(a_2)$. Similarly, let $d_1 := g_1(c)$, $c_2 := g_2^{-1}(d_1)$, then $d = g_3(c_2)$. By the definition of $\perp$ we then have

$$(a, b_1, c, d_1) \in Q, (a_2, b_1, c_2, d_1) \in Q, (a_2, b, c_2, d) \in Q.$$

Applying (P2) for the partition $\{1,3\} \cup \{2,4\}$, this implies $(a,b,c,d) \in Q$, as wanted.

Now fix an arbitrary $c \in C$ and take the corresponding d, varying $a \in A$ the observation implies that the graph of $f_3 \circ f_2^{-1} \circ f_1$ is given by the fiber $Q_{(c,d)}$. Similarly, the graph of $g_3 \circ g_2^{-1} \circ g_1$ is given by the fiber $Q_{(a,b)}$ for an arbitrary $a \in A$ and the corresponding b; and $(f_3 \circ f_2^{-1} \circ f_1)^{\perp} = g_3 \circ g_2^{-1} \circ g_1$ follows. □

Claim 3.5. *For any $f_1, f_2, f_3 \in \mathcal{F}$ we have $f_3 \circ f_2^{-1} \circ f_1 = f_1 \circ f_2^{-1} \circ f_3$, and similarly for $\mathcal{G}$.*

Proof. Let $a \in A$ be arbitrary. We define $b_1 := f_1(a)$, $a_2 := f_2^{-1}(b_1)$ and $b_3 := f_3(a_2)$, so we have $(f_3 \circ f_2^{-1} \circ f_1)(a) = b_3$. Let also $b_4 := f_3(a)$, $a_4 := f_2^{-1}(b_4)$ and $b_5 := f_1(a_4)$, so we have $(f_1 \circ f_2^{-1} \circ f_3)(a) = b_5$.

We need to show that $b_5 = b_3$.

Let $c_1 \in C$ be arbitrary. By (P1) there exists some $d_1 \in D$ such that

$$(a, b_1, c_1, d_1) \in Q. \tag{3.1}$$

Applying (P1) again, there exists some $c_2 \in C$ such that

$$(a_2, b_1, c_2, d_1) \in Q, \tag{3.2}$$

and then some $d_2 \in D$ such that

$$(a_2, b_3, c_2, d_2) \in Q. \tag{3.3}$$

Using (P2) for the partition $\{1,3\} \cup \{2,4\}$, it follows from (3.1), (3.2), (3.3) that

$$(a, b_3, c_1, d_2) \in Q. \tag{3.4}$$

On the other hand, by the choice of b_1, a_2, b_3, (3.1), (3.2), (3.3) and Claim 3.1 we have: $Q_{(c_1,d_1)}$ is the graph of f_1, $Q_{(c_2,d_1)}$ is the graph of f_2 and $Q(c_2, d_2)$ is the graph of f_3. Hence we also have

$$(a, b_4, c_2, d_2) \in Q, (a_4, b_4, c_2, d_1) \in Q, (a_4, b_5, c_1, d_1) \in Q.$$

Applying (P2) for the partition $\{1,4\} \cup \{2,3\}$ this implies

$$(a, b_5, c_1, d_2) \in Q,$$

and combining with (3.4) and (P1) we obtain $b_3 = b_5$. □

Claim 3.6. *Given an arbitrary element $f_0 \in \mathcal{F}$, for every pair $f, f' \in \mathcal{F}$ we define*

$$f + f' := f \circ f_0^{-1} \circ f'.$$

Then $(\mathcal{F}, +)$ is an abelian group, with the identity element f_0.

Proof. Note that for every $f, f' \in \mathcal{F}$, $f + f' \in \mathcal{F}$ by Claim 3.4. Associativity follows from the associativity of the composition of functions. For any $f \in \mathcal{F}$ we have $f + f_0 = f \circ f_0^{-1} \circ f_0 = f$, $f_0 \circ f^{-1} \circ f_0 \in \mathcal{F}$ by Claim 3.4 and $f + (f_0 \circ f^{-1} \circ f_0) = f \circ f_0^{-1} \circ (f_0 \circ f^{-1} \circ f_0) = f_0$, hence f_0 is the right identity and $f_0 \circ f^{-1} \circ f_0$ is the right inverse of f. Finally, by Claim 3.5 we have $f + f' = f' + f$ for any $f, f' \in \mathcal{F}$, hence $(\mathcal{F}, +)$ is an abelian group. □

Remark 3.7. Moreover, if we also fix $g_0 := f_0^{\perp}$ in $\mathcal{G}$, then similarly we obtain an abelian group on $\mathcal{G}$ with the identity element g_0, so that $(\mathcal{F}, +)$ is isomorphic to $(\mathcal{G}, +)$ via the map $f \mapsto f^{\perp}$ (it is a homomorphism as for any $f_1, f_2 \in \mathcal{F}$ we have $(f_1 \circ f_0^{-1} \circ f_2)^{\perp} = f_1^{\perp} \circ g_0^{-1} \circ f_2^{\perp}$ by Claim 3.4, and its inverse is $g \in \mathcal{G} \mapsto g^{\perp}$ by Remark 3.3).

Next we establish a connection of these groups and the relation Q. We fix arbitrary $a_0 \in A$, $b_0 \in B$, $c_0 \in C$ and $d_0 \in D$ with $(a_0, b_0, c_0, d_0) \in Q$. By Claim 3.1, let $f_0 \in \mathcal{F}$ be unique with $f_0(a_0) = b_0$, and let $g_0 \in \mathcal{G}$ be unique with $g_0(c_0) = d_0$. Then $g_0 = f_0^{\perp}$ by Claim 3.2, and by Remark 3.7 we have isomorphic groups on $\mathcal{F}$ and on $\mathcal{G}$. We will denote this common group by $G := (\mathcal{F}, +)$.

We consider the following bijections between each of A, B, C, D and G, using our identification of G with both $\mathcal{F}$ and $\mathcal{G}$ and Claim 3.1:

- let $\pi_A : A \to \mathcal{F}$ be the bijection that assigns to $a \in A$ the unique $f_a \in \mathcal{F}$ with $f_a(a) = b_0$;
- let $\pi_B : B \to \mathcal{F}$ be the bijection that assigns to $b \in B$ the unique $f_b \in \mathcal{F}$ with $f_b(a_0) = b$;
- let $\pi_C : C \to \mathcal{G}$ be the bijection that assigns to $c \in C$ the unique $g_c \in \mathcal{G}$ with $g_c(c) = d_0$;
- let $\pi_D : D \to \mathcal{G}$ be the bijection that assigns to $d \in D$ the unique $g_d \in \mathcal{G}$ with $g_d(c_0) = d$.

Claim 3.8. *For any $a \in A$ and $b \in B$, $\pi_A(a) + \pi_B(b)$ is the unique function $f \in \mathcal{F}$ with $f(a) = b$.*

Similarly, for any $c \in C$ and $d \in D$, $\pi_C(a) + \pi_D(b)$ is the unique function $g \in \mathcal{G}$ with $g(c) = d$.

Proof. Let $(a, b) \in A \times B$ be arbitrary, and let $f := \pi_A(a) + \pi_B(b) = \pi_B(b) + \pi_A(a) = \pi_B(b) \circ f_0^{-1} \circ \pi_A(a)$. Note that, from the definitions, $\pi_A(a) : a \mapsto b_0$, $f_0^{-1} : b_0 \mapsto a_0$ and $\pi_B(b) : a_0 \mapsto b$, hence $f(a) = b$. The second claim is analogous. □

Proposition 3.9. *For any $(a, b, c, d) \in A \times B \times C \times D$, $(a, b, c, d) \in Q$ if and only if $\pi_A(a) + \pi_B(b) = \pi_C(c)^{\perp} + \pi_D(d)^{\perp}$ (in G).*

Proof. Given (a, b, c, d), by Claim 3.8 we have: $\pi_A(a) + \pi_B(b)$ is the function $f \in \mathcal{F}$ with $f(a) = b$, and $\pi_C(c) + \pi_D(d)$ is the function $g \in \mathcal{G}$ with $g(c) = d$. Then, by Claim 3.2, $(a, b, c, d) \in Q$ if and only if $f = g^{\perp}$, and since $\perp$ is an isomorphism this happens if and only $f = \pi_C(c)^{\perp} + \pi_D(d)^{\perp}$. □

3.2. **Q-relation of any arity for dcl.** Now we extend the construction of an abelian group to relations of arbitrary arity ≥ 4. Assume that we are given $m \in \mathbb{N}_{\geq 4}$, sets $X_1, \ldots, X_m$ and a relation $Q \subseteq X_1 \times \cdots \times X_m$ satisfying the following two conditions (corresponding to the conditions in Section 3.1 when $m = 4$).

(P1) For any permutation of the variables of Q we have:

$$\forall x_1, \ldots, \forall x_{m-1} \exists! x_m Q(x_1, \ldots, x_m).$$

(P2) For any permutation of the variables of Q we have:

$$\forall x_1, x_2 \forall y_3, \ldots y_m \forall y'_3, \ldots, y'_m \Big(Q(\bar{x}, \bar{y}) \wedge Q(\bar{x}, \bar{y}') \rightarrow \big(\forall x'_1, x'_2 Q(\bar{x}', \bar{y}) \leftrightarrow Q(\bar{x}', \bar{y}') \big) \Big),$$

where $\bar{x} = (x_1, x_2), \bar{y} = (y_3, \ldots, y_m)$, $Q(\bar{x}, \bar{y})$ evaluates Q on the concatenated tuple $(x_1, x_2, y_3, \ldots, y_m)$, and similarly for $\bar{x}', \bar{y}'$.

We let $\mathcal{F}$ be the set of all functions $f : X_1 \to X_2$ whose graph is given by the set of pairs $(x_1, x_2) \in X_1 \times X_2$ satisfying $Q(x_1, x_2, \bar{b})$ for some $\bar{b} \in X_3 \times \ldots \times X_m$.

Remark 3.10. (1) Every $f \in \mathcal{F}$ is a bijection, by (P1).

(2) For every $a_1 \in X_1, a_2 \in X_2$ there exists a unique $f \in \mathcal{F}$ such that $f(a_1) = a_2$ (existence by (P1), uniqueness by (P2)). We will denote it as f_{a_1,a_2}.

Lemma 3.11. *For every $c_i \in X_i, 4 \leq i \leq m$ and $f \in \mathcal{F}$ there exists some $c_3 \in X_3$ such that $Q(x_1, x_2, c_3, c_4, \ldots, c_m)$ is the graph of f.*

Proof. Choose any $a_1 \in X_1$, let $a_2 := f(a_1)$. Choose $c_3 \in X_3$ such that $Q(a_1, a_2, c_3, \ldots, c_m)$ holds by (P1). Then $Q(x_1, x_2, c_3, c_4, \ldots, c_m)$ defines the graph of f by Remark 3.10(2). □

Lemma 3.12. *For any $f_1, f_2, f_3 \in \mathcal{F}$ there exists some $f_4 \in \mathcal{F}$ such that $f_1 \circ f_2^{-1} \circ f_3 = f_3 \circ f_2^{-1} \circ f_1 = f_4$.*

Proof. Choose any $c_i \in X_i, 5 \leq i \leq m$ and consider the quaternary relation $Q' \subseteq X_1 \times \cdots \times X_4$ defined by $Q'(x_1, \ldots, x_4) := Q(x_1, \ldots, x_4, \bar{c})$. Hence Q' also satisfies (P1) and (P2), and the graph of every $f \in \mathcal{F}$ is given by $Q'(x_1, x_2, b_3, b_4)$ for some $b_3 \in X_3, b_4 \in X_4$, by Lemma 3.11. Then the conclusion of the lemma follows from Claims 3.4 and 3.5 applied to Q'. □

Definition 3.13. We fix arbitrary elements $e_i \in X_i, i = 1, \ldots, m$ so that $Q(e_1, \ldots, e_m)$ holds. Let $f_0 \in \mathcal{F}$ be the function whose graph is given by $Q(x_1, x_2, e_3, \ldots, e_m)$, i.e. $f_0 = f_{e_1,e_2}$. We define $+ : \mathcal{F} \times \mathcal{F} \to \mathcal{F}$ by taking $f_1 + f_2 := f_1 \circ f_0^{-1} \circ f_2$.

As in Claim 3.6, from Lemma 3.12 we get:

Lemma 3.14. *$G := (\mathcal{F}, +)$ is an abelian group with the identity element f_0.*

Definition 3.15. We define the map $\pi_1 : X_1 \to G$ by $\pi_1(a) := f_{a,e_2}$ for all $a \in X_1$, and the map $\pi_2 : X_2 \to G$ by $\pi_2(b) := f_{e_1,b}$ for all $b \in X_2$.

Note that both π_1 and π_2 are bijections by Remark 3.10.

Lemma 3.16. *For any $a \in X_1$ and $b \in X_2$ we have $\pi_1(a) + \pi_2(b) = f_{a,b}$.*

Proof. Let $f_1 := \pi_1(a), f_2 := \pi_2(b)$. Note that $f_1(a) = e_2$, $f_0^{-1}(e_2) = e_1$ and $f_2(e_1) = b$, hence $(f_1 + f_2)(a) = (f_2 + f_1)(a) = f_2 \circ f_0^{-1} \circ f_1(a) = b$, so $f_1 + f_2 = f_{a,b}$. □

Definition 3.17. For any set $S \subseteq \{3, \ldots, m\}$, we define the map $\pi_S : \prod_{i\in S} X_i \to G$ as follows: for $\bar{a} = (a_i : i \in S) \in \prod_{i\in S} X_i$, let $\pi_S(\bar{a})$ be the function in $\mathcal{F}$ whose graph is given by $Q(x_1, x_2, c_3, \ldots, c_m)$ with $c_j := a_j$ for $j \in S$ and $c_j := e_j$ for $j \notin S$. We write π_j for $\pi_{\{j\}}$.

Remark 3.18. For each $i \in \{3, \ldots, m\}$, the map $\pi_i : X_i \to G$ is a bijection (by (P2)).

Lemma 3.19. *Fix some $S \subsetneq \{3, \ldots, m\}$ and $j_0 \in \{3, \ldots, m\} \setminus S$. Let $S_0 := S \cup \{j_0\}$. Then for any $\bar{a} \in \prod_{i\in S} X_i$ and $a_{j_0} \in X_{j_0}$ we have $\pi_S(\bar{a}) + \pi_{j_0}(a_{j_0}) = \pi_{S_0}(\bar{a}^\frown a_{j_0})$.*

Proof. Without loss of generality we have $S = \{3, \ldots, k\}$ and $j_0 = k+1 \le m$ for some k. Take any $\bar{a} = (a_3, \ldots, a_k) \in \prod_{3\le i\le k} X_i$ and $a_{k+1} \in X_{k+1}$. Then, from the definitions:

- the graph of $f_1 := \pi_S(\bar{a})$ is given by $Q(x_1, x_2, a_3, \ldots, a_k, e_{k+1}, \bar{e}')$, where $\bar{e}' := (e_{k+2}, \ldots, e_m)$;
- the graph of $f_2 := \pi_{k+1}(a_{k+1})$ is given by $Q(x_1, x_2, e_3, \ldots, e_k, a_{k+1}, \bar{e}')$;
- the graph of $f_3 := \pi_{S_0}(\bar{a}^\frown a_{k+1})$ is given by $Q(x_1, x_2, a_3, \ldots, a_k, a_{k+1}, \bar{e}')$.

Let $c_1 \in X_1$ be such that $f_1(c_1) = e_2$ and let $c_2 \in X_2$ be such that $f_2(e_1) = c_2$. Then $(f_1 + f_2)(c_1) = (f_2 + f_1)(c_1) = f_2 \circ f_0^{-1} \circ f_1(c_1) = c_2$. On the other hand, the following also hold:

- $Q(c_1, e_2, a_3, \ldots, a_k, e_{k+1}, \bar{e}')$;
- $Q(e_1, e_2, e_3, \ldots, e_k, e_{k+1}, \bar{e}')$;
- $Q(e_1, c_2, e_3, \ldots, e_k, a_{k+1}, \bar{e}')$.

Applying (P2) with respect to the coordinates $(2, k+1)$ and the rest, this implies that $Q(c_1, c_2, a_3, \ldots, a_k, a_{k+1}, \bar{e}')$ holds, i.e. $f_3(c_1) = c_2$. Hence $f_1 + f_2 = f_3$ by Remark 3.10(2), as wanted.

□

Proposition 3.20. *For any $\bar{a} = (a_1, \ldots, a_m) \in \prod_{i\in[m]} X_i$ we have*

$$Q(a_1, \ldots, a_m) \iff \pi_1(a_1) + \pi_2(a_2) = \pi_3(a_3) + \ldots + \pi_m(a_m).$$

Proof. Let $\bar{a} = (a_1, \ldots, a_m) \in \prod_{i\in[m]} X_i$ be arbitrary. By Lemma 3.16, $\pi_1(a_1) + \pi_2(a_2) = f_{a_1,a_2}$. Applying Lemma 3.19 inductively, we have

$$\pi_{3,\ldots,m}(a_3, \ldots, a_m) = \pi_3(a_3) + \ldots + \pi_m(a_m).$$

And by definition, the graph of the function $\pi_{3,\ldots,m}(a_3, \ldots, a_m)$ is given by $Q(x_1, x_2, a_3, \ldots, a_m)$. Combining and using Remark 3.10(2), we get $Q(a_1, \ldots, a_m) \iff \pi_1(a_1) + \pi_2(a_2) = \pi_{3,\ldots,m}(a_3, \ldots, a_m) \iff \pi_1(a_1) + \pi_2(a_2) = \pi_3(a_3) + \ldots + \pi_m(a_m)$. □

We are ready to prove the main theorem of the section.

Theorem 3.21. *Given $m \in \mathbb{N}_{\geq 4}$, sets $X_1, \ldots, X_m$ and $Q \subseteq \prod_{i\in[m]} X_i$ satisfying (P1) and (P2), there exists an abelian group $(G, +, 0_G)$ and bijections $\pi'_i : X_i \to G$ such that for every $(a_1, \ldots, a_m) \in \prod_{i\in[m]} X_i$ we have*

$$Q(a_1, \ldots, a_m) \iff \pi'_1(a_1) + \ldots + \pi'_m(a_m) = 0_G.$$

Moreover, if we have first-order structures $\mathcal{M} \preceq \mathcal{N}$ so that $\mathcal{N}$ is $|\mathcal{M}|^+$-saturated, each $X_i, i \in [m]$ is type-definable (respectively, definable) in $\mathcal{N}$ over $\mathcal{M}$ and $Q = F \cap \prod_{i\in[m]} X_i$ for a relation F definable in $\mathcal{N}$ over $\mathcal{M}$, then given an arbitrary tuple $\bar{e} \in Q$, we can take G to be type-definable (respectively, definable) and the bijections $\pi'_i, i \in [m]$ to be definable in $\mathcal{N}$, in both cases only using parameters from $\mathcal{M}$ and $\bar{e}$, so that $\pi'_i(e_i) = 0_G$ for all $i \in [m]$.

Proof. By Proposition 3.20, for any $\bar{a} = (a_1, \ldots, a_m) \in \prod_{i\in[m]} X_i$ we have

$$\begin{gathered} Q(a_1, \ldots, a_m) \iff \\ \pi_1(a_1) + \pi_2(a_2) = \pi_3(a_3) + \ldots + \pi_m(a_m) \iff \\ \pi_1(a_1) + \pi_2(a_2) + (-\pi_3(a_3)) + \ldots + (-\pi_m(a_m)) = 0_G, \end{gathered} \tag{3.5}$$

hence the bijections $\pi'_1 := \pi_1, \pi'_2 := \pi_2$ and $\pi'_i : X_i \to G, \pi'_i(x) := -\pi_i(x)$ for $3 \leq i \leq m$ satisfy the requirement.

Assume now that, for each $i \in [m]$, X_i is type-definable in $\mathcal{N}$ over $\mathcal{M}$, i.e. X_i is the set of solutions in $\mathcal{N}$ of some partial type $\mu_i(x_i)$ over $\mathcal{M}$; and that $Q = F \cap \prod_{i\in[m]} X_i$ for some $\mathcal{M}$-definable relation F. Then from (P1) and (P2) for Q, for any permutation of the variables of Q we have in $\mathcal{N}$:

$$\begin{gathered} \mu_m(x_m) \wedge \mu_m(x'_m) \wedge \bigwedge_{1\leq i\leq m-1} \mu_i(x_i) \wedge \\ \wedge F(x_1, \ldots, x_{m-1}, x_m) \wedge F(x_1, \ldots, x_{m-1}, x'_m) \to x_m = x'_m, \\ \bigwedge_{i\in[m]} \mu_i(x_i) \wedge \bigwedge_{i\in[m]} \mu_i(x'_i) \wedge F(x_1, x_2, x_3, \ldots, x_m) \wedge F(x_1, x_2, x'_3, \ldots, x'_m) \wedge \\ \wedge F(x'_1, x'_2, x_3, \ldots, x_m) \to F(x'_1, x'_2, x'_3, \ldots, x'_m). \end{gathered}$$

By $|\mathcal{M}|^+$-saturation of $\mathcal{N}$, in each of these implications μ_i can be replaced by a finite conjunction of formulas in it. Hence, taking a finite conjunction over all permutations of the variables, we conclude that there exist some $\mathcal{M}$-definable sets $X'_i \supseteq X_i, i \in [m]$ so that $Q' := F \cap \prod_{i\in[m]} X'_i$ satisfies (P2) and

(P1′) For any permutation of the variables of Q', for any $x_i \in X'_i, 1 \leq i \leq m-1$, there exists at most one (but possibly none) $x_m \in X'_m$ satisfying $Q'(x_1, \ldots, x_m)$.

We proceed to type-definability of G. Let $(e_1, \ldots, e_m) \in Q$ (so in $\mathcal{N}$) be as above (see Definition 3.13). We identify X_2 with $\mathcal{F}$, the domain of G, via the bijection π_2 above mapping $a_2 \in X_2$ to f_{e_1,a_2} (in an analogous manner we could identify the domain of G with any of the type-definable sets $X_i, i \in [s]$). Under this identification, the graph of addition in G is given by

$$R_+ := \left\{(a_2, a_2', a_2'') \in X_2 \times X_2 \times X_2 : a_2'' = f_{e_1,a_2} \circ f_{e_1,e_2}^{-1} \circ f_{e_1,a_2'}(e_1)\right\}$$
$$= \left\{(a_2, a_2', a_2'') \in X_2 \times X_2 \times X_2 : a_2'' = f_{e_1,a_2} \circ f_{e_1,e_2}^{-1}(a_2')\right\}.$$

We have the following claim.

Claim 3.22. • *For any $a_1 \in X_1, a_2 \in X_2$ and $\bar{b} \in \prod_{3 \le i \le m} X_i'$, if $F(a_1, a_2, \bar{b})$ holds then $F_{\bar{b}} \restriction_{X_1 \times X_2}$ defines the graph of f_{a_1,a_2} (since Q' satisfies (P2)).*

- *For any $\bar{b} \in \prod_{3 \le i \le m} X_i'$, if $F_{\bar{b}} \restriction_{X_1 \times X_2}$ coincides with the graph of some function $f \in \mathcal{F}$, then using that Q' satisfies (P1$'$) we have:*
  - *for any $a_1 \in X_1$, $f(a_1)$ is the unique element in X_2' satisfying $F(a_1, x_2, \bar{b})$;*
  - *for any $a_2 \in X_2$, $f^{-1}(a_2)$ is the unique element in X_1' satisfying $F(x_1, a_2, \bar{b})$.*

Using Claim 3.22, we have

$$R_+ = R_+' \restriction \prod_{i \in [m]} X_i,$$

where R_+' is a definable relation in $\mathcal{N}$ (with parameters in $\mathcal{M} \cup \{e_1, e_2\}$) given by

$$R_+'(x_2, x_2', x_2'') : \iff \exists \bar{y}, \bar{y}', z\Big(\bar{y} \in \prod_{3 \le i \le m} X_i' \wedge \bar{y}' \in \prod_{3 \le i \le m} X_i' \wedge z \in X_1' \wedge$$
$$F(e_1, e_2, \bar{y}') \wedge F(z, x_2', \bar{y}') \wedge F(e_1, x_2, \bar{y}) \wedge F(z, x_2'', \bar{y})\Big).$$

This shows that $(G, +)$ is type-definable over $\mathcal{M} \cup \{e_1, e_2\}$. It remains to show definability of the bijections $\pi_i' : X_i \to \mathcal{F}$, where $\mathcal{F}$ is identified with X_2 as above (i.e. to show that the graph of π_i' is given by some $\mathcal{N}$-definable relation $P_i(x_i, x_2)$ intersected with $X_i \times X_2$).

We have $\pi_1' : a_1 \in X_1 \mapsto f_{a_1,e_2} \in \mathcal{F}$, hence we need to show that the relation

$$\{(a_1, a_2) \in X_1 \times X_2 : f_{a_1,e_2}(e_1) = a_2\}$$

is of the form $P_1(x_1, x_2) \restriction X_1 \times X_2$ for some relation P_1 definable in $\mathcal{N}$. Using Claim 3.22, we can take

$$P_1(x_1, x_2) : \iff \exists \bar{y} \left(\bar{y} \in \prod_{3 \le i \le m} X_i' \wedge F(x_1, e_2, \bar{y}) \wedge F(e_1, x_2, \bar{y})\right).$$

We have $\pi_2' : a_2 \in X_2 \mapsto f_{e_1,a_2} \in \mathcal{F}$, hence the corresponding definable relation $P_2(x_2, x_2)$ is just the graph of the equality.

Finally, given $3 \le i \le m$, π_i maps $a_i \in X_i$ to the function in $\mathcal{F}$ with the graph given by $Q(x_1, x_2, e_3, \ldots, e_{i-1}, a_i, e_{i+1}, \ldots, e_m)$. Hence, remembering that the identity of G is f_{e_1,e_2}, which corresponds to $e_2 \in X_2$, and using

Claim 3.22, the graph of $\pi'_i : a_i \in X_i \mapsto -\pi_i(a_i)(e_1) \in X_2$ is given by the intersection of $X_i \times X_2$ with the definable relation

$$P_i(x_i, x_2) :\iff \exists z\Big(z \in X'_2 \wedge F(e_1, z, e_3, \ldots, e_{i-1}, x_i, e_{i+1}, \ldots, e_m)\wedge R'_+(x_2, z, e_2)\Big). \qquad \square$$

4. Reconstructing an abelian group from an abelian m-gon

Let $T = T^{\mathrm{eq}}$ be a stable theory in a language $\mathcal{L}$ and $\mathbb{M}$ a monster model of T. By "independence" we mean independence in the sense of forking, unless stated otherwise, and write $a \downarrow_c b$ to denote that $\mathrm{tp}(a/bc)$ does not fork over c. We assume some familiarity with the properties of forking in stable theories (see e.g. [44] for a concise introduction to model-theoretic stability, and [46] for a detailed treatment). We say that a subset A of $\mathcal{M}$ is *small* if $|A| \leq |\mathcal{L}|$.

4.1. **Abelian m-gons.** For a small set A, as usual by its acl_A-*closure* we mean the algebraic closure over A, i.e. for a set X its acl_A-closure is $\mathrm{acl}_A(X) := \mathrm{acl}(A \cup X)$.

Definition 4.1. [1] We say that a tuple $(a_1, \ldots, a_m)$ is *an m-gon over a set A* if each type $\mathrm{tp}(a_i/A)$ is not algebraic, any $m-1$ elements of the tuple are independent over A, and every element is in the acl_A-closure of the rest. We refer to a 3-gon as a *triangle*.

Definition 4.2. We say that an m-gon $(a_1, \ldots, a_m)$ over A with $m \geq 4$ is *abelian* if for any $i \neq j \in [m]$, taking $\bar{a}_{ij} := (a_k)_{k \in [m] \setminus \{i,j\}}$, we have

$$a_i a_j \mathop{\downarrow}_{\mathrm{acl}_A(a_i a_j) \cap \mathrm{acl}_A(\bar{a}_{ij})} \bar{a}_{ij}.$$

Example 4.3. Let A be a small set and let $(G, \cdot, 1_G)$ be an abelian group type-definable over A. Let $g_1, \ldots, g_{m-1} \in G$ be independent generic elements over A, and let g_m be such that $g_1 \cdot \ldots \cdot g_m = 1_G$. Then $(g_1, \ldots, g_m)$ is an abelian m-gon over A *associated to* G.

Indeed, by assumption we have $g_1 \cdot g_2 \in \mathrm{dcl}(g_1, g_2) \cap \mathrm{dcl}(g_3, \ldots, g_m)$. Also $g_1 g_2 \downarrow_A g_3 \cdots g_{m-1}$, hence $g_1 g_2 \downarrow_{A, g_1 \cdot g_2} g_3 \cdots g_{m-1}$, which together with $g_m \in \mathrm{dcl}(g_1 \cdot g_2, g_3, \ldots, g_{m-1})$ implies $g_1 g_2 \downarrow_{A, g_1 \cdot g_2} g_3 \cdots g_m$. As the group G is abelian, the same holds for any $i \neq j \in [m]$ instead of $i = 1, j = 2$.

Definition 4.4. Given two tuples $\bar{a} = (a_1, \ldots, a_m)$, $\bar{a}' = (a_1, \ldots, a_m)$ and a small set A we say that $\bar{a}$ and $\bar{a}'$ are acl-*equivalent over A* if $\mathrm{acl}_A(a_i) = \mathrm{acl}_A(a'_i)$ for all $i \in [m]$. As usual if $A = \emptyset$ we omit it.

Remark 4.5. Note that the condition "$\bar{a}, \bar{a}'$ are acl-equivalent" is stronger than "the tuples $\bar{a}, \bar{a}'$ are inter-algebraic", as it requires inter-algebraicity component-wise.

---

[1] An analogous notion in the context of geometric theories was introduced in [10] under the name of an *algebraic m-gon*, and it was also used in [16, Section 7].

In this section we prove the following theorem.

Theorem 4.6. *Let $\bar{a} = (a_1, \dots, a_m)$ be an abelian m-gon, over some small set A. Then there is a finite set C with $\bar{a} \mathop{\smile\!\!\!\!\!\vert}_A C$, a type-definable (in $\mathbb{M}^{\mathrm{eq}}$) over $\operatorname{acl}(C \cup A)$ connected (i.e. $G = G^0$) abelian group $(G, \cdot)$ and an abelian m-gon $\bar{g} = (g_1, \dots, g_m)$ over $\operatorname{acl}(C \cup A)$ associated to G such that $\bar{a}$ and $\bar{g}$ are acl-equivalent over $\operatorname{acl}(C \cup A)$.*

Remark 4.7. After this work was completed, we have learned that independently Hrushovski obtained a similar (but incomparable) result [31, 32].

Remark 4.8. In the case $m = 4$, Theorem 4.6 follows from the Abelian Group Configuration Theorem (see [9, Theorem C.2]).

In the rest of the section we prove Theorem 4.6, following the presentation of Hrushovski's Group Configuration Theorem in [7, Theorem 6.1] with appropriate modifications.

First note that, adding to the language new constants naming the elements of $\operatorname{acl}(A)$, we may assume without loss of generality that $A = \emptyset$ in Theorem 4.6, and that all types over the empty set are stationary.

Given a tuple $\bar{a} = (a_1, \dots a_m)$ we will often modify it by applying the following two operations:

- for a **finite** set B with $\bar{a} \mathop{\smile\!\!\!\!\!\vert} B$ we expand the language by constants for the elements of $\operatorname{acl}(B)$, and refer to this as *"base change to B"*.
- we replace $\bar{a}$ with an acl-equivalent tuple $\bar{a}'$ (over $\emptyset$), and refer to this as *"inter-algebraic replacing"*.

It is not hard to see that these two operations transform an (abelian) m-gon to an (abelian) m-gon, and we will freely apply them to the m-gon $\bar{a}$ in the proof of Theorem 4.6.

Definition 4.9. We say that a tuple $(a_1, \dots, a_m, \xi)$ is an *expanded abelian m-gon* if $(a_1, \dots, a_m)$ is an abelian m-gon, $\xi \in \operatorname{acl}(a_1, a_2) \cap \operatorname{acl}(a_3, \dots, a_m)$ and $a_1 a_2 \mathop{\smile\!\!\!\!\!\vert}_\xi a_3 \dots a_m$.

We remark that the tuple ξ might be infinite even if all of the tuples a_i's are finite. Similarly, base change and inter-algebraic replacement transform an expanded abelian m-gon to an expanded abelian m-gon.

From now on, we fix an abelian m-gon $\vec{a} = (a_1, \dots, a_m)$. We also fix $\xi \in \operatorname{acl}(a_1, a_2) \cap \operatorname{acl}(a_3, \dots, a_m)$ such that $a_1 a_2 \mathop{\smile\!\!\!\!\!\vert}_\xi a_3 \dots a_m$ (exists by the definition of abelianity).

Claim 4.10. *(a_1, a_2, ξ) is a triangle and $(\xi, a_3, \dots, a_m)$ is an $(m-1)$-gon.*

Proof. For $i = 1, 2$, since $a_i \mathop{\smile\!\!\!\!\!\vert} a_3, \dots, a_m$ and $\xi \in \operatorname{acl}(a_3, \dots, a_m)$ we have $a_i \mathop{\smile\!\!\!\!\!\vert} \xi$. Also $a_1 \mathop{\smile\!\!\!\!\!\vert} a_2$. Thus the set $\{a_1, a_2, \xi\}$ is pairwise independent. We also have $\xi \in \operatorname{acl}(a_1, a_2)$. From $a_1 a_2 \mathop{\smile\!\!\!\!\!\vert}_\xi a_3 \dots a_m$ we obtain $a_1 \mathop{\smile\!\!\!\!\!\vert}_{\xi a_2} a_3 \dots a_m$. Since $a_1 \in \operatorname{acl}(a_2, \dots, a_m)$ we obtain $a_1 \in \operatorname{acl}(\xi, a_2)$. Similarly $a_2 \in \operatorname{acl}(\xi, a_1)$, thus (a_1, a_2, ξ) is a triangle.

The proof that $(\xi, a_3, \dots, a_m)$ is an $(m-1)$-gon is similar. $\square$

4.2. **Step 1. Obtaining a pair of interdefinable elements.** *After applying finitely many base changes and inter-algebraic replacements we may assume that a_1 and a_2 are interdefinable over ξ, i.e. $a_1 \in \mathrm{dcl}(\xi, a_2)$ and $a_2 \in \mathrm{dcl}(\xi, a_1)$.*

Our proof of Step 1 follows closely the proof of the corresponding step in the proof of [7, Theorem 6.1], but in order to keep track of the additional parameters we work with enhanced group configurations.

Definition 4.11. *An enhanced group configuration* is a tuple

$$(a, b, c, x, y, z, d, e)$$

satisfying the following diagram.

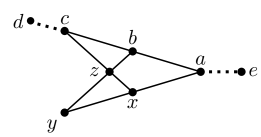

That is,

- (a, b, c) is a triangle over de;
- (c, z, x) is a triangle over d;
- (y, x, a) is a triangle over e;
- (y, z, b) is a triangle;
- for any non-collinear triple in (a, b, c, x, y, z), the set given by it and de is independent over $\emptyset$.

If $e = \emptyset$ we omit it from the diagram:

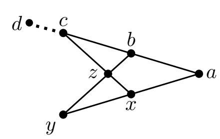

In order to complete Step 1 we first show a few lemmas.

Lemma 4.12. *Let (a, b, c, x, y, z, d, e) be an enhanced group configuration. Let $\tilde{z} \in \mathbb{M}^{\mathrm{eq}}$ be the imaginary representing the finite set $\{z_1, \ldots, z_k\}$ of all conjugates of z over $bcxyd$. Then $\tilde{z}$ is inter-algebraic with z.*

Proof. It suffices to show that $\mathrm{acl}(z_i) = \mathrm{acl}(z_j)$ for all $1 \leq i, j \leq k$. Indeed, then $\tilde{z} \in \mathrm{acl}(z_1, \ldots, z_k) = \mathrm{acl}(z)$, and $z \in \mathrm{acl}(\tilde{z})$ as it satisfies the algebraic formula "$z \in \tilde{z}$".

We have $cd \mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}} yz$, so $cd \mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}_z y$, so $cdx \mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}_z by$. Let $B := \mathrm{acl}(cdx) \cap \mathrm{acl}(by)$, then $B \mathop{\smile\hspace{-0.9em}\mid\hspace{0.5em}}_z B$, so $B \subseteq \mathrm{acl}(z)$. But $z \in B$, so $B = \mathrm{acl}(z)$. Then we also have $\mathrm{acl}(z_i) = B$ since for each z_i there is an automorphism σ of $\mathbb{M}$ with $\sigma(z) = z_i$ and $\sigma(B) = B$. □

Lemma 4.13. *Assume that (a, b, c, x, y, z, d, e) is an enhanced group configuration. Then after a base change it is* acl*-equivalent to an enhanced group configuration $(a, b_1, c, x, y_1, z_1, d, e)$ such that $z_1 \in \operatorname{dcl}(b_1 y_1)$. Moreover, $b \in \operatorname{dcl}(b_1)$ and $y \in \operatorname{dcl}(y_1)$.*

Proof. Recall that by our assumption all types over the empty set are stationary.

Let $a'd'e' \models \operatorname{tp}(ade)|_{abcdexyz}$. We have $ade \mathop{\smile\!\!\!\!\!\mid}\, yz$, hence $ade \mathop{\smile\!\!\!\!\!\mid}\, yzb$. Then by stationarity we have $a'd'e' \equiv_{yzb} ade$. Let x', c' be such that $a'd'e'x'c' \equiv_{yzb} adexc$. So $(a', b, c', x', y, z, d', e')$ is also an enhanced group configuration. Applying Lemma 4.12 to it, the set $\tilde{z}'$ of conjugates of z over $ybx'c'd'$ is inter-algebraic with z, and $\tilde{z}' \in \operatorname{dcl}(ybx'c'd')$.

We add $\operatorname{acl}(a'd'e')$ to the base, and take $y_1 := yx'$, $b_1 := bc'$, $z_1 := \tilde{z}'$. Then $(a, b_1, c, x, y_1, z_1, d, e)$ is an enhanced group configuration satisfying the conclusion of the lemma. □

Lemma 4.14. *Let (a, b, c, x, y, z, d, e) be an enhanced group configuration with $e \in \operatorname{dcl}(\emptyset)$. Then, applying finitely many base changes and inter-algebraic replacements, it can be transformed to a configuration*

$$(a_1, b_1, c_1, x_1, y_1, z_1, d, e)$$

such that y_1 and z_1 are interdefinable *over b_1. (Notice that d and e remain unchanged.)*

Proof. Applying Lemma 4.13, after a base change and an inter-algebraic replacement we may assume $z \in \operatorname{dcl}(by)$.

Next observe that, since $e \in \operatorname{dcl}(\emptyset)$, the tuple (b, a, c, z, y, x, d, e) is also an enhanced group configuration.

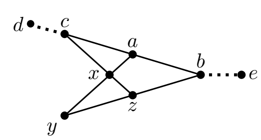

By Lemma 4.13, after a base change, it is acl-equivalent to a configuration $(b, a_1, c, z, y_1, x_1, d, e)$ with $x_1 \in \operatorname{dcl}(a_1, y_1)$ and $y \in \operatorname{dcl}(y_1)$. Thus after an inter-algebraic replacement we may assume that $x \in \operatorname{dcl}(ay)$ and $z \in \operatorname{dcl}(by)$.

Finally, observe that (c, b, a, x, z, y, e, d) is an enhanced group configuration.

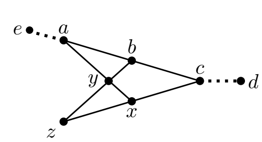

Applying the proof of Lemma 4.13 to it, after base change to an independent copy $c'd'e'$ of cde, let $a'x'c'd'e' \equiv_{ybz} axcde$, let $\tilde{y}'$ be the set of conjugates of y over $ba'zx'e'$, equivalently over $ba'zx'$ since $e' \in \operatorname{dcl}(\emptyset)$. So $y' \in \operatorname{dcl}(ba'zx')$.

Now since $x' \in \operatorname{dcl}(a'y)$ and $z \in \operatorname{dcl}(by)$ (since this was satisfied on the previous step), we have $zx' \in \operatorname{dcl}(ba'y)$. But then $zx' \in \operatorname{dcl}(ba'y')$ for any y' a conjugate of y over $ba'zx'$, and so $zx' \in \operatorname{dcl}(ba'\tilde{y}')$. We take $b_1 := ba'$, $z_1 := zx'$ and $y_1 := \tilde{y}'$. Then $y_1 \in \operatorname{dcl}(b_1 z_1)$, and also $z_1 \in \operatorname{dcl}(b_1 y_1)$, and the tuple $(a, b_1, c, x, y_1, z_1, d, e)$ satisfies the conclusion of the lemma. □

We can now finish Step 1.

Let $(a_1, \ldots, a_m, \xi)$ be an expanded abelian m-gon. Let $\tilde{a} := a_5 \ldots a_m$ and $\eta := \operatorname{acl}(a_1 a_3) \cap \operatorname{acl}(a_2 a_4 \ldots a_m)$

It is easy to check that $(a_3, \xi, a_4, \eta, a_1, a_2, \tilde{a}, \emptyset)$ is an enhanced group configuration.

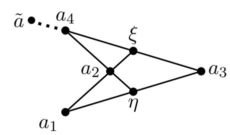

Applying Lemma 4.14, after a base change it is acl-equivalent to an enhanced group configuration $(a_3', \xi', a_4', \eta', a_1', a_2', \tilde{a}, \emptyset)$ such that a_1' and a_2' are interdefinable over ξ'. Replacing a_1, a_2, a_3, a_4 with a_1', a_2', a_3', a_4', respectively, and ξ with ξ' we complete Step 1.

Reduction 1. *From now on we assume that in the expanded abelian m-gon $(a_1, \ldots, a_m, \xi\,)$ we have that a_1 and a_2 are interdefinable over ξ.*

4.3. **Step 2. Obtaining a group from an expanded abelian m-gon.** As in Hrushovski's Group Configuration Theorem, we will construct a group using germs of definable functions. We begin by recalling some definitions (see e.g. [7, Section 5.1]).

Let $p(x)$ be a stationary type over a set A. By *a definable function on* $p(x)$ we mean a (partial) function $f(x)$ definable over a set B such that every element $a \models p|_{AB}$ is in the domain of f.

If f and g are two definable functions on $p(x)$, defined over sets B and C respectively, then we say that *they have the same germ at* $p(x)$, and write $f \sim_p g$, if for all (equivalently, some) $a \models p|_{ABC}$ we have $f(a) = g(a)$. We may omit p and write $f \sim g$ if no confusion arises.

The germ of a definable function f at p is the equivalence class of f under this equivalence relation, and we denote it by $\tilde{f}$.

If $p(x)$ and $q(y)$ are stationary types over $\emptyset$, we write $\tilde{f} : p \to q$ if for some (any) representative f of $\tilde{f}$ definable over B and $a \models p|_B$ we have $f(a) \models q$. We say that $\tilde{f}$ is *invertible* if there exists a germ $\tilde{g} : q \to p$ and for some (any) representative g definable over C and $a \models p|_{BC}$ we have $g(f(a)) = a$. We denote $\tilde{g}$ by $\tilde{f}^{-1}$.

By a *type-definable family of functions from p to q* we mean an $\emptyset$-definable family of functions f_z and a stationary type $s(z)$ over $\emptyset$ such that for any $c \models s(z)$ the function f_c is a definable function on p, and for any $a \models p|_c$

we have $f_c(a) \models q(y)|_c$. We will denote such a family as $f_s \colon p \to q$, and the family of the corresponding germs as $\tilde{f}_s \colon p \to q$.

Let p, q, s be stationary types over $\emptyset$ and $f_s \colon p \to q$ a type-definable family of functions. This family is *generically transitive* if $f_c(a) \mathop{\smile\!\!\!\!\!\!|}\, a$ for any (equivalently, some) $c \models s$ and $a \models p|c$. This family is *canonical* if for any $c, c' \models s$ we have $f_c \sim f_{c'} \Leftrightarrow c = c'$.

We now return to our expanded abelian m-gon $(\vec{a}, \xi)$.

Let $p_i(x_i) := \operatorname{tp}(a_i/\emptyset)$ for $i \in \{1, 2\}$, and let $q(y) := \operatorname{tp}(\xi/\emptyset)$.

Since a_1 and a_2 are interdefinable over ξ and $\xi \in \operatorname{acl}(a_1, a_2)$, there exists a formula $\varphi(x_1, x_2, y) \in \operatorname{tp}(a_1, a_2, \xi)$ such that

$$\models \forall y \forall x_1 \exists^{\leq 1} x_2 \varphi(x_1, x_2, y), \models \forall y \forall x_2 \exists^{\leq 1} x_1 \varphi(x_1, x_2, y),$$
$$\models \forall x_1 \forall x_2 \exists^{\leq d} \varphi(x_1, x_2, y),$$

for some $d \in \mathbb{N}$, and also

$$\varphi(a_1, a_2, y) \vdash \operatorname{tp}(\xi/a_1 a_2).$$

It follows that $\varphi(x_1, x_2, r), r \models q$ gives a type-definable family of invertible germs $\tilde{f}_q \colon p_1 \to p_2$ with $f_\xi(a_1) = a_2$.

Remark 4.15. Let $r \models q$, $b_1 \models p_1|r$ and $b_2 := f_r(b_1)$. By stationarity of types over $\emptyset$ we then have $b_1 r \equiv a_1 \xi$, and as $\varphi(b_1, x_2, r)$ has a unique solution this implies $b_1 b_2 r \equiv a_1 a_2 \xi$, so $b_1 \mathop{\smile\!\!\!\!\!\!|}\, b_2$, $b_1 \in \operatorname{dcl}(b_2, r)$ and $r \in \operatorname{acl}(b_1, b_2)$.

In particular $\tilde{f}_q \colon p_1 \to p_2$ is a generically transitive invertible family.

Consider the equivalence relation $E(y, y')$ on the set of realizations of q given by $rEr' \Leftrightarrow f_r \sim f_{r'}$. By the definability of types it is relatively definable, i.e. it is an intersection of an $\emptyset$-definable equivalence relation with $q(y) \cup q(y')$. Assume $\xi' \models q$ with $\xi E \xi'$. We choose $b_1 \models p_1|\xi\xi'$ and let $b_2 := f_\xi(b_1) = f_{\xi'}(b_1)$. By the choice of φ we have $\xi, \xi' \in \operatorname{acl}(b_1, b_2)$, hence ξ and ξ' are inter-algebraic over b_1. Since $b_1 \mathop{\smile\!\!\!\!\!\!|}\, \xi\xi'$ it follows that ξ and ξ' are inter-algebraic over $\emptyset$: as $b_1 \mathop{\smile\!\!\!\!\!\!|}_{\xi}\, \xi'$ and $\xi' \in \operatorname{acl}(b_1 \xi)$ implies $\xi' \mathop{\smile\!\!\!\!\!\!|}_{\xi}\, \xi'$, hence $\xi' \in \operatorname{acl}(\xi)$; and similarly $\xi \in \operatorname{acl}(\xi')$. Hence the E-class of ξ is finite. Replacing ξ by ξ/E, if needed, we will assume that the family $\tilde{f}_q \colon p_1 \to p_2$ is canonical.

We now consider the type-definable family of germs $\tilde{f}_{r_1}^{-1} \circ \tilde{f}_{r_2} \colon p_1 \to p_1$, $(r_1, r_2) \models q^{(2)}$. Again let E be a relatively definable equivalence relation on $q^{(2)}$ defined as $(r_1, r_2) E (r_3, r_4)$ if and only if $f_{r_1}^{-1} \circ f_{r_2} \sim f_{r_3}^{-1} \circ f_{r_4}$. Let $s(z)$ be the type $q^{(2)}/E$. We then have (by e.g. [29, Remark 3.3.1(1)]) a canonical family of germs $\tilde{h}_s \colon p_1 \to p_1$ such that for every $(r_1, r_2) \models q^{(2)}$ there is unique $c \models s(z)$ with $\tilde{h}_c = \tilde{f}_{r_1}^{-1} \circ \tilde{f}_{r_2}$. We will denote this c as $c = \ulcorner f_{r_1}^{-1} \circ f_{r_2} \urcorner$. Clearly $c \in \operatorname{dcl}(r_1, r_2)$, $r_1 \in \operatorname{dcl}(c, r_2)$ and $r_2 \in \operatorname{dcl}(c, r_1)$.

Lemma 4.16. *For any $(r_1, r_2) \models q^{(2)}$ and $c := \ulcorner f_{r_1}^{-1} \circ f_{r_2} \urcorner$ we have $r_1 \mathop{\smile\!\!\!\!\!\!|}\, c$ and $r_2 \mathop{\smile\!\!\!\!\!\!|}\, c$.*

Proof. It is sufficient to prove the lemma for some $(r_1, r_2) \models q^{(2)}$. We take $r_1 := \xi$ from our abelian expanded m-gon $(\vec{a}, \xi)$ and let $r_2 \models q|_{a_1,\ldots,a_m}$. Let $c := \ulcorner f_{\xi}^{-1} \circ f_{r_2} \urcorner$.

Let $\tilde{a} := (a_5, \ldots, a_m)$ and $\eta := \operatorname{acl}(a_1 a_3) \cap \operatorname{acl}(a_2 a_4 \ldots a_m)$. We have an enhanced group configuration

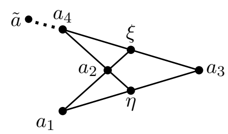

In particular $(a_3, \xi, a_4, \eta, a_1, a_2)$ form a group configuration *over* $\tilde{a}$, i.e. we have a group configuration

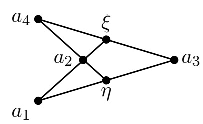

where any three distinct collinear points form a triangle over $\tilde{a}$, and any three distinct non-collinear points form an independent set over $\tilde{a}$.

It follows from the proof of the Group Configuration Theorem (e.g. see Step (II) in the proof of [7, Theorem 6.1]) that $c \perp_{\tilde{a}} \xi$ and $c \perp_{\tilde{a}} r_2$.

We also have $r_2 \perp a_1 \ldots a_m$, hence $r_2 \perp_{a_1 a_2} \tilde{a}$, and as $a_1 a_2 \perp \tilde{a}$ this implies $r_2 a_1 a_2 \perp \tilde{a}$, which together with $\xi \in \operatorname{acl}(a_1 a_2)$ implies $\xi r_2 \perp \tilde{a}$. Hence $c \perp \xi$ and $c \perp r_2$. $\square$

This shows that the families of germs $\tilde{f}_q : p_1 \to p_2, \tilde{h}_s : p_1 \to p_1$ satisfy the assumptions of the Hrushovski-Weil theorem for bijections (see [7, Lemma 5.4]), applying which we obtain the following.

(a) The family of germs $\tilde{h}_s \colon p_1 \to p_1$ is closed under generic composition and inverse, i.e. for any independent $c_1, c_2 \models s(z)$ there exists $c \models s(z)$ with $\tilde{h}_c = \tilde{h}_{c_1} \circ \tilde{h}_{c_2}$, and also there is $c_3 \models s(z)$ with $\tilde{h}_{c_3} = \tilde{h}_{c_1}^{-1}$.

(b) There is a type-definable connected group $(G, \cdot)$ and a type-definable set S with a relatively definable faithful transitive action of G on S that we will denote by $* : G \times S \to S$, so that G, S and the action are defined over the empty set.

(c) There is a definable embedding of $s(z)$ into G as its unique generic type, and a definable embedding of $p_1(x_1)$ into S as its unique generic type, such that the generic action of the family h_s on p_1 agrees with that of G on S, i.e. for any $c \models s(z)$ and $a \models p_1(x)|c$ we have $h_c(a) = c * a$.

Reduction 2. *Let r_1, r_2 be independent realizations of $q(y)$, $c := \ulcorner f_{r_1}^{-1} \circ f_{r_2} \urcorner$ and $s(z) := \operatorname{tp}(c/\emptyset)$.*

From now on we assume that $s(z)$ is the generic type of a type-definable connected group $(G, \cdot)$, the group G relatively definably acts faithfully and transitively on a type-definable set S, the type $p_1(x_1)$ is the generic type of

S, and generically the action of h_s on p_1 agrees with the action of G on S, and G, S and the action are definable over the empty set.

4.4. **Step 3. Finishing the proof.** We fix an independent copy $(\vec{e}, \xi_e)$ of $(\vec{a}, \xi)$, i.e. $(\vec{e}, \xi_e) \equiv (\vec{a}, \xi)$ and $\vec{e}\xi_e \mathop{\smile\!\!\!\!\!|}\, \vec{a}\xi$.

We denote by π the map $\pi\colon q(y)|_{\xi_e} \to s(z)|_{\xi_e}$ given by $\pi\colon r \mapsto \ulcorner f_{\xi_e}^{-1} \circ f_r \urcorner$. Note that π is relatively definable over $\operatorname{acl}(\vec{e})$. Let

$$t(x_3, \dots, x_m) := \operatorname{tp}(a_3, \dots, a_m/\emptyset),$$
$$t_\xi(y, x_3, \dots, x_m) := \operatorname{tp}(\xi, a_3, \dots, a_m/\emptyset).$$

Note that by Claim 4.10 every tuple realizing t_ξ is an $(m-1)$-gon.

Notation 4.17. For a tuple $\bar{c} = (c_3, \dots, c_m)$, $j \in \{3, \dots, m\}$ and $\square \in \{<, \leq, >, \geq\}$, we will denote by $\bar{c}_{\square j}$ the tuple $\bar{c}_{\square j} = (c_i : 3 \leq i \leq m \wedge i \square j)$. For example, $\bar{c}_{<j} = (c_3, \dots, c_{j-1})$. We will typically omit the concatenation sign: e.g., for $\bar{c} = (c_3, \dots, c_m)$, $\bar{b} = (b_3, \dots, b_m)$ and $j \in \{3, \dots, m\}$ we denote by $\bar{c}_{<j}, b_j, \bar{c}_{>j}$ the tuple $(c_3, \dots, c_{j-1}, b_j, c_{j+1}, \dots, c_m)$.

Also in the proof of the next proposition we let $\bar{a} := (a_3, \dots, a_m)$, $\bar{e} := (e_3, \dots, e_m)$, and continue using $\vec{a}$ and $\vec{e}$ to denote the corresponding m-tuples.

Proposition 4.18. *For each $j \in \{3, \dots, m\}$ there exists $r_j \models q(y)|_{\xi_e}$ such that $\models t_\xi(r_j, \bar{e}_{<j}, a_j, \bar{e}_{>j})$ and $\pi(\xi) = \pi(r_m) \cdot \pi(r_{m-1}) \cdot \ldots \cdot \pi(r_3)$.*

We will choose such r_j by reverse induction on j. Before proving Proposition 4.18 we first establish the following lemma and its corollary that will provide the induction step.

Lemma 4.19. *For $j \in \{4, \dots, m\}$ there exist $r_{<j}, r_j, r_{\leq j}$, each realizing $q(y)|_{\xi_e}$, such that*

$$\models t_\xi(r_{<j}, \bar{a}_{<j}, \bar{e}_{\geq j}), \models t_\xi(r_j, \bar{e}_{<j}, a_j, \bar{e}_{>j}), \models t_\xi(r_{\leq j}, \bar{a}_{\leq j}, \bar{e}_{>j})$$

and $\pi(r_{\leq j}) = \pi(r_j) \cdot \pi(r_{<j})$.

Proof. First we note that the condition $r_{<j}, r_j, r_{\leq j} \models q(y)|_{\xi_e}$ can be relaxed to $r_{<j}, r_j, r_{\leq j} \models q(y)$ by stationarity of q, since for $j \in \{4, \dots, m\}$ and $r \models q(y)$ satisfying one of $\models t_\xi(r, \bar{a}_{<j}, \bar{e}_{\geq j})$, $\models t_\xi(r, \bar{e}_{<j}, a_j, \bar{e}_{>j})$, $\models t_\xi(r, \bar{a}_{\leq j}, \bar{e}_{>j})$ we have $r \mathop{\smile\!\!\!\!\!|}\, \xi_e$. Indeed, assume e.g. $\models t_\xi(r, \bar{a}_{<j}, \bar{e}_{\geq j})$. We have $r \in \operatorname{acl}(\bar{a}_{<j}, \bar{e}_{\geq j})$ and $\xi_e \in \operatorname{acl}(e_3, \dots, e_m)$. By assumption

$$\{e_3, \dots, e_m, a_3, \dots, a_m\}$$

is an independent set, hence we obtain $r \mathop{\smile\!\!\!\!\!|}_{\bar{e}_{\geq j}} \xi_e$. Using $\xi_e \mathop{\smile\!\!\!\!\!|}\, \bar{e}_{\geq j}$ we conclude $r \mathop{\smile\!\!\!\!\!|}\, \xi_e$. The other two cases are similar.

Let $\eta := \operatorname{acl}_{\bar{e}_{>j}}(e_1, e_j) \cap \operatorname{acl}_{\bar{e}_{>j}}(e_2, e_3, \dots, e_{j-1})$. Note that $\operatorname{acl}(\eta) = \eta$, hence all types over η are stationary, and $\bar{e}_{>j} \in \eta$.

Then one verifies by basic forking calculus that

(4.1)
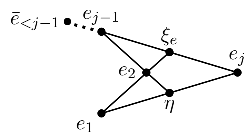

is an enhanced group configuration *over* $\bar{e}_{>j}$. Namely,

- (e_j, ξ_e, e_{j-1}) and (η, e_2, e_{j-1}) are triangles over $\bar{e}_{<j-1}, \bar{e}_{>j}$;
- (e_1, η, e_j) and (e_1, e_2, ξ_e) are triangles over $\bar{e}_{>j}$;
- for any non-collinear triple in $e_1, e_2, e_{j-1}, e_j, \eta, \xi_e$, the set given by it and $\bar{e}_{<j-1}$ is independent over $\bar{e}_{>j}$.

In addition, $e_1 e_2 \xi_e \mathop{\smile\hskip-0.9em\vert}\, \bar{e}_{>j}$ and $f_{\xi_e}(e_1) = e_2$.

The triple η, e_j, e_{j-1} is non-collinear, hence $\eta \mathop{\smile\hskip-0.9em\vert}_{\bar{e}_{>j}} e_3 \dots e_j$. Since

$$\bar{e}_{>j} \mathop{\smile\hskip-0.9em\vert}\, e_3 \dots e_j,$$

this implies $\eta \mathop{\smile\hskip-0.9em\vert}\, e_3 \dots e_j$. Since also $\eta \mathop{\smile\hskip-0.9em\vert}\, a_3 \dots a_j$, by stationarity of types over $\emptyset$ we have $a_3 \dots a_j \equiv_\eta e_3 \dots e_j$. Hence there exist $r_{\leq j}$, b_1, b_2 such that the diagram

(4.2)
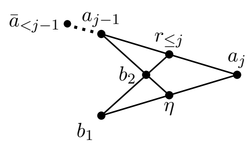

is isomorphic over η to the diagram (4.1). I.e., there is an automorphism of $\mathbb{M}$ fixing η (hence also $\bar{e}_{>j}$) and mapping (4.2) to (4.1).

It follows from the choice of the tuple $(\vec{e}, \xi_e)$, diagrams (4.1), (4.2) and their isomorphism over η that $e_1 e_j \mathop{\smile\hskip-0.9em\vert}_\eta e_2 \dots e_{j-1}$ and $b_1 a_j \equiv_\eta e_1 e_j$. Since $a_j \mathop{\smile\hskip-0.9em\vert}\, e_1 \dots e_m$ we have $a_j \mathop{\smile\hskip-0.9em\vert}_\eta e_2 \dots e_{j-1}$. As $b_1 \in \operatorname{acl}(a_j \eta)$, we have

$$b_1 a_j \mathop{\smile\hskip-0.9em\vert}_\eta e_2 \dots e_{j-1}.$$

Since all types over η are stationary, this implies

$$b_1 a_j e_2 \dots e_{j-1} \equiv_\eta e_1 e_j e_2 \dots e_{j-1},$$

hence there exists r_j such that the diagram

(4.3)
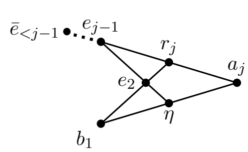

is isomorphic to the diagram (4.1) over η.

A similar argument with the roles of the a's and the e's interchanged shows that $e_1 e_j a_2 \dots a_{j-i} \equiv_\eta b_1 a_j a_2 \dots a_{j-1}$, hence there exists $r_{<j}$ such

that the diagram

(4.4)

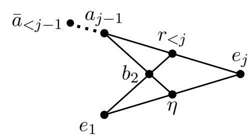

is isomorphic to the diagram (4.1) over η.

From the choice of $(\vec{e}, \xi_e)$ and the isomorphisms of the diagrams we have

$$(f_{r_{<j}} \circ f_{\xi_e}^{-1} \circ f_{r_j})(b_1) = b_2 = f_{r_{\leq j}}(b_1). \tag{4.5}$$

We claim that $b_1 \mathop{\smile}\limits^{|} r_{<j}, \xi_e, r_j, r_{\leq j}$. Indeed, as

$$r_{<j}, \xi_e, r_j, r_{\leq j} \in \operatorname{acl}(a_3, \dots, a_m, e_3, \dots e_m) \text{ and}$$
$$e_2 \mathop{\smile}\limits^{|} a_3, \dots, a_m, e_3, \dots e_m,$$

we obtain $e_2 \mathop{\smile}\limits^{|} r_{<j}, \xi_e, r_j, r_{\leq j}$, hence $e_2 \mathop{\smile}\limits^{|}_{r_j} r_{<j}, \xi_e, r_{\leq j}$. As $b_1 \in \operatorname{acl}(e_2, r_j)$ we have $b_1 \mathop{\smile}\limits^{|}_{r_j} r_{<j}, \xi_e, r_{\leq j}$. Using $b_1 \mathop{\smile}\limits^{|} r_j$ we conclude

$$b_1 \mathop{\smile}\limits^{|} r_{<j}, \xi_e, r_j, r_{\leq j}. \tag{4.6}$$

It follows from (4.5) and (4.6) that

$$\tilde{f}_{r_{<j}} \circ \tilde{f}_{\xi_e}^{-1} \circ \tilde{f}_{r_j} = \tilde{f}_{r_{\leq j}},$$

and hence

$$\Big((\tilde{f}_{\xi_e}^{-1} \circ \tilde{f}_{r_{<j}}) \circ (\tilde{f}_{\xi_e}^{-1} \circ \tilde{f}_{r_j})\Big) = \tilde{f}_{\xi_e}^{-1} \circ \tilde{f}_{r_{\leq j}}. \tag{4.7}$$

As noted at the beginning of the proof, we have that $r_j, r_{<j}, r_{\leq j} \models q(y)|_{\xi_e}$, and we define $c_0, c_1, c_2 \models s(z)|_{\xi_e}$ as follows:

$$c_0 := \pi(r_{<j}) = \ulcorner f_{\xi_e}^{-1} \circ f_{r_{<j}} \urcorner,$$
$$c_1 := \pi(r_j) = \ulcorner f_{\xi_e}^{-1} \circ f_{r_j} \urcorner,$$
$$c_2 := \pi(r_{\leq j}) = \ulcorner f_{\xi_e}^{-1} \circ f_{r_{\leq j}} \urcorner.$$

By (4.7), to conclude that $c_2 = c_0 \cdot c_1$ in G and finish the proof of the lemma it is sufficient to show that $c_0 \mathop{\smile}\limits^{|} c_1$.

As $r_{<j} \in \operatorname{acl}(\bar{a}_{<j}, \bar{e}_{\geq j})$, $r_j, \xi_e \in \operatorname{acl}(\bar{e}, a_j)$, and $\{e_3, \dots, e_m, a_j, \bar{a}_{<j}\}$ is an independent set, we have $r_{<j} \mathop{\smile}\limits^{|}_{\bar{e}_{\geq j}} r_j \xi_e$. Since $r_{<j} \mathop{\smile}\limits^{|} \bar{e}_{\geq j}$ (as $(r_{<j}, \bar{a}_{<j}, \bar{e}_{\geq j})$ is an $(m-1)$-gon) we also have $r_{<j} \mathop{\smile}\limits^{|}_{\xi_e} r_j$. It follows then that $c_0 \mathop{\smile}\limits^{|}_{\xi_e} c_1$. Since, by Lemma 4.16, $c_0 \mathop{\smile}\limits^{|} \xi_e$ we have $c_0 \mathop{\smile}\limits^{|} c_1$.

This concludes the proof of Lemma 4.19. □

Corollary 4.20. *For any $j \in \{4, \dots, m\}$, let $r_{\leq j} \models q(y)|_{\xi_e}$ with*

$$\models t_\xi(r_{\leq j}, \bar{a}_{\leq j}, \bar{e}_{>j}).$$

Then there exist $r_{<j}, r_j \models q(y)|_{\xi_e}$ such that

$$\models t_\xi(r_{<j}, \bar{a}_{<j}, \bar{e}_{\geq j}), \models t_\xi(r_j, \bar{e}_{<j}, a_j, \bar{e}_{>j})$$

and $\pi(r_{\leq j}) = \pi(r_j) \cdot \pi(r_{<j})$.

Proof. It is sufficient to show that for any r, r' with $\models t_\xi(r, \bar{a}_{\leq j}, \bar{e}_{>j})$, $\models t_\xi(r', \bar{a}_{\leq j}, \bar{e}_{>j})$ we have $r\bar{a}\bar{e} \equiv r'\bar{a}\bar{e}$. Indeed, given any $(r'_{\leq j}, r'_j, r'_{>j})$ satisfying the conclusion of Lemma 4.19, we then have an automorphism σ of $\mathbb{M}$ fixing $\bar{a}\bar{e}$ with $\sigma(r'_{\leq j}) = r_{\leq j}$; as the map π is relatively definable over $\operatorname{acl}(\bar{e})$, it then follows that $r_{<j} := \sigma(r'_{<j}), r_j := \sigma(r'_j)$ satisfy the requirements.

We have $r\bar{a}_{\leq j}\bar{e}_{>j} \equiv r'\bar{a}_{\leq j}\bar{e}_{>j}$. As $\bar{e} \mathop{\smile\!\!\!\!\!\mid}\, \bar{a}$ and each of $\bar{e}, \bar{a}$ is an $(m-2)$-tuple from the corresponding m-gon, we get $\bar{a}_{\leq j}\bar{e}_{>j} \mathop{\smile\!\!\!\!\!\mid}\, \bar{a}_{>j}\bar{e}_{\leq j}$. Also $r, r' \in \operatorname{acl}(\bar{a}_{\leq j}\bar{e}_{>j})$, as any realization of t_ξ is an $(m-1)$-gon, hence

$$rr'\bar{a}_{\leq j}\bar{e}_{>j} \mathop{\smile\!\!\!\!\!\mid}\, \bar{a}_{>j}\bar{e}_{\leq j}.$$

As all types over the empty set are stationary, we conclude $r\bar{a}\bar{e} \equiv r'\bar{a}\bar{e}$. □

We can now finish the proof of Proposition 4.18.

Proof of Proposition 4.18. We start with $r_{\leq m} := \xi$. Applying Corollary 4.20 with $j := m$, we obtain r_m and $r_{<m}$ with $\pi(\xi) = \pi(r_m) \cdot \pi(r_{<m})$.

Applying Corollary 4.20 again with $j := m-1$ and $r_{\leq m-1} := r_{<m}$ we obtain r_{m-1} and $r_{<m-1}$ with $\pi(\xi) = \pi(r_m) \cdot \pi(r_{m-1}) \cdot \pi(r_{<m-1})$.

Continuing this process with $j := m-2, \ldots, 4$ we obtain some

$$r_{m-2}, \ldots, r_4, r_{<4}$$

with $\pi(\xi) = \pi(r_m) \cdot \ldots \cdot \pi(r_4) \cdot \pi(r_{<4})$. We take $r_3 := r_{<4}$, which concludes the proof of the proposition. □

Proposition 4.21. *There exist $r_1, r_2 \models q(y)|_{\xi_e}$ such that $f_{r_1}(a_1) = e_2$, $f_{r_2}(e_1) = a_2$ and $\pi(r_2) \cdot \pi(r_1) = \pi(\xi)$.*

Proof. We choose $r_1 \models q(y)$ with $f_{r_1}(a_1) = e_2$ (possible by generic transitivity: as $a_1 \mathop{\smile\!\!\!\!\!\mid}\, e_2$, hence $a_1 e_2 \equiv a_1 a_2$ by stationarity of types over $\emptyset$; and as $f_\xi(a_1) = a_2$, we can take r_1 to be the image of ξ under the automorphism of $\mathbb{M}$ sending (a_1, a_2) to (a_1, e_2)). We also have $r_1 \mathop{\smile\!\!\!\!\!\mid}\, \xi_e$ ($a_1 \mathop{\smile\!\!\!\!\!\mid}\, \vec{e}$ and $e_2 \mathop{\smile\!\!\!\!\!\mid}\, \bar{e}$ by the choice of $\vec{e}$, so $a_1 e_2 \mathop{\smile\!\!\!\!\!\mid}\, \bar{e}$; as $r_1 \in \operatorname{acl}(a_1, e_2), \xi_e \in \operatorname{acl}(\bar{e})$, we conclude $r_1 \mathop{\smile\!\!\!\!\!\mid}\, \xi_e$), hence $r_1 \models q|_{\xi_e}$ by stationarity again.

Similarly $\xi \mathop{\smile\!\!\!\!\!\mid}\, \xi_e r_1$, hence $\xi \mathop{\smile\!\!\!\!\!\mid}\, \lceil f_{r_1}^{-1} \circ f_{\xi_e} \rceil$. By Lemma 4.16 we also have $r_1 \mathop{\smile\!\!\!\!\!\mid}\, \lceil f_{r_1}^{-1} \circ f_{\xi_e} \rceil$. By stationarity of q this implies $\xi \equiv_{\lceil f_{r_1}^{-1} \circ f_{\xi_e} \rceil} r_1$, so there exists some $r_2 \models q$ such that $\xi r_2 \equiv_{\lceil f_{r_1}^{-1} \circ f_{\xi_e} \rceil} r_1 \xi_e$. Hence

$$\tilde{f}_\xi^{-1} \circ \tilde{f}_{r_2} = \tilde{f}_{r_1}^{-1} \circ \tilde{f}_{\xi_e},$$

equivalently

$$\tilde{f}_{r_2} = \tilde{f}_\xi \circ \tilde{f}_{r_1}^{-1} \circ \tilde{f}_{\xi_e}. \tag{4.8}$$

In particular, $r_2 \in \operatorname{acl}(\xi, r_1, \xi_e)$.

We claim that $e_1 \mathop{\smile\!\!\!\!\!\mid}\, r_2 \xi r_1 \xi_e$. Since $\xi_e \in \operatorname{acl}(e_1, e_2)$, $r_1 \in \operatorname{acl}(a_1, e_2)$ and $\{a_1, e_1, e_2\}$ is an independent set, we have $r_1 \mathop{\smile\!\!\!\!\!\mid}_{e_2} e_1 \xi_e$. Using $r_1 \mathop{\smile\!\!\!\!\!\mid}\, e_2$ we deduce $r_1 \mathop{\smile\!\!\!\!\!\mid}\, e_1 \xi_e$. As $\xi_e \mathop{\smile\!\!\!\!\!\mid}\, e_1$, it implies that $\{r_1, e_1, \xi_e\}$ is an independent set. We have $r_1, e_1, \xi_e \in \operatorname{acl}(a_1, e_1, e_2)$ and $\xi \in \operatorname{acl}(a_1, a_2)$. Using independence of a_1, a_2, e_1, e_2 we obtain $\xi \mathop{\smile\!\!\!\!\!\mid}_{a_1} e_1\xi_e r_1$. Since $\xi \mathop{\smile\!\!\!\!\!\mid} a_1$, we have that $\xi \mathop{\smile\!\!\!\!\!\mid} e_1 r_1 \xi_e$, hence $\{\xi, e_1, r_1, \xi_e\}$ is an independent set and $e_1 \mathop{\smile\!\!\!\!\!\mid} \xi r_1 \xi_e$. As $r_2 \in \operatorname{acl}(\xi, r_1, \xi_e)$ we can conclude $e_1 \mathop{\smile\!\!\!\!\!\mid} r_2 \xi r_1 \xi_e$.

It then follows from (4.8) that

$$f_{r_2}(e_1) = (f_\xi \circ f_{r_1}^{-1} \circ f_{\xi_e})(e_1) = a_2,$$

so $f_{r_2}(e_1) = a_2$.

It also follows from (4.8) that

$$\Big((\tilde{f}_{\xi_e}^{-1} \circ \tilde{f}_{r_2}) \circ (\tilde{f}_{\xi_e}^{-1} \circ \tilde{f}_{r_1}) \Big) = \tilde{f}_{\xi_e}^{-1} \circ \tilde{f}_{\xi}.$$

We let

$$c_1 := \pi(r_1) = \ulcorner f_{\xi_e}^{-1} \circ f_{r_1} \urcorner \text{ and } c_2 := \pi(r_2) = \ulcorner f_{\xi_e}^{-1} \circ f_{r_2} \urcorner.$$

To show that $c_2 \cdot c_1 = \pi(\xi)$ and finish the proof of the proposition it is sufficient to show that $c_1 \mathop{\smile\!\!\!\!\!\mid} c_2$.

Since $r_1 \in \operatorname{acl}(a_1, e_2)$, $r_2 \in \operatorname{acl}(e_1, a_2)$ (by Remark 4.15, as by the above we have $r_2 \models q, e_1 \mathop{\smile\!\!\!\!\!\mid} r_2$ and $f_{r_2}(e_1) = a_2$) and $\xi_e \in \operatorname{acl}(e_1, e_2)$, we obtain $r_1 \mathop{\smile\!\!\!\!\!\mid}_{e_2} r_2 \xi_e$. Using $r_1 \mathop{\smile\!\!\!\!\!\mid} e_2$ we deduce $r_1 \mathop{\smile\!\!\!\!\!\mid} r_2 \xi_e$, hence $r_1 \mathop{\smile\!\!\!\!\!\mid}_{\xi_e} r_2$. It follows then that $c_1 \mathop{\smile\!\!\!\!\!\mid}_{\xi_e} c_2$ and, as $c_1 \mathop{\smile\!\!\!\!\!\mid} \xi_e$, we obtain $c_1 \mathop{\smile\!\!\!\!\!\mid} c_2$. ☐

Combining Propositions 4.21 and 4.18, we obtain some $r_1, \ldots, r_m \models q(y)|_{\xi_e}$ such that each r_i is inter-algebraic with a_i over $\{e_1, \ldots, e_m\}$ and

$$\pi(r_2) \cdot \pi(r_1) = \pi(r_m) \cdot \ldots \cdot \pi(r_3).$$

Obviously each r_i is also inter-algebraic over $\{e_1, \ldots, e_m\}$ with $\pi(r_i)$.

Thus, after a base change to $\{e_1, \ldots, e_m\}$ and inter-algebraically replacing a_1 with $\pi(r_1)^{-1}$, a_2 with $\pi(r_2)^{-1}$, and a_i with $\pi(r_i)$ for $i \in \{3, \ldots, m\}$, and using that permuting the elements of an abelian m-gon we still obtain an abelian m-gon, we achieve the following.

Reduction 3. *We may assume that $a_1, \ldots, a_m$ realize the generic type $s(z)$ of a connected group G that is type-definable over the empty set, with $a_1 \cdot a_2 \cdot a_m \cdot \ldots \cdot a_3 = 1_G$.*

To finish the proof of Theorem 4.6 it only remains to show that the group G is abelian. We deduce it from the Abelian Group Configuration Theorem, more precisely [9, Lemma C.1].

Claim 4.22. *Let G be a connected group type-definable over the empty set, $m \geq 4$ and $g_1, \ldots, g_m$ are generic elements of G such that $g_1, \ldots, g_m$ form an abelian m-gon and $g_1 \cdot \ldots \cdot g_m = 1_G$. Then the group G is abelian.*

Proof. Let $B := \operatorname{acl}(g_5, \ldots, g_m)$. We have that $g_1, \ldots, g_4$ are generics of G over B, and they form an abelian 4-gon over B. Since g_4 is inter-algebraic over B with $g_1 \cdot g_2 \cdot g_3$, we have that $g_1, g_2, g_3, g_1 \cdot g_2 \cdot g_3$ form an abelian 4-gon over B. Let $D := \mathrm{acl}_B(g_1, g_3) \cap \mathrm{acl}_B(g_2, g_1 \cdot g_2 \cdot g_3)$. We have $g_1, g_3 \mathop{\smile\!\!\!\!\!\!\mid}_D g_2, g_1 \cdot g_2 \cdot g_3$, hence

$$g_1 \cdot g_2 \cdot g_3 \in \mathrm{acl}_B(g_2, D) = \mathrm{acl}_B\left(g_2, \mathrm{acl}_B(g_1, g_3) \cap \mathrm{acl}_B(g_2, g_1 \cdot g_2 \cdot g_3)\right).$$

By [9, Lemma C.1], the group G is abelian. □

5. Main theorem in the stable case

Throughout the section we work in a complete theory T in a language $\mathcal{L}$. We fix an $|\mathcal{L}|^+$-saturated model $\mathcal{M} = (M, \ldots)$ of T, and also choose a large saturated elementary extension $\mathbb{M}$ of $\mathcal{M}$. We say that a subset A of $\mathcal{M}$ is *small* if $|A| \leq |\mathcal{L}|$. Given a definable set X in $\mathcal{M}$, we will often view it as a definable subset of $\mathbb{M}$, and sometimes write explicitly $X(\mathbb{M})$ to denote the set of tuples in $\mathbb{M}$ realizing the formula defining X.

5.1. **On the notion of $\mathfrak{p}$-dimension.** We introduce a basic notion of dimension in an arbitrary theory imitating the topological definition of dimension in o-minimal structures, but localized at a given tuple of commuting definable global types. We will see that it enjoys definability properties that may fail for Morley rank even in nice theories such as DCF_0.

Definition 5.1. If X is a definable set in $\mathcal{M}$ and $\mathcal{F}$ is a family of subsets of X, we say that $\mathcal{F}$ is a *definable family* (over a set of parameters A) if there exists a definable set Y and a definable set $D \subseteq X \times Y$ (both defined over A) such that $\mathcal{F} = \{D_b : b \in Y\}$, where $D_b = \{a \in X : (a, b) \in D\}$ is the fiber of D at b.

Definition 5.2. (1) By a $\mathfrak{p}$-*pair* we mean a pair $(X, \mathfrak{p}_X)$ where X is an $\emptyset$-definable set and $\mathfrak{p}_X \in S(\mathcal{M})$ is an $\emptyset$-definable stationary type on X.

(2) Given $s \in \mathbb{N}$, we say that $(X_i, \mathfrak{p}_i)_{i \in [s]}$ is a $\mathfrak{p}$-*system* if each $(X_i, \mathfrak{p}_i)$ is a $\mathfrak{p}$-pair and the types $\mathfrak{p}_1, \ldots, \mathfrak{p}_s$ commute, i.e. $\mathfrak{p}_i \otimes \mathfrak{p}_j = \mathfrak{p}_j \otimes \mathfrak{p}_i$ for all $i, j \in [s]$.

Example 5.3. Assume T is a stable theory, $(\mathfrak{p}_i)_{i \in [s]}$ are arbitrary types over $\mathcal{M}$ and $X_i \in \mathfrak{p}_i$ are arbitrary definable sets. By local character we can choose a model $\mathcal{M}_0 \preceq \mathcal{M}$ with $|\mathcal{M}_0| \leq |\mathcal{L}|$ such that each $\mathfrak{p}_i$ is definable (and stationary) over $\mathcal{M}_0$ and $X_i, i \in [s]$ are definable over $\mathcal{M}_0$. The types $(\mathfrak{p}_i)_{i \in [s]}$ automatically commute in a stable theory. Hence, naming the elements of $\mathcal{M}_0$ by constants, we obtain a $\mathfrak{p}$-system.

Assume now that $(X_i, \mathfrak{p}_i)_{i \in [s]}$ is a $\mathfrak{p}$-system. Given $u \subseteq [s]$, we let $\pi_u : \prod_{i \in [s]} X_i \to \prod_{i \in u} X_i$ be the projection map. For $i \in [s]$, we let $\pi_i := \pi_{\{i\}}$. Given $u, v \subseteq [s]$ with $u \cap v = \emptyset$, $a = (a_i : i \in u) \in \prod_{i \in u} X_i$ and $b = (b_i : i \in v) \in \prod_{i \in v} X_i$, we write $a \oplus b$ to denote the tuple $c = (c_i : i \in u \cup v) \in \prod_{i \in u \cup v} X_i$ with $c_i = a_i$ for $i \in u$ and $c_i = b_i$ for $i \in v$. Given $Y \subseteq \prod_{i \in [s]} X_i$, $u \subseteq [s]$ and $a \in \prod_{i \in u} X_i$, we write $Y_a := \{b \in \prod_{i \in [s] \setminus u} X_i : a \oplus b \in Y\}$ to denote the fiber of Y above a.

Example 5.4. If $\mathcal{F}$ is a definable family of subsets of $\prod_{i\in[s]} X_i$ and $u \subseteq [s]$, then $\{\pi_u(F) : F \in \mathcal{F}\}$ and $\left\{F_a : F \in \mathcal{F}, a \in \prod_{i\in[s]\setminus u} X_i\right\}$ are definable families of subsets of $\prod_{i\in u} X_i$ (over the same set of parameters).

Definition 5.5. Let $\bar{a} = (a_1, \ldots, a_s) \in X_1 \times \cdots \times X_s$ and A a small subset of $\mathcal{M}$.

(1) We say that $\bar{a}$ *is* $\mathfrak{p}$*-generic in* $X_1 \times \cdots \times X_s$ *over* A if $(a_1, \ldots, a_s) \models \mathfrak{p}_1 \otimes \cdots \otimes \mathfrak{p}_s{\restriction}A$.
(2) (a) For $k \leq s$ we write $\dim_{\mathfrak{p}}(\bar{a}/A) \geq k$ if for *some* $u \subseteq [s]$ with $|u| \geq k$ the tuple $\pi_u(\bar{a})$ is $\mathfrak{p}$-generic (with respect to the corresponding $\mathfrak{p}$-system $\{(X_i, \mathfrak{p}_i) : i \in u\}$).
  (b) As usual, we define $\dim_{\mathfrak{p}}(\bar{a}/A) = k$ if $\dim_{\mathfrak{p}}(\bar{a}/A) \geq k$ and it is not true that $\dim_{\mathfrak{p}}(\bar{a}/A) \geq k+1$.
(3) If $q(\bar{x}) \in S(A)$ and $q(\bar{x}) \vdash \bar{x} \in X_1 \times \ldots \times X_s$, we write $\dim_{\mathfrak{p}}(q) := \dim_{\mathfrak{p}}(\bar{a}/A)$ for some (equivalently, any) $\bar{a} \models q$.
(4) For a subset $Y \subseteq X_1 \times \cdots \times X_s$ definable over A, we define
$$\begin{aligned}\dim_{\mathfrak{p}}(Y) &:= \max\{\dim_{\mathfrak{p}}(\bar{a}/A)\colon \bar{a} \in Y\}\\ &= \max\{\dim_{\mathfrak{p}}(q)\colon q \in S(A), Y \in q\},\end{aligned}$$
note that this does not depend on the set A such that Y is A-definable.
(5) As usual, for a definable subset $Y \subseteq X_1 \times \cdots \times X_s$ we say that Y is *a* $\mathfrak{p}$*-generic subset of* $X_1 \times \cdots \times X_s$ if $\dim_{\mathfrak{p}}(Y) = s$ (equivalently, Y is contained in $\mathfrak{p}_1 \otimes \cdots \otimes \mathfrak{p}_s$.)

If $A = \emptyset$ we will omit it.

Remark 5.6. It follows from the definition that for a definable set $Y \subseteq X_1 \times \cdots \times X_s$, $\dim_{\mathfrak{p}}(Y)$ is the maximal k such that the projection of Y onto some k coordinates is $\mathfrak{p}$-generic. As usual, for a definable $Y \subseteq X_1 \times \cdots \times X_s$ and small $A \subseteq \mathcal{M}$ we say that an element $a \in Y$ is *generic in* Y *over* A if $\dim_{\mathfrak{p}}(a/A) = \dim_{\mathfrak{p}}(Y)$.

Remark 5.7. It also follows that if $\mathcal{N} \succeq \mathcal{M}$ is an arbitrary $|\mathcal{L}|^+$-saturated model and $\mathfrak{p}_i' := \mathfrak{p}_i|_{\mathcal{N}} \in S(\mathcal{N})$ is the unique definable extension, for $i \in [s]$, then $(X_i(\mathcal{N}), \mathfrak{p}_i')_{i\in[s]}$ is a $\mathfrak{p}$-system in $\mathcal{N}$, and for every definable subset $Y \subseteq X_1 \times \ldots \times X_s$ in $\mathcal{M}$ we have $\dim_{\mathfrak{p}}(Y) = \dim_{\mathfrak{p}}(Y(\mathcal{N}))$, where the latter is calculated in $\mathcal{N}$ with respect to this $\mathfrak{p}$-system.

Claim 5.8. *Let* $\mathcal{F}$ *be a definable (over* A*) family of subsets of* $X_1 \times \cdots \times X_s$ *and* $k \leq s$*. Then the family*
$$\{F \in \mathcal{F}\colon \dim_{\mathfrak{p}}(F) = k\}$$
is definable (over A *as well).*

Proof. Assume that $\mathcal{F} = \{D_b : b \in Y\}$ for some definable Y and definable $D \subseteq (X_1 \times \ldots \times X_s) \times Y$. Given $0 \leq k \leq s$, let $Y_k := \{b \in Y : \dim_{\mathfrak{p}}(D_b) = k\}$, it suffices to show that Y_k is definable. As every $\mathfrak{p}_i$ is definable, for every $u \subseteq [s]$, the type $\mathfrak{p}_u = \bigotimes_{i\in u} \mathfrak{p}_i$ is also definable. In particular, there is a definable (over any set of parameters containing the parameters of Y and D) set $Z_u \subseteq Y$ such that for any $b \in Y$, $\pi_u(D_b) \in \mathfrak{p}_u \iff b \in Z_u$. Then Y_k is definable as

$$Y_k = \left(\bigvee_{u \subseteq [s], |u|=k} b \in Z_u \right) \wedge \left(\bigwedge_{u \subseteq [s], |u|>k} b \notin Z_u \right). \qquad \square$$

The following lemma shows that $\mathfrak{p}$-dimension is "super-additive".

Lemma 5.9. *Let $Y \subseteq X_1 \times \cdots \times X_s$ be definable and $u \subseteq [s]$. Assume that $0 \leq n \leq [s]$ is such that for every $a \in \pi_u(Y)$ we have $\dim_{\mathfrak{p}}(Y_a) \geq n$. Then $\dim_{\mathfrak{p}}(Y) \geq \dim_{\mathfrak{p}}(\pi_u(Y)) + n$.*

Proof. Assume that Y is definable over a small set of parameters A, and let $m := \dim_{\mathfrak{p}}(\pi_u(Y))$. Then there is some $u^* \subseteq u, |u^*| = m$ such that

$$\pi_{u^*}(Y)(\pi_u(Y)) = \pi_{u^*}(Y) \in \mathfrak{p}_{u^*} = \bigotimes_{i \in u^*} \mathfrak{p}_i.$$

Let $b_{u^*} = (b_i : i \in u^*) \models \mathfrak{p}_{u^*}|_A$. As $b_{u^*} \in \pi_{u^*}(Y)(\pi_u(Y))$, there exist some $(b_i : i \in u \setminus u^*)$ so that $b_u := (b_i : i \in u) \in \pi_u(Y)$. Then by assumption $\dim_{\mathfrak{p}}(Y_{b_u}) \geq n$, that is for some $v^* \subseteq v := [s] \setminus u$ with $|v^*| \geq n$ we have $\pi_{v^*}(Y_{b_u}) \in \mathfrak{p}_{v^*} := \bigotimes_{i \in v^*} \mathfrak{p}_i$. Let $b_{v^*} = (b_i : i \in v^*) \models \mathfrak{p}_{v^*}|_{Ab_u}$, and let $w := u^* \sqcup v^*$. Since the types $(\mathfrak{p}_i : i \in w)$ are stationary and commuting, it follows that $b_w := (b_i : i \in w) \models \mathfrak{p}_w|_A$ for $p_w := \bigotimes_{i \in u^* \sqcup v^*} \mathfrak{p}_i$. As $b_{v^*} \in \pi_{v^*}(Y_{b_u})$, there exists some $(b_i : i \in v \setminus v^*)$ so that $(b_i : i \in v) \in Y_{b_u}$, hence $(b_i : i \in [s]) \in Y$. Thus $b_w \in \pi_w(Y)$, hence $\pi_w(Y) \in \mathfrak{p}_w$, and $|w| \geq m + n$ — which shows that $\dim_{\mathfrak{p}}(Y) \geq m + n$, as required. $\square$

5.2. Fiber-algebraic relations and $\mathfrak{p}$-irreducibility.

Definition 5.10. Given a definable set $Y \subseteq \prod_{i \in [s]} X_i$ and a small set of parameters $C \subseteq \mathcal{M}$ so that Y is defined over C, we say that Y is $\mathfrak{p}$-*irreducible over* C if there do not exist disjoint sets Y_1, Y_2 definable over C with $Y = Y_1 \cup Y_2$ and $\dim_{\mathfrak{p}}(Y_1) = \dim_{\mathfrak{p}}(Y_2) = \dim_{\mathfrak{p}}(Y)$.

We say that Y is *absolutely* $\mathfrak{p}$-*irreducible* if it is irreducible over any small set $C \subseteq \mathcal{M}$ such that Y is defined over C.

Remark 5.11. It follows from the definition of $\mathfrak{p}$-dimension that a definable set $Y \subseteq X_1 \times \ldots \times X_s$ is $\mathfrak{p}$-irreducible over C if and only if any two tuples generic in Y over C have the same type over C.

Lemma 5.12. *If $Q(\bar{x}) \subseteq X_1 \times \ldots \times X_s$ is fiber-algebraic of degree $\leq d$, then the set*

$$\{q \in S_{\bar{x}}(\mathcal{M}) : Q \in q \text{ and } \dim_{\mathfrak{p}}(q) \geq s - 1\}$$

has cardinality at most sd.

Proof. Assume towards a contradiction that $q_1, \ldots, q_{sd+1}$ are pairwise different types in this set. Then there exist some formulas $\psi_i(\bar{x})$ with parameters in $\mathcal{M}$ such that $\psi_i(\bar{x}) \in q_i$ and $\psi_i(\bar{x}) \to \neg\psi_j(\bar{x})$ for all $i \neq j \in [sd+1]$. Let $C \subseteq \mathcal{M}$ be the (finite) set of the parameters of Q and $\psi_i, i \in [sd+1]$. For each $i \in [sd+1]$, as $(\psi_i(\bar{x}) \wedge Q(\bar{x})) \in q_i$, we have $\dim_{\mathfrak{p}}(\psi_i(\bar{x}) \wedge Q(\bar{x})) \geq s-1$, which by definition of $\mathfrak{p}$-dimension implies $\exists x_k\,(\psi_i(\bar{x}) \wedge Q(\bar{x})) \in \bigotimes_{\ell \in [s]\setminus\{k\}} \mathfrak{p}_\ell$ for at least one $k \in [s]$. By pigeonhole, there must exist some $k' \in [s]$ and some $u \subseteq [sd+1]$ such that $|u| \geq d+1$ and $\exists x_{k'}\,(\psi_i(\bar{x}) \wedge Q(\bar{x})) \in \bigotimes_{\ell \in [s]\setminus\{k'\}} \mathfrak{p}_\ell$ for all $i \in u$. Now let $\bar{a} = (a_\ell : \ell \in [s] \setminus \{k'\})$ be a tuple in $\mathcal{M}$ satisfying $\bar{a} \models \left(\bigotimes_{\ell \in [s]\setminus\{k'\}} \mathfrak{p}_\ell\right)|_C$. By the choice of u, for each $i \in u$ there exists some b_i in $\mathcal{M}$ such that $(\psi_i \wedge Q)(a_1, \ldots, a_{k'-1}, b_i, a_{k'+1}, \ldots, a_s)$ holds. By the choice of the formulas ψ_i, the elements $(b_i : i \in u)$ are pairwise distinct, and $|u| > d$ — contradicting that Q is fiber-algebraic of degree d. □

Corollary 5.13. *Every fiber-algebraic $Q \subseteq X_1 \times \ldots \times X_s$ of degree $\leq d$ is a union of at most sd absolutely $\mathfrak{p}$-irreducible sets (which are then automatically fiber-algebraic, of degree $\leq d$).*

Proof. Let $(q_i : i \in [D])$ be an arbitrary enumeration of the set

$$\{q \in S_{\bar{x}}(\mathcal{M}) : Q \in q \wedge \dim_{\mathfrak{p}}(q) \geq s-1\},$$

we have $D \leq sd$ by Lemma 5.12. We can choose formulas $(\psi_i(\bar{x}) : i \in [D])$ with parameters over $\mathcal{M}$ such that $\psi_i(\bar{x}) \in q_i$ and $\psi_i(\bar{x}) \to \neg\psi_j(\bar{x})$ for all $i \neq j \in [D]$. Let $Q_i(\bar{x}) := Q(\bar{x}) \wedge \psi_i(\bar{x})$, then $Q = \bigsqcup_{i \in [D]} Q_i$ and each Q_i is absolutely $\mathfrak{p}$-irreducible (by Remark 5.11, as every generic tuple in Q_i over a small set C has the type $q_i|_C$). □

Lemma 5.14. *If $Q \subseteq \prod_{i \in [s]} X_i$ is $\mathfrak{p}$-irreducible over a small set of parameters C and $\dim_{\mathfrak{p}}(Q) = s-1$, then for any $i \in [s]$ and any tuple $\bar{a} = (a_j : j \in [s] \setminus \{i\})$ which is $\mathfrak{p}$-generic in $\prod_{j \in [s]\setminus\{i\}} X_j$ over C (i.e. $\bar{a} \models (\bigotimes_{j \in [s]\setminus\{i\}} \mathfrak{p}_j)|_C$), if $Q(a_1, \ldots, a_{i-1}, x_i, a_{i+1}, \ldots, a_s)$ is consistent then it implies a complete type over $C \cup \{a_j : j \in [s] \setminus \{i\}\}$.*

Proof. Otherwise there exist two types $r_t \in S_{x_i}(C\bar{a}), t \in \{1,2\}$ such that $r_1 \neq r_2$ and $Q(a_1, \ldots, a_{i-1}, x_i, a_{i+1}, \ldots, a_s) \in r_t$ for both $t \in \{1,2\}$. Then there exist some formulas $\varphi_t(\bar{x}), t \in \{1,2\}$ with parameters in C such that $\varphi_t(a_1, \ldots, a_{i-1}, x_i, a_{i+1}, \ldots, a_s) \in r_t$, $\varphi_1(\bar{x}) \to \neg\varphi_2(\bar{x})$ and $\varphi_2(\bar{x}) \to \neg\varphi_1(\bar{x})$. In particular, by assumption on $\bar{a}$,

$$\dim_{\mathfrak{p}}(Q(\bar{x}) \wedge \varphi_t(\bar{x})) \geq s-1$$

for both $t \in \{1,2\}$ — contradicting irreducibility of Q over C. □

5.3. **On general position.** We recall the notion of general position from Definition 1.5, specialized to the case of $\mathfrak{p}$-dimension.

Definition 5.15. Let $(X, \mathfrak{p})$ be a $\mathfrak{p}$-pair, and let $\mathcal{F}$ be a definable family of subsets of X. For $\nu \in \mathbb{N}$, we say that a set $A \subseteq X$ is in $(\mathcal{F}, \nu)$-*general position* if for every $F \in \mathcal{F}$ with $\dim_{\mathfrak{p}}(F) = 0$ we have $|A \cap F| \leq \nu$.

We extend this notion to cartesian products of $\mathfrak{p}$-pairs.

Definition 5.16. For sets $X_1 \times X_2 \times \cdots \times X_s$ and an integer $n \in \mathbb{N}$, by an n-*grid on* $X_1 \times \cdots \times X_s$ we mean a set of the form $A_1 \times A_2 \times \cdots \times A_s$ with $A_i \subseteq X_i$ and $|A_i| \leq n$ for all $i \in [s]$.

Definition 5.17. Let $s \in \mathbb{N}$ and $(X_i, \mathfrak{p}_i)$, $i \in [s]$, be $\mathfrak{p}$-pairs. Let $\vec{\mathcal{F}}$ be a definable system of subsets of (X_i), $i \in [s]$, i.e. $\vec{\mathcal{F}} = (\mathcal{F}_1, \ldots, \mathcal{F}_s)$ where each $\mathcal{F}_i$ is a definable family of subsets of X_i. For $\nu \in \mathbb{N}$, we say that a grid $A_1 \times \cdots \times A_s$ on $X_1 \times \cdots \times X_s$ is in $(\vec{\mathcal{F}}, \nu)$-*general position* if each A_i is in $(\mathcal{F}_i, \nu)$-general position.

We will need a couple of auxiliary lemmas bounding the size of the intersection of sets in a definable family with finite grids in terms of their $\mathfrak{p}$-dimension.

Lemma 5.18. *Let $s \in \mathbb{N}_{\geq 1}$, $(X_i, \mathfrak{p}_i)_{i \in [s]}$ a $\mathfrak{p}$-system, and $\mathcal{G}$ a definable family of subsets of $X_1 \times \cdots \times X_s$ such that $\dim_{\mathfrak{p}}(G) = 0$ for every $G \in \mathcal{G}$. Then there is a definable system of subsets $\vec{\mathcal{F}} = (\mathcal{F}_1, \ldots, \mathcal{F}_s)$ such that: for any finite grid $A = A_1 \times \cdots \times A_s$ on $X_1 \times \cdots \times X_s$ in $(\vec{\mathcal{F}}, \nu)$-general position and any $G \in \mathcal{G}$ we have $|G \cap A| \leq \nu^s$.*

Proof. Assume that $\mathcal{G}$ is a definable family of subsets $X_1 \times \cdots \times X_s$ with $\dim_{\mathfrak{p}}(G) = 0$ for all $G \in \mathcal{G}$. For $i \in [s]$ and $G \in \mathcal{G}$, we let $G_i := \pi_i(G)$, note that still $\dim_{\mathfrak{p}}(G_i) = 0$. Let $\mathcal{F}_i := \{G_i : G \in \mathcal{G}\}$, we claim that then $\vec{\mathcal{F}} := (\mathcal{F}_1, \ldots, \mathcal{F}_s)$ satisfies the requirements.

Indeed, let $A = A_1 \times \cdots \times A_s$ be a finite grid on $X_1 \times \cdots \times X_s$ in $(\vec{\mathcal{F}}, \nu)$-general position. Let $G \in \mathcal{G}$ be arbitrary. As $G_i \in \mathcal{F}_i$ with $\dim_{\mathfrak{p}}(G_i) = 0$, by assumption we have $|G_i \cap A_i| \leq \nu$ for every $i \in [s]$. As $G \subseteq \prod_{i \in [s]} G_i$, we have

$$G \cap \left(\prod_{i \in [s]} A_i \right) \subseteq \left(\prod_{i \in [s]} G_i \right) \cap \left(\prod_{i \in [s]} A_i \right) = \prod_{i \in [s]} (G_i \cap A_i),$$

hence $\left| G \cap \prod_{i \in [s]} A_i \right| \leq \nu^s$, as required. □

Lemma 5.19. *Let $s \in \mathbb{N}_{\geq 1}$ and $(X_i, \mathfrak{p}_i)_{i \in [s]}$ be a $\mathfrak{p}$-system, and $\mathcal{G}$ a definable family of subsets of $X_1 \times \cdots \times X_s$. Assume that for some $0 \leq k \leq s$ we have $\dim_{\mathfrak{p}}(G) \leq k$ for every $G \in \mathcal{G}$. Then there is a definable system $\vec{\mathcal{F}} = (\mathcal{F}_1, \ldots, \mathcal{F}_s)$ of subsets of $X_1 \times \ldots \times X_s$ such that: for any ν and any n-grid $A = A_1 \times \cdots \times A_s$ on $X_1 \times \cdots \times X_s$ in $(\vec{\mathcal{F}}, \nu)$-general position, for every $G \in \mathcal{G}$ we have $|G \cap A| \leq s^k \nu^{s-k} n^k$.*

Proof. Given $s \geq k$ and ν, we let $C(k,s,\nu)$ be the smallest number in $\mathbb{N}$ (if it exists) so that the bound $|G \cap A| \leq C(k,s,\nu)n^k$ holds (with respect to all possible $\mathfrak{p}$-systems $(X_i, \mathfrak{p}_i)_{i\in[s]}$ and definable families $\mathcal{G}$). We will show that $C(k,s,\nu) \leq s^k \nu^{s-k}$ for all $s \geq k \geq 0$ and ν.

For any $s \in \mathbb{N}_{\geq 1}$ and $k = 0$, the claim holds by Lemma 5.18 with $C(0,s,\nu) = \nu^s$. For any $s \in \mathbb{N}_{\geq 1}$ and $k = s$, the claim trivially holds with $C(s,s,\nu) = 1$ (and $\mathcal{F}_i = \emptyset, i \in [s]$).

We fix $s > k \geq 1$ and assume that the claim holds for all pairs $s' \geq k' \geq 0$ with either $s' < s$ or $k' < k$. Assume that $\dim_{\mathfrak{p}}(G) \leq k$ for every $G \in \mathcal{G}$. Given $G \in \mathcal{G}$, let $G' := \{g \in \pi_1(G) : \dim_{\mathfrak{p}}(G_g) \geq k\}$. Then $\mathcal{F}_1 := \{G' : G \in \mathcal{G}\}$ is a definable family of subsets of X_1 by Claim 5.8. By assumption and Lemma 5.9 we have $\dim_{\mathfrak{p}}(G') = 0$ for every $G \in \mathcal{G}$. Let

$$\mathcal{G}^* := \{G_g : G \in \mathcal{G} \wedge g \in \pi_1(G)\},$$
$$\mathcal{G}^*_{<k} := \{G_g : G \in \mathcal{G} \wedge g \in \pi_1(G) \wedge \dim_{\mathfrak{p}}(G_g) < k\}.$$

Both $\mathcal{G}^*$ and $\mathcal{G}^*_{<k}$ (by Claim 5.8) are definable families of subsets of $\prod_{2\leq i\leq s} X_i$, all sets in $\mathcal{G}^*$ have $\mathfrak{p}$-dimension $\leq k$, and all sets in $\mathcal{G}^*_{<k}$ have $\mathfrak{p}$-dimension $\leq k-1$. Applying the $(k,s-1)$-induction hypothesis, let $\vec{\mathcal{F}}^* = (\mathcal{F}^*_i : 2 \leq i \leq s)$ be a definable system of subsets of $X_2 \times \ldots \times X_s$ satisfying the conclusion of the lemma with respect to $\mathcal{G}^*$. Applying the $(k-1,s-1)$-induction hypothesis, let $\vec{\mathcal{F}}^*_{<k} = \left(\mathcal{F}^*_{<k,i} : 2 \leq i \leq s\right)$ be a definable system of subsets of $X_2 \times \ldots \times X_s$ satisfying the conclusion of the lemma with respect to $\mathcal{G}^*_{<k}$. We let $\vec{\mathcal{F}} = (\mathcal{F}_i : i \in [s])$ be a definable system of subsets of $X_1 \times \ldots \times X_s$, with $\mathcal{F}_1$ defined above and $\mathcal{F}_i := \mathcal{F}^*_i \cup \mathcal{F}^*_{<k,i}$ for $2 \leq i \leq s$.

Let now $\nu \in \mathbb{N}$ and $A = A_1 \times \cdots \times A_s$ be a finite grid on $X_1 \times \cdots \times X_s$ in $(\vec{\mathcal{F}}, \nu)$-general position. Let $G \in \mathcal{G}$ be arbitrary. As $G' \in \mathcal{F}_0$, we have in particular that $|G' \cap A_1| \leq \nu$, and by the choice of $\vec{\mathcal{F}}^*$, for every $g \in G' \cap A_1$ we have $|G_g \cap (A_2 \times \ldots \times A_s)| \leq C(k,s-1,\nu)n^k$. And by the choice of $\vec{\mathcal{F}}^*_{<k}$, for every $g \in A_1 \setminus G'$, we have $|G_g \cap (A_2 \times \ldots \times A_s)| \leq C(k-1,s-1,\nu)n^{k-1}$. Combining, we get

$$|G \cap (A_1 \times \ldots \times A_s)| \leq$$
$$\nu C(k,s-1,\nu)n^k + (n-\nu)C(k-1,s-1,\nu)n^{k-1} \leq$$
$$\big(\nu C(k,s-1,\nu) + C(k-1,s-1,\nu)\big)n^k.$$

This establishes a recursive bound on $C(k,s,\nu)$. Given $s \geq k \geq 1$, we can repeatedly apply this recurrence for $s, s-1, \ldots, k$, and using that $C(s,s,\nu) = 1$ for all s, ν we get that

$$C(k,s,\nu) \leq \nu^{s-k} + \sum_{i=1}^{s-k} \nu^{i-1} C(k-1, s-i, \nu)$$

for any $s \geq k \geq 1$. Using that $C(0, s, \nu) = \nu^s$ for all s, ν and iterating this inequality for $0, 1, \ldots, k$, it is not hard to see that $C(s, k, \nu) \leq s^k \nu^{s-k}$ for all s, k, ν. ☐

5.4. **Main theorem: the statement and some reductions.** From now on we will assume additionally that the theory T is stable and eliminates imaginaries, i.e. $T = T^{\mathrm{eq}}$ (we refer to e.g. [57] for a general exposition of stability). As before, $\mathcal{M}$ is an $|\mathcal{L}|^+$-saturated model of T, $\mathbb{M}$ is a monster model of T, and we assume that $(X_i, \mathfrak{p}_i)_{i \in [s]}$ is a $\mathfrak{p}$-system in $\mathcal{M}$, with each $\mathfrak{p}_i$ non-algebraic. "Definable" means "definable with parameters in $\mathcal{M}$". As usual, if X is a definable set, a family $\mathcal{F}$ of subsets of X is definable if there exist definable sets $Y, F \subseteq X \times Y$ so that $\mathcal{F} = \{F_b : b \in Y\}$.

Remark 5.20. Note that if $Q \subseteq X_1 \times \cdots \times X_s$ is a fiber-algebraic relation of degree d, then for any n-grid $A \subseteq \prod_{i \in [s]} X_i$ we have

$$|Q \cap A| \leq dn^{s-1} = O_d(n^{s-1}).$$

Definition 5.21. Let $\mathcal{Q}$ be a definable family of subsets of $X_1 \times \cdots \times X_s$.

(1) Given a real $\varepsilon > 0$, we say that $\mathcal{Q}$ admits *ε-power saving* if there exist definable families $\mathcal{F}_i$ on X_i, such that for $\vec{\mathcal{F}} = (\mathcal{F}_i)_{i \leq s}$ and any $\nu \in \mathbb{N}$, for any n-grid $A = A_1 \times \cdots \times A_s$ on $X_1 \times \cdots \times X_s$ in $(\vec{\mathcal{F}}, \nu)$-general position and any $Q \in \mathcal{Q}$ we have
$$|Q \cap A| = O_\nu\left(n^{(s-1)-\varepsilon}\right).$$
(2) We say that $\mathcal{Q}$ *admits power saving*[2] if it admits ε-power saving for some $\varepsilon > 0$.
(3) We say that a relation $Q \subseteq X_1 \times \ldots \times X_s$ admits (ε-)power saving if the family $\mathcal{Q} := \{Q\}$ does.
(4) We say that Q is *special* if it is fiber-algebraic and does not admit power-saving.

Lemma 5.1. Assume $\mathcal{Q}, \mathcal{Q}_1, \ldots, \mathcal{Q}_m$ are definable families of subsets of $X_1 \times \cdots \times X_s$ and $\varepsilon > 0$ is such that each $\mathcal{Q}_t$ satisfies ε-power saving. Assume that for every $Q \in \mathcal{Q}$, $Q = \bigcup_{t \in [m]} Q_t$ for some $Q_t \in \mathcal{Q}_t$. Then $\mathcal{Q}$ also satisfies ε-power saving.

Proof. Assume each $\mathcal{Q}_t, t \in [m]$ satisfies ε-power saving, i.e. there exist definable families $\mathcal{F}_{t,i}$ on X_i and functions $C_t : \mathbb{N} \to \mathbb{N}$ so that letting $\vec{\mathcal{F}}_t = (\mathcal{F}_{t,i})_{i \leq s}$, for every grid A in $\left(\vec{\mathcal{F}}_t, \nu\right)$-general position and every $Q_t \in \mathcal{Q}_t$ we have $|Q_t \cap A| \leq C_t(\nu) n^{(s-1)-\varepsilon}$. Let $\mathcal{F}_i := \bigcup_{t \in [m]} \mathcal{F}_{t,i}$, $\vec{\mathcal{F}} = (\mathcal{F}_i)_{i \leq s}$ and $C := \sum_{t \in [m]} C_t$. Then for every grid A in $\left(\vec{\mathcal{F}}, \nu\right)$-general position and every $Q \in \mathcal{Q}$ we have $|Q \cap A| \leq C(\nu) n^{(s-1)-\varepsilon}$, as required. ☐

We recall Definition 1.6, specializing to $\mathfrak{p}$-dimension.

[2] We are following the terminology in [8].

Definition 5.22. Let $Q \subseteq \prod_{i \in [s]} X_i$ be a definable relation and $(G, \cdot, 1_G)$ a type-definable group in $\mathcal{M}$ (over a small set of parameters A). We say that Q is in a $\mathfrak{p}$-*generic correspondence with* G (over A) if there exist elements $g_1, \ldots, g_s \in G(\mathcal{M})$ such that:

(1) $g_1 \cdot \ldots \cdot g_s = 1_G$;
(2) $g_1, \ldots, g_{s-1}$ are independent generics in G over A (in the usual sense of stable group theory);
(3) for each $i \in [s]$ there is a generic element $a_i \in X_i$ realizing $\mathfrak{p}_i|_A$ and inter-algebraic with g_i over A, such that $\mathcal{M} \models Q(a_1, \ldots, a_s)$.

Remark 5.23. If Q is $\mathfrak{p}$-irreducible over A, then (3) holds for all $g_1, \ldots, g_s \in G$ satisfying (1) and (2), providing a definable generic finite-to-finite correspondence between Q and the graph of the $(s-1)$-fold multiplication in G.

The following is the main theorem of the section characterizing special fiber-algebraic relations in stable reducts of distal structures.

Theorem 5.24. *Assume that $\mathcal{M}$ is an $|\mathcal{L}|^+$-saturated $\mathcal{L}$-structure, and* $\mathrm{Th}(\mathcal{M})$ *is stable and admits a distal expansion. Assume that $s \geq 3$, $(X_i, \mathfrak{p}_i)_{i \in [s]}$ is a $\mathfrak{p}$-system with each $\mathfrak{p}_i$ non-algebraic and $Q \subseteq X_1 \times \cdots \times X_s$ is a definable fiber-algebraic relation. Then at least one of the following holds.*

(1) Q admits power saving.
(2) Q is in a $\mathfrak{p}$-generic correspondence with an abelian group G type-definable in $\mathcal{M}^{\mathrm{eq}}$ over a set of parameters of cardinality $\leq |\mathcal{L}|$.

The only property of distal structures actually used is that every definable binary relation in $\mathcal{M}$ satisfies the γ-ST property (Definition 2.12) for some $\gamma > 0$, by Proposition 2.14 and Fact 2.4. In fact, Theorem 5.24 follows from the following more precise version with the additional uniformity in families and explicit bounds on power saving.

Definition 5.25. Let $\mathcal{Q}$ be a definable family of subsets $X_1 \times \cdots \times X_s$.

(1) We say that $\mathcal{Q}$ is a *fiber-algebraic* family if each $Q \in \mathcal{Q}$ is fiber-algebraic.
(2) We say that $\mathcal{Q}$ is an *absolutely* $\mathfrak{p}$-*irreducible fiber-algebraic family* if each $Q \in \mathcal{Q}$ is $\mathfrak{p}$-irreducible and fiber-algebraic

Remark 5.26. Let $\mathcal{Q}$ be a definable fiber-algebraic family. By saturation of $\mathcal{M}$ there is $d \in \mathbb{N}$ such that every $Q \in \mathcal{G}$ has degree $\leq d$. In this case we say that $\mathcal{Q}$ is *of degree* $\leq d$.

Theorem 5.27. *Assume that $\mathcal{M}$ is an $|\mathcal{L}|^+$-saturated $\mathcal{L}$-structure and* $\mathrm{Th}(\mathcal{M})$ *is stable. Assume that $s \geq 4$, $(X_i, \mathfrak{p}_i)_{i \in [s]}$ is a $\mathfrak{p}$-system with each $\mathfrak{p}_i$ non-algebraic, and let $\mathcal{Q}$ be a fiber-algebraic definable family, and fix $0 < \gamma \leq 1$.*

- *If $s \geq 4$, assume that there exist $m \in \mathbb{N}$ and definable families $\mathcal{Q}_i, i \in [m]$ of absolutely $\mathfrak{p}$-irreducible sets so that for every $Q \in \mathcal{Q}$ we have $Q = \bigcup_{i \in [m]} Q_i$ for some $Q_i \in \mathcal{Q}_i$. Assume also that for each $i \in [m], t_1 \neq t_2 \in [s]$, the family $\mathcal{Q}_i$ viewed as a definable family of subsets of $(X_{t_1} \times X_{t_2}) \times \left(\prod_{k \in [s] \setminus \{t_1, t_2\}} X_k\right)$ satisfies the γ-ST property.*

- *If $s = 3$, for each $i \in [m]$ and $\mathcal{Q}_i$ as above, we additionally consider the definable family $\mathcal{Q}_i^* := \{Q^* : Q \in \mathcal{Q}_i\}$ of subsets of $X_1 \times X_2 \times X_3 \times X_4$, where*

$$Q^* := \Big\{(x_2, x_2', x_3, x_3') \in X_2 \times X_2 \times X_3 \times X_3 : \\ \exists x_1 \in X_1 \left((x_1, x_2, x_3) \in Q \wedge (x_1, x_2', x_3') \in Q\right)\Big\}.$$

Assume moreover that there exist $m_i \in \mathbb{N}, i \in [m]$ and definable families $\mathcal{Q}_{i,j}$ for $i \in [m], j \in [m_i]$ so that for every $i \in [m], Q^* \in \mathcal{Q}_i^*$ we have $Q^* = \bigcup_{j \in [m_i]} Q_{i,j}$ for some $Q_{i,j} \in \mathcal{Q}_{i,j}$. Assume also that for each $i \in [m], j \in [m_i], t_1 \neq t_2 \in [4]$, the family $\mathcal{Q}_{i,j}$ viewed as a definable family of subsets of $(X_{t_1} \times X_{t_2}) \times \left(\prod_{k \in [4] \setminus \{t_1, t_2\}} X_k\right)$ satisfies the 2γ-ST property.

Then there is a definable subfamily $\mathcal{Q}' \subseteq \mathcal{Q}$ such that the family $\mathcal{Q}'$ admits γ-power saving, and for each $Q \in \mathcal{Q} \setminus \mathcal{Q}'$ the relation Q is in a $\mathfrak{p}$-generic correspondence with an abelian group G_Q type-definable in $\mathcal{M}^{\mathrm{eq}}$ over a set of parameters of cardinality $\leq |\mathcal{L}|$.

To see that Theorem 5.24 follows from Theorem 5.27, assume that a definable relation Q is as in Theorem 5.24, and consider the definable family $\mathcal{Q} := \{Q\}$ consisting of a single element Q. By Proposition 2.14 and Fact 2.4 every definable family of binary relations in $\mathcal{M}$ satisfies the γ-ST property (Definition 2.12) for some $\gamma > 0$. Moreover, by Corollary 5.13, if $Q \subseteq X_1 \times \ldots \times X_s$ is definable and fiber-algebraic of degree $\leq d$, we have $Q = \bigcup_{i \in [sd]} Q_i$ for some definable absolutely $\mathfrak{p}$-irreducible sets Q_i. By distality, each Q_i satisfies the γ_i-ST-property for some $\gamma_i > 0$. Hence, taking $\mathcal{Q}_i := \{Q_i\}$, $m := sd$ and $\gamma := \min\{\gamma_i : i \in [m]\} > 0$, the assumption of Theorem 5.27 is satisfied for $s \geq 4$. If $s = 3$, note that each Q_i is still fiber-algebraic of degree d, hence each $Q_i' \subseteq X_1 \times \ldots \times X_4$ is fiber-algebraic, of degree $\leq d^2$ by Lemma 5.44. By Corollary 5.13 again, for each i we have $Q_i' = \bigcup_{j \in [4d^2]} Q_{i,j}$ for some definable absolutely $\mathfrak{p}$-irreducible sets $Q_{i,j}$, each satisfying the $\gamma_{i,j}$-ST-property for some $\gamma_{i,j} > 0$. Hence, taking $m_i := 4d^2$, $\mathcal{Q}_{i,j} := \{Q_{i,j}\}$ and $\gamma := \min\{\gamma_{i,j} : i \in [m], j \in [m_i]\} > 0$, the assumption of Theorem 5.27 is satisfied for $s = 3$. In either case, let $\mathcal{Q}'$ be as given by applying Theorem 5.27. If $\mathcal{Q}' = \mathcal{Q}$, then Q is in Case (1) of Theorem 5.24. Otherwise $\mathcal{Q}' = \emptyset$, and Q is in Case (2) of Theorem 5.24.

In the rest of the section we give a proof of Theorem 5.27 (which will also establish Theorem 5.24). In fact, first we will prove a special case of Theorem 5.27 for definable families of absolutely $\mathfrak{p}$-irreducible sets and $s \geq 4$ (Theorem 5.31), and then derive full Theorem 5.27 from it in Section 5.6 (for $s \geq 4$) and Section 5.7 (for $s = 3$). We begin with some auxiliary observations and reductions.

Assumption 1. *For the rest of Section 5, we assume that $s \in \mathbb{N}_{\geq 3}$ (even though some of the results below make sense for $s \in \mathbb{N}_{\geq 1}$), $\mathcal{M}$ is $|\mathcal{L}|^+$-saturated, $(X_i, \mathfrak{p}_i)_{i\in[s]}$ is a $\mathfrak{p}$-system with each $\mathfrak{p}_i$ non-algebraic, and X_i is a $\emptyset$-definable. "Definable" will mean "definable with parameters in $\mathcal{M}$"*

Lemma 5.28. *If $Q \subseteq X_1 \times \cdots \times X_s$ is fiber-algebraic then $\dim_{\mathfrak{p}}(Q) \leq s-1$.*

Proof. Let $(a_1, \ldots, a_{s-1}) \models \bigotimes_{i\in[s-1]} \mathfrak{p}_i|_A$, where A is some finite set such that Q is A-definable. The type $\mathfrak{p}_s$ is non-algebraic by Assumption 1, and $Q(a_1, \ldots, a_{s-1}, x_s)$ has at most d solutions. Hence necessarily

$$Q(a_1, \ldots, a_{s-1}, x_s) \notin \mathfrak{p}_s,$$

so $Q(x_1, \ldots, x_s) \notin \bigotimes_{i\in[s]} \mathfrak{p}_i$. □

The following is straightforward by definition of fiber-algebraicity.

Lemma 5.29. *Let $Q \subseteq X_1 \times \cdots \times X_s$ be a fiber-algebraic relation of degree $\leq d$ and $u \subseteq [s]$ with $|u| = s-1$. Let π_u be the projection from $X_1 \times \cdots \times X_s$ onto $\prod_{i\in u} X_i$. Let $A = A_1 \times \cdots \times A_s$ be a grid on $X_1 \times \cdots \times X_s$. Then*

$$|Q \cap A| \leq d \left| \pi_u(Q) \cap \prod_{i\in u} A_i \right|.$$

Proposition 5.30. *Let $\mathcal{Q}$ be a definable family of fiber-algebraic subsets of $X_1 \times \cdots \times X_s$. Let $u \subseteq [s]$ with $u = s-1$. Assume that for every $Q \in \mathcal{Q}$ the projection $\pi_u(Q)$ onto $\prod_{i\in u} X_i$ is not $\mathfrak{p}$-generic. Then $\mathcal{Q}$ admits 1-power saving.*

Proof. By Lemma 5.19 there exists a definable system $\vec{\mathcal{F}}^*_u = (\mathcal{F}_i : i \in u)$ of subsets of $\prod_{i\in u} X_i$ such that for any $\nu \in \mathbb{N}$, for any n-grid A^* on $\prod_{i\in u} X_i$ in $(\vec{\mathcal{F}}^*_u, \nu)$-general position, for any $Q \in \mathcal{Q}$ we have $|\pi_u(Q) \cap A^*| \leq s^{s-2}\nu^2 n^{s-2}$. Let $d \in \mathbb{N}$ be such that $\mathcal{Q}$ is of degree $\leq d$. Taking $\mathcal{F}_i := \emptyset$ for $i \in [s] \setminus u$, let $\vec{\mathcal{F}}_u := \{\mathcal{F}_i : i \in [s]\}$. Then by Lemma 5.29, for any n-grid A on $\prod_{i\in[s]} X_i$ in $(\vec{\mathcal{F}}, \nu)$-general position, for any $Q \in \mathcal{Q}$ we have $|Q \cap A| \leq ds^{s-2}\nu^2 n^{s-2} = O_\nu(n^{s-2})$, hence the family $\mathcal{Q}$ admits 1-power saving. □

The following is the main theorem for definable families of absolutely irreducible sets:

Theorem 5.31. *Assume that $\mathcal{M}$ is an $|\mathcal{L}|^+$-saturated $\mathcal{L}$-structure and $\mathrm{Th}(\mathcal{M})$ is stable. Assume that $s \geq 4$, $(X_i, \mathfrak{p}_i)_{i\in[s]}$ is a $\mathfrak{p}$-system with each $\mathfrak{p}_i$ non-algebraic, and let $\mathcal{Q}$ be a fiber-algebraic definable family of absolutely $\mathfrak{p}$-irreducible subsets of $X_1 \times \cdots \times X_s$. Assume that for some $0 < \gamma \leq 1$, for each $i \neq j \in [s]$, $\mathcal{Q}$ viewed as a definable family of subsets of $(X_i \times X_j) \times \left(\prod_{k\in[s]\setminus\{i,j\}} X_k\right)$ satisfies the γ-ST property. Then there is a definable subfamily $\mathcal{Q}' \subseteq \mathcal{Q}$ such that the family $\mathcal{Q}'$ admits γ-power saving, and for each $Q \in \mathcal{Q} \setminus \mathcal{Q}'$ the relation Q is in a $\mathfrak{p}$-generic correspondence with an abelian group G_Q type-definable in $\mathcal{M}^{\mathrm{eq}}$ over a set of parameters of cardinality $\leq |\mathcal{L}|$.*

In the rest of this section we give a proof of Theorem 5.31 (and then of Theorem 5.27).

We fix a fiber-algebraic definable family $\mathcal{Q}$ of absolutely $\mathfrak{p}$-irreducible subsets of $X_1 \times \cdots \times X_s$.

Let $\mathcal{Q}_0$ be the set of all $Q \in \mathcal{Q}$ such that for some $u \subseteq [s]$ with $|u| = s-1$ for the projection $\pi_u(Q)$ of Q onto $\prod_{i \in u} X_i$ we have $\dim_{\mathfrak{p}}(\pi_u(Q) < s-1$.

By Claim 5.8, the family $\mathcal{Q}_0$ is definable and it follows from Proposition 5.30 that the family $\mathcal{Q}_0$ admits 1-power saving. Hence replacing $\mathcal{Q}$ with $\mathcal{Q} \setminus \mathcal{Q}_0$, if needed, we will assume the following:

Assumption 2. *$\mathcal{Q}$ is a fiber-algebraic definable family of absolutely $\mathfrak{p}$-irreducible subsets of $X_1 \times \cdots \times X_s$. For any $Q \in \mathcal{Q}$ the projection of Q onto any $s-1$ coordinates is $\mathfrak{p}$-generic. In particular, $\dim_{\mathfrak{p}}(Q) = s-1$ (by Lemma 5.28).*

Proposition 5.32. *Let C be a small set in $\mathcal{M}$, $Q \in \mathcal{Q}$ and let $\bar{a} = (a_1, \ldots, a_s)$ be a $\mathfrak{p}$-generic in Q over C (see Remark 5.6 for the definition). Then for any $i \in [s]$ we have*

$$(a_j : j \in [s] \setminus \{i\}) \models \bigotimes_{j \in [s] \setminus \{i\}} \mathfrak{p}_j|_C.$$

Proof. Since Q is absolutely $\mathfrak{p}$-irreducible, it has unique $\mathfrak{p}$-generic type over C. By our assumption for any $i \in [s]$ the projection of Q onto $[s] \setminus \{i\}$ is $\mathfrak{p}$-generic. Hence any realization of $\bigotimes_{j \in [s] \setminus \{i\}} \mathfrak{p}_j|_C$ can be extended to a $\mathfrak{p}$-generic of Q. ☐

Next we observe that the assumption that the projection of Q onto any $s-1$ coordinates is $\mathfrak{p}$-generic in Proposition 5.32 was necessary, but could be replaced by the assumption that Q does not admit 1-power saving (this will not be used in the proof of the main theorem).

Proposition 5.33. *Assume that Q is absolutely $\mathfrak{p}$-irreducible, $\dim_{\mathfrak{p}}(Q) = s-1$ (but no assumption on the projections of Q), and Q does not admit 1-power saving. Let C be a small set in $\mathcal{M}$ and let $\bar{a} = (a_1, \ldots, a_s)$ be a generic in Q over C. Then for any $i \in [s]$ we have*

$$(a_j : j \in [s] \setminus \{i\}) \models \bigotimes_{j \in [s] \setminus \{i\}} \mathfrak{p}_j|_C.$$

Proof. Let $\bar{a}$ be a generic in Q over C. Permuting the variables if necessary and using that the types $\mathfrak{p}_i$ commute, we may assume

$$(a_1, \ldots, a_{s-1}) \models \mathfrak{p}_1 \otimes \cdots \otimes \mathfrak{p}_{s-1}|_C.$$

We only consider the case $i = 1$, i.e. we need to show that

$$(a_2, \ldots, a_s) \models \mathfrak{p}_2 \otimes \cdots \otimes \mathfrak{p}_s|_C,$$

the other cases are analogous.

Assume this does not hold, then there is a relation $G_1 \subseteq X_2 \times \cdots \times X_s$ definable over C such that $\dim_{\mathfrak{p}}(G_1) < s-1$ and $(a_2, \ldots, a_s) \in G_1$.

Since Q is $\mathfrak{p}$-irreducible over C, the formula $Q(a_1, \ldots, a_{s-1}, x_s)$ implies a complete type over $C \cup \{a_1, \ldots, a_{s-1}\}$ by Lemma 5.14. Hence we have

$$Q(a_1, \ldots, a_{s-1}, x_s) \vdash \mathrm{tp}(a_s/C \cup \{a_1, \ldots, a_{s-1}\}),$$

so in particular

$$Q(a_1, \ldots, a_{s-1}, x_s) \to G_1(a_2, \ldots, a_{s-1}, x_s),$$

which implies

$$\{Q(x_1, \ldots, x_{s-1}, x_s)\} \cup (\mathfrak{p}_1 \otimes \ldots \otimes \mathfrak{p}_{s-1})\,|_C(x_1, \ldots, x_{s-1}) \\ \to G_1(x_2, \ldots, x_{s-1}, x_s).$$

Then, by saturation of $\mathcal{M}$, there exists some $\mathfrak{p}$-generic set $G_2 \subseteq X_1 \times \cdots \times X_{s-1}$ definable over C (given by a finite conjunction of formulas from $(\mathfrak{p}_1 \otimes \ldots \otimes \mathfrak{p}_{s-1})\,|_C$) such that

$$Q(x_1, \ldots, x_{s-1}, x_s) \wedge G_2(x_1, \ldots, x_{s-1}) \to G_1(x_2, \ldots, x_{s-1}, x_s),$$

hence

$$Q(x_1, \ldots, x_{s-1}, x_s) \to \big(\neg G_2(x_1, \ldots, x_{s-1}) \vee G_1(x_2, \ldots, x_{s-1}, x_s)\big).$$

Let $H_2 := (\neg G_2) \times X_s$ and $H_1 := X_1 \times G_1$. Then $\dim_{\mathfrak{p}}\left(\pi_{[s-1]}(H_2)\right) = \dim_{\mathfrak{p}}(\neg G_2) < s-1$ and $\dim_{\mathfrak{p}}\left(\pi_{[s]\setminus\{1\}}(H_1)\right) = \dim_{\mathfrak{p}}(\neg G_1) < s-1$. Thus Q is covered by the union of H_1 and H_2, each with 1-power saving by Proposition 5.30, which implies that Q admits 1-power-saving. ☐

Remark 5.34. The assumption that Q has no 1-power saving is necessary in Proposition 5.33, and the assumption that the projection of Q onto any $s-1$ coordinates is $\mathfrak{p}$-generic in necessary in Proposition 5.32. For example let $s = 2$ and assume $Q(x_1, x_2)$ is the graph of a bijection from X_1 to some $\emptyset$-definable set $Y_2 \subseteq X_2$ with $Y_2 \notin \mathfrak{p}_2$. Then Q is clearly fiber algebraic, absolutely $\mathfrak{p}$-irreducible, with $\dim_{\mathfrak{p}}(Q) = 1$. But for a generic $(b_1, b_2) \in Q$, b_2 does not realize $\mathfrak{p}_2|_{\emptyset}$. Note that Q has 1-power saving. Indeed, let $\vec{\mathcal{F}} := (\mathcal{F}_1, \mathcal{F}_2)$ with $\mathcal{F}_1 := \emptyset, \mathcal{F}_2 := \{Y_2\}$. Then, given any $n, \nu \in \mathbb{N}$ and an n-grid $A_1 \times A_2$ in $\left(\vec{\mathcal{F}}, \nu\right)$-general position, as $\dim_{\mathfrak{p}}(Y_2) = 0$ we must have $|A_2 \cap Y_2| \leq \nu$, hence, by definition of Q, $|Q \cap (A_1 \times A_2)| \leq \nu = O_\nu(1) = O_\nu\left(n^{(2-1)-1}\right)$. Also note that $\pi_{\{2\}}(Q)$ is not $\mathfrak{p}$-generic.

We can now state the key structural dichotomy at the core of Theorem 5.31:

Theorem 5.35. *Let $\mathcal{Q} = \{Q_\alpha \colon \alpha \in \Omega\}$ be a definable family of absolutely $\mathfrak{p}$-irreducible fiber-algebraic subsets of $\prod_{i\in s} X_i$ satisfying the Assumption 1 above. Assume the family $\mathcal{Q}$, as a family of binary relations on*

$$\left(\prod_{i\in[s-2]} X_i\right) \times (X_{s-1} \times X_s),$$

satisfies the γ-ST property for some $0 < \gamma \leq 1$.

Then there is a definable $\Omega_1 \subseteq \Omega$ such that the family $\{Q_\alpha \colon \alpha \in \Omega_1\}$, admits γ-power-saving, and for every $\alpha \in \Omega \setminus \Omega_1$, for every tuple $(a_1, \ldots, a_s) \in Q_\alpha$ generic over α there exists some tuple

$$\xi \in \operatorname{acl}(a_1, \ldots, a_{s-2}, \alpha) \cap \operatorname{acl}(a_{s-1}, a_s, \alpha)$$

of length at most $|\mathcal{L}|$ such that

$$(a_1, \ldots, a_{s-2}) \underset{\xi}{\perp} (a_{s-1}, a_s).$$

Remark 5.36. Theorem 5.35 is trivial for $s = 3$ with $\Omega_1 = \emptyset$, as $a_1 \perp_\xi (a_2, a_3)$ always holds with $\xi := a_1\alpha$.

First we show how the above theorem, combined with the reconstruction of abelian groups from abelian s-gons in Theorem 4.6, implies Theorem 5.31. Then we use theorem Theorem 5.31 to deduce Theorem 5.24 for $s \geq 4$ (along with the bound in Theorem 5.27) in Section 5.6. The case $s = 3$ of Theorem 5.24 requires a separate argument reducing to the case $s = 4$ of Theorem 5.24 given in Section 5.7.

Proof of Theorem 5.31. From the reductions described above, we assume that $\mathcal{Q}$ and $(X_i, \mathfrak{p}_i)_{i \in [s]}$ satisfy Assumptions 1 and 2, and that for some $0 < \gamma \leq 1$, for each $i \neq j \in [s]$, $\mathcal{Q}$ viewed as a definable family of subsets of $(X_i \times X_j) \times \left(\prod_{k \in [s] \setminus \{i,j\}} X_k\right)$ satisfies the γ-ST property.

It follows that for every permutation of $[s]$, the family $\mathcal{Q}$ and the $\mathfrak{p}$-system obtained from $\mathcal{Q}$ and $(X_i, \mathfrak{p}_i)_{i \in [s]}$ by permuting the variables accordingly still satisfy the assumption of Theorem 5.35. Applying Theorem 5.35 to every permutation of $[s]$, and taking (definable) $\Omega' \subseteq \Omega$ to be the union of the corresponding Ω_1's, we have that the family $\mathcal{Q}' = \{Q_\alpha \colon \alpha \in \Omega'\}$ admits γ-power saving and for any $\alpha \in \Omega \setminus \Omega'$, for every tuple $(a_1, \ldots, a_s)$ generic in Q_α over α, after any permutation of $[s]$ we have

$$a_1 a_2 \underset{\operatorname{acl}(a_1 a_2 \alpha) \cap \operatorname{acl}(a_3 \ldots a_s \alpha)}{\perp} a_3 \ldots a_s.$$

Together with fiber-algebraicity of Q_α this implies that $(a_1, \ldots, a_s)$ is an abelian s-gon over α.

Applying Theorem 4.6, we obtain that for any $\alpha \in \Omega \setminus \Omega'$ there exists a small set $A_\alpha \subseteq \mathcal{M}$ and a connected abelian group G_α type-definable over A_α and such that Q_α is in a $\mathfrak{p}$-generic correspondence with G_α over A_α. (As stated, Theorem 4.6 only guarantees the existence of an appropriate set of parameters A_α of size $\leq |\mathcal{L}|$ and G_α *in* $\mathbb{M}$, however by $|\mathcal{L}|^+$-saturation of $\mathcal{M}$ there exists a set A'_α in $\mathcal{M}$ with the same type as A_α, hence we obtain the required group applying an automorphism of $\mathbb{M}$ sending A_α to A'_α.) □

In the remainder of the section we prove Theorem 5.35.

5.5. **Proof of Theorem 5.35.** Theorem 5.35 is trivial in the case $s = 3$ by Remark 5.36, so we will assume $s \geq 4$.

Let $U := X_1 \times \ldots \times X_{s-2}$ and $V := X_{s-1} \times X_s$. We view each $Q \in \mathcal{Q}$ as a binary relation $Q \subseteq U \times V$.

We fix a formula $\varphi(u; v; w) \in \mathcal{L}$ such that for $\alpha \in \Omega$ the formula $\varphi(u; v; \alpha)$ defines Q_α, with the variables u corresponding to U and v to V.

We also fix $d \in \mathbb{N}$ such that $\mathcal{Q}$ is of degree $\leq d$.

Definition 5.37. For $\alpha \in \Omega$ and $a \in U$, let $Z_\alpha(a)$ be the set

$$Z_\alpha(a) := \left\{a' \in U \colon \dim_{\mathfrak{p}} \left(\varphi(a; v; \alpha) \cap \varphi(a'; v; \alpha)\right) = 1\right\}.$$

Claim 5.38. *The family $\{Z_\alpha(a) \colon \alpha \in \Omega, a \in U\}$ is a definable family of subsets of U.*

Proof. By Claim 5.8, the set

$$\begin{aligned} D &:= \{(a, a', \alpha) \in U \times U \times \Omega : a' \in Z_\alpha(a)\} \\ &= \left\{(a, a', \alpha) \in U \times U \times \Omega \colon \dim_{\mathfrak{p}} \left(\varphi(a; v; \alpha) \cap \varphi(a'; v; \alpha)\right) = 1\right\} \end{aligned}$$

is definable, hence the family $\{Z_\alpha(a) \colon \alpha \in \Omega, a \in U\}$ is definable. □

Claim 5.39. *For any $\alpha \in \Omega$ and $a \in U$, we have that $Z_\alpha(a) \neq \emptyset$ if and only if $a \in Z_\alpha(a)$, if and only if $\dim_{\mathfrak{p}}(\varphi(a; v; \alpha)) = 1$.*

Proof. Let $\alpha \in \Omega$ and $a \in U$. As Q_α is fiber-algebraic, we also have that the binary relation $\varphi(a; v; \alpha) \subseteq X_{s-1} \times X_s$ is fiber-algebraic, hence $\dim_{\mathfrak{p}}(\varphi(a; v; \alpha)) \leq 1$ (by Lemma 5.28). The claim follows as, by definition of $\mathfrak{p}$-dimension, $\dim_{\mathfrak{p}}(\varphi(a; v; \alpha) \cap \varphi(a; v; \alpha)) = \dim_{\mathfrak{p}}(\varphi(a; v; \alpha)) \geq \dim_{\mathfrak{p}}(\varphi(a; v; \alpha) \cap \varphi(a'; v; \alpha))$ for any $a' \in U$. □

Claim 5.40. *For every $\alpha \in \Omega$ and $a \in U$ the set $Z_\alpha(a) \subseteq X_1 \times \cdots \times X_{s-2}$ is fiber-algebraic, of degree $\leq 2d^2$.*

Proof. We fix $\alpha \in \Omega$ and $a \in U$. Assume $Z_\alpha(a) \neq \emptyset$. Since $\varphi(a; v; \alpha)$ is fiber-algebraic of degree $\leq d$ (by fiber-algebraicity of Q_α), the set S of types $q \in S_v(\mathcal{M})$ with $\varphi(a; v; \alpha) \in q$ and $\dim_{\mathfrak{p}}(q) = 1$ is finite, of size $\leq 2d$ (by Lemma 5.12); and for any $a' \in U$ we have $a' \in Z_\alpha(a)$ if and only if $\varphi(a', v; \alpha)$ belongs to one of these types (by definition of $\mathfrak{p}$-dimension). Thus

$$Z_\alpha(a) = \{a' \in U \colon \varphi(a', v, \alpha) \in q \text{ for some } q \in S\}.$$

Let $q_1, \ldots, q_t, t \leq 2d$ list all types in S. We then have $Z_\alpha(a) = \bigcup_{i \in [t]} d_\varphi(q_i)$, where $d_\varphi(q_i) = \{a' \in U \colon \varphi(a', v, \alpha) \in q_i\}$. It is sufficient to show that each $d_\varphi(q_i)$ is fiber-algebraic of degree $\leq d$. Choose a realization β_i of q_i in $\mathbb{M}$. Obviously $d_\varphi(q_i) \subseteq \varphi(\mathbb{M}, \beta_i; \alpha)$. As $\mathcal{M} \preceq \mathbb{M}$ and $\alpha \in \mathcal{M}$, the set $\varphi(\mathbb{M}; \beta_i; \alpha) \subseteq \prod_{i \in [s-2]} X_i(\mathbb{M})$ is fiber-algebraic of degree $\leq d$, hence the set $d_\varphi(q_i)$ is fiber-algebraic of degree $\leq d$ as well. □

By Claim 5.40 and Lemma 5.28, each $Z_\alpha(a)$ is not a $\mathfrak{p}$-generic subset of $X_1 \times \cdots \times X_{s-2}$, hence we have that $\dim_{\mathfrak{p}}(Z_\alpha(a)) \leq s - 3$ for any $\alpha \in \Omega$ and $a \in U$.

Definition 5.41. Let $Z_\alpha \subseteq U$ be the set

$$Z_\alpha := \{a \in U \colon \dim_{\mathfrak{p}}(Z_\alpha(a)) = s-3\}.$$

Note that the family $\{Z_\alpha : \alpha \in \Omega\}$ is definable by Claim 5.8.

Let $\Omega_1 := \{\alpha \in \Omega \colon \dim_{\mathfrak{p}}(Z_\alpha) < s-2\}$. By Claim 5.8 the set Ω_1 is definable. We will show that the family $\mathcal{Q}_1 := \{Q_\alpha \colon \alpha \in \Omega_1\}$ admits γ-power saving for the required γ.

To show that the family $\{Q_\alpha \colon \alpha \in \Omega_1\}$ admits γ-power saving, it suffices to show that both families $\{Q_\alpha \cap (Z_\alpha \times V) \colon \alpha \in \Omega_1\}$ and $\{Q_\alpha \cap (\bar{Z}_\alpha \times V) \colon \alpha \in \Omega_1\}$ admit γ-power saving, where $\bar{Z}_\alpha := U \setminus Z_\alpha$ is the complement of Z_α in U.

Since for any $\alpha \in \Omega_1$ the set Z_α is not a $\mathfrak{p}$-generic subset of $X_1 \times \cdots \times X_{s-2}$, for the projection $\pi_{[s-1]} \colon X_1 \times \cdots \times X_s \to X_1 \times \cdots \times X_{s-1}$ we have that $\pi_{[s-1]}(Q_\alpha \cap (Z_\alpha \times V))$ is not a $\mathfrak{p}$-generic subset of $X_1 \times \ldots \times X_{s-1}$. Hence, by Proposition 5.30, the family $\{Q_\alpha \cap (Z_\alpha \times V) \colon \alpha \in \Omega_1\}$ admits 1-power saving.

Next we show that the family $\{Q_\alpha \cap (\bar{Z}_\alpha \times V) \colon \alpha \in \Omega_1\}$ admits γ-power saving. By the definition of Z_α, for any $\alpha \in \Omega_1$ and $a \in \bar{Z}_\alpha$ we have $\dim_{\mathfrak{p}}(Z_\alpha(a)) \leq s-4$. By Lemma 5.19, there is a definable system of sets $\vec{\mathcal{F}}_1 = (\mathcal{F}_1, \ldots, \mathcal{F}_{s-2})$ on $X_1 \times \ldots \times X_{s-2}$ such that for any n-grid $A_1 \times \cdots \times A_{s-2}$ in $(\vec{\mathcal{F}}_1, \nu)$-general position we have

$$|Z_\alpha(a) \cap (A_1 \times \cdots \times A_{s-2})| \leq (s-2)^{s-4} \nu^2 n^{s-4},$$

for any $\alpha \in \Omega_1$ and $a \in \bar{Z}_\alpha$.

Applying Lemma 5.18 to the definable family

$$\mathcal{G} := \big\{\varphi(a_1; v; \alpha) \cap \varphi(a_2; v; \alpha) \colon \alpha \in \Omega_1, a_1, a_2 \in U, \\ \dim_{\mathfrak{p}}(\varphi(a_1; v; \alpha) \cap \varphi(a_2; v; \alpha)) = 0\big\},$$

we obtain that there is a definable system of sets $\vec{\mathcal{F}}_2 = (\mathcal{F}_{s-1}, \mathcal{F}_s)$ on $X_{s-1} \times X_s$ such that for any n-grid $A_{s-1} \times A_s$ in $(\vec{\mathcal{F}}_2, \nu)$-general position and any $G \in \mathcal{G}$ we have

$$|G \cap (A_{s-1} \times A_s)| \leq \nu^2.$$

Then $\vec{\mathcal{F}} := \vec{\mathcal{F}}_1 \cup \vec{\mathcal{F}}_2 = (\mathcal{F}_1, \ldots \mathcal{F}_s)$ is a definable system of sets on $X_1 \times \cdots \times X_s$.

Let $A = A_1 \times \cdots \times A_s$ be an n-grid on $X_1 \times \cdots \times X_s$ in $(\vec{\mathcal{F}}, \nu)$-general position and $\alpha \in \Omega_1$. We need to estimate from above $|Q_\alpha \cap (\bar{Z}_\alpha \times V) \cap A|$. Let $A_u := A_1 \times \cdots \times A_{s-2}, A'_u := A_u \cap \bar{Z}_\alpha$ and $A_v := A_{s-1} \times A_s$, then $|A'_u| \leq |A_u| \leq n^{s-2}$ and $|A_v| \leq n^2$. Let Q'_α be Q_α viewed as a binary relation on $U \times V$, we have

$$|Q_\alpha \cap (\bar{Z}_\alpha \times V) \cap A| = |Q'_\alpha \cap (\bar{Z}_\alpha \times V) \cap (A_u \times A_v)| \leq |Q'_\alpha \cap (A'_u \times A_v)|,$$

so it suffices to obtain the desired upper bound on $|Q'_\alpha \cap (A'_u \times A_v)|$.

From the $(\vec{\mathcal{F}}, \nu)$-general position assumption and the choice of $\vec{\mathcal{F}}$ we have: for any $a \in A'_u$ there are at most $(s-2)^{s-4}\nu^2 n^{s-4}$ elements $a' \in A'_u$ such that $|Q'_\alpha(a, v) \cap Q'_\alpha(a', v) \cap A_v| \geq \nu^2$.

By assumption on $\mathcal{Q}$ the definable family $\mathcal{Q}'_1 := \{Q'_\alpha : \alpha \in \Omega_1\}$ of subsets of $U \times V$ satisfies the γ-ST property, and let $C' : \mathbb{N} \to \mathbb{N}$ be as given by Definition 2.12 for $C(\nu) := (s-2)^{s-4}\nu$. Then we have $|Q'_\alpha \cap (A'_u \times A_v)| \leq C'(\nu^2)n^{(s-1)-\gamma}$ (independently of α), as required.

Thus the family $\mathcal{Q}_1 = \{Q_\alpha \colon \alpha \in \Omega_1\}$ admits γ-power saving.

We now fix $\alpha \in \Omega \setminus \Omega_1$.

By absolute irreducibility of Q_α and Remark 5.11 it is sufficient to show that there exists a tuple $(a_1, \dots a_s) \in Q_\alpha$ generic over α and some tuple $\xi \in \operatorname{acl}(a_1, \dots, a_{s-2}, \alpha) \cap \operatorname{acl}(a_{s-1}, a_s, \alpha)$ of length at most $|\mathcal{L}|$ such that

$$(a_1, \dots, a_{s-2}) \mathop{\smile\!\!\!\!\!\!|}_{\xi} (a_{s-1}, a_s).$$

We add $\operatorname{acl}(\alpha)$ *to our language if needed and assume that* $\alpha \in \operatorname{dcl}(\emptyset)$.

By $|\mathcal{L}|^+$-saturation of $\mathcal{M}$, let $e = (e_1, \dots, e_{s-2})$ be a tuple in $\mathcal{M}$ which is $\mathfrak{p}$-generic in Z_α, namely $e \in Z_\alpha$ with $\dim_{\mathfrak{p}}(e/\emptyset) = s-2$ (note that Z_α is $\emptyset$-definable). Let $\mathcal{M}_0 = (M_0, \dots) \preceq \mathcal{M}$ be a model containing e with $|\mathcal{M}_0| \leq |\mathcal{L}|$.

Let $\beta = (\beta_1, \dots, \beta_{s-2}) \in U$ be a $\mathfrak{p}$-generic point in $Z_\alpha(e)$ over M_0, i.e. $\beta \in Z_\alpha(e)$ and $\dim_{\mathfrak{p}}(\beta/M_0) = s-3$.

Let $\delta = (\delta_1, \delta_2)$ be a tuple in $\varphi(e, \mathcal{M}, \alpha) \cap \varphi(\beta, \mathcal{M}, \alpha)$ with $\dim_{\mathfrak{p}}(\delta/M_0\beta) = 1$. Without loss of generality we assume that $\dim_{\mathfrak{p}}(\delta_1/M_0\beta) = 1$, namely $\delta_1 \models \mathfrak{p}_{s-1} \restriction_{M_0\beta}$. Note that such β and δ can be chosen in $\mathcal{M}$ by $|\mathcal{L}|^+$-saturation.

We now collect some properties of β and δ.

Claim 5.42. *(1)* (e, δ) *is generic in* Q_α *over* $\emptyset$.
(2) $\delta_1 \mathop{\smile\!\!\!\!\!\!|}_{M_0} \delta_2$ *and* $(\delta_1, \delta_2) \models \mathfrak{p}_{s-1} \otimes \mathfrak{p}_s|_\emptyset$.
(3) $\beta \mathop{\smile\!\!\!\!\!\!|}_{M_0} \delta$.
(4) (β, δ) *is generic in* Q_α *over* $\emptyset$.

Proof. (1) We have, by our assumption above, that $\dim_{\mathfrak{p}}(\delta_1/M_0\beta) = 1$, hence in particular $\dim_{\mathfrak{p}}(\delta_1/e) = 1$. Since $\dim_{\mathfrak{p}}(e/\emptyset) = s-2$ we have $\dim_{\mathfrak{p}}((e, \delta)/\emptyset) \geq s-1$ (as $(e, \delta_1) \models \left(\bigotimes_{i \in [s-2]} \mathfrak{p}_i\right) \otimes \mathfrak{p}_{s-1}|_\emptyset$ using that the types $\mathfrak{p}_i, i \in [s-1]$ commute). Since Q_α is fiber-algebraic and $(e, \delta) \in Q_\alpha$, we also have $\dim_{\mathfrak{p}}((e, \delta)/\emptyset) \leq s-1$ by Lemma 5.28.

(2) Since (e, δ) is generic in Q_α over $\emptyset$ by (1), by Proposition 5.32 we have $(\delta_1, \delta_2) \models \mathfrak{p}_{s-1} \otimes \mathfrak{p}_s|_\emptyset$.

(3) As $\beta \mathop{\smile\!\!\!\!\!\!|}_{M_0} \delta_1$ and $\delta_2 \in \operatorname{acl}(e\delta_1) \subseteq \operatorname{acl}(M_0\delta_1)$, we have $\beta \mathop{\smile\!\!\!\!\!\!|}_{M_0} (\delta_1, \delta_2)$.

(4) We have $(\beta, \delta) \in Q_\alpha$. Since $\dim_{\mathfrak{p}}(\beta/M_0) = s-3$ and $\beta \mathop{\smile\!\!\!\!\!\!|}_{M_0} \delta$, we have $\dim_{\mathfrak{p}}(\beta/M_0\delta) = s-3$ (as $\beta \models \bigotimes_{i \in [s-3]} \mathfrak{p}_i|_{M_0\delta}$ by stationarity of non-forking

over models), hence in particular $\dim_{\mathfrak{p}}(\beta/\delta) \geq s-3$. Also, since $\dim_{\mathfrak{p}}(\delta/\emptyset) = 2$ by (2), we have $\dim_{\mathfrak{p}}((\beta,\delta)/\emptyset) \geq s-1$. Since Q_α is fiber-algebraic we also have $\dim_{\mathfrak{p}}((\beta,\delta)/\emptyset) \leq s-1$, hence $\dim_{\mathfrak{p}}((\beta,\delta)/\emptyset) = s-1$. □

Let $p(u) := \mathrm{tp}(\beta/M_0)$ and $q(v) := \mathrm{tp}(\delta/M_0)$, both are definable types over M_0 by stability.

We choose canonical bases ξ_p and ξ_q of p and q, respectively; i.e. ξ_p, ξ_q are tuples of length $\leq |\mathcal{L}|$ in $\mathcal{M}_0^{\mathrm{eq}}$, and for any automorphism σ of $\mathcal{M}$ we have $\sigma(p|\mathcal{M}) = p|\mathcal{M}$ if and only if $\sigma(\xi_p) = \xi_p$ (pointwise); and $\sigma(q|\mathcal{M}) = q|\mathcal{M}$ if and only if $\sigma(\xi_q) = \xi_q$.

Note that p does not fork over ξ_p and q does not fork over ξ_q.

Claim 5.43. *We have:*

(a) $\xi_q \in \mathrm{acl}(\beta)$*;*
(b) $\xi_p \in \mathrm{acl}(\delta)$*;*
(c) $\xi_q \in \mathrm{acl}(\xi_p)$*;*
(d) $\xi_p \in \mathrm{acl}(\xi_q)$*.*

Proof. (a) Assume not, then the orbit of ξ_q under the automorphisms of $\mathcal{M}$ fixing β would be infinite. Hence we can choose a model $\mathcal{N} = (N, \ldots) \preceq \mathcal{M}$ containing $M_0\beta$ with $|N| \leq |\mathcal{L}|$, and distinct types $q_i \in S_v(N), i \in \omega$, each conjugate to $q|N$ under an automorphism of $\mathcal{N}$ fixing β.

Let $\delta'_1 \models \mathfrak{p}_{s-1}|N$. For each $i \in \omega$ we choose δ_2^i such that $(\delta'_1, \delta_2^i) \models q_i$. We have that $(\beta, \delta'_1, \delta_2^i) \in Q_\alpha$, hence, by fiber-algebraicity, $|\{\delta_2^i \colon i \in \omega\}| \leq d$. But all q_i are pairwise distinct types, a contradiction.

(b) Since $\dim_{\mathfrak{p}}(\beta/M_0\delta) = s-3$, permuting variables if needed, we may assume that $(\beta_1, \ldots, \beta_{s-3}) \models \mathfrak{p}_1 \otimes \cdots \otimes \mathfrak{p}_{s-3}|_{M_0\delta}$.

Assume (b) fails. Then we can find a model $\mathcal{N} \preceq \mathcal{M}, |\mathcal{N}| \leq |\mathcal{L}|$ containing $M_0\delta$, and distinct types $p_i \in S(N), i \in \omega$, each conjugate to $p{\restriction}N$ under an automorphism of $\mathcal{N}$ fixing β. Let

$$(\beta'_1, \ldots, \beta'_{s-3}) \models \mathfrak{p}_1 \otimes \ldots \otimes \mathfrak{p}_{s-3}|N$$

in $\mathcal{M}$. For each $i \in \omega$ we choose β_{s-2}^i in $\mathcal{M}$ such that

$$(\beta'_1, \ldots, \beta'_{s-3}, \beta_{s-2}^i) \models p_i,$$

and get a contradiction as in (a).

(c) Since $\xi_q \in M_0$ and p does not fork over ξ_p, we have $\xi_q \mathop{\smile\!\!\!\!\!\!\mid}_{\xi_p} \beta$, which by part (a) implies $\xi_q \in \mathrm{acl}(\xi_p)$.

(d) Similar to (c). □

We have that the tuple (β, δ) is generic in Q_α by Claim 5.42(4). Let $\xi := \xi_p \cup \xi_q$, then $\xi \in \mathrm{acl}(\beta) \cap \mathrm{acl}(\delta)$ by Claim 5.43. Finally $\delta \mathop{\smile\!\!\!\!\!\!\mid}_{M_0} \beta$ by Claim 5.42(3), $\beta \mathop{\smile\!\!\!\!\!\!\mid}_{\xi_p} M_0$ by the choice of ξ_p, hence $\delta \mathop{\smile\!\!\!\!\!\!\mid}_{\xi_p} \beta$, and as $\xi_q \in \mathrm{acl}(\beta)$ we conclude $\beta \mathop{\smile\!\!\!\!\!\!\mid}_{\xi} \delta$.

This finishes the proof of Theorem 5.35, and hence of Theorem 5.31.

5.6. **Proof of Theorem 5.27 for $s \geq 4$.** Let $\mathcal{Q} = \{Q_\alpha : \alpha \in \Omega\}$ be a definable family of subsets of $X_1 \times \cdots \times X_s$ satisfying the assumption of Theorem 5.27, and say $\mathcal{Q}$ is fiber-algebraic of degree $\leq d$. In particular, there exist $m \in \mathbb{N}$ and definable families $\mathcal{Q}_i, i \in [m]$ of absolutely $\mathfrak{p}$-irreducible subsets of $X_1 \times \cdots \times X_s$, so that for every $Q \in \mathcal{Q}$ we have $Q = \bigcup_{i\in[m]} Q_i$ for some $Q_i \in \mathcal{Q}_i$. Note that each $\mathcal{Q}_i$ is automatically fiber-algebraic, of degree $\leq d$. By assumption each $\mathcal{Q}_i$ satisfies the γ-ST property for some fixed $\gamma > 0$, under any partition of the variables into two groups of size 2 and $s - 2$.

For each $i \in [m]$, let the definable family $\mathcal{Q}'_i$ be as given by Theorem 5.31 for $\mathcal{Q}_i$. That is, for each $i \in [m]$ the family $\mathcal{Q}'_i$ admits γ-power saving, and for each $Q_i \in \mathcal{Q}_i \setminus \mathcal{Q}'_i$ the relation Q_i is in a $\mathfrak{p}$-generic correspondence with an abelian group G_{Q_i} type-definable in $\mathcal{M}^{\mathrm{eq}}$ over a set of parameters A_i of cardinality $\leq |\mathcal{L}|$. Consider the definable family

$$\mathcal{Q}' := \left\{ Q \in \mathcal{Q} : Q = \bigcup_{i\in[m]} Q_i \text{ for some } Q_i \in \mathcal{Q}'_i \right\} \subseteq \mathcal{Q}.$$

By Lemma 5.1, $\mathcal{Q}'$ satisfies γ-power saving. On the other hand, from Definition 5.22, if $Q \in \mathcal{Q}$, $Q = \bigcup_{i\in[m]} Q_i$ with $Q_i \in \mathcal{Q}_i$, and at least one of the Q_i is in a $\mathfrak{p}$-generic correspondence with a type-definable group, then Q is also in a $\mathfrak{p}$-generic correspondence with the same group. Hence every element $Q \in \mathcal{Q} \setminus \mathcal{Q}'$ is in a $\mathfrak{p}$-generic correspondence with a group type-definable over some $A := \bigcup_{i\in[m]} A_i, |A| \leq |\mathcal{L}|$.

5.7. **Proof of Theorem 5.27 for ternary Q.** In this subsection we reduce the remaining case $s = 3$ of Theorem 5.27 to the case $s = 4$.

Let $(X_i, \mathfrak{p}_i)_{i\in[3]}$ and a definable fiber-algebraic (say, of degree $\leq d$) family $\mathcal{Q}$ of subsets of $X_1 \times X_2 \times X_3$ satisfy the assumption of Theorem 5.27 with some fixe $\gamma > 0$. In particular, there exist $m \in \mathbb{N}$ and fiber-algebraic (of degree $\leq d$) definable families $\mathcal{Q}_i, i \in [m]$ of absolutely $\mathfrak{p}$-irreducible subsets of $X_1 \times \cdots \times X_s$, so that for every $Q \in \mathcal{Q}$ we have $Q = \bigcup_{i\in[m]} Q_i$ for some $Q_i \in \mathcal{Q}_i$. By the same reduction as in Section 5.6, it suffices to establish the theorem separately for each $\mathcal{Q}_i$, so we may assume from now on that additionally all sets in $\mathcal{Q}$ are absolutely $\mathfrak{p}$-irreducible.

Consider the definable family $\mathcal{Q}^* := \{Q^* : Q \in \mathcal{Q}\}$ of subsets of $X_1 \times X_2 \times X_3 \times X_4$, where

$$Q^* := \Big\{ (x_2, x'_2, x_3, x'_3) \in X_2 \times X_2 \times X_3 \times X_3 :$$
$$\exists x_1 \in X_1 \left((x_1, x_2, x_3) \in Q \wedge (x_1, x'_2, x'_3) \in Q \right) \Big\}.$$

Lemma 5.44. *The definable family $\mathcal{Q}^*$ of subsets of $X_2 \times X_2 \times X_3 \times X_3$ is fiber algebraic, of degree $\leq d^2$.*

Proof. We consider the case of fixing the first three coordinates of $Q^* \in \mathcal{Q}^*$, all other cases are similar. Let $Q \in \mathcal{Q}$, $(a_2, a'_2) \in X_2 \times X_2$ and $a_3 \in X_3$ be fixed. As Q is fiber algebraic of degree $\leq d$, there are at most d elements $x_1 \in X_1$ such that $(x_1, a_2, a_3) \in Q$; and for each such x_1, there are at most d elements $x'_3 \in X_3$ such that $(x_1, a'_2, x'_3) \in Q$. Hence, by definition of Q^*, there are at most d^2 elements $x'_3 \in X_3$ such that $(a_2, a'_2, a_3, x'_3) \in Q^*$. □

Remark 5.45. Note that $(X'_i, \mathfrak{p}'_i)_{i\in[4]}$ with $X'_1 = X'_2 := X_2, X'_3 = X'_4 := X_3$ and $\mathfrak{p}'_1 = \mathfrak{p}'_2 := \mathfrak{p}_2, \mathfrak{p}'_3 = \mathfrak{p}'_4 := \mathfrak{p}_3$ is a $\mathfrak{p}$-system with each $\mathfrak{p}$ non-algebraic.

The following lemma will be used to show that power saving for $\mathcal{Q}^*$ implies power saving for $\mathcal{Q}$ (this is a version of [20, Proposition 3.10] for families, which in turn is essentially [50, Lemma 2.2]). We include a proof for completeness.

Lemma 5.46. *For any finite $A_i \subseteq X_i, i \in [3]$ and $Q \in \mathcal{Q}$, taking $\tilde{Q} := Q \cap (A_1 \times A_2 \times A_3)$ and $\tilde{Q}^* := Q^* \cap (A_2 \times A_2 \times A_3 \times A_3)$ we have*

$$\left|\tilde{Q}\right| \leq d\, |A_1|^{\frac{1}{2}} \left|\tilde{Q}^*\right|^{\frac{1}{2}}.$$

Proof. Consider the (definable) set

$$W := \big\{(x_1, x_2, x'_2, x_3, x'_3) \in X_1 \times X_2^2 \times X_3^2 : \\ (x_1, x_2, x_3) \in Q \wedge (x_1, x'_2, x'_3) \in Q\big\},$$

and let $\tilde{W} := W \cap \left(A_1 \times A_2^2 \times A_3^2\right)$. As usual, for arbitrary sets $S \subseteq B \times C$ and $b \in B$, we denote by S_b the fiber $S_b = \{c \in C : (b, c) \in S\}$.

Note that $|\tilde{Q}| = \sum_{a_1 \in A_1} |\tilde{Q}_{a_1}|$ and $|\tilde{W}| = \sum_{a_1\in A_1} |\tilde{Q}_{a_1}|^2$, which by the Cauchy-Schwarz inequality implies

$$|\tilde{Q}| \leq |A_1|^{\frac{1}{2}} \left(\sum_{a_1 \in A_1} |\tilde{Q}_{a_1}|^2\right)^{\frac{1}{2}} = |A_1|^{\frac{1}{2}} |\tilde{W}|^{\frac{1}{2}}.$$

For any tuple $\bar{a} := (a_2, a'_2, a_3, a'_3) \in \tilde{Q}^*$, the fiber $\tilde{W}_{\bar{a}} \subseteq A_1$ has size at most d by fiber algebraicity of Q. Hence $|\tilde{W}| \leq d|\tilde{Q}^*|$, and so $|\tilde{Q}| \leq d|A_1|^{\frac{1}{2}}|\tilde{Q}^*|^{\frac{1}{2}}$. □

Lemma 5.47. *Assume that $\gamma' > 0$ and $\mathcal{Q}^*$ admits γ'-power saving (with respect to the $\mathfrak{p}$-system $(X'_i, \mathfrak{p}'_i)_{i\in[4]}$ in Remark 5.45). Then $\mathcal{Q}$ admits γ-power saving for $\gamma := \frac{\gamma'}{2}$.*

Proof. By assumption there exist $\vec{\mathcal{F}}' = (\mathcal{F}'_i)_{i\in[4]}$ with $\mathcal{F}'_1, \mathcal{F}'_2$ definable families on X_2 and $\mathcal{F}'_3, \mathcal{F}'_4$ definable families on X_3, and a function $C' : \mathbb{N} \to \mathbb{N}$, such that for any $Q^* \in \mathcal{Q}^*, \nu, n \in \mathbb{N}$ and an n-grid $A' = \prod_{i\in[4]} A'_i$ on $X_2 \times X_2 \times X_3 \times X_3$ in $(\vec{\mathcal{F}}', \nu)$-general position we have $|Q^* \cap A'| \leq C'(\nu)n^{3-\gamma'}$.

We take $\mathcal{F}_1 := \emptyset$, $\mathcal{F}_2 := \mathcal{F}'_1 \cup \mathcal{F}'_2$, $\mathcal{F}_3 := \mathcal{F}'_3 \cup \mathcal{F}'_4$, $C(\nu) := d \cdot C'(\nu)^{\frac{1}{2}}$ and $\gamma := \frac{\gamma'}{2}$.

Assume we are given $Q \in \mathcal{Q}, \nu, n \in \mathbb{N}$ and $A_i \subseteq X_i, i \in [3]$ with $|A_i| = n$ in $(\vec{\mathcal{F}}, \nu)$-general position. By the choice of $\vec{\mathcal{F}}$ it follows that the grid $A_2 \times$

$A_2 \times A_3 \times A_3$ is in $\vec{\mathcal{F}}'$-general position, hence $|Q^* \cap (A_2^2 \times A_3^2)| \leq C'(\nu) n^{3-\gamma'}$. By Lemma 5.46 this implies

$$|Q \cap (A_1 \times A_2 \times A_3)| \leq d|A_1|^{\frac{1}{2}} |Q^* \cap (A_2^2 \times A_3^2)|^{\frac{1}{2}}$$
$$\leq d n^{\frac{1}{2}} C'(\nu)^{\frac{1}{2}} n^{\frac{3}{2} - \frac{\gamma'}{2}} \leq C(\nu) n^{2-\gamma}.$$

Hence $\mathcal{Q}$ satisfies γ-power saving. □

We are ready to finish the proof of Theorem 5.27 (and hence of Theorem 5.24), the required bound on power saving follows from the proof.

Proof of Theorem 5.27 for $s = 3$. By the reduction explained above we may assume that $\mathcal{Q}$ is a definable family of absolutely $\mathfrak{p}$-irreducible sets and does not satisfy 1-power saving. Applying the case $s = 4$ of Theorem 5.27 to the family $\mathcal{Q}^*$ (note that $\mathcal{Q}^*$ satisfies the assumption of Theorem 5.27 by the reduction above and since $\mathcal{Q}$ satisfies the $s = 3$ assumption of Theorem 5.27), we find a definable subfamily $(\mathcal{Q}^*)' \subseteq \mathcal{Q}^*$ such that the family $(\mathcal{Q}^*)'$ admits γ-power saving, and for each $Q^* \in \mathcal{Q}^* \setminus (\mathcal{Q}^*)'$ the relation Q^* is in a $\mathfrak{p}$-generic correspondence with an abelian group G_{Q^*} type-definable in $\mathcal{M}^{\mathrm{eq}}$ over a set of parameters of cardinality $\leq |\mathcal{L}|$.

Let $\mathcal{Q}_0$ be the set of all $Q \in \mathcal{Q}$ such that for some $u \subseteq [3]$ with $|u| = 2$ for the projection $\pi_u(Q)$ of Q onto $\prod_{i \in u} X_i$ we have $\dim_{\mathfrak{p}}(\pi_u(Q)) < 2$. By Claim 5.8, the family $\mathcal{Q}_0$ is definable and it follows from Proposition 5.30 that the family $\mathcal{Q}_0$ admits 1-power saving.

Consider the definable subfamily $\mathcal{Q}' := \left\{Q \in \mathcal{Q} : Q^* \in (\mathcal{Q}^*)'\right\} \cup \mathcal{Q}_0$ of $\mathcal{Q}$. By Lemma 5.47, as $\gamma \leq 1$, $\mathcal{Q}'$ admits $\frac{\gamma}{2}$-power saving. On the other hand, if $Q \in \mathcal{Q} \setminus \mathcal{Q}'$, then $Q^* \in \mathcal{Q}^* \setminus (\mathcal{Q}^*)'$, hence there exists a small set $A \subseteq M$ and an abelian group $(G, \cdot, 1_G)$ type-definable over A so that Q^* is in a $\mathfrak{p}$-generic correspondence with G.

This means that there exists a tuple $(g_2, g_2', g_3, g_3') \in G^4$ so that $g_2 \cdot g_2' \cdot g_3 \cdot g_3' = 1_G$, g_2, g_3, g_3' are independent generics over A and a tuple $(a_2, a_2', a_3, a_3') \in Q^*$ so that each of the elements a_2, a_2', a_3, a_3' is $\mathfrak{p}$-generic over A and each of the pairs $(g_2, a_2), (g_2', a_2'), (g_3, a_3), (g_3', a_3')$ is inter-algebraic over A.

By definition of Q^* there exists some $a_1 \in X_1$ such that $(a_1, a_2, a_3) \in Q$ and $(a_1, a_2', a_3') \in Q$. We let $A' := A a_3'$ and $g_1 := g_2' \cdot g_3'$, and make the following observations.

(1) $g_1 \cdot g_2 \cdot g_3 = 1_G$ (using that G is abelian).
(2) Each of the pairs $(a_1, g_1), (a_2, g_2), (a_3, g_3)$ is inter-algebraic over A'.
    The pairs $(a_2, g_2), (a_3, g_3)$ are inter-algebraic over A by assumption. Note that a_1 and a_2' are inter-algebraic over A' as Q is fiber-algebraic, so it suffices to show that a_2' and g_1 are inter-algebraic over A'. By definition $g_1 \in \mathrm{acl}(g_2', g_3') \subseteq \mathrm{acl}(a_2', a_3', A) \subseteq \mathrm{acl}(a_2', A')$. Conversely, as $g_2' \in \mathrm{acl}(g_2' \cdot g_3', g_3') \subseteq \mathrm{acl}(g_1, A')$, we have $a_2' \in \mathrm{acl}(g_2', A) \subseteq \mathrm{acl}(g_1, A')$.
(3) g_2 and g_3 are independent generics in G over A'.

By assumption $g_2 \mathop{\smile\!\!\!\!\!\!\mid}_A g_3 g_3'$ and a_3' is inter-algebraic with g_3' over A, hence $g_2 \mathop{\smile\!\!\!\!\!\!\mid}_{A'} g_3$.

(4) $a_i \models \mathfrak{p}_i|_{A'}$ for all $i \in [3]$.
For $i \in \{2,3\}$: as $g_i \mathop{\smile\!\!\!\!\!\!\mid}_A g_3'$ and g_3' is inter-algebraic with a_3' over A, we have $a_i \mathop{\smile\!\!\!\!\!\!\mid}_A a_3'$, which by stationarity of $\mathfrak{p}_i$ implies $a_i \models \mathfrak{p}_i|_{A'}$.
For $i = 1$: as $a_i \models \mathfrak{p}_i|_{A'}$ for $i \in \{2,3\}$ and $a_2 \mathop{\smile\!\!\!\!\!\!\mid}_{A'} a_3$, it follows that $(a_2, a_3) \models (\mathfrak{p}_2 \otimes \mathfrak{p}_3)|_{A'}$ and the tuple (a_1, a_2, a_3) is generic in Q over A' (as $\dim_{\mathfrak{p}}(Q) = 2$ by the choice of $\mathcal{Q}'$). But then $a_1 \models \mathfrak{p}_1|_{A'}$ by the assumption on Q and Proposition 5.32 (can be applied by absolute irreducibility of Q and the choice of $\mathcal{Q}'$).

It follows that Q is in a $\mathfrak{p}$-generic correspondence with G over A', witnessed by the tuples (g_1, g_2, g_3) and (a_1, a_2, a_3). □

5.8. **Discussion and some applications.** First we observe how Theorem 5.24, along with some standard facts from model theory of algebraically closed fields, implies a higher arity generalization of the Elekes-Szabó theorem for algebraic varieties over $\mathbb{C}$ similar to [8]. Recall from [8] that a *generically finite algebraic correspondence* between irreducible varieties V and V' over $\mathbb{C}$ is a closed irreducible subvariety $C \subseteq V \times V'$ such that the projections $C \to V$ and $C \to V'$ are generically finite and dominant (hence necessarily $\dim(V) = \dim(V')$). And assuming that W_i, W_i' and $V \subseteq \prod_{i \in [s]} W_i, V' \subseteq \prod_{i\in[s]} W_i'$ are irreducible algebraic varieties over $\mathbb{C}$, we say that V and V' are in *coordinate-wise correspondence* if there is a generically finite algebraic correspondence $C \subseteq V \times V'$ such that for each $i \in [s]$, the closure of the projection $(\pi_i \times \pi_i')(C) \subseteq W_i \times W_i'$ is a generically finite algebraic correspondence between the closure of $\pi_i(V)$ and the closure of $\pi_i'(V')$.

Corollary 5.48. *Assume that $s \geq 3$, and $X_i \subseteq \mathbb{C}^{m_i}, i \in [s]$ and $Q \subseteq \prod_{i\in[s]} X_i$ are irreducible algebraic varieties, with $\dim(X_i) = d$. Assume also that for each $i \in [s]$, the projection $Q \to \prod_{j\in[s]\setminus\{i\}} X_i$ is dominant and generically finite. Let $m := (m_1, ..., m_s)$, $t := \max\{\deg(Q), \deg(X_1), ..., \deg(X_s)\}$. Then one of the following holds.*

(1) For every ν there exist $D = D(d, s, t, m)$ and $c = c(d, s, t, m, \nu)$ such that: for any n and finite $A_i \subseteq X_i, |A_i| = n$ such that $|A_i \cap Y_i| \leq \nu$ for every algebraic subsets Y_i of X_i of dimension $< d$ and degree $\leq D$, we have
$$|Q \cap A| \leq cn^{s-1-\gamma}$$
for $\gamma = \frac{1}{16d-5}$ if $s \geq 4$, and $\gamma = \frac{1}{2(16d-5)}$ if $s = 3$.

(2) There exists a connected abelian complex algebraic group $(G, \cdot)$ with $\dim(G) = d$ such that Q is in a coordinate-wise correspondence with
$$Q' := \{(x_1, \ldots, x_s) \in G^s : x_1 \cdot \ldots \cdot x_s = 1_G\}.$$

The above Corollary 5.48 immediately follows from the following slightly more general statement:

Corollary 5.49. *Assume that $s \geq 3$, and $X_i \subseteq \mathbb{C}^{m_i}, i \in [s]$ are irreducible algebraic varieties with $\dim(X_i) = d$, and let $\mathcal{Q}$ be a definable family of subsets of $\prod_{i\in[s]} X_i$, each of Morley degree 1. Assume also that for each $Q \in \mathcal{Q}$, $i \in [s]$, the projection $Q \to \prod_{j\in[s]\setminus\{i\}} X_i$ is Zariski dense and is generically finite to one. Then there is a definable family $\mathcal{Q}' \subseteq \mathcal{Q}$ such that:*

(1) $\mathcal{Q}'$ admits γ-power saving for $\gamma = \frac{1}{16d-5}$ if $s \geq 4$, and $\gamma = \frac{1}{2(16d-5)}$ if $s = 3$.

(2) For every $Q \in \mathcal{Q} \setminus \mathcal{Q}'$ there exists a connected abelian complex algebraic group $(G, \cdot)$ with $\dim(G) = d$ such that for some independent generics $g_1, \ldots, g_{s-1} \in G$ and generic $(q_1, \ldots, q_s) \in Q$ we have that g_i is inter-algebraic with q_i for $i < s$ and q_s inter-algebraic with $(g_1 \cdot g_2 \cdot \ldots \cdot g_{s-1})^{-1}$.

It is not hard to see that Corollary 5.49 implies 5.48. Indeed, if Q is an irreducible variety then it has Morley degree one. Let $\mathcal{Q}$ be the family of *all* irreducible algebraic varieties contained in $\prod_{i\in[s]} X_i$ of degree $\deg Q$, Morley rank $\mathrm{MR}(Q)$ and with all projections Zariski dense and generically finite to one. It is a definable family in $\mathcal{M}$ by definability of Morley rank and irreducibility (see e.g. [26, Theorem A.7]), defined by a formula depending only on m, t, s, d; and $Q \in \mathcal{Q}$. Applying Corollary 5.49 we can conclude depending on whether $Q \in \mathcal{Q}'$ or $Q \in \mathcal{Q} \setminus \mathcal{Q}'$.

Proof of Corollary 5.49. Let $\mathcal{M} := (\mathbb{C}, +, \times, 0, 1) \models \mathrm{ACF}$, then $|\mathcal{L}| = \aleph_0$ and $\mathcal{M}$ is an $|\mathcal{L}|^+$-saturated structure. We recall that $\mathcal{M}$ is a strongly minimal structure, in particular it is ω-stable and has additive Morley rank MR coinciding with the Zariski dimension (see e.g. [47]).

For each i, as X_i is irreducible, i.e. has Morley degree 1, let $\mathfrak{p}_i \in S_{x_i}(\mathcal{M})$ be the unique type in X_i with $\mathrm{MR}(\mathfrak{p}_i) = \mathrm{MR}(X_i) = d$. By stability, types are definable, commute and are stationary after naming a countable elementary submodel of $\mathcal{M}$ so that all of the X_i's are defined over it.

Hence $(X_i, \mathfrak{p}_i)_{i\in[s]}$ is a $\mathfrak{p}$-system; and by the additivity of Morley rank we see that $\mathrm{MR}(Y) \geq d \dim_{\mathfrak{p}}(Y)$ for any definable $Y \subseteq \prod_{i\in[s]} X_i$.

For any $Q \in \mathcal{Q}$, since the projection of Q onto $\prod_{i=1}^{s-1} X_i$ is Zariski dense and generically finite, we have $\mathrm{MR}(Q) = d(s-1)$.

Let $Q' \subseteq Q$ be a definable set with $\mathrm{RM}(Q') = d(s-1)$. Since Q and Q' have the same generic points, the item (2) is equivalent for Q and Q'. Obviously γ-power saving for Q implies γ-power saving for Q', and we observe that γ-power saving for Q' with $0 < \gamma < 1$ implies γ-power saving for Q. Let $Q'' := Q \setminus Q'$. Then, as Q has Morley degree 1, $\mathrm{MR}(Q'') < d(s-1)$, hence $\dim_{\mathfrak{p}}(Q'') \leq s-2$. Applying Lemma 5.19 to $\mathcal{G} := \{Y''\}$ we obtain that Y'' has 1-power saving. Since $\gamma < 1$, it follows that $Y = Y' \cup Y''$ also has γ-power saving.

As by assumption every $Q \in \mathcal{Q}$ has generically finite projections, after removing a subset of smaller Morley rank we may assume that Q is fiber-algebraic. This can be done uniformly for the family by [26, Theorem A.7] (however, on this step we have to pass from a family of algebraic sets to a family of constructible sets, that is why we can only use bounds from Corollary 2.16(2) but not from Fact 2.17 in the following), hence we may assume that the family $\mathcal{Q}$ consists of fiber algebraic sets of fixed degree.

As $\dim(X_i) = d$, it follows that X_i has a generically finite-to-one projection onto $\mathbb{C}^d$, hence, after possibly a coordinate-wise correspondence, we may assume that $Q \subseteq \prod_{i \in [s]} \mathbb{C}^d$ — again, uniformly for the whole family $\mathcal{Q}$. By Corollary 2.16(2), every definable family of sets $Y \subseteq \mathbb{C}^{2d} \times \mathbb{C}^{(s-2)d}$ satisfies the $\left(\frac{1}{8d-1}\right)$-ST property. Applying Theorem 5.27 (we are using once more that irreducible components are uniformly definable in families in ACF, see [26, Theorem A.7]) we find a definable subfamily $\mathcal{Q}'$ with γ-power saving for the stated γ.

Every type-definable group in $\mathcal{M}^{\mathrm{eq}}$ is actually definable (by ω-stability, see e.g. [40, Theorem 7.5.3]), and every group interpretable in an algebraically closed field is definably isomorphic to an algebraic group (see e.g. [47, Proposition 4.12 + Corollary 1.8]). Thus, for $Q \in \mathcal{Q} \setminus \mathcal{Q}'$, there exists an abelian connected complex algebraic group $(G, \cdot)$, independent generic elements $g_1, \ldots, g_{s-1} \in G$ and $g_s \in G$ such that $g_1 \cdot \ldots \cdot g_s = 1$ and generic $a_i \in X_i$ inter-algebraic with g_i, such that $(a_1, \ldots, a_s) \in Q$. In particular, $\dim(G) = \dim(X_i) = d$. And, by irreducibility of Q, hence uniqueness of the generic type, such a_i's exist *for any* independent generics $g_1, \ldots, g_{s-1}$. As the model-theoretic algebraic closure coincides with the field-theoretic algebraic closure, by saturation of $\mathcal{M}$ this gives the desired coordinate-wise correspondence. □

Remark 5.50. Failure of power saving on arbitrary grids, not necessarily in a general position, does not guarantee coordinate-wise correspondence with an *abelian* group in Corollary 5.48. For example, let $(H, \cdot)$ be the Heisenberg group of 3×3 matrices over $\mathbb{C}$, viewed as a definable group in $\mathcal{M} := (\mathbb{C}, +, \times)$. For $n \in \mathbb{N}$, consider the subset of H given by

$$A_n := \left\{ \begin{pmatrix} 1 & n_1 & n_3 \\ 0 & 1 & n_2 \\ 0 & 0 & 1 \end{pmatrix} : n_1, n_2, n_3 \in \mathbb{N}, n_1, n_2 < n, n_3 < n^2 \right\}.$$

It is not hard to see that $|A_n| = n^4$. For the definable fiber-algebraic relation $Q(x_1, x_2, x_3, x_4)$ on H^4 given by $x_1 \cdot x_2 = x_3 \cdot x_4$ we have $|Q \cap A_n^4| \geq \frac{1}{16}(n^4)^3 = \Omega(|A_n|^3)$.

However, Q is not in a generic correspondence with an *abelian* group. Indeed, the sets $A_n \subseteq H, n \in \mathbb{N}$ are not in an $(\mathcal{F}, \nu)$-general position for any ν, with respect to the definable family $\mathcal{F} = \{u_1 - u_2 = c : c \in \mathbb{C}\}$ of subsets of H.

However, restricting further to the case $\dim(X_i) = 1$ for all $i \in [s]$, the general position requirement is satisfied automatically: for any definable set $Y \subseteq X_i$, $\dim(Y) < 1$ if and only if Y is finite; and for every definably family $\mathcal{F}_i$ of subsets of X_i there exists some ν_0 such that for any $Y \in \mathcal{F}_i$, if Y has cardinality greater than ν_0 then it is infinite. Hence (using the classification of one-dimensional connected complex algebraic groups) we obtain the following simplified statement.

Corollary 5.51. *Assume $s \geq 3$, and let $Q \subseteq \mathbb{C}^s$ be an irreducible algebraic variety so that for each $i \in [s]$, the projection $Q \to \prod_{j \in [s] \setminus \{i\}} \mathbb{C}^s$ is generically finite. Then exactly one of the following holds.*

(1) There exist c depending only on $s, \deg(Q)$ such that: for any n and $A_i \subseteq \mathbb{C}_i, |A_i| = n$ we have

$$|Q \cap A| \leq cn^{s-1-\gamma}$$

for $\gamma = \frac{1}{11}$ if $s \geq 4$, and $\gamma = \frac{1}{22}$ if $s = 3$.

(2) For G one of $(\mathbb{C}, +)$, $(\mathbb{C}, \times)$ or an elliptic curve, Q is in a coordinate-wise correspondence with

$$Q' := \{(x_1, \ldots, x_s) \in G^s : x_1 \cdot \ldots \cdot x_s = 1_G\}.$$

Remark 5.52. We expect that the two cases in Corollary 5.48 are not mutually exclusive (a potential example is suggested in [12, Remark 7.14]), however they are mutually exclusive in the 1-dimensional case in Corollary 5.51. The proof of this for $s = 3$ is given in [20, Proposition 1.7], and the argument generalizes in a straightforward manner to an arbitrary s.

We remark that the case of complex algebraic varieties corresponds to a rather special case of our general Theorem 5.24 which also applies e.g. to the theories of differentially closed fields or compact complex manifolds (see Facts 2.20 and 2.21). For example:

Remark 5.53. Given definable strongly minimal sets $X_i, i \in [s]$ and a fiber-algebraic $Q \subseteq \prod_{i \in [s]} X_i$ in a differentially closed field $\mathcal{M}$ of characteristic 0, we conclude that either Q has power saving (however, we do not have an explicit exponent here, see Problem 2.22), or that Q is in correspondence with one of the following strongly minimal differential-algebraic groups: the additive, multiplicative or elliptic curve groups over the field of constants $\mathcal{C}_{\mathcal{M}}$ of $\mathcal{M}$, or a *Manin kernel* of a simple abelian variety A that does not descend to $\mathcal{C}_{\mathcal{M}}$ (i.e. the Kolchin closure of the torsion subgroup of A; we rely here on the Hrushovski-Sokolovic trichotomy theorem, see e.g. [42, Section 2.1]).

6. Main theorem in the o-minimal case

6.1. Main theorem and some reductions.

In this section we will assume that $\mathcal{M} = (M, \ldots)$ is an o-minimal, $\aleph_0$-saturated $\mathcal{L}$-structure expanding a group (or just with definable Skolem functions). We shall use several times the following basic property of o-minimal structures:

Fact 6.1. [45, Fact 2.1] *Assume that $a \in M^n$ and $A \subseteq B \subseteq M$ are small sets. For every definable open neighborhood U of a (defined over arbitrary parameters), there exists $C \supseteq A$,* acl*-independent from aB over A, and a*

C-definable open $W \subseteq U$ containing a. In particular, $\dim(a/A) = \dim(a/C)$ *and* $\dim(aB/C) = \dim(aB/A)$.

For the rest of the section we assume that $s \geq 3$ and for $i = 1, \dots, s$, we have $\emptyset$-definable sets X_i with $\dim X_i = m$ for all $i \in [s]$ (throughout the section, dim refers to the standard notion of dimension in o-minimal structures). We also have an $\emptyset$-definable set $Q \subseteq \overline{X} := X_1 \times \cdots \times X_s$, with $\dim Q = (s-1)m$, and such that Q is fiber algebraic of degree d, for some $d \in \mathbb{N}$ (see Definition 1.4).

The following is the equivalent of Definitions 5.17 and 5.21 in the o-minimal setting.

Definition 6.2. For $\gamma \in \mathbb{R}_{>0}$, we say that $Q \subseteq \overline{X}$ satisfies *γ-power saving* if there are definable families $\overrightarrow{\mathcal{F}} = (\mathcal{F}_1, \dots, \mathcal{F}_s)$, where each $\mathcal{F}_i$ consists of subsets of X_i of dimension smaller than m, such that for every $\nu \in \mathbb{N}$ there exists a constant $C = C(\nu)$ such that: for every $n \in \mathbb{N}$ and every n-grid $\overline{A} := A_1 \times \cdots \times A_k \subseteq \overline{X}, |A_i| = n$ in $(\overrightarrow{\mathcal{F}}, \nu)$-general position (i.e. for every $i \in [s]$ and $S \in \mathcal{F}_i$ we have $|A_i \cap S| \leq \nu$) we have

$$|Q \cap \overline{A}| \leq Cn^{(s-1)-\gamma}.$$

It is easy to verify that if $Q_1, Q_2 \subseteq \overline{X}$ satisfy γ-power saving then so does $Q_1 \cup Q_2$. Before stating our main theorem in the o-minimal case, we define:

Definition 6.3. Given a finite tuple a in an o-minimal structure $\mathcal{M}$, we let $\mu_{\mathcal{M}}(a)$ be the *infinitesimal neighborhood of* a, namely the intersection of all $\mathcal{M}$-definable open neighborhoods of a. It can be viewed as a partial type over $\mathcal{M}$, or we can identify it with the set of its realizations in an elementary extension of $\mathcal{M}$.

Theorem 6.4. *Under the above assumptions, one of the following holds.*

(1) The set Q has γ-power saving, for $\gamma = \frac{1}{8m-5}$ if $s \geq 4$, and $\gamma = \frac{1}{16m-10}$ if $s = 3$.

(2) There exist (i) a tuple $\bar{a} = (a_1, \dots, a_s)$ in $\mathcal{M}$ generic in Q, (ii) a substructure $\mathcal{M}_0 := \operatorname{dcl}(\bar{a})$ of $\mathcal{M}$ of cardinality $\leq |\mathcal{L}|$ (iii) a type-definable abelian group $(G, +)$ of dimension m, defined over M_0 and (iv) M_0-definable bijections $\pi_i : \mu_{\mathcal{M}_0}(a_i) \cap X_i \to G, i \in [s]$, sending a_i to $0 = 0_G$, such that

$$\pi_1(x_1) + \cdots + \pi_s(x_s) = 0 \Leftrightarrow Q(x_1, \dots, x_s)$$

for all $x_i \in \mu_{\mathcal{M}_0}(a_i) \cap X_i, i \in [s]$.

We begin working towards a proof of Theorem 6.4.

Notation

(1) For $i, j \in [s]$, we write $\overline{X}_{i,j}$ for the set $\prod_{\ell \neq i,j} X_\ell$.

(2) For $z \in X_1 \times X_2$ and $V \subseteq \overline{X}_{1,2}$ we write

$$Q(z, V) := \{w \in V : (z, w) \in Q\}.$$

We similarly write $Q(U, w)$, for $U \subseteq X_1 \times X_2$ and $w \in \overline{X}_{1,2}$.

Lemma 6.5. *The following are easy to verify:*

(1) For every $z \in X_1 \times X_2$, $\dim Q(z, \overline{X}_{1,2}) \leq (s-3)m$.

(2) If $\alpha = (z, w) \in (X_1 \times X_2) \times \overline{X}_{1,2}$ is generic in Q then for every neighborhood $U \times V$ of α, $\dim Q(z, V) = (s-3)m$ and $\dim Q(U, w) = m$.

We will need to consider a certain *local* variant of the property (P2) from Section 3.2.

Definition 6.6. Assume that $\alpha = (z, w) \in Q \cap (X_1 \times X_2) \times \overline{X}_{1,2}$.

- We say that Q *has the* $(P2)_{1,2}$ *property near* α if for all $U' \subseteq X_1 \times X_2$ and $V' \subseteq \overline{X}_{1,2}$ neighborhoods of z, w respectively,

$$\dim Q(U', w) = m \text{ and } \dim Q(z, V') = (s-3)m, \tag{6.1}$$

  and there are open neighborhoods $U \times V \ni (z, w)$ in $(X_1 \times X_2) \times \overline{X}_{i,j}$ such that

$$Q(U, w) \times Q(z, V) \subseteq Q, \tag{6.2}$$

  (namely, for every $z_1 \in U$ and $w_1 \in V$, if $(z_1, w), (z, w_1) \in Q$ then $(z_1, w_1) \in Q$).
- We say that Q *satisfies the* $(P2)_{i,j}$*-property near* α, for $1 \leq i < j \leq s$, if the above holds under the coordinate permutation of $1, 2$ and i, j, respectively.
- We say that Q *satisfies the* $(P2)$*-property near* α if it has the $(P2)_{i,j}$-property for all $1 \leq i < j \leq s$.

Remark 6.7. Note that if U, V satisfy (6.2), then also every $U_1 \subseteq U$ and $V_1 \subseteq V$ satisfy it. Note also that under the above assumptions, we have $\dim(Q(U, w) \times Q(z, V)) = (s-2)m$.

Definition 6.8. • Let $Q^*_{i,j}$ be the set of all $\alpha \in Q$ such that Q satisfies $(P2)_{i,j}$ near α.

- Let $Q^* = \bigcap_{i \neq j} Q^*_{i,j}$ be the set of all $\alpha \in Q$ near which Q satisfies $(P2)$.

Clearly, $Q^*_{i,j}$ and Q^* are $\emptyset$-definable sets.

The main ingredient towards the proof of Theorem 6.4 is the following:

Theorem 6.9. *Assume that Q does not satisfy γ-power saving for γ as in Theorem 6.4(1). Then $\dim Q^* = \dim Q = (s-1)m$.*

6.2. **The proof of Theorem 6.9.** The following is an analog of Lemma 5.19 in the *o*-minimal setting.

Lemma 6.10. *Let $\{Z_t : t \in T\}$ be a definable family of subsets of $\overline{X}$, each fiber-algebraic of degree $\leq d$ with $\dim(Z_t) < (s-1)m$. Then there exist definable families $\mathcal{F}_i$, $i \in [s]$, each consisting of subsets of X_i of dimension smaller than m, such that for every $\nu \in \mathbb{N}$, if $\bar{A} \subseteq \overline{X}$ is an n-grid in $(\overrightarrow{\mathcal{F}}, \nu)$-general position then for every $t \in T$,*

$$|\bar{A} \cap Z_t| \leq sd(\nu - 1)n^{s-2}.$$

In particular, each $Z_t, t \in T$ satisfies 1-power saving.

Proof. For $t \in T$ and $a_1 \in X_1$ we let

$$Z_{ta_1} := \{(a_2, \dots, a_s) \in X_2 \times \cdots \times X_s : (a_1, a_2, \dots, a_s) \in Z_t\}.$$

For $i \in [s-1]$, we similarly define $Z_{ta_1 \cdots a_i} \subseteq X_{i+1} \times \cdots \times X_s$.
(1) For $t \in T$, we let

$$Y_t^1 := \{a_1 \in X_1 : \dim(Z_{ta_1}) = (s-2)m\}.$$

By our assumption on $\dim Z_t$, $\dim Y_t^1 < m$. Let $\mathcal{F}_1 := \{Y_t^1 : t \in T\}$.

(2) For $t \in T$ and $a_1 \notin Y_t^1$, let

$$Y_{ta_1}^2 := \{a_2 \in X_2 : \dim(Z_{ta_1 a_2}) = (s-3)m\}.$$

Then define $\mathcal{F}_2 := \left\{Y_{ta_1}^2 : t \in T, a_1 \notin Y_t^1\right\}$. Note that whenever $a_1 \notin Y_t^1$, $\dim(Z_{ta_1}) < (s-2)m$ and therefore the set $Y_{ta_1}^2$ has dimension smaller than m.

For $i = 1, \dots, s-2$, we continue in this way to define a family $\mathcal{F}_i = \{Y_{ta_1 \cdots a_{i-1}}^i\}$ of subsets of X_i as follows: for $a_1 \notin Y_t^1$, $a_2 \notin Y_{ta_1}^2$, $a_3 \notin Y_{ta_1 a_2}^3, \dots, a_{i-1} \notin Y_{ta_1 \cdots a_{i-2}}^{i-1}$, we let

$$Y_{ta_1 \cdots a_{i-1}}^i := \{a_i \in X_i : \dim(Z_{ta_1 \cdots a_i}) = (s-(i+1))m\},$$

and let

$$\mathcal{F}_i := \left\{Y_{ta_1 \cdots a_{i-1}}^i : t \in T, a_1 \notin Y_t^1, a_2 \notin Y_{ta_1}^2, \dots, a_{i-1} \notin Y_{ta_1 \cdots a_{i-2}}^{i-1}\right\}.$$

Finally, for $a_1, \dots, a_{s-2}$ such that $a_i \notin Y_{ta_1 \cdots a_{i-1}}^i$ for $i = 1, \dots, s-2$, we let

$$Y_{ta_1 \cdots a_{s-2}}^{s-1} := \pi_{s-1}(Z_{ta_1 \dots a_{s-2}}) \subseteq X_{s-1},$$

and let

$$\mathcal{F}_{s-1} := \left\{Y_{ta_1 \cdots a_{s-2}}^{s-1} : t \in T, a_1 \notin Y_t^1, \dots, a_{s-2} \notin Y_{ta_1 \cdots a_{s-3}}^{s-2}\right\}.$$

We provide some details on why the families $\vec{\mathcal{F}} := (\mathcal{F}_i : i \in [s])$ satisfy the requirement.

Assume that $n, \nu \in \mathbb{N}$ and $\bar{A} \subseteq \overline{X}$ is an n-grid which is in $(\overrightarrow{\mathcal{F}}, \nu)$-general position, and fix $t \in T$.

Because $|A_1 \cap Y_t^1| < \nu$ there are at most $\nu - 1$ elements $a_1 \in \pi_1(Z_t \cap \bar{A}) \cap Y_t^1$, and for each such a_1 there are at most dn^{s-2} elements in $Z \cap \bar{A}$ which project to it. Indeed, this is true because Z_{ta_1} is fiber-algebraic of degree $\leq d$, so for every tuple $(a_2, \dots, a_{s-1}) \in A_2 \times \cdots A_{s-1}$ (and there are at most n^{s-2} such tuples) there are $\leq d$ elements $a_s \in A_s$ such that $(a_2, \dots, a_{s-1}, a_s) \in (A_2 \times \cdots \times A_s) \cap Z_{ta_1}$.

So, altogether there are at most $d(\nu-1)n^{s-2}$ elements $(a_1, \dots, a_s) \in \bar{A} \cap Z_t$ for which $a_1 \in Y_1^t$. On the other hand, there are at most $n - \nu \leq n$ elements $a_1 \notin Y_t^1$. We now compute for how many $\bar{a} \in \bar{A} \cap Z_t$ we have $a_1 \notin Y_t^1$.

By definition, $\dim(Z_{ta_1}) < (s-2)m$, so now we consider two cases, $a_2 \in Y_{ta_1}^2$ and $a_2 \notin Y_{ta_1}^2$. In the first case, there are at most $\nu - 1$ such a_2, by general position, and as above, for each such a_2 there are at most dn^{s-3}

tuples $(a_3, \dots, a_s) \in A_3 \times \cdots \times A_s$ such that $(a_2, a_3, \dots, a_s) \in Z_{ta_1}$. Thus all together there are $n(\nu - 1)dn^{s-3} = d(\nu - 1)n^{s-2}$ elements $\bar{a} \in \bar{A} \cap Z_t$ such that $a_1 \notin Y_t^1$ and $a_2 \in Y_t^2$. There are at most $(n - \nu) \leq n$ elements $a_2 \in A_2$ which are not in $Y_{ta_1}^2$. Of course, there are at most n^2 pairs (a_1, a_2) such that $a_1 \notin Y_t^1$ and $a_2 \notin Y_{ta_1}^2$, and we now want to compute how many $\bar{a} \in \bar{A} \cap Z_t$ project onto such (a_1, a_2). Repeating the same process along the other coordinates, we see that there are at most $(s - 2)d(\nu - 1)n^{s-4}$ elements which project into each such (a_1, a_2), so all together there are at most $(s-2)d(\nu-1)n^{s-2}$ tuples $\bar{a} \in \bar{A} \cap Z_t$ for which $a_1 \notin Y_t^1$ and $a_2 \notin Y_{ta_1}^2$. If we add it all we get at most $sd(\nu - 1)n^{s-2}$ elements in $\bar{A} \cap Z_t$, which concludes the proof of the lemma. □

Corollary 6.11. *Assume that $Q \subseteq \overline{X}$ does not satisfy 1-power saving and that $Z \subseteq Q$ is a definable set with $\dim Z < (s-1)m$. Then $Q' := Q \setminus Z$ also does not satisfy 1-power saving.*

Proof. Indeed, Lemma 6.10 (applied to the constant family) implies that Z itself satisfies 1-power saving, and since γ-power saving is preserved under union then it fails for Q'. □

In order to prove Theorem 6.9, it is sufficient to prove the following:

Proposition 6.12. *Let $Q' \subseteq Q$ be a definable set and assume that there exist $i \neq j \in [s]$ such that $\dim(Q' \cap Q^*_{i,j}) < (s-1)m$. Then Q' satisfies γ-power saving for γ as in Theorem 6.4(1).*

Let us first see that indeed Proposition 6.12 quickly implies Theorem 6.9. Let γ be as in Theorem 6.4(1). Assuming that Q does not have γ-power saving, Proposition 6.12 with $Q' := Q$ implies that $\dim(Q^*_{1,2}) = (s-1)m$. Also, if we take $Q'' := Q \setminus Q^*_{1,2}$ then clearly $Q'' \cap Q^*_{1,2} = \emptyset$ and therefore by the same proposition Q'' satisfies γ-power saving, and therefore $Q^*_{1,2}$ does not satisfy γ-power saving. We can thus replace Q by $Q_1 := Q^*_{1,2}$ and retain the original properties of Q together with the fact that Q_1 has $(P2)_{1,2}$ at every $\alpha \in Q_1$. Next we repeat the process with respect to every $(i,j) \neq (1,2)$ and eventually obtain a definable set $Q' \subseteq Q$ of dimension $(s-1)m$ such that Q' satisfies $(P2)$ at every point — establishing Theorem 6.9.

Proof of Proposition 6.12.

Let $Q' \subseteq Q$ and γ be as in Proposition 6.12. It is sufficient to prove the proposition for $Q^*_{1,2}$ (the case of arbitrary $i \neq j \in [s]$ follows by permuting the coordinates accordingly). If $\dim Q' < (s-1)m$ then by Lemma 6.10 Q' satisfies 1-power saving, hence γ-power saving. Thus we may assume that $\dim Q' = (s-1)m$, and hence, by throwing away a set of smaller dimension, we may assume that Q' is open in Q. It is then easy to verify that $(Q')^*_{1,2} = Q^*_{1,2} \cap Q'$. Hence, without loss of generality, $Q = Q'$. We now assume that $\dim Q^*_{1,2} < (s-1)m$ and therefore, by Lemma 6.10, $Q^*_{1,2}$ has γ-power saving. Thus, in order to show that Q has γ-power saving, it is sufficient to prove that $Q \setminus Q^*_{1,2}$ has γ-power saving, so we assume from now on that $Q^*_{1,2} = \emptyset$.

We let $U := X_1 \times X_2$ and $V := \overline{X}_{1,2}$.

Claim 6.13. *For every $w \in V$, the set*

$$X_w := \{w' \in V : \dim(Q(U,w) \cap Q(U,w')) = m\}$$

has dimension strictly smaller than $(s-3)m$. Moreover, the set X_w is fiber algebraic in $X_3 \times \cdots \times X_s$.

Proof. We assume that relevant sets thus far (i.e. $X_i, Q, U, V, Q^*_{i,j}$) are defined over $\emptyset$. Now, if $\dim(X_w) = (s-3)m$ (it is not hard to see that it cannot be bigger), then by $\aleph_0$-saturation of $\mathcal{M}$ we may take w' generic in X_w over w, and then u' generic in $Q(U,w) \cap Q(U,w')$ over w, w'. Note that the fiber-algebraicity of Q implies that $\dim(Q(u',V)) \leq (s-3)m$, and since $\dim(w'/wu') = \dim(w'/w) = (s-3)m$ it follows that w' is generic in both X_w and $Q(u',V)$ over wu', so in particular, $\dim X_w = \dim Q(u',V) = (s-3)m$.

We claim that $(u',w') \in Q^*_{1,2}$. Indeed, by our assumption,

$$\dim(u'/ww') = \dim(Q(U,w) \cap Q(U,w')) = \dim Q(U,w) = m.$$

Thus, there exists an open $U_0 \subseteq U$ containing u', such that $U_0 \cap Q(U,w) = U_0 \cap Q(U,w')$, or, said differently, $Q(U_0,w) = Q(U_0,w')$. By Fact 6.1, we may assume that the tuple (w,w',u') is independent from the parameters defining U_0 over $\emptyset$. Thus, without loss of generality, U_0 is definable over $\emptyset$. The set $W_1 := \{v \in V : Q(U_0,w) \subseteq Q(U,v)\}$ is defined over w and the set $Q(u',V)$ is defined over u', and both contain w'. Since $\dim(w'/w,u') = (s-3)m$ then $\dim(W_1 \cap Q(u',V)) = (s-3)m$. We can therefore find an open $V_0 \subseteq V$ such that $Q(u',V_0) \subseteq W_1$. Now, by the definition of W_1, we have $Q(U_0,w) \times W_1 \subseteq Q$, and hence $Q(U_0,w) \times Q(u',V_0) \subseteq Q$ and therefore (since $Q(U_0,w) = Q(U_0,w')$), $Q(U_0,w') \times Q(u,V_0) \subseteq Q$. This shows that $(u',w') \in Q^*_{1,2}$, contradicting our assumption that $Q^*_{1,2} = \emptyset$.

To see that X_w is fiber algebraic, assume towards contradiction that there exists a tuple $(a_3,\ldots,a_{s-1}) \in X_3 \times \cdots \times X_{s-1}$ for which there are infinitely many $a_s \in X_s$ such that $(a_3,\ldots,a_s) \in X_w$ (the other coordinates are treated similarly). We can now pick such a_s generic over $w, a_3, \ldots, a_{s-1}$ and then pick $(a_1,a_2) \in Q(U,w) \cap Q(U,a_3,\ldots,a_s)$ generic over $w, a_3, \ldots, a_s$. Because $\dim(a_1,a_2/w) = \dim(a_1,a_2/w,a_3,\ldots,a_s)$ it follows by the additivity of dimension that for any subtuple a' of $a_3,\ldots,a_s$ we have $\dim(a'/w,a_1,a_2) = \dim(a'/w)$. It follows that

$$0 < \dim(a_s/w,a_3,\ldots,a_{s-1}) = \dim(a_s/w,a_1,a_2,a_3,\ldots,a_{s-1}).$$

Since $Q(a_1,a_2,a_3,\ldots,a_s)$ holds, it follows that $Q(a_1,a_2,a_3,\ldots,a_{s-1},X_n)$ is infinite — contradicting the fiber-algebraicity of Q. ☐

We similarly have:

Claim 6.14. *For every $u \in U$, the set*

$$X^u := \{u' \in U : \dim(Q(u,V) \cap Q(u',V)) = (s-3)m\}$$

has dimension strictly smaller than m. Moreover, the set X^u is fiber-algebraic in $X_1 \times X_2$.

Lemma 6.15. *There exist s definable families $\vec{\mathcal{F}} = (\mathcal{F}_1, \dots, \mathcal{F}_s)$ of subsets of $X_1, \dots, X_s$, respectively, each containing only sets of dimension strictly smaller than m, such that for every $\nu \in \mathbb{N}$ and every n-grid $\bar{A} \subseteq \overline{X}$ in $(\vec{\mathcal{F}}, \nu)$-general position, we have the following.*

(1) For all $w, w' \in A_3 \times \cdots \times A_s$, if $|Q(A_1 \times A_2, w) \cap Q(A_1 \times A_2, w')| \geq d\nu$ then $w' \in X_w$.

(2) For all $w \in A_3 \times \cdots \times A_s$, there are at most $C(\nu)n^{s-4}$ elements $w' \in A_3 \times \cdots \times A_s$ such that $|Q(A_1 \times A_2, w) \cap Q(A_1 \times A_2, w')| \geq d\nu$.

(3) $|\bar{A} \cap Q| \leq C'(\nu)n^{(s-1)-\gamma}$.

Proof. We choose the definable families in $\vec{\mathcal{F}}$ as follows. Let

$$\mathcal{F}_1 := \{\pi_1(Q(U,w) \cap Q(U,w')) : \\ w, w' \in V \;\&\; \dim\left(Q(U,w) \cap Q(U,w')\right) < m\},$$

and $\mathcal{F}_2 := \{\emptyset\}$. Clearly, each set in $\mathcal{F}_1$ has dimension smaller than m. Because Q is fiber algebraic of degree $\leq d$, it is easy to verify that (1) holds independently of the other families.

For the other families, we first recall that by Claim 6.13, for each $w \in \overline{X}_{1,2}$, the set $X_w \subseteq \overline{X}_{1,2}$ has dimension smaller than $(s-3)m$.

We now apply Lemma 6.10 to the family $\{X_w : w \in \overline{X}_{1,2}\}$ (note that s from Lemma 6.10 is replaced here by $s-2$), and obtain definable families $\vec{\mathcal{F}}' = (\mathcal{F}_3, , \dots, \mathcal{F}_s)$, each $\mathcal{F}_i$ consisting of subsets of X_i of dimension smaller than m, such that for every ν and every n-grid $A_3 \times \cdots \times A_s \subseteq \overline{X}_{1,2}$ in $(\vec{\mathcal{F}}', \nu)$-general position and every $w \in \overline{X}_{1,2}$ we have

$$|\,(A_3 \times \cdots \times A_s) \cap X_w| \leq C(\nu)n^{s-4}.$$

Let $\vec{\mathcal{F}} := (\mathcal{F}_1, \mathcal{F}_2, \vec{\mathcal{F}}')$ and assume that $\bar{A}$ is in $(\vec{\mathcal{F}}, \nu)$-general position. It follows that for every $w \in A_3 \times \cdots \times A_s$ there are at most $C(\nu)n^{s-4}$ elements $w' \in A_3 \times \cdots \times A_s$ such that $|Q(A_1 \times A_2, w) \cap Q(A_1 \times A_2, w')| \geq d\nu$. This proves (2).

We claim that the relation Q, viewed as a binary relation on $(X_1 \times X_2) \times \overline{X}_{1,2}$, satisfies the γ-ST property. Indeed, for $i \in [s]$, let $X_i = \bigsqcup_{\ell \in [k_i]} X_{i,\ell}$ be an o-minimal cell decomposition of X_i, for some $k_i \in \mathbb{N}$, we have $m = \dim(X_i) = \max\{\dim(X_{i,\ell}) : \ell \in [k_i]\}$. Then each (definable) cell $X_{i,\ell}$ is in a definable bijection with a definable subset of $M^{\dim\left(X_{i,\ell}\right)}$ (namely, the projection on the appropriate coordinates is a homeomorphism), hence in a definable bijection with a definable subset of M^m. For $\bar{\ell} = (\ell_1, \dots, \ell_s) \in [k_1] \times \ldots \times [k_s]$, let $Q_{\bar{\ell}} := Q \cap \prod_{i \in [s]} X_{i,\ell_i}$. Applying these definable bijections coordinate-wise, by Lemma 2.1(1) we may assume

$Q_{\bar{\ell}} \subseteq \prod_{i\in[s]} M^m$ and apply Fact 2.15 to conclude that for each $\bar{\ell}$, $Q_{\bar{\ell}}$ satisfies the γ-ST property. But then, by Lemma 2.1(2), Q also satisfies the γ-ST property. Finally, given an n-grid $\bar{A} \subseteq (X_1 \times X_2) \times \overline{X}_{1,2}$ in $(\vec{\mathcal{F}}, \nu)$-general position, we thus have by the γ-ST property that (2) implies (3). □

This shows that Q has γ-power saving, in contradiction to our assumption, thus completing the proof of Proposition 6.12, and with it Theorem 6.9.

6.3. **Obtaining a nice Q-relation.** By Theorem 6.9 we may assume that $\dim Q = \dim Q^*$. Thus, in order to prove Theorem 6.4, we may replace Q by Q^*, and assume from now on that $Q = Q^*$.

Using o-minimal cell decomposition, we may partition Q into finitely many definable sets such that each is *fiber-definable*, namely for each tuple $(a_1, \dots, a_{s-1}) \in A_1 \times \dots \times A_{s-1}$, there exists at most one

$$a_s = f(a_1, \dots, a_{s-1}) \in X_s$$

such that $(a_1, \dots, a_{s-1}, a_s) \in Q$, and furthermore f is a continuous function on its domain. We can do the same for all permutations of the variables. Since Q does not satisfy γ-power saving by assumption, one of these finitely many sets, of dimension $(s-1)m$, also does not satisfy γ-power saving.

Hence from now on we assume that Q is the graph of a continuous partial function from any of its $s-1$ variables to the remaining one.

By further partitioning Q and changing the sets up to definable bijections, we may assume that each X_i is an open subset of M^m. Fix $\bar{e} = (e_1, \dots, e_s)$ in $\mathcal{M}$ generic in Q, and let $\mathcal{M}_0 := \operatorname{dcl}(\bar{e})$. Note that for each $(a_3, \dots, a_s)$ in a neighborhood of $(e_3, \dots, e_s)$, the set $Q(x_1, x_2, a_3, \dots, a_s)$ is the graph of a homeomorphism between neighborhoods of e_1 and e_2. We let $\mu_i := \mu_{\mathcal{M}_0}(e_i)$ (see Definition 6.3) and identify these partial types over $\mathcal{M}_0$ with their sets of realizations in $\mathcal{M}$.

Lemma 6.16. *There exist $\mathcal{M}_0$-definable relatively open sets $U \subseteq X_1 \times X_2$ and $V \subseteq \bar{X}_{1,2}$, containing (e_1, e_2) and $(e_3, \dots, e_s)$, respectively, and a relatively open $W \subseteq Q$, containing $\bar{e}$, such that for every $(u, v) \in W$, $Q(u, V) \times Q(U, v) \subseteq Q$.*

In particular, for any $u, u' \in \mu_{\mathcal{M}_0}(e_1, e_2) \cap (X_1 \times X_2)$ and any $v, v' \in \mu_{\mathcal{M}_0}(e_3, \dots, e_s) \cap \bar{X}_{1,2}$ we have

$$(u, v), (u, v'), (u', v) \in Q \Rightarrow (u', v') \in Q.$$

Proof. Because the properties of U, V and W are first-order expressible over $\bar{e}$, it is sufficient to prove the existence of U, V, W in any elementary extension of $\mathcal{M}_0$.

Because $\bar{e} \in Q = Q^*$, there are definable, relatively open neighborhoods $U \subseteq X_1 \times X_2$ and $V \subseteq \bar{X}_{1,2}$ of (e_1, e_2) and $(e_3, \dots, e_s)$, respectively, such that

$$Q(U, e_3, \dots, e_s) \times Q(e_1, e_2, V) \subseteq Q.$$

By Fact 6.1, we may assume that U, V are definable over $A \subseteq M$ such that $\bar{e}$ is still generic in Q over A. It follows that there exists a relatively open $W \subseteq Q$, containing $\bar{e}$, such that for every $(u, v) \in W$ (so, $u \in X_1 \times X_2$ and $v \in \bar{X}_{1,2}$), we have $Q(U, v) \times Q(u, V) \subseteq Q$. As already noted, we now can find such U, V and W defined over $\mathcal{M}_0$.

Note that $\mu_{\mathcal{M}_0}(e_1, e_2) \cap (X_1 \times X_2) \subseteq U$ and $\mu_{\mathcal{M}_0}(e_3, \ldots, e_n) \cap \overline{X}_{1,2} \subseteq V$, and $\mu_{\mathcal{M}_0}(\bar{e}) \subseteq W$. Let us see how the last clause follows: assume that $u, u' \in \mu_{\mathcal{M}_0}(e_1, e_2) \cap (X_1 \times X_2)$, $v, v' \in \mu_{\mathcal{M}_0}(e_3, \ldots, e_n) \cap \overline{X}_{1,2}$, and $(u, v), (u, v'), (u', v) \in Q$. We have $u, u' \in U$, $v, v' \in V$ and

$$(u, v), (u, v'), (u', v) \in W.$$

By the choice of U, V, W, we thus have $(u', v') \in Q$. □

Lemma 6.17. *The definable relation Q satisfies properties (P1) and (P2) from Section 3.2 with respect to the $\mathcal{M}_0$-type-definable sets $\mu_i \cap X_i, i \in [s]$, namely:*

(P1) For any $(a_1, \ldots, a_{s-1}) \in \mu_1 \times \cdots \times \mu_{s-1}$, there exists exactly one $a_s \in \mu_s$ with $(a_1, \ldots, a_{s-1}, a_s) \in Q$, and this remains true under any coordinate permutation.

(P2) Let $\tilde{U} := \mu_1 \times \mu_2 \cap X_1 \times X_2$ and $\tilde{V} := \mu_3 \times \ldots \times \mu_s \cap \bar{X}_{1,2}$. Then for every $u, u' \in \tilde{U}$ and $w, w' \in \tilde{V}$,

$$(u; w), (u'; w), (u; w') \in Q \Rightarrow (u'; w') \in Q.$$

The same is true when $(1, 2)$ is replaced by any (i, j) with $i \neq j \in [s]$.

Proof. By continuity of the function given by Q, for every tuple

$$(a_1, \ldots, a_{s-1}) \in \mu_1 \times \cdots \times \mu_{s-1}$$

there exists a unique $a_s \in \mu_s$ such that $(a_1, \ldots, a_s) \in Q$. The same is true for any permutation of the variables. This shows (P1).

Property (P2) holds by Lemma 6.16 for the $(1, 2)$-coordinates. The same proof works for the other pairs (i, j). □

Let us see how Theorem 6.4 follows. Assume first that $s \geq 4$, and that Q does not have γ-power saving for $\gamma = \frac{1}{16s-10}$. By Theorem 6.9 and the resulting Lemma 6.17 (see also the choice of the parameters before Lemma 6.16), there is $\bar{e} = (e_1, \ldots, e_s)$ generic in Q and a substructure $\mathcal{M}_0 = \mathrm{dcl}(\bar{e})$ of cardinality $|\mathcal{L}|$ such that $Q \cap \prod_{i \in [s]} (\mu_{\mathcal{M}_0}(e_i) \cap X_i)$ satisfies $(P1)$ and $(P2)$ of Theorem 3.21. Note that $\mu_{\mathcal{M}_0}(e_i)$ is a partial type over $\mathcal{M}_0$ for $i \in [s]$, $\bar{e}$ satisfies the relation, and $\bar{e}$ is contained in $\mathcal{M}_0$. Thus, by the "moreover" clause of Theorem 3.21, there exists a type definable abelian group G over $\mathcal{M}_0$ and $\mathcal{M}_0$-definable bijections $\pi'_i : \mu_{\mathcal{M}_0}(e_i) \cap X_i \to G$ each sending e_i to 0_G and satisfying:

$$Q(a_1, \ldots, a_n) \Leftrightarrow \pi'_1(a_1) + \cdots + \pi'_m(a_m) = 0$$

for all $a_i \in \mu_{\mathcal{M}_0}(e_i) \cap X_i$. This is exactly the second clause of Theorem 6.4.

Finally, the case $s = 3$ of Theorem 6.4 reduces to the case $s = 4$ as in the stable case, Section 5.7, with the obvious modifications.

6.4. **Discussion and some applications.** We discuss some variants and corollaries of the main theorem. In particular, we will deduce a variant that holds in an arbitrary *o*-minimal structure, i.e. without the saturation assumption on $\mathcal{M}$ used in Theorem 6.4.

Definition 6.18. (see [27, Definition 2.1]) A *local group* is a tuple $(\Gamma, 1, \iota, p)$, where Γ is a Hausdorff topological space, $\iota : \Lambda \to \Gamma$ (the inversion map) and $p : \Omega \to \Gamma$ (the product map) are continuous functions, with $\Lambda \subseteq \Gamma$ and $\Omega \subseteq \Gamma^2$ open subsets, such that $1 \in \Lambda$, $\{1\} \times \Gamma, \Gamma \times \{1\} \subseteq \Omega$ and for all $x, y, z \in \Gamma$:

(1) $p(x, 1) = p(1, x) = x$;
(2) if $x \in \Lambda$ then $(x, \iota(x)), (\iota(x), x) \in \Omega$ and $p(x, \iota(x)) = p(\iota(x), x) = 1$;
(3) if $(x, y), (y, z) \in \Omega$ and $(p(x, y), z), (x, p(y, z)) \in \Omega$, then

$$p((p(x, y), z) = p(x, p(y, z)).$$

Our goal is to show that in Theorem 6.4 we can replace the type-definable group with *a definable* local group. Namely,

Corollary 6.19. *Let $\mathcal{M}$ be an $\aleph_0$-saturated o-minimal expansion of a group, $s \geq 3$, $Q \subseteq X_1 \times \cdots \times X_s$ are $\emptyset$-definable with $\dim(X_i) = m$, and Q is fiber algebraic. Then one of the following holds.*

(1) The set Q has γ-power saving, for $\gamma = \frac{1}{8m-5}$ if $s \geq 4$, and $\gamma = \frac{1}{16m-10}$ if $s = 3$.
(2) There exist (i) a finite set $A \subseteq M$ and a structure $\mathcal{M}_0 = \mathrm{dcl}(A)$ (ii) a definable local abelian group Γ with $\dim(\Gamma) = m$, defined over M_0, (iii) definable relatively open $U_i \subseteq X_i$, a definable open neighborhood $V \subseteq \Gamma$ of $0 = 0_\Gamma$, and (iv) definable homeomorphisms $\pi_i : U_i \to V$, $i \in [s]$, such that for all $x_i \in U_i$,

$$\pi_1(x_1) + \cdots + \pi_s(x_s) = 0 \Leftrightarrow Q(x_1, \ldots, x_s).$$

Proof. We assume that (1) fails and apply Theorem 6.4 to obtain $\bar{a}$ generic in Q, $\mathcal{M}_0 = \mathrm{dcl}(\bar{a})$, a type-definable abelian group G over $\mathcal{M}_0$, and bijections $\pi_i : \mu_{\mathcal{M}_0}(a_i) \to G$ sending a_i to 0, such that for all $i \in [s]$, and $x_i \in \mu_{\mathcal{M}_0}(a_i)$,

$$\pi_1(x_1) + \cdots + \pi_s(x_s) = 0 \Leftrightarrow Q(x_1, \ldots, x_s).$$

By pulling back the group operations via, say, π_1, we may assume that the domain of G is $\mu_{\mathcal{M}_0}(a_1)$. We denote this pull-back of the addition and the inverse operations by $x \oplus y$ and $\ominus y$, respectively. Let us see that $\oplus$ and $\ominus$ are continuous with respect to the induced topology on $\mu_{\mathcal{M}_0}(a_1) \subseteq X_1$. Because $\bar{a}$ is generic in Q, and Q is fiber algebraic, it follows from *o*-minimality that the set $Q(x_1, x_2, x_3, a_4, \ldots, a_s)$ defines a continuous function from any two of the coordinates x_1, x_2, x_3 to the third one, on the corresponding infinitesimal types $\mu_{\mathcal{M}_0}(a_i) \times \mu_{\mathcal{M}_0}(a_j)$.

The following is easy to verify: for $x', x'', x''' \in \mu_{\mathcal{M}_0}(a_1)$, $x' \oplus x'' = x'''$ if and only if there exist $x_2 \in \mu_{\mathcal{M}_0}(a_2)$ and $x_3, x_3' \in \mu_{\mathcal{M}_0}(a_3)$ such that

$$\begin{array}{lll} Q\left(x', x_2, x_3, a_4, \ldots, a_s\right), & Q\left(x''', a_2, x_3, a_4, \ldots, a_s\right) & \text{and} \\ Q(x'', a_2, x_3', a_4, \ldots, a_s), & Q(a_1, x_2, x_3', a_4, \ldots, a_s). & \end{array}$$

By the above comments, $\oplus$ can thus be obtained as a composition of continuous maps, thus it is continuous. We similarly show that $\ominus$ is continuous.

Applying logical compactness, we may now replace the type-definable G with an $\mathcal{M}_0$-definable $\Gamma \supseteq G = \mu_{\mathcal{M}_0}(a_1)$, with partial continuous group operations, which make Γ into a local group (we note that in general, any type-definable group is contained in a definable local group by logical compactness, except for the topological conditions). Similarly, we find $U_i \supseteq \mu_{\mathcal{M}_0}(a_i)$, $V \subseteq \Gamma$ and $\pi_i : U_i \to V$ as needed. □

Note that if $\mathbb{R}_{\text{o-min}}$ is an o-minimal expansion of the field of reals and the X_i's and Q are definable in $\mathbb{R}_{\text{o-min}}$, with Q not satisfying Clause (1) of Corollary 6.19, then taking a sufficiently saturated elementary extension $\mathcal{M} \succeq \mathbb{R}_{\text{o-min}}$, $Q(\mathcal{M})$ still does not satisfy Clause (1) in $\mathcal{M}$. Hence we may deduce that Clause (2) of Corollary 6.19 holds for Q in $\mathcal{M}$, possibly over additional parameters from $\mathcal{M}$. However, the definition of a local group is first-order in the parameters defining Γ, ι and p. Thus, by elementarity, we obtain that Clause (2) of Corollary 6.19 holds for $Q(\mathbb{R})$, with Γ and the functions π_i definable in the original structure $\mathbb{R}_{\text{o-min}}$.

By Goldbring's solution [27] to the Hilbert's 5th problem for local groups, if Γ is a *locally Euclidean* local group (i.e. there is an open neighborhood of 1 homeomorphic to an open subset of $\mathbb{R}^n$, for some n), then there is a neighborhood U of 1 such that U is isomorphic, as a local group, to an open subset of an actual Lie group G. Clearly, if the local group is abelian then the connected component of G is also abelian. Combining these observations with Corollary 6.19 we conclude:

Corollary 6.20. *Let $\mathbb{R}_{\text{o-min}}$ be an o-minimal expansion of the field of reals. Assume $s \geq 3$, $Q \subseteq X_1 \times \cdots \times X_s$ are $\emptyset$-definable with $\dim(X_i) = m$, and Q is fiber-algebraic. Then one of the following holds.*

(1) The set Q satisfies γ-power saving, for $\gamma = \frac{1}{8m-5}$ if $s \geq 4$, and $\gamma = \frac{1}{16m-10}$ if $s = 3$.
(2) There exist definable relatively open sets $U_i \subseteq X_i$, $i \in [s]$, an abelian Lie group $(G, +)$ of dimension m and an open neighborhood $V \subseteq G$ of 0, and definable homeomorphisms $\pi_i : U_i \to V$, $i \in [s]$, such that for all $x_i \in U_i, i \in [s]$

$$\pi_1(x_1) + \cdots + \pi_s(x_s) = 0 \Leftrightarrow Q(x_1, \ldots, x_s).$$

Finally, this takes a particularly explicit form when $\dim(X_i) = 1$ for all $i \in [s]$.

Corollary 6.21. *Let $\mathbb{R}_{\text{o-min}}$ be an o-minimal expansion of the field of reals. Assume $s \geq 3$ and $Q \subseteq \mathbb{R}^s$ is definable and fiber-algebraic. Then exactly one of the following holds.*

(1) There exists a constant c, depending only on the formula defining Q (and not on its parameters), such that: for any finite $A_i \subseteq \mathbb{R}$ with $|A_i| = n$ for $i \in [s]$ we have

$$|Q \cap (A_1 \times \ldots \times A_s)| \leq cn^{s-1-\gamma},$$

where $\gamma = \frac{1}{3}$ if $s \geq 4$, and $\gamma = \frac{1}{6}$ if $s = 3$.

(2) There exist definable open sets $U_i \subseteq \mathbb{R}, i \in [s]$, an open set $V \subseteq \mathbb{R}$ containing 0, and homeomorphisms $\pi_i : U_i \to V$ such that

$$\pi_1(x_1) + \cdots + \pi_s(x_s) = 0 \Leftrightarrow Q(x_1, \ldots, x_s)$$

for all $x_i \in U_i, i \in [s]$.

Proof. Corollary 6.20 can be applied to Q.

Assume we are in Clause (1). As the proof of Theorem 6.4 demonstrates, we can take any γ such that Q satisfies the γ-ST property (as a binary relation, under any partition of its variables into two and the rest) if $s \geq 4$; and such that Q' (as defined in Section 5.7) satisfies the γ-ST property if $s = 3$. Applying the stronger bound for definable subsets of $\mathbb{R}^2 \times \mathbb{R}^{d_2}$ from Fact 2.15(1), we get the desired γ-power saving. Note that in the 1-dimensional case, the general position requirement is satisfied automatically: for any definable set $Y \subseteq \mathbb{R}$, $\dim(Y) < 1$ if and only if Y is finite; and for every definable family $\mathcal{F}_i$ of subsets of $\mathbb{R}$, by *o*-minimality there exists some ν_0 such that for any $Y \in \mathcal{F}_i$, if Y has cardinality greater than ν_0 then it is infinite.

In Clause (2), we use that every connected 1-dimensional Lie group G is isomorphic to either $(\mathbb{R}, +)$ or S^1, and in the latter case we can restrict to a neighborhood of 0 and compose the π_i's with a local isomorphism from S^1 to $(\mathbb{R}, +)$.

Finally, the two clauses are mutually exclusive as in Remark 5.52. □

Remark 6.22. In the case that definable sets in $\mathbb{R}_{\text{o-min}}$ admit *analytic cell decomposition* (e.g. in the *o*-minimal structure $\mathbb{R}_{\text{an,exp}}$, see [59, Section 8]) then one can strengthen Clause (2) in Corollaries 6.20 and 6.21, so that the U_i's are analytic submanifolds and the maps π_i are analytic bijections with analytic inverses.

Remark 6.23. If Q is semialgebraic (which corresponds to the case $\mathbb{R}_{\text{o-min}} = \mathbb{R}$ of Corollary 6.21), of description complexity D (i.e. defined by at most D polynomial (in-)equalities, with all polynomials of degree at most D), then in Clause (1) the constant c depends only on s and D (as all Q's are defined by the instances of a single formula depending only on s and D).

Remark 6.24. If Q is semilinear, then by Fact 2.19 it satisfies $(1 - \varepsilon)$-ST property, for any $\varepsilon > 0$. In this case, in Clause (1) of Corollary 6.21 for

$s \geq 4$ we get $(1-\varepsilon)$-power saving — which is essentially the best possible bound. See [39] concerning the lower bounds on power saving.

References

[1] Aaron Anderson, *Combinatorial bounds in distal structures*, Preprint, arXiv:2104.07769 (2021).

[2] Matthias Aschenbrenner, Artem Chernikov, Allen Gehret, and Martin Ziegler, *Distality in valued fields and related structures*, Trans. Amer. Math. Soc. **375** (2022), no. 7, 4641–4710.

[3] Salvador Barone, *Some quantitative results in real algebraic geometry*, ProQuest LLC, Ann Arbor, MI, 2013. Thesis (Ph.D.)–Purdue University.

[4] Abdul Basit, Artem Chernikov, Sergei Starchenko, Terence Tao, and Chieu-Minh Tran, *Zarankiewicz's problem for semilinear hypergraphs*, Forum Math. Sigma **9** (2021), Paper No. e59, 23.

[5] Saugata Basu, Antonio Lerario, and Abhiram Natarajan, *Zeroes of polynomials on definable hypersurfaces: pathologies exist, but they are rare*, Q. J. Math. **70** (2019), no. 4, 1397–1409.

[6] Saugata Basu and Orit E Raz, *An o-minimal Szemerédi–Trotter theorem*, The Quarterly Journal of Mathematics **69** (2017), no. 1, 223–239.

[7] Martin Bays, *Geometric stability theory*, Lectures in model theory, 2018, pp. 29–58.

[8] Martin Bays and Emmanuel Breuillard, *Projective geometries arising from Elekes-Szabó problems*, Ann. Sci. Éc. Norm. Supér. (4) **54** (2021), no. 3, 627–681.

[9] Martin Bays, Martin Hils, and Rahim Moosa, *Model theory of compact complex manifolds with an automorphism*, Transactions of the American Mathematical Society **369** (2017), no. 6, 4485–4516.

[10] Alexander Berenstein and Evgueni Vassiliev, *Geometric structures with a dense independent subset*, Selecta Mathematica **22** (2016), no. 1, 191–225.

[11] Béla Bollobás, *Extremal graph theory*, Dover Publications, Inc., Mineola, NY, 2004. Reprint of the 1978 original.

[12] Emmanuel Breuillard, Amador Martin-Pizarro, Katrin Tent, and Frank Olaf Wagner, *Model theory: Groups, geometries and combinatorics*, Oberwolfach Reports **17** (2021), no. 1, 91–142.

[13] Boris Bukh and Jacob Tsimerman, *Sum–product estimates for rational functions*, Proceedings of the London Mathematical Society **104** (2012), no. 1, 1–26.

[14] Bernard Chazelle, Herbert Edelsbrunner, Leonidas J Guibas, and Micha Sharir, *A singly exponential stratification scheme for real semi-algebraic varieties and its applications*, Theoretical Computer Science **84** (1991), no. 1, 77–105.

[15] Artem Chernikov, David Galvin, and Sergei Starchenko, *Cutting lemma and Zarankiewicz's problem in distal structures*, Selecta Mathematica **26** (2020), no. 2, 1–27.

[16] Artem Chernikov and Nadja Hempel, *On n-dependent groups and fields II*, Forum Math. Sigma **9** (2021), Paper No. e38, 51.

[17] Artem Chernikov and Martin Hils, *Valued difference fields and NTP$_2$*, Israel Journal of Mathematics **204** (2014), no. 1, 299–327.

[18] Artem Chernikov and Pierre Simon, *Externally definable sets and dependent pairs II*, Transactions of the American Mathematical Society **367** (2015), no. 7, 5217–5235.

[19] Artem Chernikov and Sergei Starchenko, *Regularity lemma for distal structures*, Journal of the European Mathematical Society **20** (2018), no. 10, 2437–2466.

[20] ———, *Model-theoretic Elekes-Szabó in the strongly minimal case*, J. Math. Log. **21** (2021), no. 2, Paper No. 2150004, 20.

[21] Bruno Dinis and Mário J Edmundo, *On definable Skolem functions and trichotomy*, Preprint, arXiv:2207.11339 (2022).

[22] György Elekes and Lajos Rónyai, *A combinatorial problem on polynomials and rational functions*, Journal of Combinatorial Theory, Series A **89** (2000), no. 1, 1–20.
[23] György Elekes and Endre Szabó, *How to find groups? (and how to use them in Erdős geometry?)*, Combinatorica (2012), 1–35.
[24] Paul Erdős and Endre Szemerédi, *On sums and products of integers*, Studies in pure mathematics, 1983, pp. 213–218.
[25] Jacob Fox, János Pach, Adam Sheffer, Andrew Suk, and Joshua Zahl, *A semi-algebraic version of Zarankiewicz's problem*, J. Eur. Math. Soc. (JEMS) **19** (2017), no. 6, 1785–1810.
[26] James Freitag, Wei Li, Thomas Scanlon, and William Johnson, *Differential chow varieties exist*, Journal of the London Mathematical Society **95** (2017), no. 1, 128–156.
[27] Isaac Goldbring, *Hilbert's fifth problem for local groups*, Annals of Mathematics (2010), 1269–1314.
[28] WT Gowers and Jason Long, *Partial associativity and rough approximate groups*, Geometric and Functional Analysis **30** (2020), no. 6, 1583–1647.
[29] Deirdre Haskell, Ehud Hrushovski, and Dugald Macpherson, *Definable sets in algebraically closed valued fields: elimination of imaginaries*, J. Reine Angew. Math. **597** (2006), 175–236.
[30] Ehud Hrushovski, *On pseudo-finite dimensions*, Notre Dame Journal of Formal Logic **54** (2013), no. 3-4, 463–495.
[31] ______, *Unpublished manuscript* (2014).
[32] ______, *Pseudofinite dimensions: Proper intersections and modularity*, Oberwolfach Reports **13** (2016), no. 1, 27–29.
[33] Yifan Jing, Souktik Roy, and Chieu-Minh Tran, *Semialgebraic methods and generalized sum-product phenomena*, Discrete Analysis **18** (2022), 23 pp.
[34] Itay Kaplan, Alf Onshuus, and Alexander Usvyatsov, *Additivity of the dp-rank*, Transactions of the American Mathematical Society **365** (2013), no. 11, 5783–5804.
[35] S. V. Konyagin and I. D. Shkredov, *On sum sets of sets having small product set*, Proc. Steklov Inst. Math. **290** (2015), no. 1, 288–299. Published in Russian in Tr. Mat. Inst. Steklova **2**90 (2015), 304–316.
[36] Tamás Kovári, Vera Sós, and Pál Turán, *On a problem of K. Zarankiewicz*, Colloquium mathematicae, 1954, pp. 50–57.
[37] Nathan Linial and Zur Luria, *An upper bound on the number of high-dimensional permutations*, Combinatorica **34** (2014), no. 4, 471–486.
[38] ______, *Discrepancy of high-dimensional permutations*, Discrete Analysis (2016), 845.
[39] Mehdi Makhul, Oliver Roche-Newton, Audie Warren, and Frank de Zeeuw, *Constructions for the Elekes–Szabó and Elekes–Rónyai problems*, The Electronic Journal of Combinatorics (2020), P1–57.
[40] David Marker, *Model theory: An introduction*, Graduate Texts in Mathematics, vol. 217, Springer-Verlag, New York, 2002.
[41] Rahim N. Moosa, *The model theory of compact complex spaces*, Logic colloquium '01, 2005, pp. 317—349.
[42] Joel Nagloo and Anand Pillay, *On algebraic relations between solutions of a generic Painlevé equation*, J. Reine Angew. Math. **726** (2017), 1–27.
[43] János Pach and Micha Sharir, *Repeated angles in the plane and related problems*, Journal of Combinatorial Theory, series A **59** (1992), no. 1, 12–22.
[44] Daniel Palacín, *An introduction to stability theory*, Lectures in model theory, 2018, pp. 1–27.
[45] Ya'acov Peterzil, *An o-minimalist view of the group configuration*, Preprint, arXiv:1909.09994 (2019).
[46] Anand Pillay, *Geometric stability theory*, Oxford Logic Guides, vol. 32, The Clarendon Press, Oxford University Press, New York, 1996. Oxford Science Publications.

[47] ———, *Model theory of algebraically closed fields*, Model theory and algebraic geometry, 1998, pp. 61–84.
[48] Orit E. Raz, Micha Sharir, and Frank De Zeeuw, *Polynomials vanishing on Cartesian products: the Elekes-Szabó theorem revisited*, Duke Math. J. **165** (2016), no. 18, 3517–3566.
[49] Orit E. Raz, Micha Sharir, and Frank de Zeeuw, *The Elekes-Szabó theorem in four dimensions*, Israel J. Math. **227** (2018), no. 2, 663–690.
[50] Orit E. Raz, Micha Sharir, and József Solymosi, *Polynomials vanishing on grids: the Elekes-Rónyai problem revisited*, Amer. J. Math. **138** (2016), no. 4, 1029–1065.
[51] Orit E. Raz and Zvi Shem-Tov, *Expanding polynomials: a generalization of the Elekes-Rónyai theorem to d variables*, Combinatorica **40** (2020), no. 5, 721–748.
[52] Thomas Scanlon, *O-minimality as an approach to the André-Oort conjecture*, Around the Zilber-Pink conjecture/Autour de la conjecture de Zilber-Pink, 2017, pp. 111–165.
[53] Adam Sheffer, *Polynomial methods and incidence theory*, Cambridge Studies in Advanced Mathematics, vol. 197, Cambridge University Press, Cambridge, 2022.
[54] Pierre Simon, *Distal and non-distal NIP theories*, Annals of Pure and Applied Logic **164** (2013), no. 3, 294–318.
[55] József Solymosi, *Bounding multiplicative energy by the sumset*, Adv. Math. **222** (2009), no. 2, 402–408.
[56] Terence Tao, *Expanding polynomials over finite fields of large characteristic, and a regularity lemma for definable sets*, Contributions to Discrete Mathematics **10** (2015), no. 1.
[57] Katrin Tent and Martin Ziegler, *A course in model theory*, Lecture Notes in Logic, vol. 40, Association for Symbolic Logic, La Jolla, CA; Cambridge University Press, Cambridge, 2012.
[58] Lou van den Dries, *Tame topology and o-minimal structures*, London Mathematical Society Lecture Note Series, vol. 248, Cambridge University Press, Cambridge, 1998.
[59] Lou van den Dries and Chris Miller, *On the real exponential field with restricted analytic functions*, Israel Journal of Mathematics **85** (1994), no. 1-3, 19–56.
[60] Miguel N. Walsh, *The polynomial method over varieties*, Invent. Math. **222** (2020), no. 2, 469–512.