ON LINEAR CONFIGURATIONS IN SUBSETS OF COMPACT ABELIAN GROUPS, AND INVARIANT MEASURABLE HYPERGRAPHS

PABLO CANDELA, BALÁZS SZEGEDY, AND LLUÍS VENA

ABSTRACT. We prove an arithmetic removal result for all compact abelian groups, generalizing a finitary removal result of Král', Serra and the third author. To this end, we consider infinite measurable hypergraphs that are invariant under certain group actions, and for these hypergraphs we prove a symmetry-preserving removal lemma, which extends a finitary result of the same name by the second author. We deduce our arithmetic removal result by applying this lemma to a specific type of invariant measurable hypergraph. As a direct application, we obtain the following generalization of Szemerédi's theorem: for any compact abelian group G, any measurable set $A \subseteq G$ with Haar probability $\mu(A) \geq \alpha > 0$ satisfies

$$\int_G \int_G 1_A(x) \, 1_A(x+r) \cdots 1_A(x+(k-1)r) \, \mathrm{d}\mu(x) \, \mathrm{d}\mu(r) \geq c,$$

where the constant $c = c(\alpha, k) > 0$ is valid uniformly for all G. This result is shown to hold more generally for any translation-invariant system of r linear equations given by an integer matrix with coprime $r \times r$ minors.

1. INTRODUCTION

This paper concerns the general question of the extent to which linear configurations of a given type must occur in subsets of abelian groups. Given a matrix $M \in \mathbb{Z}^{r \times m}$, and a subset A of an abelian group G, we consider the set of elements $x \in A^m$ solving the system $Mx = 0$, that is the set $A^m \cap \ker_G M$. In relation to the above question, it is a well-known fruitful approach to examine what can be deduced about A if the set $A^m \cap \ker_G M$ occupies a small proportion of the total set of configurations $\ker_G M$. In this direction, useful information is provided by what are often called arithmetic removal results. The following example treats the case of simple abelian groups $G = \mathbb{Z}_p$.

Theorem 1.1. *Let m, r be positive integers, with $m \geq r$. Then for any $\epsilon > 0$ there exists $\delta > 0$ such that the following holds. Let M be a matrix of rank r in $\mathbb{Z}^{r \times m}$ and suppose that $A_1, A_2, \ldots, A_m$ are subsets of $\mathbb{Z}_p$ such that $|A_1 \times A_2 \times \cdots \times A_m \cap \ker_{\mathbb{Z}_p} M| \leq \delta |\ker_{\mathbb{Z}_p} M|$. Then there exist $R_1 \subseteq A_1, \ldots, R_m \subseteq A_m$ such that $|R_j| \leq \epsilon p$ for every $j \in [m]$, and $\left(\prod_{j \in [m]} A_j \setminus R_j \right) \cap \ker_{\mathbb{Z}_p} M = \emptyset$.*

As a consequence, if $|A^m \cap \ker_{\mathbb{Z}_p} M| \leq \delta |\ker_{\mathbb{Z}_p} M|$, then it is possible to eliminate all these solutions in A^m by removing at most ϵp elements from A. Thus A must be of the form $B \cup R$, where $|R| \leq \epsilon p$ and B is what we call an *M-free set*, that is it satisfies $B^m \cap \ker_G M = \emptyset$.

Theorem 1.1 was proved by Shapira [27] and independently by Král', Serra and the third author [17]. (Strictly speaking, the result was proved more generally for finite fields.) This result confirmed a conjecture of Green from [10]. In that paper, Green introduced the notion of such removal results as arithmetic counterparts of well-known combinatorial removal results from graph theory, and he proved a version of Theorem 1.1 for a single linear equation on an arbitrary finite abelian group. For more background on the relation between arithmetic and combinatorial removal results, the reader is referred to the survey [6], especially Section 4 therein.

One of the central consequences of Theorem 1.1 is a general form of Szemerédi's famous theorem on arithmetic progressions [28], Theorem 1.2 below. To state the result, we use the following terminology. We say that a matrix $M \in \mathbb{Z}^{r \times m}$ is *invariant* if its columns sum to zero, that is if $M(1, 1, \ldots, 1)^T = 0$; equivalently, for any abelian group G, the set $\ker_G M$ is invariant under translations by constant elements $(t, t, \ldots, t)$, $t \in G$. Examples of configurations given by invariant matrices include arithmetic progressions of an arbitrary fixed length.

Theorem 1.2. *Let m, r be positive integers, with $m \geq r$. For any $\alpha > 0$ there exists $c = c(\alpha, m) > 0$ such that the following holds. Let M be an invariant matrix of rank r in $\mathbb{Z}^{r \times m}$, and let A be a subset of $\mathbb{Z}_p$ of cardinality at least αp. Then we have*

$$|A^m \cap \ker_{\mathbb{Z}_p} M| \, / \, |\ker_{\mathbb{Z}_p} M| \geq c.$$

In particular, for any positive integer k, the set A must contain a positive proportion $c(\alpha, k)$ of the total number p^2 of k-term progressions in $\mathbb{Z}_p$. The deduction of Theorem 1.2 from Theorem 1.1 is very short, we record a proof in a more general context at the end of Section 3.

If M is not invariant, then the conclusion of Theorem 1.2 fails, in that there exists $\alpha = \alpha(M) > 0$ such that in any group $\mathbb{Z}_p$ there is an M-free set of size at least αp. This can be shown using a simple adaptation of the argument from [22, Theorem 2.1].

Thus for $G = \mathbb{Z}_p$, as a direct consequence of Theorem 1.1, the question recalled at the beginning of this introduction receives a strong answer (Theorem 1.2) which is also exhaustive as far as systems of linear equations are concerned.[1] It is natural to wonder whether this picture holds for more general abelian groups.

Given $M \in \mathbb{Z}^{r \times m}$ of rank r, let us denote by $d_r(M)$ the *determinantal divisor* of M of order r, that is the greatest common divisor of the non-zero determinants of $r \times r$ submatrices of M; see [19, Chapter II, §13]. We shall not consider determinantal divisors of lower order, and will therefore refer to $d_r(M)$ simply as 'the determinantal' of M.

Under the assumption that $d_r(M) = 1$, Král', Serra and the third author generalized Theorem 1.1 to all finite abelian groups, obtaining[2] [18, Theorem 1]. This extension has found several applications. In particular it immediately implies a corresponding extension of Szemerédi's theorem to all finite abelian groups, since a matrix characterizing

---

[1]The answer is strong in a qualitative sense. The quantitative problem of obtaining optimal estimates for the function $c(\alpha, M)$ in Theorem 1.2 is a vast and very interesting one, that includes improving the bounds for Szemerédi's theorem. For the latter theorem the current best general bounds were given in [8]; see also [1, 12, 24] for the latest improvements in the cases $k = 3, 4$.

[2]Theorem 1 in [18] actually assumes that $\gcd(d_r(M), |G|) = 1$, a weaker assumption than $d_r(M) = 1$. However, the theorem itself holds equivalently for each of these two assumptions; see Remark 4.4.

arithmetic progressions of a fixed length satisfies the above assumption; other applications include those in [25, Section 10] and [26]. Assuming that $d_r(M) = 1$ is a simple way to ensure that the set of solutions has the 'expected dimension'; more precisely, we then have $\ker_G M \cong G^{m-r}$, as can be seen using the Smith normal form of M (see [19, Theorem II.9]). We shall say more about this assumption in Section 5 below.

Some recent works have made use of removal results in the setting of infinite compact abelian groups. For instance, in [5] it was shown that Theorem 1.1 implies an analogous result for the circle group $G = \mathbb{R}/\mathbb{Z}$, formulated in terms of Haar measure, which was found to be useful for certain additive-combinatorial questions studied in $\mathbb{Z}_p$ as $p \to \infty$; see also [4]. At the end of [5], the possibility of a removal result for a general compact abelian group was raised.

The main result of this paper is an extension of Theorem 1.1, for matrices of determinantal 1, to all compact abelian groups. Below we discuss further motivation for this extension, but before that let us state the result formally.

All topological groups in this paper are assumed to be Hausdorff. Any compact group G admits a unique Haar probability measure, which we denote by μ_G. A subset of G is said to be *Haar measurable* (or just *measurable*) if it is in the completion of the Borel σ-algebra on G relative to μ_G. Given a compact abelian group G and a matrix $M \in \mathbb{Z}^{r \times m}$, the kernel $\ker_G M$ of the continuous homomorphism $M : G^m \to G^r$ is a compact subgroup of G^m, with its own Haar probability $\mu_{\ker_G M}$. For a measurable set $A \subseteq G$, the quantity $\mu_{\ker_G M}(A^m \cap \ker_G M)$ gives the natural notion of the proportion (or density) of solutions contained in A^m. This makes the setting of compact abelian groups a very natural one in which to seek general versions of results such as Theorem 1.2 (note that if G is finite then $\mu_{\ker_G M}(A^m \cap \ker_G M)$ is just $|A^m \cap \ker_G M|/|\ker_G M|$). For more background on the Haar measure, we refer the reader to [7, 13, 21].

We can now state our main result.

Theorem 1.3. *Let $M \in \mathbb{Z}^{r \times m}$ satisfy $d_r(M) = 1$. For any $\epsilon > 0$, there exists $\delta = \delta(\epsilon, M) > 0$ such that the following holds. If $A_1, A_2, \ldots, A_m$ are Borel subsets of a compact Hausdorff abelian group G such that $\mu_{\ker_G M}\big(A_1 \times \cdots \times A_m \cap \ker_G M\big) \leq \delta$, then there exist Borel sets $R_1 \subseteq A_1, \ldots, R_m \subseteq A_m$ such that $\mu_G(R_j) \leq \epsilon$ for all $j \in [m]$ and $\big(\prod_{j \in [m]} A_j \setminus R_j\big) \cap \ker_G M = \emptyset$.*

We shall deduce this result from a more precise version, which holds for second countable compact abelian groups, and which gives additional information on the location of the sets R_j and on their measure; see Theorem 3.1. Note that Theorem 1.3 also implies the inhomogeneous version of itself, where instead of $\ker_G M$ we consider the set of solutions $x \in G^m$ to $Mx = b$ for some non-zero $b \in G^r$.

From Theorem 1.3, one deduces directly the following generalization of Szemerédi's theorem (for a proof see the end of Section 3).

Theorem 1.4. *Let $M \in \mathbb{Z}^{r \times m}$ be invariant and satisfy $d_r(M) = 1$. Then for any $\alpha > 0$ there exists $c = c(\alpha, M) > 0$ such that if A is a measurable subset of a compact abelian group G with $\mu_G(A) \geq \alpha$, then $\mu_{\ker_G M}(A^m \cap \ker_G M) \geq c$.*

In particular, for any positive integer k, any measurable set $A \subseteq G$ with $\mu_G(A) \geq \alpha > 0$ satisfies[3]

$$\int_G \int_G 1_A(x)\, 1_A(x+r)\, \cdots\, 1_A(x+(k-1)r)\, \mathrm{d}\mu_G(x)\, \mathrm{d}\mu_G(r) \geq c,$$

where the positive lower bound $c = c(\alpha, k)$ is independent of the particular structure of A and is in fact valid uniformly for all G.

In addition to the generality of Theorem 1.3, this extension to compact abelian groups offered us the motivation that it does not seem to follow from the known finite results by a simple measure-theoretic argument. Significant additive-combinatorial aspects had to be taken into account, requiring in particular further understanding of the relationship between combinatorial removal results for hypergraphs and their arithmetic counterparts. Let us complete this introduction by detailing these points.

In order to prove a removal result in an infinite compact abelian group, it is natural to try to deduce it from a finitary version by a discretization argument. An approach of this type was taken in [5], yielding the above-mentioned analogue of Theorem 1.1 for the circle group. However, as noted at the end of that paper, for more general compact abelian groups this approach yields a version of Theorem 1.3 with a parameter δ depending on the topological dimension of the group. By contrast, the function δ in Theorem 1.3 is independent of the compact abelian group. To obtain this, the approach in this paper consists instead in finding infinite analogues of some elements from known proofs of finite removal results, and combining those with some new elements in the infinite setting.

Most of the known proofs in the finite setting proceed by reducing the arithmetic removal result somehow to its combinatorial counterpart for uniform hypergraphs, a method which first appeared explicitly, using graph removal lemmas, in [16].

The most elaborate form of this method so far, i.e. the proof of [18, Theorem 1], is implemented in a way that makes important use of properties specific to finite abelian groups, in particular the fact that multiplication by an integer does not increase the measure of a set in such a group (these aspects are discussed in more detail in Section 4 below). This prevents a simple transfer of the whole argument from [18] to the infinite setting, although several tools from that argument do transfer and are used in this paper.

The above-mentioned method is implemented in another way in the approach to arithmetic removal results given in [29]. The main result of that paper is a so-called *symmetry-preserving* version of the removal lemma for finite hypergraphs. This version has the additional information that if the edge sets of the given hypergraph were invariant under a certain group action, then the edge sets to be removed can be guaranteed also to be invariant. This version of the hypergraph removal lemma turns out to have a useful extension to the infinite setting, which we prove in this paper; see Lemma 2.12. This extension concerns hypergraphs defined on general probability spaces and acted upon in a certain way by a compact group; see Definitions 2.8 and 2.10. This infinite symmetry-preserving removal lemma gives a convenient footing for a proof of Theorem 3.1. However, completing the proof requires finding how to associate such an invariant hypergraph with a given system of linear equations on a compact abelian group. Indeed, in [29] the finite

---

[3]The case $k = 3$ of this result, namely Roth's theorem for a general compact abelian group, can be treated using Fourier analysis; see for instance [31].

symmetry-preserving removal lemma was shown to yield finite arithmetic removal results, but this was demonstrated only for certain examples of linear configurations, and it was not clear how to handle more general systems. In this paper, to clarify this we define a notion of a *hypergraph representation* of a system of linear equations on an abelian group. This notion extends and unifies previous finitary notions of a similar kind [3, 17, 27], and it is designed to go together with the symmetry-preserving removal lemma; see Definition 3.7. More precisely, this representation is a homomorphism which enables us to associate a certain measurable invariant hypergraph to the given system of equations, in such a way that the desired arithmetic removal result can be deduced from the removal lemma for this hypergraph; see Definition 3.2.

In Section 2, we prove the symmetry-preserving removal lemma. In Section 3, we define the hypergraph representation and use it to deduce the arithmetic removal result as mentioned above. In Section 4 we show that for any matrix $M \in \mathbb{Z}^{r \times m}$ with $d_r(M) = 1$ and any compact abelian group, there exists such a hypergraph representation. In Section 5 we end with some remarks on potential further extensions of Theorem 1.3.

2. A SYMMETRY-PRESERVING REMOVAL LEMMA FOR MEASURABLE HYPERGRAPHS

In this section we establish the main result that we shall use concerning measurable hypergraphs, namely the symmetry-preserving removal lemma (Lemma 2.12). This generalizes [29, Theorem 2]. Let us set up some terminology and notation.

Let $[t] = \{1, 2, \ldots, t\}$, and let us denote the set of subsets of $[t]$ of size k by $\binom{[t]}{k}$. Given any cartesian product $\prod_{i \in [t]} V_i$, and any set $e \subseteq [t]$, we denote by p_e the projection $\prod_{i \in [t]} V_i \to \prod_{i \in e} V_i$ to the components indexed by e, thus $p_e(v) = (v(i))_{i \in e}$. (If e is a singleton $\{i\}$ we write p_i rather than $p_{\{i\}}$.) When there is no danger of confusion, we shall often use the notation V_e to refer to the product $\prod_{i \in e} V_i$.

The kind of hypergraph that we consider is the following.

Definition 2.1. A *t-partite m-colored k-uniform hypergraph*, or (t, m, k)-graph for short, is a triple (V, C, E) consisting of the following elements. The *vertex set* V is the disjoint union of labelled sets $V_1, V_2, \ldots, V_t$. The set C of *edge color-classes* is a collection of m distinct labelled sets $C_1, \ldots, C_m \in \binom{[t]}{k}$. The *edge set* E is the union of sets $E_1, \ldots, E_m$ where each E_j is a subset of $\prod_{i \in C_j} V_i$, the elements of which are the *edges of color j*.

We say that a (t, m, k)-graph is *measurable* if there is a probability space structure $(V_i, \mathcal{V}_i, \mu_i)$ on each vertex set V_i (here $\mathcal{V}_i$ denotes a σ-algebra of subsets of V_i, and μ_i a probability on $\mathcal{V}_i$), and every set E_j is in the product σ-algebra $\prod_{i \in C_j} \mathcal{V}_i$. All the (t, m, k)-graphs that we consider in this paper are assumed to be measurable.

Given probability spaces $(V_i, \mathcal{V}_i, \mu_i)$, $i \in [t]$, for any $e \subseteq [t]$ of size $|e| > 1$ we shall denote by $(V_e, \mathcal{V}_e, \mu_e)$ the product probability space $(\prod_{i \in e} V_i, \prod_{i \in e} \mathcal{V}_i, \prod_{i \in e} \mu_i)$.

Definition 2.2 ((t, m, k)-graph homomorphism)**.** Let H_1 be a (t, m, k)-graph with vertex set $U = \bigsqcup_i U_i$, and let H_2 be a (t, m, k)-graph with vertex sets $V = \bigsqcup_i V_i$. A *homomorphism* from H_1 to H_2 is a map $\phi : U \to V$ defined by $\phi(u) = \phi_i(u)$ for $u \in U_i$, where $(\phi_i)_{i \in [t]}$ is a t-tuple of measurable maps $\phi_i : U_i \to V_i$ with the following property: if $(u_i)_{i \in C_j}$ is an edge of H_1, then the image $(\phi_i(u_i))_{i \in C_j}$ is an edge of H_2.

We say that H_2 is H_1-*free* if there is no injective homomorphism $\phi : H_1 \to H_2$. A measurable (t, m, k)-graph is *finite* if the vertex sets V_i are finite and the probabilities μ_i are uniform. In this paper we will only use homomorphisms from a finite (t, m, k)-graph to a possibly infinite (t, m, k)-graph. It is helpful to view these homomorphisms as points in the space $V_1^{U_1} \times V_2^{U_2} \times \cdots \times V_t^{U_t}$. Indeed, this leads naturally to the following definition of the homomorphism density, using the product probability on this space.

Definition 2.3. Let F be a finite (t, m, k)-graph with vertex sets U_i, and let H be a (t, m, k)-graph with vertex sets V_i. The *homomorphism density* of F in H, denoted $\tau(F, H)$, is the probability that for a random t-tuple of maps $(\phi_i : U_i \to V_i)_{i \in [t]}$ the corresponding map ϕ is a homomorphism.

In particular, if H has color-classes $C_1, \ldots, C_m$ and F is the finite hypergraph with vertex set $[t]$ and edges $C_1, \ldots, C_m$, then, recalling that $(V_{[t]}, \mathcal{V}_{[t]}, \mu_{[t]})$ denotes the product of the probability spaces $(V_i, \mathcal{V}_i, \mu_i)$, we have

$$\tau(F, H) = \int_{V_{[t]}} \prod_{j \in [m]} 1_{E_j}\big(p_{C_j}(v)\big) \, \mathrm{d}\mu_{[t]}(v). \tag{1}$$

For reasons that will become clear in the following sections, in this paper we only need this type of homomorphism $\phi : F \to H$ where each vertex class of F is a singleton $U_i = \{i\}$. Note that any such homomorphism is an injective map, since the vertex classes of H are disjoint by definition. We may sometimes refer to the image $\phi(F) = (\phi(i))_{i \in [t]}$ as a *copy* of F in H. In the general case, where F may have more than one vertex per class, there is a similar but more complicated version of formula (1), but as mentioned above we shall not use this.

In the next subsection we shall obtain a removal lemma for (t, m, k)-graphs, Lemma 2.4, by deducing it from the well-known removal lemma for finite hypergraphs. We shall then add the symmetry-preserving property in subsection 2.2, obtaining the main result of this section, Lemma 2.12.

2.1. **A removal lemma for (t, m, k)-graphs.** In this subsection we establish the following result.

Lemma 2.4. *Let $t \geq k \geq 2$ and m be positive integers, and let $0 < \epsilon < 1$. There exists $\delta = \delta(t, k, \epsilon) > 0$ such that the following holds. Let H be a (t, m, k)-graph with vertex sets V_i, $i \in [t]$, and edge color-classes C_j, $j \in [m]$, let F be the (t, m, k)-graph with vertex set $[t]$ and edges C_j, and suppose that $\tau(F, H) \leq \delta$. Then for each $j \in [m]$ there exists a measurable set $R_j \subseteq E_j(H)$ with $\mu_{C_j}(R_j) \leq \epsilon$, such that removing each R_j from $E_j(H)$ yields an F-free (t, m, k)-graph.*

The finite version of this result, that is the special case in which both F and H are finite (t, m, k)-graphs, is a version of the well-known hypergraph removal lemma, given for instance in [30]. Our task here is to show that the above version for arbitrary probability spaces follows from the finite version. To prove this we use a discretization argument whereby H is approximated by a (t, m, k)-graph $H^{(1)}$ whose vertex sets are partitioned into finitely many parts, and whose edge sets are disjoint unions of products of some of these parts. Then, we model each of these parts by a finite set of vertices, the cardinality of which is chosen according to the measure of the part. This enables

us to relate $\tau(F, H)$ with $\tau(F, H^{(2)})$ for some associated finite (t, m, k)-graph $H^{(2)}$, thus reducing the proof to an application of the finite version of Lemma 2.4.

Proof of Lemma 2.4. Let $\delta' \leq \epsilon/(4m)$ be such that the finite version of Lemma 2.4 holds with parameters $\epsilon/(4m), t, k$. (As mentioned above, this finite version is known; indeed it is essentially [30, Corollary 1.14].) Suppose that $\tau(F, H) \leq \delta$ with $\delta = \delta'/2$.

For each $j \in [m]$, since the σ-algebra $\mathcal{V}_{C_j}$ on $V_{C_j} = \prod_{i \in C_j} V_i$ is generated by products of measurable subsets of the components V_i, there exist disjoint sets $B_{j,1}, B_{j,2}, \ldots, B_{j,M_j}$, each of the form $B_{j,r} = \prod_{i \in C_j} D_{i,j,r}$ with $D_{i,j,r} \in \mathcal{V}_i$, satisfying

$$\mu_{C_j}\left(E_j(H) \, \triangle \, \bigsqcup_{r=1}^{M_j} B_{j,r}\right) \leq \delta/m \leq \epsilon/2. \tag{2}$$

Let $H^{(1)}$ be the (t, m, k)-graph obtained from H by replacing the edge sets $E_j(H)$ with $E_j^{(1)} := \bigsqcup_{r=1}^{M_j} B_{j,r}$. By (2) and a simple telescoping argument using multilinearity of the function $(1_{E_1}, \ldots, 1_{E_m}) \mapsto \prod_{j \in [m]} 1_{E_j} \circ p_{C_j}$, we have[4]

$$\tau(F, H^{(1)}) \leq \tau(F, H) + m\frac{\delta}{m} \leq \delta'.$$

We shall now show that $H^{(1)}$ can be made F-free by removing a set of measure at most $\epsilon/2$ from each set $E_j^{(1)}$.

For each $i \in [t]$ we define a partition of V_i generated by all the sets $D_{i,j,r}$. More precisely, let $\mathcal{P}_i$ denote the partition of V_i into the atoms of the finite σ-algebra generated by the collection of sets $\bigcup_{j \in [m]: C_j \ni i}\{D_{i,j,r} : r \in [M_j]\}$. Let $K_i = |\mathcal{P}_i|$, thus $\mathcal{P}_i = \{P_{i,1}, P_{i,2}, \ldots, P_{i,K_i}\}$. Each set $E_j^{(1)}$ is a disjoint union of sets of the form $\prod_{i \in C_j} P_{i,\ell_i}$ for some $\ell = (\ell_i)_{i \in C_j} \in \prod_{i \in C_j}[K_i]$. Thus $H^{(1)}$ can already be viewed as a finite hypergraph, with vertex sets $\mathcal{P}_1, \ldots, \mathcal{P}_t$ and edges these k-tuples ℓ. However, the measures of the atoms $P_{i,j}$ are not necessarily equal, so the probabilities on the vertex sets of this hypergraph may fail to be uniform. In order to apply the finite version of the removal lemma, we shall now approximate this weighted hypergraph by a finite (t, m, k)-graph $H^{(2)}$.

Note that if $v = (v(1), \ldots, v(t))$ is a copy of F in $H^{(1)}$, with $v(i) \in P_{i,r_i} \subseteq V_i$ for each $i \in [t]$, then in fact every point in $P_{1,r_1} \times \cdots \times P_{t,r_t}$ is such a copy, and this product set gives us a measure $\mu_1(P_{1,r_1}) \cdots \mu_t(P_{t,r_t})$ of homomorphisms $F \to H^{(1)}$.

Let N be a large positive integer to be determined below, depending on k, t, ϵ and the measure of the atoms P_{i,r_i}.

Let $H^{(2)}$ be the finite (t, m, k)-graph defined as follows. The finite vertex sets, denoted $V_1', \ldots, V_t'$, are each of cardinality N, with uniform probability denoted μ_i'. Each set V_i' is partitioned into sets $Q_{i,0}, Q_{i,1}, \ldots, Q_{i,K_i}$, such that we have

$$\forall r \in [K_i], \; |Q_{i,r}| = q_{i,r}, \; q_{i,r}/N \leq \mu_i(P_{i,r}) < (q_{i,r}+1)/N, \; \text{and} \; |Q_{i,0}| = q_{i,0} \leq K_i. \tag{3}$$

The edge color classes of $H^{(2)}$ are the same as for $H^{(1)}$, and for each such class C_j the edge set $E_j^{(2)}$ of $H^{(2)}$ is defined as follows. A k-tuple $(v(i))_{i \in C_j} \in V_{C_j}'$ is an edge in $E_j^{(2)}$

---

[4]We illustrate the argument for $m = 3$: for any functions $f_j, g_j : V_{[t]} \to \mathbb{R}$, $j \in [3]$, we have $f_1 f_2 f_3 = (f_1 - g_1)f_2 f_3 + g_1 f_2 f_3 = (f_1 - g_1)f_2 f_3 + g_1(f_2 - g_2)f_3 + g_1 g_2 f_3 = (f_1 - g_1)f_2 f_3 + g_1(f_2 - g_2)f_3 + g_1 g_2(f_3 - g_3) + g_1 g_2 g_3$; we then apply this with $f_j = 1_{E_j^{(1)}} \circ p_{C_j}$ and $g_j = 1_{E_j} \circ p_{C_j}$.

if and only if $\prod_{i \in C_j} P_{i,r_i} \subseteq E_j^{(1)}$, where $v(i) \in Q_{i,r_i}$ for each $i \in C_j$. In other words, $E_j^{(2)}$ is the disjoint union of all the sets $\prod_{i \in C_j} Q_{i,r_i}$ satisfying $\prod_{i \in C_j} P_{i,r_i} \subseteq E_j^{(1)}$.

Since each set $Q_{i,r}$ satisfies $\mu_i'(Q_{i,r}) \le \mu_i(P_{i,r})$, we have $\tau(F, H^{(2)}) \le \tau(F, H^{(1)}) \le \delta'$. By the finite version of the removal lemma, there exist sets $R_j' \subseteq E_j^{(2)}$ with $\mu_{C_j}'(R_j') \le \epsilon/(4m)$ such that, removing each R_j' from $E_j^{(2)}$, the resulting hypergraph is F-free.

Let us now use the sets R_j' to specify which subsets to remove from $E_j^{(1)}$. To do so, we first show that each R_j' may be replaced with a set R_j'' that is a union of sets of the form $\prod_{i \in C_j} Q_{i,r_i}$, in such a way that the sets R_j'' still have small measure and preserve the removal property.

Let R_j'' be the union of sets $\prod_{i \in C_j} Q_{i,r_i}$ such that

$$\Big| R_j' \cap \prod_{i \in C_j} Q_{i,r_i} \Big| \ge m^{-1} \prod_{i \in C_j} q_{i,r_i}.$$

We have $|R_j''| \le m |R_j'|$, and so $\mu_{C_j}'(R_j'') \le m \, \mu_{C_j}'(R_j') \le \epsilon/4$.

We claim that removing R_j'' (instead of R_j') from $E_j^{(2)}$ still yields an F-free (t, m, k)-graph. Indeed, suppose that $v_0 \in \prod_{i \in [t]} Q_{i,r_i}$ is a copy of F in $H^{(2)}$. Then, by the definition of $H^{(2)}$, every element $v \in \prod_{i \in [t]} Q_{i,r_i}$ is such a copy. By the removal property of the sets R_j', for any such v there exists $j \in [m]$ such that the edge $p_{C_j}(v)$ lies in R_j'. There must therefore exist $j \in [m]$ such that there are at least $m^{-1} \prod_{i \in [t]} q_{i,r_i}$ such copies v with $p_{C_j}(v) \in R_j'$. On the other hand, an edge $w \in \prod_{i \in C_j} Q_{i,r_i}$ can satisfy $w = p_{C_j}(v)$ for at most $\prod_{i \in [t] \setminus C_j} q_{i,r_i}$ of these copies v. We therefore conclude that $\Big| R_j' \cap \prod_{i \in C_j} Q_{i,r_i} \Big| \ge m^{-1} \prod_{i \in C_j} q_{i,r_i}$. Hence all these copies (including v_0) have $p_{C_j}(v) \in R_j''$ and are therefore eliminated by removing R_j''. This proves our claim.

We can now specify the sets $R_j^{(1)}$ that we remove from $E_j^{(1)}$. Let $R_j^{(1)}$ be the union of sets $\prod_{i \in C_j} P_{i,r_i} \subseteq E_j^{(1)}$ such that $\prod_{i \in C_j} Q_{i,r_i} \subseteq R_j''$. Note that

$$\mu_{C_j}\Big(\prod_{i \in C_j} P_{i,r_i} \Big) \le \prod_{i \in C_j} \Big(\mu_i'(Q_{i,r_i}) + \frac{1}{N} \Big) \le \mu_{C_j}'\Big(\prod_{i \in C_j} Q_{i,r_i} \Big) + \frac{2^k}{N}. \qquad (4)$$

Choosing $N > \frac{2 \cdot 2^k}{\epsilon} \max_{j \in [m]} \Big(\prod_{i \in C_j} K_i \Big)$, we deduce from (4) that

$$\mu_{C_j}\Big(R_j^{(1)} \Big) \le \mu_{C_j}'\Big(R_j'' \Big) + \Big(\prod_{i \in C_j} K_i \Big) \frac{2^k}{N} \le \frac{\epsilon}{2}, \quad \text{for each } j \in [m]. \qquad (5)$$

If there was a copy v left in $\bigcap_{j \in [m]} p_{C_j}^{-1} \Big(E^{(1)} \setminus R_j^{(1)} \Big)$, then there would have to be in fact a measure $\mu_{[t]} \Big(\prod_{i \in [t]} P_{i,r_i} \Big)$ of such copies, where $v(i) \in P_{i,r_i}$ for each $i \in [t]$. Therefore, by an analogue of (4), there would be a measure at least $\mu_{[t]} \Big(\prod_{i \in [t]} P_{i,r_i} \Big) - \frac{2^t}{N}$ of copies

of F in $\bigcap_{j\in[m]} p_{C_j}^{-1}\left(E_j^{(2)} \setminus R_j^{(2)}\right)$. If

$$N > 2^t / \min\left\{\mu_{[t]}\left(\prod_{i\in[t]} P_{i,r_i}\right) : (r_i) \in \prod_{i\in C_j} [K_i], \ j \in [m]\right\},$$

then there is at least one such copy of F, contradicting the removal property of the sets $R_j^{(2)}$.

We now set $R_j = R_j^{(1)} \cup \left(E_j(H) \setminus \bigsqcup_{\ell=1}^{M_j} B_{j,\ell}\right)$, which by (2) and (5) has measure at most ϵ for each $j \in [m]$, and the proof is complete. $\qquad\square$

Remark 2.5. Lemma 2.4 concerns the so-called 'partite hypergraph version' of the removal lemma (as it is called in [30]), which corresponds to the case of formula (1) in which F has one vertex per class. This case suffices for our purposes in this paper, as we shall see in the next sections. Let us mention that there is a version of Lemma 2.4 where F may have more than one vertex in each part U_i, and that in fact this extension can be deduced using Lemma 2.4.

2.2. **Preserving symmetries.** We now move on to the main result of this section, Lemma 2.12. This is a version of Lemma 2.4 which preserves certain symmetries of the given hypergraph. The symmetries of (t,m,k)-graphs that we shall consider are described in terms of a type of group action on the product of the vertex sets, that we call a *t-partite action* (see Definition 2.8). To build up to this notion, we first recall the definition of a measurable group action (see for instance [32, §3]). We denote the identity element of a group G by id_G.

Definition 2.6 (Group action on a probability space)**.** Let $(V, \mathcal{V}, \mu)$ be a probability space, and let $\mathcal{G}$ be a group. An *action* of $\mathcal{G}$ on V is a map $\Phi : \mathcal{G} \times V \to V$ satisfying the following properties:

(i) $\forall\, v \in V$, $\forall\, g, h \in \mathcal{G}$ we have $\Phi(gh, v) = \Phi(g, \Phi(h, v))$, and $\Phi(\mathrm{id}_\mathcal{G}, v) = v$.
(ii) For each $g \in \mathcal{G}$ the invertible map $\Phi_g : v \mapsto \Phi(g, v)$ is measurable and preserves μ, that is for any set $A \in \mathcal{V}$, we have $\Phi_g^{-1}(A) \in \mathcal{V}$ and $\mu(\Phi_g^{-1}(A)) = \mu(A)$.

In other words, the map $g \mapsto \Phi_g$ is a homomorphism from $\mathcal{G}$ into the group of measure-preserving automorphisms of V. If $\mathcal{G}$ is a topological group, with Borel σ-algebra denoted $\mathcal{B}_\mathcal{G}$, then we say that the action Φ is *measurable* if the map Φ is measurable from $(\mathcal{G} \times V, \mathcal{B}_\mathcal{G} \times \mathcal{V})$ to $(V, \mathcal{V})$.

We shall often use the simpler notation $g \cdot v$ for $\Phi(g, v)$.

Given an action of $\mathcal{G}$ on $(V, \mathcal{V}, \mu)$, a set $B \in \mathcal{V}$ is said to be $\mathcal{G}$-*invariant* if $g \cdot B = B$ for all $g \in \mathcal{G}$. These sets form a sub-σ-algebra of $\mathcal{V}$ that we denote by $\mathcal{E}_\mathcal{G}$. A measurable function $f : V \to \mathbb{R}$ is said to be $\mathcal{G}$-*invariant* if, for every $g \in \mathcal{G}$, we have $f(g \cdot v) = f(v)$ for all $v \in V$. This is equivalent to f being measurable with respect to $\mathcal{E}_\mathcal{G}$.

In this paper we consider measurable actions mainly of *compact* groups. We shall use the following simple notion of the average of a measurable function with respect to such an action. (We shall only need to take the average of non-negative functions.)

Definition 2.7. Let $(V, \mathcal{V}, \mu)$ be a probability space, let $\mathcal{G}$ be a compact group with Haar probability measure $\mu_\mathcal{G}$, and let $\Phi : \mathcal{G} \times V \to V$ be a measurable action. Then, for

any non-negative measurable function $f : V \to \mathbb{R}$, we denote by $\vartheta_{\mathcal{G}}(f)$ the non-negative measurable function defined by $\vartheta_{\mathcal{G}}(f)(v) = \int_{\mathcal{G}} f(g^{-1} \cdot v) \, \mathrm{d}\mu_{\mathcal{G}}(g)$.

From our assumptions we have that the function $(g, v) \mapsto f(g^{-1} \cdot v)$ is $(\mathcal{B}_{\mathcal{G}} \times \mathcal{V})$-measurable. By Fubini's theorem [20, Theorem 8.8], we therefore have that $\vartheta_{\mathcal{G}}(f)$ is indeed a $\mathcal{V}$-measurable function, and satisfies

$$\int_V \vartheta_{\mathcal{G}}(f)(v) \, \mathrm{d}\mu(v) = \int_{\mathcal{G}} \left(\int_V f(g^{-1} \cdot v) \, \mathrm{d}\mu(v) \right) \mathrm{d}\mu_{\mathcal{G}}(g) = \int_{\mathcal{G} \times V} f(g^{-1} \cdot v) \, \mathrm{d}(\mu_{\mathcal{G}} \times \mu). \quad (6)$$

Note also that for any non-negative measurable functions f, g on V we have $\vartheta_{\mathcal{G}}(f + g) = \vartheta_{\mathcal{G}}(f) + \vartheta_{\mathcal{G}}(g)$, and in particular if $f \geq g$ then $\vartheta_{\mathcal{G}}(f) \geq \vartheta_{\mathcal{G}}(g)$.

A more general notion of averaging can be given in terms of the conditional expectation relative to the σ-algebra $\mathcal{E}_{\mathcal{G}}$, but the above definition is more convenient for us. (We discuss this in Remark 2.13.)

Definition 2.8 (*t*-partite action). Let $(V_i, \mathcal{V}_i, \mu_i), i \in [t]$, be probability spaces, and let $\mathcal{G}$ be a topological group. We say that an action $\Phi : \mathcal{G} \times V_{[t]} \to V_{[t]}$ is a *t-partite action* if it is of the following form: for each $i \in [t]$ there is a topological group G_i with a measurable action $\Phi_i : G_i \times V_i \to V_i$, such that $\mathcal{G}$ is a closed subgroup of $G_1 \times \cdots \times G_t$ (in the product topology) and for every $g \in \mathcal{G}, v \in V_{[t]}$ we have $\Phi(g, v)(i) = \Phi_i(g(i), v(i))$ for each $i \in [t]$.

In the next section we shall focus on *t*-partite actions where each V_i is a second-countable compact abelian group G_i acting on itself by addition. For the main results of this section, however, we can work with more general *t*-partite actions of compact groups. Let us record the following basic fact.

Lemma 2.9. *A t-partite action is a measurable action.*

Proof. The fact that a *t*-partite action Φ is indeed an action is straightforward. To see that the measurability of each map Φ_i implies measurability of Φ, it suffices to check this for an arbitrary product set $A = A_1 \times \cdots \times A_t, A_i \in \mathcal{V}_i$. To this end we note that $\Phi^{-1}A = (\mathcal{G} \times V_{[t]}) \cap \mathcal{R}\left(\prod_i \Phi_i^{-1}A_i \right)$, where $\mathcal{R} : \prod_i (G_i \times V_i) \to \left(\prod_i G_i \right) \times V_{[t]}$ is the map permuting the coordinates appropriately. We can then use the fact that each $\Phi_i^{-1}A_i$ lies in $\mathcal{B}_{G_i} \times \mathcal{V}_i$ to deduce that $\Phi^{-1}A$ lies in $\mathcal{B}_{\mathcal{G}} \times \mathcal{V}_{[t]}$. $\square$

Given a *t*-partite action of a compact group $\mathcal{G}$ on $V_{[t]}$, and a non-empty set $e \subseteq [t]$, we denote by $\mathcal{G}_e$ the closed subgroup $p_e(\mathcal{G})$ of $\prod_{i \in e} G_i$. Recall that the map p_e is the coordinate projection corresponding to e. On the direct product $G_1 \times \cdots \times G_t$, this map is a continuous homomorphism onto $\mathcal{G}_e$. We can then define a measurable action $\Phi_e : \mathcal{G}_e \times V_e \to V_e$ by $\Phi_e(g, v)(i) = \Phi_i(g(i), v(i))$.

Definition 2.10 (Invariant (t, m, k)-graph). Let $(V_i, \mathcal{V}_i, \mu_i), i \in [t]$, be probability spaces, and let $\mathcal{G}$ be a topological group with a *t*-partite action $\mathcal{G} \times V_{[t]} \to V_{[t]}$. A (t, m, k)-graph H with vertex sets V_i is said to be $\mathcal{G}$-*invariant* if for each $j \in [m]$, the edge set $E_j(H)$ is $\mathcal{G}_{C_j}$-invariant.

We shall use the following fact that relates averaging over $\mathcal{G}$ to averaging over $\mathcal{G}_e$, for each projection p_e.

Lemma 2.11. *Let $(V_i, \mathcal{V}_i, \mu_i), i \in [t]$, be probability spaces, and let $\mathcal{G} \times V_{[t]} \to V_{[t]}$ be a t-partite action by a compact group $\mathcal{G}$ with Haar probability. Then for any $e \in \binom{[t]}{k}$, for any non-negative measurable function $f : V_e \to \mathbb{R}$, we have*

$$\vartheta_{\mathcal{G}}(f \circ p_e) = (\vartheta_{\mathcal{G}_e}(f)) \circ p_e. \tag{7}$$

Proof. The actions Φ, Φ_e commute with p_e, that is we have

$$p_e(g \cdot v) = p_e(g) \cdot p_e(v) \text{ for every } g \in \mathcal{G}, v \in V_{[t]}. \tag{8}$$

Moreover, the map $p_e : \mathcal{G} \to \mathcal{G}_e$ is a surjective continuous homomorphism. We therefore have $\mu_{\mathcal{G}_e} = \mu_{\mathcal{G}} \circ p_e^{-1}$, where $\mu_{\mathcal{G}}, \mu_{\mathcal{G}_e}$ are the Haar probabilities on $\mathcal{G}, \mathcal{G}_e$. Thus for any $v \in V_{[t]}$ we have

$$
\begin{aligned}
\vartheta_{\mathcal{G}}(f \circ p_e)(v) &= \int_{\mathcal{G}} f\big(p_e(g^{-1} \cdot v)\big) \, \mathrm{d}\mu_{\mathcal{G}}(g) &= \int_{\mathcal{G}} f\big(p_e(g)^{-1} \cdot p_e(v)\big) \, \mathrm{d}\mu_{\mathcal{G}}(g) \\
&= \int_{\mathcal{G}_e} f\big(g_e^{-1} \cdot p_e(v)\big) \, \mathrm{d}\mu_{\mathcal{G}_e}(g_e) &= \vartheta_{\mathcal{G}_e}(f)\big(p_e(v)\big). \qquad \square
\end{aligned}
$$

We can finally establish the main result of this section.

Lemma 2.12 (Symmetry-preserving removal lemma).
Let $t \geq k \geq 2$ and m be positive integers, and let $\epsilon > 0$. There exists $\delta = \delta(t, k, \epsilon) > 0$ such that the following holds. Let $(V_i, \mathcal{V}_i, \mu_i), i \in [t]$, be probability spaces, let H be a (t, m, k)-graph with vertex sets V_i and edge color-classes C_j, let $\mathcal{G} \times V_{[t]} \to V_{[t]}$ be a t-partite action by a compact group $\mathcal{G}$ such that H is $\mathcal{G}$-invariant, let F be the (t, m, k)-graph on $[t]$ with edges $C_1, \ldots, C_m$, and suppose that $\tau(F, H) \leq \delta$. Then for each $j \in [m]$ there exists a measurable set $S_j \subseteq E_j(H)$ with $\mu_{C_j}(S_j) \leq \epsilon$, such that removing S_j from $E_j(H)$ for each $j \in [m]$ yields an F-free (t, m, k)-graph that is still $\mathcal{G}$-invariant.

An equivalent version of the conclusion is that for each $j \in [m]$ there exists a $\mathcal{G}_{C_j}$-invariant set $S_j \subseteq E_j(H)$ with $\mu_{C_j}(S_j) \leq \epsilon$, such that removing S_j from $E_j(H)$ for each $j \in [m]$ yields an F-free (t, m, k)-graph.

Proof. Let $R_j \subseteq E_j(H), j \in [m]$, be the removal sets given by Lemma 2.4 applied with parameter δ such that $\mu_{C_j}(R_j) \leq \epsilon/(2|E(F)|) = \epsilon/(2m)$.

We define a new removal set $S_j \subseteq E_j(H)$ as follows:

$$S_j := \{v \in V_{C_j} : h_j(v) > 1/(2m)\}, \text{ where } h_j := \vartheta_{\mathcal{G}_{C_j}}(1_{R_j}). \tag{9}$$

Note that h_j is a $\mathcal{G}_{C_j}$-invariant function, whence S_j is a $\mathcal{G}_{C_j}$-invariant measurable set. Moreover, we have $\mu_{C_j}(S_j) \leq \epsilon$. Indeed, by Markov's inequality and (6) we have

$$\frac{\mu_{C_j}(S_j)}{2m} \leq \int_{V_{C_i}} h_j(v) \, \mathrm{d}\mu_{C_j}(v) = \int_{\mathcal{G}_{C_j}} \left(\int_{V_{C_j}} 1_{R_j}(g^{-1} \cdot v) \, \mathrm{d}\mu_{C_j}(v) \right) \mathrm{d}\mu_{\mathcal{G}_{C_j}}(g) = \mu_{C_j}(R_j).$$

We now show that removing S_j from E_j for each $j \in [m]$ yields an F-free (t, m, k)-graph, i.e. that we have

$$\bigcap_{j \in [m]} p_{C_j}^{-1}(E_j \setminus S_j) = \emptyset. \tag{10}$$

In other words, we show that the function $\prod_{j \in [m]} 1_{E_j} \circ p_{C_j}$ is 0 everywhere on the region $\bigcap_{j \in [m]} p_{C_j}^{-1}(V_{C_j} \setminus S_j)$.

By a telescoping argument using multilinearity (similar to the one used in the proof of Lemma 2.4), we have

$$\prod_{j \in [m]} 1_{E_j} \circ p_{C_j} = \sum_{j \in [m]} 1_{R_j} \circ p_{C_j} \prod_{\ell \neq j} f_\ell \; + \; \prod_{j \in [m]} (1_{E_j} - 1_{R_j}) \circ p_{C_j},$$

where for each j we have $f_\ell = (1_{E_\ell} - 1_{R_\ell}) \circ p_{C_\ell}$ if $\ell < j$, and $f_\ell = 1_{E_\ell} \circ p_{C_\ell}$ if $\ell > j$. It follows that

$$\prod_{j \in [m]} 1_{E_j} \circ p_{C_j} \leq \sum_{j \in [m]} 1_{R_j} \circ p_{C_j} \; + \; \prod_{j \in [m]} (1_{E_j} - 1_{R_j}) \circ p_{C_j} \qquad \text{everywhere on } V_{[t]}.$$

Let us now apply $\vartheta_{\mathcal{G}}$ to both sides. Since each set E_j is $\mathcal{G}_{C_j}$-invariant, by (8) each set $p_{C_j}^{-1} E_j$ is in the σ-algebra of $\mathcal{G}$-invariant sets, and therefore so is their intersection. By the removal property of the sets R_j, we also have $\prod_{j \in [m]} (1_{E_j} - 1_{R_j}) \circ p_{C_j} = 0$ everywhere. Finally, by linearity and (7) we have $\vartheta_{\mathcal{G}} \left(\sum_{j \in [m]} 1_{R_j} \circ p_{C_j} \right) = \sum_{j \in [m]} h_j \circ p_{C_j}$. Combining these facts, we conclude that

$$\prod_{j \in [m]} 1_{E_j} \circ p_{C_j} - \sum_{j \in [m]} h_j \circ p_{C_j} \leq 0 \quad \text{everywhere on } V_{[t]}. \tag{11}$$

Now, on the region $\bigcap_{j \in [m]} p_{C_j}^{-1}(V_{C_j} \setminus S_j)$, the function $\sum_{j \in [m]} h_j \circ p_{C_j}$ takes values at most $1/2$, by definition of the sets S_j. Therefore, if $\prod_{j \in [m]} 1_{E_j} \circ p_{C_j}(v)$ were positive for some v in this region, then it would have to take value 1 at v and then the left side of (11) would be positive at v, a contradiction. $\square$

Remark 2.13. Recall that the conditional expectation relative to a sub-σ-algebra $\mathcal{E}$ of $\mathcal{V}$ can be defined on the Hilbert space $L^2(V, \mathcal{V}, \mu)$ as the orthogonal projection to the closed subspace $L^2(V, \mathcal{E}, \mu|_{\mathcal{E}})$; the conditional expectation of $f \in L^2(V, \mathcal{V}, \mu)$ relative to $\mathcal{E}$ is denoted $\mathbb{E}(f | \mathcal{E})$. If a compact group $\mathcal{G}$ with Haar probability has a measurable action on $(V, \mathcal{V}, \mu)$ then one can show that $\vartheta_{\mathcal{G}}$ agrees with the conditional expectation relative to the σ-algebra $\mathcal{E}_{\mathcal{G}}$ of $\mathcal{G}$-invariant sets. More precisely, letting f be any function class in $L^2(V, \mathcal{V}, \mu)$, and letting f' be any function in this class, we have that $\vartheta_{\mathcal{G}}(f')$ is in the class $\mathbb{E}(f | \mathcal{E}_{\mathcal{G}})$ (this can be proved by showing that ϑ yields an orthogonal projection $L^2(V, \mathcal{V}, \mu) \to L^2(V, \mathcal{E}_{\mathcal{G}}, \mu|_{\mathcal{E}_{\mathcal{G}}})$). This conditional expectation relative to $\mathcal{E}_G$ is defined even for actions that are not necessarily measurable. Thus one can obtain analogues of the results in this subsection for possibly non-measurable actions. However, $\mathbb{E}(f | \mathcal{E}_{\mathcal{G}})$ defines a function only up to a null-set, and this introduces several additional technicalities. Arguments using $\vartheta_{\mathcal{G}}$, as above, are therefore more convenient for our purposes, in addition to being more explicit.

In a similar vein, one can obtain analogues of the results in this section when each set $E_j(H)$ is only assumed to lie in the completion $\mathcal{V}_{C_j}^*$ of $\mathcal{V}_{C_j}$ relative to μ_{C_j}. One can also define a group action $\Phi : \mathcal{G} \times V \to V$ to be measurable in the weaker sense that Φ^{-1} takes values in the *completion* of $\mathcal{B}_{\mathcal{G}} \times \mathcal{V}$ relative to $\mu_{\mathcal{G}} \times \mu$. One can then use the version of Fubini's theorem for completed product measures [20, Theorem 8.12], but again this is less convenient for us.

3. Cayley (t, m, k)-graphs and systems of linear equations

Our aim now is to apply the results from the previous section to prove Theorem 1.3. We shall in fact prove the following version first.

Theorem 3.1. *Let $M \in \mathbb{Z}^{r \times m}$ satisfy $d_r(M) = 1$. For any $\epsilon > 0$, there exists $\delta = \delta(\epsilon, M)$, $0 < \delta < 1$, such that the following holds. Let $A_1, \ldots, A_m$ be Borel subsets of a second-countable compact Hausdorff abelian group G such that $\mu_{\ker_G M}\big(A_1 \times \cdots \times A_m \cap \ker_G M\big) \leq \delta$, and for each $j \in [m]$ let $G^{(j)}$ denote the closed subgroup $p_j(\ker_G M)$ of G. Then for each $j \in [m]$ there exists a Borel set $R_j \subseteq A_j \cap G^{(j)}$, such that $\mu_{G^{(j)}}(R_j) \leq \epsilon$ for all $j \in [m]$, and $\big(\prod_{j \in [m]} A_j \setminus R_j \big) \cap \ker_G M = \emptyset$.*

The added information here is firstly that we only need to remove elements from A_j that are in the projection $p_j(\ker_G M)$ of the solution space (this is quite clear intuitively and is also the case in Theorem 1.3). Secondly, each set R_j is small not just in the measure μ_G but in the possibly larger measure $\mu_{G^{(j)}}$. Indeed, note that if the index $\kappa_j = |G : G^{(j)}|$ is finite then we must have $\mu_{G^{(j)}} = \kappa_j \cdot \mu_G|_{G^{(j)}}$, where $\mu_G|_{G^{(j)}}$ denotes the restriction of μ_G to $G^{(j)}$. Thus, while the conclusion $\mu_{G^{(j)}}(R_j) \leq \epsilon$ above is roughly equivalent to the conclusion $\mu_G(R_j) \leq \epsilon$ in Theorem 1.3 if $\kappa_j = O(\epsilon^{-1})$, the former conclusion is stronger otherwise. We explain the use of second countability in Remark 3.3.

To prove Theorem 3.1, we want to find, given a system of linear equations of determinantal 1 on G, a certain invariant hypergraph that represents the system in such a way that the theorem follows from Lemma 2.12. In [29], a notion of a finite *Cayley hypergraph* was introduced and shown to give a representation of the desired kind for certain systems of equations on certain finite abelian groups, but it was not clear how far this method could be extended. The main objective for the remainder of this paper is to show that there is a general version of this framework that can handle all systems of determinantal 1.

We begin with a definition analogous to [29, Definition 2.1].

Definition 3.2 (Cayley (t, m, k)-graph)**.** We call a (t, m, k)-graph H a *Cayley (t, m, k)-graph* if it has the following properties. For each $i \in [t]$, the i-th vertex set is a compact group G_i with Borel σ-algebra and Haar probability, and there is a closed subgroup $\mathcal{G}$ of the direct product $\prod_{i \in [t]} G_i$ such that H is invariant under the t-partite action of $\mathcal{G}$ on $G_{[t]}$, where each G_i acts on itself by left-multiplication. To specify these properties, we write $H \in H_{k,t}((G_i), C, \mathcal{G})$, where C is the set of edge color-classes of H.

To illustrate this, let us note briefly how as a special case one finds "bipartite Cayley graphs" (as they are called in [9], for instance). Let $t = k = 2$, $m = 1$, $C = \{\{1, 2\}\}$, let $G_1 = G_2 = G$ be a finite group, and let $\mathcal{G} = \{(g, g) : g \in G\}$ be the diagonal subgroup of $G \times G$. Then $H \in H_{2,2}(G, C, \mathcal{G})$ means that the edge set $E(H)$ is a union of right cosets of $\mathcal{G}$. Using the map $(g_1, g_2) \mapsto g_2^{-1} g_1$, we can identify the quotient $\mathcal{G} \setminus E$ with a set $A \subseteq G$, and thus see that H is the bipartite Cayley graph on $G_1 \sqcup G_2$ generated by A (that is we have $(g_1, g_2) \in E$ if and only if $g_2^{-1} g_1 \in A$).

Remark 3.3. By Lemma 2.9, the t-partite action in Definition 3.2 is measurable (in the sense of Definition 2.6) if the action of each G_i on itself by left-multiplication is measurable. The latter measurability of the group operation holds for any second-countable

group G (that is a topological group such that the underlying topological space has a countable base). Indeed, by continuity of multiplication the preimage of a Borel set $A \subseteq G$ is Borel in $G \times G$, i.e. it lies in the Borel σ-algebra $\mathcal{B}_{G \times G}$. By second countability, we have that $\mathcal{B}_{G \times G}$ equals the product σ-algebra $\mathcal{B}_G \times \mathcal{B}_G$ (see [2, Lemma 6.4.2]), so the action is measurable. Without second countability, the σ-algebra $\mathcal{B}_{G \times G}$ may be strictly larger than $\mathcal{B}_G \times \mathcal{B}_G$ (see [2, Example 6.4.3]). These facts, together with other aspects (such as Lemma 3.5 below), make second countability a useful assumption in Theorem 3.1. Moreover, once this theorem has been proved, Theorem 1.3 can be deduced using an inverse limit argument. This is done in Appendix A. Thus, from now on we shall consider such invariant hypergraphs only on *second-countable* compact groups.

Defining Cayley (t, m, k)-graphs in terms of invariance, as above, relates them clearly to the previous section. To relate them to arithmetic removal results, it is useful to describe the edge sets of such hypergraphs in terms of generating sets.

Lemma 3.4. *Let $G_1, G_2, \ldots, G_t$ be second-countable compact groups, let $\mathcal{G}$ be a closed subgroup of $G_{[t]}$, and let $H \in H_{k,t}((G_i), C, \mathcal{G})$. For each $j \in [m]$, let ψ_{C_j} denote the canonical map from $G_{C_j} := \prod_{i \in C_j} G_i$ to the quotient topological space $p_{C_j}(\mathcal{G}) \backslash G_{C_j}$. Then for each j we have $E_j(H) = \psi_{C_j}^{-1}(A_j)$, where A_j is the Borel set $\psi_{C_j}(E_j(H))$.*

Thus, the edge set of H has the following form: $E(H) = \bigsqcup_{j \in [m]} \psi_{C_j}^{-1}(A_j)$. We call the sets A_j the *generators* of H. When the groups G_i are labelled copies of the same group G, we write $H \in H_{k,t}(G, C, \mathcal{G})$. If we wish to specify the generators, we shall write $H = H_{k,t}(G, C, \mathcal{G}, (A_j))$.

The only thing there is to prove in Lemma 3.4 is that each generator A_j is indeed a Borel set in $p_{C_j}(\mathcal{G}) \backslash G_{C_j}$. This fact is not trivial, since a priori the σ-algebra of Borel sets on this quotient could be smaller than the σ-algebra obtained by pushing forward, via ψ_{C_j}, the Borel subsets of G_{C_j}. In other words, we are using the following fact.

Lemma 3.5. *Let G be a Hausdorff second-countable compact group, let K be a closed subgroup of G, and let $\pi : G \to K \backslash G$ be the quotient map. Then for any K-invariant Borel set $E \subseteq G$, the set $\pi(E)$ is Borel.*

This follows from results in descriptive set theory, for instance combining [15, Theorem 12.17 and Corollary 15.2].

We now focus on *abelian* groups, and for these we shall now relate Cayley (t, m, k)-graphs to systems of linear equations. From now on, given $M \in \mathbb{Z}^{r \times m}$ and an abelian group G, we shall write (M, G) to refer to the system $Mx = 0$ with $x \in G^m$. Recall that our aim is to construct some invariant hypergraph H such that Theorem 3.1 for (M, G) can be deduced from the symmetry-preserving removal lemma for H.

One of the simplest examples of such a construction, the idea of which can be traced back to Ruzsa and Szemerédi [23], concerns Schur's equation $x_1 + x_2 = x_3$. Let us revisit this example in order to motivate our general construction.

Example 3.6 (Schur's equation, $M = (1 \quad 1 \quad -1)$). Consider the homomorphism $G^3 \to G^3$ given by the following matrix:

$$\Psi = \begin{pmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \\ 1 & 0 & -1 \end{pmatrix}. \tag{12}$$

This homomorphism has image equal to $\ker_G M$. Moreover, the row structure of Ψ allows us to define a very convenient tripartite Cayley graph, given Borel sets $A_1, A_2, A_3 \subseteq G$. Indeed, let H be the $(3, 3, 2)$-graph with three vertex sets equal to G, with edge color classes $C_1 = \{1, 2\}$, $C_2 = \{2, 3\}$, $C_3 = \{1, 3\}$, and with j-th edge-set $E_j = \psi_{C_j}^{-1} A_j$, where the map $\psi_{C_j} : G^{C_j} \to G$ is given by the j-th row of Ψ. (Thus for instance $\psi_{\{1,2\}}$ takes a couple $(v(1), v(2))$ from the product of the first two vertex sets to $v(1) - v(2)$.) Letting F be the triangle graph with vertices $1, 2, 3$, it can be checked easily that for each homomorphism $v = (v(1), v(2), v(3))$ of F in H, the image $x = \Psi(v)$ is an element of $A_1 \times A_2 \times A_3 \cap \ker_G M$, and that we have in fact $\mu_{\ker_G M}(A_1 \times A_2 \times A_3 \cap \ker_G M) = \tau(F, H)$. Moreover, H is a Cayley graph invariant under the 3-partite action of $\mathcal{G} = \ker_G \Psi$, which means here that each Borel set E_j is a union of cosets of $\ker_G \psi_{C_j}$. Therefore, if S_j is a set of small measure that is also a union of such cosets, then removing it from E_j corresponds to removing a subset of small measure from A_j. We can thus establish Theorem 3.1 for (M, G) using Lemma 2.12.

In order to generalize the argument above, we shall now define a type of group homomorphism Ψ that will enable us to associate a useful invariant hypergraph with a given system (M, G). The definition uses the following notation.

For a group G and a subset e of $[t]$, we denote by γ_e the homomorphism embedding the direct power G^e into G^t, defined by letting $\gamma_e(g')$ be the element g such that $g(i) = g'(i)$ for $i \in e$ and $g(i) = 0_G$ otherwise.

Given any abelian groups G_1, G_2, and $m, t \in \mathbb{N}$, any homomorphism $\Psi : G_1^t \to G_2^m$ can be viewed as an $m \times t$ matrix of homomorphisms $G_1 \to G_2$, namely for each $(j, k) \in [m] \times [t]$ the entry $\Psi_{j,k}$ is the homomorphism $p_j \circ \Psi \circ \gamma_k : G_1 \to G_2$. We denote by ψ_j the j-th row of this matrix, that is the homomorphism $\psi_j := p_j \circ \Psi = \sum_{k \in [t]} \Psi_{j,k} \circ p_k : G_1^t \to G_2$.

We write $\operatorname{Supp} \psi_j$ for the set of $k \in [t]$ such that $\Psi_{j,k}$ is not the 0-homomorphism $G_1 \to \{0_{G_2}\}$. When $\operatorname{Supp} \psi_j$ is a proper subset C_j of $[t]$, we will often want to work with the homomorphism $\psi_j \circ \gamma_{C_j} : G_1^{C_j} \to G_2$ rather than with $\psi_j : G_1^t \to G_2$. To simplify the notation, we shall denote $\psi_j \circ \gamma_{C_j}$ by ψ_{C_j}.

We can now give the main definition of this section.

Definition 3.7 (Hypergraph representation). Let G be an abelian group, and let $M \in \mathbb{Z}^{r \times m}$. A (t, m, k)-*representation* of the system (M, G) is a homomorphism $\Psi : G_*^t \to G^m$, for some abelian group G_*, such that the following conditions hold:

  (i) There are distinct sets $C_1, C_2, \ldots, C_m \in \binom{[t]}{k}$ such that $\forall j \in [m]$, $\operatorname{Supp} \psi_j \subseteq C_j$.
  (ii) $\Psi(G_*^t) = \ker_G M$.
  (iii) For each $j \in [m]$, we have $p_{C_j}(\ker_{G_*} \Psi) = \ker_{G_*} \psi_{C_j}$.

When G is second-countable compact, we require that the same be true for G_*, and that Ψ be continuous.

A simple example is given by the matrix Ψ in (12), which gives a $(3, 3, 2)$-representation for Schur's equation on any abelian group G, where we can take $G_* = G$.

The following proposition is the main result of this section.

Proposition 3.8. *Let $M \in \mathbb{Z}^{r \times m}$, let G be a second-countable compact abelian group, and suppose that the system (M, G) has a (t, m, k)-representation for some positive integers t, k. Then Theorem 3.1 holds for (M, G) with $\delta(\epsilon) = \delta(t, k, \epsilon)$, where $\delta(t, k, \epsilon)$ is given by Lemma 2.12.*

Proof. Let Ψ be a (t, m, k)-representation for (M, G). Fix $\epsilon > 0$ and let $\delta = \delta(t, k, \epsilon) > 0$ be such that Lemma 2.12 holds.

Let $A_1, A_2, \ldots, A_m$ be Borel subsets of G such that

$$\mu_{\ker_G M}(A_1 \times \cdots \times A_m \cap \ker_G M) \leq \delta.$$

We may assume that each A_j is a subset of $G^{(j)} = p_j(\ker_G M)$, since the part of A_j outside the latter subgroup does not contribute to the above measure.

Let $H = H_{k,t}(G_*, C, \mathcal{G}, (A_j))$ be the Cayley (t, m, k)-graph given by Ψ, that is the hypergraph with vertex sets $V_i = G_*$, with j-th edge-color-class C_j, with $\mathcal{G} = \ker_{G_*} \Psi$, and with generators A_j, $j \in [m]$. Let F be the k-uniform hypergraph on $[t]$ with edges $C_1, C_2, \ldots, C_m$.

Note the following fact concerning the quotient group $G_*^{C_j}/p_{C_j}(\mathcal{G})$ from Definition 3.2:

$$\forall j \in [m], \quad G_*^{C_j}/p_{C_j}(\ker_{G_*} \Psi) \;\cong\; \psi_j(G_*^t) \;=\; p_j(\ker_G M) =: G^{(j)}, \tag{13}$$

this being an isomorphism of compact abelian groups. Indeed, by condition (iii) of Definition 3.7, we have $G_*^{C_j}/p_{C_j}(\ker_{G_*} \Psi) = G_*^{C_j}/\ker_{G_*} \psi_{C_j}$. By the first isomorphism theorem, this is isomorphic as a compact abelian group to $\psi_{C_j}(G_*^{C_j})$. Since $\operatorname{Supp} \psi_j \subseteq C_j$, we have $\psi_{C_j}(G_*^{C_j}) = \psi_j(G_*^t)$. By definition of ψ_j, the latter group is $p_j \circ \Psi(G_*^t)$, and by condition (ii) this is $p_j(\ker_G M)$.

Now, since the map Ψ is measure-preserving from G_*^t onto $\ker_G M$ (as a continuous surjective homomorphism between compact abelian groups), we have

$$\tau(F, H) \;=\; \int_{G_*^t} \prod_{j \in [m]} 1_{E_j}(p_{C_j}(g)) \, \mathrm{d}g \;=\; \int_{G_*^t} \prod_{j \in [m]} 1_{A_j}(\psi_j(g)) \, \mathrm{d}g$$

$$=\; \mu_{\Psi(G_*^t)}\Big(A_1 \times \cdots \times A_m \cap \Psi(G_*^t)\Big)$$

$$=\; \mu_{\ker_G M}(A_1 \times \cdots \times A_m \cap \ker_G M) \leq \delta.$$

By Lemma 2.12, for each $j \in [m]$ there exists a Borel set $S_j \subseteq E_j(H)$, such that by removing S_j from $E_j(H)$ for each $j \in [m]$ we obtain a (t, m, k)-graph H' that is F-free and $\mathcal{G}$-invariant. In particular, each set S_j is invariant under the action of $p_{C_j}(\mathcal{G}) = \ker_{G_*} \psi_{C_j}$, so by Lemma 3.4 there is a Borel set $R_j \subseteq G^{(j)}$ such that $E_j \setminus S_j = \psi_{C_j}^{-1}(A_j \setminus R_j)$. We also have $\mu_{G^{(j)}}(R_j) = \mu_{C_j}(S_j) \leq \epsilon$. Moreover, the set $\left(\prod_{j \in [m]} A_j \setminus R_j\right) \cap \ker_G M$ must be empty, for if it contained some element $x = (x_1, \ldots, x_m)$ then there would be $g \in G_*^t$ such that $\Psi(g) = x$ and such that $g_{C_j} \in \psi_{C_j}^{-1}(A_j \setminus R_j) = E_j \setminus S_j$ for each $j \in [m]$, contradicting the removal property of the sets S_j. $\qquad\square$

We close this section by recording the deduction of Theorem 1.4 from Theorem 1.3.

Proof of Theorem 1.4. Suppose that $A \subseteq G$ is measurable with $\mu_G(A) \geq \alpha > 0$. Apply Theorem 1.3 with $\epsilon = \alpha/2m$. Let $c = \delta(\epsilon)$ and suppose that $\mu_{\ker_G M}(A^m \cap \ker_G M) < c$. Then by Theorem 1.3 there exists a measurable set $R \subseteq A$ of measure at most $\alpha/2$ such that $A \setminus R$ is M-free. However, $A \setminus R$ has measure at least $\alpha/2 > 0$, so it is non-empty, therefore it is not M-free (by invariance of M), a contradiction. $\square$

4. Finding a hypergraph representation for a given linear system

Having established Proposition 3.8, the proof of Theorem 3.1 is reduced to the following task: given an integer matrix M with determinantal 1 and any abelian group G, show that the system (M, G) admits a (t, m, k)-representation, with t, m, k depending only on M. In fact, we shall need to complete this task only for matrices that do not satisfy the following property.

Definition 4.1. We say that $M \in \mathbb{Z}^{r \times m}$ is *plain* if there exists $\ell \in [m]$ such that $p_\ell(\ker_G M) = \{0_G\}$ for every abelian group G.

This notion is a special case of that of a 'thin system' from [18]. Examples include any square matrix $M \in \mathbb{Z}^{r \times r}$ with $d_r(M) = \det M = 1$, since this has $\ker_G M = \{0_{G^r}\}$.

The following result allows us to restrict the above-mentioned task to non-plain matrices. Recall from the previous section the definition of the embedding homomorphism $\gamma_e : G^e \to G^m$ for a given $e \subseteq [m]$.

Lemma 4.2. *Let $M \in \mathbb{Z}^{r \times m}$ be a plain matrix satisfying $d_r(M) = 1$. Then either Theorem 3.1 holds for M, or for some $s \in [r-1]$ there exists a matrix $M' \in \mathbb{Z}^{(r-s) \times (m-s)}$ that is not plain and such that, for some set $C \subseteq [m]$ of size $m - s$, the map γ_C yields a (measure-preserving) isomorphism from $\ker_G M'$ to $\ker_G M$.*

In other words, any element $x' \in \ker_G M'$ can be extended uniquely to an element $x \in \ker_G M$ by adding coordinates equal to 0_G with indices in $[m] \setminus C$ (i.e. corresponding with columns from M missing in M'). As a consequence, if Theorem 3.1 holds for M' then it holds for M.

Proof. From Definition 4.1 we have $p_\ell(\ker_\mathbb{Q} M) = \{0\}$ for some $\ell \in [m]$. We claim that then there exists a unimodular matrix $U \in \mathbb{Z}^{r \times r}$ such that the matrix $M_0 = UM$ has first row equal[5] to the standard basis element e_ℓ. To see this, note that e_ℓ must be in the row space over $\mathbb{Q}$ of the rows of M, that is $e_\ell \in \text{Span}_\mathbb{Q}\{M_1, \ldots, M_r\}$. (Indeed, our assumption is that $\{e_\ell\}^\perp \supseteq \ker_\mathbb{Q} M$, so $e_\ell \in (\ker_\mathbb{Q} M)^\perp = \{M_1, \ldots, M_r\}^{\perp\perp}$.) Thus $e_\ell \in \text{Span}_\mathbb{Q}\{M_1, \ldots, M_r\} \cap \mathbb{Z}^m$. But this set equals $\text{Span}_\mathbb{Z}\{M_1, \ldots, M_r\}$ because $d_r(M) = 1$, as can be seen using the Smith normal form $M = V(I_r|0)W$ (where $V \in \mathbb{Z}^{r \times r}, W \in \mathbb{Z}^{m \times m}$ are unimodular, and I_r denotes the identity matrix of order r). Thus we have $e_\ell = n_1 M_1 + \cdots + n_r M_r$ where the n_i are coprime integers. By [18, Lemma 9], there exists a unimodular matrix $U \in \mathbb{Z}^{r \times r}$ with first row equal to $(n_1, n_2, \ldots, n_r)$. Thus we have $M_0 = UM$ as claimed, and so $\ker_G M = \ker_G M_0$. Now, with the notation from Theorem 3.1, if $0_G \notin A_\ell$, then Theorem 3.1 holds as there are no solutions x with $x_i \in A_i$, for all $i \in [m]$. Otherwise, we remove the first row of M_0 as well as the ℓ-th column, obtaining

---

[5]Incidentally, this claim also implies that in Definition 4.1, if $d_r(M) = 1$, then the case $G = \mathbb{Q}$ of the definition (i.e. $p_\ell(\ker_\mathbb{Q} M) = \{0\}$) implies the general case.

a matrix M'. Note that the embedding homomorphism $\gamma_{[m]\setminus\{\ell\}}$ is a measure-preserving isomorphism $\ker_G M' \to \ker_G M$. If M' is plain, we repeat the same procedure.

This iteration must produce the desired matrix M' before all the rows of M_0 are removed, for otherwise we would have that 0_{G^m} is the only solution and that $0_G \in A_j$ for each $j \in [m]$, which implies that $\mu_{\ker_G M}(\prod_j A_j \cap \ker_G M) = 1$, contradicting the assumption in Theorem 3.1. $\qquad\square$

Thus, our objective in this section is to prove the following result.

Proposition 4.3 (Existence of a hypergraph representation)**.** *Let* $M \in \mathbb{Z}^{r\times m}$ *with* $d_r(M) = 1$, *and suppose that* M *is not plain. Then there exist positive integers* t, k *such that, for any abelian group* G, *the system* (M, G) *has a* (t, m, k)-*representation.*

The combination of this result with Proposition 3.8 (via Lemma 4.2 if M is plain) establishes Theorem 3.1.

Remark 4.4. Note that if two matrices $M, M' \in \mathbb{Z}^{r\times m}$ satisfy $\ker_G M = \ker_G M'$, then a (t, m, k)-representation for (M', G) is also a (t, m, k)-representation for (M, G). In particular, for *finite* abelian groups G, in Proposition 4.3 the assumption $d_r(M) = 1$ can be relaxed to $\gcd(d_r(M), |G|) = 1$ (and the same holds for the finite case of Theorem 1.3). Indeed, the Smith normal form of M is then $U(D|0)V$ where $U \in \mathbb{Z}^{r\times r}, V \in \mathbb{Z}^{m\times m}$ are unimodular, and $D \in \mathbb{Z}^{r\times r}$ is a diagonal matrix with non-zero entries coprime with $|G|$, so the endomorphism $D : G^r \to G^r$ is invertible. Therefore, letting $M' = (I_r|0)V$, we have that $\ker_G M = \ker_G M'$, whence M has a (t, m, k)-representation if and only if M' does.

We shall prove Proposition 4.3 in several steps that constitute the subsections below. One of the main tools that we shall use is a notion of extension for integer matrices, which will enable us to replace the given matrix M by a simpler one at each step of the argument. To define this notion of extension, we use the following notation. Given a set $J \subseteq [m']$ of size m, and a group G, recall that we denote by p_J the coordinate projection $G^{m'} \to G^J$. Instead of the image group G^J, we shall often want to work with the group G^m, isomorphic to G^J. To avoid a possibly confusing abuse of notation, we shall denote by π_J the homomorphism $G^{m'} \to G^m$ that takes $(g_j)_{j\in[m']}$ to $(g_{\sigma_J(j)})_{j\in[m]}$, where σ_J is the order-preserving bijection $[m] \to J$.

Definition 4.5 (Matrix extension)**.** Let $r' \geq r$, $m' \geq m$ and let $M \in \mathbb{Z}^{r\times m}, M' \in \mathbb{Z}^{r'\times m'}$. We say that M' is an *extension* of M if the following holds. There is a subset $J \subseteq [m']$ of size m such that, for any abelian group G, the homomorphism $\pi_J : G^{m'} \to G^m$ restricts to an isomorphism $\ker_G M' \to \ker_G M$.

Note that if M and M' both have full rank, then we must have $m' - r' = m - r$, since this is the dimension of their isomorphic kernels over $G = \mathbb{Q}$. The extensions that we shall consider will always be given by a matrix M' having M as a submatrix in such a way that π_J has the required property.

The key fact that makes extensions useful for us is that they preserve the property of having a hypergraph representation, in the following sense.

Lemma 4.6. *Let $M \in \mathbb{Z}^{r \times m}$, and let $M' \in \mathbb{Z}^{r' \times m'}$ be an extension of M with corresponding index set $J \in \binom{[m']}{m}$. Let G be an abelian group, and suppose that Ψ' is a (t, m', k)-representation for (M', G). Then $\Psi := \pi_J \circ \Psi'$ is a (t, m, k)-representation for (M, G). Moreover if Ψ' is given by an integer matrix, then so is Ψ.*

Proof. First note that we can express the projection π_J as left-multiplication by the $m \times m'$ integer matrix whose j-th row is the vector with entry $\sigma_J(j)$ equal to 1 and all other entries 0, for each $j \in [m]$. Thus Ψ is an $m \times t$ homomorphism matrix with j-th row equal to the $\sigma_J(j)$-th row of Ψ', with support $C_j = C'_{\sigma_J(j)}$, where the latter is the support of the $\sigma_J(j)$-th row of Ψ'. In particular, the claim in the last sentence of the lemma is clear. Let us now check that the conditions of Definition 3.7 are satisfied.

Condition (i) is inherited by Ψ from Ψ', since the m rows of Ψ form a subset of the m' rows of Ψ'.

Condition (ii) is also satisfied, indeed we have

$$\Psi\big(G_*^{m'}\big) = \pi_J\big(\Psi'\big(G_*^{m'}\big)\big) = \pi_J(\ker_G M') = \ker_G M,$$

where the last equality follows from Definition 4.5.

To check condition (iii), fix $j \in [m]$. Then, given $y \in \ker_{G_*} \Psi \leq G_*^t$, we must have in particular the j-th coordinate of $\Psi(y)$ equal to 0, and this coordinate equals $\psi_{C_j}(p_{C_j}(y))$ by condition (i), whence $p_{C_j}(\ker_{G_*} \Psi) \subseteq \ker_{G_*} \psi_{C_j}$. To see the opposite containment, let $y' \in \ker_{G_*} \psi_{C_j}$. Since $C_j = C'_{\sigma_J(j)}$, we have $y' \in \ker_{G_*} \psi'_{C'_{\sigma_J(j)}}$, and since condition (iii) holds for Ψ', there exists $y \in \ker_{G_*} \Psi' \subseteq \ker_{G_*} \Psi$ such that $p_{C'_{\sigma_J(j)}}(y) = p_{C_j}(y) = y'$, so we have indeed $\ker_{G_*} \psi_{C_j} \subseteq p_{C_j}(\ker_{G_*} \Psi)$.

If G, G_* are topological groups and Ψ' is continuous, then so is Ψ. $\square$

4.1. **A reduction to matrices of the form $(I_r|B)$.** Our first application of matrix extensions consists in showing that to establish Proposition 4.3 it suffices to prove it for matrices $M = (I_r|B)$. To that end we shall use the following result, the role of which is analogous to [18, Lemma 10].

Lemma 4.7. *Let $M \in \mathbb{Z}^{r \times m}$ and suppose that $d_r(M) = 1$ and that M is not plain. Then M has an extension $M' \in \mathbb{Z}^{m \times (2m-r)}$, with $J = [m]$ and $M' = U(I_m|B)$, where $U \in \mathbb{Z}^{m \times m}$ is unimodular and every row of B is non-zero.*

Proof. By [18, Lemma 9] there exists an $m \times m$ matrix $U = \begin{pmatrix} M \\ E \end{pmatrix}$ satisfying $\det(U) = d_r(M) = 1$. We set M' to be the following $m \times (2m - r)$ matrix:

$$M' = \begin{pmatrix} M & 0 \\ E & I_{m-r} \end{pmatrix} = (U|B_0).$$

This is an extension of M with $J = [m]$. Letting $B = U^{-1}B_0$, we have $U^{-1}M' = (I_m|B)$ as required.

Since U is unimodular, we have $\ker_G M' = \ker_G(I_m|B)$ for any abelian group G. Therefore, if for some $j \in [m]$ the row B_j of B is 0, then

$$p_j(\ker_G M') = p_j(\ker_G(I_m|B)) = \{0_G\}.$$

Then, since $\pi_J : \ker_G M' \to \ker_G M$ is an isomorphism, we must also have $p_j(\ker_G M) = \{0_G\}$, whence M is plain. $\square$

We can now reduce the proof of Proposition 4.3 to establishing the following result.

Proposition 4.8. *Suppose that $M \in \mathbb{Z}^{r \times m}$ is of the form $M = (I_r|B)$, where $m \geq r+1$ and all rows of B are non-zero. Then there exist positive integers t, k such that, for any abelian group G, the system (M, G) has a (t, m, k)-representation.*

Lemma 4.9. *Proposition 4.8 implies Proposition 4.3.*

Proof. Suppose that $M_0 \in \mathbb{Z}^{r_0 \times m_0}$ satisfies the assumptions in Proposition 4.3. We have $m_0 \geq r_0 + 1$, otherwise M_0 is plain. Applying Lemma 4.7, we obtain a matrix $M' = U(I_r|B) \in \mathbb{Z}^{r \times m}$, with $r = m_0$, $m = 2m_0 - r_0 \geq r+1$, and all rows of B non-zero, such that M' is an extension of M_0 with $J = [m_0]$. Let $M = (I_r|B)$, let t, k be the integers given by Proposition 4.8, and let G be an abelian group. By Proposition 4.8 there is a (t, m, k)-representation Ψ for (M, G). Then, since U is unimodular, we have $\ker_G M' = \ker_G M$, and so Ψ is also a (t, m, k)-representation for (M', G), just by Definition 3.7. Hence, by Lemma 4.6, the map $\pi_{[m_0]} \circ \Psi$ is a (t, m_0, k)-representation for the original system (M_0, G). $\qquad\square$

Our goal now is to prove Proposition 4.8. To begin with, in the next subsection we deal with a special case consisting of what we call *simple* matrices.

4.2. **Simple matrices.** Given a non-zero element $v \in \mathbb{Z}^m$, we denote by $\gcd(v)$ the greatest common divisor of the integers $v(i), i \in [m]$.

Definition 4.10. We say that a matrix $M \in \mathbb{Z}^{r \times m}$ of the form $(I_r|B)$ is *simple* if $m \geq r+2$, and for each $i \in [r]$ the i-th row of B, denoted B_i, is non-zero and satisfies $\gcd(B_i) = 1$.

Our main result concerning these matrices is the following.

Proposition 4.11. *Suppose that $M = (I_r|B) \in \mathbb{Z}^{r \times m}$ is simple. Then, for some positive integers t, k, there exists $\Psi \in \mathbb{Z}^{m \times t}$ such that, for any abelian group G, the homomorphism $\Psi : G^t \to G^m$ is a (t, m, k)-representation for the system (M, G).*

Given $M \in \mathbb{Z}^{r \times m}$, for each $j \in [m]$ let $M_{(j)}$ denote the square matrix formed by the columns of M with indices $j - r \mod m, j - (r-1) \mod m, \ldots, j-1 \mod m$.

The main part of the proof of Proposition 4.11 consists in showing that any simple matrix has an extension with the particularly convenient property of being what we call a circular matrix.

Definition 4.12 (Circular matrix)**.** We say that a matrix $M \in \mathbb{Z}^{r \times m}$ is *circular* if for each $j \in [m]$ the matrix $M_{(j)}$ is unimodular.

Proposition 4.13. *Let $M \in \mathbb{Z}^{r \times m}$ be a simple matrix. Then there is an extension $M' \in \mathbb{Z}^{r' \times m'}$ such that M' is circular.*

Before we turn to the proof, let us motivate this proposition by briefly discussing circular matrices. One of the simplest examples of a circular matrix is the one corresponding to Schur's equation, that is $M = (1 \quad 1 \quad -1)$. In Example 3.6 we saw that this has a nice representation, given in (12), having the triangle as its corresponding graph F. Circular matrices are very convenient in that they provide simple generalizations of this construction, as shown by the following result.

Lemma 4.14. *Suppose that $M \in \mathbb{Z}^{r \times m}$ is circular and that $m \geq r + 2$. Then there exists $\Psi \in \mathbb{Z}^{m \times m}$ such that for any abelian group G, the homomorphism $\Psi : G^m \to G^m$ is an $(m, m, r + 1)$-representation of (M, G), with $C_j = \{j, j + 1 \mod m, \ldots, j + r \mod m\}$ for each $j \in [m]$.*

Thus, the triangle graph corresponding to Schur's equation is generalized here to the 'cyclic' $(r + 1)$-uniform hypergraph on $[m]$ with edges C_j. Analogues of this construction have been used in previous works (though not in relation to hypergraph representations as defined here), specifically in [3, 18]. In particular, Definition 4.12 is an analogue of the notion of 'n-circular matrix' used in [18].

Proof of Lemma 4.14. We construct Ψ as follows: the j-th column Ψ^j is an element of $\mathbb{Z}^m$ lying in $\ker_{\mathbb{Q}} M$, with support inside $\{j - r, j - (r - 1), \ldots, j - 1, j\}$ (subtractions mod m), and with j-th entry equal to -1. More precisely, let $y = M_{(j)}^{-1} M^j \in \mathbb{Z}^r$; then $M_{(j)} y = M^j$. We then define Ψ^j by $\Psi^j(i) = y(i)$ for $i \in \{j - r, j - (r - 1), \ldots, j - 1\}$, $\Psi^j(j) = -1$, and $\Psi^j(i) = 0$ otherwise. Note that we have indeed $\Psi^j \in \mathbb{Z}^m$ and $M \Psi^j = 0$. Note also that the resulting matrix Ψ has row j with support indeed contained in the set $C_j = j + [0, r] \mod m$, and that these sets C_j are distinct since $m \geq r + 2$. Hence, condition (i) from Definition 3.7 is satisfied with $k = r + 1$.

Let us check condition (ii), i.e. that $\Psi(G^m) = \ker_G M$. Since $M\Psi = 0$, we clearly have $\Psi(G^m) \subseteq \ker_G M$. To see equality, fix any $x \in \ker_G M$. Observe that x is uniquely determined by any sequence of $m - r$ consecutive coordinates mod m, because the submatrix formed by the remaining r columns of M, being unimodular, gives a bijection on G^r. Hence, if we find $y \in G^m$ such that $\Psi(y)$ agrees with x on such a sequence of $m - r$ coordinates, then this together with the fact that $M\Psi(y) = 0$ will imply that $x = \Psi(y) \in \Psi(G^m)$. Now note that the top-left square submatrix of Ψ of order $m - r$ is upper triangular with entries -1 in the diagonal, so we can indeed find the desired element y.

To check condition (iii), let $\mathcal{G} = \ker_G \Psi$. We have to check that for each $j \in [m[$, the map Ψ_{C_j} (the restriction of the j-th row of Ψ to G^{C_j}) satisfies $\ker \psi_{C_j} = p_{C_j}(\mathcal{G})$. Clearly $p_{C_j}(\mathcal{G}) \subseteq \ker \psi_{C_j}$, since if $\Psi(y) = 0$ then in particular the j-th entry, which equals $\psi_{C_j}\big(p_{C_j}(y)\big)$, is 0. To see the opposite containment, suppose that $y' \in G^{C_j}$ satisfies $\psi_{C_j}(y') = 0$. We want to show that there exists $y \in \mathcal{G}$ such that $p_{C_j}(y) = y'$. Using the row structure of Ψ, we can find successively elements $y_{j-1}, y_{j-2}, \ldots, y_{j-(m-r-1)} \in G$ such that for each $\ell \in \{j, j - 1, \ldots, j - (m - r - 1)\}$ we have $\psi_{C_\ell}(y_\ell, y_{\ell+1}, \ldots, y_{\ell+r}) = 0$. We use these elements y_ℓ to extend y' to an element $y \in G^m$, defined by $p_{C_j}(y) = y'$ and $p_\ell(y) = y_\ell$ for each $\ell \in [m] \setminus C_j$. By construction, Ψy has $m - r$ coordinates equal to 0. Since 0 and Ψy are both in $\ker_G M$ and agree on these $m - r$ coordinates, by the observation in the previous paragraph we must have $\Psi y = 0$, so $y \in \mathcal{G}$ as required.

If G is a topological group then Ψ is clearly continuous. We have thus shown that all the conditions in Definition 3.7 are satisfied. $\square$

Let us now turn to establishing Proposition 4.13. The proof is an adaptation of an argument from [18]. Given a matrix B, we shall denote by $B_{[i_1, i_2]}$ the submatrix of B formed by consecutive rows with indices $i_1, i_1 + 1, \ldots, i_2$. First we adapt [18, Lemma 11], to obtain the following.

Lemma 4.15. *Let $B \in \mathbb{Z}^{r \times r}$ be a unimodular matrix. Then for some integer $s = O_B(1)$, there exist integer matrices S, T such that the $s \times r$ matrix*

$$\overline{B} = \begin{pmatrix} I_r \\ S \\ B \\ T \\ I_r \end{pmatrix} \tag{14}$$

satisfies the following property: for each $i \in [1, s - (r-1)]$, the $r \times r$ submatrix $\overline{B}_{[i, i+r-1]}$ is unimodular.

In other words, each submatrix of $\overline{B}$ formed by r consecutive rows is unimodular.

Proof. We shall say that an integer matrix with r columns is *good* if each of its square submatrices formed by r consecutive rows is unimodular.

We first claim that there exists a matrix T such that the matrix $L = \begin{pmatrix} B \\ T \\ I_r \end{pmatrix}$ is good and has $O_B(1)$ rows. (The upper part of $\overline{B}$ will be found analogously.)

This holds for $r = 1$ since we can set $L = \begin{pmatrix} \pm 1 \\ 1 \end{pmatrix}$. For $r > 1$, we can suppose by induction that the claim holds for $r - 1$. The matrix L is constructed by repeatedly adding an appropriate new row at the bottom of B while ensuring that the new bottom $r \times r$ submatrix is unimodular. The idea is that each new row essentially captures a step in an application of Euclid's algorithm to the entries in the first column of B.

Thus we first form the matrix $\begin{pmatrix} B_2 \\ B_3 \\ \vdots \\ B_r \\ \sum_{i=1}^{r} \lambda_i B_i \end{pmatrix}$, where $\lambda_1 \in \{-1, 1\}$, and such that $|\lambda_1 B_{1,1} + \sum_{i=2}^{r} \lambda_i B_{i,1}|$ is minimized. This coefficient λ_1 having magnitude 1 ensures that the new row is the result of an elementary row operation on B, so that the above matrix is still unimodular. If $|B_{1,1}| = \max_{i \in [r]} |B_{i,1}| > 1$, then we can find $\{\lambda_i : i \in [r]\}$ with $|\lambda_1| = 1$ such that $|\lambda_1 B_{1,1} + \sum_{i=2}^{r} \lambda_i B_{i,1}| < |B_{1,1}|$. Otherwise, note that we can certainly make the left side here at most $|B_{1,1}|$. It follows that after repeating this addition of a new row at most r times, we have decreased the greatest magnitude of the entries in the first column (provided that this magnitude was greater than 1). We then iterate this process, denoting by T_i the i-th new bottom row. By Euclid's algorithm, after $\ell = O_B(1)$ steps, we obtain a good matrix $\begin{pmatrix} B \\ T' \end{pmatrix}$ where $T' = \begin{pmatrix} T_1 \\ T_2 \\ \vdots \\ T_\ell \end{pmatrix}$ and the first entry of T_ℓ is 1 (the entries in the first column of B are coprime by assumption). Now we can carry out $r - 1$ further steps consisting in subtracting integer multiples of T_ℓ from previous rows, to obtain the $r \times r$ matrix $T' = \begin{pmatrix} 1 & * \\ 0 & B' \end{pmatrix}$, with top row T_ℓ, and with $B' \in \mathbb{Z}^{(r-1) \times (r-1)}$

being unimodular.

Now we apply the induction hypothesis to B', obtaining an $s' \times (r-1)$ good matrix $L' = \begin{pmatrix} B' \\ T'' \\ I_{r-1} \end{pmatrix}$, where $s' = O_{B'}(1) = O_B(1)$. We then add to L' a first column of zeros, and we insert in the resulting matrix the row $(1, 0, \ldots, 0)$ of length r between the positions $j(r-1)$ and $j(r-1) + 1$, for each $j \in [1, s'-1]$. The resulting matrix L'' has $O_B(1)$ rows and, by construction, the following matrix is good: $L = \begin{pmatrix} B \\ L'' \end{pmatrix} = \begin{pmatrix} B \\ T \\ I_r \end{pmatrix}$, for some matrix T with $O_B(1)$ rows.

The proof is completed by a similar argument adding top rows to B, yielding the desired matrix S. $\qquad\square$

To complete the proof of Proposition 4.13, we adapt the argument from [18, Lemma 12], using Lemma 4.15 instead of [18, Lemma 11].

Proof of Proposition 4.13. Since each row B_i of B satisfies $\gcd(B_i) = 1$, by [18, Lemma 9] there exists an $(m-r) \times (m-r)$ unimodular matrix U_i with top row equal to B_i. Applying Lemma 4.15 to each such matrix U_i we obtain $\overline{U_i}$ as given by (14). We then form the following $r' \times (m-r)$ matrix, where $r' = O_B(1)$: $B' = \begin{pmatrix} \overline{U_1} \\ \overline{U_2} \\ \vdots \\ \overline{U_r} \end{pmatrix}$. Let $M' = (I_{r'} | B')$, and note that $M' = \begin{pmatrix} I_{r'} & \begin{matrix} I_{m-r} \\ X \\ I_{m-r} \end{matrix} \end{pmatrix}$, for some matrix X.

We claim that M' is circular. To see this, let $M'_{(i)}$ denote the square submatrix formed by r' consecutive columns of M' in the circular order, starting with the i-th column. Then, for the first $m-r$ and last $m-r$ values of $i \in [m']$ it is clear that $M'_{(i)}$ is unimodular; for example, for the first $m-r$ values, $M'_{(i)}$ is unimodular because its columns form a circular permutation of the columns of a lower triangular matrix with diagonal entries equal to 1. For $i \in (m-r, m'-(m-r)] = (m-r, r']$, note that $\det M'_{(i)} = \pm \det B'_{[i-(m-r), i-1]} = \pm 1$ so $M'_{(i)}$ is indeed unimodular.

To complete the proof, let us specify the index set $J \subseteq [m']$ of size m showing that M' is an extension of M. Let $J_1 \subseteq [r']$ be the set of size r containing the subscript of each row of B' that is the first row of a submatrix U_i (recall that U_i is a submatrix of $\overline{U_i}$). Since this first row is B_i by construction, we thus have that the order-preserving bijection $\sigma_{J_1} : [r] \to J_1$ satisfies $B_i = B'_{\sigma_{J_1}(i)}$. We set $J = J_1 \cup [r'+1, m'] \subseteq [m']$. From the structure of M', it then follows that the homomorphism π_J restricted to $\ker_G M'$ gives an isomorphism $\ker_G M' \to \ker_G M$, as required. Indeed, the submatrix of M' formed by the rows indexed by J_1 is equal to M (up to relabelling rows and columns) so if $x' \in \ker_G M'$ then $\pi_J(x') \in \ker_G M$; moreover, given $x \in \ker_G M$, the element of G^J with j-th coordinate $x(\sigma_J^{-1}(j))$ is in the kernel of the above submatrix, and using the structure of M' we then extend this element uniquely to an element $x' \in \ker_G M'$

such that $\pi_J(x') = x$ (since an element $x' \in \ker_G M'$ is uniquely determined by its last $m' - r' = m - r$ coordinates). $\qquad\square$

Let us finally combine the above ingredients to obtain the main result of this subsection.

Proof of Proposition 4.11. By Proposition 4.13 there is a circular matrix $M' = (I_{r'}|B') \in \mathbb{Z}^{r' \times m'}$ extending M. By Lemma 4.14 we have an $(m', m', r' + 1)$-representation for M' given by an integer matrix Ψ'. Hence, by Lemma 4.6, we obtain a (t, m, k)-representation $\Psi = \pi_J \circ \Psi'$ for M, also given by an integer matrix, with $t = m'$ and $k = r' + 1$. $\qquad\square$

Remark 4.16. Note that Proposition 4.11 establishes Proposition 4.8 for simple matrices in the strong sense that the conditions in Definition 3.7 of hypergraph representability are satisfied with $G_* = G$ and Ψ being just an integer matrix (rather than a more general homomorphism matrix). In the next subsection, the full generality of Definition 3.7 will be used to handle all remaining matrices of the form $(I_r|B)$. For this purpose, instead of matrix extensions, we shall use a different construction.

4.3. **General matrices of the form** $(I_r|B)$. In this subsection we complete the proof of Proposition 4.8, by using Proposition 4.11 to construct a hypergraph representation for any system (M, G) with $M \in \mathbb{Z}^{r \times m}$ a non-plain matrix of the form $(I_r|B)$. This will establish Proposition 4.3, and thereby we shall have completed the proof of Theorem 3.1.

Let us first explain briefly the main difficulty, and in particular why the rest of the argument from the finite setting of [18] does not work in our setting of general compact abelian groups.

We want to find a hypergraph representation for any given system with a non-simple $r \times m$ matrix $(I_r|B)$. There are two cases to treat: in the main case we have $m \geq r + 2$ and for some row B_i of B we have $\gcd(B_i) = s > 1$; in the second case we have $m = r+1$.

In [18], an analogue of the main case is addressed using a notion of 'system extension', which differs from the matrix extensions used in this paper. In particular, the extensions in [18] allow one to multiply group elements by s and thus reduce the task to the case of simple matrices; see for instance the proof of [18, Lemma 10]. This multiplication is allowed in the setting of finite abelian groups because it does not increase the measure of sets, which is important to ensure that the measures of the removal sets are kept small. Our general setting includes groups in which multiplication by an integer may increase measures (e.g. the circle group), so we cannot use this argument. Instead, we shall construct a certain 'covering' of the original kernel $\ker_G M$ by kernels of systems given by *simple* matrices associated with the original matrix.[6] Proposition 4.11 gives us a representation for each of these simple matrices, and we shall then combine these to obtain a representation for the original system. This will address the first case stated above.

The second case will then be simpler to handle and will be treated at the end of this section.

Before we go into the details of the main case, let us briefly illustrate the idea of the argument.

---

[6]This construction plays a role, relative to Definition 3.7, somewhat analogous to the role played by the blowup construction described in [33] relative to the arguments in [18].

Example 4.17. Consider a system (M, G) with $M = (1\ 2\ 2)$, a non-simple matrix of the form $(I_1|B)$. Consider then the following two systems: $(M^{(0)}, G_0)$ with $M^{(0)} = (1\ 1\ 1)$ and $G_0 = G$, and $(M^{(1)}, G_1)$ with $M^{(1)} = (1\ 2\ 2\ 1)$ and G_1 the subgroup of G consisting of the preimages of 0_G under multiplication by 2. Note that $M^{(0)}$ and $M^{(1)}$ are simple matrices, so we have representations given by matrices $\Psi^{(0)} \in \mathbb{Z}^{3 \times 3}$ and $\Psi^{(1)} \in \mathbb{Z}^{4 \times t}$ for $(M^{(0)}, G_0)$, $(M^{(1)}, G_1)$ respectively.

Let us denote an element of $\ker_{G_0} M^{(0)}$ by $\mathbf{x}^{(0)} = (x_{0,1}, x_{0,2}, x_{0,3})$ and similarly an element of $\ker_{G_1} M^{(1)}$ by $\mathbf{x}^{(1)} = (x_{1,1}, x_{1,2}, x_{1,3}, x_{1,4})$. Consider the following map:

$$K \ : \ \ker_{G_0} M^{(0)} \times \ker_{G_1} M^{(1)} \to \ker_G M$$
$$(\mathbf{x}^{(0)}, \mathbf{x}^{(1)}) \mapsto \ \mathbf{x} = (2(x_{0,1} + x_{1,1}), x_{0,2} + x_{1,2}, x_{0,3} + x_{1,3}).$$

This map is surjective. Moreover, the preimages of any solution $\mathbf{x} \in \ker_G M$ have convenient 'covering' properties when fixing any given coordinate $\mathbf{x}(j)$ (see part (iii') of Definition 4.21), properties which are obtained essentially by using the component $x_{1,4}$ as a free variable. These properties are then used to construct a representation Ψ, consisting essentially in a $3 \times (3 + t)$ matrix in which the left 3×3 submatrix is given by $\Psi^{(0)}$ and the right submatrix is given by $\Psi^{(1)}$ without the last row.

To define our construction formally, we shall use the following simple fact.

Lemma 4.18. *Let $M \in \mathbb{Z}^{r \times m}$, let G be an abelian group, and suppose that $\Psi' : G_*'^t \to G^m$ is a (t, m, k)-representation for (M, G). Let G_* be another abelian group and let $\pi : G_* \to G_*'$ be a surjective homomorphism (if G is compact second countable then we assume that the same holds for G_*, and that π is continuous). Then letting π^t denote the homomorphism $G_*^t \to G_*'^t$ mapping $g = (g_1, \ldots, g_t)$ to $(\pi(g_1), \ldots, \pi(g_t))$, we have that $\Psi := \Psi' \circ \pi^t : G_*^t \to G^m$ is also a (t, m, k)-representation for (M, G).*

Proof. Condition (i) from Definition 3.7 holds clearly for Ψ.

Condition (ii) is also clear: by assumption we have $\ker_G M = \Psi'(G_*'^t)$ and this equals $\Psi'(\pi^t(G_*^t))$ by surjectivity of π.

For condition (iii), fix any $j \in [m]$ and note that we certainly have $p_{C_j}(\ker_{G_*} \Psi) \subseteq \ker_{G_*} \psi_{C_j}$. To see the opposite containment, suppose that $g' \in G_*^{C_j}$ satisfies $\psi_{C_j}(g') = 0_G$ and note that since $\psi_{C_j}(g') = \psi_{C_j}'(\pi^{C_j}(g'))$ and condition (iii) holds for Ψ', there exists $g_0' \in \ker_{G_*'} \Psi'$ such that $p_{C_j}(g_0') = \pi^{C_j}(g')$. By surjectivity of π there exists $g_0 \in G_*^t$ such that $\pi^t(g_0) = g_0'$. We may not have $p_{C_j}(g_0) = g'$, but we do have this equality mod π^{C_j}, that is we have

$$\pi^{C_j}(p_{C_j}(g_0)) = p_{C_j}(\pi^t(g_0)) = p_{C_j}(g_0') = \pi^{C_j}(g').$$

Hence there exists $h' \in \ker_{G_*} \pi^{C_j}$ such that $p_{C_j}(g_0) = g' + h'$. Adding coordinates equal to 0_{G_*} to h', we obtain $h \in G_*^t$ such that $p_{C_j}(h) = h'$ and $\pi^t(h) = 0$. Letting $g = g_0 - h$, we have $\Psi(g) = \Psi'(g_0') - \Psi'(\pi^t(h)) = 0_{G^m}$ and $p_{C_j}(g) = g'$. Hence condition (iii) holds.

Finally, if Ψ' is continuous on the compact abelian group $G_*'^t$, then Ψ is also continuous on the compact abelian group G_*^t. □

Let us now describe the construction for the main case in detail.

Let $M = (I_r|B) \in \mathbb{Z}^{r \times m}$ be the given matrix with $m \geq r + 2$ and some row B_i satisfying $\gcd(B_i) > 1$. Let B' denote the matrix obtained from B by dividing, for each $i \in [r]$,

each coordinate of B_i by $\gcd(B_i)$.

Let $M^{(0)}$ be the simple matrix $M^{(0)} = (I_r|B')$, and for each $i \in [r]$ let $M^{(i)}$ be the following $r \times (m+1)$ simple matrix

$$M^{(i)} = \begin{pmatrix} I_{i-1} & 0 & 0 & B'_{[1,i-1]} & 0 \\ \mathbf{0} & 1 & 0 & B_i & 1 \\ \mathbf{0} & 0 & I_{r-i} & B'_{[i+1,r]} & 0 \end{pmatrix},$$

where $B'_{[i_1,i_2]}$ denotes the submatrix of B' formed by rows $i_1, i_1 + 1, \ldots, i_2$.

Let G be an arbitrary abelian group. For each $i \in [r]$ let G_i denote the preimage of 0_G under multiplication by $\gcd(B_i)$ (thus $G_i \le G$), and let $G_0 = G$.

Let $G_* = G_0 \times G_1 \times \cdots \times G_r$. (Note that G_* is compact second countable if G is.) For each $i \in [0, r]$, let π_i be the projection homomorphism $G_* \to G_i$, and let $\Psi'^{(i)}$ be the matrix in $\mathbb{Z}^{m_i \times t_i}$ given by Proposition 4.11, thus $\Psi'^{(i)} : G_i^{t_i} \to G_i^{m_i}$ is a (t_i, m_i, k_i)-representation for the system $(M^{(i)}, G_i)$. Then, by Lemma 4.18, the homomorphism

$$\Psi^{(i)} := \Psi'^{(i)} \circ \pi_i^{t_i} : G_*^{t_i} \to G_i^{m_i}$$

is also a (t_i, m_i, k_i)-representation for the system $(M^{(i)}, G_i)$.

Let $t = t_0 + t_1 + \cdots + t_r$. We shall now combine these representations $\Psi^{(i)}$ to define a map $\Psi : G_*^t \to G^m$.

Each element $g \in G_*^t$ may be written in the form $g = \left(g^{(0)}, g^{(1)}, \ldots, g^{(r)} \right)$, where $g^{(i)} \in G_*^{t_i}$ for each $i \in [0, r]$. We then define the homomorphism

$$\begin{aligned} \Phi : G_*^t &\to G_0^m \times G_1^{m+1} \times \cdots \times G_r^{m+1} \le G^{m+r(m+1)} \\ g &\mapsto \left(\Psi^{(0)}(g^{(0)}), \Psi^{(1)}(g^{(1)}), \ldots, \Psi^{(r)}(g^{(r)}) \right). \end{aligned} \tag{15}$$

Note that Φ can be viewed as an $(m + r(m+1)) \times t$ matrix of homomorphisms, where the top-left $m \times t_0$ submatrix is $\Psi^{(0)}$, the submatrix on the next $m + 1$ rows and t_1 columns is $\Psi^{(1)}$, and so on, and every other entry is the zero homomorphism.

By condition (ii) from Definition 3.7 for each $\Psi^{(i)}$, we have

$$\Phi\left(G_*^t \right) = \ker_{G_0} M^{(0)} \times \ker_{G_1} M^{(1)} \times \cdots \times \ker_{G_r} M^{(r)}.$$

We shall denote an element of this group by $\mathbf{x} = \left(\mathbf{x}^{(0)}, \mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(r)} \right)$, where $\mathbf{x}^{(i)} \in \ker_{G_i} M^{(i)}$.

To complete the definition of Ψ, we shall now compose Φ with another homomorphism, denoted K, which will combine the entries of $\Phi(g)$ appropriately to produce an element of $\ker_G M$.

We define

$$K := P \circ \Sigma \; : \; \prod_{i=0}^r \ker_{G_i} M^{(i)} \; \to \; G^m, \tag{16}$$

where Σ is the following coordinate-summation map:

$$\begin{aligned} \Sigma : \quad \prod_{i=0}^r \ker_{G_i} M^{(i)} &\to G^m \\ \mathbf{x} = \left(\mathbf{x}^{(0)}, \mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(r)} \right) &\mapsto \left(\sum_{i=0}^r \mathbf{x}^{(i)}(j) \right)_{j \in [m]}, \end{aligned}$$

and where P is the following coordinate-multiplication map:

$$
\begin{aligned}
P: \quad G^m &\to G^m \\
y &\mapsto \left(\gcd(B_1)\, y_1,\ \gcd(B_2)\, y_2,\ \ldots,\ \gcd(B_r)\, y_r,\ y_{r+1},\ \ldots,\ y_m \right).
\end{aligned}
$$

Note that Σ (and hence K) ignores the $(m+1)$-st component of $\mathbf{x}^{(i)}$ for each $i \in [r]$.

Definition 4.19. Let $M = (I_r | B) \in \mathbb{Z}^{r \times m}$ with $m \geq r + 2$, let G be an abelian group, and suppose that M is not simple. Then we define the following homomorphism:

$$
\begin{aligned}
\Psi : G_*^t &\to G^m \\
g &\mapsto K \circ \Phi(g),
\end{aligned}
\tag{17}
$$

where G_*, K, Φ, t are as defined above.

One can view Ψ as an $m \times t$ matrix of homomorphisms, with $r+1$ submatrices formed by sets of consecutive columns, where the i-th submatrix is formed by t_i such columns and is equal to $P \circ \Psi^{(i)}$, for each $i \in [0, r]$.

Proposition 4.20. *Let $M = (I_r | B) \in \mathbb{Z}^{r \times m}$ with $m \geq r + 2$ and let G be an abelian group. Then the homomorphism Ψ in (17) is a (t, m, k)-representation for (M, G), with $k = k_0 + k_1 + \cdots + k_r$.*

Thus, the coordinates of an element of G_*^t with indices in $\left[\sum_{j=0}^{i-1} t_j + 1\ ,\ \sum_{j=0}^{i} t_j \right]$ are used by $\Psi^{(i)}$ to represent $M^{(i)}$, and these representations are then combined by K to make Ψ a representation for (M, G).

To prove Proposition 4.20, we first record an equivalent definition of a hypergraph representation, where condition (iii) from Definition 3.7 is replaced with a variant that is convenient for our arguments below.

Definition 4.21 (Hypergraph representation, equivalent formulation). Let $M \in \mathbb{Z}^{r \times m}$ and let G be an abelian group. A (t, m, k)-*representation* of the system (M, G) is a homomorphism $\Psi : G_*^t \to G^m$, for some abelian group G_*, such that the following conditions hold:

(i) There are distinct sets $C_1, C_2, \ldots, C_m \in \binom{[t]}{k}$ such that $\forall\, j \in [m]$, $\mathrm{Supp}\, \psi_j \subseteq C_j$.
(ii) $\Psi(G_*^t) = \ker_G M$.
(iii') For every $j \in [m]$, for every $x = (x_1, \ldots, x_m) \in \ker_G M$, if $g' \in G_*^{C_j}$ satisfies $\psi_{C_j}(g') = x_j$, then there exists $g \in G_*^t$ such that $\Psi(g) = x$ and $p_{C_j}(g) = g'$.

When G is a *compact* abelian topological group, we require that G_* also be compact, and that Ψ be continuous.

Lemma 4.22. *Definitions 3.7 and 4.21 are equivalent.*

Proof. Let us recall condition (iii) from Definition 3.7:

(iii) For each $j \in [m]$, we have $p_{C_j}(\ker_{G_*} \Psi) = \ker_{G_*} \psi_{C_j}$.

To see that (iii') implies (iii), note that by definition we have $p_{C_j}(\ker_{G_*} \Psi) \subseteq \ker_{G_*} \psi_{C_j}$, and that the opposite containment also holds, by (iii') applied with $x = 0_{G^m}$.

To see that (iii) implies (iii'), suppose that $x \in \ker_G M$ and $g' \in G_*^{C_j}$ satisfy $\psi_{C_j}(g') = x_j$. Note that by condition (ii) we have $x = \Psi(g_0)$ for some $g_0 \in G_*^t$. Then

$\psi_{C_j}(g') = \psi_{C_j}\big(p_{C_j}(g_0)\big) = x_j$, so $g' - p_{C_j}(g_0) \in \ker_{G_*} \psi_{C_j}$. Therefore, by (iii), there exists $g_1 \in \ker_{G_*} \Psi$ such that $p_{C_j}(g_1) = g' - p_{C_j}(g_0)$. Letting $g = g_0 + g_1$, we thus have $p_{C_j}(g) = g'$ and $\Psi(g) = \Psi(g_0) = x$, so (iii') holds. $\square$

We now turn to the proof of Proposition 4.20. A central fact that we shall use is that K ignores the $(m+1)$-st coordinates of each $\mathbf{x}^{(i)}$, $i \in [r]$, indeed this provides the additional degrees of freedom sufficient for establishing the required properties of Ψ.

Let $C_{j,i} \subseteq [t_i]$ denote the support of the j-th row of $\Psi^{(i)}$. Then the support of the j-th row of Ψ is
$$C_j = C_j(0) \sqcup C_j(1) \sqcup \ldots \sqcup C_j(r) \subseteq [t],$$
where $C_j(i)$ is the shifted set $(t_0 + t_1 + \cdots + t_{i-1}) + C_{j,i} \subseteq [t]$.

It is then clear that condition (i) from Definition 4.21 is inherited by Ψ from the $\Psi^{(i)}$.

We prove the other two conditions separately.

Lemma 4.23. *The map Ψ in* (17) *satisfies $\Psi(G_*^t) = \ker_G M$.*

Proof. We first check that $\Psi(G_*^t) \subseteq \ker_G M$. We have to ensure that for any $g \in G_*^t$, for each $j \in [r]$ we have $M_j(\Psi(g)) = 0$. Note that $\Psi(g) = K(\mathbf{x})$ for $\mathbf{x} = \Phi(g) \in \prod_{i=0}^r \ker_{G_i} M^{(i)}$. Let $\mathbf{x}'^{(i)} = p_{[m]}(\mathbf{x}^{(i)})$, for each $i \in [r]$. Observe that, letting P' denote the map on G^r that multiplies the i-th coordinate by $\gcd(B_i)$ for each $i \in [r]$, we have $M \circ P = P' \circ M^{(0)}$. It follows that for each $j \in [r]$ we have

$$M_j(\Psi(g)) = M_j \circ P(\Sigma\mathbf{x}) = \gcd(B_j) M_j^{(0)}(\Sigma\mathbf{x}) = \gcd(B_j)\Big(M_j^{(0)}\mathbf{x}^{(0)} + \sum_{i=1}^r M_j^{(0)}\mathbf{x}'^{(i)}\Big).$$

Here we have firstly that $M_j^{(0)}\mathbf{x}^{(0)} = 0_G$, since $\mathbf{x}^{(0)} \in \ker_G M^{(0)}$ by definition of Φ. For $i \in [r] \setminus \{j\}$, since $M_j^{(i)}$ restricted to $[m]$ equals $M_j^{(0)}$, and the $(m+1)$-st component of $M_j^{(i)}$ is 0, we have $M_j^{(0)}\mathbf{x}'^{(i)} = M_j^{(i)}\mathbf{x}^{(i)} = 0_G$. Finally, for $i = j$, we have $\gcd(B_j) M_j^{(0)}\mathbf{x}'^{(j)} = 0_G$, since $\mathbf{x}'^{(j)}$ has coordinates in G_j. We have thus shown that $M_j(\Psi(g)) = 0_G$ for each $j \in [r]$, hence $\Psi(g) \in \ker_G M$.

We now show that $\Psi(G_*^t) \supseteq \ker_G M$. Let $\mathbf{y} = (y_1, y_2, \ldots, y_m) \in \ker_G M$. For each $i \in [r]$, let $x_i = -B_i'(y_{r+1}, \ldots, y_m) \in G$. Then the element
$$\mathbf{x}^{(0)} := (x_1, \ldots, x_r, y_{r+1}, \ldots, y_m)$$
satisfies $M^{(0)}\mathbf{x}^{(0)} = 0$. Setting $\mathbf{x} = (\mathbf{x}^{(0)}, 0_{G^{m+1}}, 0_{G^{m+1}}, \ldots, 0_{G^{m+1}})$, we have $K(\mathbf{x}) = P(\mathbf{x}^{(0)}) = \mathbf{y}$. The map Φ is surjective (since each map $\Psi^{(i)}$ is), so there exists $g \in G_*^t$ such that $\Phi(g) = \mathbf{x}$, and so $\Psi(g) = K \circ \Phi(g) = \mathbf{y}$ as required. $\square$

We now check condition (iii).

Lemma 4.24. *The map Ψ in* (17) *satisfies condition* (iii) *from Definition 3.7.*

Proof. Fix any $j \in [m]$. From the definitions, it is clear that $\ker_{G_*} \psi_{C_j} \supseteq p_{C_j}(\ker_{G_*} \Psi)$. To prove the opposite containment, we shall use the fact that, by Lemma 4.22, condition (iii') holds for each map $\Psi^{(i)}$.

Recall the notations: $C_{j,i} \subseteq [t_i]$ is the support of the j-th row of $\Psi^{(i)}$, and $C_j(i) \subseteq [t]$ for each i is such that $C_j = \bigsqcup_{i=0}^r C_j(i)$ is the support of the j-th row of Ψ.

Suppose that $g' \in G_*^{C_j}$ is given such that $\psi_{C_j}(g') = 0$. We want to find $g \in \ker_{G_*} \Psi$ such that $p_{C_j}(g) = g'$. We identify the groups $G_*^{C_{j,i}}$ and $G_*^{C_j(i)}$ the obvious way via the order-preserving bijection $C_{j,i} \to C_j(i)$.

For each $i \in [0, r]$, let $g'^{(i)} \in G_*^{C_{j,i}}$ be the element such that we can identify g' as

$$g' = \left(g'^{(0)}, \ldots, g'^{(r)} \right) \in G_*^{C_{j,0}} \times G_*^{C_{j,1}} \times \cdots \times G_*^{C_{j,r}}.$$

Our task is to show that there exists $g \in G_*^t$ such $\Psi(g) = 0$ and $p_{C_j(i)}(g) = g'^{(i)}$ for each i (modulo the above identification).

<u>Case 1</u>: $j \in [r]$.

For each $i \in [r] \setminus \{j\}$, let $\mathbf{x}^{(i)}$ be an element of $\ker_{G_i} M^{(i)}$ having the right j-th coordinate, i.e. such that the j-th row of $\Psi^{(i)}$, denoted $\psi_{C_{j,i}}^{(i)}$, satisfies $\psi_{C_{j,i}}^{(i)}(g'^{(i)}) = \mathbf{x}_j^{(i)}$. Note that such a solution $\mathbf{x}^{(i)}$ can be obtained just by extending $g'^{(i)}$ arbitrarily to an element $g_0^{(i)}$ of $G_*^{t_i}$, for example by adding 0 coordinates; indeed we then have that $\Psi^{(i)}(g_0^{(i)}) = \mathbf{x}^{(i)}$ lies in $\ker_{G_i} M^{(i)}$, as $\Psi^{(i)}$ satisfies condition (ii).

Define

$$x_{\text{def}} := -\psi_{C_{j,0}}^{(0)}\big(g'^{(0)}\big) - \sum_{i \in [r] \setminus \{j\}} \mathbf{x}_j^{(i)} = - \sum_{i \in [0,r] \setminus \{j\}} \psi_{C_{j,i}}^{(i)}(g'^{(i)}). \tag{18}$$

Note the important fact that $x_{\text{def}} \in G_j$. Indeed, letting $d_j = \gcd(B_j)$, and using that $d_j \psi_{C_{j,j}}^{(j)}(g'^{(j)}) = 0$ (since $d_j G_j = \{0_G\}$), we have

$$d_j x_{\text{def}} = - \sum_{i \in [0,r] \setminus \{j\}} d_j \psi_{C_{j,i}}^{(i)}(g'^{(i)}) = - \sum_{i \in [0,r]} d_j \psi_{C_{j,i}}^{(i)}(g'^{(i)}) = - \psi_{C_j}(g') = 0.$$

We can therefore find $\mathbf{x}'^{(j)} \in \ker_{G_j} M^{(0)}$ such that $\mathbf{x}'^{(j)}_j = -x_{\text{def}}$ (using the fact that the rows of B' in $M^{(0)}$ are coprime).

Define $\mathbf{x}^{(j)}$ to be the element of $\ker_{G_j} M^{(j)}$ that restricts to $\mathbf{x}'^{(j)}$ on its first m coordinates. (Thus $\mathbf{x}^{(j)}$ is $\mathbf{x}'^{(j)}$ with an extra $(m+1)$-st coordinate equal to x_{def}.)

For each $i \in [r] \setminus \{j\}$, since $\Psi^{(i)}$ is a representation, by condition (iii') from Definition 4.21 we can extend $g'^{(i)}$ to some $g^{(i)} \in G_i^{t_i}$ such that $\Psi^{(i)}(g^{(i)}) = \mathbf{x}^{(i)}$.

Now we want to extend $g'^{(j)}$ to an element $g^{(j)}$ such that $\Psi^{(j)}(g^{(j)})$ agrees with $\mathbf{x}^{(j)}$ at each of its first m coordinates except perhaps the j-th one; equivalently, we want $(\Psi^{(j)}(g^{(j)}))_u = \mathbf{x}_u^{(j)}$ for each $u \in [r+1, m]$. (This implies equality also for $u \in [r] \setminus \{j\}$ since $\Psi^{(j)}(g^{(j)})$ and $\mathbf{x}^{(j)}$ are both in $\ker_{G_j} M^{(j)}$.)

We can find this element $g^{(j)}$ thanks to the freedom in the $(m+1)$-st variable in the system $(M^{(j)}, G_j)$. In other words, we are using condition (iii') for $\Psi^{(j)}$ to extend $g'^{(j)}$, but we are doing so with target-solution the element of $\ker_{G_j} M^{(j)}$ that has j-th coordinate $\psi_{C_{j,j}}^{(j)}(g'^{(j)})$ and uth coordinate $\mathbf{x}_u^{(j)}$ for $u \in [m] \setminus \{j\}$, and we are using the freedom in the $(m+1)$-st coordinate to claim that such a target-solution exists.

We finally come to extending $g'^{(0)}$, and to do so we first have to choose an appropriate element $\mathbf{x}^{(0)} \in \ker_{G_0} M^{(0)}$. Consider the element $\mathbf{x}^{(0)}$ such that the restriction $\mathbf{x}^{(0)}|_{[r+1,m]}$ of $\mathbf{x}^{(0)}$ to coordinates indexed in $[r+1, m]$ satisfies

$$\mathbf{x}^{(0)}|_{[r+1,m]} = - \sum_{i \in [r]} \mathbf{x}^{(i)}|_{[r+1,m]}, \tag{19}$$

for the $\mathbf{x}^{(i)}$ defined above.

The key claim now is that this solution $\mathbf{x}^{(0)} \in \ker_{G_0} M^{(0)}$, determined by (19), satisfies $\mathbf{x}_j^{(0)} = \psi_{C_{j,0}}^{(0)}(g'^{(0)})$. If this holds then we may use condition (iii') to obtain the desired extension $g^{(0)}$ of $g'^{(0)}$ such that $\Psi^{(0)}(g^{(0)}) = \mathbf{x}^{(0)}$.

To prove the claim, note that on one hand by (18) we have

$$\psi_{C_{j,0}}^{(0)}(g'^{(0)}) = -x_{\text{def}} - \sum_{i \in [r] \setminus \{j\}} \mathbf{x}_j^{(i)}. \tag{20}$$

On the other hand, letting $B_j^{(0)}$ denote the restriction of the row $M_j^{(0)}$ to the entries indexed by $[r+1, m]$, we deduce from $M_j^{(0)}(\mathbf{x}^{(0)}) = 0$ that

$$\mathbf{x}_j^{(0)} = -B_j^{(0)}\mathbf{x}^{(0)}|_{[r+1,m]} = \sum_{i \in [r]} B_j^{(0)}\mathbf{x}^{(i)}|_{[r+1,m]} = \sum_{i \in [r]} B_j^{(0)}(\mathbf{x}_{r+1}^{(i)}, \ldots, \mathbf{x}_m^{(i)}).$$

Here the summand with index $i = j$ is $\mathbf{x}_j^{(j)} = -x_{\text{def}}$, and for each $i \in [r] \setminus \{j\}$ the i-th summand is $B_j^{(i)}(\mathbf{x}_{r+1}^{(i)}, \ldots, \mathbf{x}_m^{(i)}) = -\mathbf{x}_j^{(i)}$, since $M^{(i)}$ has same j-th row as $M^{(0)}$. Hence

$$\mathbf{x}_j^{(0)} = -x_{\text{def}} - \sum_{i \in [r] \setminus \{j\}} \mathbf{x}_j^{(i)}. \tag{21}$$

Combining (20) and (21) we deduce that $\psi_{C_{j,0}}^{(0)}(g'^{(0)}) = \mathbf{x}_j^{(0)}$ as claimed.

We have thus obtained $g = (g^{(0)}, g^{(1)}, \ldots, g^{(r)})$ such that $p_{C_j}(g) = g'$, and such that for each $u \in [r+1, m]$ we have $\Psi_u(g) = \sum_{i \in [0,r]} \Psi_u^{(i)}(g^{(i)}) = \sum_{i \in [0,r]} \mathbf{x}_u^{(i)} = 0$, by definition of $\mathbf{x}^{(0)}$. Since an element of $\ker_G M$ is determined by its last $m-r$ coordinates, we must have $\Psi(g) = 0$. This completes Case 1.

Case 2: $j \in [r+1, m]$.

The argument is similar but simpler. Suppose that we are given

$$g' = (g'^{(0)}, \ldots, g'^{(r)}) \in G_*^{C_{j,0}} \times \cdots \times G_*^{C_{j,r}}$$

such that $\sum_{i \in [0,r]} \psi_{C_{j,i}}^{(i)}(g'^{(i)}) = 0$. For each $i \in [r]$, let $g^{(i)} \in G_*^{t_i}$ be any element satisfying $p_{C_{j,i}}(g^{(i)}) = g'^{(i)}$ (e.g. obtained by extending $g'^{(i)}$ by 0-coordinates) and let $\mathbf{x}^{(i)} = \Psi^{(i)}(g^{(i)}) \in \ker_{G_i} M^{(i)}$. (Note that $\psi_{C_{j,i}}^{(i)}(g'^{(i)}) = \mathbf{x}_j^{(i)}$.) For each $i \in [r]$, we define

$$\mathbf{y}^{(i)} = \left(-\mathbf{x}_1^{(i)}, \ldots, -\mathbf{x}_{i-1}^{(i)}, B_i'(\mathbf{x}_{r+1}^{(i)}, \ldots, \mathbf{x}_m^{(i)}), -\mathbf{x}_{i+1}^{(i)}, \ldots, -\mathbf{x}_m^{(i)}\right).$$

Note that $\mathbf{y}^{(i)} \in \ker_G M^{(0)}$. Therefore, the element $\mathbf{x}^{(0)} := \sum_{i \in [r]} \mathbf{y}^{(i)}$ lies in $\ker_G M^{(0)}$. We also have

$$\mathbf{x}_j^{(0)} = -\sum_{i \in [r]} \mathbf{x}_j^{(i)} = -\sum_{i \in [r]} \psi_{C_{j,i}}^{(i)}(g'^{(i)}) = \psi_{C_{j,0}}^{(0)}(g'^{(0)}),$$

where the first equality follows from the definition of the $\mathbf{y}^{(i)}$, the second from the definition of the $\mathbf{x}^{(i)}$, and the third equality follows by assumption on g'. By condition (iii'), there exists $g^{(0)}$ extending $g'^{(0)}$ such that $\Psi^{(0)}(g^{(0)}) = \mathbf{x}^{(0)}$. The element $g = (g^{(0)}, \ldots, g^{(r)})$ that we have thus obtained extends g', and for each $u \in [r+1, m]$ we have $\Psi_u(g) = \sum_{i \in [0,r]} \Psi_u^{(i)}(g^{(i)}) = \sum_{i \in [0,r]} \mathbf{x}_u^{(i)} = 0$, so just like in the previous case we must have $\Psi(g) = 0$. This completes Case 2. $\square$

The proof of Proposition 4.20 is now complete.

It remains only to address the second case described at the beginning of this subsection, namely that of matrices of the form $(I_r|B) \in \mathbb{Z}^{r\times(r+1)}$. We shall do so by using a matrix extension slightly different from those used in earlier sections. This replaces the given matrix by an $(r+1) \times (r+3)$ matrix. We show that this extension also conserves hypergraph-representability, so that we can then just apply Proposition 4.20 to the latter matrix.

Lemma 4.25. *Let $M = (I_r|B) \in \mathbb{Z}^{r\times(r+1)}$. Let*

$$M' = \left(\quad I_{r+1} \quad \left| \begin{array}{cc} B & \mathbf{0} \\ 0 & -1 \end{array} \right. \right) \in \mathbb{Z}^{(r+1)\times(r+3)}.$$

Suppose that Ψ' is a $(t, r+3, k)$-representation for M', and let $J = [r+2] \setminus \{r+1\}$. Then $\Psi := \pi_J \circ \Psi'$ is a $(t, r+1, k)$-representation for M.

Proof. Condition (i) from Definition 3.7 is clearly inherited by Ψ from Ψ'.

For condition (ii), note that $\Psi'(G_*^t) = \ker_G M'$ by assumption. Thus for any $x' \in \Psi'(G_*^t)$ we have $M'x' = 0_{G^{r+3}}$, which implies by construction that $x := \pi_J(x')$ satisfies $Mx = 0_{G^{r+1}}$. Hence $\Psi(G_*^t) \subseteq \ker_G M$. To see the opposite containment, let $x \in \ker_G M$, and note that $x' := (x_1, x_2, \ldots, x_r, 0, x_{r+1}, 0)$ lies in $\ker_G M'$, so there exists $g \in G_*^t$ such that $\Psi'(g) = x'$ and so $\Psi(g) = \pi_J(x') = x$.

To check condition (iii), fix any $j \in [r+1]$ and note that by assumption we have $p_{C'_{\sigma_J(j)}}(\ker_{G_*} \Psi') = \ker_{G_*} \psi'_{C'_{\sigma_J(j)}}$. By construction, the j-th row of Ψ has support $C_j = C'_{\sigma_J(j)}$, where the latter is the support of the $\sigma_J(j)$th row of Ψ'. We also have that the corresponding maps $\psi_{C_j}, \psi'_{C'_{\sigma_J(j)}}$ are equal, whence $p_{C_j}(\ker_{G_*} \Psi') = \ker_{G_*} \psi_{C_j}$. Therefore it suffices to check that $p_{C_j}(\ker_{G_*} \Psi') = p_{C_j}(\ker_{G_*} \Psi)$.

The rows of Ψ form a subset of those of Ψ', so we certainly have $\ker_{G_*} \Psi' \subseteq \ker_{G_*} \Psi$ and so $p_{C_j}(\ker_{G_*} \Psi') \subseteq p_{C_j}(\ker_{G_*} \Psi)$. For the opposite containment, suppose that $g' \in G_*^{C_j}$ equals $p_{C_j}(g_0)$ for some $g_0 \in \ker_{G_*} \Psi$, so in particular $\psi_{C_j}(g') = 0$. Applying (iii') from Definition 4.21 to Ψ', with $x = 0_{G^{r+3}}$ and g' satisfying $\psi'_{C'_{\sigma(j)}}(g') = x_{\sigma(j)} = 0$, we obtain that there exists $g \in G_*^t$ such that $\Psi'(g) = 0_{G^{r+3}}$ and $p_{C_j}(g) = p_{C'_{\sigma(j)}}(g) = g'$. $\square$

5. Remarks

There are several ways in which one could try to extend Theorem 1.3 further.

To begin with, one may want to remove the assumption $d_r(M) = 1$. To achieve this, the arguments in this paper would have to be modified in a non-trivial way, especially those in Section 4, starting with Lemma 4.7, and including the proofs of Lemmas 4.23 and 4.24.

One may also want to bring Theorem 1.3 more in line with the $\mathbb{Z}_p$ version (Theorem 1.1) by making sure that the parameter δ depends only on the dimensions m, r of the matrix M and not on the entries themselves. Note that the function δ in Theorem 1.3 is currently not guaranteed to be independent of the entries of M, because of the argument involving Euclid's algorithm in the proof of Lemma 4.15. One would therefore need at least to modify Lemma 4.15.

Thus, obtaining the above two improvements of Theorem 1.3 via our approach requires handling several technical difficulties of a purely algebraic nature, and we have therefore preferred not to pursue these matters in this paper.

One may also seek extensions of these removal results to noncommutative settings. It seems plausible, for instance, that there is an analogue of Theorem 1.3 for nilmanifolds. One possible such result would say, roughly speaking, that if a product of measurable subsets of a nilmanifold G/Γ has an intersection of small-measure with the so-called 'Leibman nilmanifold' associated with a system of linear forms (see [11, §3]), then this intersection can be eliminated by removing small-measure subsets from the given sets.

Appendix A. Reduction of the main theorem

In this appendix we show that Theorem 3.1 implies Theorem 1.3. In fact, as we shall see, it is not hard to establish the following stronger reduction.

Proposition A.1. *If Theorem 3.1 holds for every compact abelian Lie group, then Theorem 1.3 holds.*

As is well-known, every compact Hausdorff abelian group G is a strict projective limit of compact abelian Lie groups (see [14, Rem. 2.35, Cor. 2.43]). We shall use this to prove Proposition A.1, by approximating the given Borel sets $A_j \subseteq G$ in Theorem 1.3 by Borel subsets coming from a Lie quotient of G. More precisely, we use the following approximation result.

Lemma A.2. *Let G be a compact abelian group, let A be a Borel subset of G, and let $0 < \delta < 1$. There exists a compact abelian Lie group G_0, a continuous surjective homomorphism $q : G \to G_0$, and a Borel set $A_0 \subseteq G_0$, such that $\mu_G\big(A \,\Delta\, (q^{-1}A_0)\big) < \delta$.*

Proof. By Lusin's theorem there exists a continuous function h on G with $\|h\|_{L^\infty(G)} \le 1$ such that $\|h - 1_A\|_{L^1(G)} < \delta^3/2^{10}$; see [21, Appendix E8]. By the Stone-Weierstrass theorem, the trigonometric polynomials are dense in the set of continuous functions on G, relative to the $L^\infty(G)$-norm; see [21, p. 24]. Thus there exists a trigonometric polynomial $P(x)$ such that $\|h - P\|_{L^\infty(G)} < \delta^3/2^{10}$, whence $\|1_A - P\|_{L^1(G)} < \delta^3/2^9$. We also have $\|P\|_{L^\infty(G)} < \|h\|_{L^\infty(G)} + \delta^3/2^{10} < 2$, and by taking real parts we can also suppose that P is real-valued.

Let $\widehat{G}$ be the dual group of G and let $\widehat{G_0}$ be the subgroup of $\widehat{G}$ generated by the spectrum of P, i.e. by the finite set $\{\gamma \in \widehat{G} : \widehat{P}(\gamma) \ne 0\}$. Then $\widehat{G_0}$ is a finitely generated (discrete) abelian group, and is thus the dual of a compact abelian Lie group G_0. Letting Λ denote the annihilator of $\widehat{G_0}$ (Λ is a closed subgroup of G), we have that G_0 is isomorphic as a compact abelian group to G/Λ (see [21, §2.1]), and so the quotient map $G \to G/\Lambda$ gives a continuous surjective homomorphism $q : G \to G_0$. There exists a trigonometric polynomial P_0 on G_0 with $P = P_0 \circ q$. We then have $\|P_0\|_{L^\infty(G_0)} \le 2$. Moreover,

$$\left\|P_0 - P_0^2\right\|_{L^1(G_0)} = \left\|P - P^2\right\|_{L^1(G)} \le \int_G |1_A - P| \,\mathrm{d}\mu_G + \int_G |1_A - P| \,|1_A + P| \,\mathrm{d}\mu_G < \delta^3/2^7.$$

This implies that the set $D = \{x \in G_0 : |P_0(x) - P_0^2(x)| > \delta^2/2^4\}$ has measure at most $\delta/8$. On the complement $D^c = G \setminus D$, we must have $|P_0(x)| \le \delta/4$ or $|1 - P_0(x)| \le \delta/4$.

Now let $A_0 = \{x \in G_0 : |P_0(x) - 1| \leq \delta/4\}$. We have that $\|1_{A_0} - P_0\|_{L^1(G_0)}$ is at most

$$3 \int_{G_0} 1_D \, \mathrm{d}\mu_{G_0} + \int_{G_0} 1_{A_0 \cap D^c}(x)|1 - P_0(x)| + 1_{A_0^c \cap D^c}(x)|P_0(x)| \, \mathrm{d}\mu_{G_0} < 7\delta/8.$$

Hence $\mu_G\big(A \triangle (q^{-1}A_0)\big) \leq \|1_A - P\|_{L^1(G)} + \|P - 1_{A_0} \circ q\|_{L^1(G)} = \|1_A - P\|_{L^1(G)} + \|P_0 - 1_{A_0}\|_{L^1(G_0)} < \delta$. $\qquad\square$

By iterating the main argument in this proof we can simultaneously approximate any finite number of Borel sets $A_1, A_2, \ldots, A_m \subseteq G$, that is we can find a single Lie group G_0 in which there are Borel sets $A_{j,0}$ such that $\mu_G\big(A_j \triangle (q^{-1}A_{j,0})\big) < \delta$ for every $j \in [m]$.

We shall also use the following basic fact.

Lemma A.3. *Let $M \in \mathbb{Z}^{r \times m}$ satisfy $d_r(M) = 1$. Then for any abelian group G, and any surjective homomorphism $\theta : G \to H$, the homomorphism $\theta^m : G^m \to H^m, x \mapsto (\theta(x_1), \ldots, \theta(x_m))$ is surjective from $\ker_G M$ to $\ker_H M$.*

Proof. The homomorphism $M : G^m \to G^r$ is surjective for any abelian group G; this is immediate from the Smith normal form $M = U\left(I_r | 0^{r \times (m-r)}\right) V$, where $U \in \mathbb{Z}^{r \times r}, V \in \mathbb{Z}^{m \times m}$ are unimodular matrices. On $\ker_G M$, the homomorphism θ^m takes values in $\ker_H M$. Now given $x_H \in \ker_H M$, there exists $x \in G^m$ such that $\theta^m(x) = x_H$. We have $0 = M(x_H) = M(\theta^m(x)) = \theta^r(M(x))$, so $M(x) \in (\ker \theta)^r \leq G^r$. By surjectivity of M, there exists $z \in (\ker \theta)^m$ such that $M(z) = M(x)$. Thus $x - z$ is an element of $\ker_G M$ satisfying $\theta^m(x - z) = x_H$, so θ^m is indeed onto $\ker_H M$. $\qquad\square$

Finally, we shall also use the fact that the integral of bounded functions across a kernel $\ker_G M$ can be controlled in terms of their $L^1(G)$ norms, in the following sense.

Lemma A.4. *Let $M \in \mathbb{Z}^{r \times m}$, and let G be a compact abelian group. For $j \in [m]$ let $p_j : G^m \to G$ denote the projection homomorphism to the j-th component, and let G_j denote the closed subgroup $p_j(\ker_G M) \leq G$. Suppose that each G_j has finite index $\kappa_j = |G : G_j|$ in G. Then, for any measurable functions $f_1, \ldots, f_m : G \to \mathbb{C}$ with $\|f_j\|_{L^\infty(G)} \leq 1$ for all j, we have*

$$\left| \int_{\ker_G M} f_1(x_1) \cdots f_m(x_m) \, \mathrm{d}\mu_{\ker_G M}(x) \right| \leq \min_{j \in [m]} \kappa_j \, \|f_j 1_{G_j}\|_{L_1(G)}. \tag{22}$$

Proof. Fix $j \in [m]$. By the triangle inequality and the bounds $\|f_j\|_{L^\infty(G)} \leq 1$, the left side of (22) is at most

$$\int_{\ker_G M} |f_j(x_j)| \, \mathrm{d}\mu_{\ker_G M}(x) = \int_{\ker_G M} |f_j \circ p_j(x)| \, \mathrm{d}\mu_{\ker_G M}(x) = \int_{G_j} |f_j(y)| \, \mathrm{d}\mu_{G_j}(y).$$

On the other hand, by the quotient integral formula [7, Theorem 1.5.2], we have

$$\int_G |f_j(x) 1_{G_j}(x)| \, \mathrm{d}\mu_G(x) = \kappa_j^{-1} \sum_{z \in G/G_j} \int_{G_j} |f_j(y+z)| 1_{G_j}(y+z) \, \mathrm{d}\mu_{G_j}(y)$$

$$= \kappa_j^{-1} \int_{G_j} |f_j(y)| \, \mathrm{d}\mu_{G_j}(y).$$

The result follows. $\qquad\square$

Proof of Proposition A.1. Given $M \in \mathbb{Z}^{r \times m}$ with $d_r(M) = 1$, let $\delta' > 0$ be such that Theorem 3.1 holds for any compact abelian Lie group with initial parameter $\epsilon/2$, and let $\delta = \min(\delta'/2, \epsilon/2)$. Suppose that $A_1, \ldots, A_m$ are Borel subsets of a compact abelian group G satisfying $\mu_{\ker_G M}(A_1 \times \cdots \times A_m \cap \ker_G M) \leq \delta$. Note that we may assume without loss of generality that each closed subgroup $G_j = p_j(\ker_G M)$ has positive measure in G, for otherwise the conclusion of Theorem 1.3 holds just by removing the null-set $G_j \cap A_j$. It follows that each G_j has finite index κ_j in G. Indeed, by [13, Corollary 20.17], there is a non-empty open set U contained in the difference set $G_j - G_j = G_j$, and by compactness G can be covered by finitely many translates of U and therefore of G_j.

It follows from Lemma A.2 (or rather a repeated application of its proof) that there exists a compact abelian Lie group G_0, a continuous surjective homomorphism $q : G \to G_0$, and a Borel set $A_{j,0} \subseteq G_0$ for each $j \in [m]$, such that $\mu_G\big(A_j \, \Delta \, (q^{-1} A_{j,0})\big) \leq \delta/\kappa_j m$ for every $j \in [m]$. We claim that we therefore have

$$\mu_{\ker_{G_0} M}(A_{1,0} \times \cdots \times A_{m,0} \cap \ker_{G_0} M) \leq \delta'.$$

Indeed, firstly by Lemma A.3 the homomorphism $q^m : \ker_G M \to \ker_{G_0} M$ is surjective, whence

$$\mu_{\ker_{G_0} M}\big(A_{1,0} \times \cdots \times A_{m,0} \cap \ker_{G_0} M\big) = \mu_{\ker_G M}\big(q^{-1} A_{1,0} \times \cdots \times q^{-1} A_{m,0} \cap \ker_G M\big).$$

Now, the map $(f_1, \ldots, f_m) \mapsto \int_{\ker_G M} f_1(x_1) \cdots f_m(x_m) \, d\mu_{\ker_G M}(x)$ is multilinear (for measurable functions f_i), and reduces to $\mu_{\ker_G M}(A_1 \times \cdots \times A_m \cap \ker_G M)$ when $f_i = 1_{A_i}$. A multilinearity argument (similar to that in the proof of Lemma 2.4) gives us

$$\mu_{\ker_G M}\big(\prod_{j \in [m]} q^{-1} A_{j,0} \cap \ker_G M\big) = \sum_{j \in [m+1]} \int_{\ker_G M} f_{1,j}(x_1) \cdots f_{m,j}(x_m) \, d\mu_{\ker_G M}(x),$$

where for each $j \in [m+1]$ we have $f_{i,j} = 1_{A_i}$ if $i < j$, we have $f_{j,j} = 1_{q^{-1} A_{j,0}} - 1_{A_j}$, and $f_{i,j} = 1_{q^{-1} A_{i,0}}$ if $i > j$. By Lemma A.4, it follows that $\mu_{\ker_G M}(q^{-1} A_{1,0} \times \cdots \times q^{-1} A_{m,0} \cap \ker_G M)$ is at most

$$\mu_{\ker_G M}(A_1 \times \cdots \times A_m \cap \ker_G M) + \sum_{j \in [m]} \kappa_j \, \mu_G\big(A_j \, \Delta \, (q^{-1} A_{j,0})\big) \leq \delta',$$

as we claimed.

We now apply Theorem 3.1 on G_0, obtaining sets $R_{j,0}$ of measure at most $\epsilon/2$ such that $\prod_{j \in [m]} A_{j,0} \setminus R_{j,0}$ is M-free. Then $\prod_{j \in [m]} q^{-1}(A_{j,0}) \setminus q^{-1}(R_{j,0})$ is M-free, whence, setting $R_j = q^{-1}(R_{j,0}) \cup (A_j \setminus q^{-1}(A_{j,0}))$, we are done. $\square$

Remark A.5. One can obtain a version of Lemma A.2 in which the approximating group G_0 is just second countable, arguing along the following lines. By Plancherel's theorem, the Fourier transform of 1_A is square-summable, hence supported on a countable subset of the Pontryagin dual $\widehat{G}$, hence supported on a countable subgroup of $\widehat{G}$. Taking duals then yields an approximation of A, up to a null set, given by a subset A_0 of a second countable quotient of G.

References

[1] T. Bloom, *A quantitative improvement for Roth's theorem on arithmetic progressions*, preprint. arXiv:1405.5800

[2] V. Bogachev, *Measure theory* Vol. II, Springer-Verlag, Berlin, 2007.

[3] P. Candela, *Developments at the interface between combinatorics and Fourier analysis*, PhD thesis, University of Cambridge, 2009.

[4] P. Candela, O. Sisask, *On the maximal density of a set avoiding linear equations modulo a prime*, Acta Math. Hungar. **132** (2011) (3), 223–243. arXiv:1109.2934

[5] P. Candela, O. Sisask, *A removal lemma for linear configurations in subsets of the circle*, Proc. Edinburgh Math. Soc. **56** (2013) (3), 657–666. arXiv:1109.2939

[6] D. Conlon, J. Fox, *Graph removal lemmas*, Surveys in Combinatorics, Cambridge University Press, 2013, 1–50. arXiv:1211.3487

[7] A. Deitmar, S. Echterhoff, *Principles of harmonic analysis*, (Springer, 2009).

[8] W. T. Gowers, *A new proof of Szemerédi's theorem*, GAFA **11** (2001), 465–588.

[9] W. T. Gowers, *Quasirandom groups*, Combin. Probab. Comput. **17** no. 3 (2008), 363–387.

[10] B. Green, *A Szemerédi-type regularity lemma in abelian groups, with applications*, GAFA **15** no. 2 (2005), 340–376. arXiv:math/0310476

[11] B. Green, T. Tao, *An arithmetic regularity lemma, an associated counting lemma, and applications*, An irregular mind, 261–334, Bolyai Soc. Math. Stud., 21, Budapest, 2010. arXiv:1002.2028

[12] B. Green, T. Tao, *New bounds for Szemerédi's theorem. II. A new bound for $r_4(N)$*, Analytic number theory, 180–204, Cambridge Univ. Press, Cambridge, 2009. arXiv:math/0610604

[13] E. Hewitt, K. Ross, *Abstract harmonic analysis*, vol. 1, 2nd edition, (Springer, 1979).

[14] K. H. Hofmann, S.A. Morris, *The structure of compact groups*, Berlin: de Gruyter (1998).

[15] S. Kechris, *Classical descriptive set theory*, Graduate Texts in Mathematics, **156**, Springer-Verlag, New York, 1995.

[16] D. Král', O. Serra, L. Vena, *A combinatorial proof of the removal lemma for groups*, J. Combin. Theory Ser. A **116** (2009), 971–978. arXiv:0804.4847

[17] D. Král', O. Serra, L. Vena, *A removal lemma for systems of linear equations over finite fields*, Israel J. Math. **187** (2012), 193–207. arXiv:0809.1846

[18] D. Král', O. Serra, L. Vena, *On the removal lemma for linear systems over abelian groups*, European Journal of Combinatorics **34** (2013), no. 2, 248–259. arXiv:1106.4243

[19] M. Newman, *Integral matrices*, Academic Press, New York, 1972. Pure and Applied Mathematics, Vol. 45.

[20] W. Rudin, *Real and complex analysis*, McGraw-Hill, New York, 1987.

[21] W. Rudin, *Fourier analysis on groups*, Interscience Tracts in Pure and Applied Mathematics, no. 12, Interscience Publishers, New York-London 1962.

[22] I. Z. Ruzsa, *Solving a linear equation in a set of integers, II*, Acta Arith. **72** (1995), 385–397.

[23] I. Z. Ruzsa, E. Szemerédi, *Triple systems with no six points carrying three triangles*, Combinatorics (Proc. Fifth Hungarian Colloq., Keszthely, 1976), Vol. II, pp. 939–945, Colloq. Math. Soc. János Bolyai, 18 (1978) North-Holland, Amsterdam-New York.

[24] T. Sanders, *On Roth's theorem on progressions*, Ann. of Math. **174** (2011) (2) , no. 1, 619–636. arXiv:1011.0104

[25] D. Saxton, A. Thomason, *Hypergraph containers*, preprint. arXiv:1204.6595

[26] O. Serra, L. Vena, *On the number of monochromatic solutions of integer linear systems on abelian groups*, European J. Combin. **35** (2014), 459–473. arXiv:1203.2383

[27] A. Shapira, *A proof of Green's conjecture regarding the removal properties of sets of linear equations*, J. London Math. Soc. **81** (2010), 355–373. arXiv:0807.4901

[28] Szemerédi, *On sets of integers containing no k elements in arithmetic progression*, Acta Arithmetica **27** (1975) 199–245.

[29] B. Szegedy, *The symmetry preserving removal lemma*, Proc. Amer. Math. Soc. **138** (2010), no. 2, 405–408. arXiv:0809.2626

[30] T. Tao, *A variant of the hypergraph removal lemma*, J. Combin. Thy. A **113** (2006), 1257–1280. arXiv:math/0503572

[31] T. Tao, *A proof of Roth's theorem*, blog post `http://terrytao.wordpress.com/2014/04/24/a-proof-of-roths-t`

[32] V.S. Varadarajan, *Groups of automorphisms of Borel spaces*, Trans. Amer. Math. Soc. **109** (1963), 191–220.

[33] L. Vena. The removal lemma for products of systems. Midsummer Combinatorial Workshop 2012, KAM-DIMATIA Series. Available at `http://kam.mff.cuni.cz/workshops/work18/mcw2012booklet.pdf`

ALFRÉD RÉNYI MATHEMATICAL RESEARCH INSTITUTE
BUDAPEST, HUNGARY

E-mail address: `candela83@gmail.com`

ALFRÉD RÉNYI MATHEMATICAL RESEARCH INSTITUTE
BUDAPEST, HUNGARY

E-mail address: `szegedyb@gmail.com`

DEPARTMENT OF MATHEMATICS
UNIVERSITY OF TORONTO, TORONTO, CANADA

E-mail address: `lluis.vena@utoronto.ca`