Sparse Partition Regularity

Imre Leader*† Paul A. Russell*‡

June 3, 2005

Abstract

Our aim in this paper is to prove Deuber's conjecture on sparse partition regularity, that for every m, p and c there exists a subset of the natural numbers whose (m, p, c)-sets have high girth and chromatic number. More precisely, we show that for any m, p, c, k and g there is a subset S of the natural numbers that is sufficiently rich in (m, p, c)-sets that whenever S is k-coloured there is a monochromatic (m, p, c)-set, yet is so sparse that its (m, p, c)-sets do not form any cycles of length less than g.

Our main tools are some extensions of Nešetřil-Rödl amalgamation and a Ramsey theorem of Bergelson, Hindman and Leader. As a sideline, we obtain a Ramsey theorem for products of trees that may be of independent interest.

1 Introduction

The notion of a 'sparse' Ramsey theorem goes back to Erdős. The starting case of Ramsey's theorem is the assertion that whenever the edges of a complete graph on six points (a K_6) are 2-coloured, there is a monochromatic triangle. It is natural to ask the converse: if G is a graph such that whenever its edges are 2-coloured there is a monochromatic triangle, then must G contain a K_6? Some simple examples show that the answer is no: in fact, it is easy to construct such a graph G that does not even contain a K_5. The question of whether or not such a graph must contain a K_4 remained open for some time (it was a question of Erdős), until Folkman [4] answered it in the negative. This was extended by Nešetřil and Rödl [13] to any (finite) number of colours, using their important 'amalgamation' method.

This result is a typical 'sparse' theorem: it says that a graph can be very non-dense in edges and yet still have enough edges to have the required Ramsey property. (More precisely, it is a 'restricted' Ramsey theorem; the 'sparse' version, which was also proved by Nešetřil and Rödl, asserts that there is such a graph G whose triangles are so spread out that there is no short cycle of them. Here, as usual, a *cycle* in a hypergraph is a sequence A_1, x_1, A_2, x_2, …, A_n, x_n where A_1, A_2, …, A_n are edges of the hypergraph and x_1, x_2, …, x_n are

*Department of Pure Mathematics and Mathematical Statistics, Centre for Mathematical Sciences, Wilberforce Road, Cambridge CB3 0WB, England.

†I.Leader@dpmms.cam.ac.uk

‡P.A.Russell@dpmms.cam.ac.uk

vertices of the hypergraph, with all of the A_i and all of the x_i distinct, and satisfying $x_i \in A_i \cap A_{i+1}$ for all i, where A_{n+1} means A_1. In this case, we consider the hypergraph whose vertices are the edges of G and whose edges are the triangles of G.) Spencer [20] observed that a restricted version of van der Waerden's theorem holds—he gave a set that is Ramsey for arithmetic progressions of length m, yet contains no arithmetic progression of length $m+1$. Rödl [18] proved a *sparse* version of van der Waerden's theorem—here the set of arithmetic progressions of length m in the set is not allowed to have any short cycles. For other approaches, see Frankl, Graham and Rödl [5], and Prömel and Voigt [16]. Around the same time, Nešetřil and Rödl [14] proved a sparse version of the Finite Sums theorem. (The *Finite Sums theorem* states that for any n, whenever the positive integers are finitely coloured there exist positive integers $x_1, x_2, \ldots, x_n$ such that all of the sums $\sum_{i \in A} x_i$ for non-empty $A \subset [n]$ have the same colour.) All of these results were placed in a more general context, and greatly extended, by Nešetřil and Rödl (see for example [15]).

However, despite these results, one conjecture remained open: a sparse version of Rado's theorem, which we now describe.

A (finite) matrix A over the integers is said to be *partition regular (PR)* if whenever the positive integers are finitely coloured, there exists a vector x with all entries the same colour and satisfying $Ax = 0$. We may also refer to 'the system of linear equations $Ax = 0$' as being PR. Rado [17] provided a characterization of all PR matrices. We shall be concerned with a version of this characterization given by Deuber [2] in his refinement of Rado's theorem.

To describe this characterization, it is necessary to introduce the notion of an (m, p, c)-set. These can be thought of as 'iterated arithmetic progressions'. Given positive integers m, p and c, and positive integers $x_1, x_2, \ldots, x_m$, the *(m, p, c)-set generated by* $x_1, x_2, \ldots, x_m$ is the set consisting of all sums of the form $cx_k + \sum_{i=1}^{k-1} \lambda_i x_i$ for $k = 1, 2, \ldots, m$ and λ_i $(1 \leq i \leq k-1)$ integers with $0 \leq \lambda_i \leq p$. So, in particular, a $(2, p, 1)$-set is simply an arithmetic progression of length $p + 1$ together with its common difference. Deuber [2] proved that a system of linear equations is PR if and only if there exist some m, p and c such that every (m, p, c)-set contains a solution to the system. So (m, p, c)-sets are the 'building-blocks' of partition regularity, and are in a sense the only important PR systems.

Returning to the sparse version of Rado's theorem, Deuber [3] conjectured that for any m, p and c, there should exist a set S such that whenever S is finitely coloured there is a monochromatic (m, p, c)-set, but the (m, p, c)-sets in S form a hypergraph of large girth. Our main aim in this paper is to prove this conjecture.

One of the main obstacles to proving Deuber's conjecture is the fact that there is really no 'abstract' version of Rado's theorem—as opposed to, say, van der Waerden's theorem, where the Hales-Jewett theorem may be viewed as its abstract version. Rather, to prove Rado's theorem one has to *iterate* some Ramsey results (such as the Hales-Jewett theorem). This appears to be the reason why even the very general results of Nešetřil and Rödl [15] do not seem to help with Deuber's conjecture.

We shall discuss more about the ideas of the proof once we have defined the notion of a 'picture' in §3. For the moment, let us say that the proof is largely based on two notions. One is an extension of Nešetřil-Rödl amalgamation. The other is a Ramsey theorem of Bergelson, Hindman and Leader [1]—it was used in [1] to show that for every n, there is a subset S of the positive integers which is partition regular for every PR system of n equations (meaning that for any PR matrix A with n rows, whenever S is finitely coloured there is a monochromatic vector x with $Ax = 0$), but with S not partition regular for some PR system of $n+1$ equations. We shall explain the relevance of this Ramsey theorem when we define a 'picture'; it will be needed as a 'starting point' for all of the later amalgamation, and will provide a framework within which the amalgamation can take place.

We begin by considering the simplest interesting case, namely the triangle-free case for arithmetic progressions with common difference. It is convenient to refer to a $(2, p, 1)$-set (or arithmetic progression of length $p+1$ with common difference) as a *p-line*. So, in other words, we wish to find a subset S of the natural numbers such that whenever S is k-coloured there is a monochromatic p-line, but such that the p-lines in S do not form a triangle. All of the main ideas necessary for the general result are present in the proof of this case, but the notation is less impenetrable. In §2, we present the Ramsey theorem of [1] in this case, and give a proof in our language. As a digression, we explain how the result can be extended to give a Ramsey theorem for products of trees. This result is not required in our proof of the sparse Rado theorem, but may be of independent interest—in the language of Nešetřil and Rödl [12], we show that the class of tree-products has the vertex-Ramsey property. Next, in §3, we carry out the amalgamation to complete the proof of our main result in this case. Finally, in §4, we explain how our methods can be extended to prove the sparse version of Rado's theorem in full generality.

We generally use standard notation throughout the paper. We denote by $\mathbb{N}$ the set $\{0, 1, 2, \dots\}$ of natural numbers, and by $\mathbb{N}_+$ the set $\mathbb{N} - \{0\} = \{1, 2, 3, \dots\}$ of positive integers. For $n \in \mathbb{N}_+$, we write $[n]$ to denote the finite set $\{1, 2, \dots, n\}$.

2 Ramsey results

Our main aim in this section is to present a proof of Theorem 3. This is a result from [1]; we need it here because it will provide a 'framework' within which the later amalgamation can take place. We provide a proof to familiarize the reader with the language we are using—the concepts in the proof will be important later.

Our other aim is to generalize Theorem 3 to a Ramsey theorem for products of trees. While this more general result will play no part in the proof of the sparse Rado theorem (and so the reader should feel free to omit §2.2 if so desired), it may be of independent interest as it shows that the class of tree-products has a natural Ramsey structure.

We make extensive use of the Hales-Jewett theorem, of which we shall remind the reader after some necessary definitions. This theorem can be thought of as an 'abstract' version of van der Waerden's theorem.

Let A be a finite set and d a positive integer. We work in A^d, the *d-dimensional Hales-Jewett cube on alphabet A*. A *combinatorial line* in the cube A^d is a set L of the form

$$L = \{(x_1, x_2, \dots, x_d) \in A^d : x_i = x_j \text{ for } i, j \in I,\ x_i = c_i \text{ for } i \in [d] - I\}$$

where I is a non-empty subset of $[d]$ and the c_i $(i \in [d] - I)$ are elements of the alphabet A. We call I the set of *active coordinates* of L and $[d] - I$ the set of *inactive coordinates* of L. We are now ready to state the Hales-Jewett theorem.

Theorem 1 (The Hales-Jewett theorem [7]). *Let A be a finite set and k a positive integer. Then there exists a positive integer d such that whenever A^d is k-coloured, it contains a monochromatic line.*

Observe that van der Waerden's theorem follows easily from the Hales-Jewett theorem. It is possible to map the Hales-Jewett cube $[n]^d$ into the positive integers $\mathbb{N}_+$ in such a way that each combinatorial line in $[n]^d$ is taken to an arithmetic progression of length n in $\mathbb{N}_+$: for example, define $\phi\colon [n]^d \to \mathbb{N}_+$ by

$$\phi(x_1, x_2, \dots, x_d) = x_1 + x_2 + \cdots + x_d.$$

Then a k-colouring of $\mathbb{N}_+$ induces a k-colouring of $[n]^d$, which, assuming d is sufficiently large, gives a monochromatic line in $[n]^d$, which in turn gives a monochromatic arithmetic progression of length n in $\mathbb{N}_+$.

We shall also require a multi-dimensional extension of this theorem. An *m-dimensional combinatorial subspace* of a Hales-Jewett cube A^d ($m = 1, 2, 3, \dots$) is a set L of the form

$$L = \left\{ (x_1, x_2, \dots, x_d) \in A^d \; : \; \begin{array}{l} \text{for each } j = 1, 2, \dots, m,\ x_i = x_h \text{ for } i, h \in I_j, \\ x_i = c_i \text{ for } i \in [d] - \bigcup_{j=1}^m I_j \end{array} \right\}$$

where $I_1, I_2, \dots, I_m$ are disjoint non-empty subsets of $[d]$ and the c_i $(i \in [d] - \bigcup_{j=1}^m I_j)$ are elements of the alphabet A. We call $I_1, I_2, \dots, I_m$ the *active coordinate sets* of L and $[d] - \bigcup_{j=1}^m I_j$ the set of *inactive coordinates* of L. Note that a 1-dimensional combinatorial subspace is simply a combinatorial line.

Theorem 2 (The multi-dimensional Hales-Jewett theorem [7]). *Let A be a finite set, and let k and m be positive integers. Then there exists a positive integer d such that whenever A^d is k-coloured, it contains a monochromatic m-dimensional subspace.*

While this result does not appear explicitly in [7], it follows immediately from Theorem 1 by applying it to the alphabet A^m (see, for example, [6]).

2.1 Products of trees

We begin by constructing, for each p and k, a particular subset F of the positive integers such that whenever F is k-coloured it contains a monochromatic $(2, p, 1)$-set. When constructing the set F, we try to avoid endowing it with unnecessary structure. In particular, F will have no 2-cycles of $(2, p, 1)$-sets (i.e. any pair of $(2, p, 1)$-sets in F will intersect in at most one point), but it will turn out to contain triangles of $(2, p, 1)$-sets. This sparseness in the structure of F is crucial when we come to use F to index the sets in our amalgamation in §3.

To that end, let p be a positive integer.

A *p-line* is a set of the form $\{a, x, x + a, x + 2a, \dots, x + pa\}$, where a and x are non-zero elements of $\mathbb{N}$, i.e. a $(2, p, 1)$-set. We also consider more general

p-lines, where we allow x and a to be non-zero elements of $\mathbb{N}^e$ for any $e \geq 1$. A set of this form but with a or x allowed to be zero will be called a *p-pseudo-line*. A map $\phi\colon L \to L'$ between p-pseudo-lines L and L' is called a *homomorphism* if there exist integers a, a', x and x' with $L = \{a, x, x+a, x+2a, \ldots, x+pa\}$ and $L' = \{a', x', x'+a', x'+2a', \ldots, x'+pa'\}$ such that $\phi(a) = a'$ and $\phi(x+\lambda a) = x' + \lambda a'$ for $\lambda = 0, 1, \ldots, p$. An *isomorphism* between p-lines is a bijective homomorphism.

We note that there are two isomorphism classes of p-lines: those with $x \neq a$, which have order $p+2$, and those with $x = a$, which have order $p+1$. In general we shall only need to deal with those of the first class; we refer to those in the second class as *degenerate* p-lines, and we shall construct our structures in such a way that they do not arise.

A *p-tree of height 0* is a set of the form $\{a\}$ for some non-zero $a \in \mathbb{N}$. A *p-tree of height 1* is a non-degenerate p-line $T \subset \mathbb{N}$; in other words, $T = \{a, x, x+a, x+2a, \ldots, x+pa\}$ for non-zero a, $x \in \mathbb{N}$ with $a \neq x$. We say that the p-tree $\{a\}$ of height 0 is a *pre-tree* of T. A *p-tree of height h* ($h \geq 2$) is a set T of the form

$$T = R \cup \bigcup_{a \in S-R} \{a, x_a, x_a+a, x_a+2a, \ldots, x_a+pa\}$$

where S is a p-tree of height $h-1$ with pre-tree R and the x_a ($a \in S-R$) are non-zero elements of $\mathbb{N}$ chosen so that all the $x_a + \lambda a$ ($a \in S-R$, $0 \leq \lambda \leq p$) are distinct and not contained in S and so that no unnecessary p-lines are created in T: in other words, the only p-lines in T are those in S together with those of the form $\{a, x_a, x_a+a, x_a+2a, \ldots, x_a+pa\}$ for $a \in S-R$. In particular, T contains no degenerate p-lines. We say that S is a *pre-tree* of T. Note that it is possible to find a p-tree of any given height h: simply select each x_a, in turn, sufficiently large.

If T is a tree of height h, we say that $(T_0, T_1, \ldots, T_h)$ is a *tree-sequence* for T if T_i is a tree of height i ($0 \leq i \leq h$), T_i is a pre-tree of T_{i+1} ($0 \leq i \leq h-1$) and $T_h = T$. (The reader may check that, apart from in certain trivial cases, the tree-sequence for a given tree is unique.)

Two p-trees T and T' are said to be *isomorphic* if there is a bijection $\phi\colon T \to T'$ such that

- for $L \subset T$, $\phi(L)$ is a p-line in T' precisely when L is a p-line in T; and
- for each p-line L in T, the restriction of ϕ to L is an isomorphism (of p-lines).

It is clear that two p-trees are isomorphic precisely when they have the same height. We refer to an isomorphic image of a p-tree T as a *copy* of T.

A *p-tree-product of dimension d* is a set $F \subset \mathbb{N}^d$ of the form

$$F = \{(t_1, t_2, \ldots, t_d) : t_i \in T_i \cup \{0\}, t_i \text{ not all } 0\}$$

where T_1, T_2, $\ldots$, T_d are p-trees. We say that F is the *tree-product* of the trees T_1, T_2, $\ldots$, T_d, and we sometimes write $F = T_1 \diamondsuit T_2 \diamondsuit \cdots \diamondsuit T_d$, or $F = \diamondsuit_{i=1}^d T_i$. A p-tree-product is said to be *of height h* if each p-tree in the definition is of height h. Observe that a p-tree-product of dimension 1 is simply a p-tree.

We define isomorphisms and copies of p-tree-products exactly as for p-trees. In general, we allow the image of an isomorphism to be any subset of $\mathbb{N}^e$ for any $e \geq 1$, and we refer to any copy of a p-tree-product as a p-tree-product. So, in particular, we may encounter p-trees as subsets of $\mathbb{N}^e$ for $e > 1$.

Note that it is possible to embed any p-tree-product of any dimension in $\mathbb{N}$. In other words, given a p-tree product $F = T_1 \diamondsuit T_2 \diamondsuit \cdots \diamondsuit T_d$, there is an isomorphism ϕ from F to some subset of $\mathbb{N}$; for example, we may take the map

$$(t_1, t_2, \cdots, t_d) \mapsto a_1 t_1 + a_2 t_2 + \cdots + a_d t_d,$$

where a_1, a_2, $\ldots$, a_d are positive integers selected, in turn, sufficiently large that no new p-lines are introduced.

Note also that p-tree-products do not contain any degenerate p-lines. For suppose that $\{x, 2x, 3x, \ldots, (p+1)x\}$ is a degenerate p-line in the tree-product $F = T_1 \diamondsuit T_2 \diamondsuit \cdots \diamondsuit T_d$. Choose some i such that x has ith coordinate $x' \neq 0$. Then $\{x', 2x', 3x', \ldots, (p+1)x'\}$ is a degenerate p-line in T_i, a contradiction.

We generally regard p as being a fixed positive integer and suppress the p-dependence in the notation, referring simply to lines, trees, tree-products, etc.

We are now ready to give a proof of the Ramsey theorem of Bergelson, Hindman and Leader. (While this result does not appear explicitly in [1], it may be read out of Theorem 2.5 of [1] using Lemma 2.6 and Theorem 2.7 of [1].)

Theorem 3 ([1]). *Let p and k be positive integers. Then there exists a p-tree-product F such that whenever F is k-coloured, it contains a monochromatic p-line.*

Proof. For the remainder of this proof, 'line', 'tree' and 'tree-product' shall mean 'p-line', 'p-tree' and 'p-tree-product'. (In future, we shall often use such terminology without comment.)

Let T be a tree of height $k+1$ with tree sequence $(T_0, T_1, \ldots, T_{k+1})$. Define a finite sequence d_0, d_1, d_2, $\ldots$, d_{k+1} of positive integers inductively as follows:

- $d_0 = 1$;
- for $1 \leq n \leq k+1$, take d_n sufficiently large that whenever $T_n^{d_n}$ is k-coloured, there exists a monochromatic combinatorial subspace of dimension d_{n-1}.

Note that d_n is guaranteed to exist by the Hales-Jewett theorem [7]. We take F to be the tree-product of d_{k+1} copies of T.

Suppose F is k-coloured. This induces a k-colouring of the subset $T_{k+1}^{d_{k+1}}$ and so, by our choice of d_{k+1}, we may find a monochromatic d_k-dimensional subspace S_k. We may assume without loss of generality that the active coordinates of S are $[d_k]$, i.e. that there exist $z_{d_k+1}, z_{d_k+2}, \ldots, z_{d_{k+1}} \in T_{k+1}$ such that

$$S_k = \{(t_1, t_2, \ldots, t_{d_k}, z_{d_k+1}, z_{d_k+2}, \ldots, z_{d_{k+1}}) : t_1, t_2, \ldots, t_{d_k} \in T_{k+1}\}.$$

[The conscientious reader may be concerned at this point that some of the active coordinate sets of S may contain two or more coordinates varying together. But this does not cause a problem—we may simply identify such coordinates by a suitable isomorphism. This will result in a smaller number of inactive coordinates in the set S, but the number of inactive coordinates has no bearing on the remainder of the proof.]

Now, write

$$F_k = \{(t_1, t_2, \dots, t_{d_k}, \underbrace{0, 0, \dots, 0}_{d_{k+1}-d_k}) : t_1, t_2, \dots, t_{d_k} \in T_k \cup \{0\}, t_i \text{ not all } 0\}.$$

Note that we may think of F_k as a tree-product of height k by considering it as the tree product of d_k copies of T_k; i.e. we identify F_k with the set

$$\{(t_1, t_2, \dots, t_{d_k}) : t_1, t_2, \dots, t_{d_k} \in T_k \cup \{0\}, t_i \text{ not all } 0\}.$$

Now, our original colouring induces a k-colouring of F_k, which in turn gives a k-colouring of the subset $T_k^{d_k}$. By our choice of d_k, we may find a monochromatic d_{k-1}-dimensional subspace S_{k-1}. We may assume that the active coordinate set of S is $[d_{k-1}]$. So there exist $z_{d_{k-1}+1}, z_{d_{k-1}+2}, \dots, z_{d_k} \in T_k$ such that

$$S_{k-1} = \{(t_1, t_2, \dots, t_{d_{k-1}}, z_{d_{k-1}+1}, z_{d_{k-1}+2}, \dots, z_{d_k}) : t_1, t_2, \dots, t_{d_{k-1}} \in T_k\}.$$

Now, write

$$F_{k-1} = \{(t_1, t_2, \dots, t_{d_{k-1}}, \underbrace{0, 0, \dots, 0}_{d_k - d_{k-1}}) : t_1, t_2, \dots, t_{d_{k-1}} \in T_{k-1} \cup \{0\}, t_i \text{ not all } 0\}.$$

Note that we may think of F_{k-1} as a tree-product of height $k-1$ by considering it as the tree product of d_{k-1} copies of T_{k-1}; i.e. we identify F_{k-1} with the set

$$\{(t_1, t_2, \dots, t_{d_{k-1}}) : t_1, t_2, \dots, t_{d_{k-1}} \in T_{k-1} \cup \{0\}, t_i \text{ not all } 0\}.$$

And so we continue. After $k+1$ applications of Hales-Jewett, we have obtained sequences $F_0, F_1, \dots, F_k$ of subsets of F and $z_1, z_2, \dots, z_{d_{k+1}}$ of elements of T satisfying:

- $F_i = \{(t_1, t_2, \dots, t_{d_i}, \underbrace{0, 0, \dots, 0}_{d_{k+1}-d_i}) : t_1, t_2, \dots, t_{d_i} \in T_i \cup \{0\}, t_i \text{ not all } 0\}$;
- $S_i = \{(t_1, \dots, t_{d_i}, z_{d_i+1}, \dots, z_{d_{i+1}}, \underbrace{0, 0, \dots, 0}_{d_{k+1}-d_{i+1}}) : t_1, \dots, t_{d_i} \in T_{i+1}\}$ is monochromatic, with colour c_i, say;
- $z_i \in T_j$ for $i \leq d_j$.

Now, by the pigeonhole principle, some two of the sets $S_0, S_1, \dots, S_k$ must have the same colour; say $c_m = c_n$ for some $0 \leq m < n \leq k$. Choose arbitrarily

$$a = (a_1, a_2, \dots, a_{d_{m+1}}, \underbrace{0, 0, \dots, 0}_{d_{k+1}-d_{m+1}}) \in S_m.$$

Note that for each i, $1 \leq i \leq d_{m+1}$, we have $a_i \in T_{m+1}$ and so there is some x_i such that $x_i, x_i + a_i, x_i + 2a_i, \dots, x_i + pa_i \in T_{m+2} \subset T_{n+1}$. So, choosing $x_{d_{m+1}+1}, x_{d_{m+1}+2}, \dots, x_{d_n} \in T_{n+1}$ arbitrarily, and setting $x_i = z_i$ for $d_n + 1 \leq i \leq d_{n+1}$, we may take

$$x = (x_1, x_2, \dots, x_{d_{n+1}} \underbrace{0, 0, \dots, 0}_{d_{k+1}-d_{n+1}}) \in S_n.$$

We now have $a \in S_m$ and $x, x+a, x+2a, \dots, x+pa \in S_n$, and so the line $\{a, x, x+a, x+2a, \dots, x+pa\}$ is monochromatic with colour $c_m = c_n$. □

Theorem 3 asserts that given any p and k, there exists a tree-product G such that whenever G is k-coloured, it contains a monochromatic p-line. In fact, much more than this is true: given any *tree-product* F and positive integer k, there is some tree-product G such that whenever G is k-coloured, it contains a monochromatic copy of F. It turns out that this result is not needed in the proof of the sparse Rado theorem. Nevertheless, it may be found to be of independent interest as it shows that the class of tree-products has a natural Ramsey structure. The reader who is interested only in the proof of the sparse Rado theorem may safely skip to the beginning of §3.

2.2 A full Ramsey theorem for products of trees

Our aim in this section is to extend Theorem 3 to a full Ramsey theorem for products of trees. The result of §2.1 showed that whenever a sufficiently large tree-product is k-coloured, it contains a monochromatic line. Here, we prove that we can find not only a monochromatic line, but a monochromatic copy of any tree-product we desire: given any tree-product F, there exists a tree-product G such that whenever G is k-coloured, it contains a monochromatic copy of F.

We begin by considering the case where F is a single tree. Note that Theorem 3 is simply this result with $h = 1$. Our proof is a slight extension of the proof of Theorem 3.

Lemma 4. *Let p, h and k be positive integers. Then there exists a p-tree-product F such that whenever F is k-coloured, it contains a monochromatic p-tree of height h.*

Proof. This time, we begin by taking a tree T of height $kh+1$ with tree sequence $(T_0, T_1, \dots, T_{kh+1})$ and a finite sequence d_0, d_1, …, d_{kh+1} defined inductively by:

- $d_0 = 1$;
- for $1 \leq n \leq kh+1$, take d_n sufficiently large that whenever $T_n^{d_n}$ is k-coloured, there exists a monochromatic combinatorial subspace of dimension d_{n-1}.

We take F to be the tree product of d_{kh+1} copies of T.

Proceeding exactly as in the proof of Theorem 3, we apply Hales-Jewett $kh+1$ times to obtain sequences F_0, F_1, …, F_{kh} of subsets of F and z_1, z_2, …, $z_{d_{kh+1}}$ of elements of T satisfying:

- $F_i = \{(t_1, t_2, \dots, t_{d_i}, \underbrace{0, 0, \dots, 0}_{d_{kh+1}-d_i}) : t_1, t_2, \dots, t_{d_i} \in T_i \cup \{0\}, t_i \text{ not all } 0\}$;

- $S_i = \{(t_1, \ldots, t_{d_i}, z_{d_i+1}, \ldots, z_{d_{i+1}}, \underbrace{0, 0, \ldots, 0}_{d_{kh+1}-d_{i+1}}) : t_1, \ldots, t_{d_i} \in T_{i+1}\}$ is monochromatic;
- $z_i \in T_j$ for $i \leq d_j$.

Now, by the pigeonhole principle, we may find $n_0 < n_1 < \cdots < n_h$ such that the sets S_{n_0}, S_{n_1}, ..., S_{n_h} all have the same colour, c, say.

We shall inductively construct trees U_0, U_1, ..., U_h such that

- U_i is of height i and has tree sequence $(U_0, U_1, \ldots, U_i)$;
- $U_0 \subset S_{n_0}$ and $U_i - U_{i-1} \subset S_{n_i}$ for $1 \leq i \leq h$.

In particular, U_h will be a tree of height h contained in the monochromatic set $\bigcup_{j=0}^{h} S_{n_j}$, precisely as required.

First, we take U_0 to be any singleton subset of S_{n_0}.

Now, suppose that we have already constructed U_{i-1} as required for some i, $1 \leq i \leq h$. Suppose $a \in U_{i-1} - U_{i-2}$ (where we interpret U_{-1} to be the empty set). Then $a \in S_{n_{i-1}}$ and so

$$a = (a_1, a_2, \ldots, a_{d_{n_{i-1}+1}}, \underbrace{0, 0, \ldots, 0}_{d_{kh+1}-d_{n_{i-1}+1}})$$

for some $a_1, a_2, \ldots, a_{d_{n_{i-1}+1}} \in T_{n_{i-1}+1}$.

For each $j \leq d_{n_{i-1}+1}$, we have $a_j \in T_{n_{i-1}+1}$ and so there is some x_j such that $x_j, x_j + a_j, x_j + 2a_j, \ldots, x_j + pa_j \in T_{n_{i-1}+2} \subset T_{n_i+1}$. So, choosing $x_{d_{n_{i-1}+1}+1}, x_{d_{n_{i-1}+1}+2}, \ldots, x_{d_{n_i}} \in T_{n_i+1}$ arbitrarily, and setting $x_j = z_j$ for $d_{n_i} + 1 \leq j \leq d_{n_i+1}$, we may take

$$x_a = (x_1, x_2, \ldots, x_{d_{n_{i+1}}}, \underbrace{0, 0, \ldots, 0}_{d_{kh+1}-d_{n_{i+1}}}) \in S_{n_i}.$$

Then set

$$U_i = U_{i-2} \cup \bigcup_{a \in U_{i-1} - U_{i-2}} \{a, x_a, x_a + a, x_a + 2a, \ldots, x_a + pa\}$$

(where we interpret U_{-1} as $\emptyset$). □

We now consider general tree-products $F = \diamondsuit_{i=1}^{d} T_i$. The proof proceeds in two stages. First, by a product argument, we can find a tree-product G such that whenever G is k-coloured, G contains a copy of F in which each of the $2^d - 1$ possible Cartesian products of the form $\times_{i=1}^{d} U_i$ with $U_i = T_i$ or $U_i = \{0\}$ for each i (not all U_i zero) is monochromatic. The result then follows by an application of the Finite Unions theorem, that given any positive integers d and k, there exists some positive integer D such that whenever the non-empty subsets of $[D]$ are k-coloured, there exists a collection $\mathcal{S} = \{S_1, S_2, \ldots, S_d\}$ of pairwise disjoint subsets of $[D]$ with all non-empty unions of sets in $\mathcal{S}$ the same colour.

Theorem 5. *Let p, k, d and h be positive integers. Then there exists a p-tree-product F such that whenever F is k-coloured, it contains a monochromatic p-tree-product of d trees of height h. In particular, given any p-tree-product F, there exists a p-tree-product G such that whenever G is k-coloured, it contains a monochromatic copy of F.*

Proof. It is enough to prove the following claim:

> For all positive integers p, d, k and h, there exists a tree-product F such that whenever F is k-coloured, it contains a tree-product $\diamondsuit_{i=1}^{d} T_i$ of trees of height h with every set of the form $\times_{i=1}^{d} U_i$ ($U_i = T_i$ or $U_i = \{0\}$ for each i, not all U_i zero) is monochromatic.

We can then deduce the theorem immediately by use of the Finite Unions theorem. For suppose that we have proved this claim, and that we are given p, k, d and h. Pick D large enough that whenever $\mathbb{P}([D]) - \emptyset$ is k-coloured, there are disjoint non-empty sets $S_1, S_2, \ldots, S_d \subset [D]$ with all non-empty unions of the S_i having the same colour. Now apply the claim, but with d replaced by D: we obtain a tree-product $\diamondsuit_{i=1}^{D} T'_i$, with each $\times_{i=1}^{D} U_i$ ($U_i = T'_i$ or $U_i = \{0\}$ for each i, not all U_i zero) monochromatic. This induces a k-colouring of $\mathbb{P}([D]) - \emptyset$: we simply give $S \subset [D]$ the colour of the monochromatic set $\times_{i=1}^{D} U_i$ where U_i is taken to be T'_i for $i \in S$ and $\{0\}$ otherwise. So we can now find S_1, S_2, ..., S_d as in the Finite Unions theorem, and set

$$T_i = \bigcup_{t \in T} \{(u_1, u_2, \ldots, u_D) : u_j = t \text{ if } j \in S_i, u_j = 0 \text{ otherwise}\},$$

where T is a tree of height h which we identify with each T'_i by a suitable isomorphism. Then each T_i is a tree of height h, and their tree-product satisfies the conclusion of the claim.

So it remains to prove the claim. We shall do so by induction on d. The case $d = 1$ is simply Lemma 4, so assume $d > 1$ and fix p, k and h.

Let G be a tree-product which satisfies the conclusion of the claim, but with d replaced by $d - 1$ and k replaced by k^2; this is of course possible by the induction hypothesis. Then take a tree-product H such that whenever H is $k^{|G|+1}$-coloured, it contains a monochromatic tree of height h; we may do this by Lemma 4. Define $F = G \diamondsuit H$.

Now suppose that we are given a k-colouring c of F. This induces a $k^{|G|+1}$-colouring of H: simply colour each $t \in H$ by the entire colouring of $(G \cup \{0\}, t)$. By our choice of H, we can find a monochromatic tree T_d of height h for this colouring.

We now have two k-colourings c_0 and c_1 of G, defined by

$$c_0(t) = c(t, 0)$$

and

$$c_1(t) = c(t, t_0),$$

where t_0 is some arbitrary element of T_d; note that c_1 does not depend on the choice of t_0.

Consider the k^2-colouring (c_0, c_1) of G. By our choice of G, there is a tree-product $\diamondsuit_{i=1}^{d-1} T_i$ of trees of height h which satisfies the conclusions of the claim when applied to this colouring, and so also satisfies the conclusions of the claim when applied to the colourings c_0 and c_1 separately.

We now claim that the set $\diamondsuit_{i=1}^{d} T_i$ will do. This is easy to check, as follows. Firstly, the fact that $\diamondsuit_{i=1}^{d-1} T_i$ satisfies the conclusions of the claim for the colouring c_0 means that each of the sets

$$\left(\times_{i=1}^{d-1} S_i\right) \times \{0\} \qquad (S_i = T_i \text{ or } S_i = \{0\} \text{ for each } i, \text{ not all } S_i \text{ zero})$$

is monochromatic; and for the colouring c_1 that each of the sets

$$\left(\times_{i=1}^{d-1} S_i\right) \times T_d \qquad (S_i = T_i \text{ or } S_i = \{0\} \text{ for each } i, \text{ not all } S_i \text{ zero})$$

is monochromatic. Finally, the fact that T_d is monochromatic for our $k^{|G|+1}$-colouring of H gives us that the set

$$\{0\}^{d-1} \times T_d$$

is monochromatic. □

3 The triangle-free extended van der Waerden theorem

We now begin the proof of our main result in the $(2, p, 1)$-set case. We start with the 'triangle-free' version:

Theorem 6. *Let p and k be positive integers. Then there exists a p-tree-product $P \subset \mathbb{N}$ and a subset $S \subset P$ such that*

- *whenever S is k-coloured, it contains a monochromatic p-line; and*
- *S contains no triangle of p-lines.*

This result could be described as a 'triangle-free extended van der Waerden theorem'.

We begin by developing some machinery which will be necessary for the proof. Roughly speaking, we will apply a kind of Nešetřil-Rödl amalgamation. We emphasize that the proof of the triangle-free result presented in this section is entirely self-contained, although for background the reader may wish to see one of the original uses of amalgamation by Nešetřil and Rödl, to construct graphs of large girth and chromatic number [11], and follow-up results of Frankl, Graham and Rödl [5]. We mention that in our proof of the general sparse Rado theorem in §4, we will require some extra machinery in the form of a sparse Hales-Jewett theorem from [18] and [16].

The main idea is that the 'indexing' of the amalgamation (the index-set for a 'picture' as defined below) will be carried out by Theorem 3. This seems to give us the flexibility we need to make the amalgamation work. We remark that if we indexed by (m, p, c)-sets themselves (which would be the 'conventional' approach), there would be no way to control the presence of superfluous p-lines. This is why it is so important to realise that this result from [1] gives us precisely what we need to start the amalgamation.

Now, fix p and k. We begin by fixing a tree-product F_0 such that whenever F_0 is k-coloured, it contains a monochromatic line. (We can of course do this by Theorem 3.) F_0 will be used to index the sets in our amalgamation.

A *picture* S in a tree-product F will consist of disjoint sets $S_v \subset F$ for each $v \in F_0$. The *underlying set* of S is the set $\bigcup_{v \in F_0} S_v \subset F$; we often simply refer to this *set* as S. A line $L = \{a, x, x+a, x+2a, \ldots, x+pa\} \subset S$ will be called a *picture-line* if there exist $b, y \in F_0$ such that $a \in S_b$, $x \in S_y$, $x + a \in S_{y+b}$, $x + 2a \in S_{y+2b}$, $\ldots$, $x + pa \in S_{y+pb}$.

An *isomorphism* between pictures S and S' is a bijection $\phi \colon S \to S'$ such that

- for $L \subset S$, $\phi(L)$ is a picture-line in S' precisely when L is a picture-line in S; and
- whenever L is a picture-line in S, the restriction of ϕ to L is an isomorphism (of lines).

3.1 Starting Picture

To start, find a picture S inside some large tree-product F such that:

- for any line $\{a, x, x+a, x+2a, \ldots, x+pa\}$ in F_0, there exist b and y such that $b \in S_a$, $y \in S_x$, $y + b \in S_{x+a}$, $y + 2b \in S_{x+2a}$, $\ldots$, $y + pb \in S_{x+pa}$— in other words, "for each line in F_0 there is a corresponding line in the picture S";
- the collection of lines in $\bigcup_{v \in F_0} S_v$ contains no triangle; and
- every line in S is a picture-line.

This is of course possible. For each line L in F_0, we choose disjointly a line L' in F, making sure that the union of the L''s contains no lines other than the L''s themselves. We then take our picture S to have underlying set the union of the L''s, with the points assigned appropriately to the S_v's.

Observe that whenever S is k-coloured with each S_v monochromatic, it contains a monochromatic line. For such a k-colouring of S induces a k-colouring of F_0, giving a monochromatic line in F_0, which in turn gives a monochromatic line in S.

3.2 Amalgamation

Now we "try to force the S_v to be monochromatic".

Specifically, given a picture S satisfying the above conditions, and a fixed $u \in F_0$, we want to find a picture S' in some tree-product which still satisfies our conditions, *and* such that whenever S' is k-coloured, it contains a copy of S with S_u monochromatic. This 'amalgamation' is at the heart of the proof; the difficulty lies in somehow finding a way to construct the set S' which preserves the properties we need. Once this has been successfully carried out, we will have proved our result; for we obtain the required tree-product by applying this to each $u \in F_0$ in turn.

We shall require the notion of a homomorphism of tree-products. So suppose F and G are tree-products. A *homomorphism* π from F to G is a function $\pi \colon F \to G$ such that

- whenever $L \subset F$ is a line, $\pi(L) \subset G$ is a pseudo-line; and
- whenever L is a line in F, the restriction of π to L is a homomorphism (of pseudo-lines).

If F is a tree-product, we write $\bar{F} = F \cup \{0\}$. We use the exact same conditions to define homomorphisms from F to $\bar{G}$ and from $\bar{F}$ to $\bar{G}$. When $\pi\colon F \to \bar{G}$ is a homomorphism, we shall, without comment, also denote by π the extension of this homomorphism to the domain $\bar{F}$ given by setting $\pi(0) = 0$.

It is vital to the amalgamation that there is a large supply of such homomorphisms. Specifically, we need the following lemma.

Lemma 7. *Let F be a p-tree-product, and suppose u and v are distinct elements of F. Then there exists a homomorphism $\pi\colon F \to \overline{F \Diamond F}$ such that $\pi(u) = (u, 0)$ but $\pi(v) \neq (v, 0)$.*

Proof. We first deal with the case where F is a single tree T. Let $(T_0, T_1, \ldots, T_h)$ be a tree sequence for T. For each $x \in T$, we define the *height* $h(x)$ of x to be the least i such that $x \in T_i$. If $a \in T$ with $h(a) < h$ then there is a unique line in T of the form $\{a, x, x+a, \ldots, x+pa\}$ with $h(x) = h(a)+1$. We write $x_a = x$ and $L_a = \{a, x_a, x_a + a, \ldots, x_a + pa\}$.

The set of *descendants* of x is defined to be the smallest set D_x such that

- $x \in D_x$; and
- if $y \in D_x$ with $h(y) < h$ then $L_y \subset D_x$.

Note that D_x is itself a tree. The set G_w of *siblings* of w is the collection of all elements of T at the same height as w which lie in a common line with w (including w).

We now proceed to the definition of π. There are two cases to consider:

Case (i): v is not a descendant of u.

For each sibling u' of u, there is a unique isomorphism $\pi_{u'}\colon D_{u'} \to D_u$ taking u' to u. We define $\pi\colon T \to \overline{T \Diamond T}$ by

$$\pi(w) = \begin{cases} (\pi_{u'}(w), 0) & \text{if } w \in D_{u'} \text{ for some } u' \in G_u \\ (0, 0) & \text{otherwise} \end{cases}.$$

Case (ii): v is a descendant of u.

For each sibling v' of v, there is a unique isomorphism $\pi_{v'}\colon D_{v'} \to D_v$ taking v' to v. We define $\pi\colon T \to \overline{T \Diamond T}$ by

$$\pi(w) = \begin{cases} (w, \pi_{v'}(w)) & \text{if } w \in D_{v'} \text{ for some } v' \in G_v \\ (w, 0) & \text{otherwise} \end{cases}.$$

In each case, it is a routine matter to check that π is indeed a homomorphism; and it is obvious that $\pi(u) = (u, 0)$ but that $\pi(v) \neq (v, 0)$.

We now proceed to consider a general tree-product $F = \Diamond_{i=1}^{d} T_i$. Suppose first that u and v differ in some coordinate where they are both supported; without loss of generality, we assume that $u_1 \neq v_1$ with u_1 and v_1 both non-zero. By our result for a single tree, there is some homomorphism $\pi'\colon T_1 \to \overline{T_1 \Diamond T_1}$

with $\pi'(u_1) = (u_1, 0)$ but $\pi'(v_1) \neq (v_1, 0)$. We may then define $\pi\colon F \to \overline{F\Diamond F}$ by

$$\pi(x_1, x_2, \ldots, x_d) = (\pi'_1(x_1), x_2, x_3, \ldots, x_d, \pi'_2(x_1), \underbrace{0, 0, \ldots, 0}_{d-1})$$

where π'_1, $\pi'_2\colon T \to \bar{T}$ are homomorphisms such that $\pi' = (\pi'_1, \pi'_2)$.

Suppose instead that v is supported on some coordinate where u is not, say $v_1 \neq 0 = u_1$. Then we may simply define

$$\pi(x_1, x_2, \ldots, x_d) = (0, x_2, x_3, \ldots, x_d, \underbrace{0, 0, \ldots, 0}_{d}).$$

Finally, if neither of the above holds then we must have u supported on some coordinate where v is not. There must also be some coordinate on which v is supported, and u must agree with v on this coordinate. So we may assume without loss of generality that $u_1 \neq 0$, $v_1 = 0$ and $u_2 = v_2 \neq 0$. Let D_1 and D_2 denote the sets of descendants of u_1 in T_1 and of u_2 in T_2 respectively.

If the height of the tree D_1 does not exceed the height of the tree D_2, then there is a homomorphism $\pi'\colon D_1 \to D_2$ taking u_1 to u_2. Writing D_x for the set of descendants of a point x in T_1, we know that for each sibling u' of u_1 there is an isomorphism $\pi_{u'}\colon D_{u'} \to D_{u_1} = D_1$ taking u' to u_1. So we may extend π' to the whole of T_1 by setting

$$\pi'(x) = \begin{cases} \pi'(\pi_{u'}(x)) & \text{if } x \in D_{u'} \text{ for some } u' \in G_{u_1} \\ 0 & \text{otherwise} \end{cases}.$$

(Note that this agrees with our previous definition on D_{u_1} as the isomorphism $\pi_{u_1}\colon D_{u_1} \to D_{u_1}$ must be the identity.) Then define $\pi\colon F \to \overline{F\Diamond F}$ by

$$\pi(x_1, x_2, \ldots, x_d) = (x_1, \pi'(x_1), x_3, x_4, \ldots, x_d, \underbrace{0, 0, \ldots, 0}_{d}).$$

If instead the height of the tree D_1 *does* exceed the height of the tree D_2, then we may find a homomorphism $\pi'\colon T_2 \to T_1$ taking u_2 to u_1. We can then set

$$\pi(x_1, x_2, \ldots, x_d) = (\pi'(x_2), x_2, x_3, x_4, \ldots, x_d, \underbrace{0, 0, \ldots, 0}_{d}).$$

□

In what follows, it will often be convenient to think of F as being embedded in $\overline{F\Diamond F}$ by means of the injective homomorphism $x \mapsto (x, 0)$. We shall do so without comment; when we do, we shall, by an abuse of notation, refer to $(x, 0) \in \overline{F\Diamond F}$ simply as x. In this language, Lemma 7 says that for any u, $v \in F$ with $u \neq v$ we can find a homomorphism from F to $\overline{F\Diamond F}$ which fixes u and moves v.

We also require the following trivial observation:

Lemma 8. *Let F_0 and F be p-tree-products, and let $u \in F_0$. Then there exists a p-tree-product F' and an injective homomorphism $\theta\colon F \to F'$ such that, given any $c \in F$, there exists a homomorphism $\theta_c\colon F_0 \to \overline{F'}$ with $\theta_c(u) = \theta(c)$.*

Proof. We may assume without loss of generality that F_0 is a product of n copies of a tree T_0 of height h, and that F is a product of N copies of a tree T of height H. Let F' be a product of N copies of a tree T' of height $H+h$. Let $\phi\colon T \to T'$ be an injective homomorphism taking the root of T to the root of T'. Then define $\theta\colon F \to F'$ by $\theta(x_1, x_2, \dots, x_N) = (\phi(x_1), \phi(x_2), \dots, \phi(x_N))$.

Now, suppose that $(u_1, u_2, \dots, u_n) \in F_0$ and $(c_1, c_2, \dots, c_N) \in F$. We may assume without loss of generality that $u_1 \neq 0$. For $1 \leq i \leq N$ and for each $u' \in G_{u_1}$, let $\theta_i^{(u')}\colon D_{u'} \to \overline{T'}$ be a homomorphism taking u' to $\phi(c_i)$. (Note that this is possible as $D_{u'}$ is a tree of height at most h, and $\phi(c_i)$ is either zero or a point at height at most H in a tree of height $H+h$.) Then define $\theta_{c_i}\colon F_0 \to \overline{F'}$ by

$$\theta_{c_i}(x_1, x_2, \dots, x_n) = \begin{cases} (\theta_1^{(u')}(x_1), \dots, \theta_N^{(u')}(x_1)) & \text{if } x_1 \in D_{u'} \text{ for some } u' \in G_{u_1} \\ (\underbrace{0, 0, \dots, 0}_{N}) & \text{otherwise} \end{cases} .$$

□

We now return to the amalgamation. Recall that we are given a picture S, inside some large tree-product F, satisfying the conditions of §3.1, and $u \in F_0$. With F', θ and θ_c ($c \in F$) as given by Lemma 8, we may regard S as a picture inside the tree-product F' (i.e. we replace S with the copy $\theta(S)$). Making this identification, we see that for each $c \in S_u$ we have $\theta_c(u) = c$. Our goal is a picture S' which also satisfies the conditions of §3.1, and such that whenever S' is k-coloured, it contains a copy of S with S_u monochromatic. We call S' the *amalgamation of* S *over* S_u and define it as follows:

First, fix some e such that whenever the Hales-Jewett cube $(S_u)^e$ is k-coloured, it contains a monochromatic S_u-line. (Here, by 'S_u-line', we of course mean a line in the usual sense of the Hales-Jewett theorem. We can find such an e by applying Hales-Jewett on alphabet S_u.) We list all of the S_u-lines in $(S_u)^e$ as $S_u^{(1)}$, $S_u^{(2)}$, …, $S_u^{(D)}$. Note that $(S_u)^e$ lies inside some large tree-product $G = F'^e$.

Next, fix j with $1 \leq j \leq D$. Let $I \subset [e]$ be the set of active coordinates of the line $S_u^{(j)}$ in $(S_u)^e$. Then there exist $c_i \in S_u$ ($i \in [e] - I$) such that

$$S_u^{(j)} = \{x \in G : x_i = c_i \text{ for } i \in [e] - I \text{ and } x_i = x_h \in S_u \text{ for all } i, h \in I\}\,.$$

For each $v \in F_0$, define

$$S_v^{(j)} = \{x \in G : x_i = \theta_{c_i}(v) \text{ for } i \in [e] - I \text{ and } x_i = x_h \in S_v \text{ for all } i, h \in I\}\,.$$

We observe that this is consistent with our original definition of $S_u^{(i)}$.

Now, let π_i ($1 \leq i \leq n$) be all the homomorphisms from F_0 to $\overline{F_0 \diamondsuit F_0}$ which fix u, with π_1 being the identity. Let f_j ($1 \leq j \leq N = n^D$) be the collection of all functions from $[D]$ to $[n]$, with f_1 being the constant function with value 1. We define our goal picture S' in $\overline{F_0 \diamondsuit F_0}^N \diamondsuit G$ by

$$S'_v = \bigcup_{i=1}^{D} (\pi_{f_1(i)}(v), \pi_{f_2(i)}(v), \dots, \pi_{f_N(i)}(v), S_v^{(i)}).$$

The S'_v are indeed disjoint as each S'_v has first coordinate v.

Note in particular that

$$S'_u = \big(u, u, \dots, u, \bigcup_{i=1}^{D} S_u^{(i)}\big).$$

Note also that we have D copies $S^{(1)}$, $S^{(2)}$, $\dots$, $S^{(D)}$ of S in S', with $(S^{(i)})_v = (\pi_{f_1(i)}(v), \pi_{f_2(i)}(v), \dots, \pi_{f_N(i)}(v), S_v^{(i)}) \subset S'_v$. Moreover, these copies intersect only in S'_u. For if $v \neq u$ then, by Lemma 7, we can find some homomorphism π which fixes u and moves v. Now, given i and j with $i \neq j$, there is some coordinate in which $(S^{(i)})_v$ has v and $(S^{(j)})_v$ has $\pi(v) \neq v$.

We must now check that S' does indeed have the properties we claim. Specifically, we need the following lemma.

Lemma 9. *Let S be a picture satisfying the conditions of §3.1, let $u \in F_0$, and let S' be the amalgamation of S over S_u. Then (i) whenever S' is k-coloured, it contains a copy of S with S_u monochromatic; and (ii) S' satisfies the conditions of §3.1.*

We shall use the following result on lines in S':

Lemma 10. *Let S be a picture satisfying the conditions of §3.1, let $u \in F_0$, and let S' be the amalgamation of S over S_u. Then any line in S' is a picture-line and is entirely contained in one of the copies $S^{(i)}$ of S.*

Proof. Let $L = \{a, x, x + a, x + 2a, \dots, x + pa\}$ be a line in S'. Suppose $a \in S'_{a'}$ and $x \in S'_{x'}$. Then a has first coordinate a' and x has first coordinate x'. Hence for $\lambda = 1, 2, \dots, p$, the first coordinate of $x + \lambda a$ is $x' + \lambda a'$, and so $x + \lambda a \in S_{x'+\lambda a'}$. So L is a picture-line.

We now come to the proof of the second assertion of the lemma, beginning with the simpler case where $u \notin L' = \{a', x', x' + a', x' + 2a', \dots, x' + pa'\}$. Suppose that x and a are in different copies of S, say $x \in S^{(i)}$ and $a \in S^{(j)}$ with $i \neq j$. Suppose first that $x + a \notin S^{(j)}$. Then let π be a homomorphism which fixes u and moves a'. By looking at an appropriate column, we find $x' + \pi(a') = x' + a'$, a contradiction. If instead $x + a \in S^{(j)}$ then we may similarly find that, for some homomorphism π' which fixes u and moves x', $\pi'(x') + a' = x' + a'$, again a contradiction.

Hence x and a are in the same copy of S, say x, $a \in S^{(i)}$. Now suppose that for some λ with $1 \leq \lambda \leq p$ we have $x + \lambda a \notin S^{(i)}$. Then for some homomorphism π which fixes u and moves $x' + \lambda a'$ we have $x' + \lambda a' = \pi(x' + \lambda a')$, a contradiction. So in fact $x + \lambda a \in S^{(i)}$ for all λ, which completes the proof.

We must now deal with the case where $u \in L'$. Let $w \in L$ be the unique point of L which lies in S'_u. By a precisely similar argument to the foregoing, we see that there is some copy of S, say $S^{(i)}$, which contains every point of $L - \{w\}$. It is impossible to obtain any better result than this by considering the first N coordinates, as $\pi_j(u) = u$ for every j. Hence we are forced to examine what happens in the final coordinate.

The projection of L onto this final coordinate forms a line L'' in the tree product G. For each $y \in L$, denote the corresponding points in L' and L'' by y' and y'' respectively. So $y \in S'_{y'}$ and y'' is the projection of y onto G. In particular, $u = w'$.

Let $I \subset [e]$ denote the set of active coordinates of the line $S_u^{(i)}$ in S_u^e, and for each $j \in [e] - I$, let $c_j \in S_u$ be the constant for position j in $S_u^{(i)}$ (i.e. such that given any $x \in S_u^{(i)}$ we have $x_j = c_j$ for every $j \in [e] - I$). Now, for each $y \in L - \{w\}$, we know that $y'' \in S_{y'}^{(i)}$. Hence we can find $d_y \in S_{y'}$ such that $y''_j = \theta_{c_j}(y')$ for $j \in [e] - I$ and $y''_j = d_y$ for $j \in I$. Now, choose $d_w \in G$ such that $\{d_y : y \in L\}$ is a line (for example, take $d_w = w''_j$ for an arbitrary $j \in I$). Consider the point W defined by $W_j = c_j = \theta_{c_j}(u)$ for $j \in [e] - I$ and $W_j = d_w$ for $j \in I$. Then $L'' - \{w'\} \cup \{W\}$ is a line in G. But a line in a tree-product is completely determined by specifying all but one of its points, and so $W = w''$. Finally, as $w'' \in S_u$, we have $w''_j \in S_u$ for every $j \in [e]$, and in particular $d_w \in S_u$. Hence $w' \in S_u^{(i)}$ and so $w \in S^{(i)}$ as required, and we are done. □

We proceed to the verification of the desired properties of S':

Proof (of Lemma 9). (i) Suppose that S' is k-coloured; say we have some k-colouring c of S'. This induces a k-colouring c' of $(S_u)^e$ given by $c'(t) = c(u, u, \ldots, u, t)$. By our choice of e, the Hales-Jewett cube $(S_u)^e$ must contain a monochromatic S_u-line, say $S_u^{(i)}$ for some i with $1 \leq i \leq D$. But then $S^{(i)}$ is a copy of S with $S_u^{(i)}$ monochromatic.

(ii) We now verify that S' satisfies the conditions of §3.1. We have already shown in Lemma 10 that every line in S' is a picture line, so there are only two remaining conditions to be checked.

First, we must check the condition that for each line in F_0 there is a corresponding line in the picture S'. But this is easy. For if L is a line in F_0 then there is a corresponding line in S, and S' contains copies $S^{(i)}$ of S with $S_v^{(i)} \subset S'_v$ for each $v \in F_0$.

Finally, we must ensure that S' is triangle-free. So suppose we have some triangle $L_1 x_1 L_2 x_2 L_3 x_3$ in S', i.e. suppose that L_1, L_2 and L_3 are distinct lines and x_1, x_2 and x_3 are distinct points with $x_1 \in L_1 \cap L_2$, $x_2 \in L_2 \cap L_3$ and $x_3 \in L_3 \cap L_1$. By Lemma 10, each of L_1, L_2 and L_3 must be a picture-line and be contained entirely within some copy of S; say $L_1 \subset S^{(i_1)}$, $L_2 \subset S^{(i_2)}$ and $L_3 \subset S^{(i_3)}$. Each copy of S is triangle-free, and so i_1, i_2 and i_3 cannot all be equal. As copies of S intersect only in S'_u, at least two of the points of intersection, say x_1 and x_2, must lie in S'_u. But then L_2 has two points in S'_u, which is impossible as it must be a picture-line and F_0 contains no degenerate lines. So S' is triangle-free as required. □

This is enough to establish our main result:

Proof (of Theorem 6). Take the starting picture of §3.1, and amalgamate over each $u \in F_0$ in turn. By the preceding lemmas, the resulting set S has the properties we require and the theorem is proved. □

4 The general sparse Rado theorem

We have so far been able to show that given any p and k, there is a subset S of the natural numbers such that

- whenever S is k-coloured, it contains a monochromatic $(2, p, 1)$-set ('line'); but

- no triangle is formed by the lines in S.

In fact, the bulk of the work is now behind us and it is relatively easy to complete the proof of the general result. For the moment, we shall stay with the case of $(2, p, 1)$-sets but attempt to force the set S to have arbitrarily large girth.

4.1 Large girth

Recall that we constructed our set S to contain no triangle of lines. In fact, it also contains no 4-cycle of lines.

Clearly our starting picture contains no 4-cycle of lines—no two lines in the starting picture intersect, so it contains no cycles at all.

Let us consider how a 4-cycle of lines could arise while carrying out the amalgamation. Suppose that we have some picture S containing no 4-cycle of lines, but that when we amalgamate over the set S_u the picture S' thus formed does contain such a 4-cycle, say $L_1x_1L_2x_2L_3x_3L_4x_4$. By our earlier work, we know that each of L_1, L_2, L_3 and L_4 must be a picture-line and be contained entirely within some copy of S. But since none of the copies of S contains a 4-cycle and copies of S intersect only in S'_u, the only way that this can happen is if L_1 and L_2 are contained in the same copy of S, say $S^{(i)}$, and also L_3 and L_4 are contained in the same copy of S, say $S^{(j)}$, with $i \neq j$. Then $S^{(i)}$ and $S^{(j)}$ intersect in at least two points, namely x_2 and x_4. This means that the distinct (S_u)-lines $S_u^{(i)}$ and $S_u^{(j)}$ must intersect in at least two points. But this is impossible—distinct lines in a Hales-Jewett cube can intersect in at most one point.

By a similar argument, we shall see (in the proof of Theorem 12) that a g-cycle of lines in S' can only arise if we have a cycle of (S_u)-lines in $(S_u)^e$ of length at most $g/2$. This suggests that what is needed is precisely the sparse Hales-Jewett theorem of Rödl [18] (see also Prömel and Voigt [16]).

Theorem 11 ([18]). *Let A be a finite alphabet with $|A| \geq 3$, and let k and g be positive integers. Then there exists a Hales-Jewett cube A^e and a subset $R \subset A^e$ such that*

- *whenever R is k-coloured it contains a monochromatic line; and*
- *there is no cycle of lines in R of length $\leq g$.*

Armed with this result, we are now able to prove the following theorem.

Theorem 12. *Let p, k and g be positive integers. Then there exists a set $S \subset \mathbb{N}$ such that*

- *whenever S is k-coloured it contains a monochromatic p-line; and*
- *there is no cycle of p-lines in S of length $\leq g$.*

Proof. We carry out precisely the construction of §3 *except* that when defining the amalgamation S' of S over S_u, we replace the Hales-Jewett cube $(S_u)^e$ by some subset R of an appropriate Hales-Jewett cube such that

- whenever R is k-coloured it contains a monochromatic line; and
- there is no cycle of lines in R of length $\leq g/2$.

We know that R of course exists by Theorem 11.

The proof carries through almost identically to the proof of Theorem 6. All we must check is that if S is a picture with no cycle of length $\leq g$ and S' is the (newly-defined) amalgamation of S over S_u then also S' has no cycle of length $\leq g$.

So suppose that $L_1 x_1 L_2 x_2 \dots L_r x_r$ $(r \leq g)$ is a cycle of lines in S'. We know that each L_i must be a picture-line and contained entirely within some copy of S. Moreover, the copies of S intersect only in S'_u. Hence the x_i which belong to S'_u must induce a cycle of (S_u)-lines in R; for if $i < j$ with x_i, $x_j \in S'_u$ but x_{i+1}, x_{i+2}, …, $x_{j-1} \notin S'_u$ then, by our construction, x_i and x_j must lie in the same copy of S. (As each copy of S has no cycle of length $\leq g$, we cannot have all the x_i lying in the same copy of S and so this must involve more than one (S_u)-line; i.e. we genuinely do obtain a cycle.) Moreover, no two consecutive x_i can belong to S'_u as each L_i is a picture line and so intersects S'_u at most once. Hence the cycle of lines we obtain in R has length $\leq g/2$. But this contradicts our choice of R, and so no such cycle of lines in S' can exist, which is enough to establish our theorem. □

4.2 (m, p, c)-sets

We now proceed to generalize the preceding results to general (m, p, c)-sets. Throughout this section, we shall regard m, p and c as fixed positive integers, and refer to an (m, p, c)-set as a *line*. At one point in the proof, we need to assume that p and c are not both 1—but fortunately this case is precisely the sparse Finite Sums theorem of Nešetřil and Rödl [14].

4.2.1 Ramsey results

Our first task is to give an appropriate definition of a tree such that an analogue of Theorem 3 holds. The following definitions and result are essentially a paraphrase of work of Bergelson, Hindman and Leader [1].

Recall that now a line is an (m, p, c)-set generated by, say, non-zero x_1, x_2, …, $x_m \in \mathbb{N}$. We shall find it convenient to denote this set by $[x_1, x_2, \dots, x_m]$. As before, we also consider more generally lines of the form $[x_1, x_2, \dots, x_m]$ for non-zero x_1, x_2, …, $x_m \in \mathbb{N}^e$ for any $e \geq 1$. We define a *pseudo-line* in exactly the same way, but allowing any or all of the x_i to be zero. A map $\phi\colon L \to L'$ between pseudo-lines L and L' is called a *homomorphism* if we can find x_1, x_2, …, x_m such that $L = [x_1, x_2, \dots, x_m]$ and $L' = [\frac{1}{c}\phi(cx_1), \frac{1}{c}\phi(cx_2), \dots, \frac{1}{c}\phi(cx_m)]$. Again, an isomorphism between lines is a bijective homomorphism.

We now define an *(m, p, c)-tree of height 0* to be a set of the form $\{cx_1\}$ for some non-zero $x_1 \in \mathbb{N}$. An *(m, p, c)-tree of height 1* is an (m, p, c)-set $T = [x_1, x_2, \dots, x_m] \subset \mathbb{N}$, with the property that all of the sums $cx_k + \sum_{i=1}^{k-1} \lambda_i x_i$ $(1 \leq k \leq m$ and $0 \leq \lambda_i \leq p$ for $1 \leq i \leq k-1)$ are distinct; we say that the tree $\{cx_1\}$ of height 0 is a *pre-tree* of T. An *(m, p, c)-tree of height h* $(h > 0)$ is a set T of the form

$$T = R \cup \bigcup_{ca \in S - R} [a, x_2^{(a)}, \dots, x_m^{(a)}]$$

where S is an (m, p, c)-tree of height $h - 1$ with pre-tree R and the $x_i^{(a)}$ $(ca \in S - R,\ 2 \leq i \leq m)$ are non-zero elements of $\mathbb{N}$ chosen so as to create no unnecessary lines in T and so that all of the new sums required are distinct and do not already appear in S. So, in particular, our (m, p, c)-trees contain no degenerate lines.

Observe that a $(2, p, 1)$-tree is simply a p-tree in the sense of §2.1.

To ensure that short cycles cannot be created by the amalgamation process described below, it will be necessary to work with tree-products containing no degenerate lines. Unless $p = 1$, this can be achieved by insisting that no tree in the product contains both x and $2x$ for any $x \in \mathbb{N}_+$. (It is, of course, possible to construct such an (m, p, c)-tree of any given height, simply by taking each new variable sufficiently large.) For if $[x_1, x_2, \ldots, x_m]$ is a line in such a tree-product then $x_2, x_3, \ldots, x_m$ must all be supported on any coordinate where x_1 is supported, and so a degenerate line in the product gives rise to a degenerate line in some single coordinate. If $p = 1$ then $c \neq 1$ and we may instead take trees which do not contain x and cx for any $x \in \mathbb{N}_+$. Hence we may assume that none of the tree-products with which we work contains a degenerate line.

Again, it is clear that (m, p, c)-trees of arbitrary height exist in $\mathbb{N}$: we can construct an (m, p, c)-tree of any given height in $\mathbb{Q}_+$ by successively choosing elements sufficiently large to ensure that there are no unnecessary lines, then multiply up by an appropriate constant to make each element an integer. When m, p and c are fixed, we shall often refer to (m, p, c)-trees simply as trees.

The remainder of our definitions are precisely as in §2.

We are now ready to prove an analogue of Theorem 3. The proof is a minor adaptation of that of the $(2, p, 1)$-set result.

Theorem 13 ([1]). *Let m, p, c and k be positive integers. Then there exists an (m, p, c)-tree-product F such that whenever F is k-coloured, it contains a monochromatic (m, p, c)-set.*

Proof. Let T be a tree of height $K = (m-1)k + 1$ with tree sequence $(T_0, T_1, \ldots, T_K)$. Define a finite sequence $d_0, d_1, \ldots, d_K$ inductively, exactly as before, using the Hales-Jewett theorem:

- take $d_0 = 1$;
- for $1 \leq n \leq K$, take d_n sufficiently large that whenever $T_n^{d_n}$ is k-coloured, there exists a monochromatic combinatorial subspace of dimension d_{n-1}.

Take F to be the tree product of d_K copies of T and assume that we are given a k-colouring of F.

As in the proof of Theorem 3, after K applications of Hales-Jewett, we obtain sequences $F_0, F_1, \ldots, F_{K-1}$ of subsets of F and $z_1, z_2, \ldots, z_{d_K}$ of elements of T satisfying:

- $F_i = \{(t_1, t_2, \ldots, t_{d_i}, \underbrace{0, 0, \ldots, 0}_{d_K - d_i}) : t_1, t_2, \ldots, t_{d_i} \in T_i \cup \{0\}, t_i \text{ not all } 0\}$;
- $S_i = \{(t_1, \ldots, t_{d_i}, z_{d_i+1}, \ldots, z_{d_{i+1}}, \underbrace{0, 0, \ldots, 0}_{d_K - d_{i+1}}) : t_1, t_2, \ldots, t_{d_i} \in T_{i+1}\}$ is monochromatic, with colour c_i, say;
- $z_i \in T_j$ for $i \leq d_j$.

Now, as $K = (m-1)k + 1$, we can find some m of the sets S_0, S_1, S_2, …, S_{K-1} with the same colour; say $c_{n_1} = c_{n_2} = \cdots = c_{n_m}$ for some $0 \leq n_1 < n_2 < \cdots < n_m \leq K-1$. Choose arbitrarily

$$cx_1 = (cx_1^{(1)}, cx_1^{(2)}, \ldots, cx_1^{(d_{n_1+1})}, \underbrace{0, 0, \ldots, 0}_{d_K - d_{n_1+1}}) \in S_{n_1}.$$

Now, for each i, $1 \leq i \leq d_{n_1+1}$, we have $cx_1^{(i)} \in T_{n_1+1}$ and so there exist $x_2^{(i)}$, $x_3^{(i)}$, …, $x_m^{(i)}$ such that $[x_1^{(i)}, x_2^{(i)}, \ldots, x_m^{(i)}] \subset T_{n_1+2} \subset T_{n_j+1}$ for all $j > 1$. So, choosing $cx_2^{(d_{n_1+1}+1)}$, $cx_2^{(d_{n_1+1}+2)}$, …, $cx_2^{(d_{n_2})} \in T_{n_2+1}$ arbitrarily, and setting $cx_2^{(i)} = z_i$ for $d_{n_2} + 1 \leq i \leq d_{n_2+1}$, we may take

$$x_2 = (x_2^{(1)}, x_2^{(2)}, \ldots, x_2^{(d_{n_2+1})}, \underbrace{0, 0, \ldots, 0}_{d_K - d_{n_2+1}}).$$

Observe that $cx_2 + \lambda x_1 \in S_{n_2}$ for $0 \leq \lambda \leq p$.

Now, for each i with $d_{n_1+1} + 1 \leq i \leq d_{n_2+1}$, we have $cx_2^{(i)} \in T_{n_2+1}$ and so there exist $x_3^{(i)}$, $x_4^{(i)}$, …, $x_m^{(i)}$, y such that $[x_2^{(i)}, x_3^{(i)}, \ldots, x_m^{(i)}, y] \subset T_{n_2+2} \subset T_{n_j+1}$ for all $j > 2$. [Note that we are not interested in y; indeed, our interest in the line $[x_2^{(i)}, x_3^{(i)}, \ldots, x_m^{(i)}, y]$ only goes so far as to ensure that the appropriate linear combinations of $x_2^{(i)}$, $x_3^{(i)}$, …, $x_m^{(i)}$ are in all the T_{n_j+1} for $j > 2$.] So, choosing $cx_3^{(d_{n_2+1}+1)}$, $cx_3^{(d_{n_2+1}+2)}$, …, $cx_3^{(d_{n_3})} \in T_{n_3+1}$ arbitrarily, and setting $cx_3^{(i)} = z_i$ for $d_{n_3} + 1 \leq i \leq d_{n_3+1}$, we may take

$$x_3 = (x_3^{(1)}, x_3^{(2)}, \ldots, x_3^{(d_{n_3+1})}, \underbrace{0, 0, \ldots, 0}_{d_K - d_{n_3+1}}).$$

Observe that $cx_3 + \lambda x_2 + \mu x_1 \in S_{n_3}$ for $0 \leq \lambda$, $\mu \leq p$.

Continuing in this way, we define x_4, x_5, …, x_m in such a way that for all i with $1 \leq i \leq m$, we have $cx_i + \sum_{j=1}^{i-1} \lambda_j x_j \in S_{n_i}$ for all $\lambda_1, \lambda_2, \ldots, \lambda_{i-1}$ with $0 \leq \lambda_1, \lambda_2, \ldots, \lambda_{i-1} \leq p$. Hence the line $[x_1, x_2, \ldots, x_m]$ is monochromatic. □

We observe that this again generalizes to give a Ramsey theorem for products of trees, which again will not be necessary for the argument to follow:

Theorem 14. *Let m, p, c and k be positive integers, and let F be an (m, p, c)-tree-product. Then there exists an (m, p, c)-tree-product G such that whenever G is k-coloured, it contains a monochromatic copy of F.*

Proof. This result is deduced from Theorem 13 in exactly the same way as Theorem 5 was deduced from Theorem 3. □

4.2.2 Main Result

We now come to the proof of our main result for (m, p, c)-sets. Again, the proof is almost identical to that for $(2, p, 1)$-sets.

The basic definitions are essentially the same as those given at the start of §3. The only definition deserving of comment is that of a *picture-line:* we now say that a line L in a picture S indexed by a tree-product F_0 is a picture-line if there is some line L' in F_0 and an isomorphism $\phi\colon L \to L'$ such that for all $x \in L$ we have $x \in S_{\phi(x)}$. Note that in the case $m = 2$ and $c = 1$ this reduces to the definition given in §3.

As before, we take a starting picture S inside some large tree-product F. The conditions that we now require S to satisfy are:

- for any line L in F_0, there is a line L' in S and an isomorphism $\phi\colon L \to L'$ such that for all $x \in L$, $\phi(x) \in S_x$;
- the collection of lines in $\bigcup_{v \in F_0} S_v$ contains no cycle of length $\leq g$; and
- every line in S is a picture-line.

We define homomorphisms as before. To carry out our amalgamation, we will need an analogue of Lemma 7 to tell us that there are enough homomorphisms. However, we now need a larger codomain available for our homomorphisms.

We define the *extension* F^+ of a tree-product F to be the tree-product obtained from F by 'growing each tree in F one extra level'. Formally, if $F = \Diamond_{i=1}^n T_i$ where T_i is a tree of height h_i, we define $F^+ = \Diamond_{i=1}^n T_i^+$, where T_i^+ is a tree of height $h_i + 1$; we identify F with a suitable subset of F^+ via an isomorphism mapping each T_i into T_i^+ with the root of T_i being taken to the root of T_i^+.

Lemma 15. *Let F be an (m,p,c)-tree-product, and suppose u and v are distinct elements of F. Then there exists a homomorphism $\pi\colon F \to \overline{F^+ \Diamond F^+}$ such that $\pi(u) = (u,0)$ but $\pi(v) \neq (v,0)$.*

Proof. The proof is similar to that of Lemma 7, but with some additional complications. Again, we start with the case where F is a single tree T with tree-sequence $(T_0, T_1, \ldots, T_h)$. Again, for each $x \in T$, we define the *height* $h(x)$ of x to be the least i such that $x \in T_i$. If $a \in T$ with $h(a) < h$ then there is a unique line L in T of the form $L = [a, x_2, x_3, \ldots, x_m]$ with $h(x_2) = h(x_3) = \cdots = h(x_m) = h(a) + 1$. We write $L_a = L$ and $x_i^{(a)} = x_i$ for $2 \leq i \leq m$. We define the set of descendants D_x and the set of siblings G_x of a point $x \in T$ as in the proof of Lemma 7. We shall denote the set of descendants of a point $x \in T^+$ by D_x^+. Given $x \in T$ with $h(x) > 0$, there is a unique $x_1 \in T$ with $x \in L_{x_1} = [x_1, x_2, \ldots, x_m]$ and $h(x_1) = h(x) - 1$. We may then express x uniquely in the form $x = cx_r + \sum_{i=1}^{r-1} \lambda_i x_i$ for some r, $2 \leq r \leq m$ and $0 \leq \lambda_1, \lambda_2, \ldots, \lambda_{r-1} \leq p$. We call r the *rank* of x and denote it by $r(x)$. We define the set of *younger siblings* of x to be the set $G_x^- = \{y \in G_x : r(y) \geq r(x)\}$.

The definition of π is again divided into two cases.

Case (i): v is not a descendant of u.

As v is not a descendant of u, we must have $h(u) > 0$, and $h \in L_{x_1} = [x_1, x_2, \ldots, x_m]$ for some x_1 with $h(x_1) = h(u) - 1$. Write $r = r(u)$, and $L_u = [y_1, y_2, \ldots, y_m] \subset T^+$ (so $y_1 = u$). For each $t \in G_u^-$, we may write $t = cx_{r(t)} + \sum_{i=1}^{r(t)-1} \mu_i x_i$ for some $0 \leq \mu_1, \mu_2, \ldots, \mu_{r(t)-1} \leq p$. Define $t' = cy_{r(t)-r+1} + \sum_{i=1}^{r(t)-r} \mu_{i+r-1} y_i \in T^+$. Then there is an isomorphism $\pi_t\colon D_t \to D_{t'}^+$ taking t to t'. We define $\pi\colon T \to \overline{T^+ \Diamond T^+}$ by

$$\pi(w) = \begin{cases} (\pi_t(w), 0) & \text{if } w \in D_t \text{ for some } t \in G_u^- \\ (0,0) & \text{otherwise} \end{cases} .$$

Case (ii): v is a descendant of u.

Similarly to the above, by taking for each $t \in G_v^-$ a suitable $t' \in T^+$ and an isomorphism $\pi_t \colon D_t \to D_{t'}^+$, we may define a homomorphism $\pi \colon T \to \overline{T^+ \Diamond T^+}$ by

$$\pi(w) = \begin{cases} (w, \pi_t(w)) & \text{if } w \in D_t \text{ for some } t \in G_v^- \\ (w, 0) & \text{otherwise} \end{cases}.$$

We have now dealt with the case where F is a single tree. Suppose instead $F = \Diamond_{i=1}^d T_i$. If u and v differ in some coordinate where they are both supported, or if v is supported somewhere where u is not, then we can finish exactly as in the proof of Lemma 7. In the only remaining case, we may assume without loss of generality that $u_1 \neq 0$, $v_1 = 0$ and $u_2 = v_2 \neq 0$. Denote by D_1 and D_2 the sets of desendants of u_1 in T_1 and u_2 in T_2 respectively.

Suppose that the height of the tree D_1 does not exceed the height of the tree D_2. Then, similarly to the constructions earlier in this proof, for each $t \in G_{u_1}^-$, we may choose a point $t' \in D_2^+$, with $u_1' = u_2$, and an isomorphism $\pi_t \colon D_t \to D_{t'}^+$ taking t to t', in such a way that we can define a homomorphism $\pi \colon T_1 \to T_2^+$ by

$$\pi'(x) = \begin{cases} \pi_t(x) & \text{if } x \in D_t \text{ for some } t \in G_{u_1}^- \\ 0 & \text{otherwise} \end{cases}.$$

We may then define $\pi \colon F \to \overline{F^+ \Diamond F^+}$ by

$$\pi(x_1, x_2, \dots, x_d) = (x_1, \pi'(x_1), x_3, x_4, \dots, x_d, \underbrace{0, 0, \dots, 0}_{d}).$$

If instead the height of the tree D_1 *does* exceed the height of the tree D_2 then we may find a homomorphism $\pi' \colon T_2 \to T_1^+$ taking u_2 to u_1. We can then set

$$\pi(x_1, x_2, \dots, x_d) = (\pi'(x_2), x_2, x_3, x_4, \dots, x_d, \underbrace{0, 0, \dots, 0}_{d}).$$

□

Again, we embed F in $\overline{F^+ \Diamond F^+}$ by means of the injective homomorphism $x \mapsto (x, 0)$. We also have an analogue of Lemma 8:

Lemma 16. *Let F_0 and F be (m,p,c)-tree products, and let $u \in F_0$. Then there exists an (m,p,c)-tree-product F' and an injective homomorphism $\theta \colon F \to F'$ such that, given any $c \in F$, there exists a homomorphism $\theta_c \colon F_0 \to \overline{F'}$ with $\theta_c(u) = \theta(c)$.*

Proof. Assume without loss of generality that F_0 is a product of n copies of a tree T_0 of height h and that F is a product of N copies of a tree T of height H, and let F' be a product of N copies of a tree T' of height $H + h + 1$. Define $\theta \colon F \to F'$ as in Lemma 8.

To construct the $\theta_c \colon F_0 \to \overline{F'}$ as we did in Lemma 8, all we need is that for any $w \in T_0$ and $b \in \overline{T}$, we can find a homomorphism $\psi \colon T_0 \to \overline{T'}$ taking w to $\phi(b)$. But this is easy—if $\phi(b) = 0$, take $\psi(x) = 0$ for all $x \in T_0$; otherwise, by the same procedure as used in the proof of Lemma 15, we can construct such a ψ with $\psi(x) \neq 0$ precisely for $x \in \cup_{t \in G_w^-} D_t$. □

We can now define the amalgamation S' as before, using the sparse Hales-Jewett theorem. It only remains to check that S' has the required properties. As long as we can prove an analogue of Lemma 10, the remainder of the proof will go through exactly as for $(2, p, 1)$-sets. Given the work that we have already done on homomorphisms, the proof of this lemma is only a slight extension of the proof of Lemma 10.

Lemma 17. *Let S be a picture (in the sense of this section, i.e. for (m, p, c)-sets), let $u \in F_0$ and let S' be the amalgamation of S over S_u. Then any line in S' is a picture-line and is entirely contained in one of the copies $S^{(i)}$ of S.*

Proof. Let $L = [x_1, x_2, \ldots, x_m]$ be a line in S'. Suppose for each $i = 1, 2, \ldots, m$ that $cx_i \in S'_{cx'_i}$. Then cx_i has first coordinate cx'_i, and so x_i has first coordinate x'_i. Hence for any k with $1 \leq k \leq m$ and λ_i $(1 \leq i \leq k-1)$ with $0 \leq \lambda_i \leq p$, we see that $cx_k + \sum_{i=1}^{k-1} \lambda_i x_i$ has first coordinate $cx'_k + \sum_{i=1}^{k-1} \lambda_i x'_i$ and so $cx_k + \sum_{i=1}^{k-1} \lambda_i x_i \in S'_{cx'_k + \sum_{i=1}^{k-1} \lambda_i x'_i}$. So L is a picture-line.

Assume that $u \notin L$, and suppose that cx_i and cx_j $(i < j)$ are in different copies of S, say $cx_i \in S^{(a)}$ and $cx_j \in S^{(b)}$ with $a \neq b$. Suppose first that $cx_j + x_i \in S^{(e)}$ for some $e \neq b$. Then let π be a homomorphism which fixes u and moves cx'_j. By taking some column in which $S^{(a)}$ and $S^{(e)}$ have the identity map but $S^{(b)}$ has π, we get $\pi(cx'_j) + x'_i = cx'_j + x'_i$, a contradiction. If instead $cx'_j + x'_i \in S^{(b)}$ then by taking a homomorphism π which fixes u and moves cx'_i and looking at an appropriate column, we get $cx'_j + \frac{1}{c}\pi(cx'_i) = cx'_j + x'_i$, again a contradiction. So all of the cx_i $(1 \leq i \leq m)$ must be in the same copy of S, say $S^{(a)}$.

Now, if L is not contained entirely within $S^{(a)}$ then we can choose i with $1 \leq i \leq k$ and $\lambda_1, \lambda_2, \ldots, \lambda_{i-1}$ with $0 \leq \lambda_j \leq p$ for $1 \leq j \leq i-1$ such that $y = cx_i + \sum_{k=1}^{i-1} \lambda x_k \notin S^{(a)}$, say $y \in S^{(b)}$. Then choose a homomorphism which fixes u but moves $y' = cx'_i + \sum_{k=1}^{i-1} \lambda x'_k$, the first coordinate of y. Picking a column in which $S^{(a)}$ has the identity but $S^{(b)}$ has π, we get $y = cx'_i + \sum_{k=1}^{i-1} \lambda x'_k = \pi(y)$, a contradiction.

Hence L is contained entirely within $S^{(a)}$.

Finally, if $u \in L$ then, precisely as in the proof of Lemma 10, we show that all but one point of L is contained in some copy of $S^{(a)}$ of S, and deduce that the whole of L is contained in $S^{(a)}$. □

The remainder of the proof goes through exactly as before, and we have thus established our sparse version of Rado's theorem:

Theorem 18. *Let m, p, c, k and g be positive integers. Then there exists a subset $S \subset \mathbb{N}$ such that*

- *whenever S is k-coloured it contains a monochromatic (m, p, c)-set; and*
- *the (m, p, c)-sets in S form no cycle of length $< g$.*

This settles the Deuber conjecture.

5 Concluding remarks

Let us end by mentioning that all of the work in this paper relates to *finite* partition regular systems. In contrast, in the infinite case nothing at all is known about sparseness. For example, recall the Finite Sums theorem: for any k and n, whenever $\mathbb{N}_+$ is k-coloured there exist $x_1, x_2, \ldots, x_n \in \mathbb{N}_+$ such that the set $\mathrm{FS}(x_1, x_2, \ldots, x_n)$ of all non-zero finite sums of the x_i is monochromatic. As we mentioned above, a sparse version of this result was proved by Nešetřil and Rödl [14]: for any k, g and n, there exists a subset $S \subset \mathbb{N}_+$ such that whenever S is k-coloured it contains a monochromatic set of the form $\mathrm{FS}(x_1, x_2, \ldots, x_n)$, but such that S contains no cycle of length $\leq g$ of sets of this form. This immediately gives us also a restricted Finite Sums theorem: for any k and n, there exists a subset $S \subset \mathbb{N}_+$ such that whenever S is k-coloured it contains a monochromatic set of the form $\mathrm{FS}(x_1, x_2, \ldots, x_n)$, but such that S contains no set of the form $\mathrm{FS}(x_1, x_2, \ldots, x_{n+1})$.

An infinite version of the Finite Sums Theorem was proved by Hindman [8]: whenever $\mathbb{N}_+$ is finitely coloured, there exists a monochromatic set of the form $\mathrm{FS}(x_1, x_2, x_3, \ldots)$. A sensible first question to ask is whether we can produce some restricted Hindman theorem.

For k a positive integer, denote by $\mathrm{FS}_{\leq k}(x_1, x_2, x_3, \ldots)$ the set of all non-zero sums of at most k elements from x_1, x_2, x_3, ... An obvious question to ask is the following.

Question 19 ([10]). *Let k be a positive integer. Does there exist a subset $S \subset \mathbb{N}_+$ such that whenever S is finitely coloured it contains a monochromatic set of the form $\mathrm{FS}_{\leq k}(x_1, x_2, x_3, \ldots)$ but such that S contains no set of the form $\mathrm{FS}_{\leq k+1}(x_1, x_2, x_3, \ldots)$?*

Nešetřil and Rödl [15] have conjectured that the answer to this question is yes; while Hindman [9] has conjectured that the answer is no. Even the simplest case, $k = 2$, is unknown. Indeed, it is not even known if there is a set $S \subset \mathbb{N}$ such that whenever S is finitely coloured there is a sequence x_1, x_2, x_3, ..., with all x_i and all $x_i + x_j$ $(i \neq j)$ the same colour, but such that S contains no set of the form $\mathrm{FS}(y_1, y_2, y_3, \ldots)$. See [10] for related questions and discussion.

References

[1] Bergelson, V., Hindman, N., and Leader, I., Sets partition regular for n equations need not solve $n + 1$, *Proc. London Math. Soc. (3)* **73** (1996), 41–68.

[2] Deuber, W.A., Partitionen und lineare Gleichungssysteme, *Math. Zeit.* **133** (1973), 109–123.

[3] Deuber, W.A., Developments based on Rado's dissertation "Studien zur Kombinatorik" in *Surveys in combinatorics, 1989* (Siemons, J., ed.), Cambridge Univ. Press, Cambridge, 1989, 52–74.

[4] Folkman, J., Graphs with monochromatic complete subgraphs in every edge colouring, *SIAM J. Appl. Math.* **18** (1970), 19–24.

[5] Frankl, P., Graham, R., and Rödl, V., Induced restricted Ramsey theorems for spaces, *J. Combin. Theory Ser. A.* **44** (1987), 120–128.

[6] Graham, R.L., Rothschild, B.L., and Spencer, J.H., Ramsey theory, John Wiley & Sons, Inc., New York, 1980.

[7] Hales, A.W., and Jewett, R.I., Regularity and positional games, *Trans. Amer. Math. Soc.* **106** (1963), 222–229.

[8] Hindman, N., Finite sums from sequences within cells of a partition of $\mathbb{N}$, *J. Combin. Theory Ser. A.* **17** (1974), 1–11.

[9] Hindman, N., Problems and new results in the algebra of βS and Ramsey Theory in *Unsolved Problems in Mathematics for the 21st Century* (Abe, J., and Tanaka, S., eds), IOS Press, Amsterdam, 2001, 295–305.

[10] Hindman, N., Leader, I., and Strauss, D., Open problems in partition regularity, *Combin. Probab. Comput.* **12** (2003), 571–583.

[11] Nešetřil, J., and Rödl, V., A short proof of the existence of highly chromatic hypergraphs without short cycles, *J. Combin Thoery Ser. B* **27** (1979), 225–227.

[12] Nešetřil, J., and Rödl, V., Partition theory and its application, in *Surveys in combinatorics (Proc. Seventh British Combinatorial Conf., Cambridge, 1979)*, Cambridge Univ. Press, Cambridge-New York, 1979.

[13] Nešetřil, J., and Rödl, V., Sparse Ramsey graphs, *Combinatorica* **4** (1984), 71–78.

[14] Nešetřil, J., and Rödl, V., Finite union theorems with restrictions, *Graphs and Combinatorics* **2** (1980), 357–361.

[15] Nešetřil, J., and Rödl, V., Partite constructions and Ramsey space systems, in *Mathematics of Ramsey Theory* (Nešetřil, J. and Rödl, V., eds), Springer-Verlag, 1990, 98–112.

[16] Prömel, H.-J., and Voigt, B., A sparse Graham-Rothschild theorem, *Trans. Amer. Math. Soc.* **309** (1985), 113–137.

[17] Rado, R., Studien zur Kombinatorik, *Math. Zeit.* **36** (1933), 242–280.

[18] Rödl, V., On Ramsey families of sets, *Graphs and Combinatorics* **6** (1990), 187–195.

[19] Schur, I., Über die Kongruenz $x^m + y^m \equiv z^m \pmod{p}$, *Jber. Deutsch. Math.-Verein.* **25** (1916), 114–117.

[20] Spencer, J.H., Restricted Ramsey configurations, *J. Combin. Theory Ser. A* **19** (1975), 278–286.

[21] van der Waerden, B.L., Beweis einer Baudet'schen Vermutung, *Nieuw. Arch. Wisk.* **15** (1927), 212–216.