Sparse Ramsey Theory

Dr. Imre Leader

Summer 2003

Introduction

It is well known that, whenever the edges of the complete graph K_6 on 6 vertices are 2-coloured, there is always a monochromatic copy of K_3. Indeed, in general, for any k there is an n such that, if $E(K_n)$ is k-coloured, there exists a monochromatic copy of K_3.

We illustrate these two statements symbolically as follows:

$$
\begin{array}{lcll}
K_6 & \longrightarrow & K_3 & \text{(to be read "}K_6\text{ is Ramsey for }K_3\text{".)} \\
K_n & \stackrel{k}{\longrightarrow} & K_3 & \text{(to be read "}K_n\text{ is }k\text{-Ramsey for }K_3\text{".)}
\end{array}
$$

Now, we have a natural question to ask: What does $K_6 \longrightarrow K_3$ tell us about the copies of K_3 in K_6: in what sense are they dense? In what sense do they fit together well?

For example: **Question:** if a graph G is such that $G \longrightarrow K_3$, must G contain K_6?

Answer: The answer to this is "no". There is an elementary construction: Let C be K_9 – *9-cycle* – i.e.:

$$C = K_9 - y_1y_2 - y_2y_3 - \cdots - y_9y_1$$

Form G by adding a new point x, joined to every vertex in C. Clearly G does not contain K_6 (or else C would have to contain K_5 – it doesn't.)

Now, given any 2-colouring of $E(G)$, we have at least five edges of the same colour emanating from x. Call them $xy_{i_1}, xy_{i_2}, xy_{i_3}, xy_{i_4}, xy_{i_5}$. Without loss of generality, we may assume $i_1 < \cdots < i_5$, and that y_{i_1} and y_{i_5} are adjacent in C.

Consider the triangle $y_{i_1}, y_{i_3}, y_{i_5}$. If any edge is red, we have formed a red triangle with x. If all the edges are blue, we have formed a blue triangle. Either way, we're finished. $\square$

Now we can start being a bit more ambitious – we can ask for a graph $G \not\supseteq K_5$ with $G \longrightarrow K_3$? Such a G does actually exist, although the construction is a bit harder.

Even more daringly, is there a $G \not\supseteq K_4$ with $G \longrightarrow K_3$? This question turns out to be much harder. Folkman found a construction for 2-colourings only. Finally Nešetřil and Rödl solved the problem for the general case of k-colourings:

For all k, there exists a graph $G \not\supseteq K_4$ such that $G \xrightarrow{k} K_3$.

This is a mildly surprising statement, since the condition $G \not\supseteq K_4$ says that the copies of K_3 in G are not dense and closely linked, whereas the condition that $G \xrightarrow{k} K_3$ says just the opposite.

In Chapter 1, "Building Sparse Systems", we'll prove the above result of Nešetřil and Rödl, and much, much more. We'll first introduce the key idea of *amalgamation*.

In Chapter 2, "Sparse Arithmetic Structures", we'll be trying to strengthen Van der Waerden's theorem in a similar way.

For example, we know that there exists an n such that, whenever $[n] = \{1, \ldots, n\}$ is 2-coloured, there exists a monochromatic arithmetic progression of length 3. Can we replace $[n]$ with a 'sparser' set: for example, a set containing no arithmetic progression of length 4, and get the same result for it?

1 Building Sparse Systems

1.1 Graphs with high girth and chromatic number

Recall that the girth $g(G)$ of a graph G is the length of a shortest cycle in the graph, and that the chromatic number $\chi(G)$ is the least number of colours required to colour the vertices so no two adjacent vertices share the same colour.

Erdös proved that, for all g and k, there is a graph G with $g(G) > g$ and $\chi(G) > k$. Of course making either $g(G)$ or $\chi(G)$ large on their own is easy: we can take C_n or K_n respectively, for large n. Making them both simultaneously large is much harder.

Erdös used a proof involving random graphs. Later on, Lovász gave a constructive proof. However, we will present a later proof yet: that of Nešetřil and Rödl, since it serves as an excellent introduction to the idea of amalgamation.

We'll start by solving an easier problem: that of finding a triangle-free graph G with a large chromatic number. In what follows, we shall use **Δ-free** as an abbreviation for "triangle-free".

This is just the $g = 3$ case of the problem above: we seek G with $g(G) > 3$ (i.e., G is triangle-free) such that $\chi(G) > k$ for given k. So we need a triangle-free graph which has a monochromatic edge whenever it is k-coloured.

Here's our first idea: for triangle-free graphs, forcing *edges* to be monochromatic could be hard; forcing *independent sets* (i.e., sets which span no edges) to be monochromatic should be easier.

So, as a start, can we find a Δ-free graph G_0, containing certain named independent sets $V_0, \ldots, V_r$ such that any k-colouring of $V(G)$ in which each of $V_0, \ldots, V_r$ are monochromatic must contain a monochromatic edge?

Of course we can. Some two V_i must have the same colour if $r = k$. Thus we ensure that all V_i, V_j are connected by an edge. Furthermore, we choose these edges disjointly – this, too, is clearly possible.

So we are in possession of a $(k+1)$-partite, Δ-free G_0, such that if G_0 is k-coloured with each V_i monochromatic then there is a monochromatic edge.

From here on, we regard all $(k+1)$-partite graphs as coming with a fixed partition; all the standard graph-theoretic ideas we use shall be required to respect this structure. For example, if G is $(k+1)$-partite on classes $(V_0, \ldots, V_k)$ and H is $(k+1)$-partite on classes $(W_0, \ldots, W_k)$, then **"G contains H"** means $E(H) \subset E(G)$ and $W_i \subset V_i$ for all i.

As another example, the **union** $G \cup H$ has edge set $E(G) \cup E(H)$ and partition $(V_0 \cup W_0, \ldots, V_k \cup W_k)$.

Our dream is now the following: that we should be able to find a partite (i.e. k-partite) graph G that is Δ-free, such that any k-colouring of the vertices gives a copy of G_0 with each partition class monochromatic.

Let G be a partite graph, and let $0 \leq i \leq n$. We define the **amalgamation** of G on V_i, written $A_i(G)$, as follows:

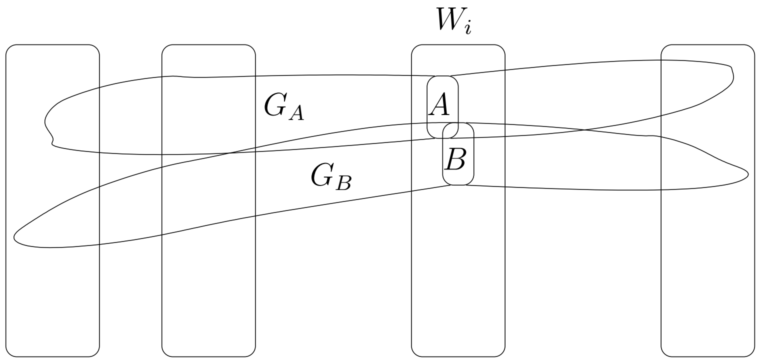

Figure 1: Amalgamation – showing two sets A and B and their corresponding disjoint copies G_A and G_B of G.

Fix a set W_i such that $|W_i| = k|V_i|$. For each set $A \subset W_i$ with $|A| = |V_i|$, let G_A be a copy of G with i-th vertex class $V_i(G_A) = A$, and all other classes disjoint. Then we put

$$A_i(G) = \bigcup_{\substack{A \subset W_i \\ |A|=|V_i|}} G_A.$$

Proposition 1 *Let G be partite, and $0 \leq i \leq k$. Then:*

1. *Whenever $A_i(G)$ is k-coloured, there exists a copy of G whose i-th class is monochromatic.*
2. *If G is Δ-free, then $A_i(G)$ is Δ-free.*

Proof:

1. We have W_i, the i-th vertex class of $A_i(G)$, k-coloured. So there exists a monochromatic $A \subset W_i$ with $|A| = |V_i|$. Then G_A will do. ✓
2. If $A_i(G) \supset \Delta$, where Δ is some triangle, then the edges $E(\Delta)$ are not contained in just one G_A, because G_A is Δ-free.

   But then at least two vertices are incident with edges from more than one G_A, and hence both belong to W_i. This is a contradiction, since W_i is independent. ✓ □

Now, by repeating this we can get:

Theorem 2 *For any k, there is a graph G which is Δ-free and satisfies $\chi(G) > k$.*

Proof: Set $G = A_k(A_{k-1}(\cdots A_0(G_0)\cdots))$. Then G is Δ-free, and whenever G is k-coloured, there exists a copy of G_0 with each class monochromatic. But, by definition of G_0, this means that $\chi(G) > k$. □

Remarks:

1. The above method is called **amalgamation**. Sometimes the idea of working only with $(k+1)$-partite graphs is called the **partite construction**, and Proposition 1 is called the **partite lemma**.
2. In finding our G, we have used the fact that there exists a graph with chromatic number greater than k: we were dependent on K_{k+1} having this property when we chose the columns and edges of G_0.

Now we conquer the problem of graphs of large girth and chromatic number. What if we want girth greater than 4? G_0 is fine: it contains no cycles whatsoever.

But $A_0(G_0)$ can contain adjacent edges: then $A_1(A_0(G_0))$ contains 4-cycles.

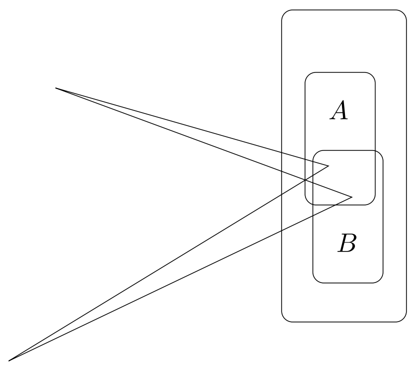

Figure 2: A 4-cycle in $A_1(A_0(G))$.

We get 4-cycles exactly when we have $A, B \subset W_i$ such that $|A| = |B| = |V_i|$, with $|A \cap B| \geq 2$.

Here's the solution: when we are forming $A_i(G)$, we mustn't take the whole of

$$\bigcup_{\substack{A \subset W_i \\ |A| = |V_i|}} G_A,$$

but only some of them:

$$\bigcup_{A \in \mathcal{A}} G_A$$

for some cleverly chosen subset $\mathcal{A} \subset [W_i]^{(V_i)}$.

So our collection $\mathcal{A}$ of subsets of W_i of size V_i must satisfy:

1. Whenever W_i is k-coloured, there must exist a monochromatic $A \in \mathcal{A}$.
2. If we have $A, B \in \mathcal{A}$ such that $A \neq B$, then we have $|A \cap B| \leq 1$.

It is not clear if such an $\mathcal{A}$ even exists.

How can we forbid bigger cycles: 6-cycles, for example?

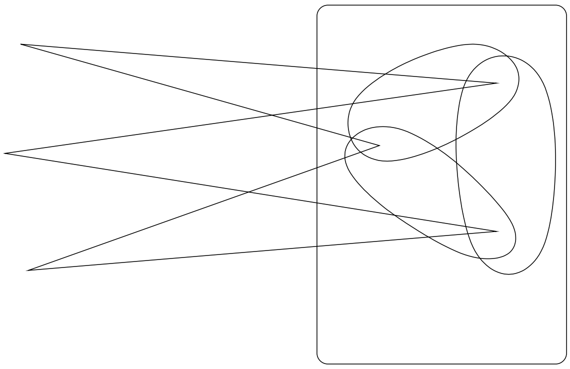

Figure 3: A 6-cycle – to be prevented.

We say that an n**-cycle** in a hypergraph $\mathcal{A}$ is a sequence $A_1, x_1, \dots, A_n, x_n$ where $A_1, \dots, A_n \in \mathcal{A}$ and $x_1, \dots, x_n \in V(\mathcal{A})$, with all $A_1, \dots, A_n$ and all $x_1, \dots, x_n$ distinct, and satisfying $x_i \in A_i \cap A_{i+1}$ for all i (where A_{n+1} means A_1).

So a 2-cycle is two sets A, B with $|A \cap B| \geq 2$. A 3-cycle is three sets A, B, C as in figure 4.

Note: If $\mathcal{A}$ is a hypergraph of 2-sets (i.e. a graph), then this reduces to the usual definition of cycle.

Equivalently: an n-cycle in $\mathcal{A}$ corresponds to a $2n$-cycle in the incidence graph of $\mathcal{A}$. (This is the bipartite graph on vertex classes $\mathcal{A}$ and $V(\mathcal{A})$ with A joined to x if $x \in A$.) Parenthetically, we should also note that it does *not* correspond to an n-cycle in the intersection graph (the graph with vertices $\mathcal{A}$, with A joined to B if $A \cap B \neq \emptyset$).

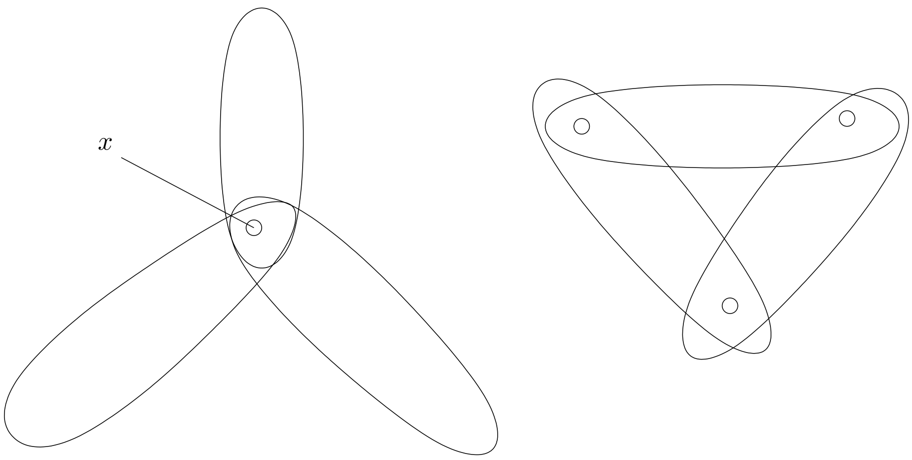

Figure 4: The left-hand set system isn't a cycle (where each pair of sets have intersection $\{x\}$). The right-hand system is a cycle.

The **girth** of $\mathcal{A}$ is the length of a shortest cycle. It looks as if, to make amalgamation preserve girth at least g, we'd need $\mathcal{A}$ of girth at least $g/2$, and such that whenever W_i is k-coloured, we'd get a monochromatic $A \in \mathcal{A}$.

We need: For all r, g, k, there exists an r-graph (i.e., a family of r-sets) with girth at least g, and chromatic number $\chi(\mathcal{A}) > k$ (this means whenever $V(\mathcal{A})$ is k-coloured, there exists a monochromatic $A \in \mathcal{A}$).

There is promising news: for $r = 2$, this is what we're trying to prove. And to prove it for girth g, it looks like we only need it for girth $g/2$, thus paving the way for an induction proof.

But there is also bad news: such an $\mathcal{A}$ looks much harder to construct than for just $r = 2$. Indeed, take an r-graph $\mathcal{A}$ with $g(\mathcal{A}) > g$ and $\chi(\mathcal{A}) > k$. For each $A \in \mathcal{A}$, select some 2-set $l_A \subset A$. Then the graph G with edges $\{l_A : A \in \mathcal{A}\}$ has $g(G) > g$ and $\chi(G) > k$. This is terrifying: *however* we choose these l_A we get a hard-to-find graph (i.e. one with large girth and chromatic number)!

Say that an r-graph $\mathcal{A}$ is t-partite on vertex classes $V_1, \ldots, V_t$ if, for all $A \in \mathcal{A}$, and for all i, we have $|A \cap V_i| \leq 1$.

Given r, g, k, can we find, for some t, a t-partite r-graph $\mathcal{A}_0$ on classes $V_1, \ldots, V_t$ such that whenever $V(\mathcal{A})$ is k-coloured with each V_i monochromatic, we have a monochromatic $A \in \mathcal{A}$, and such that $g(\mathcal{A}) > g$?

Yes, we can: let $t = kr$, and for each r-set $k \subset [1, \ldots, t]$ choose disjointly an

r-set A_r. Make sure A_r contains 1 point of each V_i for $i \in R$. ✓

Let $\mathcal{A}$ be a t-partite r-graph on vertex classes $V_1, \ldots, V_t$, and let $\mathcal{B}$ be a $|V_i|$-graph. We define the **amalgamation** of $\mathcal{A}$ over $\mathcal{B}$, written $\mathcal{A} \star \mathcal{B}$, or $\mathcal{A} \star_i \mathcal{B}$, as follows: for each $B \in \mathcal{B}$, let $\mathcal{A}_B$ be a copy of $\mathcal{A}$ with i-th class B, and disjoint in all the other classes. We then set

$$\mathcal{A} \star_i \mathcal{B} = \bigcup_{B \in \mathcal{B}} \mathcal{A}_B.$$

Proposition 3 *With $\mathcal{A}$ and $\mathcal{B}$ as above:*

1. *If $\chi(\mathcal{B}) > k$ then whenever $\mathcal{A} \star_i \mathcal{B}$ is k-coloured then there is a copy of $\mathcal{A}$ in $\mathcal{A} \star_i \mathcal{B}$ with the i-th class monochromatic.*
2. *If $g(\mathcal{A}) > g$ and $g(\mathcal{B}) > g/2$, then $g(\mathcal{A} \star \mathcal{B}) > g$.*

Proof:

1. When $\mathcal{A} \star \mathcal{B}$ is k-coloured, there is a $B \in \mathcal{B}$ which is monochromatic, because $\chi(\mathcal{B}) > k$. Then $\mathcal{A}_B$ will do. ✓
2. Suppose $R_1, x_1, \ldots, R_g, x_g$ is a g-cycle in $\mathcal{A} \star \mathcal{B}$. Then no two consecutive x_i can belong to W_i (the i-th class of $\mathcal{A} \star_i \mathcal{B}$) as $\mathcal{A} \star_i \mathcal{B}$ is t-partite.

   But the x_j belonging to W_i induce a cycle in $\mathcal{B}$: if we have $a < b$ with $x_a, x_b \in W_i$ but $x_{a+1}, \ldots, x_{b-1} \notin W_i$, then $x_a, x_b \in \mathcal{A}_B$ for some B (by construction). This is a contradiction unless no more than one of the x_j are in W_i. In this case, the g-cycle lives in some $\mathcal{A}_B$: a contradiction. ✓ □

At last we can prove the result we seek:

Theorem 4 *For all r, k, g, there is an r-graph $\mathcal{A}$ with $g(\mathcal{A}) > g$ and $\chi(\mathcal{A}) > k$. In particular, there exist graphs with girth more than g and chromatic number more than k.*

Proof: Fix any k. The proof is by induction on g (for all r).

The $g = 1$ case is easy. ✓

Given r and g, choose some t and a t-partite $\mathcal{A}_0$ as before. Let $\mathcal{A}_0$ have 1st vertex class $V_1(\mathcal{A}_0)$, and let $\mathcal{B}_1$ be a $V_1(\mathcal{A}_0)$-graph with $g(\mathcal{B}_1) > g/2$ and $\chi(\mathcal{B}_1) > k$.

Set $\mathcal{A}_1 = \mathcal{A}_0 \star_1 \mathcal{B}_1$. Now choose $\mathcal{B}_2$ a $|V_2(\mathcal{A}_1)|$-graph with $g(\mathcal{B}_2) > g/2$ and $\chi(\mathcal{B}_2) > k$. Then set $\mathcal{A}_2 = \mathcal{A}_1 \star_2 \mathcal{B}_2$. Keep going, until we obtain

$$\mathcal{A} = (\ldots((\mathcal{A}_0 \star_1 \mathcal{B}_1) \star_2 \mathcal{B}_2) \ldots).$$

Then $g(\mathcal{A}) > g$ (by Proposition 3), and whenever $V(\mathcal{A})$ is k-coloured, there is a copy of $\mathcal{A}_0$ with each vertex class monochromatic. Now we're done, by the very definition of $\mathcal{A}_0$. ✓ □

Remark: The other known constructions for large girth and chromatic number (for graphs), by Lovász and Kriz, also go outside the world of graphs (e.g. by using hypergraphs or similar).

1.2 The Restricted Ramsey Theorem

We now want a G, not containing K_4, such that whenever $E(G)$ is k-coloured, there exists a monochromatic triangle.

As a first step: can we find, for some t, a t-partite graph G_0 on vertex classes $V_1, \ldots, V_t$ such that $G_0 \not\supseteq K_4$, with the property that in any k-colouring of $E(G_0)$ with each $G_0[V_i, V_j] = G_0(V_i \cup V_j)$ (i.e. the induced subgraph spanned by $V_i \cup V_j$) monochromatic, we have a monochromatic triangle?

Yes: let $t = R_k(3)$ (i.e. t is such that whenever $E(K_t)$ is k-coloured, there is a monochromatic triangle). For each l, m, n take a disjoint triangle on classes V_l, V_m, V_n – see figure 5. ✓

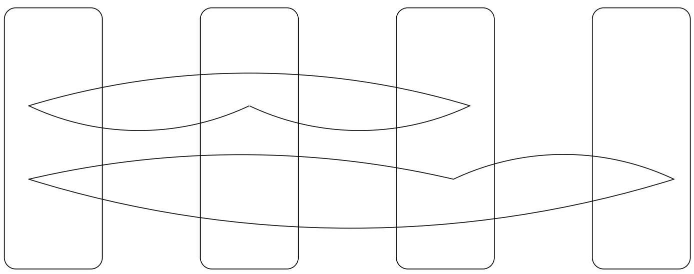

Figure 5: The construction of G_0.

How do we do amalgamation on this? We will be needing a bipartite graph B such that whenever the edges of B are 2-coloured, there exists a monochromatic copy of $G[V_i, V_j]$. A problem stems from needing to have no copies of K_4 in B: if we were unlucky, we could create a K_4 from different copies of $G[V_i, V_j]$.

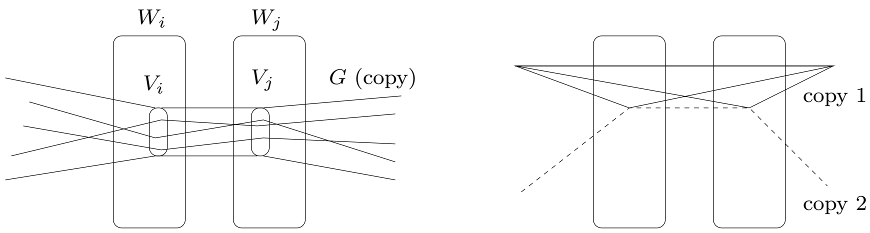

Figure 6: Left: Amalgamation for Restricted Ramsey. Right: A copy of K_4 using two copies of $G[V_i, V_j]$.

So, when $E(B)$ is k-coloured, we want there to exist a monochromatic *induced* copy of $G[V_i, V_j]$. (Recall that a subgraph $H \subset G$ is **induced** if H contains all the edges that G contains between any two of its vertices.)

Thus, we'd like, for any bipartite graph A, a bipartite graph B such that when $E(B)$ is k-coloured, we have a monochromatic induced copy of A. Such a result is a *bipartite induced* Ramsey theorem.

The appendix contains a digression on the Hales-Jewett theorem: a key result from Ramsey theory. Hales-Jewett gives us an instant proof of what we want.

Theorem 5 (Bipartite Induced Ramsey Theorem) *Let A be a bipartite graph. Then there exists a bipartite B such that whenever B is k-coloured, there exists a monochromatic induced copy of A.*

Proof: Let A have edge set E, and vertex classes X and Y. For some n large, let B have edge set E^n, and vertex classes X^n, Y^n. This works as follows: edge $(e_1, \ldots, e_n)$ joins $(x_1, \ldots, x_n)$ to $(y_1, \ldots, y_i)$ if e_i joins x_i to y_i for all i.

When $E^n = E(B)$ is k-coloured, by Hales-Jewett there exists a monochromatic line L, say:

$$L = \{(l_1, \ldots, l_n) : l_i = f_i (\forall i \notin I), l_i = l_j (\forall i, j \in I)\}$$

for some non-empty $I \subset [n]$ and some f_i for $i \notin I$.

Then L is the edge-set of a copy of A – indeed, we take as our copies of X and Y the sets:

$$\begin{aligned} X' &= \{(x_1, \ldots, x_n) : x_i = x(f_i)(\forall i \notin I), x_i = x_j (\forall i, j \in I)\} \\ Y' &= \{(y_1, \ldots, y_n) : y_i = y(f_i)(\forall i \notin I), y_i = y_j (\forall i, j \in I)\} \end{aligned}$$

(where the edge e is adjacent to $x(e) \in X$ and $y(e) \in Y$.

Also, this copy of A is induced. For suppose $(e_1, \ldots, e_n)$ connects $(x_1, \ldots, x_n) \in X'$ to $(y_1, \ldots, y_n) \in Y'$. Then for all $i \notin I$, $x_i = x(f_i)$ and $y_i = y(f_i)$, and e_i joins x_i to y_i. So $e_i = f_i$. Furthermore, for all $i, j \in I$, $x_i = x_j$ and $y_i = y_j$ so $e_i = e_j$. Thus $(e_1, \ldots, e_n) \in L$. $\square$

Remarks:

1. Hales-Jewett is a natural thing to use, because the lines in A^d don't interfere with one another.
2. One can also prove Theorem 5 directly (i.e. without Hales-Jewett), by using Ramsey's theorem. The proof is, however, longer and quite a bit harder.

Now we assault our target theorem: we take G a t-partite graph, with vertex classes $V_1, \ldots, V_t$, B a bipartite graph, and $\mathcal{A}$ a collection of copies of $G[V_i, V_j]$ inside B.

We define the **amalgamation** of G over $\mathcal{A}$, written $G \star \mathcal{A}$ (or $G \star_{i,j} \mathcal{A}$) as follows: for each $A \in \mathcal{A}$ take a copy G_A of G with $G_A[i, j] = A$ (where $G_A[i, j]$ is the subgraph of G_A spanned by classes i and j). Otherwise (i.e. on classes $\neq i, j$ the G_A are disjoint.

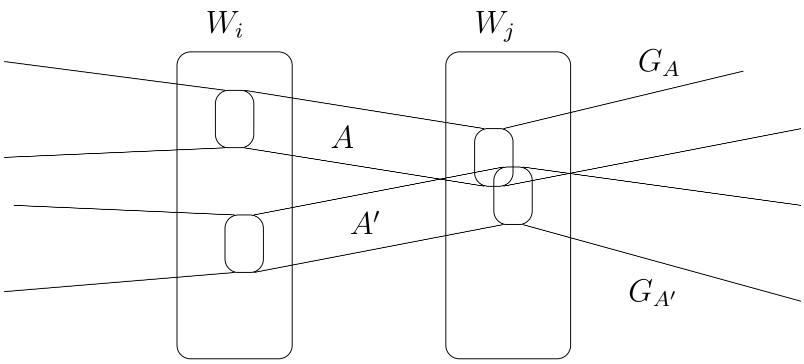

Figure 7: Amalgamation for Sparse Ramsey.

Then set $G \star \mathcal{A}$ to be $\bigcup_{A \in \mathcal{A}} G_A$.

Lemma 6 *Let G be a k-partite graph. Then:*

1. *If $\chi(\mathcal{A}) > k$ (i.e. whenever $E(\mathcal{A})$ is k-coloured, there exists a monochromatic $A \in \mathcal{A}$), then whenever $E(G \star \mathcal{A})$ is k-coloured, there exists a*

copy G_A of G with $G_A[i,j]$ monochromatic.

2. *If each $A \in \mathcal{A}$ is an induced subgraph of B, and $G \not\supset K_4$, then $G \star \mathcal{A} \not\supset K_4$.*

Proof:

1. Obvious. ✓

2. Suppose we have some $K_4 \subset G \star \mathcal{A}$. Say $V(K_4) = \{a, b, c, d\}$, where a is in the i-th class and b is in the j-th. (If it didn't meet V_i and V_j, then we'd have $K_4 \subset G_A$ for some $A \in \mathcal{A}$: a contradiction).

   Since the G_A are disjoint for all $A \in \mathcal{A}$ outside the i-th and j-th classes, we must have $K_4 - ab \subset G_A$, for some $A \in \mathcal{A}$. Then $ab \notin G_A$ (or else $G_A \supset K_4$). So $ab \in G_{A'}$ for some $A' \in \mathcal{A}$. But this contradicts A being an induced subgraph of B. ✓ □

Now we are ready for the main result:

Theorem 7 *For all k, there is a G such that $G \xrightarrow{k} K_3$ but $G \not\supset K_4$.*

Proof: Start with G_0 t-partite, for some t. Let it have vertex classes $v_1, \ldots, V_t$ such that $G_0 \not\supset K_4$, and whenever $E(G_0)$ is k-coloured with $G[V_i, V_j]$ monochromatic for all i and j, there exists a monochromatic triangle.

Choose $i < j$ and choose a bipartite B such that whenever $E(B)$ is k-coloured, there exists a monochromatic induced $G_0[V_i, V_j]$. Let $\mathcal{A}$ consist of all those induced copies of $G_0[V_i, V_j]$ and form $G_1 = G_0 \star_{i,j} \mathcal{A}$.

Repeat $\binom{t}{2}$ times – once for each pair (i, j). □

Corollary 8 *For all k and s, there is a G with $G \xrightarrow{k} K_s$ but $G \not\supset K_{s+1}$.*

Proof: Take G_0 to be a disjoint union of copies of K_s instead of copies of K_3, with t being not $R_k(3)$ but $R_k(s)$. Then the proof is as above. □

Define the **clique number** $\mathrm{Cl}(H)$ to be the largest s with $K_s \subset H$.

Corollary 9 *Let H be a graph. For all k, there is a graph G with $G \xrightarrow{k} H$ but $\mathrm{Cl}(G) = \mathrm{Cl}(H)$.*

Proof: Let H have n points and $\mathrm{Cl}(H) = s$. Take G_0 to be a disjoint union of copies of H (with t being $R_k(n)$). Then the proof is as above (if $G \not\supset K_{s+1}$ then $G \star \mathcal{A} \not\supset K_{s+1}$). □

Remark: Clearly, Corollary 9 is a strengthening of Corollary 8, which in turn is a strengthening of Theorem 7.

Amusingly, Corollary 9 immediately implies the Induced Ramsey Theorem:

Theorem 10 *For all H and k there is a G such that whenever $E(G)$ is k-coloured, there is a monochromatic induced copy of H.*

Proof: Let $\mathrm{Cl}(H) = s$. Form H' as follows: for each $ab \notin E(H)$ disjointly add a copy of $K_{s+1} - \mathit{edge}$ to H at a, b.

Find G with $G \xrightarrow{k} H'$, such that $\mathrm{Cl}(G) = \mathrm{Cl}(H') = s$. We claim such a G will do.

When $E(G)$ is k-coloured, there exists a monochromatic copy of H'. This contains a copy of H, which must be induced (or else $G \supset K_{s+1}$). ✓ □

Remarks:

1. Alternatively, one can prove Theorem 10 by mimicking the proof of Corollary 9.
2. Often, in the literature, "*induced*" is coded as having a *fixed* colouring of a complete graph. For example, a graph H on n points would be thought of as a red-blue colouring of K_n.
3. We could combine Corollary 9 and Theorem 10 to get the following: for all H and k, there is a G such that $\mathrm{Cl}(G) = \mathrm{Cl}(H)$, and whenever $E(G)$ is k-coloured, there is a monochromatic induced copy of H.

1.3 The Sparse Ramsey Theorem

In general, a *restricted* theorem says that we can find a special subobject of rank s inside an object containing no subobjects of rank $s+1$. So Theorem 7, Corollary 8 and Corollary 9 are all restricted. So is Theorem 2; it says that when we k-colour a certain graph we get a monochromatic K_2 despite having no K_3 whatsoever.

From that, we went on to prove that there is a graph G with $\chi(G) > k$ and $g(G) > g$. This is a *sparse* theorem: it states that we can find a subobject of rank s, but the subobjects of rank s have no short cycles.

In general, we have the slogan *"sparse implies restricted, which implies induced"*. We can now ask, what would a sparse version of Theorem 7 say?

We'd want a G such that $G \xrightarrow{k} K_3$, but such that the set $K_3(G)$ of triangles in G have no short cycles, viewing the triangles as forming a hypergraph: letting any triangle be represented by the set of its three edges. This is an important point: note that we are *not* representing a triangle by its vertex set. Trivially, there would be short cycles in any such graph if we did.

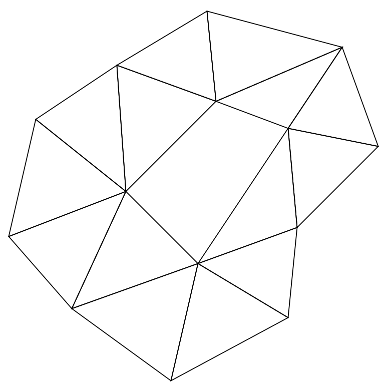

Figure 8: A 14-cycle of triangles.

From our work on amalgamation (from Lemma 6), we have that if the girth of $K_3(G)$ is more than g, and the girth of $\mathcal{A}$ is at least $g/2$, then the girth of $K_3(G \star \mathcal{A})$ is also at least g. Just as in the second part of Proposition 3, this is because of how we transfer from a triangle in some G_A to a triangle in some different $G_{A'}$: two such triangles cannot share an edge. Indeed, figure 9 shows an impossible configuration on the left (no x not in $V(B)$ can belong to G_A and $G_{A'}$ since they are disjoint except on B); as opposed to a possible configuration on the right.

So we find ourselves needing a sparse version of Hales-Jewett. We want that, for all n, k and g there is a d and a family $\mathcal{L}$ of lines in $[n]^d$ such that whenever $[n]^d$ is k-coloured there is a monochromatic line $L \in \mathcal{L}$, but $\mathcal{L}$ has no short cycles (as a hypergraph).

This was proved by Spencer; but we will give a simple proof due to Rödl, using ideas that have many other applications. It is an elegant random argument.

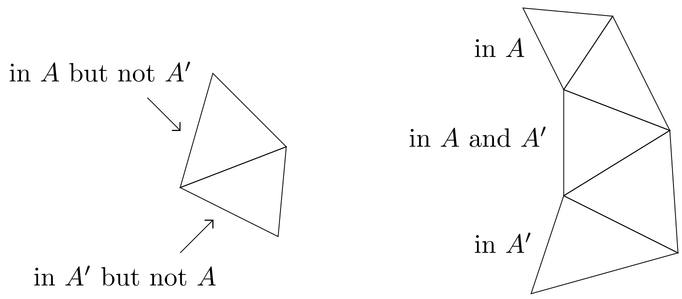

Figure 9: Avoiding cycles of triangles.

We'll do some work towards setting it up now. Fix a d_0 such that whenever $[n]^{d_0}$ is k-coloured, there is a monochromatic line.

Consider the set $\mathcal{K}$ of all lines in $[n]^d$ with at most d_0 active coordinates, for some d_0. The number of such lines is

$$|\mathcal{K}| = \sum_{i=1}^{d_0} \binom{d}{i} n^{d-i} \geq \binom{d}{d_0} n^d n^{-d_0}.$$

It can easily be checked that at least $\alpha = 1/(n+1)^{d_0}$ of all such lines are monochromatic in any k-colouring of a cube $[n]^d$ for $d > d_0$.

Now we want an upper bound for the number of lines in $\mathcal{K}$ a point can be on. In fact, any point is on at most D lines, where

$$D = \sum_{i=1}^{d_0} \binom{d}{i} \leq 2\binom{d}{d_0} \qquad \text{(for } d \text{ large enough).}$$

Now we are ready to state the theorem:

Theorem 11 (Weak Sparse Hales-Jewett) *For all n, k and g, there is a d and a family $\mathcal{L}$ of lines in $[n]^d$ such that whenever $[n]^d$ is k-coloured, there is a monochromatic line $L \in \mathcal{L}$ but such that $g(\mathcal{L}) > g$.*

Remarks:

1. This result is a weak sparse version of Hales-Jewett, as our set $\mathcal{L}$ was not the set of all lines in some $S \subset [n]^d$.

2. If you are unsure about random-type arguments, then relax: some much stronger results will be proved by other methods in the next chapter, independently of this proof.

Proof: We'll choose lines iteratively for $\mathcal{L}$, so that each is monochromatic with respect to many k-colourings, but so that we keep the girth small.

To ensure that we can do this, we wish to keep a handle on the maximum degree of $\mathcal{L}$. Fix some d large – we'll decide exactly how large later. Choose lines $L_1, L_2, \ldots$ inductively as follows:

Suppose we have chosen $L_1, \ldots, L_r$ such that:

1. There is no point in more than c of them (for some c which we'll choose later).
2. We have $g(\{L_1, \ldots, L_r\}) > g$.

Let $S = \{\Phi : \Phi \text{ is a } k\text{-colouring of } [n]^d \text{ with no } L_i \text{ monochromatic}\}$.

If S is empty, we're finished. Otherwise, we need to ask how many edges there could be to use as L_{r+1}, without violating properties 1 and 2.

How many lines violate property 1? There are at most rn/c points of degree c. So not more than rnD/c lines are illegal.

How many lines violate property 2? For each x. there are at most $(cn)^{g/2}$ points at a distance less than $g/2$ from x (measuring distance along the lines $L_1, \ldots, L_r$). So fewer than $n^d(nc)^g$ lines break property 2.

So, what we want is

$$\frac{rnD}{c} \leq \frac{\alpha}{4}\#(\mathcal{K}) \qquad (1)$$

$$n^d(nc)^g \leq \frac{\alpha}{4}\#(\mathcal{K}) \qquad (2)$$

Then we'd have more than $1 - \alpha/2$ of all lines in $\mathcal{K}$ legal to be added. Then some one of these lines is monochromatic for at least $\alpha/2$ of the k-colourings in S. Let this be L_{r+1}.

So after r steps, the number of k-colourings with none of $L_1, \ldots, L_r$ monochromatic is at most

$$k^{n^d}\left(1 - \frac{\alpha}{2}\right)^r$$

(the first term is the number of k-colourings; the second the proportion that can be left after r steps).

We may stop when this is less than 1. This happens when

$$\left(1 - \frac{\alpha}{2}\right)^r < \frac{1}{k^{n^d}},$$

i.e. when

$$r = \frac{n^d \log k}{-\log(1 - \alpha/2)} = Cn^d$$

(for some C).

To satisfy inequality 1, then, we need

$$\frac{2Cn^d n}{c} \binom{d}{d_0} \leq \frac{\alpha}{4} \binom{d}{d_0} n^d n^{-d_0}, \qquad \text{i.e.} \qquad \frac{2Cn}{c} \leq \frac{\alpha}{4} n^{-d_0}.$$

This can be done by taking $c = C'$, for some large C' (independent of d).

Now, to satisfy inequality 2, we need

$$n^d (nc)^g \leq \frac{\alpha}{4} \binom{d}{d_0} n^d n^{d_0}.$$

This is easy now: we just take d very large. □

Now we can prove the following:

Theorem 12 (Sparse Ramsey Theorem) *For all k and g, there is some G such that $G \xrightarrow{k} K_3$, but with G having no cycle of triangles of length less than g.*

Proof: In the proof of Theorem 10, just replace our use of Hales-Jewett with Theorem 11. □

Remark: In terms of the hypergraph $K_3(G)$ on $E(G)$, this says $\chi(K_3(G)) > k$ and $g(K_3(G)) > g$. So this powerfully extends Theorem 4 on hypergraphs with high girth and chromatic number.

The same proof also gives:

Corollary 13 *For all k, g and r, there exists a G such that $G \xrightarrow{k} K_r$, but with $g(K_r(G)) > g$: G has no short cycles of copies of K_r.* □

What if we want to replace K_r with some general graph H? Do we get, for any k and g, a graph $G \xrightarrow{k} H$, with $g(H(G)) > g$? (Where $H(G)$ is the hypergraph of copies of H in G.)

We need a graph with somewhat delicate properties for the proof of Theorem 12: we needed there to be no $K_3 \subset B$ (where B is our bipartite graph). We also needed each $K_3 - V(B)$ to be connected.

Say that a graph H is **3-chromatically connected** if for any $A \subset V(H)$ with $H[A]$ bipartite, we have $H - A$ connected and non-empty. Then the same proof as for Theorem 12 gives us:

Corollary 14 *Let H be 3-chromatically connected. Then, for all k and g, there is some G such that $G \xrightarrow{k} H$, but $g(H(G)) > g$.* □

Unfortunately, one cannot extend this corollary to all graphs H. For example, suppose the graph H consists of two disjoint edges. Then $g(H(G)) > 3$ says exactly that G contains no set of three disjoint edges. But then it is easy to check that such a G cannot satisfy $G \xrightarrow{k} H$ (for k at least 4).

However, we can get the following, by the same proof as Theorem 12:

Corollary 15 (Weak sparse Ramsey theorem for general H) *For every graph H, and for all k and g, there is a graph G and some $\mathcal{H} \subset H(G)$ such that $\chi(\mathcal{H}) > k$ and $g(\mathcal{H}) > g$.* □

We may now ask what happens when we are asking for no short cycles of *edges*, i.e. for some large g(G). Given H, if we have some G with $G \longrightarrow H$, then it is clearly true that $g(G) \geq g(H)$. Can we have equality? In general we don't know:

Open problem 1 *Given any H with $g(H) < \infty$, is there a G such that $G \longrightarrow H$ but $g(G) = g(H)$?*

2 Sparse Arithmetic Structures

2.1 The Sparse Hales-Jewett Theorem

We shall start our consideration of arithmetic structures by looking for a restricted form of Van der Waerden's theorem: given some k and m, we seek a set $S \subset \mathbb{N}$ such that:

1. If S is k-coloured, there is a monochromatic arithmetic progression of m terms.

2. S does not contain an arithmetic progression of $m+1$ terms.

We can find such an S easily – we prove this result exactly as we deduce Van der Waerden from Hales-Jewett (see Appendix). We embed $[m]^d$ linearly into $\mathbb{N}$ according to:

$$(x_1, \ldots, x_d) \mapsto \sum_i r_i x_i$$

where we choose $r_1 \ll r_2 \ll \cdots \ll r_d$.

What about "restricted Hales-Jewett"? We want some $S \subset A^d$ such that:

1. If S is k-coloured, then there exists a monochromatic combinatorial line.
2. S does not contain any 2-dimensional subspaces (such as $\{(1, x, y, 3, 4, y, y, x) | x, y \in A\}$).

We can do this fairly easily too. It comes out directly from the usual focusing argument proof of Hales-Jewett.

Similarly, for any r we can get an $S \subset A^d$ such that S has no $(r+1)$-dimensional subspace, but such that there is a monochromatic r-dimensional subspace whenever S is k-coloured.

Remark: For the Graham-Rothschild theorem, where we colour lines instead of points, a restricted version *is* genuinely harder.

So, the most interesting theorem to aim for must be *sparse* Hales-Jewett. For all finite sets A, and for all k and g, we want some $S \subset A^d$ such that if S is k-coloured, there is a monochromatic line, but such that S has no cycles of lines of length less than g.

In other words, writing $L(S)$ for the collection of lines in S, we want $\chi(L(S)) > k$ and $g(L(S)) > g$.

Remark: If this is satisfied, then we can get sparse Van der Waerden by the usual method of mapping a cube linearly into $\mathbb{N}$.

As is familiar, we'll start this one with the case $g = 3$: we want a triangle-free version of Hales-Jewett.

We'd expect to start as follows: find, for some t and A, sets $S_1, \ldots, S_t \subset A^d$ such that

1. When $\bigcup S_i$ is k-coloured, with each S_i monochromatic, then there is a monochromatic line.

2. $\bigcup S_i$ has no triangle.

This is easy to find: set $t = nk$ (where $A = [n] = \{1, \ldots, n\}$), and for each n-set $\{i_1, \ldots, i_n\} \subset [t]$ choose (disjointly) a line in A^d meeting each of $S_{i_1}, \ldots, S_{i_n}$.

For $S \subset A^d$, a **copy** of S in A^e is a set S' that is the image of S under a linear embedding $\pi : A^d \to A^e$ (i.e. π is the canonical map from A^d to a canonical d-dimensional subspace of A^e).

For example, for $S \in A^2$ we could have

$$S' = \{(1, 2, 7, x, x, 5, y, x, y, 3) | (x, y) \in S\}.$$

Thus, a copy of A^k means the same as a k-dimensional subspace.

Now, in graphs we can say things like "take a copy with the i-th class equal to A, and do it disjointly". But here, in Hales-Jewett, there's no method for extending a copy of S_i to a copy of $\bigcup_{i=1}^t S_j$ in many ways. For example, if S_i spans A^d (i.e. A^d is the smallest subspace containing S_i), then there is a unique copy of $\bigcup S_j$ with given S_i'. So there is no way to "keep copies disjoint" or anything like that.

One key idea is this: we can weaken the notion of a "copy of $\bigcup S_i$", as long as lines are preserved.

The main ingredient is that we will index our family $\{S_i\}$ now not by $1, \ldots, t$ but by the elements of A^d – thus our indexing will rely on Hales-Jewett itself!

We fix d_0 such that if A^{d_0} is k-coloured, there is a monochromatic line. Then a **picture** S in A^d is defined to be a collection of disjoint sets $S_v \subset A^d$, one for each $v \in A^{d_0}$.

The **underlying set** on S is $\bigcup_{v \in A^{d_0}} S_v \subset A^d$. By an abuse of notation, this is commonly referred to merely as S.

To start, we find a picture S in A^d (for some d) such that:

1. For every line $v_1 < \cdots < v_n$ in A^{d_0}, there is a line $x_1 < \cdots < x_n$ in A^d with $x_i \in S_{v_i}$ for all i.
2. $\bigcup_{v \in A^{d_0}}$ has no triangles.

Remarks:

1. Clearly, such an S exists. For example: for each line in A^{d_0} just choose, disjointly, a line in A^d and assign points accordingly.

2. Thus, if $\bigcup_{v\in A^{d_0}} S_v$ is k-coloured, with each S_v monochromatic, then there is a line L in A^{d_0} with $\bigcup_{v\in L} S_v$ monochromatic (by definition of d_0), so by remark 1 there is a monochromatic line in
$$\bigcup_{v\in L} S_v \subset \bigcup_{v\in A^{d_0}} S_v.$$

3. It is often useful to have no S_v contain a line of A^d. One easy way of ensuring this is to make each S_v an antichain. This means that, for all x, y in S_v, with $x \neq y$, we don't have $x < y$ (where $x < y$ means all of the coordinates of x are less than the corresponding coordinates of y).

4. In choosing our starting S, it is easy to ensure that every S_v is an antichain.

To define a *copy* of a picture S in A^d, we want one motivating example: in $A^{d_0} \times A^d$, let $S'_v = (v, S_v) = \{(v, x) : x \in S_v\}$. Now, $\bigcup S'_v$ is *not* a copy of $\bigcup S_v$, but each S'_v is a copy of S_v.

So we define a **copy** of S to be a picture S' in A^e such that:

1. For all v, S'_v is a copy of S_v.

2. "Lines coming from lines in A^{d_0} are preserved": for any line $v_1, \dots, v_n$ in A^{d_0}, and points $x_1, \dots, x_n$ in A^d that form a line, with $x_i \in S_{v_i}$ for all i, then we also have $x'_1 < \cdots < x'_n$ forming a line in A^e (where x'_i is the image of x_i under the embedding $S_{v_i} \to S'_{v_i}$).

Examples:

1. Fix $u \in A^{d_0}$, then let $S'_v = (u, S_v)$. Here, $\bigcup S'_v$ actually is a copy of $\bigcup S_v$.

2. Let $S_v = (v, S_v)$.

Note: In both these examples, no new lines are created.

Remarks:

1. A copy of a copy of S is itself a copy of S.

2. Let S be a copy of our starting picture $S^{(0)}$. Then, if $\bigcup S_v$ is k-coloured, with each S_v monochromatic, then there is a monochromatic line (by definition of $S^{(0)}$, and by definition of copy).

Given a picture S in A^d, and $u \in A^{d_0}$, we define $S \star S_u$, the **amalgamation** of S over S_u, as follows:

Choose e such that if $(S_u)^e$ is k-coloured, then there is a monochromatic S_u-line (which is a copy of S_u, of course) – we do this by Hales-Jewett on the alphabet S_u.

Naturally, we have $(S_u)^e \subset A^{de}$. Let the S_u-lines in $(S_u)^e$ be $S_u^{(1)}, \ldots, S_u^{(t)}$.

For each $v \in A^{d_0}$, and each $1 \leq i \leq t$, let $S_v^{(i)}$ denote the corresponding copy of S_v (i.e., the image of S_v under the embedding $\pi : A^e \to A^{de}$ that mapped S_u to $S_u^{(i)}$.

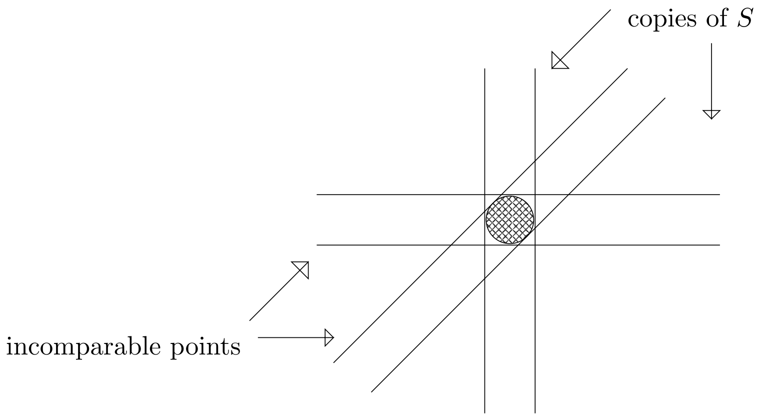

Figure 10: Ensuring no unwanted combinatorial lines.

Inside $(A^{d_0})^t \times A^{de}$, define a copy $S^{(i)}$ of S (for any $1 \leq i \leq t$), as follows:

$$S_v^{(i)} = (u, u, \ldots, u, \underbrace{v}_{i\text{-th coord.}}, u, \ldots, u, S_v).$$

Now let $S \star S_u = \bigcup_{i=1}^t S^{(i)}$; thus $(S \star S_u)_v = \bigcup_{i=1}^t S_v^{(i)}$. Note that

$$(S \star S_u)_u = (u, u, \ldots, u, S_u^e).$$

Proposition 16 *Let S be a picture in A^d, with $u \in A^{d_0}$. Write $S' = S \star S_u$. Then:*

1. *When $\bigcup S'_v$ is k-coloured, there is a copy of S whose u-th class is monochromatic.*

2. *If each S_v is an antichain, and S is Δ-free, then each S'_v is an antichain, and S' is Δ-free.*

Proof:

1. By choice of e and construction of S'. ✓

2. Clearly, each S_v is an antichain (even for $v = u$, since if S_u is an antichain then so is S_u^e).

   Furthermore, the only lines in S' are lines in $S^{(i)}$, for some i. (There are no lines in S_u, since it's an antichain).

   But $S^{(i)}$ is Δ-free, so a triangle in S would be of the form shown in Diagram 11, where i, j, k are not all equal. So some two of a, b, c belong to S'_u, which is contrary to the fact that S'_u is an antichain. ✓ □

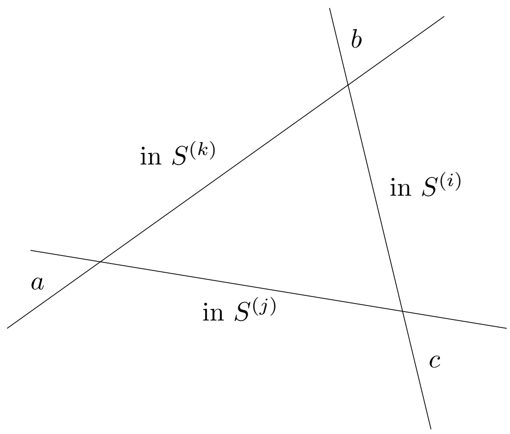

Figure 11: A putative triangle in S.

Theorem 17 *For all k and n, there exists a set $S \in [n]^d$, for some d, such that if S is k-coloured there is a monochromatic line, but such that S is Δ-free.*

Proof: Start with $S^{(0)}$ and amalgamate over each $u \in A^{d_0}$ in turn. We obtain a Δ-free S such that if S is k-coloured, there is a copy of $S^{(0)}$ with all edges monochromatic. □

What about larger girth? In fact, our S_0 above doesn't have any 4-cycles. If it did there would be a line in both $S_u^{(i)}$ and $S_u^{(j)}$, which is forbidden. However, there are 6-cycles in it.

But, magically, our Theorem 17 strengthens itself to Theorem 18:

Theorem 18 *For all k, n and g, there exists a set $S \in [n]^d$, for some d, such that if S is k-coloured then there exists a monochromatic line, but such that S has no cycles of lines of length $\leq g$.*

Remark: Writing $L(S)$ for the set of lines in S, this says that $\chi(L(S)) > k$ and $g(L(S)) > g$. So we get a significant strengthening of our Theorem 4.

Proof: This is by induction on g. We follow the proof of Theorem 17, except we replace S_u^e by a set $T \subset S_u^{e'}$ such that whenever T is k-coloured, we get a monochromatic S_u^e line, with the girth of the S_u-lines in T more than $g/2$.

Note: This was proved by Rödl, and Prömel and Voigt, and Nešetřil and Rödl.

2.2 Other results and open problems

Consider the Finite Sums Theorem: this says that, for all n, whenever $\mathbb{N}$ is finitely coloured there is a sequence $x_1, \ldots, x_n$ such that $\mathrm{FS}(x_1, \ldots, x_n)$, the set of all finite sums of subsets of $\{x_1, \ldots, x_n\}$, is monochromatic.

There is a *restricted* version: that, for all n, there is some set $S \subset \mathbb{N}$ such that whenever S is k-coloured, there is a monochromatic $\mathrm{FS}(x_1, \ldots, x_n)$, but such that S contains no $\mathrm{FS}(x_1, \ldots, x_{n+1})$.

There is also a *sparse* version: where for all n and g, there is some $S \subset \mathbb{N}$ such that S always contains a monochromatic $\mathrm{FS}(x_1, \ldots, x_n)$, but such that S contains no cycles of length less than g of sets $\mathrm{FS}(y_1, \ldots, y_n)$.

The proof of these is similar to the proof of sparse Hales-Jewett (Theorem 18).

For infinite structures, we have Hindman's Theorem: if $\mathbb{N}$ is finitely coloured, then there is a sequence $x_1, x_2, \ldots$ with the set of all finite sums monochromatic.

The fact that it's about *all* finite sums makes it hard. If we just wanted $x_1, x_2, \ldots$ monochromatic, that would be a simple application of the pigeonhole principle. If we just wanted all $x_i + x_j$ monochromatic, that would merely need Ramsey's theorem.

Since any set $\mathrm{FS}(x_1, \ldots, x_2)$ properly contains many other such sets, the natural restricted form would be:

Open problem 2 *Given d, is there a set $S \subset \mathbb{N}$ such that:*

1. *When S is finitely coloured, there is a monochromatic $\mathrm{FS}_{\leq d}(x_1, x_2, \ldots)$ (this means the set of all sums of no more than d terms), but*
2. *S contains no $\mathrm{FS}(y_1, \ldots, y_{d+1})$?*

Indeed, much weaker statements than this are completely unknown: we can't even prove this for $d = 2$ – even if we weaken criterion 2 above to "S contains no $\mathrm{FS}(y_1, y_2, \ldots)$". In other words:

Open problem 3 *Is there a set $S \in \mathbb{N}$ such that:*

1. *Whenever S is finitely coloured, there are $x_1, x_2, \ldots$ such that $\{x_i | i \in \mathbb{N}\} \cup \{x_i + x_j | i \neq j\}$ is monochromatic.*
2. *S contains no $\mathrm{FS}(y_1, y_2, \ldots)$?*

Remark: This has been conjectured true by Nešetřil and Rödl, and conjectured false by Hindman.

To prove it true, we'd need a new proof that when $\mathbb{N}$ is finitely coloured, there is a monochromatic $\mathrm{FS}_{\leq 2}(x_1, x_2, \ldots)$: a proof that doesn't just use Hindman's Theorem.

Appendix - The Hales-Jewett Theorem

We present here a diversion on the Hales-Jewett theorem, since familiarity with it is essential to understanding much of the course.

Van der Waerden's theorem says that, for all k and m, whenever $\mathbb{N}$ is k-coloured, there exists a monochromatic arithmetic progression of m terms.

Hales-Jewett is an abstract version of this: let X be a finite set. A subset L of X^n ("the n-dimensional cube on alphabet X") is called a **line** if there exists a nonempty set $I \subset [n]$, and a_i for each $i \notin I$ such that:

$$L = \{x \in X^n : x_i = a_i \quad (\forall i \notin I), \text{ and } x_i = x_j \quad (\forall i, j \in I).\}$$

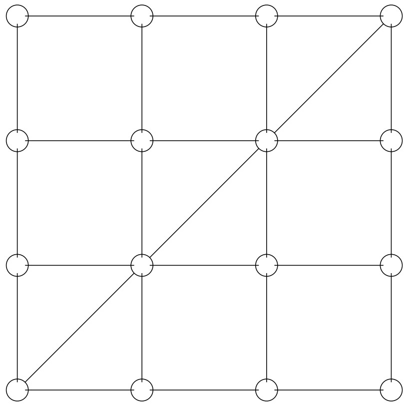

Figure 12: The combinatorial lines in $[4]^2$.

I is called the set of **active coordinates** of the line. Note that the definition of a line does not depend on any ordering on the ground set X. Now we are ready to state the following:

Theorem 19 (Hales-Jewett) *For any m and k, there exists n such that whenever $[m]^n$ is k-coloured, there exists a monochromatic line.*

Remarks:

1. We shall denote the smallest such n in the statement of the problem as $\mathrm{HJ}(m, k)$.
2. Hales-Jewett easily implies Van der Waerden. All we need do is embed a Hales-Jewett cube of sufficiently large dimension linearly into $\mathbb{N}$ so that the embedding is injective on lines. By Hales-Jewett, there is a monochromatic line, and this corresponds to a monochromatic arithmetic progression.

Notation: If L is a line in $[m]^n$, we write L^- and L^+ for its first and last points (i.e. where the active coordinates are 1 and m respectively). We say that lines $L_1, \ldots, L_k$ are **focused** at f if $L_i^+ = f$ for all i.

We say they are **colour-focused** (for a given colouring) if, in addition, each $L_i \backslash \{L_i^+\}$ is monochromatic, with no two the same colour.

Proof: (of Theorem 19). The proof is by induction on m. It is trivial for $m = 1$.

So, given $m > 1$, we may assume $\mathrm{HJ}(m-1, k)$ exists for all k.

Claim: For all $r \leq k$, there is an n such that, when $[m]^n$ is k-coloured, then there is either a monochromatic line of r colour-focused lines.

The result follows immediately from the claim – put $r = k$ and then look at the focus.

Proof of claim: The proof is by induction on r.

For $r = 1$ we can just take $n = \mathrm{HJ}(m-1, k)$.

So suppose n is suitable for r; we'll show that $n + \mathrm{HJ}(m-1, k^{m^n})$ is suitable for $r+1$. We will write $n' = \mathrm{HJ}(m-1, k^{m^n})$.

Given a k-colouring of $[m]^{n+n'}$ with no monochromatic line, identify $[m]^{n+n'}$ with $[m]^n \times [m]^{n'}$. There are k^{m^n} ways to colour a copy of $[m]^n$. So, by our choice of n', we have a line L in $[m]^{n'}$ (say with active coordinates I) such that, for all $a \in [m]^n$ and $b, b' \in L \backslash \{L^+\}$ we have $c(a, b) = c(a, b') = c'(a)$, say.

Now by definition of n, there exist r colour-focused lines for c', say $L_1, \ldots, L_r$ with active coordinates $I_1, \ldots, I_r$ respectively, and focus f. But now let L'_i be the line through the point (L'_i, L^-) with active coordinates $I_i \cup I$, $i = 1, \ldots, r$.

Then $L'_1, \ldots, L'_r$ are colour-focused at (f, L^+). What's more, the line through (f, L^-) with active coordinates I gives us $r+1$ colour-focused lines. Thus our induction is complete – the claim follows. $\square$

Notes by James Cranch.