

ON RAMSEY PROPERTIES OF CLASSES WITH FORBIDDEN TREES

JAN FONIOK

Manchester Metropolitan University, School of Computing, Mathematics and Digital Technology,
Chester Street, Manchester M1 5GD, England
e-mail address: j.foniok@mmu.ac.uk

ABSTRACT. Let \mathcal{F} be a set of relational trees and let $\text{Forb}_h(\mathcal{F})$ be the class of all structures that admit no homomorphism from any tree in \mathcal{F} ; all this happens over a fixed finite relational signature σ . There is a natural way to expand $\text{Forb}_h(\mathcal{F})$ by unary relations to an amalgamation class. This expanded class, enhanced with a linear ordering, has the Ramsey property. Both forbidden trees and Ramsey properties have previously been linked to the complexity of constraint satisfaction problems.

1. INTRODUCTION

Put vaguely, in Ramsey theory one looks for monochromatic subobjects in colourings of large objects. For instance, one might want to prove a statement like this:

Let A, B be digraphs and r an integer. Then there exists a digraph C such that whenever the copies of A in C are coloured with r colours, then there exists a copy B' of B in C such that all the copies of A in B' have the same colour.

It is, however, not hard to show that this statement is false: Let A be a single arc and B the directed 4-cycle. Given any digraph C , the adversary can colour the arcs of C with $r \geq 2$ colours as follows: First, fix an arbitrary linear ordering of the vertex set of C ; then colour every arc of C “red” if it goes “forward” with respect to the ordering on its endpoints, and “blue” if it goes “backward”. Now, no matter how we order the vertices of B it will contain both a forward and a backward arc – thus no copy of B in C can have all its arcs in the same colour class.

This issue can be fixed by considering *ordered* structures: in this case we would consider digraphs with an additional linear ordering of its vertices. Then a “forward” arc and a “backward” arc are distinct, non-isomorphic ordered digraphs and, in fact, the above statement becomes true.

Theorem 1.1 (Nešetřil–Rödl [22]). *Let A, B be ordered digraphs and r an integer. Then there exists an ordered digraph C such that whenever the copies of A in C are coloured with*

2012 ACM CCS: [Mathematics of computing]: Discrete mathematics.

2010 Mathematics Subject Classification: 05C55.

Key words and phrases: forbidden substructure; amalgamation; Ramsey class; partite method.

r colours, then there exists a copy B' of B in C such that all the copies of A in B' have the same colour.

The aim of this paper is to prove analogous results for A, B, C belonging to specific classes of structures. Perhaps the simplest example of such a result is the analogue of Theorem 1.1 where we replace “ordered digraphs” with “ordered K_n -free undirected graphs”, proved by Folkman [11] for $A = K_2, B = K_{n-1}$, as well as for $A = K_1$ and any K_n -free B , and by Nešetřil and Rödl [22, 25] in general. Here we study classes of ordered digraphs (and, more generally, relational structures) obtained not by forbidding one subgraph, such as the complete graph in the example above, but by forbidding all homomorphic images of a given set \mathcal{F} of oriented trees. In this context a *homomorphism* is a mapping that preserves the arcs but it need not preserve the linear ordering.

Thus for a (possibly infinite) set \mathcal{F} of oriented trees, let $\text{Forb}_h(\mathcal{F})$ be the class of all ordered digraphs that admit no homomorphism from any tree in \mathcal{F} . These classes are interesting in the context of constraint satisfaction problems. For a finite digraph H , $\text{CSP}(H)$ denotes the class of all digraphs that admit a homomorphism to H . The case where $\text{CSP}(H) = \text{Forb}_h(\mathcal{F})$ for a set \mathcal{F} of trees corresponds to constraint satisfaction problems with *tree duality* (also known as *width-one* constraint satisfaction problems). Ramsey theory provides a way to recognise digraphs H that define CSPs with tree duality (see Section 8).

Now, however, another issue arises that can be illustrated with this example: Let P_3 be the directed path with three arcs and consider $\mathcal{C} = \text{Forb}_h(\mathcal{F})$ for $\mathcal{F} = \{P_3\}$. For $A = P_1$, an arc ordered forward, and $B = P_2$, the directed path with two arcs $0 \rightarrow 1 \rightarrow 2$ ordered $0 \prec 1 \prec 2$, there can be no P_3 -free C with the Ramsey property for A and B : In any C we can colour an arc “red” if there is another arc going out from its head, and “blue” otherwise. In any copy of $B = P_2$ in C the first arc will be red; if there is a monochromatic copy B' of B , then its second arc must also be coloured red. This implies, however, that there is a homomorphic image of P_3 in C .

The way to tackle this problem is to introduce new unary relations on the vertices in a clever way (determined by the trees in \mathcal{F}). In our example, we would impose a unary relation on a vertex v (let us call the unary relation “square”) whenever there is an arc leaving v , and another unary relation (“circle”) whenever there is a copy of P_2 leaving v . Then there is always both a square and a circle on the starting vertex of P_2 , but never a circle on its middle vertex. Hence the two arcs of P_2 are no longer isomorphic induced subgraphs and we never colour them both: A cannot be an arc both with a circle on its tail and without one.

The precise fashion in which the unary relations are introduced is described in Section 4. Not always is it possible to use only finitely many unary relations. Interestingly, it turns out that a finite number of unary relations suffice if and only if $\text{Forb}_h(\mathcal{F}) = \text{CSP}(H)$ for some finite H .

The main result of this paper is the Ramsey property of any class $\text{Forb}_h(\mathcal{F})$ of ordered relational structures expanded by a number of unary relations, with \mathcal{F} being a set of relational trees. The setting is properly defined in Section 2; the unary relations are introduced in Section 4. Section 5 presents the main result, which is then proved in Sections 6 and 7. Section 8 describes a link between constraint satisfaction problems with tree duality and our Ramsey classes; the paper then concludes with a number of final comments.

2. BASIC DEFINITIONS

Relational structures. A *signature* σ is a set of relation symbols; each of the symbols has an associated *arity*; the arity of R is $\text{ar}(R)$. A σ -*structure* A is a set of *elements*, called the *domain* of A and denoted $\text{dom } A$, together with a relation R^A of arity $\text{ar}(R)$ on the domain for every relation symbol $R \in \sigma$. Unless specifically stated otherwise, all structures we deal with in this paper have finite domain. We allow the domain to be empty. An *ordered* σ -*structure* is a $(\sigma \cup \{\preceq\})$ -structure A such that \preceq^A is a linear ordering. A σ -structure A is a *substructure* of a σ -structure B if $\text{dom } A \subseteq \text{dom } B$ and for each $R \in \sigma$ we have $R^A = R^B \cap (\text{dom } A)^{\text{ar}(R)}$. We write $A \subseteq B$ if A is a substructure of B . Note that our substructure would be called an *induced substructure* in some literature.

An *embedding* of A into B is a one-to-one mapping $f : \text{dom } A \rightarrow \text{dom } B$ such that for any $R \in \sigma$ and any tuple \bar{x} we have $\bar{x} \in R^A$ iff $f(\bar{x}) \in R^B$, where f is applied on \bar{x} component-wise. We write $f : A \xrightarrow{e} B$ to indicate that f is an embedding. A bijective embedding is called an *isomorphism*. We write $A \cong B$ to indicate the existence of an isomorphism between A and B .

If $\sigma \subset \tau$, the σ -*reduct* of a τ -structure A is the σ -structure A^* obtained from A by leaving out all the relations R^A for $R \in \tau \setminus \sigma$. Then A is a τ -*expansion* of A^* . (In some literature a reduct is called a *shadow* and an expansion is called a *lift*.)

Ramsey classes. For any structures A, B , let $\binom{B}{A}$ denote the set of all embeddings of A into B . The partition arrow $C \rightarrow \binom{B}{A}_r$ means that for any $\chi : \binom{C}{A} \rightarrow \{1, \dots, r\}$ (called a *colouring* with r colours) there exists $g \in \binom{C}{B}$ and $j \leq r$ such that $\chi(h) = j$ for all $h \in \binom{g[B]}{A}$. In this case we call g (or $g[B]$) a *monochromatic copy* of B in C . Here $g[B]$ is the substructure of C induced by the range of g , which is isomorphic to B because g is an embedding. Hence $\binom{g[B]}{A} = \{g \circ f : f \in \binom{B}{A}\}$.

Let \mathcal{C} be a class of finite structures and let $A \in \mathcal{C}$. The class \mathcal{C} has the *A-Ramsey property* if for any $B \in \mathcal{C}$ and any natural number r there exists $C \in \mathcal{C}$ such that $C \rightarrow \binom{B}{A}_r$. The class \mathcal{C} is called a *Ramsey class* if it has the *A-Ramsey property* for all $A \in \mathcal{C}$.

Using the above terminology, Theorem 1.1 is a special case of the following:

Theorem 2.1 (Nešetřil–Rödl [22]). *Let σ be a finite relational signature. Then the class of all finite ordered σ -structures is a Ramsey class.*

The presence of orderings is indeed essential; cf. the discussion in [19]. Here we only note that for ordered structures A, B there is a one-to-one correspondence between the embeddings of A into B and substructures of B isomorphic to A because an ordered structure has no non-trivial automorphisms.

Classes with forbidden homomorphic images. Let A, B be σ -structures. A *homomorphism* of A to B is a mapping $f : \text{dom } A \rightarrow \text{dom } B$ such that for any $R \in \sigma$ and any $\bar{x} \in R^A$ we have $f(\bar{x}) \in R^B$. We write $f : A \xrightarrow{h} B$ to indicate that f is a homomorphism.

The interest of this paper lies in classes of finite σ -structures that can be defined by forbidding the existence of a homomorphism from a given set of structures. More explicitly, for a set \mathcal{F} of σ -structures let $\text{Forb}_h(\mathcal{F})$ be the class of all finite σ -structures A such that

whenever $F \in \mathcal{F}$, there exists no homomorphism of F to A . In this case we say that A is \mathcal{F} -free.

Conventions.

1. A tuple has a bar, so $\bar{x} = (x_1, x_2, \dots, x_k)$ for some k . If M is the domain of some function f and $\bar{x} \in M^k$, then $f(\bar{x}) = (f(x_1), f(x_2), \dots, f(x_k))$.
2. Instead of “substructure of X generated by M ” I write “substructure of X induced by M ” with the intended connotation that the domain of such a substructure is actually M . This is the case because our structures have no operations.
3. For a $(\sigma \cup \tau)$ -structure A , A^* almost always denotes the σ -reduct of A .
4. Usually $R \in \sigma$ and $S \in \tau$, but sometimes $R \in \sigma \cup \tau$.
5. I treat mappings and homomorphisms in a more set-theoretic rather than category-theoretic way. For example, if $f : A \xrightarrow{h} B$ is a homomorphism and $B \subseteq C$, then also $f : A \xrightarrow{h} C$. Similarly, if $f : A \xrightarrow{h} B$ is a homomorphism of $(\sigma \cup \tau)$ -structures, then the same f is a homomorphism of their σ -reducts; $f : A^* \xrightarrow{h} B^*$.

3. AMALGAMATION AND OTHER CONSTRUCTIONS

A Ramsey class of structures always has the *amalgamation property* defined below (see [19]). Most classes of the form $\text{Forb}_h(\mathcal{F})$ do not have the amalgamation property, but following Hubička–Nešetřil [15, 16] there is a *canonical way* to add new relations to the signature σ in order to obtain it. It is this *expanded class*, enhanced with a linear ordering, which is a Ramsey class.

Amalgamation. A class \mathcal{C} of finite σ -structures has the *joint-embedding property* if for any structures $A_1, A_2 \in \mathcal{C}$ there exists $B \in \mathcal{C}$ such that both A_1 and A_2 admit an embedding into B . A class \mathcal{C} of finite σ -structures has the *amalgamation property* if for any $A, B_1, B_2 \in \mathcal{C}$ and any embeddings $f_1 : A \xrightarrow{e} B_1$ and $f_2 : A \xrightarrow{e} B_2$ there exists $C \in \mathcal{C}$ and embeddings $g_1 : B_1 \xrightarrow{e} C$ and $g_2 : B_2 \xrightarrow{e} C$ such that $g_1 \circ f_1 = g_2 \circ f_2$. The amalgamation is *free* if $\text{dom } C = g_1[\text{dom } B_1] \cup g_2[\text{dom } B_2]$ and $R^C = g_1[R^{B_1}] \cup g_2[R^{B_2}]$ for all $R \in \sigma$. The amalgamation property implies the joint embedding property if \mathcal{C} contains the empty structure.

Sum. For two σ -structures A, B , their *sum* $A + B$ is defined by

$$\begin{aligned} \text{dom}(A + B) &= (\{A\} \times \text{dom } A) \cup (\{B\} \times \text{dom } B), \\ R^{A+B} &= (\{A\} \otimes R^A) \cup (\{B\} \otimes R^B), \end{aligned}$$

where

$$\{X\} \otimes R^X = \{((X, x_1), (X, x_2), \dots, (X, x_k)) : (x_1, x_2, \dots, x_k) \in R^X\}.$$

Assuming that the domains of A and B are disjoint, we could take the union $\text{dom } A \cup \text{dom } B$ to be $\text{dom}(A + B)$ and each relation R^{A+B} to be the union of the respective relations R^A and R^B . However, it will be convenient explicitly to mark which summand an element of the sum originates from. The definition can be extended to arbitrary finite sums in the obvious

way. We may also write $\coprod\{A_1, A_2, \dots, A_k\}$ for $A_1 + A_2 + \dots + A_k$. Note that the sum of σ -structures is the coproduct in the category of σ -structures and their homomorphisms.

Note further that if a class is closed under taking sums, then it has the joint-embedding property. The converse is not true; for instance, the class of complete graphs has the joint-embedding property but is not closed under taking sums.

Connected structures. A σ -structure A is *connected* if, whenever $A \cong A_1 + A_2$, either $\text{dom } A_1 = \emptyset$ or $\text{dom } A_2 = \emptyset$. (This corresponds to *weak connectedness* of digraphs.)

Incidence graph. The *incidence graph* $\text{Inc}(A)$ of a σ -structure A is the bipartite undirected multigraph whose vertex set is $\text{dom } A \cup \bigcup\{R^A \times \{R\} : R \in \sigma\}$, and which contains for every $R \in \sigma$, every $\bar{x} \in R^A$, and every i , an edge joining (\bar{x}, R) and x_i .

Gaifman graph. The *Gaifman graph* $\text{Gai}(A)$ of a σ -structure A is the graph whose vertex set is $\text{dom } A$ and there is an edge joining x and y if x, y are distinct elements of A that appear in a common tuple of some relation of A , that is,

$$E(\text{Gai}(A)) = \{\{x, y\} : x \neq y \text{ and } \exists R \in \sigma \exists \bar{v} \in R^A : x, y \in \bar{v}\}.$$

Thus every tuple of every R^A is represented by a clique in $\text{Gai}(A)$.

Lemma 3.1 (see [12]). *For a σ -structure A , the following are equivalent:*

- (a) A is connected.
- (b) Whenever $A \cong A_1 + A_2$, then $A \cong A_1$ or $A \cong A_2$.
- (c) $\text{Inc}(A)$ is connected in the graph-theoretic sense.
- (d) $\text{Gai}(A)$ is connected in the graph-theoretic sense.

Moreover, if A is connected and there is a homomorphism $f : A \xrightarrow{h} A_1 + A_2$, then either $f : A \xrightarrow{h} A_1$ or $f : A \xrightarrow{h} A_2$. □

Trees. A σ -structure A is a σ -tree (or just a tree) if $\text{Inc}(A)$ is a tree. (Thus in particular A is *not* a tree if some tuple of some relation of A contains the same element more than once.)

Factor structure. If A is a σ -structure and \sim is an equivalence relation on $\text{dom } A$, let the *factor structure* A/\sim be defined on $\text{dom}(A/\sim) = (\text{dom } A)/\sim$ (the set of all equivalence classes of \sim) by letting $(X_1, X_2, \dots, X_k) \in R^{A/\sim}$ if and only if there exist $x_1 \in X_1, x_2 \in X_2, \dots, x_k \in X_k$ such that $(x_1, x_2, \dots, x_k) \in R^A$. Informally, A/\sim is formed from A by identifying all the elements in any one \sim -equivalence class into one element of A/\sim .

Join of rooted structures. A *rooted σ -structure* is a couple (A, a) where A is a σ -structure and $a \in \text{dom } A$. Let (A, a) and (B, b) be rooted σ -structures. The *join* $(A, a) \oplus (B, b)$ is the factor structure $(A + B)/\sim$, where \sim is the equivalence relation on $\text{dom}(A + B)$ such that $(A, a) \sim (B, b)$, and $(X, x) \sim (X', x')$ iff $(X, x) = (X', x')$ otherwise. (Please do not get confused that, as a coincidence, (A, a) denotes a rooted structure on one occasion, and an element of the sum $A + B$ on another.) In other words, the join is obtained from the disjoint union of A and B by identifying a and b . In the obvious way, this definition can be extended to joins of more than two structures, too.

4. THE EXPANDED CLASS

From now on, we work in the following setting: σ is a finite relational signature; \mathcal{F} is a possibly infinite set of finite σ -trees. Recall that $\text{Forb}_h(\mathcal{F})$ is the class of all σ -structures that admit no homomorphism from any $F \in \mathcal{F}$. We are aiming to add new unary relations, possibly infinitely many, to get a class \mathcal{C} of $(\sigma \cup \tau)$ -structures such that

- (1) $\text{Forb}_h(\mathcal{F})$ would be the class of the σ -reducts of all structures in \mathcal{C} ;
- (2) \mathcal{C} would have countably many isomorphism classes;
- (3) \mathcal{C} would be closed under taking substructures (this is called the *hereditary property*);
- (4) \mathcal{C} would have the amalgamation property;
- (5) \mathcal{C} would have the Ramsey property.

This can be achieved in a trivial way, by taking a new unary relation for every element of every member of $\text{Forb}_h(\mathcal{F})$ (only considering one from each isomorphism class). We should therefore strive to

- (6) make τ “as small as possible”; in particular, τ should be finite whenever possible.

Pieces of trees. Let $F \in \mathcal{F}$. An element m of F is a *cut* of F if it is a vertex cut of $\text{Gai}(F)$. Note that, as F is a tree, m is a cut of F iff m belongs to more than one tuple of the relations of F .

Let m be a cut of F and let $D \subset \text{dom } F$ be the vertex set of some connected component of $\text{Gai}(F) \setminus \{m\}$. The rooted σ -structure (M, m) , where M is the substructure of F induced by $D \cup \{m\}$, is called a *piece* of F .

Remarks 4.1. 1. A piece of F is a non-empty connected substructure of F , $M \neq F$, and $\{m\} \neq \text{dom } M$. Moreover, M is a σ -tree.

2. For any given cut m of F , the corresponding pieces cover $\text{dom } F$. In other words,

$$F = \bigoplus \{(M, m) : (M, m) \text{ is a piece of } F\}$$

for any fixed cut m of F .

Equivalence of pieces. Let (A, a) be a rooted σ -structure. Following [16], let

$$\mathfrak{I}(A, a) = \{(B, b) : \text{there exists } F \in \mathcal{F} \text{ s.t. } (A, a) \oplus (B, b) \cong F\}$$

be the set of all rooted σ -structures that are *incompatible* with (A, a) . For two pieces (M, m) and (M', m') we say that they are *equivalent* and write $(M, m) \approx (M', m')$ if $\mathfrak{I}(M, m) = \mathfrak{I}(M', m')$. Let $\mathcal{P}(\mathcal{F})$ be the set of all \approx -equivalence classes of all pieces of all trees in \mathcal{F} , that is,

$$\mathcal{P}(\mathcal{F}) = \{(M, m) : \exists F \in \mathcal{F} \text{ s.t. } (M, m) \text{ is a piece of } F\} / \approx.$$

Example 4.2. Oriented paths can be encoded by words over the alphabet $\{0, 1\}$, so that 0 denotes a forward arc and 1 denotes a backward arc. Using this encoding, let \mathcal{F} be the set of oriented paths encoded by the words 000, 00100, 0010100, 001010100, ... (see Figure 1). Let us call these paths *thunderbolts*. On the pieces of thunderbolts, the equivalence relation \approx has four equivalence classes, as shown in the figure.

For oriented paths, it can be shown that the number of \approx -equivalence classes is finite if the set \mathcal{F} can be described by a regular language (as above): see [9].




Figure 1: Thunderbolts; all arcs are oriented upwards.

Subpieces. A *subpiece* of a piece (M, m) of F is any piece (M', m') of F such that M' is a substructure of M .

Lemma 4.3 ([16]). *Let $F_1 \in \mathcal{F}$. Suppose that (M_1, m_1) is a piece of F_1 , (M'_1, m'_1) is a subpiece of (M_1, m_1) , and (M'_2, m'_2) is a piece such that $(M'_1, m'_1) \approx (M'_2, m'_2)$. Create (M_2, m_2) from (M_1, m_1) by replacing the subpiece (M'_1, m'_1) with (M'_2, m'_2) , identifying m'_1 with m'_2 ; see Figure 2. Then (M_2, m_2) is isomorphic to a piece of some $F_2 \in \mathcal{F}$ and $(M_1, m_1) \approx (M_2, m_2)$. \square*

Expansion. Now we define an expanded signature $\sigma \cup \tau$, aiming to get a class \mathcal{C} of $(\sigma \cup \tau)$ -structures with amalgamation, such that the σ -reducts of all the structures in \mathcal{C} form exactly the class $\text{Forb}_h(\mathcal{F})$.

Definition 4.4. Let τ contain a unary relation symbol $S_{\mathfrak{M}}$ for each $\mathfrak{M} \in \mathcal{P}(\mathcal{F})$. Let $\tilde{\mathcal{C}}$ be the class of finite $(\sigma \cup \tau)$ -structures such that A belongs to $\tilde{\mathcal{C}}$ if and only if the σ -reduct A^* of A is in $\text{Forb}_h(\mathcal{F})$ and for any $\mathfrak{M} \in \mathcal{P}(\mathcal{F})$ and any $x \in \text{dom } A$ we have

$$x \in S_{\mathfrak{M}}^A \iff \exists (M, m) \in \mathfrak{M}, \exists f : M \xrightarrow{h} A^* \text{ with } f(m) = x. \quad (4.1)$$

Let \mathcal{C} be the class of all substructures of the structures in $\tilde{\mathcal{C}}$. The class \mathcal{C} is called the *expanded class* for $\text{Forb}_h(\mathcal{F})$. The structures in $\tilde{\mathcal{C}}$ are called *canonical*. We can also say that A is \mathcal{F} -free if $A^* \in \text{Forb}_h(\mathcal{F})$; so being \mathcal{F} -free is a necessary but not sufficient condition for membership in \mathcal{C} .

Lemma 4.5. *Every structure in \mathcal{C} satisfies the right-to-left implication in (4.1).*

Proof. Let $A \in \mathcal{C}$, $x \in \text{dom } A$, $(M, m) \in \mathfrak{M} \in \mathcal{P}(\mathcal{F})$, $f : M \xrightarrow{h} A^*$ with $f(m) = x$. As $A \in \mathcal{C}$, A is a substructure of some canonical $\tilde{A} \in \tilde{\mathcal{C}}$. Then the same mapping f is a homomorphism of M to \tilde{A}^* . By definition, \tilde{A} satisfies (4.1), so $x \in S_{\mathfrak{M}}^{\tilde{A}}$. Hence $x \in S_{\mathfrak{M}}^A$. \square




Figure 2: Lemma 4.3

Canonising. Given a $(\sigma \cup \tau)$ -structure A , we want to find a superstructure \tilde{A} of A that satisfies the left-to-right implication of (4.1). This is possible assuming that

$$\text{every one-element substructure of } A \text{ is in } \mathcal{C}. \quad (4.2)$$

Lemma 4.6. *If A is a $(\sigma \cup \tau)$ -structure that satisfies (4.2), then there is a $(\sigma \cup \tau)$ -structure \tilde{A} such that A is a substructure of \tilde{A} and \tilde{A} satisfies (4.2) as well as the left-to-right implication of (4.1).*

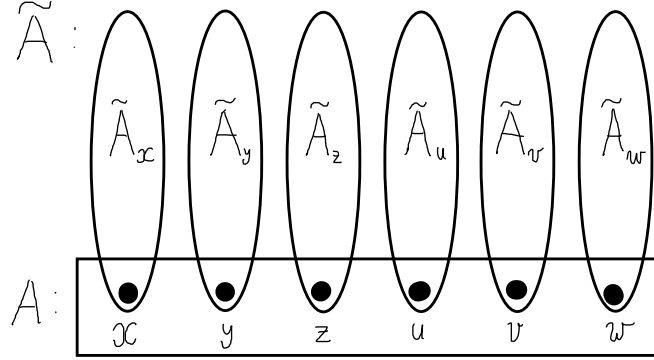
Proof. For every $x \in \text{dom } A$, let A_x be the one-element substructure of A induced by $\{x\}$. By assumption (4.2), for every x we have $A_x \in \mathcal{C}$; so there exists $\tilde{A}_x \in \tilde{\mathcal{C}}$ containing A_x . Let

$$A' = A + \coprod\{\tilde{A}_x : x \in \text{dom } A\}$$

and let \sim be the smallest equivalence relation on $\text{dom } A'$ such that $(A, x) \sim (\tilde{A}_x, x)$ for all $x \in \text{dom } A$. Let $\tilde{A} = A'/\sim$. Informally, \tilde{A} is obtained from A by gluing the corresponding \tilde{A}_x on each element x of A (see Figure 3).

Now each \tilde{A}_x is isomorphic to a substructure of \tilde{A} . Moreover, any element of \tilde{A} is of the form (\tilde{A}_x, y) for some $x \in \text{dom } A$ and some $y \in \text{dom } \tilde{A}_x$. Thus each one-element substructure of \tilde{A} is isomorphic to a substructure of $\tilde{A}_x \in \tilde{\mathcal{C}}$; hence it belongs to \mathcal{C} .

Let $(\tilde{A}_x, y) \in S_{\mathfrak{M}}^{\tilde{A}}$ for some $\mathfrak{M} \in \mathcal{P}(\mathcal{F})$. Then $y \in S_{\mathfrak{M}}^{\tilde{A}_x}$. Since $\tilde{A}_x \in \tilde{\mathcal{C}}$, there exists $(M, m) \in \mathfrak{M}$ and a homomorphism $f : M \xrightarrow{h} \tilde{A}_x^*$ with $f(m) = y$. For any $n \in \text{dom } M$, put $f'(n) = (\tilde{A}_x, f(n))$. Then $f' : M \xrightarrow{h} \tilde{A}^*$ is a homomorphism with $f'(m) = (\tilde{A}_x, f(m)) = (\tilde{A}_x, y)$. Therefore \tilde{A} satisfies the left-to-right implication of (4.1) as we wanted to show. \square

Figure 3: \tilde{A} .

Proving membership in \mathcal{C} . By definition, the class \mathcal{C} (but not $\tilde{\mathcal{C}}$) is hereditary; thus it is defined by a list of forbidden substructures. The goal in the remainder of this section is to find an implicit description of these forbidden substructures: we give a description in terms of one-element and one-tuple structures. This description is going to prove very useful in the sequel: it provides an interface between the forbidden trees, pieces and unary relations on the one hand, and the heavy machinery of Ramsey theory on the other.

Lemma 4.7. *Let $E = E(F, m)$ be the $(\sigma \cup \tau)$ -structure obtained from some $F \in \mathcal{F}$ with a cut $m \in \text{dom } F$ so that*

$$\text{dom } E = \{1\};$$

for each $S_{\mathfrak{M}} \in \tau$, $1 \in S_{\mathfrak{M}}^E$ iff \mathfrak{M} contains a piece (M, m) of F ;

all the σ -relations of E are empty.

If A is a $(\sigma \cup \tau)$ -structure such that there exists a homomorphism $f : E \xrightarrow{h} A$, then $A \notin \mathcal{C}$.

Note. In an informal way this lemma says that as soon as there are “too many” τ -relations on an element of a structure, this structure is not in \mathcal{C} . The forbidden substructures it describes are the one-element structures with all the unary σ - and τ -relations corresponding to a fixed tree $F \in \mathcal{F}$ and a fixed cut m of F , as well as one-element structures with a superset of these relations.

Proof. Let $F \in \mathcal{F}$, let $m \in \text{dom } F$ be a cut of F , and consider $E = E(F, m)$. Let the corresponding pieces of F be $(M_1, m), (M_2, m), \dots, (M_k, m)$.

Let A be a $(\sigma \cup \tau)$ -structure. For the sake of contradiction, suppose that there is a homomorphism $f : E \xrightarrow{h} A$ but $A \in \mathcal{C}$. Then there is a canonical superstructure $\tilde{A} \in \tilde{\mathcal{C}}$ of A .

For each $i = 1, 2, \dots, k$ let $\mathfrak{M}_i \in \mathcal{P}(\mathcal{F})$ be the \approx -equivalence class of the piece (M_i, m) . By definition, $1 \in S_{\mathfrak{M}_i}^E$. As f is a homomorphism, we have $f(1) \in S_{\mathfrak{M}_i}^A \subseteq S_{\mathfrak{M}_i}^{\tilde{A}}$. Since \tilde{A} is canonical, by (4.1) there exists a piece $(N_i, n_i) \in \mathfrak{M}_i$ and a homomorphism $g_i : N_i \xrightarrow{h} \tilde{A}^*$ with $g_i(n_i) = f(1)$. The union of all the homomorphisms g_i , $i = 1, \dots, k$, is a homomorphism g of $(N_1, n_1) \oplus \dots \oplus (N_k, n_k)$ to \tilde{A}^* .

Now let

$$F_0 = (M_1, m) \oplus \dots \oplus (M_k, m) \cong F.$$

For $i \geq 1$, put

$$F_i = (N_1, n_1) \oplus (N_2, n_2) \oplus \cdots \oplus (N_i, n_i) \oplus (M_{i+1}, m) \oplus \cdots \oplus (M_k, m).$$

We prove by induction on i that each F_i is isomorphic to a member of \mathcal{F} . In the rest of this proof, let “ $\in \mathcal{F}$ ” mean “is isomorphic to a member of \mathcal{F} ”, to simplify the notation. Clearly, $F_0 \in \mathcal{F}$ because $F_0 \cong F$. For $i \geq 1$, let

$$\bar{M}_i = (N_1, n_1) \oplus \cdots \oplus (N_{i-1}, n_{i-1}) \oplus (M_{i+1}, m) \oplus \cdots \oplus (M_k, m),$$

whence $F_{i-1} = (\bar{M}_i, m) \oplus (M_i, m)$ and $F_i = (\bar{M}_i, m) \oplus (N_i, n_i)$. Assuming that $F_{i-1} \in \mathcal{F}$, we thus get that $(\bar{M}_i, m) \in \mathfrak{J}(M_i, m) = \mathfrak{J}(N_i, n_i)$ because $(M_i, m) \approx (N_i, n_i)$. Hence $F_i \in \mathcal{F}$.

We conclude that $F_k \in \mathcal{F}$ and $g : F_k \xrightarrow{h} \tilde{A}^*$ is a homomorphism; hence \tilde{A} is not \mathcal{F} -free – a contradiction with the assumption that $\tilde{A} \in \tilde{\mathcal{C}}$. \square

Tuple traces. Let A be a $(\sigma \cup \tau)$ -structure, and $R \in \sigma$. The *tuple trace* of some $\bar{x} = (x_1, x_2, \dots, x_k) \in R^A$ is the structure $T = T(A, \bar{x}, R)$ with $\text{dom } T = \{1, 2, \dots, k\}$; $R^T = \{(1, 2, \dots, k)\}$; $\check{R}^T = \{j : x_j \in \check{R}^A\}$ for all unary $\check{R} \in \sigma$; $R'^T = \emptyset$ for any other $R' \in \sigma \setminus \{R\}$; $S^T = \{j : x_j \in S^A\}$ for $S \in \sigma$.

Example 4.8. For σ containing a quaternary relation symbol R and a binary R' , and τ containing two unary relation symbols \diamond and \odot , let A be the structure with $\text{dom } A = \{a, b, c\}$, $R^A = \{(a, b, b, c)\}$, $R'^A = \{(a, c)\}$, $\diamond^A = \{a, b\}$ and $\odot^A = \{c\}$. The tuple trace of $(a, b, b, c) \in R^A$ is the structure T with $\text{dom } T = \{1, 2, 3, 4\}$, $R^T = \{(1, 2, 3, 4)\}$, $R'^T = \emptyset$, $\diamond^T = \{1, 2, 3\}$, $\odot^T = \{4\}$ (see Figure 4).




Figure 4: Tuple trace for the example just before Lemma 4.9. Think of “unfolding” the tuple $\bar{x} = (a, b, b, c) \in R^A$, while keeping all the unary relations but dropping all the non-unary ones.

Lemma 4.9. *Let A be a $(\sigma \cup \tau)$ -structure. Then $A \in \mathcal{C}$ if and only if each one-element substructure of A belongs to \mathcal{C} , and for any $R \in \sigma$ and any $\bar{x} \in R^A$, the tuple trace of \bar{x} belongs to \mathcal{C} .*

Proof. If $A \in \mathcal{C}$, then each one-element substructure of A is in \mathcal{C} as well because \mathcal{C} is hereditary. Let $A \subseteq \tilde{A} \in \tilde{\mathcal{C}}$. Consider any $R \in \sigma$ and $\bar{x} = (x_1, \dots, x_k) \in R^A \subseteq R^{\tilde{A}}$. Let $T = T(A, \bar{x}, R)$ be the tuple trace of $\bar{x} \in R^A$, which is, in fact, equal to $T(\tilde{A}, \bar{x}, R)$. Let T' be the sum of T and k copies of \tilde{A} ; let \sim be the smallest equivalence relation that identifies $j \in \text{dom } T$ with x_j in the j th copy of \tilde{A} . Let $\tilde{T} = T'/\sim$. There is an obvious “projection” or “folding” homomorphism $p : \tilde{T} \xrightarrow{h} \tilde{A}$; the image under p of T is the substructure of \tilde{A} induced by $\{x_1, x_2, \dots, x_k\}$ and p restricted to any of the k copies of \tilde{A} in \tilde{T} is an isomorphism. Note that the same mapping p is a homomorphism of the σ -reduct \tilde{T}^* to \tilde{A}^* . See Figure 5.

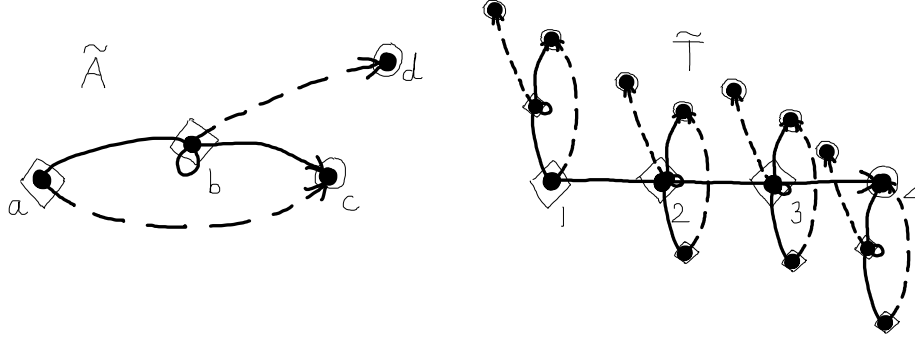


Figure 5: \tilde{T} for $\bar{x} \in R^A$ from Figure 4. \tilde{T} is obtained from the tuple trace T by gluing a copy of \tilde{A} on each element of T .

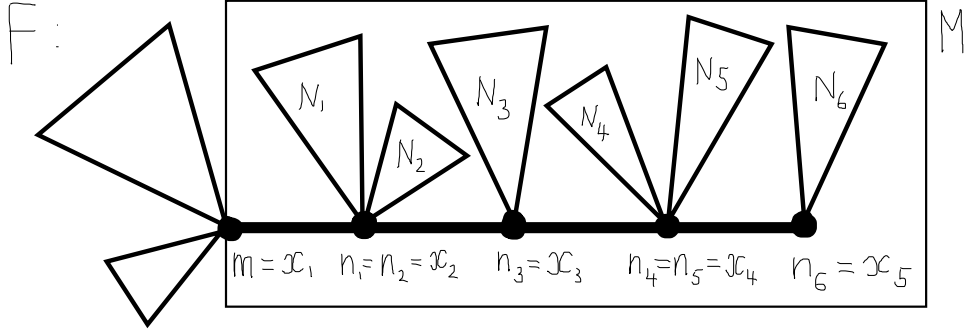
If $F \in \mathcal{F}$ and $f : F \xrightarrow{h} \tilde{T}^*$, then $p \circ f : F \xrightarrow{h} \tilde{A}^*$, a contradiction with $A \in \mathcal{C}$. Thus \tilde{T} is \mathcal{F} -free. To show that \tilde{T} satisfies (4.1), first let $(M, m) \in \mathfrak{M} \in \mathcal{P}(\mathcal{F})$ and let $g : M \xrightarrow{h} \tilde{T}^*$. Since \tilde{A} satisfies (4.1) and $p \circ g : M \xrightarrow{h} \tilde{A}^*$ is a homomorphism, we have $p(g(m)) \in S_{\mathfrak{M}}^{\tilde{A}}$. Hence, by the definition of \tilde{T} , we have $g(m) \in S_{\mathfrak{M}}^{\tilde{T}}$. Conversely, if $x \in \text{dom } \tilde{T}$ satisfies $x \in S_{\mathfrak{M}}^{\tilde{T}}$ for some $\mathfrak{M} \in \mathcal{P}(\mathcal{F})$, then $p(x) \in S_{\mathfrak{M}}^{\tilde{A}}$, thus there exist $(M, m) \in \mathfrak{M}$ and a homomorphism $h : M \xrightarrow{h} \tilde{A}^*$ such that $h(m) = p(x)$. Mapping each element a of M to the element corresponding to $h(a)$ in the copy of \tilde{A} within \tilde{T} that contains x provides a homomorphism from M to \tilde{T} that takes m to x . Therefore not only \tilde{T} is \mathcal{F} -free but it satisfies (4.1) as well, so $\tilde{T} \in \tilde{\mathcal{C}}$. The tuple trace T , which is a substructure of \tilde{T} , then belongs to \mathcal{C} .

The converse implication: Suppose that A satisfies (4.2) and all its tuple traces belong to \mathcal{C} . By Lemma 4.6 there is a $(\sigma \cup \tau)$ -structure \tilde{A} such that A is a substructure of \tilde{A} and \tilde{A} satisfies (4.2) and the left-to-right implication of (4.1). Observe that any tuple trace of \tilde{A} , as described in the proof of Lemma 4.6, is equal to a tuple trace of A (hence in \mathcal{C} by assumption) or to a tuple trace of some $\tilde{A}_x \in \tilde{\mathcal{C}}$ (hence in \mathcal{C} by the first implication of this lemma, which we have just proved). Thus we may assume that the tuple traces of \tilde{A} belong to \mathcal{C} .

To prove that \tilde{A} satisfies the right-to-left implication of (4.1), let \tilde{A}^* be the σ -reduct of \tilde{A} , let $\mathfrak{M} \in \mathcal{P}(\mathcal{F})$ and let $(M, m) \in \mathfrak{M}$ be a piece of some $F \in \mathcal{F}$ and consider any homomorphism $f : M \xrightarrow{h} \tilde{A}^*$. We want to show that $f(m) \in S_{\mathfrak{M}}^{\tilde{A}}$. For the sake of contradiction, assume that $f(m) \notin S_{\mathfrak{M}}^{\tilde{A}}$ and that M is a minimal such piece, that is, we assume that whenever (N, n) is a subpiece of (M, m) , $(N, n) \in \mathfrak{N} \in \mathcal{P}(\mathcal{F})$, then $f'(n) \in S_{\mathfrak{N}}^{\tilde{A}}$ for any homomorphism $f' : N \xrightarrow{h} \tilde{A}^*$.

Because (M, m) is a piece, m belongs to a unique tuple $\bar{x} \in R^M$ with $k = \text{ar}(R) > 1$. Let $(N_1, n_1), (N_2, n_2), \dots, (N_\ell, n_\ell)$ be all the pieces of F corresponding to all cuts x_i of F , $x_i \neq m$, such that $m \notin \text{dom } N_j$ for any j . Each (N_j, n_j) is a subpiece of (M, m) and each $n_j \in \bar{x}$ (see Figure 6). Let $\iota(j)$ be the index for which $n_j = x_{\iota(j)}$, and let $\iota(0)$ be the index for which $m = x_{\iota(0)}$. Moreover, let $\mathfrak{N}_j \in \mathcal{P}(\mathcal{F})$ be the \approx -equivalence class of (N_j, n_j) .

Consider the tuple trace $T = T(\tilde{A}, f(\bar{x}), R)$ of $f(\bar{x}) \in R^{\tilde{A}}$. As we have observed, $T \in \mathcal{C}$, hence there exists canonical $\tilde{T} \in \tilde{\mathcal{C}}$ such that T is a substructure of \tilde{T} . By definition, $\text{dom } T = \{1, 2, \dots, k\} \subseteq \text{dom } \tilde{T}$, and for each j we have $i \in S_{\mathfrak{N}_j}^{\tilde{T}}$ if and only if $f(x_i) \in S_{\mathfrak{N}_j}^{\tilde{A}}$.

Figure 6: In this example, $\ell = 6$, $k = 5$.

By minimality of counterexample, however, we have $\iota(j) \in S_{\mathfrak{M}_j}^{\tilde{T}}$ for each $j \in \{1, \dots, \ell\}$, as each $f_j = f \upharpoonright N_j : N_j \xrightarrow{h} \tilde{A}^*$ is a homomorphism with $f_j(n_j) = f(x_{\iota(j)})$. Because \tilde{T} is canonical, that is, it satisfies (4.1), for each j there exists a homomorphism g_j of some N'_j to \tilde{T}^* with $(N'_j, n'_j) \in \mathfrak{N}_j$ and $g_j(n'_j) = \iota(j)$.

Let (M', m) be obtained from (M, m) by replacing each subpiece (N_j, n_j) with the piece (N'_j, n'_j) . By Lemma 4.3, (M', m) is a piece and $(M', m) \approx (M, m)$, thus $(M', m) \in \mathfrak{M}$. Let $g : \text{dom } M' \rightarrow \text{dom } \tilde{T}^*$ be defined by

$$\begin{aligned} g(x_i) &= i && \text{for all } i = 1, \dots, k; \\ g(u) &= g_j(u) && \text{for each } u \in \text{dom } N'_j \text{ s.t. } u \notin \bar{x}. \end{aligned}$$

Clearly $g : M' \xrightarrow{h} \tilde{T}^*$ is a homomorphism with $g(m) = \iota(0)$. Hence $\iota(0) \in S_{\mathfrak{M}}^{\tilde{T}}$ and also $\iota(0) \in S_{\mathfrak{M}}^{\tilde{A}}$. By the definition of a tuple trace, $f(m) = f(x_{\iota(0)}) \in S_{\mathfrak{M}}^{\tilde{A}}$, a contradiction.

Thus we have shown that \tilde{A} satisfies (4.1). Next we show that \tilde{A} is \mathcal{F} -free. Suppose there is some $F \in \mathcal{F}$ and a homomorphism $f : F \xrightarrow{h} \tilde{A}^*$. If F has only one element, then the one-element substructure of \tilde{A} induced by $f[F]$ is not \mathcal{F} -free, hence not in \mathcal{C} , a contradiction. If F has more than one element but it is irreducible (that is, if it contains exactly one tuple \bar{x} of a relation R^F of arity more than one), then the tuple trace of $f(\bar{x}) \in R^{\tilde{A}}$ is not in \mathcal{C} , again a contradiction. Hence F has a cut m . Also, for any piece $(M, m) \in \mathfrak{M}$ of F the restriction $g = f \upharpoonright M$ is a homomorphism $g : M \xrightarrow{h} \tilde{A}^*$ such that $g(m) = f(m)$. Thus $f(m) \in S_{\mathfrak{M}}^{\tilde{A}}$ for any such piece $(M, m) \in \mathfrak{M}$. Let $E = E(F, m)$ be obtained from F with the cut m as in Lemma 4.7. Then the one-element substructure of \tilde{A} induced by $\{f(m)\}$ admits a homomorphism from E , so by Lemma 4.7 it is not in \mathcal{C} , once again a contradiction. We conclude that $\tilde{A}^* \in \text{Forb}_{\text{h}}(\mathcal{F})$.

Therefore $\tilde{A} \in \tilde{\mathcal{C}}$, and so $A \in \mathcal{C}$. □

Next is the amalgamation property of \mathcal{C} . The following theorem is proved in [16] for the case of τ being finite, but it does not require \mathcal{F} to contain only trees (the arities of τ -relations will then in general be greater than one, and the amalgamation will not be free).

Theorem 4.10. *Let σ be a finite relational signature, let \mathcal{F} be a set of finite σ -trees and let \mathcal{C} be the expanded class for $\text{Forb}_{\text{h}}(\mathcal{F})$. Then*

- (1) *the class of all σ -reducts of the structures in \mathcal{C} is $\text{Forb}_{\text{h}}(\mathcal{F})$;*

- (2) \mathcal{C} is closed under isomorphism;
- (3) \mathcal{C} is closed under taking substructures;
- (4) \mathcal{C} has only countably many isomorphism classes;
- (5) \mathcal{C} has the free amalgamation property.

Proof. (1), (2), (3) follow immediately from Definition 4.4 and from the fact that $\text{Forb}_h(\mathcal{F})$ is closed under both isomorphism and taking substructures. The class $\text{Forb}_h(\mathcal{F})$ has only countably many isomorphism classes because it is a class of finite σ -structures over a finite relational signature σ . The canonical class $\tilde{\mathcal{C}}$ contains exactly one structure for each structure in $\text{Forb}_h(\mathcal{F})$ (the one given by (4.1)), thus $\tilde{\mathcal{C}}$ also has only countably many isomorphism classes. Finally, \mathcal{C} contains finitely many structures for each structure in $\tilde{\mathcal{C}}$, hence \mathcal{C} has only countably many isomorphism classes.

To prove (5), let $A, B_1, B_2 \in \mathcal{C}$ and let $f_1 : A \xrightarrow{e} B$, $f_2 : A_2 \xrightarrow{e} B$ be embeddings. Without loss of generality we may assume that f_1, f_2 are inclusion mappings, that is, $\text{dom } A \subseteq \text{dom } B_1$ and $\text{dom } A \subseteq \text{dom } B_2$, and that $\text{dom } B_1 \cap \text{dom } B_2 = \text{dom } A$. Let C be the $(\sigma \cup \tau)$ -structure defined as follows:

$$\begin{aligned} \text{dom } C &= \text{dom } B_1 \cup \text{dom } B_2, \\ R^C &= R^{B_1} \cup R^{B_2} \quad \text{for any } R \in \sigma \cup \tau. \end{aligned}$$

Now, every one-element substructure of C is a substructure of either B_1 or B_2 (or both), thus by Lemma 4.9 C satisfies (4.2). Moreover, whenever $\bar{x} \in R^C$ for some $R \in \sigma$, then $\bar{x} \in R^{B_1}$ or $\bar{x} \in R^{B_2}$, hence by the same lemma the tuple trace of $\bar{x} \in R^C$ belongs to \mathcal{C} . Using the converse implication of Lemma 4.9 we get that $C \in \mathcal{C}$. It is easy to see that the inclusion mappings are embeddings of B_1 and B_2 to C . \square

After all this preparation, we are ready for the main result, presented in the next section.

5. MAIN RESULT

Orderings. Recall that an *ordered v -structure* is a $(v \cup \{\preceq\})$ -structure A such that the relation \preceq^A is a linear ordering.

Definition 5.1. Let σ be a finite relational signature and let \mathcal{F} be a set of finite σ -trees. The *ordered expanded class* for $\text{Forb}_h(\mathcal{F})$ is the class $\vec{\mathcal{C}}$ of ordered $(\sigma \cup \tau)$ -structures such that $A \in \vec{\mathcal{C}}$ if and only if \preceq^A is a linear ordering and the $(\sigma \cup \tau)$ -reduct of A is in the expanded class \mathcal{C} for $\text{Forb}_h(\mathcal{F})$.

Note. As a consequence of Theorem 4.10, the class $\vec{\mathcal{C}}$ is closed under isomorphism and taking substructures. It also has the amalgamation property: Take the amalgam of the $(\sigma \cup \tau)$ -reducts; the union of the orders \preceq^{B_1} and \preceq^{B_2} is a reflexive anti-symmetric relation on $\text{dom } C$ whose transitive closure is a partial ordering, and any of its linear extensions can be taken as \preceq^C .

Theorem 5.2. *Let σ be a finite relational signature and let \mathcal{F} be a set of finite σ -trees. Then the ordered expanded class for $\text{Forb}_h(\mathcal{F})$ has the Ramsey property.*

Remark 5.3. It has recently been announced by Nešetřil [21] that the ordered expanded class is a Ramsey class if \mathcal{F} is a *finite* set of finite connected σ -structures. Our Theorem 5.2 allows *infinite* \mathcal{F} , but requires that all elements of \mathcal{F} are σ -trees.

Idea of proof. The proof of this theorem, which spreads over the following two sections, is based on the ideas of *partite lemma* and *partite construction*, developed by Nešetřil and Rödl [23, 25, 26, 27].

The principal idea is the notion of a *partite structure*. A partite structure is an ordered $(\sigma \cup \tau)$ -structure whose domain is split into several parts. The parts of a partite structure X are indexed by elements of some ordered σ -structure P ; formally, the part of an element is determined by a mapping $\iota_X : \text{dom } X \rightarrow \text{dom } P$. Furthermore, we want ι_X to be a homomorphism of the $(\sigma \cup \{\preceq\})$ -reduct of X to P . Informally, we want the tuples of the σ -relations only to sit across those parts of X where P also has a corresponding tuple for the same relation symbol. The ordering \preceq^A preserves the ordering of parts given by \preceq^P ; within one part, the ordering can be arbitrary.

Another desired property of partite structures is somewhat peculiar: We do not want a tuple \bar{x} of a σ -relation of X to contain two different elements from the same part. This applies if a tuple of P contains an element of P more than once. For instance, let $R \in \sigma$ and $(a, b, b) \in R^P$. Then we allow a tuple $(x, y, y) \in R^X$ if $\iota_X(x) = a$ and $\iota_X(y) = b$, but we do not allow $(x, y, y') \in R^X$ for $y \neq y'$. Formally, we require ι_X to be injective on the set of elements of any tuple of a relation of X .

Finally, we want all the elements of any part to belong to exactly the same unary σ -relations as the corresponding element of P (the τ -relations are not prescribed by P , which is a σ -structure, and can vary within one part of a partite structure). In addition, if there is a “loop” $(a, a, \dots, a) \in R^P$ for some $R \in \sigma$, we want that $(x, x, \dots, x) \in R^X$ for each x with $\iota_X(x) = a$.

Partite lemma. The partite lemma is often proved by an application of the Hales–Jewett theorem (as in [26, 27, 28]). Our proof, however, is inspired by that of Prömel and Voigt [29]. We prove the Ramsey property for a class of very special structures: *rectified structures*. A rectified structure is similar to a partite structure in that its domain is split into parts, indexed this time by the elements of some $A \in \vec{\mathcal{C}}$. Furthermore, if we choose one element from each part, the induced substructure of X is isomorphic to A . Conversely, any tuple of a relation of X lies within some such copy of A in X . Thus a rectified structure X is actually fully determined (up to isomorphism) by A and the size of each part $\iota^{-1}[a]$ of X , $a \in \text{dom } A$ (see Figure 7). In this case, the proof of the Ramsey property is relatively straightforward by induction on the number of elements of A (i.e., the number of parts), and the base step as well as each induction step follow from the pigeon-hole principle.

Partite construction. The partite lemma is then applied repeatedly in the partite construction (also known as the *amalgamation method*). The idea is as follows: Given $A, B \in \vec{\mathcal{C}}$, we want to find $C \in \vec{\mathcal{C}}$ such that $C \rightarrow (B)_r^A$. By Theorem 2.1 we know that there exists an ordered $(\sigma \cup \tau)$ -structure C with $C \rightarrow (B)_r^A$, but there is no guarantee that $C \in \vec{\mathcal{C}}$. Indeed, typically such C will be far from being \mathcal{F} -free and satisfying (4.1). However, it will be a good starting point. In fact, we will use a σ -structure P such that $P \rightarrow (B^*)_r^{A^*}$, where A^* , B^* are the (ordered) σ -reducts of A , B . This P serves as the indexing structure for the parts of our partite structures.

The partite construction then works inductively, starting from C_0 , which is basically a suitable sum (disjoint union) of partite structures isomorphic to B . In each induction step (one for every occurrence of a copy of A^* in P), it uses the partite lemma for rectified

structures. To this end, each application of the partite lemma must be preceded by adding new tuples to the relations in order to make the structure rectified. This construction is called *rectification*; we can show that this can be done without losing membership in $\vec{\mathcal{C}}$ (Lemma 7.1).

The inductive construction produces a structure C with the following property: Whenever $\binom{C}{A}$ is r -coloured, there exists a copy C'_0 of C_0 within C such that the colour of each copy A' of A within C'_0 depends solely on the parts this copy sits on (i.e., on the image $\iota_C[A']$). Finally, the way we construct C_0 and Theorem 2.1 then guarantee the existence of a monochromatic copy of B within this C'_0 .

6. PARTITE LEMMA

Throughout this section, \mathcal{F} is a fixed set of finite σ -trees and $\vec{\mathcal{C}}$ is the ordered expanded class for $\text{Forb}_h(\mathcal{F})$.

Rectified structures. Let $A \in \vec{\mathcal{C}}$. An A -rectified structure is a pair (X, ι_X) such that X is an ordered $(\sigma \cup \tau)$ -structure, $\iota_X : \text{dom } X \rightarrow \text{dom } A$ is a mapping, $x \preceq^X x'$ implies that $\iota_X(x) \preceq^A \iota_X(x')$, and for any $R \in \sigma \cup \tau$ and any $\bar{x} \in (\text{dom } X)^{\text{ar}(R)}$ we have

$$\bar{x} \in R^X \iff \iota_X \text{ is injective on } \bar{x} \text{ and } \iota_X(\bar{x}) \in R^A. \quad (6.1)$$

(The mapping ι_X is injective on $\bar{x} = (x_1, x_2, \dots, x_k)$ if $x_i = x_j$ whenever $\iota_X(x_i) = \iota_X(x_j)$.) Observe that X is uniquely determined by A , $\text{dom } X$ and ι_X via (6.1).

A mapping $e : \text{dom } X \rightarrow \text{dom } Y$ is an embedding of A -rectified structure (X, ι_X) into (Y, ι_Y) if $e : X \xrightarrow{e} Y$ is an embedding of $(\sigma \cup \tau \cup \{\preceq\})$ -structures and $\iota_X = \iota_Y \circ e$.

Proposition 6.1. *Let $A \in \vec{\mathcal{C}}$.*

- (1) *If (X, ι_X) is A -rectified, then $X \in \vec{\mathcal{C}}$.*
- (2) *If (X, ι_X) is A -rectified, then the mapping ι_X is a homomorphism of X to A .*
- (3) *(A, id_A) is A -rectified.*
- (4) *For any A -rectified (X, ι_X) , any mapping $e : \text{dom } A \rightarrow \text{dom } X$ such that $\iota_X \circ e = \text{id}_A$ is an embedding of A into X , as well as an embedding of (A, id_A) into (X, ι_X) .*
- (5) *If $|\text{dom } A| = 1$ and (X, ι_X) is an A -rectified structure, then X is the disjoint union of several (possibly 0 or 1) copies of A , endowed with some linear ordering of its elements, and ι_X is constant. Conversely, any such (X, ι_X) is A -rectified.*

Proof.

(1) Let \hat{A}, \hat{X} be the $(\sigma \cup \tau)$ -reducts of A, X , respectively. We apply Lemma 4.9. Let $x \in \text{dom } \hat{X}$. By (6.1), the one-element substructure \hat{X}_x of \hat{X} induced by x is isomorphic to the substructure of \hat{A} induced by $\iota_X(x)$. Thus $\hat{X}_x \in \mathcal{C}$ because $\hat{A} \in \mathcal{C}$. Next, for any $R \in \sigma$ and $\bar{x} \in R^{\hat{X}}$, the tuple trace $T = T(\hat{X}, \bar{x}, R)$ is equal to the tuple trace $T(\hat{A}, \iota_X(\bar{x}), R)$ of $\iota_X(\bar{x}) \in R^{\hat{A}}$ because of (6.1). Hence $T \in \mathcal{C}$. By Lemma 4.9, $\hat{X} \in \mathcal{C}$. Therefore $X \in \vec{\mathcal{C}}$.

(2) All σ - and τ -relations are preserved by (6.1), and \preceq is preserved by definition.

(3) Obvious.

(4) Since $\iota_X \circ e$ is injective, e is injective and ι_X is injective on any tuple of $e[A]$. Let $R \in \sigma \cup \tau$. If $\bar{a} \in R^A$, then $\iota_X(e(\bar{a})) = \bar{a}$, so by (6.1), $e(\bar{a}) \in R^X$. Conversely, if $\bar{a} \in (\text{dom } A)^{\text{ar}(R)}$ and $e(\bar{a}) \in R^X$, then $\bar{a} = \iota_X(e(\bar{a})) \in R^A$ because ι_X is a homomorphism. An analogous argument applies to preservation of \preceq .

(5) If $\iota_X : \text{dom } X \rightarrow \text{dom } A$ and $|\text{dom } A| = 1$, then certainly ι_X is constant. But then if $\bar{x} \in R^X$ for some $R \in \sigma \cup \tau$, then \bar{x} is of the form (x, x, \dots, x) because ι_X is injective on \bar{x} . Hence the only relation of X that spans distinct elements of X is the linear ordering \preceq^X . It follows from (6.1) that each one-element substructure of X is isomorphic to A .

The converse is obvious. \square

Lemma 6.2 (Partite Lemma). *Let \mathcal{F} be a set of finite σ -trees and let $\vec{\mathcal{C}}$ be the ordered expanded class for $\text{Forb}_h(\mathcal{F})$; let $A \in \vec{\mathcal{C}}$. Let (B, ι_B) be A -rectified; let $r \geq 1$. Then there exists A -rectified (E, ι_E) such that $(E, \iota_E) \rightarrow (B, \iota_B)_r^{(A, \text{id}_A)}$.*

Proof. By induction on $|\text{dom } A|$. If $|\text{dom } A| = 1$, take E to be the disjoint union of $r \cdot (|\text{dom } B| - 1) + 1$ copies of A with an arbitrary linear ordering \preceq^E ; ι_E is constant. This is an A -partite structure by Proposition 6.1(5). Since – by the same Proposition – B is an ordering of the disjoint union of $|\text{dom } B|$ copies of A , any substructure of E on $|\text{dom } B|$ elements is isomorphic to B . In any r -colouring of $\text{dom } E$ there exist $|\text{dom } B|$ elements of the same colour, inducing a monochromatic copy of B .

If $|\text{dom } A| \geq 2$, assume that $\text{dom } A = \{0, 1, \dots, n\}$. Let A' be the substructure of A induced by the subset $\{1, \dots, n\}$; let B' be the substructure of B induced by $\iota_B^{-1}[\{1, \dots, n\}]$, and $\iota_{B'} = \iota_B \upharpoonright \text{dom } B'$. Then $(B', \iota_{B'})$ is A' -rectified. Apply induction to get A' -rectified $(E', \iota_{E'})$ such that $(E', \iota_{E'}) \rightarrow (B', \iota_{B'})_r^{(A', \text{id}_{A'})}$, where $k = r \cdot (|\iota_B^{-1}(0)| - 1) + 1$. Assuming that $\text{dom } E' \cap \{1, 2, \dots, k\} = \emptyset$ let $\text{dom } E = \text{dom } E' \cup \{1, 2, \dots, k\}$ and define $\iota_E(x) = 0$ if $x \in \{1, 2, \dots, k\}$ and $\iota_E(x) = \iota_{E'}(x)$ otherwise. Let all $(\sigma \cup \tau)$ -relations of E be defined by (6.1); let \preceq^E be an extension of $\preceq^{E'}$ that is preserved by ι_E . Thus E' is the substructure of E on $\iota_E^{-1}[\{1, \dots, n\}]$. See Figure 7. Clearly (E, ι_E) is A -rectified.




Figure 7: Lemma 6.2.

To prove that $(E, \iota_E) \rightarrow (B, \iota_B)_r^{(A, \text{id}_A)}$, consider any r -colouring χ of $\left(\begin{smallmatrix} (E, \iota_E) \\ (A, \text{id}_A) \end{smallmatrix}\right)$. Define $\chi' : \left(\begin{smallmatrix} (E', \iota_{E'}) \\ (A', \text{id}_{A'}) \end{smallmatrix}\right) \rightarrow \{1, \dots, r\}^{\iota_E^{-1}(0)}$ by $\chi'(e') = (c \mapsto \chi(e' \cup (0 \mapsto c)))$. That is, the χ' -colour of a copy of A' in E' is a vector of χ -colours, one for each of the k extensions of the copy of A' by an element in the 0th part to a copy of A in E . By the definition of $(E', \iota_{E'})$, there is a monochromatic $g' \in \left(\begin{smallmatrix} (E', \iota_{E'}) \\ (B', \iota_{B'}) \end{smallmatrix}\right)$. Hence for any fixed $c \in \iota_E^{-1}(0) = \{1, 2, \dots, k\}$, the mapping $\varphi_c : h' \mapsto \chi((g' \circ h') \cup (0 \mapsto c))$ is constant on $\left(\begin{smallmatrix} (B', \iota_{B'}) \\ (A', \text{id}_{A'}) \end{smallmatrix}\right)$. Define $\psi : \iota_E^{-1}(0) \rightarrow \{1, \dots, r\}$

by setting $\psi(c)$ to be the constant value of φ_c . Since $|\iota_E^{-1}(0)| = k > r (|\iota_B^{-1}(0)| - 1)$, there exists a subset $M \subseteq \iota_E^{-1}(0)$ with $|M| = |\iota_B^{-1}(0)|$ such that ψ is constant on M . Define $g \in \binom{(E, \iota_E)}{(B, \iota_B)}$ to be an extension of g' by the \preceq -preserving bijection of $\iota_B^{-1}(0)$ and M . Then g is monochromatic. \square

7. PARTITE CONSTRUCTION

This section is devoted to finishing the proof of Theorem 5.2. Again, \mathcal{F} is a fixed set of finite σ -trees and $\vec{\mathcal{C}}$ is the ordered expanded class for $\text{Forb}_h(\mathcal{F})$.

Partite structures. Let P be an ordered σ -structure. A P -partite $\vec{\mathcal{C}}$ -structure is a pair (A, ι_A) where $A \in \vec{\mathcal{C}}$ and $\iota_A : \text{dom } A \rightarrow \text{dom } P$ is a homomorphism of the $(\sigma \cup \{\preceq\})$ -reduct A^* of A to P that is injective on any tuple of the relation R^A for any $R \in \sigma$, and such that the restriction of ι_A to any one-element substructure of A^* is an embedding of this one-element $(\sigma \cup \{\preceq\})$ -structure into P . For an element a of A or a tuple \bar{a} , the image $\iota_A(a)$ or $\iota_A(\bar{a})$ is called the *trace* of a or \bar{a} . A P -partite $\vec{\mathcal{C}}$ -structure (A, ι_A) is *transversal* if ι_A is an embedding of A^* to P .

A mapping $e : \text{dom } A \rightarrow \text{dom } B$ is an embedding of a P -partite $\vec{\mathcal{C}}$ -structure (A, ι_A) into (B, ι_B) if $e : A \xrightarrow{e} B$ is an embedding of $(\sigma \cup \tau \cup \{\preceq\})$ -structures and $\iota_A = \iota_B \circ e$.

Lemma 7.1 (“rectification”). *Let $\vec{\mathcal{C}}$ be the ordered expanded class for $\text{Forb}_h(\mathcal{F})$, where \mathcal{F} is a set of finite σ -trees. Let (C, ι_C) be a P -partite $\vec{\mathcal{C}}$ -structure for some σ -structure P . If (D, ι_D) is defined by setting*

$$\text{dom } D = \text{dom } C,$$

$$\iota_D = \iota_C,$$

$$S^D = S^C \text{ for } S \in \tau,$$

$$\preceq^D = \preceq^C,$$

for $R \in \sigma$:

$$\bar{x} \in R^D \iff \iota_D \text{ is injective on } \bar{x}, \text{ and}$$

$$\exists \bar{y} \in R^C : \iota_C(\bar{y}) = \iota_D(\bar{x}) \text{ and } \forall i, \{S \in \tau : x_i \in S^D\} = \{S \in \tau : y_i \in S^C\}, \quad (7.1)$$

then (D, ι_D) is a P -partite $\vec{\mathcal{C}}$ -structure. Moreover, $R^C \subseteq R^D$ for any $R \in \sigma$; if R is unary, then $R^C = R^D$.

Proof. It is straightforward that ι_D is a homomorphism of the $(\sigma \cup \{\preceq\})$ -reduct D^* to P because ι_C is a homomorphism of C^* to P . By definition, ι_D is injective on any tuple of any σ -relation of D , and every one-element substructure of D is isomorphic to the corresponding one-element substructure of C .

To show that $D \in \vec{\mathcal{C}}$, first apply the “only if” direction of Lemma 4.9 to prove that the tuple trace of any $\bar{y} \in R^C$ is in $\vec{\mathcal{C}}$ because $C \in \vec{\mathcal{C}}$. Then observe that the tuple trace of any $\bar{x} \in R^D$ is equal to the tuple trace of some $\bar{y} \in R^C$. Also, every one-element substructure of D is equal to the corresponding one-element substructure of C . Finally apply the “if” direction of Lemma 4.9.

If $\bar{x} \in R^C$ then $\bar{y} = \bar{x}$ can be taken to show that $\bar{x} \in R^D$. Thus $R^C \subseteq R^D$. If R is unary and $x \in R^D$, then there is $y \in R^C$ with $\iota_C(y) = \iota_C(x)$; hence $x \in R^C$ because ι_C restricted

to any one-element structure is an embedding (by the definition of a P -partite structure). Therefore $R^C = R^D$ for all unary $R \in \sigma$. \square

Observe that the P -partite \vec{C} -structure (D, ι_D) from Lemma 7.1 is “*rectified*” in the following sense:

For any $R \in \sigma$ and any $\bar{y} \in R^D$, if \bar{x} is a tuple such that $\iota_D(\bar{x}) = \iota_D(\bar{y})$,
 ι_D is injective on \bar{x} , and $\{S \in \tau: y_i \in S^D\} = \{S \in \tau: x_i \in S^D\}$ for all i , then $\bar{x} \in R^D$.
(7.2)

Lemma 7.1 asserts that any P -partite \vec{C} -structure (C, ι_C) can be transformed into (D, ι_D) that satisfies (7.2) by adding tuples to (non-unary) σ -relations. Note that if (C, ι_C) already satisfies (7.2) and (D, ι_D) is defined by (7.1), then no tuples will be added and $(D, \iota_D) = (C, \iota_C)$. In particular, this is the case if (C, ι_C) is transversal.

Rectified substructures. The next lemma will apply in the proof of Theorem 5.2 in the following situation: Start with (D, ι_D) which is rectified in the above sense, that is, it satisfies (7.2), and an ordered structure $A \in \vec{C}$. Split the elements of A into parts so that A would be a transversal P -partite \vec{C} -structure ($\iota_A: \text{dom } A \rightarrow \text{dom } P$ is an embedding of A^* into P). Select those elements of D

- (i) whose trace lies in the trace of A , and
- (ii) whose unary τ -relations are exactly the same as the unary τ -relations of the corresponding element of A .

The selected elements induce a substructure B of D . There is a natural way to define $\iota_B: \text{dom } B \rightarrow \text{dom } A$ so that an element of B would be mapped to the element of A in the same P -part (that is, $\iota_A(\iota_B(b)) = \iota_D(b)$). Lemma 7.2 claims that such (B, ι_B) is A -rectified.

Lemma 7.2. *Let (D, ι_D) be a P -partite \vec{C} -structure satisfying (7.2), and let (A, ι_A) be a transversal P -partite \vec{C} -structure. Suppose there is a P -partite embedding of (A, ι_A) into (D, ι_D) . Define*

$$\text{dom } B = \{x \in \text{dom } D: \iota_D(x) \in \iota_A[\text{dom } A]\}$$

$$\text{and } \{S \in \tau: x \in S^D\} = \{S \in \tau: \iota_A^{-1}(\iota_D(x)) \in S^A\} \quad (7.3)$$

and let B be the substructure of D induced by $\text{dom } B$. Set $\iota_B = \iota_A^{-1} \circ (\iota_D \upharpoonright \text{dom } B)$. Then (B, ι_B) is A -rectified.

Proof. If $x \preceq^B x'$, then $x \preceq^D x'$ because B is a substructure of D ; thus $\iota_D(x) \preceq^P \iota_D(x')$ because D is P -partite; hence $\iota_B(x) = \iota_A^{-1}(\iota_D(x)) \preceq^A \iota_A^{-1}(\iota_D(x')) = \iota_B(x')$ because ι_A is a $(\sigma \cup \{\preceq\})$ -embedding. If $S \in \tau$ and $x \in \text{dom } B$, then $x \in S^B$ iff $\iota_B(x) = \iota_A^{-1}(\iota_D(x)) \in S^A$ by (7.3). Let $R \in \sigma$ and $\bar{x} \in R^B \subseteq R^D$. Then ι_D is injective on \bar{x} because D is P -partite; hence also ι_B is injective on \bar{x} . Moreover, $\iota_B(\bar{x}) \in R^A$ because ι_D is a homomorphism and ι_A an embedding.

Conversely, suppose that $R \in \sigma$, $\bar{x} \in (\text{dom } B)^{\text{ar}(R)}$, ι_B is injective on \bar{x} and $\iota_B(\bar{x}) \in R^A$. Let $e: A \xrightarrow{e} D$ be an embedding such that $\iota_D \circ e = \iota_A$; let $\bar{y} = e(\iota_B(\bar{x}))$. Then $\bar{y} \in R^D$ and $\iota_D(\bar{y}) = \iota_A(\iota_B(\bar{x})) = \iota_D(\bar{x})$. Moreover, for any i , $\{S \in \tau: y_i \in S^D\} = \{S \in \tau: \iota_B(x_i) \in S^A\}$ because e is an embedding, and $\{S \in \tau: \iota_B(x_i) \in S^A\} = \{S \in \tau: x_i \in S^D\}$ by (7.3). Therefore $\bar{x} \in R^D$ by (7.2), whence $\bar{x} \in R^B$ because B is a substructure of D . \square

Proof of Theorem 5.2. Let \mathcal{F} be a set of finite σ -trees and let \mathcal{C} be the expanded class and $\vec{\mathcal{C}}$ the ordered expanded class for $\text{Forb}_h(\mathcal{F})$. Consider $A, B \in \vec{\mathcal{C}}$ and a positive integer r . We construct $C \in \vec{\mathcal{C}}$ such that $C \rightarrow (B)_r^A$.

Let A^*, B^* be the $(\sigma \cup \{\preceq\})$ -reducts of A, B , respectively. By Theorem 2.1 there exists an ordered σ -structure P such that $P \rightarrow (B^*)_r^{A^*}$. Define (C_0, ι_{C_0}) by

$$\text{dom } C_0 = \binom{P}{B^*} \times \text{dom } B,$$

for any k -ary $R \in \sigma \cup \tau$:

$$R^{C_0} = \left\{ ((f, x_1), (f, x_2), \dots, (f, x_k)) : f \in \binom{P}{B^*} \text{ and } (x_1, x_2, \dots, x_k) \in R^B \right\},$$

$\iota_{C_0} : \text{dom } C_0 \rightarrow \text{dom } P$ is defined by $\iota_{C_0} : (f, x) \mapsto f(x)$,

\preceq^{C_0} is any linear ordering that is preserved by ι_{C_0} .

Thus C_0 (without the ordering) is isomorphic to a sum of structures, and each of the summands is isomorphic to B . See Figure 8. Observe that (C_0, ι_{C_0}) is a P -partite $\vec{\mathcal{C}}$ -structure because \mathcal{C} is closed under taking sums.

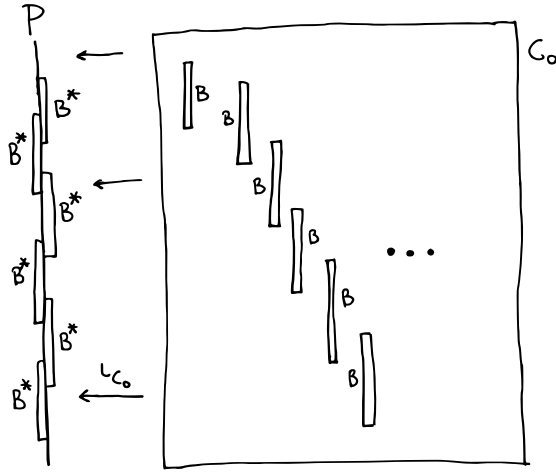


Figure 8: C_0 .

Unless B is connected, C_0 may contain other copies of B , however. Therefore we call the embeddings $c_f : B \xrightarrow{e} C_0$ with $c_f(x) = (f, x)$ for some $f \in \binom{P}{B^*}$ *distinguished*; the corresponding substructures of C_0 are called *distinguished copies* of B . Each of the distinguished copies of B in C_0 forms a transversal structure. Hence if (D_0, ι_{D_0}) is obtained from (C_0, ι_{C_0}) by (7.1), then each of the distinguished embeddings of B to C_0 is also an embedding of B to D_0 . In other words, for any $R \in \sigma$ none of the new R -tuples added by rectification lie within a distinguished copy of B .

Fix some numbering of $\binom{P}{A^*} = \{e_1, \dots, e_N\}$, the set of all embeddings of A^* into P . We will inductively construct P -partite $\vec{\mathcal{C}}$ -structures $(C_1, \iota_{C_1}), \dots, (C_N, \iota_{C_N})$.

Let $k \in \{1, \dots, N\}$ and suppose $(C_{k-1}, \iota_{C_{k-1}})$ has been constructed. If there is no P -partite embedding of (A, e_k) into $(C_{k-1}, \iota_{C_{k-1}})$,¹ let $(C_k, \iota_k) = (C_{k-1}, \iota_{C_{k-1}})$. Otherwise

¹There is a copy $e_k[A^*]$ of A^* in P , but it may not lie within any copy of B^* in P . Even if there is a copy of A^* in C_{k-1}^* with the same trace, however, the τ -relations may not be the right ones, so that these elements

let $(D_{k-1}, \iota_{D_{k-1}})$ be defined from $(C_{k-1}, \iota_{C_{k-1}})$ by the rectification construction (7.1). Let (B_k, ι_{B_k}) be the substructure of $(D_{k-1}, \iota_{D_{k-1}})$ obtained as in (7.3) (Lemma 7.2), using (A, e_k) in place of (A, ι_A) . Then (B_k, ι_{B_k}) is A -rectified and we can apply the Partite Lemma, Lemma 6.2, in order to get A -rectified (E_k, ι_{E_k}) such that $(E_k, \iota_{E_k}) \rightarrow (B_k, \iota_{B_k})_r^{(A, \text{id}_A)}$ (w.r.t. embeddings of A -rectified structures). Therefore $(E_k, e_k \circ \iota_{E_k}) \rightarrow (B_k, e_k \circ \iota_{B_k})_r^{(A, e_k)}$ (w.r.t. embeddings of P -partite structures).

Now we proceed to construct C_k from E_k and several copies of D_{k-1} by amalgamation. The construction described below gives the result explicitly. For each P -partite copy of B_k in E_k , we glue a copy of D_{k-1} onto E_k , overlapping on that copy of B_k . Formally, put

$$\text{dom } C_k = \text{dom } E_k \cup \left(\begin{matrix} (E_k, \iota_{E_k}) \\ (B_k, \iota_{B_k}) \end{matrix} \times (\text{dom } D_{k-1} \setminus \text{dom } B_k) \right).$$

Define $\lambda_k : \begin{matrix} (E_k, \iota_{E_k}) \\ (B_k, \iota_{B_k}) \end{matrix} \times \text{dom } D_{k-1} \rightarrow \text{dom } C_k$ by

$$\lambda_k : (g, x) \mapsto \begin{cases} g(x) & \text{if } x \in \text{dom } B_k, \\ (g, x) & \text{otherwise,} \end{cases}$$

(so $\lambda_k(g, x)$ gives the name of the element of C_k corresponding to the element x in the “ g th” copy of D_{k-1} within C_k). For any ℓ -ary $R \in \sigma \cup \tau$, let

$$R^{C_k} = \left\{ (\lambda_k(g, x_1), \dots, \lambda_k(g, x_\ell)) : g \in \begin{matrix} (E_k, \iota_{E_k}) \\ (B_k, \iota_{B_k}) \end{matrix}, (x_1, \dots, x_\ell) \in R^{D_{k-1}} \right\}.$$

Furthermore define $\iota_{C_k} : \text{dom } C_k \rightarrow \text{dom } P$ by


$$\begin{aligned} \iota_{C_k} : y &\mapsto e_k(\iota_{E_k}(y)) & \text{if } y \in \text{dom } E_k, \\ \iota_{C_k} : (g, x) &\mapsto \iota_{D_{k-1}}(x) & \text{otherwise.} \end{aligned}$$

Note that $\iota_{C_k}(\lambda_k(g, x)) = \iota_{D_{k-1}}(x)$ for any $x \in \text{dom } D_{k-1}$ and $g \in \begin{matrix} (E_k, \iota_{E_k}) \\ (B_k, \iota_{B_k}) \end{matrix}$. Finally, let \preceq^{C_k} be a linear ordering such that $y \preceq^{C_k} y'$ if $y \preceq^{E_k} y'$, $\lambda_k(g, x) \preceq^{C_k} \lambda_k(g, x')$ if $x \preceq^{D_{k-1}} x'$, and $z \preceq^{C_k} z'$ if $\iota_{C_k}(z) \preceq^P \iota_{C_k}(z')$. See Figure 9.

Notice that for a fixed g , the mapping $\lambda_k(g, -) : x \mapsto \lambda_k(g, x)$ is an embedding of $(D_{k-1}, \iota_{D_{k-1}})$ to (C_k, ι_{C_k}) . By definition of D_{k-1} , $\lambda_k(g, -)$ is an injective homomorphism of $(C_{k-1}, \iota_{C_{k-1}})$ to (C_k, ι_{C_k}) . The inclusion mapping is an embedding of E_k to C_k because (E_k, ι_{E_k}) is A -rectified.

Now we claim that (C_k, ι_{C_k}) is a P -partite $\vec{\mathcal{C}}$ -structure. Every one-element substructure of C_k is isomorphic to a one-element substructure of D_{k-1} , and every tuple of some relation R^{C_k} , $R \in \sigma$, corresponds to some tuple of $R^{D_{k-1}}$ with the same tuple trace. Since $D_{k-1} \in \vec{\mathcal{C}}$, by Lemma 4.9 we have $C_k \in \vec{\mathcal{C}}$. To show that ι_{C_k} is a homomorphism, let $\bar{y} \in R^{C_k}$ for some $R \in \sigma$. Then $\bar{y} = \lambda_k(g, \bar{x})$ for some $g \in \begin{matrix} (E_k, \iota_{E_k}) \\ (B_k, \iota_{B_k}) \end{matrix}$ and $\bar{x} \in R^{D_{k-1}}$, and $\iota_{C_k}(\bar{y}) = \iota_{D_{k-1}}(\bar{x})$. Thus $\iota_{C_k}(\bar{y}) \in R^P$ because $\iota_{D_{k-1}}$ is a homomorphism. By definition, ι_{C_k} also preserves \preceq^{C_k} . Hence $\iota_{C_k} : C_k^* \xrightarrow{h} P$ is a homomorphism of $(\sigma \cup \{\preceq\})$ -structures.

induce a substructure of C_{k-1} that is not isomorphic to A . In these cases, it can happen that no copy of A in C_{k-1} has trace $e_k[A^*]$.

Figure 9: C_k .

Next, to show that ι_k is injective on any $\bar{y} \in R^{C_k}$ (that is, \bar{y} does not contain two distinct elements in the same part), again use the fact that $\bar{y} = \lambda_k(g, \bar{x})$ for some $g \in \binom{(E_k, \iota_{E_k})}{(B_k, \iota_{B_k})}$ and $\bar{x} \in R^{D_{k-1}}$ and that $\lambda_k(g, -)$ preserves the parts. Finally, let $y \in \text{dom } C_k$. Then $y = \lambda_k(g, x)$ for some $g \in \binom{(E_k, \iota_{E_k})}{(B_k, \iota_{B_k})}$ and $x \in \text{dom } D_{k-1}$. The one-element substructure of C_k^* induced by y is isomorphic to the one-element substructure of D_{k-1}^* induced by x , which in turn is isomorphic to the one-element substructure of P induced by $\iota_{D_{k-1}}(x)$ because D_{k-1} is P -partite. As $\iota_{C_k}(y) = \iota_{D_{k-1}}(x)$, we have that the one-element substructure of C_k^* induced by y is isomorphic to the one-element substructure of P induced by $\iota_{C_k}(y)$. Therefore (C_k, ι_{C_k}) is a P -partite \vec{C} -structure.

Let $C = C_N$. We show that $C \rightarrow (B)_r^A$. Consider any colouring $\chi : (C) \rightarrow \{1, \dots, r\}$. By downward induction we find injective homomorphisms $h_k : (C_{k-1}, \iota_{C_{k-1}}) \xrightarrow{h_k} (C_k, \iota_{C_k})$ for $k = N, N-1, \dots, 1$ that have certain monochromatic properties. Each $h_k = \lambda(g_k, -)$ for some $g_k \in \binom{(E_k, \iota_{E_k})}{(B_k, \iota_{B_k})}$.

Suppose h_i is known for $i = N, \dots, k+1$ (possibly for no i yet). If $(C_k, \iota_{C_k}) = (C_{k-1}, \iota_{C_{k-1}})$, let h_k be the identity mapping.² Otherwise define the colouring $\chi_k : \binom{(E_k, \iota_{E_k})}{(A, \text{id}_A)} \rightarrow \{1, \dots, r\}$ by setting $\chi_k(q) = \chi(h_N \circ h_{N-1} \circ \dots \circ h_{k+1} \circ q)$. (Observe that the composed mapping is indeed an embedding.) Since $(E_k, \iota_{E_k}) \rightarrow (B_k, \iota_{B_k})_r^{(A, \text{id}_A)}$, there exists a χ_k -monochromatic embedding $g_k : (B_k, \iota_{B_k}) \xrightarrow{g_k} (E_k, \iota_{E_k})$. Let $h_k = \lambda(g_k, -)$.

Let $h = h_N \circ h_{N-1} \circ \dots \circ h_1 : (C_0, \iota_{C_0}) \xrightarrow{h} (C_N, \iota_{C_N})$. By definition, h is in fact the following composed mapping:

$$C_0 \xrightarrow{\text{id}} D_0 \xrightarrow{h_1} C_1 \xrightarrow{\text{id}} D_1 \xrightarrow{h_2} C_2 \xrightarrow{\text{id}} D_2 \xrightarrow{h_3} \dots \xrightarrow{\text{id}} D_{N-1} \xrightarrow{h_N} C_N,$$

where each of the identity mappings is a bijective homomorphism obtained implicitly from Lemma 7.1 and each $h_k = \lambda_k(g_k, -)$ is an embedding. In general, h is not an embedding because new tuples of σ -relations are added during rectification. We want to show, however,

²This is the case if there was no P -partite embedding of (A, e_k) into $(C_{k-1}, \iota_{C_{k-1}})$; see previous footnote.

that no new tuples are added to the distinguished copies of B in C_0 . In other words, for any distinguished embedding $c_f : B \xrightarrow{e} C_0$, the mapping $h \circ c_f$ is an embedding: By definition, $h \circ c_f$ is injective. For $R \in \sigma$, if $\bar{x} \in R^B$, then $h(c_f(\bar{x})) \in R^C$ because h is a homomorphism and c_f an embedding. If $h(c_f(\bar{x})) \in R^C$, then $f(\bar{x}) = \iota_{C_0}(c_f(\bar{x})) = \iota_{C_N}(h(c_f(\bar{x}))) \in R^P$ because ι_{C_N} is a homomorphism; hence $\bar{x} \in R^B$ because f is an embedding. For $S \in \tau$, we have $x \in S^B$ iff $h(c_f(x)) \in S^C$ because of (7.1) and because c_f and each h_k is an embedding.

Consider any $e_j \in \binom{P}{A^*}$. Any embedding d_j of A to C_0 such that $\iota_{C_0} \circ d_j = e_j$ is also a P -partite embedding of (A, e_j) to (C_0, ι_{C_0}) . Moreover, $h \circ d_j$ is a P -partite embedding of (A, e_j) to (C_N, ι_{C_N}) . By definition of h_j , all such embeddings take the same colour under χ . Thus we define $\chi_0 : \binom{P}{A^*} \rightarrow \{1, \dots, r\}$ by $\chi_0(e_j) = \chi(h \circ d_j)$ if there exists $d_j \in \binom{C_0}{A}$ such that $\iota_{C_0} \circ d_j = e_j$, and arbitrarily otherwise. By definition of P there exists χ_0 -monochromatic $f \in \binom{P}{B^*}$. Let $c_f : B \xrightarrow{e} C_0$ be the distinguished embedding given by $c_f : x \mapsto (f, x)$.

Conclude the proof by observing that $h \circ c_f$ is a χ -monochromatic embedding of B to C . \square

8. A NOTE ON DATALOG AND CONSTRAINT SATISFACTION

This section contains a brief description of constraint satisfaction problems (CSPs), Datalog programs and their connection to Ramsey theory, which was the original motivation for this research. A more thorough introduction to CSPs and their complexity is given, e.g., in [6, 7]. More details on Datalog can be found in [18]; a concise exposition, relevant to our setting, is given in [4].

Tree Datalog. Let σ, τ be disjoint finite relational signatures such that τ contains only unary relation symbols and a special nullary relation symbol **goal**. A *tree Datalog program* is a finite set of rules of the form

$$S(x) \leftarrow t_1, \dots, t_n$$

or

$$\mathbf{goal} \leftarrow t_1, \dots, t_n,$$

where $S \in \tau$ and each t_i is an atomic formula $R_i(x_{i_1}, \dots, x_{i_k})$ with $R_i \in \sigma \cup \tau$, so that at most one of the R_i 's belongs to σ . The part of the rule to the left of the arrow is called the *head* of the rule; the part to the right is called the *body* of the rule. In the context of a Datalog program, the predicates appearing in the head of a rule are called *IDBs* (*intensional database predicates*), whereas the predicates from σ are called *EDBs* (*extensional database predicates*).

A Datalog program can be “executed” on a σ -structure A^* to recursively construct the τ -relations S^A , $S \in \tau$, and consequently a $(\sigma \cup \tau)$ -expansion A of A^* , by repeatedly adding elements of A^* to the unary relations S^A following the program’s rules. The execution terminates when the application of any rule does not result in adding an element into an IDB. The **goal** predicate is initially set to **false**, and we say that the Datalog program *accepts* A^* if its **goal** predicate evaluates to **true** on A^* .

A Datalog program can be used to provide a finite description of an infinite regular set of forbidden trees, as the following example illustrates.

Example 8.1. We revisit the example of \mathcal{F} consisting of thunderbolts from Section 4 (Figure 1). The context is digraphs, so σ contains one binary relation symbol A . The signature τ obtained from Definition 4.4 contains four relation symbols: S_i for each \mathfrak{M}_i , $i = 1, 2, 3, 4$, which will be IDBs of the corresponding Datalog program. In addition to these unary IDBs, the Datalog will have a nullary IDB **goal**. The rules of the program are:

$$\begin{aligned} S_1(a) &\leftarrow A(b, a); \\ S_1(a) &\leftarrow A(a, b), S_2(b); \\ S_2(a) &\leftarrow A(b, a), S_1(b); \\ S_3(a) &\leftarrow A(a, b); \\ S_3(a) &\leftarrow A(b, a), S_4(b); \\ S_4(a) &\leftarrow A(a, b), S_3(b); \\ \mathbf{goal} &\leftarrow S_1(a), S_4(a); \\ \mathbf{goal} &\leftarrow S_2(a), S_3(a). \end{aligned}$$

This Datalog program accepts any given σ -structure A^* if and only if A^* admits a homomorphism from some element of \mathcal{F} . If the program rejects A^* , then it constructs a $(\sigma \cup \tau)$ -expansion A of A^* , which will be the canonical structure given by (4.1).

Datalog provides an explanation for the membership tests for the expanded class given in Section 4. Lemma 4.7 corresponds precisely to the situation where the **goal** predicate is set to **true** by the Datalog program (meaning that A^* is not \mathcal{F} -free). In Lemma 4.9, the tuple trace condition corresponds to reaching a fixed point of the Datalog program, that is, the application of no rule results in adding an element into an IDB.

Constraint satisfaction problems. For a σ -structure H , let $\text{CSP}(H) = \{A : \exists f : A \xrightarrow{h} H\}$. Given a fixed σ -structure H , the *non-uniform constraint satisfaction problem* is to decide, for an input σ -structure A , whether $A \in \text{CSP}(H)$ or not. The problem's computational complexity depends on H ; many polynomial-time cases can be explained by the existence of a “nice” *obstruction set* \mathcal{F} such that $\text{CSP}(H) = \text{Forb}_h(\mathcal{F})$. Following [14], we say that H has *tree duality* if $\text{CSP}(H) = \text{Forb}_h(\mathcal{F})$ for

$$\mathcal{F} = \{F : F \text{ is a } \sigma\text{-tree and there is no } f : F \xrightarrow{h} H\}. \quad (8.1)$$

For \mathcal{F} given by (8.1), the \approx -equivalence class of any piece (M, m) of some $F \in \mathcal{F}$ is fully determined by the set

$$\mathfrak{H}(M, m) = \{f(m) : f : M \xrightarrow{h} H\},$$

because $\mathfrak{I}(M, m) = \{(N, n) : (N, n) \text{ is a rooted } \sigma\text{-tree s.t. } \mathfrak{H}(M, m) \cap \mathfrak{H}(N, n) = \emptyset\}$. Thus the expanded signature will always be finite and we can attempt to index the τ -relations by subsets of the domain of H (in correspondence with the sets $\mathfrak{H}(M, m)$).

What we get will be the canonical tree Datalog program: The EDBs of the program are – as always – the relations in σ . There is an IDB S_X for every proper subset $X \subset \text{dom } H$ (in the end we will only use the subsets which are definable in H by a positive existential first-order formula); S_X is unary unless $X = \emptyset$: S_\emptyset is nullary and we identify it with the **goal** predicate. Moreover, to simplify the description of the program's rules, identify $S_{\text{dom } H}$

with **true**. Given any $R \in \sigma$ of arity r , $j \in \{1, 2, \dots, r\}$ and nonempty sets $X_i \subseteq \text{dom } H$ for all $i \in \{1, 2, \dots, r\} \setminus \{j\}$, put

$$X_j = \{x_j \in \text{dom } H : \exists(x_1, x_2, \dots, x_r) \in R^H \text{ s.t. } x_i \in X_i \text{ for each } i \neq j\}.$$

If $X_j \neq \text{dom } H$, introduce the rule

$$S_{X_j}(a_j) \leftarrow R(a_1, a_2, \dots, a_r), S_{X_1}(a_1), \dots, S_{X_{j-1}}(a_{j-1}), S_{X_{j+1}}(a_{j+1}), \dots, S_{X_r}(a_r).$$

(At this point, subsets X not definable in H by a positive existential first-order formula will not appear in the head of any rule and we can drop them – as well as any rules containing them in their body.)

Example 8.2. In our example (thunderbolts), it is well known that $\text{Forb}_h(\mathcal{F}) = \text{CSP}(H)$ for $H = P_2$, the directed path $0 \rightarrow 1 \rightarrow 2$. The description above results in the following Datalog program:

$$\begin{aligned} S_{\{0,1\}}(a) &\leftarrow A(a, b); \\ S_{\{1,2\}}(a) &\leftarrow A(b, a); \\ S_{\{0\}}(a) &\leftarrow A(a, b), S_{\{0,1\}}(b); \\ S_{\{1,2\}}(a) &\leftarrow A(b, a), S_{\{0,1\}}(b); \\ S_{\{0,1\}}(a) &\leftarrow A(a, b), S_{\{1,2\}}(b); \\ S_{\{2\}}(a) &\leftarrow A(b, a), S_{\{1,2\}}(b); \\ \mathbf{goal} &\leftarrow A(a, b), S_{\{0\}}(b); \\ S_{\{1\}}(a) &\leftarrow A(b, a), S_{\{0\}}(b); \\ S_{\{0\}}(a) &\leftarrow A(a, b), S_{\{1\}}(b); \\ S_{\{2\}}(a) &\leftarrow A(b, a), S_{\{1\}}(b); \\ S_{\{1\}}(a) &\leftarrow A(a, b), S_{\{2\}}(b); \\ \mathbf{goal} &\leftarrow A(b, a), S_{\{2\}}(b). \end{aligned}$$

You may notice that the program is different to the one we derived from the thunderbolts. This is because the obstruction set \mathcal{F} has changed: it now contains not only the thunderbolts, but also all other trees that do not admit a homomorphism to P_2 .

Pigeonhole classes. For $(\sigma \cup \tau)$ -structures B, C and a positive integer r , let $C \rightarrow (B)_r^1$ denote the following statement: Whenever the elements of C are coloured with r colours, there exists an embedding $g : B \xrightarrow{e} C$ such that for any $b_1, b_2 \in \text{dom } B$, if $g(b_1)$ and $g(b_2)$ induce isomorphic one-element substructures of C , then $g(b_1)$ and $g(b_2)$ have the same colour. We say that a class \mathcal{C} of $(\sigma \cup \tau)$ -structures is a *pigeonhole class* if for any structure $B \in \mathcal{C}$ and positive integer r there exists $C \in \mathcal{C}$ such that $C \rightarrow (B)_r^1$.

Consider a set \mathcal{F} of σ -trees and let \mathcal{C} be the expanded class for $\text{Forb}_h(\mathcal{F})$. Suppose that \mathcal{C} contains finitely many non-isomorphic one-element structures: this is certainly the case if the expanded signature $\sigma \cup \tau$ is finite, in particular, if \mathcal{F} is given by (8.1). Then it follows by repeated application of Theorem 5.2 that the expanded class is a pigeonhole class (in fact, it follows for the ordered expanded class, but if we only colour one-element substructures, the ordering is not needed). It has already been proved by Atserias and Weyer [1] that any class with free amalgamation is a pigeonhole class.

Characterising tree duality. Following the approach of [1] further, let H be a σ -structure, let \mathcal{F} be given by (8.1) and let \mathcal{C} be the expanded class for $\text{Forb}_h(\mathcal{F})$. We need to construct $V \in \mathcal{C}$ such that any tuple trace appearing in \mathcal{C} will be an induced substructure of V : we can take the sum (disjoint union) of all such tuple traces. Now let $W \in \mathcal{C}$ satisfy $W \rightarrow (V)_r^1$ for $r = |\text{dom } H|$. Define \sim on $\text{dom } V$ by putting $v_1 \sim v_2$ if and only if v_1 and v_2 induce isomorphic one-element substructures of V and put $U = V/\sim$. Finally, let U^* , V^* and W^* be the σ -reducts of U , V and W , respectively. Using the pigeonhole property of W , one gets the following:

Proposition 8.3 ([1]). *The following conditions are equivalent:*

- (a) *There exists a homomorphism $W^* \xrightarrow{h} H$.*
- (b) *There exists a homomorphism $U^* \xrightarrow{h} H$.*
- (c) $\text{CSP}(H) = \text{Forb}_h(\mathcal{F})$. □

Thus we have found for any H a σ -structure $U^* = U^*(H)$ such that H has tree duality if and only if there is a homomorphism $U^*(H) \xrightarrow{h} H$. This $U^*(H)$ appears to be intimately related to the power structure of [10].

Lastly, we only mention in passing that another connection between constraint satisfaction problems and Ramsey classes is studied in [2].

9. FINAL COMMENTS

Universal structures. If \mathcal{F} is a set of finite σ -trees, then by Fraïssé's Theorem [13], Theorem 4.10 implies that the expanded class \mathcal{C} for $\text{Forb}_h(\mathcal{F})$ has a Fraïssé limit: a countable homogeneous $(\sigma \cup \tau)$ -structure U such that \mathcal{C} is the class of all finite substructures of U . The σ -reduct U^* of U is a universal structure for $\text{Forb}_h(\mathcal{F})$. For finite \mathcal{F} this universal structure U^* is ω -categorical; the existence of such a universal ω -categorical structure (and much more) was proved by Cherlin, Shelah and Shi [5]. If \mathcal{F} is infinite, U^* is no longer necessarily ω -categorical (see [16] and the next paragraph); however, it is model-complete.

Regular classes of trees. Recall that two pieces (M, m) , (M', m') are \approx -equivalent if their incompatible sets are equal, that is, if $\mathfrak{I}(M, m) = \mathfrak{I}(M', m')$. By Definition 4.4, the signature τ is finite if and only if \approx has finitely many equivalence classes on the pieces of the trees contained in \mathcal{F} . In this case, we call \mathcal{F} a *regular class* of σ -trees; the term is motivated by a connection to regular languages, highlighted in [9]. This definition of regularity coincides with the one from [16]. In [8], however, a set \mathcal{F} of trees is defined to be regular if \approx has finitely many equivalence classes on all rooted σ -forests. Let us call such a set \mathcal{F} *EPTT-regular*. Obviously, every EPTT-regular set is regular, but the converse does not hold.

Let $\text{UP}(\mathcal{F}) = \{F : F \text{ is a } \sigma\text{-tree and there exists } F' \in \mathcal{F} \text{ s.t. } F' \xrightarrow{h} F\}$. Obviously, $\text{Forb}_h(\text{UP}(\mathcal{F})) = \text{Forb}_h(\mathcal{F})$. For a (not necessarily finite) σ -structure H , define $\text{CSP}(H) = \{A : A \xrightarrow{h} H\}$. From [8, 16] and the results of this paper, we can conclude:

Theorem 9.1. *Let σ be a finite relational signature and let \mathcal{F} be a set of finite σ -trees. Then the following are equivalent:*

- (a) $\text{UP}(\mathcal{F})$ is regular;
- (b) $\text{UP}(\mathcal{F})$ is EPTT-regular;

- (c) $\text{Forb}_h(\mathcal{F}) = \text{CSP}(H)$ for some finite σ -structure H ;
- (d) there is a countable ω -categorical σ -structure U^* universal for $\text{Forb}_h(\mathcal{F})$;
- (e) there is a countable homogeneous $(\sigma \cup \tau)$ -structure U over a finite signature $\sigma \cup \tau$ such that the σ -reduct of U is universal for $\text{Forb}_h(\mathcal{F})$;
- (f) there is a Ramsey class $\vec{\mathcal{C}}$ of ordered $(\sigma \cup \tau)$ -structures over a finite signature $\sigma \cup \tau$ such that $\text{Forb}_h(\mathcal{F})$ is the class of the σ -reducts of the structures in $\vec{\mathcal{C}}$. \square

Extreme amenability. By a theorem of Kechris, Pestov and Todorćević [17], the automorphism group of a Ramsey structure is extremely amenable. Thus Theorem 5.2 provides a continuum of examples of structures with an extremely amenable automorphism group: take \mathcal{F}' to be an infinite antichain of σ -trees; then the Fraïssé limit of the ordered expanded class for $\text{Forb}_h(\mathcal{F})$ provides such an example for any subset \mathcal{F} of \mathcal{F}' .

Problem. It would be interesting to classify all sets \mathcal{F} of σ -structures for which the corresponding ordered expanded class for $\text{Forb}_h(\mathcal{F})$ is a Ramsey class. In particular, is it the case for any set \mathcal{F} of connected finite σ -structures?

Limits of the partite method. Nešetřil [21] asked whether one can prove all Ramsey classes by a variant of the partite (amalgamation) construction. This is certainly a question worth considering. It is not very satisfactory that the definition of a partite structure is rather different each time: compare [3, 20, 23, 24, 25, 26, 27, 28]. Also, the partite lemma is sometimes proved by induction (as in [3, 29] and here), sometimes by an application of the Hales–Jewett theorem (as in [26, 27, 28]).

ACKNOWLEDGEMENTS

I thank Albert Atserias, Manuel Bodirsky, Jarik Nešetřil and Claude Tardif for many stimulating discussions, intriguing questions and useful comments. Many thanks go to the anonymous referees whose questions and objections have led to this paper being much better than it would otherwise have been.

I am fully supported by a fellowship from the Swiss National Science Foundation. Part of this work was done at Laboratoire d’Informatique (LIX), CNRS UMR 7161, Ecole Polytechnique, Paris, where I was supported by ERC Starting Grant “CSP-Complexity” no. 257039 of Manuel Bodirsky. Another part was done when I was visiting Fields Institute in Toronto within the Summer Thematic Program on the Mathematics of Constraint Satisfaction. Finally, a part was done at Coffee&Company in Kingston, Ontario. I also acknowledge the support of the Institute for Theoretical Computer Science (ITI), project 1M0545 of the Ministry of Education of the Czech Republic, and EPSRC grant EP/I01795X/1.

REFERENCES

- [1] A. Atserias and M. Weyer. Decidable relationships between consistency notions for constraint satisfaction problems. In E. Grädel and R. Kahle, editors, *Computer Science Logic*, volume 5771 of *Lecture Notes in Comput. Sci.*, pages 102–116, 2009.
- [2] M. Bodirsky and M. Pinsker. Reducts of Ramsey structures. In M. Grohe and J. A. Makowsky, editors, *Model Theoretic Methods in Finite Combinatorics*, volume 558 of *Contemporary Mathematics*. AMS, 2011.
- [3] J. Böttcher and J. Foniok. Ramsey properties of permutations. *Electron. J. Combin.*, 20(1), 2013. Paper 2, 10 pp.
- [4] A. A. Bulatov, A. Krokhin, and B. Larose. Dualities for constraint satisfaction problems. In Creignou et al. [7], pages 93–124.
- [5] G. Cherlin, S. Shelah, and N. Shi. Universal graphs with forbidden subgraphs and algebraic closure. *Adv. in Appl. Math.*, 22(4):454–491, 1999.
- [6] D. Cohen and P. Jeavons. The complexity of constraint languages. In F. Rossi, P. van Beek, and T. Walsh, editors, *Handbook of Constraint Programming*, volume 2 of *Foundations of Artificial Intelligence*, chapter 8. Elsevier, 2006.
- [7] N. Creignou, P. G. Kolaitis, and H. Vollmer, editors. *Complexity of Constraints: An Overview of Current Research Themes [Result of a Dagstuhl Seminar]*, volume 5250 of *Lecture Notes in Computer Science*. Springer, 2008.
- [8] P. L. Erdős, D. Pálvölgyi, C. Tardif, and G. Tardos. Regular families of forests, antichains and duality pairs of relational structures. arXiv: 1207.4402v2 [math.CO], 2014.
- [9] P. L. Erdős, C. Tardif, and G. Tardos. Caterpillar dualities and regular languages. *SIAM J. Discrete Math.*, 27(3):1287–1294, 2013.
- [10] T. Feder and M. Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory. *SIAM J. Comput.*, 28(1):57–104, 1999.
- [11] J. Folkman. Graphs with monochromatic complete subgraphs in every edge coloring. *SIAM J. Appl. Math.*, 18(1):19–24, 1970.
- [12] J. Foniok. *Homomorphisms and Structural Properties of Relational Systems*. PhD thesis, Charles University, Prague, 2007.
- [13] R. Fraïssé. Sur certaines relations qui généralisent l’ordre des nombres rationnels. *C. R. Acad. Sci. Paris*, 237:540–542, 1953.
- [14] P. Hell, J. Nešetřil, and X. Zhu. Duality and polynomial testing of tree homomorphisms. *Trans. Amer. Math. Soc.*, 348(4):1281–1297, 1996.
- [15] J. Hubička and J. Nešetřil. Homomorphism and embedding universal structures for restricted classes. arXiv: 0909.4939v1 [math.CO], 2009.
- [16] J. Hubička and J. Nešetřil. Universal structures with forbidden homomorphisms. arXiv: 0907.4079v4 [math.CO], 2013.
- [17] A. S. Kechris, V. G. Pestov, and S. Todorčević. Fraïssé limits, Ramsey theory, and topological dynamics of automorphism groups. *Geom. Funct. Anal.*, 15(1):106–189, 2005.
- [18] P. G. Kolaitis. On the expressive power of logics on finite models. In *Finite model theory and its applications*, Texts Theoret. Comput. Sci. EATCS Ser., pages 27–123. Springer, Berlin, 2007.
- [19] J. Nešetřil. Ramsey classes and homogeneous structures. *Combin. Probab. Comput.*, 14(1–2):171–189, 2005.
- [20] J. Nešetřil. Metric spaces are Ramsey. *European J. Combin.*, 28(1):457–468, 2007.
- [21] J. Nešetřil. First order definable Ramsey classes. Workshop on Homogeneous Structures, July 2011. School of Mathematics, University of Leeds.
- [22] J. Nešetřil and V. Rödl. Partitions of finite relational and set systems. *J. Combinatorial Theory Ser. A*, 22(3):289–312, 1977.
- [23] J. Nešetřil and V. Rödl. Simple proof of the existence of restricted Ramsey graphs by means of a partite construction. *Combinatorica*, 1(2):199–202, 1981.
- [24] J. Nešetřil and V. Rödl. Two proofs of the Ramsey property of the class of finite hypergraphs. *European J. Combin.*, 3(4):347–352, 1982.
- [25] J. Nešetřil and V. Rödl. Ramsey classes of set systems. *J. Combin. Theory Ser. A*, 34(2):183–201, 1983.
- [26] J. Nešetřil and V. Rödl. Strong Ramsey theorems for Steiner systems. *Trans. Amer. Math. Soc.*, 303(1):183–192, 1987.

- [27] J. Nešetřil and V. Rödl. The partite construction and Ramsey set systems. *Discrete Math.*, 75(1–3):327–334, 1989.
- [28] J. Nešetřil. Ramsey theory. In R. L. Graham, M. Grötschel, and L. Lovász, editors, *Handbook of Combinatorics*, volume 2, chapter 25, pages 1331–1403. Elsevier, 1995.
- [29] H. J. Prömel and B. Voigt. A short proof of the restricted Ramsey theorem for finite set systems. *J. Combin. Theory Ser. A*, 52(2):313–320, 1989.