A characterization of idempotent strong Mal'cev conditions for congruence meet-semidistributivity in locally finite varieties

Nemanja Draganić, Petar Marković, Vlado Uljarević, Samir Zahirović
Department of Mathematics and Informatics
University of Novi Sad
Serbia

March 12, 2018

Abstract

We prove a characterization of all idempotent, linear, strong Mal'cev conditions in two variables which hold in all locally finite congruence meet-semidistributive varieties. This is an alternative proof to the one previously given by Z. Brady in [4], and has some advantages, some disadvantages, to his approach. Along the way we prove that such a strong Mal'cev condition holds in all locally finite congruence meet-semidistributive varieties iff it is realized in a certain four-element algebra.

To Bjarni Jónsson,
who invented congruence meet-semidistributivity,
who made an area out of Mal'cev conditions,
who taught us all how to write down a pretty proof.

1 Introduction

The various conditions which are equivalent to congruence meet-semidistributivity in locally finite varieties of algebras have been explored in several previous papers and books [6], [8], [13], [17], [24], [15], [10] and in the recent work of Z. Brady [4]. The reason for this activity is that congruence meet-semidistributive varieties are a very general, and yet well-behaved class of varieties. Congruence meet-semidistributivity is equivalent to congruence neutrality (trivialization of the commutator) [13] and [17]; in locally finite varieties it is characterized by omitting tame congruence theory types **1** and **2** [8]; the Park conjecture is true in congruence meet-semidistributive varieties [24], as is the Restricted Quackenbush conjecture [14]; it characterizes the algebraic duals of the finite relational structures $\mathbb{A}$ such that the constraint satisfaction problem with template $\mathbb{A}$ can be accurately solved by using only the local consistency checking [16], [2], see also [1].

In this paper we characterize all linear strong Mal'cev conditions in two variables which hold in all locally finite congruence meet-semidistributive varieties. M. Siggers proved in [22] that the weaker property, having a Taylor term (characterized in locally finite varieties by omitting type **1**) is a strong Mal'cev property, when restricted to locally finite varieties. Siggers' result was a big surprise at the time of publication and spurred an investigation of what other properties, previously known to have a Mal'cev characterization, also have a strong Mal'cev characterization in locally finite varieties. The paper [12] settled the question of optimal (syntactically simplest) strong Mal'cev characterizations for having a Taylor term (= omitting type **1**) in locally finite varieties.

Three recent papers have influenced our work. The paper [15] by M. Kozik, A. Krokhin, M. Valeriote and R. Willard invented the first strong Mal'cev condition which is equivalent, in locally finite varieties, to congruence meet-semidistributivity, while many other natural properties were proved not to have a strong Mal'cev characterization in the same paper. It also posed two problems about some non-strong Mal'cev conditions for congruence meet-semidistributivity, whether they can be made strong by limiting the arity of the pertinent operations. Willard reiterated those questions at a 2016 workshop. The easier of the two questions is answered positively by an application of a theorem from the paper [10] by J. Jovanović, P. Marković, R. McKenzie and M. Moore. The harder one needed a different proof which provided immediate motivation to our work.

The paper [10] proved the syntactically easiest strong Mal'cev condition for congruence meet-semidistributivity by modifying the technique of a proof from [15], namely, introducing an auxiliary structure on the variables of the Constraint Satisfaction Problem, which determines which variables are in which constraint relations. Thus a simple pigeonhole argument from [15] becomes a more complicated Ramsey proof, proving that some substructure is monochromatic, but in exchange, this enables proofs of more powerful results. Our main theorem is proved using the same idea, using a more general language and stronger Ramsey arguments.

Finally, the paper [4] by Z. Brady, which was available to us after we proved initial results of our investigation (just after we answered the questions from [15]), gave a list of two-variable linear strong Mal'cev conditions in one operation symbol which are maximal in a sense and which all have the same derived binary operation, with additional properties of that binary operation. This also answers the same questions from [15], and proves much more. Our main theorem gives an alternative proof of that Brady's result, and adds a polynomial-time algorithm for checking whether a strong Mal'cev condition of the type under consideration holds in all locally finite congruence meet-semidistributive varieties. Also, we believe that the Ramsey techniques of our paper, while much less innovative, may be used for proofs of Mal'cev conditions in more than two variables, while Brady's major ideas like the Semilattice Preparation Lemma, seem to be limited to two variables. The initial investigation into the Mal'cev conditions in more than two variables revealed that there is no inherent difficulty in generalizing our ideas there, but the Ramsey statements to which the Mal'cev conditions reduce seem to be quite complicated.

A more detailed history of the progress of Mal'cev characterizations of congruence meet-semidistributivity is available in [10].

2 Background

By a **strong Mal'cev condition** we mean a finite set of identities in some language. Informally, a strong Mal'cev condition is realized in an algebra $\mathbf{A}$ (or variety $\mathcal{V}$) if there is a way to interpret the function symbols appearing in the condition as term operations of $\mathbf{A}$ (or $\mathcal{V}$) so that the identities in the Mal'cev condition become true equations in $\mathbf{A}$ (or $\mathcal{V}$). A **Mal'cev condition** is a sequence $\{C_n : n \in \omega\}$ of strong Mal'cev conditions such that any variety which realizes C_n must also realize C_{n+1} for all $n \in \omega$. We say that the variety $\mathcal{V}$ realizes the Mal'cev condition $\{C_n : n \in \omega\}$ if there exists an $n \in \omega$ such that $\mathcal{V}$ realizes C_n. We say that a varietal property P is a **(strong) Mal'cev property** if there exists a (strong) Mal'cev condition C such that for any variety $\mathcal{V}$, C is realized in $\mathcal{V}$ iff $\mathcal{V}$ has the property P. Also, the previous sentence is commonly relativized to locally finite varieties, so we say that some varietal property is a (strong) Mal'cev property of locally finite varieties if there exists a (strong) Mal'cev condition C such that for any locally finite variety $\mathcal{V}$, C is realized in $\mathcal{V}$ iff $\mathcal{V}$ has the property P.

A (strong) Mal'cev condition Σ is linear if it contains no composition of operations in any of its identities. Σ is idempotent if for all operations f in Σ the identities $f(x, x, \ldots, x) \approx x$ can be derived from Σ. Σ is trivial if it is realized in every variety (equivalently in the variety of sets, equivalently realized in some variety by interpreting each operation symbol as some projection).

W. Taylor proved in [23] that a variety realizes a nontrivial idempotent strong Mal'cev condition iff it realizes a nontrivial idempotent linear strong Mal'cev condition in one-operation language with identities in only two variables. This condition can be rewritten (by writing equations repeatedly) to be of the form

$$
\begin{aligned}
t(x, x, \ldots, x) &\approx x \\
t(a_{1,1}, a_{1,2}, ..., a_{1,n}) &\approx t(b_{1,1}, b_{1,2}, ..., b_{1,n}), \\
t(a_{2,1}, a_{2,2}, ..., a_{2,n}) &\approx t(b_{2,1}, b_{2,2}, ..., b_{2,n}), \\
&\vdots \\
t(a_{n,1}, a_{n,2}, ..., a_{n,n}) &\approx t(b_{n,1}, b_{n,2}, ..., b_{n,n}),
\end{aligned}
$$

where all for all i, j $a_{i,j}, b_{i,j} \in \{x, y\}$, $a_{i,i} = x$ and $b_{i,i} = y$. We will call such strong Mal'cev condition a **Taylor condition**. If the variety $\mathcal{V}$ realizes a Taylor condition C, then a term which interprets the only operation symbol used in C is called a **Taylor term** for the variety $\mathcal{V}$.

We say that an algebra is **congruence meet-semidistributive** if for any congruences $\alpha, \beta, \gamma \in \operatorname{Con} \mathbf{A}$, the following implication holds:

$$\alpha \wedge \beta = \alpha \wedge \gamma \Rightarrow \alpha \wedge \beta = \alpha \wedge (\beta \vee \gamma).$$

A variety $\mathcal{V}$ is congruence meet-semidistributive if every algebra in $\mathcal{V}$ is congruence meet-semidistributive.

Congruence meet-semidistributivity can be characterized in general varieties by a Mal'cev condition, see [24] and in locally finite varieties by a strong Mal'cev condition, see [15] and [10]. We need to state just one of these for our purpose, so we select the one having a nonempty intersection of authors with our paper.

Theorem 2.1 (Theorem 3.2 in [10])**.** *A locally finite variety $\mathcal{V}$ is congruence meet semidistributive iff $\mathcal{V}$ realizes the strong Mal'cev condition*

$$
\begin{aligned}
t(x, x, x, x) &\approx x \\
t(y, x, x, x) \approx t(x, y, x, x) &\approx t(x, x, y, x) \approx t(x, x, x, y) \\
&\approx t(y, y, x, x) \approx t(y, x, y, x) \approx t(x, y, y, x).
\end{aligned} \tag{2.1}
$$

Let us now introduce definitions of the constraint satisfaction problem and a $(2, 3)$-minimal instance of it. We follow [1] as we will use the main result of that paper a lot.

Definition 2.2. An **instance of the constraint satisfaction problem** (CSP) is triple $(V; A; \mathcal{C})$ with

- V a nonempty, finite set of variables,
- A a nonempty, finite domain,
- $\mathcal{C}$ a finite nonempty set of constraints, where each constraint is a subset C of A^W for some $W \subseteq V$. Here the subset W of V called the **scope** of C, and the cardinality $|W|$ of W is referred to as the arity of C.

A **solution** of the instance $(V; A; \mathcal{C})$ is a function $f : V \to A$ such that, for each constraint $C \in \mathcal{C}$, such that the scope of C is $W \subseteq V$, the restriction $f\restriction_W$ is in C. We say that an instance of CSP is **trivial** if it contains the empty constraint. Next we define a 2-consistent and a $(2,3)$-minimal instance.

Definition 2.3. An instance of $(V; A; \mathcal{C})$ is 2**-consistent**, if for every $U \subseteq V$ such that $|U| \leq 2$ and every pair of constraints $C, D \in \mathcal{C}$ such that U is contained in the scopes of both C and D, holds $C\restriction_U = D\restriction_U$. An instance of CSP $(V; A; \mathcal{C})$ is a 3**-dense** if every at most 3-element subset of V is contained in the scope of some constraint in $\mathcal{C}$. An instance is a $(\mathbf{2}, \mathbf{3})$**-minimal** instance if it is 2-consistent and 3-dense.

Before moving on with the following definition, let's note that we can observe any constraint with a scope W as a relation $R \subseteq A^{|W|}$ by linearly ordering the elements of W.

Definition 2.4. Let $\mathbb{A} = \langle A; \Gamma \rangle$ be a relational structure. An instance of the constraint satisfaction problem $CSP(\mathbb{A})$ is any instance of the CSP $(V; A; \mathcal{C})$ such that for each constraint $C \in \mathcal{C}$, there exists a relation $C' \in \Gamma$ such that C and C' are equal up to a permutation of coordinates. The structure $\mathbb{A}$ is called the **template** of $CSP(\mathbb{A})$.

We silently assume that all Γ contain the equality relation (to allow using the relations obtained by identification of variables).

Let $\mathbf{A}$ be an algebra. When $\Gamma \subseteq SP_{fin}(\mathbf{A})$, then we say that $CSP(\langle A; \Gamma \rangle)$ is compatible with $\mathbf{A}$. The following result is from [1]; it has as a consequence the main theorem of [1], and it will play an important role in our paper.

Theorem 2.5 (Corollary 6.5 of [1])**.** *Let $\mathbf{A}$ be an idempotent finite algebra which generates a congruence meet-semidistributive variety. Then for every $CSP(\langle A; \Gamma \rangle)$ which is compatible with $\mathbf{A}$, every nontrivial $(2,3)$-minimal instance of $CSP(\langle A; \Gamma \rangle)$ has a solution.*

The final area we need to introduce in this section is the basic Ramsey theory. In the seminal paper [21] F. P. Ramsey proved that for any positive integers k, N and n, there exists a number m such that for any set X of size m, if all k-element subsets of X are colored in n colors there exists a subset of size N whose all k-element subsets are of the same color. As usually, $R_k^n(N)$ denotes the smallest number m such that the previous sentence holds.

We recall some notation and terminology from partially ordered sets (posets): $X{\downarrow} = \{y : (\exists x \in X)\ y \leq x\}$ and $a{\downarrow} := \{a\}{\downarrow}$. The set X is called a down-set (order-ideal) when $X{\downarrow} = X$. $X{\uparrow}$, $a{\uparrow}$ and an up-set (order-filter) are defined dually.

We assume the reader is familiar with the basic notions of universal algebra, and refer those who might need some clarification to the standard textbooks [5], [19] and [3]. We will also use a little bit of Tame Congruence Theory developed in the seminal monograph [8], but only in the penultimate section, when we prove a rephrasing of our main results.

3 Decent Mal'cev conditions

All strong Mal'cev conditions which we consider in this paper will be idempotent, linear and on the set of variables $\{x, y\}$. The following proposition is folklore; it was probably first mentioned in [23], without proof.

Proposition 3.1. *For any linear idempotent strong Mal'cev condition Σ there exists a linear idempotent strong Mal'cev condition Σ' which has only one operation symbol such that for any variety $\mathcal{V}$, $\mathcal{V}$ realizes Σ iff $\mathcal{V}$ realizes Σ'.*

Proof. Assume that Σ involves operations f and g of arities k and n, respectively. We construct a strong Mal'cev condition Σ_1 in the following way: replace f and g with a new operation symbol h of arity kn so that

1. in any equation which contains f, we replace $f(x_1, \dots, x_k)$ with the term $h(x_1, \dots, x_1, x_2, \dots, x_2, \dots, x_k, \dots, x_k)$ (each variable repeated n times);

2. for any equation which involves g, we replace $g(x_1, \dots, x_n)$ with the term $h(x_1, \dots, x_n, x_1, \dots, x_n, \dots, x_1, \dots, x_n)$.

If $\mathcal{V}$ realizes Σ, the realization of Σ_1 in $\mathcal{V}$ is obtained by interpreting h as the term $f(g(x_1, \dots, x_n), g(x_{n+1}, \dots, x_{2n}), \dots, g(x_{(k-1)n+1}, \dots, x_{kn}))$. On the other hand, if $\mathcal{V}$ realizes Σ_1, the realization of Σ in $\mathcal{V}$ uses the h-terms from (1) and (2) as interpretations for f and g, respectively. The details are left to the reader.

Thus we removed one operation symbol and inductively we can keep doing it until we are left with just one. $\square$

Therefore, we may assume that any strong Mal'cev condition Σ we will consider involves only one operation symbol. Note that the procedure for obtaining Σ' described in the proof of Proposition 3.1 is deterministic up to the selection of ordered pairs of operations which get replaced by a single operation; the impact of the order of selection can be neutralized by permuting the variables of the only operation of Σ'.

Definition 3.2. We will call a strong Mal'cev condition which is linear, idempotent, and such that only two variables and one operation symbol occur in it, a *decent* Mal'cev condition.

In our deliberations on decent Mal'cev conditions we will need to represent them in different ways. First we notice that all terms involved in a decent Mal'cev condition with n-ary operation symbol f are either substitution instances of $f(x_1, x_2, \dots, x_n)$, or the variables x or y. Modulo idempotence, we may replace each x with $f(x, x, \dots, x)$ and y with $f(y, y, \dots, y)$, so now we assume that all terms involved in the identities of any decent Mal'cev condition, except for the idempotence, are not variables, but substitution instances of $f(x_1, x_2, \dots, x_n)$. We will use various sets of n distinct variables, not just $\{x_1, x_2, \dots, x_n\}$, to describe a decent Mal'cev condition.

We fix the countably infinite set of variables $Var = \{x_1, x_2, \dots\}$. Given a subset $U \subseteq Var$, we define the notation $x_j^U := y$ if $x_j \in U$, while if $x_j \notin U$, then $x_j^U := x$.

Definition 3.3. Let Σ be a decent Mal'cev condition with the only operation symbol f such that $n = ar(f)$, let the identities in Σ other than idempotence involve two operation symbols f (one in each term of the identity), and let $X = \{x_{i_1}, x_{i_2}, \dots, x_{i_n}\}$ be some n-element subset of Var, where $i_1 < i_2 < \dots < i_n$. We define the *representation of* Σ *on* X, written as $r_X(\Sigma)$, by

$$r_X(\Sigma) := \{(U, V) \in \mathcal{P}(X) \times \mathcal{P}(X) : f(x_{i_1}^U, \dots, x_{i_n}^U) \approx f(x_{i_1}^V, \dots, x_{i_n}^V) \in \Sigma\}.$$

Note that a decent Mal'cev condition has many representations, one for each n-element subset of Var, while each representation uniquely determines the decent Mal'cev condition. Since our manipulations of decent Mal'cev conditions will be reflected on their representations and invariants derived from them, whenever we introduce a decent Mal'cev condition, we will always include the set of variables X together with it (though X is not intrinsic to the condition itself, it is just an invariant of our description).

Definition 3.4. If Σ is a decent Mal'cev condition in the language with one operation f and let $X = \{x_{i_1}, \ldots, x_{i_n}\}$ be an n-element subset of Var. We define the binary relation $\epsilon(\Sigma)$ on the set $\mathcal{P}(X)$ to be the least equivalence relation which contains $r_X(\Sigma) \cup \{(U, V) \in (\mathcal{P}(X))^2 : (X \backslash U, X \backslash V) \in r_X(\Sigma)\}$.

The following easy lemma describes the impact of dummy variables on $\epsilon(\Sigma)$.

Lemma 3.5. *Let* $\mathbf{A}$ *be an algebra and* Σ *a decent Mal'cev condition represented on* $X = \{x_1, \ldots, x_n\}$. $\mathbf{A}$ *realizes* Σ *by interpreting its only operation symbol* f *as some* $\mathbf{A}$*-term* $t(x_{i_1}, \ldots, x_{i_k})$*, where* $E = \{x_{i_1}, \ldots, x_{i_k}\} \subseteq \{x_1, \ldots, x_n\}$ *iff* $\mathbf{A}$ *realizes the decent Mal'cev condition*

$$\Sigma' = \{f'(x_{i_1}^{U \cap E}, \ldots, x_{i_k}^{U \cap E}) \approx f'(x_{i_1}^{V \cap E}, \ldots, x_{i_k}^{V \cap E}) : (U, V) \in \epsilon(\Sigma)\}.$$

(together with the idempotence of f'*). Hence* $\epsilon(\Sigma')$ *is the equivalence relation on* $\mathcal{P}(E)$ *generated by* $\{(U \cap E, V \cap E) : (U, V) \in \epsilon(\Sigma)\}$.

Proof. Let $t'(x_{i_1}, \ldots, x_{i_k})$ be an $\mathbf{A}$-term such that interpreting f as t' induces a realization of Σ in $\mathbf{A}$. Define the $\mathbf{A}$-term $t(x_1, \ldots, x_n) := t'(x_{i_1}, \ldots, x_{i_k})$, the term syntactically equal to t' but with additional dummy variables. If $(U, V) \in \epsilon(\Sigma)$ then $\mathbf{A} \models t(x_1^U, \ldots, x_n^U) \approx t(x_1^V, \ldots, x_n^V)$ or, equivalently, $\mathbf{A} \models t'(x_{i_1}^U, \ldots, x_{i_k}^U) \approx t'(x_{i_1}^V, \ldots, x_{i_k}^V)$. Observe that all for all $1 \leq j \leq k$, $x_{i_j} \in E$, implying $\mathbf{A} \models t'(x_{i_1}^{U \cap E}, \ldots, x_{i_k}^{U \cap E}) \approx t'(x_{i_1}^{V \cap E}, \ldots, x_{i_k}^{V \cap E})$. From this and the definition of Σ' follows that $\mathbf{A}$ realizes Σ'.

On the other hand, assume that $\mathbf{A}$ realizes Σ' by interpreting f' as the $\mathbf{A}$-term t'. Consider any $f(x_1^U, \ldots, x_n^U) \approx f(x_1^V, \ldots, x_n^V) \in \Sigma$. Thus $(U, V) \in \epsilon(\Sigma)$, and $\mathbf{A} \models t(x_{i_1}^{U \cap E}, \ldots, x_{i_k}^{U \cap E}) \approx t(x_{i_1}^{V \cap E}, \ldots, x_{i_k}^{V \cap E})$ by the definition of Σ'. By interpreting f as $t(x_1, \ldots, x_n)$, the same term as t' with the additional dummy variables, we see that $t(x_1^U, \ldots, x_n^U) = t'(x_{i_1}^{U \cap E}, \ldots, x_{i_k}^{U \cap E})$ and $t(x_1^V, \ldots, x_n^V) = t'(x_{i_1}^{V \cap E}, \ldots, x_{i_k}^{V \cap E})$, since the pairs of terms, as sequences of symbols, are identical. Hence, $\mathbf{A} \models t(x_1^U, \ldots, x_n^U) \approx t(x_1^V, \ldots, x_n^V)$, so Σ is thus realized in $\mathbf{A}$. $\square$

Convention. To shorten notation we adopt the following convention about decent Mal'cev conditions: The operation symbol of Σ will be called f, the operation symbol of Σ' will be called f' and the operation symbol of Π will be called g, where Σ, Σ' and Π are decent Mal'cev conditions.

3.1 Decent Mal'cev conditions in the majority algebra

Let $\mathbf{A} = (\{0,1\}; m)$ be the unique two-element algebra with ternary majority operation m, i.e.

$$\mathbf{A} \models m(x,x,y) \approx m(x,y,x) \approx m(y,x,x) \approx x.$$

We investigate which decent Mal'cev conditions are realized in $\mathbf{A}$.

Lemma 3.6. $\mathbf{A}$ *realizes a decent Mal'cev condition* Σ *represented on* $X = \{x_1, \ldots, x_n\}$ *iff* $\mathbf{A}$ *realizes some decent Mal'cev condition* Π *represented on* X *such that, where* $\rho = \epsilon(\Pi)$*:*

1. $\epsilon(\Sigma) \subseteq \rho$

2. ρ *has exactly two equivalence classes* $[\emptyset]_\rho$ *and* $[X]_\rho$*, and* $[\emptyset]_\rho = \{X \setminus U : U \in [X]_\rho\}$ *(hence,* $|[\emptyset]_\rho| = |[X]_\rho|$*);*

3. $[\emptyset]_\rho$ *is a down-set (order-ideal) and* $[X]_\rho$ *is an up-set (order-filter);*

Also, the same term which interprets f *in the realization of* Σ *can be taken as the term which interprets* g *in the realization of* Π.

Proof. The implication "if" follows from $\epsilon(\Sigma) \subseteq \epsilon(\Pi)$, so we only need to prove "only if". Let $t(x_1, \ldots, x_n)$ be a term on the language $\{m\}$ such that interpreting f as t induces a realization of Σ in $\mathbf{A}$. We recall that, for every term $s(x_1, \ldots, x_n)$ in the language $\{m\}$ and tuple $(a_1, \ldots, a_n)$ such that $\{a_1, \ldots, a_n\} \subseteq \{x, y\}$, either $\mathbf{A} \models s(a_1, \ldots, a_n) \approx x$ or $\mathbf{A} \models s(a_1, \ldots, a_n) \approx y$.

The decent Mal'cev condition Π represented on X is defined as $\Pi := \{g(x_1^U, \ldots, x_n^U) \approx x : \mathbf{A} \models t(x_1^U, \ldots, x_n^U) \approx x\}$. Clearly, $\mathbf{A}$ realizes Π by interpreting g as t. Also, $\epsilon(\Sigma) \subseteq \rho$ and $\rho = \epsilon(\Pi)$ has exactly two equivalence classes, $[\emptyset]_\rho$ and $[X]_\rho$ (note that $[\emptyset]_\rho \neq [X]_\rho$ by idempotence and since $\mathbf{A} \not\models x \approx y$). Hence for all $U \in \mathcal{P}(X)$, $[U]_\rho \neq [X \setminus U]_\rho$, as $(\emptyset, U) \in \rho$ implies $(X \setminus \emptyset, X \setminus U) = (X, X \setminus U) \in \rho$. Thus also $|[\emptyset]_\rho| = |[X]_\rho|$.

Now we will prove that $[\emptyset]_\rho$ is a down-set. Let $U \in [\emptyset]_\rho$ and $V \subseteq U$. Thus $\mathbf{A} \models t(x_1^U, \ldots, x_n^U) \approx x$ and we have to prove $\mathbf{A} \models t(x_1^V, \ldots, x_n^V) \approx x$. We prove it by an induction on the complexity of t. If $t(x_1, \ldots, x_n)$ is a variable or $m(x_i, x_j, x_k)$ then it is obvious. Assume that the claim holds for every term with fewer operation symbols than $t(x_1, \ldots, x_n)$ and let $\mathbf{A} \models t(x_1^U, \ldots, x_n^U) \approx x$. Then $t = m(t_1, t_2, t_3)$ where the inductive hypothesis holds for $t_1(x_1, \ldots, x_n)$, $t_2(x_1, \ldots, x_n)$ and $t_3(x_1, \ldots, x_n)$. Since $\mathbf{A} \models t_j(x_1^U, \ldots, x_n^U) \approx x$ or $\mathbf{A} \models t_j(x_1^U, \ldots, x_n^U) \approx y$ for each $j = 1, 2, 3$, and

m is the majority, then for at least two j, $\mathbf{A} \models t_j(x_1^U, \dots, x_n^U) \approx x$. Without loss of generality we may assume $\mathbf{A} \models t_1(x_1^U, \dots, x_n^U) \approx t_2(x_1^U, \dots, x_n^U) \approx x$. By the induction hypothesis, $\mathbf{A} \models t_1(x_1^V, \dots, x_n^V) \approx t_2(x_1^V, \dots, x_n^V) \approx x$, so $\mathbf{A} \models t(x_1^V, \dots, x_n^V) \approx m(x, x, ?) \approx x$ where $? \in \{x, y\}$. $[X]_\rho$ is an up-set, being the complement of $[\emptyset]_\rho$ in $\mathcal{P}(X)$. $\square$

Remark 3.7. Assume that Σ is realized in $\mathbf{A}$ by interpreting f as the $\mathbf{A}$-term $t(x_{i_1}, \dots, x_{i_k})$, where all x_{i_j} occur in the syntactic expression for t. Then for any E such that $\{x_{i_1}, \dots, x_{i_k}\} \subseteq E \subseteq \{x_1, \dots, x_n\}$ there exists an $\mathbf{A}$-term t such that all variables in E occur in the syntactical expression for t and $\mathbf{A}$ realizes a decent Mal'cev condition Π represented on E such that $\rho = \epsilon(\Pi)$ satisfies (2)-(3) of Lemma 3.6 and $\{(U \cap E, V \cap E) : (U, V) \in \epsilon(\Sigma)\} \subseteq \rho$ by interpreting g as t. In the case when $E = \{x_{i_1}, \dots, x_{i_k}\}$, this is clear from Lemmas 3.5 and 3.6. For bigger E, we can keep adding new variables x which act as dummies though they occur in the term by repeatedly using $t' = m(t, t, x)$.

The following definition plays the role of *intersecting families* from [4]. We use it for an algorithmic result near the end of this Section.

Definition 3.8. An equivalence relation ϵ on $\mathcal{P}(X)$ is *0-1 distinguishing* if there exist no sets U, V, W, Z in $\mathcal{P}(X)$ such that:

(i) $U \epsilon V \epsilon W \epsilon Z$;

(ii) $U \cap V = \emptyset$;

(iii) $W \cup Z = X$.

Lemma 3.9. *Any equivalence relation ρ on $\mathcal{P}(X)$ which satisfies* (2) *and* (3) *of Lemma 3.6 is 0-1 distinguishing.*

Proof. Assume that ρ is not 0-1 distinguishing. Then there exist U, V, W, Z in $\mathcal{P}(X)$ such that conditions $(i) - (iii)$ from Definition 3.8 hold. Assume that $U, V, W, Z \in [\emptyset]_\rho$. From (iii) we have that $(X \setminus W) \subseteq Z$ and then by (3) of Lemma 3.6 we conclude $(X \setminus W) \in [\emptyset]_\rho$, so $W \rho \, (X \setminus W)$, a contradiction. The case $U, V, W, Z \in [X]_\rho$ leads to an analogous contradiction with condition (ii) of Definition 3.8. $\square$

Definition 3.10. Given an equivalence relation ϵ on $\mathcal{P}(X)$, we define the binary relation $\preceq_\epsilon$ on $\mathcal{P}(X)/\epsilon$ by $[U]_\epsilon \preceq_\epsilon [V]_\epsilon$ iff there exist $U' \in [U]_\epsilon$ and $V' \in [V]_\epsilon$ such that $U' \subseteq V'$. Let $\leq_\epsilon$ be the transitive closure of $\preceq_\epsilon$. If $[U]_\epsilon \leq_\epsilon [V]_\epsilon$ and $[V]_\epsilon \leq_\epsilon [U]_\epsilon$, then we say $[U]_\epsilon \sim_\epsilon [V]_\epsilon$.

By its definition, $\leq_\epsilon$ is reflexive and transitive, while $\sim_\epsilon$ is an equivalence relation on $\mathcal{P}(X)/\epsilon$.

Definition 3.11. Given an equivalence relation ϵ on $\mathcal{P}(X)$, its *closure* $\overline{\varepsilon}$ will denote $\{(U,V) : U,V \in \mathcal{P}(X) \text{ and } [U]_\epsilon \sim_\epsilon [V]_\epsilon\}$.

Definition 3.12. Also, given an equivalence relation ϵ on $\mathcal{P}(X)$, we introduce the notation $C(\epsilon) := \{(X \setminus U, X \setminus V) : (U,V) \in \epsilon\}$. Clearly, the equivalence $(U,V) \in \epsilon$ iff $(X \setminus U, X \setminus V) \in \epsilon$ holds if and only if $\epsilon = C(\epsilon)$.

Lemma 3.13. *For any equivalence relation ϵ on $\mathcal{P}(X)$, its closure $\overline{\epsilon}$ is convex, i.e. if $U \subseteq V \subseteq W \subseteq X$ and $(U,W) \in \overline{\epsilon}$, then $(U,V) \in \overline{\epsilon}$. Moreover, if $\epsilon = C(\epsilon)$, then for all $U,V \subseteq X$, $[U]_\epsilon \leq_\epsilon [V]_\epsilon$ iff $[X \setminus V]_\epsilon \leq_\epsilon [X \setminus U]_\epsilon$. Hence, if $\epsilon = C(\epsilon)$, then $\overline{\epsilon} = C(\overline{\epsilon})$.*

Proof. We leave it to the reader, it is routine. □

The connection of the closure to $\mathbf{A}$ is given by

Lemma 3.14. *Let Σ be a decent Mal'cev condition represented on $X = \{x_1, \dots, x_n\}$ which is realized in $\mathbf{A}$ by interpreting its operation symbol as the $\mathbf{A}$-term t. Then for any $U,V \in \mathcal{P}(X)$, if $(U,V) \in \overline{\epsilon(\Sigma)}$, then $\mathbf{A} \models t(x_1^U, \dots, x_n^U) \approx t(x_1^V, \dots, x_n^V)$.*

Proof. We claim that if $[U]_{\epsilon(\Sigma)} \leq_{\epsilon(\Sigma)} [V]_{\epsilon(\Sigma)}$ and $\mathbf{A} \models t(x_1^V, \dots, x_n^V) \approx x$, then $\mathbf{A} \models t(x_1^U, \dots, x_n^U) \approx x$. Since $[U]_{\epsilon(\Sigma)} \leq_{\epsilon(\Sigma)} [V]_{\epsilon(\Sigma)}$, there exist $U = W_0, W_1, W_2, \dots, W_{2k-1}, W_{2k} = V$ in $\mathcal{P}(X)$ such that for all even $0 \leq i < 2k$, $(W_i, W_{i+1}) \in \epsilon(\Sigma)$, while for all odd $0 \leq i < 2k$, $W_i \subseteq W_{i+1}$.

We prove by an induction on $k-i$ that $\mathbf{A} \models t(x_1^{W_{2(k-i)}}, \dots, x_n^{W_{2(k-i)}}) \approx x$. The base case when $i = 0$ is $\mathbf{A} \models t(x_1^V, \dots, x_n^V) \approx x$, which we assumed. Assuming the claim holds for i, and since $(W_{2(k-i)-1}, W_{2(k-i)}) \in \epsilon(\Sigma)$, we obtain $\mathbf{A} \models t(x_1^{W_{2(k-i)-1}}, \dots, x_n^{W_{2(k-i)-1}}) \approx t(x_1^{W_{2(k-i)}}, \dots, x_n^{W_{2(k-i)}}) \approx x$. Moreover, in the proof of Lemma 3.6, we proved that $[\emptyset]_\rho$ is a down-set, where $[\emptyset]_\rho = \{Z \subseteq X : \mathbf{A} \models t(x_1^Z, \dots, x_n^Z) \approx x\}$. As $W_{2(k-i)-1} \in [\emptyset]_\rho$ and $W_{2(k-1)-2} \subseteq W_{2(k-1)-1}$, thus $W_{2(k-i)-2} \in [\emptyset]_\rho$, which means that $\mathbf{A} \models t(x_1^{W_{2(k-i-1)}}, \dots, x_n^{W_{2(k-i-1)}}) \approx x$, completing the inductive proof. The case $i = k$ gives $\mathbf{A} \models t(x_1^U, \dots, x_n^U) \approx y$ finishing off the claim.

Now from $(U,V) \in \overline{\epsilon(\Sigma)}$ follows $[U]_{\epsilon(\Sigma)} \leq_{\epsilon(\Sigma)} [V]_{\epsilon(\Sigma)}$ and $[V]_{\epsilon(\Sigma)} \leq_{\epsilon(\Sigma)} [U]_{\epsilon(\Sigma)}$, so the claim we just proved implies $\mathbf{A} \models t(x_1^U, \dots, x_n^U) \approx x$ iff $\mathbf{A} \models t(x_1^V, \dots, x_n^V) \approx x$. Since for any $W \subseteq X$, either $\mathbf{A} \models t(x_1^W, \dots, x_n^W) \approx x$, or $\mathbf{A} \models t(x_1^W, \dots, x_n^W) \approx y$, what we proved establishes $\mathbf{A} \models t(x_1^U, \dots, x_n^U) \approx t(x_1^V, \dots, x_n^V)$. □

The following lemma is the goal of all the work in this subsection following Lemma 3.6. It is an easily verifiable criterion for realization of decent Mal'cev conditions in $\mathbf{A}$.

Lemma 3.15. *Let Σ be a decent Mal'cev condition represented on X. $\mathbf{A}$ realizes Σ iff $\overline{\epsilon(\Sigma)}$ is a $0-1$ distinguishing equivalence relation on $\mathcal{P}(X)$.*

Proof. Let $X = \{x_1, \ldots, x_n\}$ and denote $\epsilon := \epsilon(\Sigma)$ and $\varepsilon := \overline{\epsilon(\Sigma)}$.

Assume that ε is not a $0-1$ distinguishing equivalence relation on $\mathcal{P}(X)$. Then there exist $U, V, W, Z \in \mathcal{P}(X)$ such that $[U]_\varepsilon = [V]_\varepsilon = [W]_\varepsilon = [Z]_\varepsilon$, $U \cap V = \emptyset$ and $W \cup Z = X$. Let $t(x_1, \ldots, x_n)$ be any $\mathbf{A}$-term and suppose $\mathbf{A}$ realizes Σ by interpreting its operation as t. If $\mathbf{A} \models t(x_1^U, \ldots, x_n^U) \approx y$, from $U \cap V = \emptyset$ follows $V \subseteq (X \setminus U)$, hence $t(x_1^V, \ldots, x_n^V) \approx x$, contradicting $\mathbf{A} \models t(x_1^U, \ldots, x_n^U) \approx t(x_1^V, \ldots, x_n^V)$ which we get from $(U, V) \in \varepsilon$ and Lemma 3.14. On the other hand, if $\mathbf{A} \models t(x_1^U, \ldots, x_n^U) \approx x$, then $(U, W) \in \varepsilon$, $(U, Z) \in \varepsilon$ and Lemma 3.14 imply $\mathbf{A} \models t(x_1^W, \ldots, x_n^W) \approx t(x_1^Z, \ldots, x_n^Z) \approx x$. But $W \cup Z = X$ implies that $(X \setminus W) \subseteq Z$, so $\mathbf{A} \models t(x_1^{X\setminus W}, \ldots, x_n^{X\setminus W}) \approx x$, while $t(x_1^W, \ldots, x_n^W) \approx x$ implies $\mathbf{A} \models t(x_1^{X\setminus W}, \ldots, x_n^{X\setminus W}) \approx y$, a contradiction. Thus there can be no realization of Σ in $\mathbf{A}$.

Now assume that <u>**(1)** ε **is a 0 − 1 distinguishing equivalence relation.**</u> We will prove some other properties of ε, which will be bolded and underlined, since we will be proving that all those properties are carried over to the next, larger equivalence relation. If $[U]_\varepsilon \preceq_\varepsilon [V]_\varepsilon$, then there exist $U', V' \subseteq X$ such that $[U]_\varepsilon = [U']_\varepsilon$, $[V']_\varepsilon = [V]_\varepsilon$ and $U' \subseteq V'$. This means that $[U]_\epsilon \sim_\epsilon [U']_\epsilon$, so $[U]_\epsilon \leq_\epsilon [U']_\epsilon$, and similarly $[V']_\epsilon \leq_\epsilon [V]_\epsilon$, and taken together with $U' \subseteq V'$ we obtain $[U]_\epsilon \leq_\epsilon [V]_\epsilon$. Thus for the transitive closure $\leq_\varepsilon$ of the relation $\preceq_\varepsilon$ we must have that also $[U]_\varepsilon \leq_\varepsilon [V]_\varepsilon$ implies $[U]_\epsilon \leq_\epsilon [V]_\epsilon$. From this we conclude that, if $[U]_\varepsilon \leq_\varepsilon [V]_\varepsilon$ and $[V]_\varepsilon \leq_\varepsilon [U]_\varepsilon$, then $[U]_\epsilon \sim_\epsilon [V]_\epsilon$, so $[U]_\varepsilon = [V]_\varepsilon$. Since reflexivity and transitivity of $\leq_\varepsilon$ are given, we have proved that <u>**(2)** $(\mathcal{P}(X)/\varepsilon; \leq_\varepsilon)$ **is a partially ordered set.**</u> From the antisymmetry of $\leq_\varepsilon$ follows that $[U]_\varepsilon \sim_\varepsilon [V]_\varepsilon$ iff $[U]_\varepsilon \leq_\varepsilon [V]_\varepsilon$ and $[V]_\varepsilon \leq_\varepsilon [U]_\varepsilon$ iff $[U]_\varepsilon = [V]_\varepsilon$, so <u>**(3)** $\overline{\varepsilon} = \varepsilon$**.**</u> The least element of $\mathbb{P}$ is $[\emptyset]_\varepsilon$, while the greatest element of $\mathbb{P}$ is $[X]_\varepsilon$. From Definition 3.4 follows that $\epsilon = C(\epsilon)$, so Lemma 3.12 implies that <u>**(4)** $\varepsilon = \mathbf{C}(\varepsilon)$**.**</u> Applying Lemma 3.13 to ε, we prove that <u>**(5) for all** $U, V \subseteq X$**,** $[U]_\varepsilon \leq_\varepsilon [V]_\varepsilon$ **iff** $[X \setminus V]_\varepsilon \leq_\varepsilon [X \setminus U]_\varepsilon$**.**</u>

Since ε is $0-1$ distinguishing, $[\emptyset]_\varepsilon \neq [X]_\varepsilon$. Assume that there are more ε-classes than just these two. Select $U_0 \subseteq X$ such that $[U_0]_\varepsilon$ is minimal in the poset $(\mathcal{P}(X)/\varepsilon \setminus \{[\emptyset]_\varepsilon\}; \leq_\varepsilon)$. Just using the properties $(1)-(5)$ above we prove that the equivalence relation ε' obtained from ε by merging the class $[\emptyset]_\varepsilon$ with $[U_0]_\varepsilon$ and also merging $[X]_\varepsilon$ with $[X \setminus U_0]_\varepsilon$, but keeping all other classes the same, also satisfies $(1)-(5)$.

(**4**) $\varepsilon' = C(\varepsilon')$ follows from $\varepsilon = C(\varepsilon)$ and from $\{X \setminus U : U \in [\emptyset]_{\varepsilon'}\} = \{X \setminus U : U \in [\emptyset]_\varepsilon \cup [U_0]_\varepsilon\} = [X]_\varepsilon \cup [X \setminus U_0]_\varepsilon = [X]_{\varepsilon'}$. (**5**) follows since ε' satisfies (**4**) by applying Lemma 3.13 to ε'. To prove that ε' is 0-1 distinguishing, our property (**1**), we need only check the conditions for the classes $[\emptyset]_{\varepsilon'}$ and $[X]_{\varepsilon'}$, as the rest are ε-classes. Since $\varepsilon' = C(\varepsilon')$, it suffices to check only $[\emptyset]_{\varepsilon'}$. Assume that there exist $W, Z \in [\emptyset]_{\varepsilon'}$ such that $W \cup Z = X$. Since ε is 0-1 distinguishing, at least one of them must be in $[U_0]_\varepsilon$, say $W \in [U_0]_\varepsilon$. From $W \cup Z = X$ follows that $(X \setminus W) \subseteq Z$, so $[X \setminus W]_\varepsilon \leq_\varepsilon [Z]_\varepsilon$. By the minimality of $[U_0]_\varepsilon$ it follows that either $[X \setminus W]_\varepsilon = [U_0]_\varepsilon = [W]_\varepsilon$, contradicting that ε is $0-1$ distinguishing, or $[X \setminus W]_\varepsilon = [\emptyset]_\varepsilon$, implying that $[W]_\varepsilon = [X \setminus \emptyset]_\varepsilon$, i.e. $[U_0]_\varepsilon = [X]_\varepsilon$. Since $(\mathcal{P}(X)/\varepsilon; \leq_\varepsilon)$ is a partially ordered set with its greatest element $[X]_\varepsilon$, and by the minimality of $[U_0]_\varepsilon = [X]_\varepsilon$, this means that $\mathcal{P}(X)/\varepsilon = \{[\emptyset]_\varepsilon, [X]_\varepsilon\}$, which we assumed was not the case. To prove (**2**), we need to prove the antisymmetry of $\leq_{\varepsilon'}$. Assume that $[U]_{\varepsilon'} \leq_{\varepsilon'} [V]_{\varepsilon'}$ and $[V]_{\varepsilon'} \leq_{\varepsilon'} [U]_{\varepsilon'}$. This means that there exist $W_i, Z_j \subseteq X$ such that $U =: W_0\ \varepsilon'\ W_1 \subseteq W_2\ \varepsilon'\ W_3 \subseteq \cdots \subseteq W_{2k-1}\ \varepsilon'\ W_{2k} := V$ and $V =: Z_0\ \varepsilon'\ Z_1 \subseteq Z_2\ \varepsilon'\ Z_3 \subseteq \cdots \subseteq Z_{2l-1}\ \varepsilon'\ Z_{2l} := U$. If none of the $[W_i]_{\varepsilon'}$ or $[Z_j]_{\varepsilon'}$ is either $[\emptyset]_{\varepsilon'}$, or $[X]_{\varepsilon'}$, the conclusion follows by the antisymmetry of $\leq_\varepsilon$. Without loss of generality, assume that $[W_{2i}]_{\varepsilon'} = [\emptyset]_{\varepsilon'}$. Inductively, we prove that For all $j < 2i$, $[W_j]_{\varepsilon'} = [\emptyset]_{\varepsilon'}$. $W_{2i-1} \subseteq W_{2i}$ implies that $[W_{2i-1}]_\varepsilon \leq_\varepsilon [W_{2i}]_\varepsilon \in \{[\emptyset]_\varepsilon, [U_0]_\varepsilon\}$. The minimality of $[U_0]_\varepsilon$ implies that $[W_{2i-1}]_\varepsilon \in \{[\emptyset]_\varepsilon, [U_0]_\varepsilon\}$, hence $[W_{2i-2}]'_\varepsilon = [W_{2i-1}]_{\varepsilon'} = [\emptyset]_{\varepsilon'}$. Proceeding inductively, we conclude that $[U]_{\varepsilon'} = [\emptyset]_{\varepsilon'}$. But then $[Z_{2l}]_{\varepsilon'} = [\emptyset]_{\varepsilon'}$, and the same inductive argument applied to Z_j proves $[V]_{\varepsilon'} = [Z_0]_{\varepsilon'} = [\emptyset]_{\varepsilon'}$. So $[U]_{\varepsilon'} = [V]_{\varepsilon'}$, as desired. The antisymmetry of $\leq_{\varepsilon'}$ proves $\overline{\varepsilon'} = \varepsilon'$, property (**3**), just like in the case of ε.

Since ε' is $0-1$ distinguishing, $[\emptyset]_{\varepsilon'} \neq [X]_{\varepsilon'}$, and we can keep merging classes like this until there are only those two. Let the relation obtained at the end of this process be ρ. Since our process started from $\epsilon = \epsilon(\Sigma)$ and kept increasing the equivalence relation, we have $\epsilon \subseteq \rho$. Moreover, ρ has exactly two equivalence classes $[\emptyset]_\rho$ and $[X]_\rho$, and since ρ satisfies property (**4**) $\rho = C(\rho)$, hence $[\emptyset]_\rho = \{X \setminus U : U \in [X]_\rho\}$. If $U \in [\emptyset]_\rho$ and $V \subseteq U$, then $[\emptyset]_\rho \leq [V]_\rho \leq [U]_\rho$, so by property (2) of $\leq_\rho$ we obtain that $[V]_\rho = [U]_\rho$. Hence, $[\emptyset]_\rho$ is a down-set with respect to inclusion, and since $[X]_\rho = \{X \setminus U : U \in [\emptyset]_\rho\}$, it follows that $[X]_\rho$ is an up-set. All conditions of Lemma 3.6 are fulfilled, so we conclude that $\mathbf{A} \models \Sigma$. □

3.2 Decent Mal'cev conditions in semilattices

Let $\mathbf{B} = (\{0,1\}, \wedge)$ be the two-element semilattice. Again, we discuss decent Mal'cev conditions, but those realized in the algebra $\mathbf{C} = (\{0,1\}, s)$, where $s(x,y,z) = x \wedge y \wedge z$ for all $x,y,z \in \{0,1\}$, and $\wedge$ is the semilattice operation from $\mathbf{B}$. $\mathbf{B}$ and $\mathbf{C}$ are clearly term equivalent since also $x \wedge y = s(x,x,y)$, and the following is well known about $\mathbf{B}$:

Proposition 3.16. *For every term $t(x_1,\dots,x_n)$ on the language $\{s\}$ there exists nonempty $\{i_1,\dots,i_k\} \subseteq \{1,2,\dots,n\}$ such that*

$$\mathbf{C} \models t(x_1,\dots,x_n) \approx x_{i_1} \wedge \cdots \wedge x_{i_k}.$$

$\{x_{i_1},\dots,x_{i_k}\}$ is the set of all variables which actually occur in the syntactic expression of the term t.

Incidentally, the positive integer k from Proposition 3.16 is the *essential arity* of the term operation $t^{\mathbf{C}}$ (we omit the definition since we don't use the concept). Proposition 3.16 implies

Proposition 3.17. *For every term $t(x_1,\dots,x_n)$ on the language $\{s\}$ and $U \subseteq \{x_1,\dots,x_n\}$:*

$$\begin{aligned}
&\mathbf{C} \models t(x_1^U,\dots,x_n^U) \approx x, \text{ when } U \cap E = \emptyset\\
&\mathbf{C} \models t(x_1^U,\dots,x_n^U) \approx y, \text{ when } E \subseteq U, \text{ or}\\
&\mathbf{C} \models t(x_1^U,\dots,x_n^U) \approx x \wedge y, \text{ when } E \cap U \neq \emptyset \neq E \setminus U.
\end{aligned}$$

Here $\emptyset \neq E \subseteq \{x_1,\dots,x_n\}$ is the unique set such that $\mathbf{C} \models t \approx \bigwedge_{x_i \in E} x_i$

From this follows a corollary similar to Lemma 3.6, but for the algebra $\mathbf{C}$:

Corollary 3.18. *Let Σ be a decent Mal'cev condition represented on X. The following conditions are equivalent:*

(1) $\mathbf{C}$ *realizes Σ.*

(2) *There exists a subset $\emptyset \neq E \subseteq X$ such that interpreting f as $\bigwedge_{x_i \in E} x_i$ induces a realization of Σ in $\mathbf{C}$.*

(3) *There exists a subset $\emptyset \neq E \subseteq X$ such that $\mathbf{C}$ realizes a decent Mal'cev condition Π represented on E and such that:*

(a) $\{(U \cap E, V \cap E) : (U, V) \in \epsilon(\Sigma)\} \subseteq \epsilon(\Pi)$ *and*

(b) $\rho = \epsilon(\Pi)$ *has exactly three equivalence classes and two of them are* $[\emptyset]_\rho = \{\emptyset\}$ *and* $[E]_\rho = \{E\}$*, or* $|E| = 1$ *and* ρ *has two equivalence classes,* $\{\emptyset\}$ *and* $\{E\}$.

Proof. (1) $\Leftrightarrow$ (2) follows from Proposition 3.16.

(2) $\Rightarrow$ (3) Let $E = \{x_{i_1}, \ldots, x_{i_k}\}$. (2) implies that $\mathbf{A}$ realizes a decent Mal'cev condition Σ' obtained from Σ by considering all variables in $X \setminus E$ as dummy variables. Lemma 3.5 implies that (a) holds for $\epsilon(\Sigma')$ in place of ρ. By Proposition 3.17, $\mathbf{C} \models x_{i_1}^U \wedge \cdots \wedge x_{i_k}^U \approx y$ if and only if $E \subseteq U$, and $\mathbf{C} \models x_{i_1}^U \wedge \cdots \wedge x_{i_k}^U \approx x$ if and only if $U \cap E = \emptyset$, so $[\emptyset]_{\epsilon(\Sigma')} = \{\emptyset\}$ and $[E]_{\epsilon(\Sigma')} = \{E\}$. Since for all $U, V \in \mathcal{P}(E) \setminus \{\emptyset, E\}$, $\mathbf{C} \models x_1^U \wedge \cdots \wedge x_n^U \approx x_1^V \wedge \cdots \wedge x_n^V \approx x \wedge y$, thus $\mathbf{C}$ realizes the decent Mal'cev condition Π such that (1) and (2) hold (Π is obtained from Σ' by uniting all $\epsilon(\Sigma')$-classes into one, except $\{\emptyset\}$ and $\{E\}$).

(3) $\Rightarrow$ (1) Let t be the $\mathbf{C}$-term which interprets g in the realization of Π in $\mathbf{C}$. From Proposition 3.16, there exists a set of variables $E' \subseteq E$ such that $\mathbf{C} \models t \approx \bigwedge_{x_i \in E'} x_i$. The condition $[E]_\rho = \{E\}$ implies that $E' = E$. For the decent Mal'cev condition Σ', obtained from Σ by considering all variables outside E as dummy variables, from Lemma 3.5 and (a) follows that $\epsilon(\Sigma') \subseteq \rho = \epsilon(\Pi)$. Therefore, interpreting f' as the same term t realizes Σ', and by Lemma 3.5, (1) also holds. $\square$

So which E to take in Corollary 3.18? It is easy to come up with the conditions which EXCLUDE variables from E, i.e. for $U \subseteq X$ to be disjoint with E. We denote $\epsilon := \epsilon(\Sigma)$. Clearly, if a set $(U, V) \in \epsilon$ and $V \cap E = \emptyset$, then $U \cap E = \emptyset$, i.e. if V consists entirely of dummy variables and $U \,\epsilon\, V$ then U consists entirely of dummy variables. Another criterion is if $U \subseteq V_1 \cup \cdots \cup V_k$ and for each V_i, $V_i \cap E = \emptyset$, then $U \cap E = \emptyset$. So, we define $D(\Sigma)$ (D stands for "dummy") to be the least subset of X such that:

i) $D(\Sigma){\downarrow}$ contains $[\emptyset]_\epsilon$ and

ii) $D(\Sigma){\downarrow}$ is a union of ϵ-classes.

If E satisfies Corollary 3.18 (3), then conditions (3) (a) and (3) (b) imply that $X \setminus E$ satisfies *i)* and *ii)*. Moreover, X satisfies *i)* and *ii)*, and if D_1 and D_2 both satisfy *i)* and *ii)*, then so does $D_1 \cap D_2$, so $D(\Sigma)$ is well-defined. Finally, both criteria we came up with for disjointness with E are met if $V \subseteq D(\Sigma)$.

The next lemma shows that Corollary 3.18 works with $E = X \setminus D(\Sigma)$.

Lemma 3.19. *If* $\mathbf{C}$ *realizes a decent Mal'cev condition* Σ *represented on* $X = \{x_1, \ldots, x_n\}$*, then* $\mathbf{C}$ *realizes* Σ *by interpreting* f *as* $\bigwedge\limits_{x_i \in X \setminus D(\Sigma)} x_i$.

Proof. Denote $E := X \setminus D(\Sigma)$ and let $(U, V) \in \epsilon(\Sigma)$. Firstly, $U \cap E = \emptyset$ iff $U \subseteq D(\Sigma)$ iff $V \subseteq D(\Sigma)$ (since $D(\Sigma)$ is a union of ϵ-classes) iff $V \cap E = \emptyset$. Moreover, from $(U, V) \in \epsilon(\Sigma)$ follows that $(X \setminus U, X \setminus V) \in \epsilon(\Sigma)$, so $U \cap E = E$ iff $X \setminus U \subseteq D(\Sigma)$ iff $X \setminus V \subseteq D(\Sigma)$ iff $V \cap E = E$. So, the equivalence relation ρ on E whose only classes are $\{\emptyset\}$, $\{E\}$ and $\mathcal{P}(E) \setminus \{\emptyset, E\}$ contains the relation $\{(U \cap E, V \cap E) : (U, V) \in \epsilon(\Sigma)\}$. According to Corollary 3.18 (3) $\Rightarrow$ (2), this means that $\mathbf{C}$ realizes Σ by interpreting f as $\bigwedge\limits_{x_i \in E} x_i$. $\square$

When Σ is a decent Mal'cev condition represented on X, we denote $E(\Sigma) := X \setminus D(\Sigma)$.

3.3 Decent Mal'cev conditions in the algebra D

Let us denote $\mathbf{D} = \mathbf{A} \times \mathbf{C}$. The next lemma is significant for our purpose.

Lemma 3.20. *If* $\mathbf{D}$ *realizes a decent Mal'cev condition* Σ *represented on* $X = \{x_1, \ldots, x_n\}$*, then for* $E = E(\Sigma)$*, there exists a decent Mal'cev condition* Π *represented on* E *such that, where* $\rho = \epsilon(\Pi)$*:*

(1) $\{(U \cap E, V \cap E) : (U, V) \in \epsilon(\Sigma)\} \subseteq \rho$*;*

(2) ρ *has exactly four equivalence classes,* I*,* J*,* $\{\emptyset\}$ *and* $\{E\}$*, where* $I = \{E \setminus U : U \in J\}$ *(hence* $|I| = |J|$*), or* $|E| = 1$*, and the* ρ*-classes are* $\{\emptyset\}$ *and* $\{E\}$*;*

(3) $I \cup \{\emptyset\}$ *is a down-set in* $\mathcal{P}(E)$ *and* $J \cup \{E\}$ *is up-set in* $\mathcal{P}(E)$*;*

Proof. Let Σ be realized in $\mathbf{D}$ by interpreting f as some $\mathbf{D}$-term t' so that the set of all variables which occur in t' is E'. Then both $\mathbf{A}$ and $\mathbf{C}$ realizes Σ by interpreting f as t' so, according to remarks preceding Lemma 3.19, $E' \subseteq E$. Using the term which adds dummy variables in $\mathbf{A}$ as in Remark 3.7, we can construct the term t such that $\mathbf{A} \models t \approx t'$ and the set of variables which occur in t is E. So $\mathbf{A}$ realizes Σ by interpreting f as t. Moreover, $\mathbf{C} \models t \approx \bigwedge\limits_{x_i \in E} x_i$, so according to Lemma 3.19, $\mathbf{C}$ also realizes Σ by interpreting f as t. Therefore, $\mathbf{D} = \mathbf{A} \times \mathbf{C}$ realizes Σ by interpreting f as t.

From Lemma 3.6 follows that there exists a decent Mal'cev condition Π_1 represented on E which is realized in $\mathbf{A}$ by interpreting its operation as t and such that $\rho_1 := \epsilon(\Pi_1)$ satisfies conditions $(2) - (3)$ with ρ_1 in the place of ρ of Lemma 3.6 and also, by combining Lemma 3.6 and Lemma 3.5

(1′) $\{(U \cap E, V \cap E) : (U, V) \in \epsilon(\Sigma)\} \subseteq \rho_1$.

Since the set of variables which occur in the syntactical expression of t is E, thus $\mathbf{C} \models t \approx \bigwedge_{x_i \in E} x_i$. Therefore, Lemma 3.19 and Corollary 3.18 provide that there exists a decent Mal'cev condition Π_2 represented on E which is realized in $\mathbf{C}$ by interpreting its operation as t and conditions (a) and (b) stated in Corollary 3.18 are satisfied (by replacing ρ in their statements with $\rho_2 := \epsilon(\Pi_2)$).

Since the term t satisfies in $\mathbf{A}$ the identities implied by ρ_1 and in $\mathbf{C}$ the identities implied by ρ_2, then in $\mathbf{D}$ this term satisfies all identities implied by $\rho_1 \cap \rho_2$. Now denote $I := [\emptyset]_{\rho_1} \setminus \{\emptyset\}$ and $J := [E]_{\rho_1} \setminus \{E\}$ and the statement of the Lemma follows. □

Definition 3.21. We say that a decent Mal'cev condition Σ represented on $X = \{x_1, \ldots, x_n\}$, $n > 1$, such that:

(1) $\epsilon(\Sigma)$ has exactly four equivalence classes, I_Σ, J_Σ, $\{\emptyset\}$ and $\{X\}$, where $I_\Sigma = \{X \setminus U : U \in J_\Sigma\}$ (hence $|I_\Sigma| = |J_\Sigma|$),

(2) $I_\Sigma \cup \{\emptyset\}$ is a down-set in $\mathcal{P}(X)$ and $J_\Sigma \cup \{X\}$ is an up-set in $\mathcal{P}(X)$ and

(3) $\epsilon(\Sigma)$ is 0-1 distinguishing on $\mathcal{P}(X)$

is a *canonical* decent Mal'cev condition.

Remark 3.22. Note that each canonical decent Mal'cev condition Σ is syntactically equivalent to $\Sigma_1 = \{f(x_1^U, \ldots, x_n^U) \approx f(x_1^V, \ldots, x_n^V) : U, V \in I_\Sigma\}$ (and idempotence), also to $\Sigma_2 = \{f(x_1^U, \ldots, x_n^U) \approx f(x_1^V, \ldots, x_n^V) : U, V \in J_\Sigma\}$ (and idempotence). This follows from $I_\Sigma = \{(X \setminus U, X \setminus V) : (U, V) \in J_\Sigma\}$, so the identities in Σ_1 are equivalent to the identities in Σ_2 via interchanging the variables x and y. J_Σ is a maximal family of subsets of X such that any pair of sets in J_Σ has a nonempty intersection (recall again the intersecting families of [4]). For us, however, it will be more expedient to assume that each canonical decent Mal'cev condition is of the form Σ_1, so we make this assumption now.

Lemma 3.20 may be restated as

Corollary 3.23. *If* $\mathbf{D}$ *realizes a decent Mal'cev condition* Σ *represented on* $X = \{x_1, \ldots, x_n\}$ *then for* $E = E(\Sigma)$*, either* $|E| = 1$*, or there exists a canonical decent Mal'cev condition* Π *represented on* E *such that* $\{(U \cap E, V \cap E) : (U, V) \in \epsilon(\Sigma)\} \subseteq \epsilon(\Pi)$.

Also, for the purposes of the next subsection, we state the following

Corollary 3.24. *Let Σ be a decent Mal'cev condition represented on X. Denote by ϵ' the equivalence relation on $\mathcal{P}(E(\Sigma))$ generated by $\{(U \cap E(\Sigma), V \cap E(\Sigma)) : (U,V) \in \epsilon(\Sigma)\}$. Then $\mathbf{D}$ realizes Σ iff $\overline{\epsilon'}$ is a $0-1$ distinguishing equivalence relation on $\mathcal{P}(E(\Sigma))$.*

Proof. According to the proof of Lemma 3.20, if $\mathbf{D}$ realizes Σ, then it realizes Σ by interpreting its operation as some term t such that the set of variables which occurs in t is $E(\Sigma)$. According to Lemma 3.5, this means that $\mathbf{D}$ realizes a decent Mal'cev condition Σ' represented on $E(\Sigma)$ such that $\epsilon(\Sigma') = \epsilon'$. Now Lemma 3.15 and the fact that $\mathbf{A}$ realizes Σ' imply that $\overline{\epsilon'}$ is $0-1$ distinguishing on $\mathcal{P}(E(\Sigma))$.

On the other hand, if $\overline{\epsilon'}$ is $0-1$ distinguishing on $\mathcal{P}(E(\Sigma))$, Lemma 3.15 implies that the decent Mal'cev condition Σ' represented on $E(\Sigma)$ which is defined by ϵ' is realized in $\mathbf{A}$. According to Remark 3.7, there exists an interpretation of the operation symbol as an $\mathbf{A}$-term t which uses all variables in $E(\Sigma)$ and realizes Σ'. Lemma 3.5 and the definition of ϵ' imply that interpreting its operation as t realizes Σ in $\mathbf{A}$. Also, since the set of all variables which occur in t is $E(\Sigma)$, Proposition 3.17 and Lemma 3.19 imply that interpreting its operation as t realizes Σ in $\mathbf{C}$, too. Hence $\mathbf{D} = \mathbf{A} \times \mathbf{C}$ realizes Σ. $\square$

3.4 An algorithm for checking whether a decent Mal'cev condition is realized in D

The work done so far in this Section allows one to efficiently algorithmically determine whether $\mathbf{D}$ realizes a decent Mal'cev condition Σ. Now we describe this algorithm.

Procedure 1. Encoding $\epsilon(\Sigma)$. Let Σ be a decent Mal'cev condition on the set of variables X. First we linearly order the elements of X. This allows us to encode all subsets $U \subseteq X$ as elements of $\{0,1\}^{|X|}$, i.e. U becomes a word of length $|X|$ with 1 as its ith letter iff the ith element of X is in U. We order all subsets in the following way: $U \leq V$ iff either $|U| < |V|$, or $|U| = |V|$ and the word encoding U is lexicographically before than or equal to the word encoding V. Note that $U \leq V$ iff $(X \setminus V) \leq (X \setminus U)$ and also that $\leq$ contains (extends) the inclusion order. Now we define a directed graph $\Gamma(\Sigma)$ by its edge relation $U \to V$ iff $U < V$, $(U,V) \in \epsilon(\Sigma)$ and V is the greatest element with respect to $\leq$ in its $\epsilon(\Sigma)$-class. Thus $\epsilon(\Sigma)$ is the weak connectedness relation of the digraph $\Gamma(\Sigma)$, so all we need to do is effectively and efficiently construct $\Gamma(\Sigma)$ from Σ.

First, for each identity $t(x_1^U,\dots,x_n^U)\approx t(x_1^V,\dots,x_n^V)\in\Sigma$ we add to Γ the edges $U\to V$ and $(X\setminus V)\to(X\setminus U)$, if $U<V$, or the two edges going the other way if $V<U$. Now Γ consists of at most $2|\Sigma|$ edges and its weak connectedness relation is equal to $\epsilon(\Sigma)$.

Next, whenever $U\to V$ and $U\to W$ in Γ and $V<W$, we delete the edge $U\to V$ and replace it with $V\to W$. When this process terminates, there are still $2|\Sigma|$ edges, the weak connectedness of Γ is unchanged, and the out-degree of each vertex is at most 1. The process will terminate, and in a quadratic number of steps, provided that we always use this step with the least set U for which such V and W exist. This way the out-degree of U reduces until it equals 1, and it will never change in the remainder of the process, so we proceed to the next least such U until we can find no more.

Finally, whenever $U\to V\to W$ in Γ, we delete the edge $U\to V$ and replace it with $U\to W$. This step can be applied to the edge coming out of U and the edge coming out of V at most once, after it is performed, there is still one edge coming out of U and one coming out of V, so the number of such steps is at most quadratic in $|\Gamma|$. When the process terminates, the number of edges and the weak connectedness are unchanged, the out-degree of each vertex is still at most 1, but now the edge going from any U points to the largest set in its $\epsilon(\Sigma)$-class. This way we constructed $\Gamma(\Sigma)$.

Procedure 2. Finding $D(\Sigma)$. Given Γ, we are implicitly given a closure operation on $\mathcal{P}(X)$: $U^*=V$ iff V is the greatest element (with respect to $\leq$) in $[U]_{\epsilon(\Sigma)}$ iff either $U\to V$ or the out-degree of U is 0 and $V=U$.

If $\emptyset^*=\emptyset$, then $D(\Sigma)=\emptyset$ and we can stop. Otherwise, we start with $\Delta:=\Gamma(\Sigma)$ and expand the class $[\emptyset]_\Delta$ using two kinds of steps until we reach $D(\Sigma)$. We fix $V:=\emptyset^*$, i.e. $\emptyset\to V$ in Δ. When Δ is a digraph on $\mathcal{P}(X)$, by $[U]_\Delta$ we denote the weakly connected component which contains U.

In the first kind of step, we close $[\emptyset]_\Delta$ under unions. First we check whether $\bigcup\{U:U\to V\}\subseteq V$. This requires checking if $U\subseteq V$ whenever $U\to V$, so the number of checks equals the number of edges in Δ. If $\bigcup\{U:U\to V\}\not\subseteq V$, we define $W:=V\cup\bigcup\{U:U\to V\}$, the new Δ is obtained by replacing all edges $U\to V$ with $U\to W^*$ and adding the edge $V\to W^*$. Note that Δ is now increased by one edge, $[\emptyset]_\Delta$ and $[W]_\Delta$ are now merged. The second class may be singleton, so we can't bound the number of times this step is performed with the number of non-singleton $\Gamma(\Sigma)$-classes, but since $|W^*|\geq|W|>|V|$, hence this kind of step can be performed at most $|X|$ times. Finally define V to be W^* and return to checking whether $V=\bigcup\{U:U\to V\}$. When the first kind of step can no longer be performed, then $V=\emptyset^*$ satisfies that if $U\to V$, then $U\subseteq V$. Moreover, the number of non-singleton Δ-classes does not increase.

We perform the second kind of step when the first kind of step can no longer be performed. This step starts with searching through all edges $U \to U^*$ in Δ such that $U^* \neq V$. If we find such an edge which satisfies $U \subseteq V$ or $U^* \subseteq V$, we merge $[\emptyset]_\Delta$ and $[U]_\Delta$. In particular, if $\max\{U^*, V\} = V$, then in the new graph of Δ we replace all edges $W \to U^*$ with $W \to V$ and add the edge $U^* \to V$. On the other hand, if $\max\{U^*, V\} = U^*$, then we replace all edges $W \to V$ with $W \to U^*$, add the edge $V \to U^*$ and set the new V to be U^*. By analyzing all the substeps of the second kind of step we conclude that the second kind of step can be performed in polynomial time in $|\Sigma|$. Moreover, it reduces the number of non-singleton Δ-classes by one, and since $\Gamma(\Sigma)$ had at most $2|\Sigma|$ many such classes, thus we have a linear bound on the number of times the second kind of step can be performed. Once we have merged $[\emptyset]_\Delta$ and $[U]_\Delta$, we return to checking whether the first kind of step can be performed.

In the end we reached, in polynomial time in $|\Sigma|$ and $|X|$, a graph Δ which satisfies the following: $[\emptyset]_\Delta$ is a union of weakly connected components of $\Gamma(\Sigma) = \epsilon(\Sigma)$-classes, $\emptyset^* = \bigcup[\emptyset]_\Delta$ and whenever $U \to V$ in Δ and either $U \subseteq \emptyset^*$ or $V \subseteq \emptyset^*$, then $[U]_\Delta = [\emptyset]_\Delta$. So $V := \emptyset^*$ satisfies that $V\downarrow$ (with respect to inclusion) is a union of $\epsilon(\Sigma)$-classes and contains $[\emptyset]_{\epsilon(\Sigma)}$. Hence $D(\Sigma) \subseteq V$.

On the other hand, by an induction on the number of steps we prove that $\bigcup[\emptyset]_\Delta \subseteq D(\Sigma)$. The base is clear since $[\emptyset]_{\epsilon(\Sigma)} \subseteq D(\Sigma)\downarrow$. Assume that Δ is obtained from the previous Δ' by merging $[\emptyset]_{\Delta'}$ with $[U]_{\Delta'} = [U]_{\epsilon(\Sigma)}$. If Δ is obtained from Δ' by performing a step of the first kind, we merged $[\emptyset]_{\Delta'}$ with $[U]_{\Delta'} = [U]_{\epsilon(\Sigma)}$, where $U = \bigcup[\emptyset]_{\Delta'}$. We know that $[\emptyset]_{\Delta'} \subseteq D(\Sigma)\downarrow$ and also $U = \bigcup[\emptyset]_{\Delta'} \subseteq D(\Sigma)$, by the inductive assumption. Since $D(\Sigma)\downarrow$ is a union of $\epsilon(\Sigma)$-classes and $[U]_{\epsilon(\Sigma)}$ is an $\epsilon(\Sigma)$-class which intersects $D(\Sigma)\downarrow$, we conclude that $[U]_{\Delta'} \subseteq D(\Sigma)\downarrow$. Hence $[\emptyset]_\Delta \subseteq D(\Sigma)\downarrow$, i.e. $\bigcup[\emptyset]_\Delta \subseteq D(\Sigma)$. If Δ is obtained from Δ' by performing a step of the second kind, we merged $[\emptyset]_{\Delta'}$ with $[U]_{\Delta'} = [U]_{\epsilon(\Sigma)}$, where $U \subseteq \bigcup[\emptyset]_{\Delta'} \in [\emptyset]_{\Delta'}$. By the inductive assumption, $\bigcup[\emptyset]_\Delta \subseteq D(\Sigma)$, so $U \subseteq D(\Sigma)$. Since $D(\Sigma)\downarrow$ is a union of $\epsilon(\Sigma)$-classes, this implies $[U]_{\Delta'} = [U]_{\epsilon(\Sigma)} \subseteq D(\Sigma)\downarrow$, and as we know $[\emptyset]_{\Delta'} \subseteq D(\Sigma)\downarrow$, we obtain again $[\emptyset]_\Delta \subseteq D(\Sigma)\downarrow$, i.e. $\bigcup[\emptyset]_\Delta \subseteq D(\Sigma)$.

Procedure 3. Restricting to $E(\Sigma)$. We use the just computed $D(\Sigma)$ to get $E(\Sigma) = X \setminus D(\Sigma)$. Then we replace $\Gamma(\Sigma)$ with the graph Γ' on $\mathcal{P}(E(\Sigma))$ so that $U \to V$ in Γ' iff there exist $U' \to V'$ in $\Gamma(\Sigma)$ such that $U = U' \cap E(\Sigma)$ and $V = V' \cap E(\Sigma)$. Γ' has at most the number of edges of $\Gamma(\Sigma)$, so it can be constructed in linear time. Since $(X \setminus U) \cap E(\Sigma) = E(\Sigma) \setminus (E(\Sigma) \cap U)$, we conclude that $U \to V$ in Γ' iff $(E(\Sigma) \setminus V) \to (E(\Sigma) \setminus U)$ in Γ'. Also, the weak connectedness relation of Γ' is equal to the equivalence relation generated by $\{(U \cap E(\Sigma), V \cap E(\Sigma)) : (U, V) \in \epsilon\}$. However, the graph Γ' might have $U \to V \to W$ or $U \to V$ and $U \to W$ such that $V \neq W$. So we conclude the procedure by performing Procedure 1 on Γ'.

After Procedure 3 is completed, we have computed the graph $\Gamma' = \Gamma(\epsilon')$, where ϵ' is the equivalence relation on $\mathcal{P}(E(\Sigma))$ generated by $\{(U \cap E(\Sigma), V \cap E(\Sigma)) : (U, V) \in \epsilon(\Sigma)\}$. According to Corollary 3.24, to decide whether $\mathbf{D}$ realizes Σ we need to determine whether $\overline{\epsilon'}$ is $0-1$ distinguishing.

Procedure 4. Computing $\overline{\epsilon'}$. The problem facing us with $\overline{\epsilon'}$ is that its classes are probably too large, i.e. exponential in $|\Sigma|$. However, we know that the $\overline{\epsilon'}$-classes are convex with respect to inclusion, so we will compute an equivalence relation ϵ'' on $\mathcal{P}(E(\Sigma))$ such that the least equivalence relation on $\mathcal{P}(E(\Sigma))$ which contains ϵ'' and whose classes are convex with respect to inclusion is $\overline{\epsilon'}$. In fact, ϵ''-classes will consist of the union of all ϵ'-classes which are not singleton and which lie in the same $\overline{\epsilon'}$-class.

We first prove that if $[U]_{\epsilon'}$ and $[V]_{\epsilon'}$ are two ϵ'-classes, then $[U]_{\epsilon'} \leq_{\epsilon'} [V]_{\epsilon'}$ iff there exist some finite sequence $[W_1]_{\epsilon'}, \ldots, [W_{k-1}]_{\epsilon'}$ of non-singleton ϵ'-classes such that $[U]_{\epsilon'} \preceq_{\epsilon'} [W_1]_{\epsilon'}$, $[W_{k-1}]_{\epsilon'} \preceq_{\epsilon'} [V]_{\epsilon'}$ and for all $1 \leq i < k-1$, $[W_i]_{\epsilon'} \preceq_{\epsilon'} [W_{i+1}]_{\epsilon'}$. The definition of $\leq_{\epsilon'}$ is the same, just it allows some $[W_i]_{\epsilon'}$ to be singleton classes. However, if $[W_i]_{\epsilon'} = \{W_i\}$, then from $[W_{i-1}]_{\epsilon'} \preceq_{\epsilon'} [W_i]_{\epsilon'} \preceq_{\epsilon'} [W_{i+1}]_{\epsilon'}$ (one of the outside classes among these three might be $[U]_{\epsilon'}$ or $[V]_{\epsilon'}$) implies that there exist Z_{i-1} and Z_{i+1} in $\mathcal{P}(E(\Sigma))$ such that $[W_{i-1}]_{\epsilon'} = [Z_{i-1}]_{\epsilon'}$, $[W_{i+1}]_{\epsilon'} = [Z_{i+1}]_{\epsilon'}$ and $Z_{i-1} \subseteq W_i \subseteq Z_{i+1}$. But then $[W_{i-1}]_{\epsilon'} \preceq_{\epsilon'} [W_{i+1}]_{\epsilon'}$ and $[W_i]_{\epsilon'}$ can be omitted from the sequence which witnesses that $[U]_{\epsilon'} \leq_{\epsilon'} [V]_{\epsilon'}$. Thus when $[U]_{\epsilon'}$ and $[V]_{\epsilon'}$ are two ϵ'-classes, $[U]_{\epsilon'} \leq_{\epsilon'} [V]_{\epsilon'}$ and consequently also $[U]_{\epsilon'} \sim_{\epsilon'} [V]_{\epsilon'}$ can be computed using only non-singleton ϵ'-classes. Note that we allowed $[U]_{\epsilon'}$ and $[V]_{\epsilon'}$ to be singleton.

Next, note that when $[U]_{\epsilon'}$ is singleton and $[U]_{\overline{\epsilon'}}$ is not a singleton, then there exist $V, W \in [U]_{\overline{\epsilon'}}$ such that $V \subseteq U \subseteq W$ and that $[V]_{\epsilon'}$ and $[W]_{\epsilon'}$ are non-singleton classes. To see that, note that $[U]_{\epsilon'}$ must be $\sim_{\epsilon'}$-related to some other class $[U']_{\epsilon'}$, so $U \subseteq W_1 \ \epsilon' \ W_1' \subseteq \ldots \ \epsilon' \ W_k' \subseteq U' \ \epsilon' \ U'' \subseteq Z_1 \ \epsilon' \ Z_1' \subseteq \ldots \ \epsilon' \ Z_l' \subseteq U$. Thus $Z_l' \subseteq U \subseteq W_1$, as desired. The degenerate cases when $[U']_{\epsilon'}$ is singleton and when U' and when U are comparable by inclusion are handled similarly.

Hence, if we compute which non-singleton ϵ'-classes are $\sim_{\epsilon'}$-related to $[U]_{\epsilon'}$, the class $[U]_{\overline{\epsilon'}}$ is the convex closure of their union. So all we need to do is to merge all non-singleton ϵ'-classes which are $\sim_{\epsilon'}$-related and the relation ϵ'' which is obtained will have as its convex closure $\overline{\epsilon'}$.

Now that we know what to do, it is easy to describe the procedure for doing it. We note that $[U]_{\epsilon'}$ is a non-singleton class iff there exists an edge $V \to W$ in the graph Γ' such that $U \in \{V, W\}$ and that in this case $U^* = W$. We first encode $\preceq_{\epsilon'}$ among non-singleton ϵ'-classes by a directed graph $U^* \rightsquigarrow V^*$ iff there exist some $U' \in [U]_{\epsilon'}$ and $V' \in [V]_{\epsilon'}$ such that $U' \subseteq V'$. This is achieved by checking for inclusion among arbitrary pairs of edges of Γ' which have distinct out-vertices.

Then for any two non-singleton ϵ'-classes $[U]_{\epsilon'}$ and $[V]_{\epsilon'}$, we merge them iff U^* and V^* are in the same strong component with respect to $\rightsquigarrow$ (determining this is known to be tractable by any number of classical algorithms, such as the breadth-first search). Merging the classes $[U]_{\epsilon'}$ and $[V]_{\epsilon'}$ is done by the same method we have been accustomed to already, if $U^* < V^*$, we modify Γ' by replacing all edges of the form $W \to U^*$ with $W \to V^*$ and also adding the edge $U^* \to V^*$. This can be done at most as many times as there are non-singleton $|\epsilon'|$-classes, so the procedure terminates with the graph $\Gamma'' = \Gamma(\epsilon'')$ in polynomial time.

Procedure 5. Checking whether $\overline{\epsilon'}$ is $0-1$ distinguishing. We note that if there exist U, V, W, Z such that $U \ \overline{\epsilon'} \ V \ \overline{\epsilon'} \ W \ \overline{\epsilon'} \ Z$, and that $U \cap V = \emptyset$ and $W \cup Z = E(\Sigma)$, then we may assume that $[U]_{\epsilon'}$, $[V]_{\epsilon'}$, $[W]_{\epsilon'}$ and $[Z]_{\epsilon'}$ are all non-singleton classes. Indeed, if $[U]_{\epsilon'}$ is a singleton class, then there exists a non-singleton class $[U']_{\epsilon'} \sim_{\epsilon'} [U]_{\epsilon'}$ such that $U' \subseteq U$ (and similarly for V), while if $[W]_{\epsilon'}$ is a singleton class, then there exists a non-singleton class $[W']_{\epsilon'} \sim_{\epsilon'} [W]_{\epsilon'}$ such that $W \subseteq W'$ (and similarly for Z). All those follow from our remarks in the proof of Procedure 4.

So, we have just proved that $\overline{\epsilon'}$ is $0-1$ distinguishing iff ϵ'' is $0-1$ distinguishing. It remains to check this. We just run through all edges in Γ'' to check whether such U, V, W and Z can be found among any five sets U_1, U_2, U_3, U_4 and U_5 such that $U_i \to U_5$ in Γ'' for all $1 \leq i \leq 4$. Thus, according to Corollary 3.24, we have efficiently decided whether $\mathbf{D}$ realizes Σ.

4 Ramsey lemma on posets with disjointness

4.1 Posets with disjointness and their representation

This subsection is rather easy and tangential to our main interest. It is included since it is a different language in which some of our results may be expressed.

Let's denote the set of all nonempty subsets of a set X with $\mathcal{P}^+(X)$. There is a disjointness relation naturally defined on $\mathcal{P}^+(X)$ by $A||B$ iff $A \cap B = \emptyset$. Whenever $\mathcal{F} \subseteq \mathcal{P}^+(X)$, the disjointness relation on $\mathcal{F}$ is a symmetrical relation which satisfies that $A||B$ implies $A{\downarrow} \cap B{\downarrow} = \emptyset$ and that, if $A||B$, $A' \subseteq A$ and $B' \subseteq B$, then $A'||B'$. So we may abstractly define a partially ordered set (poset) with disjointness:

Definition 4.1. Let $\mathbb{P} = (P; \leq, ||)$ be a set with two binary relations. $\mathbb{P}$ is a poset with disjointness if $\leq$ is a partial order on P, while $||$ is a symmetric relation on P which satisfies:

(1) $a||b$ implies $a\downarrow \cap\, b\downarrow = \emptyset$ and

(2) if $a||b$, $a' \leq a$ and $b' \leq b$, then $a'||b'$.

It is easy to prove a representation of $(P; \leq, ||)$ as families of sets (the finite case suffices for our purpose, but we may as well prove the infinite one). The following proof is due to the referee, replacing our longer and less elegant one.

Proposition 4.2. *For any poset with disjointness $\mathbb{P}$, there exists a set X and a family $\mathcal{F} \subseteq \mathcal{P}^+(X)$ such that $\mathbb{P} \cong (\mathcal{F}; \subseteq, ||)$ (disjointness in the right-hand model interprets as the actual disjointness of subsets).*

Proof. Let $\mathbb{A} = (\{0,1\}; \leq, D)$, where $D = \{(0,0), (1,0), (0,1)\}$ and let $X := Hom(\mathbb{P}, \mathbb{A})$. By definition, the characteristic function of any up-set U in $\mathbb{P}$ is compatible with the order, while the compatibility with the second relation hinges on the up-set not containing elements $x, y \in U$ such that $x||y$. When $a, b \in P$, by $s_a, s_{a,b} \in A^P$ we denote the characteristic functions of $a\uparrow$ and $a\uparrow \cup\, b\uparrow$, respectively. Hence, $s_a \in X$ for all $a \in P$, while $s_{a,b} \in X$ whenever $a \not|| b$, or $a = b$, using the converse of Definition 4.1 (2).

Define $\varphi : P \to \mathcal{P}(X)$ by $\varphi(a) = \{f \in X : f(a) = 1\}$. Since $s_a \in \varphi(a)$, $\varphi(a) \in \mathcal{P}(X)^+$. If $a \leq b$, then for any $a \in \varphi(a)$, $s(b) \geq s(a) = 1$, so $s(b) = 1$ and $\varphi(a) \subseteq \varphi(b)$. If $a \not\leq b$, then $s_a \in \varphi(a) \setminus \varphi(b)$, so $\varphi(a) \not\subseteq \varphi(b)$. Hence φ is injective and both φ and its inverse are monotone. If $s \in \varphi(a) \cap \varphi(b)$, then $(s(a), s(b)) = (1,1) \notin D$, so $a \not|| b$, since s is a homomorphism. If $a \not|| b$, then $s_{a,b} \in \varphi(a) \cap \varphi(b)$, so $\varphi(a) \cap \varphi(b) \neq \emptyset$. So, $\varphi(P)$ is the desired family $\mathcal{F}$. $\square$

4.2 Monochromatic representation of finite posets with disjointness

Lemma 4.3. *For every finite poset with disjointness $\mathbb{P}$ and every positive integer n there exists an integer N so that for every coloring of $\mathcal{P}^+(N)$ in n colors there exists a monochromatic family $\mathcal{F} \subseteq \mathcal{P}^+(N)$ such that $\mathbb{P} \cong (\mathcal{F}; \subseteq, ||)$.*

Proof. Recall that natural numbers are defined by $M = \{0, 1, \ldots, M-1\}$. For short, we call the sets of size k as k-sets, a coloring in n colors as n-coloring, denote the set of all k-subsets of X as $X^{[k]}$ the set of all nonempty subsets of X of size at most $k-1$ as $X^{[<k]}$. Also, we say that a set is k-monochromatic if all of its k-subsets are of the same color, and for any $K \subseteq \omega$ we say that a set is K-monochromatic if it is k-monochromatic for all $k \in K$. Recall also the definition of the Ramsey numbers $R^n_k(N)$ from Section 2

According to the proof of Proposition 4.2, there exists a positive integer M and a family $\mathcal{G} \subseteq \mathcal{P}(M)^+$ such that $\mathbb{P} \cong (\mathcal{G}; \subseteq, ||)$. We will pick a positive integer N so that for any n-coloring of $\mathcal{P}^+(N)$, we will prove that there exists a monochromatic family $\mathcal{F} \subseteq \mathcal{P}^+(N)$ such that $(\mathcal{P}(M)^+; \subseteq, ||) \cong (\mathcal{F}; \subseteq, ||)$. This would suffice, as $\mathbb{P}$ embeds into $(\mathcal{P}(M)^+; \subseteq, ||)$.

Let t_i be defined as $t_i = 2^{i \cdot M}$ for all $0 < i$. We will show that as the number N from the statement of the lemma we can use

$$N_{n,M} = R^n_{t_1}(R^n_{t_2}(...(R^n_{t_{(M-1)n-1}}(R^n_{t_{(M-1)n}}(t_{(M-1)n+1}))...))).$$

Let A be a set of size $N_{n,M}$ and fix an n-coloring of $\mathcal{P}(A)^+$. $N_{n,M}$, by construction, is large enough for the set A to have a t_1-monochromatic subset A_1 of size $R^n_{t_2}(R^n_{t_3}(...(R^n_{t_{(M-1)n}}(t_{(M-1)n+1}))...))$. Similarly we obtain a t_2-monochromatic subset A_2 of A_1 of size $R^n_{t_3}(...(R^n_{t_{(M-1)n}}(t_{(M-1)n+1}))...)$, and continuing like this we obtain sets $A_1 \supseteq A_2 \supseteq ... \supseteq A_{(M-1)n+1} =: B'$. The set B' is $\{t_1, t_2, ..., t_{(M-1)n+1}\}$-monochromatic (it is $t_{(M-1)n+1}$-monochromatic since $|B'| = t_{(M-1)n+1}$). Using the pigeonhole principle we obtain a subset $\{t_{i_1}, t_{i_2}, ..., t_{i_M}\} \subseteq \{t_1, t_2, ..., t_{(M-1)n+1}\}$ such that $i_1 < i_2 < \cdots < i_M$ and all subsets of B' of sizes t_{i_1}, t_{i_2}, ..., t_{i_M} are of the same color. We choose a t_{i_M}-sized subset of B' and denote it by B. So B has the same color as any of its subsets of size t_{i_j}, for any $1 \leq j < M$.

Let us denote $B^{[t_{i_j}]}$ as $\mathcal{B}_j$ for all $1 \leq j \leq M$. For any $I \in \mathcal{P}(M)^+$, we will choose $D_I \in \mathcal{B}_{|I|}$ (so $D_M = B$) and make $\mathcal{F} = \{D_I : I \in \mathcal{P}(M)^+\}$. Note that any t_{i+1}-set is large enough to have $\binom{M}{k} < 2^M$ disjoint subsets of size t_i (for any $0 \leq k \leq M$). We choose D_I inductively, starting from $|I| = 1$ where we just pick any M disjoint subsets in $\mathcal{B}_1$. If D_J are selected for all $J \in M^{[k-1]}$, for each $I \in M^{[k]}$ we define subsets $E_I = \bigcup\{D_J : J \in I^{[k-1]}\}$. Next we select subsets $E'_I \subseteq B$ such that E'_I is disjoint from any D_J for $|J| < k$ and that for distinct k-subsets I_1 and I_2 of M, E'_{I_1} and E'_{I_2} are disjoint and that for each $I \in M^{[k]}$, $|E'_I| = t_{i_k} - |E_I|$. If we make E'_I disjoint from $\bigcup\{E_J : J \in M^{[k]}\}$, then it will automatically be disjoint from $\bigcup\{D_J : J \in M^{[<k]}\}$, as each such D_J, where $J \in M^{[<k]}$, can be found as a subset within some E_K with

$K \in M^{[k]}$. The following sequence of inequalities shows that the set B is large enough to allow us to select such pairwise disjoint sets E'_I:

$$|B| \geq t_{i_k+1} = 2^M t_{i_k} > \binom{M}{k} t_{i_k} = \sum_{I \in M^{[k]}} t_{i_k} =$$

$$\sum_{I \in M^{[k]}} (|E_I| + (t_{i_k} - |E_I|)) \geq \left| \bigcup \{E_I : I \in M^{[k]}\} \right| + \sum_{I \in M^{[k]}} (t_{i_k} - |E_I|)$$

Finally, make $D_I = E_I \cup E'_I$. Note for future use that $E'_I \cap D_J = \emptyset$ whenever $J \neq I$, and $|J| \leq |I|$.

For notational reasons denote $D_\emptyset := \emptyset$. We want to show that $D_I \cap D_J = D_{I \cap J}$. This would suffice to show that $(\mathcal{P}(M)^+; \subseteq, ||) \cong (\mathcal{F}; \subseteq, ||)$. If $J \subseteq I$, then there exists a sequence of sets $J = J_1 \subseteq J_2 \subseteq \cdots \subseteq J_s = I$ such that $|J_{k+1}| = |J_k| + 1$ and by construction $D_{J_k} \subseteq D_{J_{k+1}}$ for all $1 \leq k < s$. Thus $D_J \subseteq D_I$. In the general case this means that $D_{I \cap J} \subseteq D_I \cap D_J$. We prove the reverse inclusion by an induction on $|I| + |J|$. Assume that $|I| \geq |J|$ and $J \not\subseteq I$. The base case when $|I| = |J| = 1$ is true since then $D_I \cap D_J = \emptyset$. If $|I| = k > 1$, then E'_I is disjoint from D_J by construction, so $D_I \cap D_J = E_I \cap D_J = \bigcup\{D_K : K \in I^{[k-1]}\} \cap D_J = \bigcup\{D_K \cap D_J : K \in I^{[k-1]}\} = \bigcup\{D_{K \cap J} : K \in I^{[k-1]}\} \subseteq D_{I \cap J}$. The last inclusion follows from the observation that for all $K \in I^{[k-1]}$, $D_{K \cap J} \subseteq D_{I \cap J}$. □

5 Canonical decent Mal'cev conditions are realized in all locally finite congruence meet-semidistributive varieties

The following theorem generalizes Theorem 3.2 of [10] and its proof is analogous. We include the full proof to make the paper more self-contained.

Theorem 5.1. *Let Σ be a canonical decent Mal'cev condition. Every locally finite congruence meet-semidistributive variety realizes Σ.*

Proof. Let $\mathcal{V}$ be a locally finite congruence meet-semidistributive variety. Let $\mathcal{W}$ be the idempotent reduct of $\mathcal{V}$, which is the variety whose clone is the clone of idempotent term operations of $\mathcal{V}$ and whose fundamental operations are the distinct elements of this clone. Since congruence meet-semidistributivity can be characterized by an idempotent Mal'cev condition, see Theorem 2.1, $\mathcal{W}$ is a locally finite, idempotent, congruence meet-semidistributive variety. Since all term operations in $\mathcal{W}$ are idempotent and all idempotent term operations of $\mathcal{V}$ are term operations of $\mathcal{W}$, it follows that $\mathcal{V}$ realizes Σ iff $\mathcal{W}$ realizes Σ', which consists of all identities of Σ, except for idempotence.

Now, consider Σ' and let it be represented on $X = \{x_1, \dots, x_m\}$. According to Remark 3.22, Σ' is syntactically equivalent to $\{f(x_1^U, \dots, x_m^U) \approx f(x_1^V, \dots, x_m^V) : U, V \in I_\Sigma\}$, where $I_\Sigma \subseteq \mathcal{P}(X)$ satisfies $(1) - (3)$ of Definition 3.21. Note that each pair of sets $U, V \in I_\Sigma$ satisfies that $U \cup V \neq X$, otherwise $\emptyset \neq X \setminus V \subseteq U$, so $X \setminus V \in I_\Sigma$ by (2). On the other hand, by (1), from $V \in I_\Sigma$ follows that $X \setminus V \in J_\Sigma$. Since I_Σ and J_Σ are equivalence classes of the relation $\epsilon(\Sigma)$, $I_\Sigma \cap J_\Sigma = \emptyset$, so this is impossible. Hence, there must always be at least one $x_i \in X \setminus (U \cup V)$ such that $x_i^U = x_i^V = x$. Define the poset with disjointness $\mathbb{P} = (I_\Sigma; \subseteq, ||)$. We fix $k = |I_\Sigma|$ and $m = |X|$. Let $\mathbf{F}$ be the two-generated free algebra in $\mathcal{W}$, freely generated by x and y and let us fix the number N provided by Lemma 4.3 applied to $\mathbb{P}$ and $n = |F|$.

We define some subuniverses of $\mathbf{F}^2$ (compatible binary relations of $\mathbf{F}$) the same way as in [10] (as with other relations we will define, we permuted variables from their versions in [10] for aesthetic reasons):

$$E = \mathrm{Sg}^{\mathbf{F}^2}\left(\begin{bmatrix} x \\ x \end{bmatrix}, \begin{bmatrix} x \\ y \end{bmatrix}, \begin{bmatrix} y \\ x \end{bmatrix}\right),$$
$$\geq = \mathrm{Sg}^{\mathbf{F}^2}\left(\begin{bmatrix} x \\ x \end{bmatrix}, \begin{bmatrix} y \\ x \end{bmatrix}, \begin{bmatrix} y \\ y \end{bmatrix}\right),$$
$$G = \mathrm{Sg}^{\mathbf{F}^2}\left(\begin{bmatrix} x \\ x \end{bmatrix}, \begin{bmatrix} x \\ y \end{bmatrix}, \begin{bmatrix} y \\ x \end{bmatrix}, \begin{bmatrix} y \\ y \end{bmatrix}\right).$$

We note that, because of idempotence, G is the full product $F \times F$ (it is easy to prove and written up in [10]). The converse of $\geq$ is denoted as $\leq$, while E and G are clearly symmetric. Next we recall the subuniverses $R_1 - R_{11}$ of $\mathbf{F}^3$ defined in [10]. R_i (with permuted variables) can be alternatively defined as:

$$\begin{aligned}
R_1(x_1,x_2,x_3) &:= E(x_1,x_2) \wedge E(x_1,x_3) \wedge E(x_2,x_3),\\
R_2(x_1,x_2,x_3) &:= x_1 \geq x_2 \wedge x_1 \geq x_3 \wedge E(x_2,x_3),\\
R_3(x_1,x_2,x_3) &:= x_1 \geq x_2 \wedge E(x_1,x_3) \wedge E(x_2,x_3),\\
R_4(x_1,x_2,x_3) &:= x_1 \geq x_2 \wedge x_1 \geq x_3 \wedge x_2 \geq x_3,\\
R_5(x_1,x_2,x_3) &:= x_1 \geq x_2 \wedge x_1 \geq x_3 \wedge G(x_2,x_3),\\
R_6(x_1,x_2,x_3) &:= G(x_1,x_2) \wedge x_1 \geq x_3 \wedge x_2 \geq x_3,\\
R_7(x_1,x_2,x_3) &:= x_1 \geq x_2 \wedge G(x_1,x_3) \wedge E(x_2,x_3),\\
R_8(x_1,x_2,x_3) &:= E(x_1,x_2) \wedge E(x_1,x_3) \wedge G(x_2,x_3),\\
R_9(x_1,x_2,x_3) &:= E(x_1,x_2) \wedge G(x_1,x_3) \wedge G(x_2,x_3),\\
R_{10}(x_1,x_2,x_3) &:= x_1 \geq x_2 \wedge G(x_1,x_3) \wedge G(x_2,x_3),\\
R_{11}(x_1,x_2,x_3) &:= G(x_1,x_2) \wedge G(x_1,x_3) \wedge G(x_2,x_3).
\end{aligned}$$

The definitions of R_i as subuniverses of $\mathbf{F}^3$ generated by x, y-valued vector columns are given in [10] and too cumbersome to include here again. However, they are significant as they show that the projection to each pair of coordinates of each R_i is the relation from the above definition, and not smaller. For instance $\pi_{\{2,3\}}(R_8) = G$ and $\pi_{\{1,3\}}(R_2) = \geq$. This is not obvious in general, for instance, take $R'(x_1,x_2,x_3) := x_1 \geq x_2 \wedge E(x_1,x_3) \wedge G(x_2,x_3)$. It is not hard to show that $\pi_{\{2,3\}}(R') = E$, so $\pi_{\{2,3\}}(R')$ does not have to equal G, it can be smaller (if $\mathcal{W}$ has a Mal'cev term then $E = G$, so sometimes $\pi_{\{2,3\}}(R') = G$, but not always). We will use the restrictions of the relations R_i to show 2-consistency of a CSP instance.

The final compatible relation we define is R_I which is the subuniverse of $\mathbf{F}^k$ generated by $\{\overline{a}_1, \ldots, \overline{a}_m\}$ (recall $m = |X|$ and $k = |I_\Sigma|$), where $\overline{a}_i$ are defined in the following way: We fix an arrangement of the sets in I_Σ in a linear order which extends $\supseteq$, so $I_\Sigma = \{U_1, U_2, \ldots, U_k\}$ and $U_i \supseteq U_j \Rightarrow i \leq j$ (in other words, if $i < j$, then $U_i \backslash U_j \neq \emptyset$). Next we define $\overline{a}_i(j) = x$ if $x_i \notin U_j$ and $\overline{a}_i(j) = y$ if $x_i \in U_j$, i.e. $\overline{a}_i(j) = x_i^{U_j}$. Finally, $\overline{a}_i := [\overline{a}_i(1), \ldots, \overline{a}_i(k)]^T$, so we think of $\overline{a}_i$ as vector columns of elements of F.

We consider what are the possibilities for $\pi_{\{i,j\}}(R_I)$, where $1 \leq i < j \leq k$. Since $U_i \cup U_j \neq X$ and $U_i \setminus U_j \neq \emptyset$, as we said, there must be some $1 \leq r, s \leq m$ such that $\overline{a}_r(i) = x = \overline{a}_r(j)$, $\overline{a}_s(i) = y$ and $\overline{a}_s(i) = x$.

- If U_i and U_j are comparable, then $U_i \supsetneq U_j \supsetneq \emptyset$, so there exists $x_t \in U_i \cap U_j$ for some $1 \leq t \leq m$, but $U_j \setminus U_i = \emptyset$. Thus $\overline{a}_t(i) = y = \overline{a}_t(j)$, but for no $1 \leq l \leq m$, $\overline{a}_l(i) = x$ and $\overline{a}_l(j) = y$. Hence, the projection $\pi_{\{i,j\}}(\{\overline{a}_1, \ldots, \overline{a}_m\}) = \{[x,x]^T, [y,x]^T, [y,y]^T\}$ and $\pi_{\{i,j\}}(R_I) = \geq$.

- If U_i and U_j are disjoint, then $U_j \cap U_j = \emptyset$, while $U_j \setminus U_i$ is nonempty (since $U_j \neq \emptyset$), so there exists $1 \leq t \leq m$ such that $\overline{a}_t(i) = x$ and $\overline{a}_t(j) = y$, but for no $1 \leq l \leq m$, $\overline{a}_l(i) = y = \overline{a}_l(j)$. Hence, the projection $\pi_{\{i,j\}}(\{\overline{a}_1, \ldots, \overline{a}_m\}) = \{[x,x]^T, [y,x]^T, [x,y]^T\}$ and $\pi_{\{i,j\}}(R_I) = E$.

- If U_i and U_j are neither comparable, nor disjoint, then $U_j \setminus U_i \neq \emptyset \neq U_i \cap U_j$, so there exist $1 \leq t, l \leq m$ such that $\overline{a}_t(i) = x$, $\overline{a}_t(j) = y$ and $\overline{a}_l(i) = y = \overline{a}_l(j)$. Hence, the projection $\pi_{\{i,j\}}(\{\overline{a}_1, \ldots, \overline{a}_m\}) = \{[x,x]^T, [y,x]^T, [x,y]^T, [y,y]^T\}$ and $\pi_{\{i,j\}}(R_I) = G$.

Next, we define $\mathbb{F} = (F; E, \geq, G, R_1, \ldots, R_{11}, R_I)$ and an instance $(V; F; \mathcal{C})$ of $CSP(\mathbb{F})$. First, let Y be a finite set such that $|Y| = N$ and define $V = \mathcal{P}(Y)^+$.

We proceed with defining $\mathcal{C}$. Let $Z_1, Z_2 \in V$ be distinct. If $Z_1 \supseteq Z_2$, add the constraint $Z_1 \geq Z_2$ to $\mathcal{C}$. If $Z_1 \cap Z_2 = \emptyset$, add the constraint $E(Z_1, Z_2)$ to $\mathcal{C}$. Finally, if Z_1 and Z_2 are neither comparable, nor disjoint, add the constraint $G(Z_1, Z_2)$ to $\mathcal{C}$.

Let $Z_1, Z_2, Z_3 \in V$ be pairwise distinct and such that $i \leq j$ implies $Z_i \nsubseteq Z_j$. Depending on the inclusion and disjointness relations between these sets Z_1, Z_2 and Z_3, the possible cases are:

- If one set contains the other two, then $Z_1 \supseteq Z_2 \cup Z_3$, and we add the constraint $R_2(Z_1, Z_2, Z_3)$, $R_4(Z_1, Z_2, Z_3)$, or $R_5(Z_1, Z_2, Z_3)$ to $\mathcal{C}$, when Z_2 and Z_3 are disjoint, comparable or neither (respectively).

- If one set is contained in the other two, then $Z_3 \subseteq Z_1 \cap Z_2$ and Z_1 and Z_2 are not disjoint, while when Z_1 and Z_2 are comparable, we are in the above case, so the remaining case is when Z_1 and Z_2 are neither disjoint nor comparable. Then we add $R_6(Z_1, Z_2, Z_3)$ to $\mathcal{C}$,

- If there is only one comparability between the three sets, say $Z_1 \supseteq Z_2$, then the possibilities are that Z_3 intersects each of Z_1 and Z_2, that Z_3 intersects Z_1, but not Z_2 and that Z_3 is disjoint from each of Z_1 and Z_2. In those cases, we add $R_{10}(Z_1, Z_2, Z_3)$, $R_7(Z_1, Z_2, Z_3)$, or $R_3(Z_1, Z_2, Z_3)$ to $\mathcal{C}$ (respectively).

- Finally, assume that there is no comparability between the three sets. If all three sets are pairwise disjoint, we add $R_1(Z_1, Z_2, Z_3)$ to $\mathcal{C}$. If Z_1 is disjoint from the other two, but they intersect, then we add $R_8(Z_1, Z_2, Z_3)$ to $\mathcal{C}$. If Z_1 and Z_2 are disjoint, but the other two pairs intersect, then we add $R_9(Z_1, Z_2, Z_3)$ to $\mathcal{C}$. If any pair among Z_1, Z_2

and Z_3 has a nonempty intersection, then we add $R_{11}(Z_1, Z_2, Z_3)$ to $\mathcal{C}$.

As we have covered all possibilities and each $\{Z_1, Z_2, Z_3\}$ has a unique constraint imposed on them, depending only on inclusion and disjointness between the sets Z_1, Z_2 and Z_3, thus our instance is 3-dense. Moreover, since the projections of the ternary constraints onto pairs of coordinates are precisely the binary constraints imposed on those pairs of coordinates, thus our instance is also 2-consistent.

Finally, we add one more type of constraint to $\mathcal{C}$. When $\mathbb{P} \cong (\mathcal{F}; \subseteq, ||)$, where $\mathcal{F} = \{Z_1, \dots, Z_k\} \subseteq \mathcal{P}(Y)^+$, then we select one such isomorphism ϕ[1]. We impose the constraint $R_I(\phi(U_1), \dots, \phi(U_k))$ on $\mathcal{F}$. Of course, adding constraints doesn't violate 3-density of the instance. Since we proved that $\pi_{\{i,j\}}(R_I)$ is either E, $\geq$, $\leq$ or G, depending only on the inclusion and disjointness between U_i and U_j, and since φ is an isomorphism with respect to inclusion and disjointness, thus the 2-consistency is also not violated by adding these constraints to $\mathcal{C}$.

We just proved that the instance $(V; F; \mathcal{C})$ of $CSP(\mathbb{F})$ is $(2,3)$-minimal and nontrivial, while $\mathbb{F}$ is compatible with $\mathbf{F}$ which generates a congruence meet-semidistributive variety. We apply Theorem 2.5 to conclude that $(V; F; \mathcal{C})$ has a solution g.

This solution $g : V \to F$ is a coloring of $V = \mathcal{P}(Y)^+$ into $|F| = n$ colors, so since $|Y| = N$, Lemma 4.3 provides us with family $\mathcal{F} \subseteq \mathcal{P}(Y)^+$ such that $\mathbb{P} \cong (\mathcal{F}; \subseteq, ||)$ and that $g \restriction_{\mathcal{F}}$ is constant, say $g(Z) = u(x,y) \in F$ for all $Z \in \mathcal{F}$. We know that the constraint R_I is imposed on $\mathcal{F}$ in some order, so we have proved that R_I must contain a constant vector $[u, u, \dots, u]^T$. But this means that there exists some $\mathcal{W}$-term $t(x_1, \dots, x_m)$ such that $t^{\mathbf{F}^k}(\overline{a}_1, \dots, \overline{a}_m) = [u, u, \dots, u]^T$. Since the terms in $\mathbf{F}^k$ are computed coordinatewise, thus for each $1 \leq i \leq k$, $t^{\mathbf{F}}(x_1^{U_i}, \dots, x_m^{U_i}) = t^{\mathbf{F}}(\overline{a}_1(i), \dots, \overline{a}_m(i)) = u(x,y)$. Hence, for all $1 \leq i, j \leq k$, $\mathbf{F} \models t^{\mathbf{F}}(x_1^{U_i}, \dots, x_m^{U_i}) = t^{\mathbf{F}}(x_1^{U_j}, \dots, x_m^{U_j})$. But $\mathbf{F}$ is the $\mathcal{W}$-free algebra on $\{x, y\}$ and all equalities $t^{\mathbf{F}}(x_1^{U_i}, \dots, x_m^{U_i}) = t^{\mathbf{F}}(x_1^{U_j}, \dots, x_m^{U_j})$ use just the variables x and y, so for all $U, V \in I_\Sigma$, $\mathcal{W} \models t(x_1^U, \dots, x_m^U) \approx t(x_1^V, \dots, x_m^V)$. Thus, $\mathcal{W}$ realizes Σ', and as we noted in the first paragraph of this proof, this means that $\mathcal{V}$ realizes Σ. $\square$

The rest of this section is devoted to odds and ends we conclude after proving Theorem 5.1, which is the main result of the section.

[1] This detail is missing in the proof of Theorem 3.2 in [10], imposing a constraint for each isomorphic copy of $\mathbb{P}$ on $\mathcal{F}$ could make the intersecting constraint too small.

Proposition 5.2. *A canonical decent Mal'cev condition* Σ *represented on* X *is a Taylor condition iff for any* $x_i \in X$, $I_\Sigma \neq \mathcal{P}(X \setminus \{x_i\})^+$.

Proof. Let $X = \{x_1, \dots, x_n\}$. If $I_\Sigma = \mathcal{P}(X \setminus \{x_i\})^+$, then Σ is realized by $f(x_1, \dots, x_n) = x_i$. Now assume that for any $x_i \in X$, $I_\Sigma \neq \mathcal{P}(X \setminus \{x_i\})^+$. Then for all $1 \leq i, j \leq n$, Σ includes the identities $f(x_1^{\{x_i\}}, \dots, x_n^{\{x_i\}}) \approx f(x_1^{\{x_j\}}, \dots, x_n^{\{x_j\}})$, which, together with idempotence, are the identities which claim that f is a weak near-unanimity operation. Since every weak near-unanimity term is a Taylor term, thus any realization of Σ is a Taylor term. □

Corollary 5.3. *Let* $\mathcal{V}$ *be a locally finite variety and* Σ *a canonical decent Mal'cev condition represented on the set* X *such that* $|X| \geq 4$ *and* Σ *is a Taylor condition. Then* $\mathcal{V}$ *realizes* Σ *iff* $\mathcal{V}$ *is congruence meet-semidistributive. Those are all characterizations of locally finite congruence meet-semidistributive varieties by canonical decent Mal'cev conditions.*

Proof. If Σ is a canonical decent Mal'cev condition represented on the set X such that $|X| \geq 4$ and Σ is a Taylor condition, then Σ is realized in all locally finite congruence meet-semidistributive varieties, according to Theorem 5.1. Moreover, according to Proposition 5.2, I_Σ contains all one-element subsets of X, and since $|X| \geq 4$, I_Σ contains at least one two-element set, say $\{x_k, x_l\}$.

Let $\mathbf{R}$ be any associative ring with unit and $\mathbf{M}$ any right $\mathbf{R}$-module. Assume that $\mathbf{M}$ realizes Σ. We may assume without loss of generality that $\mathbf{M}$ is faithful, since $\mathbf{M}$ is term-equivalent to an $\mathbf{R}/I$ module, where the ideal I is defined by $I = \{\alpha \in R : (\forall x \in M)\alpha x = 0\}$. Hence, there is a term $t(x_1, \dots, x_n) = \sum\limits_{i=1}^{n} \alpha_i x_i$ which realizes all identities in Σ. $t(x_1^{\{x_i\}}, \dots, x_n^{\{x_i\}}) = (\sum\limits_{\substack{j=1 \\ j\neq i}}^{n} \alpha_j)x + \alpha_i y$, so evaluating $x = 0$, the identities $t(x_1^{\{x_i\}}, \dots, x_n^{\{x_i\}}) \approx t(x_1^{\{x_j\}}, \dots, x_n^{\{x_j\}})$ become $\mathbf{M} \models \alpha_i y \approx \alpha_j y$. Because $\mathbf{M}$ is faithful, we conclude that for all $1 \leq i \leq n$, $\alpha_i = \alpha \in R$. The identity $f(x_1^{\{x_1\}}, \dots, x_n^{\{x_1\}}) \approx f(x_1^{\{x_k, x_l\}}, \dots, x_n^{\{x_k, x_l\}})$, again in the case $x = 0$, implies that $\mathbf{M} \models 2\alpha y \approx \alpha y$, so again using faithfulness, $\alpha = 0$. But then $\mathbf{M} \models t(x_1, \dots, x_n) \approx 0$, and idempotence of t implies $\mathbf{M} \models x \approx 0$, so the module $\mathbf{M}$ is trivial. Since any module which realizes Σ is trivial, according to [11], Theorem 8.1 (1) $\Leftrightarrow$ (10), Σ is realized only in congruence meet-semidistributive varieties.

The final sentence follows from the fact that any such characterization Σ is not realized by the Abelian group $\mathbf{Z}_2$. However, all those Σ which are not Taylor conditions are realized by a projection (in any nontrivial algebra, hence also in $\mathbf{Z}_2$), while those Σ that are Taylor conditions and have less than four variables must be the ternary weak near-unanimity term, which is realized in $\mathbf{Z}_2$ by $t(x, y, z) = x + y + z$. $\square$

Another aspect in which Theorem 5.1 can be improved is the following result, also proved by Z. Brady in [4]:

Theorem 5.4. *Let $\mathcal{V}$ be a locally finite congruence meet-semidistributive variety. There exists a binary $\mathcal{V}$-term $p(x, y)$ such that for every $n > 2$, every $X = \{x_1, \ldots, x_n\}$ and every canonical decent Mal'cev condition Σ represented on X, there exists a realization of Σ in $\mathcal{V}$ in which f is interpreted as some $\mathcal{V}$-term t so that $\mathcal{V} \models t(x_1^U, \ldots, x_n^U) \approx p(x, y)$ for all $U \in I_\Sigma$.*

Moreover, $p(x, y)$ can be chosen so that $\mathcal{V} \models p(p(x, y), p(y, x)) \approx p(x, y)$

Proof. Such a term $p(x, y)$, if it exists, would have to be idempotent because decent Mal'cev conditions are all idempotent. So, like in the proof of Theorem 5.1, we turn to the idempotent reduct of $\mathcal{V}$, the variety $\mathcal{W}$. Let $\mathbf{F} = \mathbf{F}_\mathcal{W}(x, y)$ be the free algebra and $|F| = n$.

We assume the opposite. Let the set of all pairwise distinct binary $\mathcal{W}$-terms be $\{p_1(x, y), \ldots, p_n(x, y)\}$ (this is a set of representatives for the elements of $\mathbf{F}$). Then for every $1 \leq j \leq n$ there exists a canonical decent Mal'cev condition Σ_j represented on $X_j = \{x_{j,1}, \ldots, x_{j,m_j}\}$, determined by $I_{\Sigma_j} = \{U_1^j, \ldots, U_{k_j}^j\} \subseteq \mathcal{P}^+(X_j)$, such that for every realization of Σ_j in $\mathcal{V}$, where some $\mathcal{V}$-term t interprets the operation of Σ_j, and for any $U \in I_{\Sigma_j}$, $\mathcal{V} \not\models t(x_{j,1}^U, \ldots, x_{j,m_j}^U) \approx p_j(x, y)$.

Moreover, for any canonical decent Mal'cev condition Σ represented on the set of variables $X = \{x_1, \ldots, x_k\}$ there exists another canonical decent Mal'cev condition Σ, represented on a set of variables $Y = \{x_1, \ldots, x_l\}$, such that $k < l$ and the set of binary $\mathcal{W}$-terms

$$\{p(x, y) : (\exists t)(\forall U \in I_\Sigma)\mathcal{W} \models t(x_1^U, \ldots, x_k^U) \approx p(x, y)\} \supseteq$$
$$\{p(x, y) : (\exists t)(\forall U \in I_{\Sigma'})\mathcal{W} \models t(x_1^U, \ldots, x_l^U) \approx p(x, y)\}.$$

Proof for the case $l = k + 1$: Define $I_{\Sigma'}$ and $J_{\Sigma'}$ by:

- when $U \in I_\Sigma$, make $U, U \cup \{x_{k+1}\} \in I_{\Sigma'}$,
- when $U \in J_\Sigma$, make $U, U \cup \{x_{k+1}\} \in J_{\Sigma'}$,

- $\{x_{k+1}\} \in I_{\Sigma'}$ and $X \in J_{\Sigma'}$.

Now it is routine to verify that $I_{\Sigma'} \cup \{\emptyset\}$ is a down-set since $I_{\Sigma} \cup \{\emptyset\}$ is, and that $U \in I_{\Sigma'}$ iff $(Y \setminus U) \in J_{\Sigma'}$, again using the same property for I_{Σ}, J_{Σ} and X. Hence $J_{\Sigma'} \cup \{Y\}$ is an up-set. If there exist $U, V \in I_{\Sigma'}$ such that $U \cup V = Y$, then $U \setminus \{x_{k+1}\}$ and $V \setminus \{x_{k+1}\}$ are both in I_{Σ} and their union is X, contradicting the assumption that Σ is a canonical decent Mal'cev condition. Finally, note that if $x_k \in U$ and $\emptyset \subsetneq U \subsetneq X$, then $U \in I_{\Sigma}$ iff $U \cup \{x_{k+1}\} \in I_{\Sigma'}$ (since the conclusion works even if $x_k \notin U$). Hence, for any realization t of Σ', $t(x_1, x_2, \ldots, x_k, x_k)$ is a realization of Σ, proving the desired containment for the case $l = k + 1$.

For $l > k + 1$ we proceed inductively adding one variable at a time. In conclusion, we may assume that $i < j$ implies $|X_i| < |X_j|$ by adding more variables to the set on which Σ_j is represented, as needed.

Without loss of generality, we may assume that all X_j are pairwise disjoint and select X such that $X_j \subseteq X$ for all $1 \leq j \leq n$ and such that $|X| = 1 + 2 \sum_{j=1}^{n} m_j =: m$. We have all sets $I_{\Sigma_j} \subseteq \mathcal{P}(X_j) \subseteq \mathcal{P}(X)$. Define the canonical decent Mal'cev condition Σ represented on X by $I_{\Sigma} = \{U \subseteq X : 1 \leq |U| \leq \sum_{j=1}^{n} m_j = \frac{m-1}{2}\}$.

Now we modify the proof that $\mathcal{V}$ realizes Σ from Theorem 5.1 in the following way: we keep all the constraint relations and constraints we had there, and we add the compatible relations R_{I_j}, $1 \leq j \leq n$, to the constraint language. For each $1 \leq j \leq n$, R_{I_j} is the subalgebra of $\mathbf{F}^{k_j}$ generated by $\{\overline{a}_1^j, \ldots, \overline{a}_{m_j}^j\}$ so that for all $1 \leq r \leq m_j$ and $1 \leq s \leq k_j$, $\overline{a}_r^j(s) = y$ if $x_{j,r} \in U_s^j$, while $\overline{a}_r^j(s) = x$ otherwise.

Now we impose additional constraints with constraint relations R_{I_j}. Let T be any set of variables on which we imposed R_I defined by the mapping ϕ_T which is an isomorphism between $(I_{\Sigma}; \subseteq, ||)$ and $(T; \subseteq, ||)$. For each such T and each $1 \leq j \leq n$ we will select one subset $S_j^T \subseteq T$ and impose R_{I_j} on S_j^T. We select the subsets $S_j^T \subseteq T$ so that $S_j^T = \{\phi(U) : U \subseteq X_j \subseteq X$ and $U \in I_{\Sigma_j}\}$. We impose R_{I_j} on S_j^T in some way which is determined by a fixed isomorphism between $(I_{\Sigma_j}; \subseteq, ||)$ and $(S_j^T; \subseteq, ||)$. The sets S_j^T and $S_{j'}^{T'}$ never coincide for $j \neq j'$, since $|S_j| = 2^{|X_j|-1} - 1$, so from $i < j \Rightarrow |X_i| < |X_j|$ follows that the constraint relation imposed on a set S, if any, is always the same. If there exists T such that $S = S_j^T$, then we always select the same isomorphism from $(I_{\Sigma_j}; \subseteq, ||)$ and $(S; \subseteq, ||)$ for any T' such that $S = S_j^{T'}$. Therefore the constraint relation on S_j, doesn't depend on T.

Now there is no set of variables which is subject to more than one constraint, the instance is still 3-dense, and the 2-consistency follows from the fact that the restriction to pairs of variables is always the binary constraint determined by inclusion and disjointness, which does not depend on the larger constraint, just on the relationship between those two variables. So we know from Theorem 2.5 that there is a solution g. As before, using Lemma 4.3 we conclude that for some set of variables T on which the constraint relation R_I is imposed, all variables in T are mapped by g to the same element $u \in F$. This element u is $\mathcal{V}$-equal to one of $p_j(x, y)$, for some $1 \leq j \leq n$. But each R_{I_r} is imposed on some subset of T. Hence, the constraint R_{I_j} is imposed on the subset $S_j \subseteq T$. We conclude that there exists some $\mathcal{W}$-term $t(x_1, \ldots, x_{m_j})$ such that $t^{\mathbf{F}^{k_j}}(\overline{a}_1^j, \ldots, \overline{a}_{m_j}^j) = [u, u, \ldots, u]^T$. Computed coordinatewise, for each $U_i^j \in I_{\Sigma_j}$, $t^{\mathbf{F}}(x_{j,1}^{U_i^j}, \ldots, x_{j,m_j}^{U_i^j}) = t^{\mathbf{F}}(\overline{a}_1(i), \ldots, \overline{a}_{m_j}(i)) = u(x, y)$. Thus $\mathcal{W}$ realizes Σ_j, take away idempotence, but in such a way that $\mathcal{W} \models t(x_{j,1}^{U_i^j}, \ldots, x_{j,m_j}^{U_i^j}) \approx p_j(x, y)$ for all $U_i^j \in I_{\Sigma_j}$. Since $\mathcal{W}$ is the idempotent reduct of $\mathcal{V}$, this implies that $\mathcal{V} \models t(x_{j,1}^{U_i^j}, \ldots, x_{j,m_j}^{U_i^j}) \approx p_j(x, y)$ for all $U_i^j \in I_{\Sigma_j}$ and t is an idempotent term of $\mathcal{V}$. This is a contradiction with the choice of Σ_j.

It remains to prove the final sentence. We define a sequence of binary terms $p^{(1)}(x, y) = p(x, y)$ and $p^{(k+1)}(x, y) = p(p^{(k)}(x, y), p^{(k)}(y, x))$. We aim to prove that for all $k \geq 1$, $p^{(k)}$ can replace p in the first paragraph of our theorem. The base case is proved, so we proceed by an induction on k.

Let Σ be any decent Mal'cev condition on $X_n := \{x_1, \ldots, x_n\}$, where $n > 2$. Also let I_Σ be the family of subsets of X_n associated with Σ and $t(x_1, \ldots, x_n)$ a realization of Σ in $\mathcal{V}$ such that for all $U \in I_\Sigma$, $\mathcal{V} \models t(x_1^U, \ldots, x_n^U) \approx p(x, y)$. Denote $t^{(1)} := t$. Define the canonical decent Mal'cev condition $\Sigma_{<n}$ represented on the set $X_{2n-1} = \{x_1, x_2, \ldots, x_{2n-1}\}$ by $I_{\Sigma_{<n}} := \{U \subseteq X_{2n-1} : 1 \leq |U| < n\}$. By the inductive assumption applied to the condition $\Sigma_{<n}$, there exists a realization of $\Sigma_{<n}$ by a $(2n-1)$-ary $\mathcal{V}$-term $q_n^{(k)}$ such that for all $U \in I_{\Sigma_{<n}}$, $\mathcal{V} \models q_n^{(k)}(x_1^U, x_2^U, \ldots, x_{2k-1}^U) \approx p^{(k)}(x, y)$.

We introduce a little notation, x_i^s denotes $x_i, x_i, \ldots, x_i$, where x_i repeats s times. We define $t^{(k+1)}$ inductively:

$$\begin{aligned} t^{(k+1)}(x_1, \ldots, x_n) := t(q_n^{(k)}(x_1^n, x_2, x_3, \ldots, x_n), \\ q_n^{(k)}(x_1, x_2^n, x_3, \ldots, x_n), \ldots, q_n^{(k)}(x_1, x_2, \ldots, x_{n-1}, x_n^n)) \end{aligned} \tag{5.1}$$

Fix an arbitrary $\emptyset \neq U \subsetneq X_n$. We obtain that

$$\mathcal{V} \models q_n^{(k)}(x_1^U, \ldots, x_{i-1}^U, (x_i^U)^n, x_{i+1}^U, \ldots, x_n^U) \approx p^{(k)}(x,y)$$
$$\text{iff } |U \cap (X_n \setminus \{x_i\})| + n|U \cap \{x_i\}| \leq n \text{ iff } x_i \notin U.$$

On the other hand, analogously

$$\mathcal{V} \models q_n^{(k)}(x_1^U, \ldots, x_{i-1}^U, (x_i^U)^n, x_{i+1}^U, \ldots, x_{n+1}^U) \approx p^{(k)}(y,x) \text{ iff } x_i \in U.$$

From (5.1) follows that $t_n^{(k+1)}(x_1^U, \ldots, x_n^U)$ equals in $\mathcal{V}$ to a substitution instance of $t(x_1^U, \ldots, x_n^U)$ where x is substituted by $p^{(k)}(x,y)$ and y is substituted by $p^{(k)}(y,x)$. Since t is a realization of Σ in $\mathcal{V}$, with the derived operation $p(x,y)$, it follows that for any $U \subseteq X_k$, $U \in I_\Sigma$ iff $\mathcal{V} \models t^{(k+1)}(x_1^U, \ldots, x_n^U) \approx p(p^{(k)}(x,y), p^{(k)}(y,x)) = p^{(k+1)}(x,y)$, finishing the inductive proof. $\square$

Our (messy) proof of Theorem 5.4 is included to demonstrate that our framework is equally powerful as Z. Brady's from [4]. For now we may prove the most elegant characterization of congruence meet-semidistributive varieties from [4]:

Theorem 5.5. *A locally finite variety $\mathcal{V}$ is congruence meet-semidistributive iff it realizes the following Mal'cev condition:*

$$\begin{gathered} f(x,x,x) \approx x \\ f(f(x,x,y), f(x,x,y), f(y,y,x)) \approx f(x,x,y) \approx f(x,y,x) \approx f(y,x,x). \end{gathered} \tag{5.2}$$

Proof. The implication ($\Rightarrow$) follows directly from Theorem 5.4 applied to the canonical decent Mal'cev condition W represented on $\{x_1, x_2, x_3\}$ given by $I_W = \{\{x_1\}, \{x_2\}, \{x_3\}\}$. Any realization of W is a ternary weak near-unanimity term, while the final sentence of Theorem 5.4 implies

$$\mathcal{V} \models f(f(x,x,y), f(x,x,y), f(y,y,x)) \approx f(x,x,y).$$

Now we prove the implication ($\Leftarrow$). Let $\mathbf{R}$ be any associative ring with unit and $\mathbf{M}$ any right $\mathbf{R}$-module. Assume that $\mathbf{M}$ realizes (5.2). We may assume without loss of generality that $\mathbf{M}$ is faithful. Hence, there is a term $t(x_1, x_2, x_3) = \sum_{i=1}^{3} \alpha_i x_i$ which realizes all identities in (5.2). By plugging in $x = 0$, we get $\alpha_1 y = \alpha_2 y = \alpha_3 y$ for all y, and the faithfulness of $\mathbf{M}$ implies $\alpha_1 = \alpha_2 = \alpha_3$ in $\mathbf{R}$ (denote $\alpha := \alpha_1$). Applying $x = y$ in (5.2), we get $3\alpha x =$

x, so $3\alpha = 1$ by faithfulness of $\mathbf{M}$. Plugging $x = 0$ again into (5.2) we get $\alpha^2 y + \alpha^2 y + 2\alpha^2 y = \alpha y$. Hence $\alpha(3\alpha+\alpha)y = y$, which implies $\alpha(\alpha+1)y = \alpha y$, from which $\alpha^2 y + \alpha y = \alpha y$ and finally $\alpha^2 y = 0y$, which means $\alpha^2 = 0$ by the faithful property. But then $0 = 9\alpha^2 = (3\alpha)(3\alpha) = 1 \cdot 1 = 1$. Hence for all $x \in M$, $x = 1x = 0x = 0$, so $\mathbf{M}$ is trivial. Since any module which realizes Σ is trivial, according to [11], Theorem 8.1 (1) $\Leftrightarrow$ (10), Σ is realized only in congruence meet-semidistributive varieties. $\square$

6 Putting it all together

The following theorem summarizes the main line of the results of our paper:

Theorem 6.1. *Let Σ be a decent Mal'cev condition on the set of variables X. The following conditions are equivalent:*

(1) *Every locally finite congruence meet-semidistributive variety realizes Σ.*

(2) $\mathbf{D}$ *realizes Σ.*

(3) *There exists a canonical decent Mal'cev condition Π represented on $E \subseteq X$ such that $\{(U \cap E, V \cap E) : (U, V) \in \epsilon(\Sigma)\} \subseteq \epsilon(\Pi)$.*

Moreover, whether these three conditions are satisfied can be checked in polynomial time in $|\Sigma|$ and $|X|$.

Proof. (1) $\Rightarrow$ (2) $\mathcal{V}(\mathbf{D})$ is locally finite since $\mathbf{D}$ is finite. Also, $\mathbf{D}$ generates a congruence meet-semidistributive variety, as was noticed in [9] (we could also see this directly, it is easy to check that $\mathbf{D}$ realizes all known strong Mal'cev characterizations of congruence meet-semidistributivity).

(2) $\Rightarrow$ (3) This follows from Lemma 3.20 by taking $E := E(\Sigma)$.

(3) $\Rightarrow$ (1) Let $\mathcal{V}$ be a locally finite congruence meet-semidistributive variety. By Theorem 5.1, $\mathcal{V}$ realizes Π. Hence, where $E = \{x_{i_1}, \ldots, x_{i_k}\}$, $\mathcal{V}$ realizes the condition $\Sigma' = \{g(x_{i_1}^{U\cap E}, \ldots, g(x_{i_k}^{U\cap E}) \approx g(x_{i_1}^{V\cap E}, \ldots, g(x_{i_k}^{V\cap E}) : (U, V) \in \epsilon(\Sigma)\}$. According to Lemma 3.5, $\mathcal{V}$ realizes Σ by interpreting its operation as some term on the set of variables E.

Finally, the algorithm for efficiently checking (2) was described in Subsection 3.4. $\square$

For the next corollary, which is another restatement of our main results, we ask the reader to recall a little notation and some results from [8].

Corollary 6.2. *Let Σ be a linear, idempotent strong Mal'cev condition in which only two variables occur. Then the following are equivalent:*

(1) *Σ is realized in $\mathbf{D}$, but not in any nontrivial finite module.*

(2) *For any locally finite variety $\mathcal{V}$, $\mathcal{V}$ realizes Σ iff $\mathcal{V}$ is congruence meet-semidistributive.*

Proof. (1) $\Rightarrow$ (2) If Σ is realized in $\mathbf{D}$, then we can, according to Proposition 3.1, select the decent Mal'cev condition Σ', so that Σ' is realized in any variety which realizes Σ. By Theorem 6.1, Σ' is realized in every locally finite congruence meet-semidistributive variety and using Proposition 3.1 again, we conclude that Σ is realized in every locally finite congruence meet-semidistributive variety.

On the other hand, assume that a locally finite variety $\mathcal{V}$ realizes Σ. Let $\mathcal{E} := \mathrm{Mod}(\Sigma)$. By the terminology introduced in Definition 9.1 of [8], $\mathcal{E}$ and by assumption (1), for every finite field $\mathbf{F}$, the one-dimensional vector space ${}_{\mathbf{F}}\mathbf{V}$ over $\mathbf{F}$ does not realize Σ. The variety ${}_{\mathbf{F}}\mathcal{V}$ of all vector spaces over $\mathbf{F}$ is generated by ${}_{\mathbf{F}}\mathbf{V}$, so by the notation from the beginning of Chapter 9 of [8], $\mathcal{E} \not\leq {}_{\mathbf{F}}\mathbf{V}$. Finally, $\mathcal{V}$ realizes Σ, so $\mathcal{E} \leq \mathcal{V}$ (again using the notation of [8]). Hence, by the implication (2) $\Rightarrow$ (5) of Theorem 9.10 of [8], $\mathcal{V}$ is congruence meet-semidistributive.

(2) $\Rightarrow$ (1) The easy direction is true since $\mathcal{V}(\mathbf{D})$ is a congruence meet-semidistributive locally finite variety, but $\mathcal{V}(\mathbf{M})$ is not, where $\mathbf{M}$ is any nontrivial finite module. $\square$

7 Problems and remarks

One obvious way to improve on our results would be to remove the condition that there are only two variables involved in any identity. This is not prescribed by our technique (and that is its chief selling point as opposed to Brady's from [4], his technique seems to be inherently limited to two variables); for instance, if we used three variables, then we would encode $f(x, y, y, z, x, z)$ as the equivalence relation with named classes $C_x = \{x_1, x_5\}$, $C_y = \{x_2, x_3\}$ and $C_z = \{x_4, x_6\}$. Since we were dealing with two-class equivalence relations only, we were able to encode them by just considering the sets C_y; C_x was understood to be its complement. However, while we needed consider only disjointness, containment in one of the directions, and "neither" in the two-variable case, in the situation with only three variables, we are already faced with many more possible interactions between two equivalence relations (evaluations). This makes the Ramsey argument much more complex. Moreover, consider the following examples:

Example 7.1. The strong Mal'cev conditions

$$\begin{aligned} t(x,x,x,x) &\approx x \\ t(x,x,y,z) \approx t(y,z,y,x) &\approx t(x,z,z,y) \end{aligned} \tag{BD}$$

and

$$\begin{aligned} t(x,x,x,x) &\approx x \\ t(x,x,y,z) \approx t(y,x,z,x) &\approx t(y,z,x,y) \end{aligned} \tag{LS}$$

are realized in $\mathbf{D}$, but not in all locally finite congruence meet-semidistributive varieties.

The above examples were discovered by a computer search in [10] but it was left as an open question whether they are realized in all locally finite congruence meet-semidistributive varieties. Z. Brady in [4], also independently M. Maróti (circulated by email to the participants of the 2016 Nashville workshop), proved they are not realized in certain locally finite congruence meet-semidistributive varieties. Hence, we don't have a good candidate for the set of all linear idempotent strong Mal'cev conditions realized by all locally finite congruence meet-semidistributive varieties. In the two-variable case we were able to just check which ones are realized in $\mathbf{D}$, but in the more general seting this is not the same set of conditions and we don't know how to narrow it down.

Still, though the above discussion indicates it is probably difficult, we believe that the following problem may be attempted:

Problem 7.2. Characterize all linear idempotent strong Mal'cev conditions which are realized in all locally finite congruence meet-semidistributive varieties.

Another, perhaps even more challenging, and certainly more attractive, problem lies in shedding the local finiteness. In that direction, the only known result is the following shocker by M. Olšák from [20]:

Theorem 7.3. *A variety $\mathcal{V}$ realizes some Taylor condition iff $\mathcal{V}$ realizes*

$$\begin{aligned} t(x,x,x,x,x,x) &\approx x \\ t(x,y,x,y,x,y) \approx t(y,y,y,x,x,x) &\approx t(x,x,y,y,y,x) \end{aligned} \tag{O}$$

Since there exists a strong Mal'cev characterization of having a Taylor term, it is conceivable that congruence meet-semidistributivity also admits such a strong Mal'cev characterization. The problem was attempted ever since Olśak circulated a draft of his result in 2016, but it remains open:

Problem 7.4. Does there exist a strong Mal'cev condition Σ such that for any variety $\mathcal{V}$, $\mathcal{V}$ is congruence meet-semidistributive iff $\mathcal{V}$ realizes Σ?

Acknowledgement

The authors thank Z. Brady for suggesting we "just make the full binary composition of $p(x, y)$ with itself" which allowed us to prove the last sentence of Theorem 5.4 and Theorem 5.5. The anonymous, but very knowledgeable, referee also contributed a lot in an exceptionally detailed review. Along with several other improvements to our paper, the referee's proofs of Proposition 4.2 and Corollary 6.2 replaced our inferior ones.

References

[1] Barto, L.: The collapse of the bounded width hierarchy. J. Logic Comput. **26**, 923–943 (2016)

[2] Barto, L., Kozik, M.: Constraint satisfaction problems solvable by local consistency methods. J. ACM **61** 1:03, 19 pp. (2014)

[3] Bergman, C.: Universal Algebra: Fundamentals and Selected Topics. Chapman & Hall/CRC Press, Boca Raton (2011)

[4] Brady, Z.: Examples, Counterexamples, and Structure in Bounded Width Algebras. `http://math.stanford.edu/~notzeb/bounded-width.pdf` (manuscript, 2017)

[5] Burris, S., Sankappanavar, H.P.: A course in universal algebra. Graduate Texts in Mathematics, vol. 78. Springer, New York (1981)

[6] Czedli, G.: A characterization for congruence semi-distributivity. In: Proc. Conf. Universal Algebra and Lattice Theory, Puebla (Mexico, 1982), Lecture Notes in Math., vol. 1004, pp. 104–110. Springer, New York (1983)

[7] Freese, R., McKenzie, R.: Commutator Theory for congruence modular varieties. London Math. Soc. Lecture Note Series, vol. 125. Cambridge University Press, Cambridge (1987)

[8] Hobby, D., McKenzie, R.: The structure of finite algebras. Contemporary Mathematics, vol. 76. American Mathematical Society, Providence (1988)

[9] Jovanović, J.: On terms describing omitting unary and affine types. Filomat **27**, 183–199 (2013)

[10] Jovanović, J., Marković, P., McKenzie, R., Moore, M.: Optimal strong Mal'cev conditions for congruence meet-semidistributivity in locally finite varieties. Algebra Universalis **76**, 305-325 (2016).

[11] Kearnes, K., Kiss, E.: The shape of congruence lattices. Mem. Amer. Math. Soc., vol. 222, no. 1046. American Mathematical Society, Providence (2013)

[12] Kearnes, K., Marković, P., McKenzie, R.: Optimal strong Mal'cev conditions for omitting type **1** in locally finite varieties. Algebra Universalis **72**, 91–100 (2014)

[13] Kearnes, K., Szendrei, Á.: The relationship between two commutators. Internat. J. Algebra Comput. **8**, 497–531 (1998)

[14] Kearnes, K., Willard, R.: Residually finite, congruence meet-semidistributive varieties of finite type have a finite residual bound. Proc. Amer. Math. Soc. **127**, 2841–2850 (1999)

[15] Kozik, M., Krokhin, A., Valeriote, M. Willard, R.: Characterizations of several Maltsev conditions. Algebra Universalis **73**, 205–224 (2015)

[16] Larose, B., Zádori, L.: Bounded width problems and algebras. Algebra Universalis **56**, 439-466 (2007)

[17] Lipparini, P.: A characterization of varieties with a difference term, II: Neutral = meet semi-distributive. Canad. Math. Bull. **41**, 318–327 (1998)

[18] Maróti, M., McKenzie, R.: Existence theorems for weakly symmetric operations. Algebra Universalis **59**, 463–489 (2008)

[19] McKenzie, R., McNulty, G., Taylor, W.: Algebras, lattices, varieties. Vol. I. The Wadsworth & Brooks/Cole Mathematics Series, Wadsworth & Brooks/Cole Advanced Books & Software, Monterey (1987)

[20] Olšák, M.: The weakest nontrivial idempotent equations. `https://arxiv.org/abs/1609.00531` (manuscript, 2016)

[21] Ramsey, F. P.: On a problem of formal logic. Proc. London Math. Soc 30, Series 2, 264-286 (1930)

[22] Siggers, M.: A strong Mal'cev condition for locally finite varieties omitting the unary type. Algebra Universalis **64**, 15–20 (2010)

[23] Taylor, W.: Varieties obeying homotopy laws. Canad. J. Math. **29** 498-527 (1977)

[24] Willard, R.: A finite basis theorem for residually finite, congruence meet-semidistributive varieties. J. Symb. Log. **65**, 187–200 (2000)