Optimal strong Mal'cev conditions for omitting type 1 in locally finite varieties

Keith Kearnes, Petar Marković, and Ralph McKenzie

Abstract. We show that the class of locally finite varieties omitting type **1** has the following properties. This class

(1) is definable by an idempotent, linear, strong Mal'cev condition in a language with one 4-ary function symbol.
(2) is not definable by an idempotent, linear, strong Mal'cev condition in a language with only one function symbol of arity strictly less than 4.
(3) is definable by an idempotent, linear, strong Mal'cev condition in a language with two 3-ary function symbols.
(4) is not definable by an idempotent, linear, strong Mal'cev condition in a language with function symbols of arity less than 4 unless at least two of the symbols have arity 3.

1. Introduction

In [12], W. Taylor identified the weakest nontrivial idempotent Mal'cev condition, namely the one we now call "existence of a Taylor term". The class of locally finite varieties having a Taylor term was characterized in other ways by Hobby and McKenzie in [7], namely it is the class of locally finite varieties omitting type **1** and it is also the class of locally finite varieties with a weak difference term. More recently [9] and [1] have useful new Mal'cev conditions equivalent to Taylor's for locally finite varieties.

The most surprising recent result on this topic was the discovery of M. Siggers [11] that the property of omitting type **1** is equivalent to a *strong* Mal'cev condition for locally finite varieties. Siggers proved that the class of locally finite varieties omitting type **1** is definable within the class of all locally finite varieties as the subclass of varieties with a 6-variable Taylor term. To do this, he used a result of P. Hell and J. Nešetřil from [5]. In this note we strengthen Siggers' result to obtain optimal strong Mal'cev conditions for omitting type **1** using the generalization of the Hell-Nešetřil theorem that was proved by L. Barto, M. Kozik and T. Niven in [2]. The strong Mal'cev conditions expressing the omission of type **1** are those of Theorem 2.2 and Corollaries 3.1 and 3.2, while the optimality of the conditions is established in Theorem 3.4.

Key words and phrases: Mal'cev condition, variety, Taylor term, weak near-unanimity term, tame congruence theory, unary type.

The second author was supported by the grant no. 174018 of the Ministry of Education and Science of Serbia.

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 $\mathcal{V}$).

More formally, let $\mathcal{L}$ be a language and let $X = (x_1, x_2, \ldots)$ be a sequence of distinct variables. For each function symbol f of $\mathcal{L}$ of arity n, call the term $f(x_1, \ldots, x_n)$, whose variables are the first n variables of X, an *elementary $\mathcal{L}$-term*. Let Σ be a strong Mal'cev condition in the language $\mathcal{L}$. We say that Σ is *realized* in a clone $\mathcal{C}$ (of an algebra or variety) if for each n-ary function symbol f of $\mathcal{L}$ it is possible to assign to the elementary term $f(x_1, \ldots, x_n)$ an n-ary element of $\mathcal{C}$ in such a way that the assignment extends to a homomorphism of the $\mathcal{L}$-term algebra over X into $\mathcal{C}$ such that Σ is a subset of the kernel of that homomorphism. (In the case where $\mathcal{C}$ is the clone of a variety $\mathcal{V}$ we will write $f(x_1, \ldots, x_n) \mapsto t^{\mathcal{V}}$ to indicate that the elementary term $f(x_1, \ldots, x_n)$ has been assigned $t^{\mathcal{V}}$, where $t^{\mathcal{V}}$ denotes the $\mathcal{V}$-equivalence class of the n-ary $\mathcal{V}$-term t.)

An operation f on a set A is *idempotent* if $f(x, x, \ldots, x) = x$ for all $x \in A$. A strong Mal'cev condition Σ is *idempotent* if it asserts the idempotence of each of its function symbols, i. e., if $\Sigma \models f(x, x, \ldots, x) \approx x$ for each function symbol f appearing in Σ. A term t is *linear* if it involves at most one function symbol. An identity $s \approx t$ is *linear* if both s and t are linear. A strong Mal'cev condition is *linear* if its identities are linear. The strong Mal'cev conditions considered in this paper are idempotent and linear.

An idempotent strong Mal'cev condition Σ is called a *Taylor condition* unless it can be realized in every variety, equivalently, if it cannot be realized in the variety of sets. (In practice the property of being a Taylor condition is established by showing that Σ cannot be realized in a way where all the function symbols of Σ interpret as projection operations.) If, moreover, the language $\mathcal{L}$ of Σ has exactly one operation symbol, say f, then in a realization defined by $f(x_1, \ldots, x_n) \mapsto t^{\mathcal{V}}$ of Σ in a variety $\mathcal{V}$, the term t is called a *Taylor term* for $\mathcal{V}$.

The definitions and basic results of universal algebra can be found in the textbooks [4] and [10], while type **1** and similar notions of tame congruence theory are defined and developed in the monograph [7].

2. 'Omitting type 1' expressed with one 4-ary term

To prove the result from the title of this section, we need some graph-theoretic definitions which we take from [6] and modify to allow for symmetric edges. A directed graph, or digraph, is just a relational structure with signature consisting of one binary relation; elements of the universe are called *vertices*, while elements of the relation are called *edges*. A vertex x is a *source* when $x \in V - \pi_2(E)$, and x is a *sink* when $x \in V - \pi_1(E)$. An *oriented path* p in a directed graph $G = (V, E)$ is a pair of finite sequences, one of vertices $v_0, v_1, \dots, v_n$ and one of edges $e_1, e_2, \dots, e_n$, $n \geq 0$, so that for all i, $e_i = v_{i-1}v_i$ or $e_i = v_i v_{i-1}$. The edges of the former kind are called *forward edges* of p and the edges of the latter kind are *backward edges* of p. A path p is *closed* if $v_0 = v_n$. The equivalence relation on V which is the reflexive closure of the relation "u and v are endpoints of some oriented path" is called *weak connectedness.* The difference between the number of forward edges and the number of backward edges of an oriented path p is called *algebraic length* of p, written $Al(p)$. The algebraic length of a component C of weak connectedness of a digraph is the least $n > 0$ so that there exists a closed oriented path p in C such that $Al(p) = n$, or 0 if no such path exists.

A useful special case of the main result of [2] is the Loop Lemma:

Lemma 2.1. *(Theorem 8.1 of* [2]*) If a digraph without sources and sinks is compatible with a weak near-unanimity operation and has a weakly connected component of algebraic length* 1*, then this component contains a loop.* □

Theorem 2.2. *A locally finite variety $\mathcal{V}$ omits type* **1** *iff there exists a* 4*-ary term t such that*

(1) $\mathcal{V} \models t(x, x, x, x) \approx x$ *and*
(2) $\mathcal{V} \models t(x, y, z, y) \approx t(y, z, x, x)$.

Proof. ($\Leftarrow$) Item (2) is not satisfied by any projection operation, hence together with (1) implies that t is a Taylor term. This is enough to establish that $\mathcal{V}$ omits type **1** (see [7], Lemma 9.4 and Theorem 9.6).

($\Rightarrow$) 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. $\mathcal{W}$ is a locally finite, idempotent variety. If there is a 4-ary term t for $\mathcal{W}$ such that $\mathcal{W} \models t(x, y, z, y) \approx t(y, z, x, x)$, then there is also a term t for $\mathcal{V}$ such that $\mathcal{V} \models t(x, x, x, x) \approx x$ and $\mathcal{V} \models t(x, y, z, y) \approx t(y, z, x, x)$.

According to [9], $\mathcal{V}$ has a weak near-unanimity term $w(x_1, x_2, \dots, x_n)$, and since w is an idempotent $\mathcal{V}$-term, then $\mathcal{W}$ also has a weak near-unanimity term. Now let $\mathbf{F} = \mathbf{F}_{\mathcal{W}}(x, y, z)$ and let $\mathbf{E}$ be the subalgebra of $\mathbf{F}^2$ generated by $\{\langle x, y\rangle, \langle y, z\rangle, \langle z, x\rangle, \langle y, x\rangle\}$. Clearly, the first and second projections of $\mathbf{E}$ are both equal to all of $\mathbf{F}$, as we can get any element of F by applying the appropriate 3-variable $\mathcal{W}$-term to the first three generators of $\mathbf{E}$. Therefore, $\mathbf{E}$ is subdirect in $\mathbf{F}^2$, and this means that the digraph $G = (F, E)$ with vertex set F and edge set E is without sources and sinks. Moreover, as the weak near-unanimity term w is among the fundamental operations of the variety $\mathcal{W}$, E is invariant under the weak near-unanimity operation $w^{\mathbf{F}}$ applied coordinatewise. Finally, the (weakly connected) component of G containing x, y and z has algebraic length 1 as it contains the path $x \leftarrow y \rightarrow z \rightarrow x$. According to Lemma 2.1, the component of G containing x, y and z must contain a loop $a \to a$. So, there must exist a term operation t of $\mathcal{W}$ such that $t^{\mathbf{E}}(\langle x, y\rangle, \langle y, z\rangle, \langle z, x\rangle, \langle y, x\rangle) = \langle a, a\rangle$, which means that $t^{\mathbf{F}}(x, y, z, y) = a = t^{\mathbf{F}}(y, z, x, x)$. But this forces the 3-variable identity $t(x, y, z, y) \approx t(y, z, x, x)$ to hold in $\mathcal{W}$. □

3. Optimal strong Mal'cev conditions

If one desires to express the property of omitting type **1** in locally finite varieties via a Taylor condition that uses only 2-variable identities, a number of such conditions can be derived from the condition of Theorem 2.2. Probably the most intriguing is the following

$$t(x, x, x, x) \approx x,\ t(x, y, x, y) \approx t(y, x, x, x) \quad \text{and} \quad t(y, x, x, x) \approx t(x, x, y, y),$$

which, after reversing the order of variables $e(z, y, x, w) := t(w, x, y, z)$, becomes

$$e(x, x, x, x) \approx x \quad \text{and} \quad e(y, y, x, x) \approx e(y, x, y, x) \approx e(x, x, x, y). \tag{1}$$

If the equations above stipulated that furthermore $e(y, y, x, x) \approx x$, then this would be the 3-edge term of [3], a condition equivalent to all subuniverses of $\mathbf{A}^n$ having at most $O(n^2)$ many generators. In accordance with terminology of [9], we call the condition (1) the *weak 3-edge term*. We state the observation of this paragraph as a corollary to Theorem 2.2:

Corollary 3.1. *A locally finite variety $\mathcal{V}$ omits type* **1** *iff it has a weak* 3*-edge term.*

Proof. The paragraph preceding the statement of the corollary explains how to derive from Theorem 2.2 that a locally finite variety omitting type **1** has a weak 3-edge term. To prove the converse, just note that the equations (1) can not be satisfied by any projection. □

Corollary 3.1 may be viewed as a refinement of Theorem 2.2, in that the identities involve fewer variables. A different refinement of Theorem 2.2, involving terms of smaller arity, was observed by Miklós Maróti [8]:

Corollary 3.2. *A locally finite variety $\mathcal{V}$ omits type* **1** *iff it has ternary terms p and q such that $\mathcal{V} \models p(x, x, x) \approx x \approx q(x, x, x)$ and*

(1) $\mathcal{V} \models p(x, y, y) \approx q(y, x, x) \approx q(x, x, y)$ *and*
(2) $\mathcal{V} \models p(x, y, x) \approx q(x, y, x)$.

Proof. If $\mathcal{V}$ omits type **1**, then it has a 4-ary term t as described in Theorem 2.2. Then $p(x, y, z) := t(x, x, y, z)$ and $q(x, y, z) := t(x, y, z, z)$ satisfy the required identities. Conversely, the given identities for p and q cannot be realized by projections, so they force $\mathcal{V}$ to omit type **1**. (Argument: If p and q are projections, then the identities in (1) force q to be second projection and p to be first projection. But this is incompatible with the identity in (2).) □

Now we turn to proving that there are no idempotent, linear, strong Mal'cev conditions for omitting type **1** in strictly simpler languages than those we have given in Theorem 2.2 and Corollaries 3.1 and 3.2.

Lemma 3.3. *Let Σ be an idempotent, linear, strong Mal'cev condition in a language $\mathcal{L}$, and let Σ_0 be the set of all linear consequences of Σ that involve no function symbols of arity strictly less than 3. Either*

(1) *Σ and Σ_0 are realized by the same varieties, or*

(2) *$\Sigma \models q(x, y) \approx q(y, x)$ for some binary function symbol q of $\mathcal{L}$.*

(Or both.)

Proof. We may assume that the symbols of $\mathcal{L}$ are only those that appear in Σ. No 0-ary function symbols appear in Σ, since Σ is idempotent, so the language $\mathcal{L}$ has finitely many function symbols, all of arity at least 1.

Let $\mathcal{F}_0$ be the set of function symbols of arity at least 3 and let f_i, $0 \leq i < m$, be the remaining binary and unary function symbols. Let $\mathcal{F}_{i+1} = \mathcal{F}_i \cup \{f_i\}$. Let Σ_i be the linear consequences of Σ that involve only the function symbols in $\mathcal{F}_i$. (When $i = 0$ we are repeating the original definition of Σ_0.) Then

$$\Sigma_0 \subseteq \Sigma_1 \subseteq \cdots \subseteq \Sigma_m,$$

and Σ_m is logically equivalent to Σ. If Σ_0 is inconsistent (i. e. if $\Sigma_0 \models x \approx y$), then so are the supersets Σ_i. Conversely, if Σ_m ($\equiv \Sigma$) is inconsistent, then the definitions of the Σ_i guarantee that $x \approx y$ belongs to each of them. This situation, where Σ and all Σ_i are inconsistent, is a case where Item (1) of the theorem holds, so for the rest of the proof we assume that all Σ's are consistent.

Any realization of Σ is a realization of Σ_m and hence also a realization of each of the subsets Σ_i. The proof of the theorem will be effected by showing that, conversely, if the obstruction $\Sigma \models q(x, y) \approx q(y, x)$ is avoided, then for each i any realization of Σ_i in a variety $\mathcal{V}$ can be lifted to a realization of Σ_{i+1} in the same variety. Consequently, if the obstruction is avoided, then any realization of Σ_0 can be lifted to a realization of Σ_m ($\equiv \Sigma$). There are only two essentially different cases: one where $\mathcal{F}_{i+1}$ is obtained from $\mathcal{F}_i$ by adding a unary function symbol f_i and the other where it is obtained by adding a binary function symbol f_i.

We alert the reader that in order to lift a realization of Σ_i to a realization of Σ_{i+1} we shall first define a subset $\Sigma_{i+1}^\circ \subseteq \Sigma_{i+1}$ that is formally simpler yet logically equivalent to the whole set Σ_{i+1}. We shall lift a given realization of Σ_i to a realization of the subset Σ_{i+1}°, then use the logical equivalence of Σ_{i+1}° with Σ_{i+1} to declare the lift to Σ_{i+1}° to be a realization of Σ_{i+1}.

Case 1. ($\mathcal{F}_{i+1} = \mathcal{F}_i \cup \{f_i\}$ where $f_i =: p$ is unary.) Each realization of Σ_i in a variety $\mathcal{V}$ can be lifted to a realization of Σ_{i+1} in $\mathcal{V}$.

Any identity in $\Sigma_{i+1} - \Sigma_i$ must involve p on one side or on both sides, hence must be of one of the forms:

(1.a) $p(v) \approx t$ where t is a linear term whose function symbols are from $\mathcal{F}_i$ or

(1.b) $p(u) \approx p(v)$ where u, v are (not necessarily distinct) variables.

Moreover, since the original set Σ was idempotent, it follows that each Σ_j is also idempotent, so $p(v) \approx v$ must be in Σ_{i+1} for each variable v.

Suppose that some identity $p(v) \approx t$ of the form (1.a) belongs to $\Sigma_{i+1} - \Sigma_i$, so that both $p(v) \approx t$ and $p(v) \approx v$ belong to $\Sigma_{i+1} - \Sigma_i$. Since the set $\{p(v) \approx t, p(v) \approx v\}$ is logically equivalent to $\{v \approx t, p(v) \approx v\}$, and $v \approx t$ already belongs to Σ_i (since $\Sigma \models \Sigma_{i+1} \models v \approx t$ and all symbols of $v \approx t$ belong to $\mathcal{F}_i$), we may delete the identity $p(v) \approx t$ from the set Σ_{i+1} as part of the process of creating the subset Σ°_{i+1}. After such deletion Σ°_{i+1} remains logically equivalent to Σ_{i+1}.

Now suppose that some (1.b)-type identity $p(u) \approx p(v)$ belongs to $\Sigma_{i+1} - \Sigma_i$. Since Σ_{i+1} is consistent and contains both $p(u) \approx u$ and $p(v) \approx v$ it follows that $u = v$, and therefore that our (1.b)-type identity is the tautology $p(u) \approx p(u)$. This shows that we may further delete all (1.b)-type identities from Σ_{i+1} as part of the process of creating Σ°_{i+1}. After such deletion Σ°_{i+1} remains logically equivalent to Σ_{i+1}.

We have deleted some identities in $\Sigma_{i+1} - \Sigma_i$ from Σ_{i+1} to create Σ°_{i+1}. In the end, $\Sigma^\circ_{i+1} - \Sigma_i$ consists only of identities of the form $p(v) \approx v$, and these are the only identities in Σ°_{i+1} involving p. Hence the assignment $p(x_1) \mapsto x_1^{\mathcal{V}}$ extends the given realization of Σ_i to a realization of Σ°_{i+1}. Since Σ°_{i+1} is logically equivalent to Σ_{i+1}, this is also a realization of Σ_{i+1}.

Case 2. ($\mathcal{F}_{i+1} = \mathcal{F}_i \cup \{f_i\}$ where $f_i =: q$ is binary.) If $\Sigma_{i+1} \not\models q(x, y) \approx q(y, x)$, then each realization of Σ_i in a variety $\mathcal{V}$ can be lifted to a realization of Σ_{i+1} in $\mathcal{V}$.

Any identity in $\Sigma_{i+1} - \Sigma_i$ must involve q, hence must be of one of the forms:

(2.a) $q(u, v) \approx t$ where t is a linear term whose function symbols are from $\mathcal{F}_i$.
(2.b) $q(u, v) \approx q(w, x)$ where u, v, w, x are (not necessarily distinct) variables.

Moreover, $q(v, v) \approx v$ must be in Σ_{i+1} for each variable v, since each Σ_j is idempotent.

Using the same type of argument we used in Case 1, as part of the creation of Σ°_{i+1} we may delete from Σ_{i+1} all identities of the form $q(v, v) \approx t$ where t is a linear term that is not a variable. (By "the same type of argument" we mean an argument that involves the logical equivalence of $\{q(v, v) \approx t, q(v, v) \approx v\}$ and $\{v \approx t, q(v, v) \approx v\}$.) After this adjustment, the identities remaining in $\Sigma^\circ_{i+1} - \Sigma_i$ *at the present moment* have the forms: $q(v, v) \approx v$ for all variables v, along with

(3.a) $q(u, v) \approx t$ where $u \neq v$ and t is a linear term whose function symbols are from $\mathcal{F}_i$.
(3.b) $q(u, v) \approx q(w, x)$ where $u \neq v$ and $w \neq x$

We discuss (3.b) first. We delete tautologies of the form $q(u, v) \approx q(u, v)$ during the (continued) creation of Σ°_{i+1}. We do not have identities like $q(u, v) \approx$

$q(v,u)$ in Σ_{i+1} by the assumption of Case 2. So, if there are any more (3.b)-type identities $q(u,v) \approx q(w,x)$ in Σ_{i+1}, then some variable appears on only one side of the equation. If, for example, one such variable is w, then by variable replacement w/x we find that $q(u,v) \approx q(x,x)$ is in Σ_{i+1}. But then $q(u,v) \approx x$ and $x \approx q(w,x)$ belong to Σ_{i+1}. The consistency of Σ_{i+1} forces $v = x$, hence the set $\{q(u,v) \approx q(w,x), q(x,x) \approx x\}$ is logically equivalent to $\{q(w,x) \approx x, q(x,x) \approx x\}$. We may therefore delete the only remaining (3.b)-type identities $q(u,v) \approx q(w,x)$ from Σ_{i+1} during the creation of Σ°_{i+1}. (The identities $q(w,x) \approx x$, and $q(x,x) \approx x$ remain in Σ°_{i+1}.)

We are now at a point where the only identities left in $\Sigma^{\circ}_{i+1} - \Sigma_i$ are those expressing the idempotence of q and possibly some (3.a)-type identities $q(u,v) \approx t$ where $u \neq v$ and t is a linear term whose function symbols are from $\mathcal{F}_i$.

If $\Sigma^{\circ}_{i+1} - \Sigma_i$ fails to contain even one identity of form (3.a), $q(u,v) \approx t$, then we interpret q in $\mathcal{V}$ so that it is binary first projection, i. e. $q(x_1,x_2) \mapsto x_1^{\mathcal{V}}$.

Assume that $\Sigma^{\circ}_{i+1} - \Sigma_i$ contains at least one (3.a)-type identity, $q(u,v) \approx t$, where t is a variable or $t = f(v_1,\ldots,v_n)$ is a function symbol $f \in \mathcal{F}_i$ applied to variables v_i. By variable replacement we may even assume that the distinct variables u and v are actually x_1 and x_2 (so the identity is $q(x_1,x_2) \approx t$). Replacing each other variable appearing in t with x_1 or x_2 arbitrarily we obtain t', a term that is a variable in $\{x_1,x_2\}$ or is f applied to variables in $\{x_1,x_2\}$. Since $q(x_1,x_2) \approx t'$ is a substitution instance of $q(x_1,x_2) \approx t$, it also belongs to Σ_{i+1}, and in fact to Σ°_{i+1}. The assignment $q(x_1,x_2) \mapsto (t')^{\mathcal{V}}$ is an unambiguous interpretation of q, since if t' and s' are linear terms whose function symbols are from $\mathcal{F}_i$, whose variables are from $\{x_1,x_2\}$, and where $q(x_1,x_2) \approx t'$ and $q(x_2,x_2) \approx s'$ are both in $\Sigma^{\circ}_{i+1} \subseteq \Sigma_{i+1}$, then $t' \approx s'$ is in Σ_i, so $(t')^{\mathcal{V}} = (s')^{\mathcal{V}}$.

It remains to show that this assignment of the elementary term $q(x_1,x_2)$ to a $\mathcal{V}$-equivalence class of binary $\mathcal{V}$-terms defines a realization of Σ°_{i+1}. If $\Sigma^{\circ}_{i+1} - \Sigma_i$ contains no identity of form (3.a) or (3.b) (so it only contains those of the form $q(v,v) \approx v$), then the only restriction on the interpretation of the function symbol q is that it be idempotent. We have interpreted q as binary first projection in this case, so the restriction is satisfied. Otherwise $\Sigma^{\circ}_{i+1} - \Sigma_i$ contains some (3.a)-type identities $q(u,v) \approx t$, no (3.b)-type identities, and identities $q(v,v) \approx v$ expressing idempotence. We have interpreted q unambiguously so that the identities of the form $q(u,v) \approx t$ will hold. The identities expressing the idempotence of q are consequences of these. (If t is a variable, this is immediately clear, while if $t = f(v_1,\ldots,v_n)$ then the idempotence of f is one of the identities in Σ_i, so the identities of the form $q(v,v) \approx v$ also hold.)

This shows that any given realization of Σ_i can be lifted to Σ°_{i+1}. Such a lift is also a realization of the logically equivalent set Σ_{i+1}. □

Theorem 3.4. *If Σ is an idempotent, linear, strong Mal'cev condition defining (within the class of locally finite varieties) the class of varieties omitting type* $\mathbf{1}$*, then Σ involves at least one function symbol of arity at least* 4 *or at least two functions symbols of arity at least* 3*.*

Proof. We prove that if (i) Σ is an idempotent, linear, strong Mal'cev condition, (ii) Σ does not involve at least one function symbol of arity at least 4 or at least two functions symbols of arity at least 3, and (iii) Σ is realizable in the 2-element semilattice, in the 2-element group, and in the 3-element group, then (iv) Σ must also be realizable by the 2-element set. (I. e., Σ can be modeled by projections.)

Assume not. Our first step will be to simplify Σ.

Simplification 1. If Σ cannot be modeled by projections, then the set consisting of substitution instances of the identities from Σ in which all variables are replaced by variables from the set $\{x, y\}$, will also be a set of identities involving the same function symbols that cannot be modeled by projections. Therefore we assume that Σ uses only the variables x and y.

Simplification 2. Σ cannot contain any identity of the form $q(x, y) \approx q(y, x)$, since this identity is not true for any idempotent binary term operation of the 2-element group. Thus, according to Lemma 3.3, we may assume that all function symbols in Σ have arity at least 3. (The lemma allows us to replace Σ with Σ_0 without altering the class of realizations.) If there is more than one function symbol of arity at least 3, then the theorem is proved. If there is at least one function symbol whose arity is at least 4, then the theorem is proved. Hence we may assume that Σ involves exactly one function symbol r and that its arity is 3.

Simplification 3. In any identity of Σ other than one asserting the idempotence of r we may replace any occurrence of $r(v, v, v)$ with v. Therefore we assume that $r(v, v, v)$ only occurs in the identity $r(v, v, v) \approx v$.

Now the proof begins for simplified Σ. Our first step will be to show that Σ contains no identity of the form $r(v_1, v_2, v_3) \approx v$ other than one asserting the idempotence of r. (Without loss of generality, we assume $v = x$, in which case we are claiming that $r(x, x, x) \approx x$ is the only identity in Σ of the specified form.) Since Σ is realized by the 2-element semilattice $\mathbf{S}$, any such identity in Σ would imply that $r^{\mathbf{S}}$ does not depend on the positions where $v_i = y$. The term operation $r^{\mathbf{S}}$ cannot depend on only one variable (as it would be a projection and the set of identities Σ is not satisfied by any projection), hence at least two of the v_i must be x. On the other hand, at least one of the v_i must be y, since we have assumed that $r(v_1, v_2, v_3) \approx x$ is different from $r(x, x, x) \approx x$. This shows that (after permuting variables if necessary) the identity $r(v_1, v_2, v_3) \approx x$ may be assumed to be $r(x, x, y) \approx x$. The realization of Σ in the 2-element group $\mathbf{G}$ requires $r^{\mathbf{G}}(x, y, z) = \alpha x + \beta y + \gamma z$ for some choice of coefficients $\alpha, \beta, \gamma \in \mathbb{Z}_2$ so that $\mathbf{G} \models r^{\mathbf{G}}(x, x, y) = x = r^{\mathbf{G}}(x, x, x)$.

These identities force $\gamma = 0$ and $\{\alpha, \beta\} = \{0, 1\}$, which in turn forces $r^{\mathbf{G}}$ to be a projection, a contradiction.

This means that all identities in Σ, other than the idempotence of r, are of the form $r(u_1, u_2, u_3) \approx r(v_1, v_2, v_3)$ with $\{u_1, u_2, u_3\} = \{x, y\} = \{v_1, v_2, v_3\}$. Moreover, the parity of the number of x's is the same among u_1, u_2, u_3 as among v_1, v_2, v_3, since that is a feature of the linear identities satisfied by the unique ternary, idempotent, non-projection term operation $r(x, y, z) = x+y+z$ of the 2-element group.

The remaining case is when (without loss of generality) each identity in Σ that is not asserting the idempotence of r has exactly one y on each side. For every position i there is an identity $r(\text{variables}) \approx r(\text{variables}) \in \Sigma$ where x appears in the ith position on the left and y appears in the ith position on the right. This forces Σ to syntactically imply $r(x, x, y) \approx r(x, y, x) \approx r(y, x, x)$; that is, Σ implies the assertion that r is a 3-ary weak near-unanimity term. But now Σ cannot be realized in the 3-element group. $\square$

References

[1] L. Barto and M. Kozik, *Absorbing Subalgebras, Cyclic Terms, and the Constraint Satisfaction Problem*, Log. Meth. in Comp. Sci. **8**, no. 1 (2012), paper 7.

[2] L. Barto, M. Kozik and T. Niven, *The CSP dichotomy holds for digraphs with no sources and no sinks (a positive answer to a conjecture of Bang-Jensen and Hell)*, SIAM J. on Comput. **38**, no. 5 (2009), 1782–1802.

[3] J. Berman, P. Idziak, P. Marković, R. McKenzie, M. Valeriote and R. Willard, *Varieties with few subalgebras of powers*, Trans. Amer. Math. Soc. **362**, no. 3 (2010), 1445–1473.

[4] S. Burris and H. P. Sankappanavar, *A course in universal algebra,* Graduate Texts in Mathematics **78**, Springer-Verlag, New York-Berlin, 1981.

[5] P. Hell and J. Nešetřil, *On the complexity of H-colouring*, J. Combin. Theory B **48** (1990), 92–100.

[6] P. Hell, J. Nešetřil and X. Zhu, *Duality and Polynomial Testing of Tree Homomorphisms*, Trans. Amer. Math. Soc. **348**, no. 4 (1996), 1281–1297.

[7] D. Hobby and R. McKenzie, *The Structure of Finite Algebras,* Mem. Amer. Math. Soc. **76**, American Mathematical Society, Providence, RI, 1988.

[8] M. Maróti, unpublished, see http://www.karlin.mff.cuni.cz/ ical/presentations/Valeriote.pdf, page 17.

[9] M. Maróti and R. McKenzie, *Existence theorems for weakly symmetric operations,* Algebra Universalis **59**, no. 3-4 (2008), 463–489.

[10] R. McKenzie, G. McNulty and W. Taylor, *Algebras, Lattices, Varieties. Vol. I,* The Wadsworth & Brooks/Cole Mathematics Series, Wadsworth & Brooks/Cole Advanced Books & Software, Monterey, CA, 1987.

[11] M. Siggers, *A Strong Mal'cev Condition for Varieties Omitting the Unary Type*, Algebra Universalis **64**, no. 1 (2010), 15–20.

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

Keith Kearnes

Department of Mathematics, University of Colorado, Boulder, CO 80309-0395 USA
e-mail: `kearnes@euclid.colorado.edu`

Petar Marković

Department of Mathematics and Informatics, University of Novi Sad, Serbia
e-mail: pera@dmi.uns.ac.rs

Ralph McKenzie
Department of Mathematics, Vanderbilt University, Nashville, TN 37240, USA
e-mail: ralph.n.mckenzie@vanderbilt.edu