Oligomorphic Clones

Manuel Bodirsky and Hubie Chen

Abstract. A permutation group on a countably infinite domain is called *oligomorphic* if it has finitely many orbits of finitary tuples. We define a clone on a countable domain to be *oligomorphic* if its set of permutations forms an oligomorphic permutation group. There is a close relationship to *ω-categorical structures*, i.e., countably infinite structures with a first-order theory that has only one countable model, up to isomorphism. Every locally closed oligomorphic permutation group is the automorphism group of an ω-categorical structure, and conversely, the canonical structure of an oligomorphic permutation group is an ω-categorical structure that contains all first-order definable relations. There is a similar Galois connection between locally closed oligomorphic clones and ω-categorical structures containing all *primitive positive* definable relations.

In this article we generalise some fundamental theorems of universal algebra from clones over a finite domain to oligomorphic clones. First, we define *minimal* oligomorphic clones, and present equivalent characterisations of minimality, and then generalise Rosenberg's five types classification to minimal oligomorphic clones. We also present a generalisation of the theorem of Baker and Pixley to oligomorphic clones.

This is a updated post-print version of an article with the same title that appeared in *Algebra Universalis*, 57(1): 109-125, 2007.

1. Introduction

A *clone* is a set of finitary operations that is closed under compositions, and contains all projections. A clone is called *locally closed*, if it contains all operations that are interpolated on finite sets by the functions in the clone; for formal definitions see Section 2. A permutation group on a countably infinite set is called *oligomophic* [11] if it has finitely many orbits of n-tuples, for all $n \geq 1$. We define a clone to be *oligomophic* if its permutations form an oligomorphic permutation group.

Oligomorphic clones are closely related to *ω-categorical structures* (such structures are also called *countably categorical* or *$\aleph_0$-categorical*). A countably infinite structure is called *ω-categorical* if all countable models of its first-order theory are isomorphic. As we will see later, a locally closed clone is oligomorphic if and only if it is the set of *polymorphisms* of an ω-categorical relational structure Γ (an n-ary *polymorphism* of Γ is a homomorphism from the direct product Γ^n to Γ). Conversely, a countably infinite structure is the canonical structure of an oligomorphic clone if and only if it is ω-categorical and contains all primitive positive definable relations.

Key words and phrases: ω-categoricity, preservation theorems, oligomorphic permutation groups, clones on infinite domains, minimal locally closed clones.

Oligomorphic permutation groups and ω-categorical structures have attracted a lot of interest in various areas: in model theory (see [19], and, more specifically, [22]), in the study of infinite permutation groups (see e.g. [11, 2]), infinite combinatorics (e.g. [12, 13, 16]), and, recently, in constraint satisfaction [3, 6]. This is partly due to many equivalent characterisations of ω-categoricity, some of which are given in Section 3. Another reason might be the wealth of concrete examples, but also the wealth of general construction principles for ω-categorical structures, for example amalgamation or descriptions via their automorphism groups. There are also various methods to derive new ω-categorical structures from existing ones, e.g. by first-order interpretations (see [19]), by taking cores [4], or by expansions with constants.

The motivation of this article comes from the following research question.

Question. *Which parts of universal algebra can be generalised from clones over a finite domain to oligomorphic clones?*

Clones over a finite domain. The set of all clones on a finite domain D forms a lattice with respect to inclusion. The lattice has a top element, which is the set of all finitary operations, and a bottom element, which is the set of all projections. The clones over a Boolean domain were classified by Post [25]. For domains with three elements, the clone lattice is already very complicated (there is a continuum of such clones, see [24]). However, the lattice has a finite number of atoms and co-atoms for all finite domains. The corresponding clones are called *minimal* and *maximal*, respectively. The maximal clones over a finite domain were classified by Rosenberg [26]. The minimal clones were classified for domains of size three by Czákány [14], but the classification of the minimal clones over a larger finite domain is considered to be very difficult. However, Rosenberg showed that a minimal clone has one out of five types that are well-understood in many cases [27, 23].

Recently, the theory of minimal clones found applications in the complexity classification of constraint satisfaction problems [21, 9, 8]. Another result from universal algebra with applications in constraint satisfaction is the theorem of Baker and Pixley [1, 20]. In one of its versions (Corollary 5.1 in [1]) it can be seen as the characterisation of a quantifier elimination property by the existence of a polymorphism that satisfies certain identities. The corresponding constraint satisfaction problems are always tractable [20, 17]. See [10] for an introduction of the algebraic approach to constraint satisfaction.

Oligomorphic clones. The set of locally closed oligomorphic clones form a join-semilattice. Again, its top element is the set of all finitary operations. However, the locally closed oligomorphic clones do not form a lattice. Moreover, the semilattice contains no minimal elements: for every locally closed oligomorphic clone there is a smaller locally closed oligomorphic clone. To see this, observe that every oligomorphic permutation group contains an infinite orbit. The stabiliser group for an element from this orbit is smaller, and again oligomorphic and locally closed.

We call an operation of an oligomorphic clone F *elementary*, if it is locally generated by the projections and the permutations in F (the terminology is motivated by the model-theoretic notion of *elementary embeddings* – the connection becomes clear in Section 4). An oligomorphic clone is called *elementary* if it is locally closed and contains only elementary operations.

Minimal Oligomorphic Clones. In the semilattice of locally closed oligomorphic clones, we are now interested in the minimal elements above the elementary clones. These are the non-elementary locally closed oligomorphic clones such that every locally closed oligomorphic subclone is elementary. It turns out that an oligomorphic clone F is minimal if and only if every non-elementary operation in F locally generates F – this is a fundamental property for minimal clones over a finite domain. We prove various other equivalent characterisations of minimal oligomorphic clones, and then show that every oligomorphic clone contains a minimal oligomorphic clone (as in the case of finite domains).

Some properties of minimal clones over a finite domain fail for minimal oligomorphic clones. The most important one is *idempotency*: every at least binary operation f in a minimal oligomorphic clone over a finite domain satisfies $f(x, \ldots, x) = x$. However, every at least binary operation f in a minimal oligomorphic clone has a related property: the (elementary) unary function g defined by $g(x) = f(x, \ldots, x)$ preserves all first-order definable relations in the canonical structure of F. Such operations f we call *oligopotent*.

Results. We present various characterisations of elementary and of minimal oligomorphic clones, and show that if F is the polymorphism clone of an ω-categorical structure with a finite relational signature, then F contains a minimal oligomorphic clone. One of the main results in this article is a generalisation of Rosenberg's theorem; we distinguish four different types of minimal oligomorphic clones. Roughly speaking, the only property we lose in the statement of the classification is idempotency, but we always get oligopotency instead. Moreover, in the oligomorphic case we can exclude the analog for one of the Rosenberg types.

We also prove a generalisation of the theorem of Baker and Pixley in a model-theoretic formulation. The previously mentioned result in [1] does not only apply to finite algebras, but also to locally finite varieties (Theorem 5.2 in [1]) – however, oligomorphic clones are not locally finite. However, as we will see in Section 7, their result also generalises to oligomorphic clones, and again leads to tractable constraint satisfaction problems [5].

2. Clones

Let D be a countable set, and let O be the set of *finitary operations* on D, i.e., functions from D^k to D for finite k. An operation $f \in O$ is a *projection* (or a *trivial polymorphism*) if for all n-tuples, $f(x_1, \ldots, x_n) = x_i$ for some fixed $i \in \{1, \ldots, n\}$. The *composition* $f(g_1, \ldots, g_k)$ of a k-ary operation f and k operations $g_1, \ldots, g_k$ of arity n is an n-ary operation defined by

$$f(g_1, \dots, g_k)(x_1, \dots, x_n) := f\big(g_1(x_1, \dots, x_n), \dots, g_k(x_1, \dots, x_n)\big) \enspace .$$

A *clone* F is a set of operations from O that is closed under composition and that contains all projections. We write D_F for the *domain* D of the clone F. For a set of operations F from O we write $\langle F \rangle$ for the smallest clone containing all operations in F (which is called the clone *generated* by F). If the set of operations of a clone A is a (proper) subset of the operations of another clone B over the same domain, we say that A is a *(proper) subclone* of B.

We say that a k-ary operation $f \in O$ is *locally generated* by a subset F of O if for every finite subset A of D there is some k-ary operation $g \in \langle F \rangle$ such that $f(\overline{a}) = g(\overline{a})$ for every $\overline{a} \in A^k$. The smallest clone that contains the operations locally generated by F is called the *local closure* of F, and denoted by $\langle F \rangle_{\text{loc}}$. Clearly, $F \subseteq \langle F \rangle_{\text{loc}}$. Note that the permutations in the clone might locally generate some non-surjective unary operations. As an example, consider the set of all permutations on a set D. This set locally generates all injective unary operations on D.

We say that a k-ary operation f *depends* on an argument i iff there is no $k{-}1$-ary operation f' such that $f(x_1, \dots, x_k) = f'(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_k)$. We can equivalently characterise k-ary operations that depend on the i-th argument by requiring that there are elements $x_1, \dots, x_k$ and x_i' such that $f(x_1, \dots, x_k) \neq f(x_1, \dots, x_{i-1}, x_i', x_{i+1}, \dots, x_k)$.

Many important properties of operations in a clone can be specified with identities that are satisfied by the operations. We list some of those fundamental properties below, where the free variables in the identities are understood as universally quantified. Let $\hat{f}$ denote the operation defined by $\hat{f}(x) = f(x, \dots, x)$. A k-ary operation f is

- a *projection (to the i-th argument)* iff there is an $i \in \{1, \dots, k\}$ such that $f(x_1, \dots, x_k) = x_i$;
- *conservative* iff $f(x_1, \dots, x_k) \in \{x_1, \dots, x_k\}$;
- *idempotent* iff $\hat{f}(x) = x$;
- *essentially unary* iff there is an $i \in \{1, \dots, k\}$ and a unary operation f_0 such that $f(x_1, \dots, x_k) = f_0(x_i)$. In other words, f depends on at most one argument.
- a *quasi near-unanimity operation* (short, *qnu-operation*) iff $\hat{f}(x) = f(x, \dots, x, y) = \cdots = f(x, \dots, x, y, x, \dots, x) = \cdots = f(y, x, \dots, x)$;
- a ternary *quasi majority operation* iff f is a quasi near-unanimity operation and $k = 3$;
- a *quasi Maltsev operation* iff f is ternary and $f(y, x, x) = f(x, x, y) = \hat{f}(y)$;
- a *quasi semiprojection* iff there is an essentially unary operation g such that $f(x_1, \dots, x_k) = g(x_1, \dots, x_k)$ whenever $|\{x_1, \dots, x_k\}| < k$. If g depends on the i-th argument, we say that f is a *quasi semiprojection to the i-th argument.*

An operation f is called *near-unanimity, majority, minority*, or *semiprojection*, if f is an *idempotent* quasi near-unanimity, quasi majority, quasi minority, or quasi semiprojection, respectively. Except for the notions with the qualifier *quasi*, all notions in this section are standard [28].

A clone is called *projective, conservative, idempotent, or essentially unary*, if all operations are projections, conservative, idempotent, or essentially unary, respectively.

3. Oligomorphic Clones

The *permutations* of a clone F are the bijective unary operations in F.

Definition 1. *A permutation group is called* oligomorphic*, if it contains only finitely many orbits of n-tuples, for all $n \geq 0$. A clone is called* oligomorphic*, if its permutations form an oligomorphic permutation group.*

Let F be a clone with domain D. We say that an n-ary operation $f \in O$ *preserves* a k-ary relation $R \subseteq D^k$ iff for all tuples $(c_1^1, \dots, c_1^k), \dots, (c_n^1, \dots, c_n^k)$ in R the k-tuple $(f(c_1^1, \dots, c_n^1), \dots, f(c_1^k, \dots, c_n^k))$ is also in R. An operation that preserves all relations of a relational structure Γ with domain D is called a *polymorphism* of Γ. We write $\mathrm{Pol}(\Gamma)$ for the set of all finitary polymorphisms. Observe that $\mathrm{Pol}(\Gamma)$ is always a locally closed clone with domain D. Unary polymorphisms are the *endomorphisms*, and polymorphisms that also preserve the complements of all relations of the structure are called *strong*. The bijective strong endomorphisms are the *automorphisms* of Γ. Injective strong endomorphisms are called *embeddings*. Endomorphisms that are not embeddings are called *strict*.

We described how to associate a clone to a relational structure. We now describe how to associate a relational structure to a clone.

Definition 2 (Canonical Structure)**.** *A relation $R \subseteq D^m$* is invariant under *F, if every $f \in F$ preserves R. The relational structure on the domain D that contains all relations that are invariant under F is called the* canonical structure *for F, and denoted by $Inv(F)$.*

It is known that an ω-categorical structure admits uncountably many polymorphisms (in fact, already uncountably many automorphisms). But note that the clone of polymorphisms of an arbitrary countable structure is locally generated by a countable number of polymorphisms: choose for each finite set B and all potential images of tuples from B a polymorphism if there exists such a polymorphism. The following is a well-known fact (see e.g. [28]).

Proposition 3. *A set $F \subseteq O$ of operations is locally closed if and only if F is the set of polymorphisms of Γ for some relational structure Γ.*

In the following we use classical concepts from logic and model-theory; see e.g. [19].

Definition 4. *A countable relational structure Γ is called ω-categorical, if all countable models of the first-order theory of Γ are isomorphic to Γ.*

The following is due to Ryll-Nardzewski, Engeler, and Svenonius (see [19]).

Theorem 5. *The following properties of a structure Γ are equivalent.*

- Γ *is ω-categorical;*
- *The automorphism group of Γ is oligomorphic;*

A τ-formula is *primitive positive* if it is of the form $\exists \overline{x}.\psi_1 \wedge \cdots \wedge \psi_n$ where $\psi_1, \ldots, \psi_n$ are atomic with relation symbols from $\tau \cup \{=\}$. It is a classical result that a relation can be defined by a primitive positive formula in a finite structure Γ if and and only if it is preserved by the polymorphisms of Γ [7]. In [6] it is proven that also on ω-categorical structures Γ a relation can be defined by a primitive positive formula if and and only if it is preserved by the polymorphisms of Γ. This is equivalent to the following.

Theorem 6. *A relational structure is the canonical structure (see Definition 2) for an oligomorphic clone F if and only if it is ω-categorical and contains all primitive positive definable relations.*

An *expansion* of a τ-structure Γ is a structure with a larger signature τ' obtained from Γ by adding a relation for each relation symbol from $\tau' \setminus \tau$. The expansion is called *proper* if the inclusion $\tau \subseteq \tau'$ is strict.

Cores. The concept of a core of a finite relational structure is central in the literature on graph and structure homomorphisms (e.g., see [18]), and frequently used in constraint satisfaction. We need the following generalisation for countable relational structures.

Definition 7. *A relational structure Γ is a* core*, if every endomorphism is an embedding.*

An example of a core structure is $(\mathbb{Q}, <)$, the linear order defined on the set of rational numbers. Note that this structure contains many endomorphisms that are not automorphisms. This is different from cores on finite domains, where all endomorphisms are also surjective. An example of a structure that is not a core is the random graph, see e.g. [19, 11]. However, the random graph is *homomorphically equivalent* to an infinite complete subgraph, which is a core. Two structures Γ and Δ are called *homomorphically equivalent* if there is a homomorphism from A to B and from B to A.

An ω-categorical structure Γ is called *model-complete* if every embedding of Γ into Γ preserves all first-order formulas. Embeddings with this property are called *elementary* in model theory. We later need the following [4].

Theorem 8. *Every ω-categorical structure Γ is homomorphically equivalent to a model-complete core Δ, which is unique up to isomorphism. The core Δ is ω-categorical or finite, and the orbits of n-tuples in Γ are primitive positive definable, for all $n \geq 1$.*

4. Minimal Oligomorphic Clones

In this section, we define and study the notion of *minimal* oligomorphic clones. Recall that an embedding between two structures is called *elementary* if it preserves all first-order formulas. One can, similarly, define an operation of an oligomorphic clone F to be elementary if it preserves all first-order definable relations of the canonical structure of F (see Definition 2). We will take the following to be our official definition of an elementary operation.

Definition 9. *An operation of an oligomorphic clone F is called* elementary*, if it is locally generated by the projections and permutations of F. An oligomorphic clone is called* elementary *if it only contains elementary operations.*

Proposition 3 will show that an operation of an oligomorphic clone F is elementary if and only if it preserves all first-order definable relations in the canonical structure of F.

Definition 10. *A non-elementary oligomorphic clone F is called* minimal *if every locally closed oligomorphic proper subclone of F is elementary.*

We call the signature of the canonical structure of a minimal clone *maximal*. Note that an ω-categorical structure Γ has a maximal signature if and only if in every proper expansion of Γ every first-order definable relation has a primitive positive definition.

To deal with elementary and minimal oligomorphic clones, the following characterisation of the essentially unary clones will be useful. The characterisation does not require that the clone is oligomorphic, but holds for general clones. We need the following special relation P_4 that is defined on a domain D by $P_4 := \{(a,b,c,d) \mid a = b \text{ or } c = d;\ a,b,c,d \in D\}$.

Proposition 11. *Let F be a clone on the (finite or infinite) domain D. Then the following are equivalent.*

1. *F is essentially unary*
2. *The relation P_4 is contained in $Inv(F)$*
3. *The projections and the unary operations in F (i.e., the endomorphisms of $Inv(F)$) generate F*

Proof. The equivalence of (1) and (2) follows from [Lemma 1.3.1 in [24], page 56]. To give intuition about the expressive power of P_4, we repeat the argument. Clearly every essentially unary operation preserves P_4. Suppose a k-ary function f is not essentially unary, but depends on the i-th and j-th argument, $1 \leq i \neq j \leq k$. Hence there exist tuples $a_1, b_1, a_2, b_2 \in D_F^k$ where a_1, b_1 and a_2, b_2 only differ at the entries i and j, respectively, such that $f(a_1) \neq f(b_1)$ and $f(a_2) \neq f(b_2)$. Since $(a_1(l), b_1(l), a_2(l), b_2(l)) \in P_4$ for all $l \leq k$, but $(f(a_1), f(b_1), f(a_2), f(b_2)) \notin P_4$, we conclude that $P_4 \notin Inv(F)$.

(3) implies (2): Clearly, unary operations and projections preserve the relation P_4, which is therefore contained in $\text{Inv}(F)$.

(1) implies (3): By definition, every essentially unary operation $f \in F$ is composed out of a projection and a unary operation. $\square$

Proposition 12. *Let F be an oligomorphic clone, and $\Gamma := Inv(F)$ its canonical structure. Then the following are equivalent:*

(1) *F is elementary, i.e., all operations in F are locally generated by the projections and permutations of F*
(2) *All operations in F preserve all first-order definable relations in Γ*
(3) *Γ contains all first-order definable relations*
(4) *Γ is a model-complete core and contains the relation P_4*

Proof. Since projections and permutations preserve all first-order definable relations in Γ, and by Proposition 3, (1) implies (2). The implication from (2) to (3) is by definition of the canonical structure.

(3) implies (4). The relation P_4 clearly is first-order definable and therefore in the signature of Γ. Since this is also the case for the inequality relation and all negated atoms, every endomorphism of Γ preserves all first-order definable relations and therefore Γ is a model-complete core.

To show that (4) implies (1), first apply Proposition 11 to see that all polymorphisms of Γ are essentially unary. Since Γ is a core, Theorem 8 implies that all endomorphisms of Γ are embeddings, and since Γ is model-complete, all endomorphisms are elementary. Therefore all polymorphisms of Γ, i.e., all operations in F, are elementary. $\square$

We will often use the different characterisations of elementary oligomorphic clones without referring to Proposition 12.

Theorem 13. *Let F be a non-elementary oligomorphic clone, and $\Gamma := Inv(F)$ its canonical structure. Then the following are equivalent:*

(1) *F is minimal, i.e., every oligomorphic proper subclone of F is elementary*
(2) *Γ is maximal, i.e., in every first-order expansion of Γ all first-order formulas have a primitive positive definition*
(3) *Every proper first-order expansion of Γ is a model-complete core with a primitive positive definition of P_4*
(4) *Every non-elementary operation in F together with the permutations in F locally generates F.*

Proof. The equivalence of (1), (2), and (3) follows from Proposition 12.

(1) implies (4). Let $f \in F$ be a non-elementary operation. Let F' be the clone that is locally generated by f together with the permutations, and suppose for contradiction that there is an operation g in $F \setminus F'$. The clone F' is a proper subclone of F, and (1) implies that it is elementary. This contradicts the assumption that f is non-elementary.

(4) implies (1). Suppose that (4) holds, and that F has a proper non-elementary subclone F'. Then F' contains a non-elementary operation f. Statement (4) implies that f together with the permutations from F locally generates F, contradicting the assumption that F' is a proper subclone. $\square$

It is well-known that every non-elementary clone F over a finite domain contains a minimal clone. For oligomorphic clones, we can show the following.

Theorem 14. *Let Γ be an ω-categorical model-complete core such that $F := Pol(\Gamma)$ be non-elementary. Then F contains a minimal oligomorphic clone.*

Proof. We have to prove that Γ has a maximal expansion Δ. Observe that the relation P_4 can not have a primitive positive definition in Γ, since otherwise Proposition 12 would imply that F is elementary. Consider an enumeration of all first-order definable relations in Γ. We construct an expansion Δ of Γ with a maximal signature as follows. We start our construction with Γ. Whenever the next relation R in the enumeration together with the expansion constructed so far does not admit a primitive positive definition of P_4, we expand Γ by R. Let Δ be the expanded structure that can be obtained from this process by Zorn's Lemma. By construction, the relation P_4 does not have a primitive positive definition in Δ. Suppose Δ does not have a maximal signature. Then there is a relation R such that the expansion of Δ with R still does not admit a primitive positive definition of all first-order definable relations. But then, R should have been added in the construction of the signature of Δ, a contradiction. $\square$

5. Oligopotent Operations

For finite relational structures that are cores it is easy to see that every polymorphism is the composition of an idempotent polymorphism and an automorphism. This does not generalise to ω-categorical structures, since idempotency is a very strong condition for polymorphisms of infinite structures. There are structures with many polymorphisms, but where the only idempotent operations are the projections. This is for instance the case for the well-known homogeneous generic triangle-free graph; see the discussion in [4]. We therefore introduce the notion of *oligopotent operations*.

Definition 15. *An operation f of an oligomorphic clone F is called* oligopotent *if $\hat{f}$ is an elementary operation in F, i.e., if $\hat{f}$ is locally generated by the permutations in F.*

A core contains only oligopotent polymorphisms. On finite subsets of the domain oligopotent operations behave like idempotent operations. We formalise this in the next proposition. Let A be a finite subset of the domain D of an operation. We say that f is *idempotent on A* if $\hat{f}(x) = x$ for all $x \in A$. We say that f is *conservative on A* if $f(x_1, \ldots, x_n) \in \{x_1, \ldots, x_n\}$ for all $x_1, \ldots, x_n \in A$.

Proposition 16. *Let f be an operation of an oligomorphic clone F. Then f is oligopotent if and only if for every finite subset A of the domain D the clone F contains a permutation h such that $h(f)$ is idempotent on A.*

Proof. First suppose that f is oligopotent, i.e., $\hat{f}$ is elementary. Let A be a finite subset of the domain, and let $\overline{a}$ be an tuple that enumerates the elements of A. Since $\hat{f}$ is locally generated by the permutations in F, there exists a permutation h in F that maps $\hat{f}(\overline{a})$ back to $\overline{a}$. Then the operation $h(f) \in F$ is idempotent on A.

Now suppose that $\hat{f}$ is not elementary. Consider the oligomorphic subclone locally generated by f and the permutations in F, and let Γ be its canonical structure. Here, $\hat{f}$ is a strict endomorphism of Γ, i.e., $\hat{f}$ is not injective or strong on some finite subset of vertices A. Then there can be no automorphism h of Γ such that $h(f)$ is idempotent on A. □

6. Minimal Operations

We generalise Rosenberg's five types theorem from clones on finite domains to oligomorphic clones.

Definition 17. *Let F be an oligomorphic clone. A* minimal operation *in F is a non-elementary operation of minimal arity from F that together with the permutations in F locally generates a minimal oligomorphic clone.*

It follows from the discussions in Section 4 that every oligomorphic clone contains such a minimal operation. We show that a minimal operation has one out of the following four types.

Theorem 18. *Let f be a minimal operation in an oligomorphic clone F. Then f is of one of the following types:*

1. *A strict endomorphism of the canonical structure of F, where $f(f)$ together with the permutations in F locally generates f.*
2. *A binary operation*
3. *A ternary quasi majority operation*
4. *A k-ary quasi semi-projection, for $k \geq 3$*

Moreover, if f is at least binary, then it is oligopotent and the canonical structure of F is a model-complete core.

Proof. First consider the case that f is unary. Since f is non-elementary, it is a strict endomorphism of the canonical structure. Also the composition of f with itself is strict, and thus non-elementary. Hence, together with the permutations in F the operation $f(f)$ locally generates a non-elementary locally closed oligomorphic clone. By minimality of f, this clone can not be a proper subclone of the clone generated by f, and therefore $f(f)$ locally generates f.

If f is not unary, we show that f is oligopotent. Suppose for contradiction that $\hat{f}(x)$ and the permutations generate a non-elementary clone. This contradicts the minimal arity requirement of Definition 17 for f. Moreover, if f is not unary, the canonical structure Γ of F is in this case a model-complete core. The reason is that if Γ has a non-elementary endomorphism e, then this endomorphism locally generates with the permutations in F a non-elementary oligomorphic clone. Clearly, all operations in this clone are essentially unary, and hence this clone does not contain the elementary operation f, contradicting the choice of f.

For ternary f we have to show that f satisfies the equations of the quasi majority or the quasi semiprojection. It is easy to check that these cases are disjoint. By minimality of f, every operation $f_1(x, y) := f(y, x, x)$, $f_2(x, y) := f(x, y, x)$, and $f_3(x, y) := f(x, x, y)$ obtained by identifications of two variables yields an elementary operation, i.e., an essentially unary operation. Thus, each of f_1, f_2, f_3 is either equal to $\hat{f}(x)$ or to $\hat{f}(y)$. As in the proof of Rosenberg's theorem [27], we therefore distinguish eight cases.

(1) $f(y,x,x) = \hat{f}(x), f(x,y,x) = \hat{f}(x), f(x,x,y) = \hat{f}(x)$: In this case f satisfies the equations of the quasi majority operation.
(2) $f(y,x,x) = \hat{f}(x), f(x,y,x) = \hat{f}(x), f(x,x,y) = \hat{f}(y)$: In this case f satisfies the equations of a quasi semiprojection to the third argument.
(3) $f(y,x,x) = \hat{f}(x), f(x,y,x) = \hat{f}(y), f(x,x,y) = \hat{f}(x)$: The same for the second argument.
(4) $f(y,x,x) = \hat{f}(x), f(x,y,x) = \hat{f}(y), f(x,x,y) = \hat{f}(y)$.
(5) $f(y,x,x) = \hat{f}(x), f(x,y,x) = \hat{f}(x), f(x,x,y) = \hat{f}(y)$. In this case f satisfies the equations of a semiprojection to the first argument.
(6) $f(y,x,x) = \hat{f}(y), f(x,y,x) = \hat{f}(x), f(x,x,y) = \hat{f}(y)$.
(7) $f(y,x,x) = \hat{f}(x), f(x,y,x) = \hat{f}(y), f(x,x,y) = \hat{f}(y)$.
(8) $f(y,x,x) = \hat{f}(y), f(x,y,x) = \hat{f}(y), f(x,x,y) = \hat{f}(y)$.

We claim that cases 4,6,7,8 are impossible. For that, we first show that in these cases we can derive a quasi Maltsev operation, and we then show that a quasi Malsev operation is not possible. In case 6 and 7, the operation f satisfies the equations of the quasi Maltsev operation, i.e., $f(x,x,y) = f(y,x,x) = \hat{f}(y)$. In case 4 the operation g' defined by $g'(x,y,z) := g(x,z,y)$, and in case 7 the operation g' defined by $g'(x,y,z) := f(y,x,z)$ satisfy the identities of a quasi Maltsev operation.

We now show that F can not contain a quasi Maltsev operation. By Theorem 5, the relational structure Γ has finitely many orbits of pairs of elements. Because Γ is a model-complete core, Theorem 8 implies that every relation that consists of an orbit of k-tuples has an existential positive definition. In particular, the orbits of pairs must have primitive positive definitions. Now, let $v_1, v_2, \dots$ be an enumeration of the elements of Γ. We assign a pair (v_i, v_j), where $i < j$, a color according to the orbit of the pair. This is, two pairs have the same color if and only if they lie in the same orbit of pairs. Since Γ is ω-categorical, Theorem 5 implies that we only have to use a finite number of colors. The classical version of Ramsey's theorem then implies that Γ *contains a monochromatic triangle*, i.e., there exist three elements a, b, c such that (a,b), (a,c), and (b,c) are from the same orbit of pairs O that is distinct from the orbit of pairs of the form (x,x). Now, suppose f is a quasi Maltsev operation. We then have $f((a,b),(a,c),(b,c)) = (\hat{f}(b), \hat{f}(b))$. Because $(a,b),(a,c),(b,c)$ are from O, but $(\hat{f}(b), \hat{f}(b))$ is not, and because O has by Theorem 8 a primitive positive definition, the operation f cannot be a polymorphism of Γ.

Finally, let f be k-ary, where $k \geq 4$. By minimality of f, the operations obtained from f by identifications of arguments of g are essentially unary. We have to show that f is a quasi semiprojection. Note that if an operation f is a quasi semiprojection then there is an i, $1 \leq i \leq k$, such that f satisfies the equations $f(x_1, x_1, x_3, \dots, x_k) = f(x_1, x_2, x_1, x_4, \dots, x_k) = \cdots = f(x_1, \dots, x_{k-1}, x_{k-1}) = f(x_i, \dots, x_i)$. Conversely, every operation satisfying these equations is a semiprojection. By compactness, it suffices to show that on every finite subset A, there exists an i such that the above equations hold on A.

The lemma of Świerczkowski (the lemma was discovered independently several times; see Satz 4.4.6 in [24]) states that every at least 4-ary operation *on a finite domain* that turns into a projection whenever two arguments are the same is a semiprojection. We can use essentially the same proof for the lemma as it is given in [15] to show that the operation g above acts like a semiprojection on A. Because the proof is short, we recall it here for the convenience of the reader.

Fix a finite subset A of the domain of size at least two. By Proposition 16, there exists a permutation h from F such that $g := h(f)$ is idempotent on A. In the following, the variables x_i, for $1 \leq i \leq k$, denote values from A. It is enough to show that $g(x_1, x_1, x_3, x_4, \dots, x_k) = g(x_1, x_2, x_1, x_4, \dots, x_k)$ for all $x_1, \dots, x_k$. For other identifications of nonoverlapping pairs of arguments of f, apply permutations of arguments; for identifications of overlapping pairs of variables, apply the above two times.

Let i and j be such that $g(x_1, x_1, x_3, x_4, \dots, x_k) = x_i$ and $g(x_1, x_2, x_1, x_4, \dots, x_k) = x_j$. If $i = 1$ and $j \in \{1, 2\}$ we are done. Similarly, if $i \in \{1, 3\}$ and $j = 1$ we are also done. We claim that all other cases lead to a contradiction.

Suppose that $i \in \{1, 3\}$ and $j \notin \{1, 2\}$. Then $g(x_1, x_1, x_3, x_4, \dots, x_k) = x_i$ implies that $g(x_1, x_1, x_1, x_4, \dots, x_k) = x_i$. But $g(x_1, x_2, x_1, x_4, \dots, x_k) = x_j$ implies that $g(x_1, x_1, x_1, x_4, \dots, x_k) = x_j$, a contradiction. The same argument can be applied in case that $i \notin \{1, 3\}$ and $j \in \{1, 2\}$.

The only remaining case is that $i = 3$ and $j = 2$. Because $g(x_1, x_1, x_3, x_4, \dots, x_k) = x_3$ we also have $g(x_1, x_1, x_3, x_1, \dots, x_k) = x_3$. This in turn shows that $g(x_1, x_2, x_3, x_1, \dots, x_k) = x_3$. On the other hand, $g(x_1, x_2, x_1, x_4, \dots, x_k) = x_2$ implies that $g(x_1, x_2, x_1, x_1, \dots, x_k) = x_2$, and this in turn implies that $g(x_1, x_2, x_3, x_1, \dots, x_k) = x_2$, a contradiction. □

7. Baker and Pixley Generalised

Several results on clones over a finite domain are about idempotent operations or idempotent clones. As mentioned earlier, for operations on an infinite domain idempotency is a very strong condition. In this section, we show that the famous theorem of Baker and Pixley can be proved for oligomorphic clones if the condition of idempotency is replaced with that of oligopotency.

Theorem 19. *Let Γ be an ω-categorical relational structure. Then for every $k \geq 2$ the following is equivalent.*

- *Γ has an oligopotent $k+1$-ary polymorphism that is a quasi near-unanimity operation.*
- *Every primitive positive formula is in Γ equivalent to a conjunction of at most k-ary primitive positive formulas.*

Proof. Throughout the whole proof let τ denote the signature that contains a relation symbol for each at most k-ary primitive positive relation in Γ.

We first show that (1) implies that every primitive positive formula ϕ is equivalent to a conjunction of atomic τ-formulas. We can assume without loss of generality that ϕ has a single existentially quantified variable x_0, otherwise we eliminate the existential quantifiers one by one. This is, we can assume that ϕ has the form $\exists x_0.\Phi(x_0, x_1, \dots, x_l)$, where Φ is quantifier free.

If ϕ has at most k free variables, then we can directly eliminate the existential quantifier with the help of the corresponding k-ary relation. We now proceed by induction on the arity l of ϕ, $l > k$, inductively assuming that the claim holds for all formulas with $l-1$ free variables. Let Ψ be the set of all atomic τ-formulas with the free variables $x_1, \ldots, x_l$ that are implied by the formula ϕ. We have to prove that every tuple $B := \{b_1, \ldots, b_l\}$ satisfying all formulas in Ψ also satisfies ϕ. By inductive assumption, the formula $\exists x_i.\phi$ is equivalent to a quantifier-free τ-formula ψ_i in Ψ, for all $1 \leq i \leq l$. Hence, we find a witness b_0^i for the existentially quantified variable x_0 in $\exists x_0, x_i.\Phi(x_0, b_1, \ldots, b_{i-1}, x_i, b_{i+1}, \ldots, b_l)$. By assumption and Proposition 16 we know that Γ has a $k+1$-ary qnu-operation f that is idempotent on the elements $b_1, \ldots, b_l$. We claim that $f(b_0^1, \ldots, b_0^{k+1})$ is a witness for the existentially quantified variable x_0 in ϕ that shows that ϕ holds on $b_1, \ldots, b_l$. Let $c_1, \ldots, c_k$ be the witnesses for x in $\exists x_i.\Psi(b_0^i, \ldots, b_{i-1}, x_i, b_{i+1}, \ldots, b_k, b_{k+1}, \ldots, b_l)$. Since f is a polymorphism it preserves the primitive positive definable relation $\Phi(x_0, x_1, \ldots, x_l)$, which therefore also holds on $f(b_0^1, \ldots, b_0^{k+1})$, $f(c_1, b_1, \ldots, b_1)$, $\ldots$, $f(b_k, \ldots, b_k, c_k)$, $f(b_{k+1}, \ldots, b_{k+1})$, $\ldots$, $f(b_l, \ldots, b_l)$. Moreover, since f is a qnu-operation and idempotent on B, we have that $f(b_i, \ldots, b_i, c_i, b_i, \ldots, b_i) = b_i$ for $1 \leq i \leq l$. Therefore Φ holds on $(f(b_0^1, \ldots, b_0^{k+1}), b_1, \ldots, b_l)$; this proves that the tuple $(b_1, \ldots, b_l)$ satisfies ϕ.

Conversely, let us assume that in the signature τ we have quantifier elimination for primitive positive formulas. We have to show that Γ has an oligopotent $k+1$-ary polymorphism that is a quasi near-unanimity operation. We equip the structure $\Delta := \Gamma^{k+1}$ with the *qnu equivalence relation* $\sim$, which is a binary relation defined on Δ by $(b_1, \ldots, b_{k+1}) \sim (b, \ldots, b)$ iff all except one b_i are equal to b. Every homomorphism f from $(\Delta, \sim)$ to $(\Gamma, =)$ is by definition a qnu-operation, and a polymorphism of Γ.

We now construct such a polymorphism, and also make sure that the polymorphism is oligopotent. For that we prove that every finite substructure S of Δ has a homomorphism to Γ that maps elements from S of the form $(b, \ldots, b)$ to b. The existence of the oligopotent polymorphism is then an easy consequence of König's lemma: consider an enumeration $a_1, a_2, \ldots$ of the elements of Δ, and the infinite directed graph G whose vertices are the equivalence classes of homomorphisms from $\{a_1, \ldots, a_n\}$ to Γ, for $n \geq 0$, where two homomorphisms f_1, f_2 are equivalent if there is an automorphism h of Γ such that $h(f_1) = f_2$. There is an arc from one equivalence class of homomorphisms to another in G, if there are representatives f_1 and f_2 of the two classes such that the homomorphism f_1 extends the homomorphism f_2 by one element a_i of Δ. Theorem 5 and ω-categoricity of Γ assert that every vertex in G has a finite number of outgoing arcs, and König's lemma asserts the existence of an infinite path. This infinite path gives rise to a homomorphism f from Δ to Γ defined inductively as follows. The mapping f will be such that its restriction to $\{a_1, \ldots, a_n\}$ is from the n-th element of the infinite path in G. Initially, this is trivially true if f is restricted to the empty set. Suppose f is already defined on $a_1, \ldots, a_n$, for $n \geq 0$. By construction of the infinite path, we find representatives f_n and f_{n+1} of the n-th and the $n+1$-st element on the path such that f_n is a restriction of f_{n+1}. The inductive assumption gives us an automorphism h of Γ such that $h(f_n(x)) = f(x)$ for all $x \in \{a_1, \dots, a_n\}$. We set $f(a_{n+1})$ to be $h(f_{n+1}(a_{n+1}))$. The restriction of f to $a_1, \dots, a_n$ will therefore be a member of the $n+1$-st element of the infinite path. This shows that f is indeed a homomorphism from Δ to Γ. We have thus shown that if every finite substructure of Δ homomorphically maps to Γ, then Δ homomorphically maps to Γ.

Let S be a finite substructure of Δ, and let B be the substructure of S containing all the elements $(b_1, \dots, b_{k+1})$ where all but at most one entries are equal. Consider the primitive positive formula that corresponds to S, where all vertices that are not in B are existentially quantified, and where all other vertices correspond to free variables. By assumption, this primitive positive formula is equivalent to a conjunction of τ-atoms. We show that we can map the free variables to Γ such that these atoms are satisfied, which implies that S homomorphically maps to Γ. We map vertices $(b_1, \dots, b_{k+1})$ to b if all but one entry are equal to b. The pigeon hole principle then implies that for k vertices in B there is a position i, $1 \leq i \leq k+1$, such that the i-th entries of the k vertices are not the exceptional positions. The projection to the i-th component therefore satisfies the τ-atoms on these k vertices. □

Ackowledgement. We thank an anonymous referee for his helpful comments.

References

[1] K. A. Baker and A. F. Pixley. Polynomial interpolation and the Chinese remainder theorem for algebraic systems. *Math. Z.*, 143:165–174, 1974.

[2] M. Bhattacharjee, D. Macpherson, R. G. Moller, and P. M. Neumann. *Notes on Infinite Permutation Groups*. Springer LNCS 1698, 1998.

[3] M. Bodirsky. Constraint satisfaction with infinite domains. Dissertation, Humboldt-Universität zu Berlin, 2004.

[4] M. Bodirsky. Cores of countably categorical structures. *Logical Methods in Computer Science (LMCS)*, 2007. DOI: 10.2168/LMCS-3(1:2).

[5] M. Bodirsky and V. Dalmau. Datalog and constraint satisfaction with infinite templates. In *Proceedings of the 23rd International Symposium on Theoretical Aspects of Computer Science (STACS'06)*, LNCS 3884, pages 646–659. Springer Verlag, 2006.

[6] M. Bodirsky and J. Nešetřil. Constraint satisfaction with countable homogeneous templates. *Journal of Logic and Computation (JLC)*, 16(3):359–373, 2006.

[7] V. G. Bodnarčuk, L. A. Kalužnin, V. N. Kotov, and B. A. Romov. Galois theory for post algebras, part I and II. *Cybernetics*, 5:243–539, 1969.

[8] A. Bulatov. A graph of a relational structure and constraint satisfaction problems. *In Proceedings of 19th IEEE Symposium on Logic in Computer Science (LICS04)*, pages 448–457, 2004.

[9] A. Bulatov, A. Krokhin, and P. Jeavons. The complexity of maximal constraint languages. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC'01)*, pages 667–674, 2001.

[10] A. Bulatov, A. Krokhin, and P. G. Jeavons. Classifying the complexity of constraints using finite algebras. *SIAM Journal on Computing*, 34:720–742, 2005.

[11] P. J. Cameron. *Oligomorphic Permutation Groups*. Cambridge University Press, 1990.

[12] G. Cherlin. The classification of countable homogeneous directed graphs and countable homogeneous n-tournaments. *AMS Memoir*, 131(621), January 1998.

[13] G. Cherlin, S. Shelah, and N. Shi. Universal graphs with forbidden subgraphs and algebraic closure. *Advances in Applied Mathematics*, 22:454–491, 1999.

[14] B. Czákány. All minimal clones on the three-element set. *Acta Cybernet.*, 6:227–238, 1983.
[15] B. Czákány. Minimal clones. Preprint, 2001.
[16] M. Droste and D. Kuske. On random relational structures. *Journal of Combinatorial Theory*, 102(2):241–254, 2003.
[17] T. Feder and M. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory. *SIAM J. Comput.*, 28:57–104, 1999.
[18] P. Hell and J. Nešetřil. *Graphs and Homomorphisms.* Oxford University Press, 2004.
[19] W. Hodges. *A shorter model theory.* Cambridge University Press, 1997.
[20] P. Jeavons, D. Cohen, and M. Cooper. Constraints, consistency and closure. *AI*, 101(1-2):251–265, 1998.
[21] P. Jeavons, D. Cohen, and M. Gyssens. Closure properties of constraints. *Journal of the ACM*, 44(4):527–548, 1997.
[22] R. W. Kaye and H. D. Macpherson, editors. *Automorphisms of First-order Structures.* Oxford University Press, 1994.
[23] K. A. Kearnes and A. Szendrei. The classification of commutative minimal clones. *Discussiones Math*, 19:147–178, 1999.
[24] R. Pöschel and L. A. Kalužnin. *Funktionen- und Relationenalgebren.* DVW, 1979.
[25] E. L. Post. The two-valued iterative systems of mathematical logic. *Annals of Math. Studies*, 5, 1941.
[26] I. G. Rosenberg. The number of maximal closed classes in the set of functions over a finite domain. *J. Comb. Theory, Ser. A*, 14(1):1–7, 1973.
[27] I. G. Rosenberg. Minimal clones I: the five types. *Lectures in Universal Algebra (Proc. Conf. Szeged, 1983), Colloq. Math. Soc. J. Bolyai*, 43:405–427, 1986.
[28] A. Szendrei. *Clones in universal Algebra.* Seminaire de mathematiques superieures. Les Presses de L'Universite de Montreal, 1986.

Manuel Bodirsky, Humboldt-Universität zu Berlin, Deutschland
E-mail address: bodirsky@informatik.hu-berlin.de

Hubie Chen, Universitat Pompeu Fabra, Barcelona, Spain
E-mail address: hubie.chen@upf.edu