

THE REDUCTS OF EQUALITY UP TO PRIMITIVE POSITIVE INTERDEFINABILITY

MANUEL BODIRSKY, HUBIE CHEN, AND MICHAEL PINSKER

ABSTRACT. We initiate the study of reducts of relational structures up to primitive positive interdefinability: After providing the tools for such a study, we apply these tools in order to obtain a classification of the reducts of the logic of equality. It turns out that there exists a continuum of such reducts. Equivalently, expressed in the language of universal algebra, we classify those locally closed clones over a countable domain which contain all permutations of the domain.

CONTENTS

1. Introduction	1
1.1. Introduction for model theorists	2
1.2. Introduction for universal algebraists	2
1.3. Introduction for computer scientists	3
2. Back and forth between logic and algebra	4
2.1. Notions from logic	4
2.2. Notions from algebra	4
2.3. The translation	5
3. The result	7
3.1. Additional notation and terminology	11
4. The basis: Monoids	11
5. Clones with essential finite range operations	16
6. Clones with essential infinite range operations	19
6.1. The Horn clone	19
6.2. The bar clone	21
6.3. From the bar clone to the odd clone	23
6.4. Richard	32
6.5. Richard's many friends	34
7. Clones with essential infinite range operations plus constants	37
8. Open Problems	39
References	39

1. INTRODUCTION

Our results have impact on three fields: model theory, universal algebra, and theoretical computer science. We therefore have a three-fold introduction.

2000 *Mathematics Subject Classification.* Primary 03C40; secondary 08A40; 08A70; 03D15.

Key words and phrases. relational structure, reduct, primitive positive definition, lattice, invariant relation, Galois connection, local clone, permutations.

The third author is grateful for support through project P17812-N12 as well as through Erwin Schrödinger Fellowship N2742-N18 of the Austrian Science Fund.

1.1. Introduction for model theorists. In model theory, reducts of a relational structure Γ are usually considered *up to first-order interdefinability*. To be more precise, one considers the reducts of the expansion of Γ by all first-order definable relations, and two reducts Γ' and Γ'' are considered to be the same if and only if there is a first-order definition of Γ' in Γ'' and vice versa. It is well-known that there is a close connection between classifications of reducts up to first-order interdefinability and the theory of infinite (closed) permutation groups. In 1976, Cameron [Cam76] showed that the highly set-transitive permutation groups are exactly the automorphism groups of the reducts of $(\mathbb{Q}, <)$. He also showed that there are exactly five reducts of $(\mathbb{Q}, <)$ up to first-order interdefinability. Recently, Junker and Ziegler gave a new proof of this fact, and showed that $(\mathbb{Q}, <, a)$, the expansion of $(\mathbb{Q}, <)$ by a constant a , has 116 reducts [JZ08]. Thomas showed that the first-order theory of the random graph also has exactly five reducts [Tho91] and conjectured that every structure with quantifier elimination in a finite relational signature has a finite number of reducts, up to first-order interdefinability.

In this article, we initiate the systematic study of reducts up to *primitive positive interdefinability*. A first-order formula is *primitive positive* if and only if it is of the form $\exists \bar{x}(\phi_1 \wedge \cdots \wedge \phi_l)$ for atomic formulas ϕ_1, \dots, ϕ_l . Clearly, classifications up to primitive positive interdefinability are harder to obtain, compared to first-order interdefinability, since there are far more reducts to distinguish.

The simplest structure where such a classification can be studied is the structure that has no structure at all except for the equality relation, which is considered part of first-order logic. Our result classifies the reducts of the structure $(X, =)$, where X is a countably infinite set, up to primitive positive definability. While this might look trivial at first sight, we prove that the number of such reducts equals the continuum. It turns out that the partial order (in fact, the lattice) of these reducts is quite complex, but can be described reasonably well. We remark that the classification of the reducts of $(X, =)$ up to primitive positive interdefinability is the same as the classification for a structure of the form $(Y, =)$, where Y is a set of arbitrary infinite cardinality.

To show our classification result, we use a concept that is called *clone* in universal algebra. A *clone* is a set of operations on some fixed domain that is closed under compositions and contains all projections. The clones we are interested in are moreover *locally closed*, which means that they are, similarly to closed permutation groups used for the study of reducts up to first-order interdefinability, closed sets in the natural topology on the operations. For ω -categorical structures Γ , there is a one-to-one correspondence between the reducts of Γ and the locally closed clones that contain the automorphisms of Γ . In order to classify the reducts of $(X, =)$, we thus classify those locally closed clones which contain all permutations of the domain.

We believe that several of the results and techniques presented in this article can be used to also classify reducts of other ω -categorical structures up to primitive positive definability. Such future classifications are likely to combine existing results on first-order closed reducts, such as the above-mentioned results on the dense linear order and the random graph, with the results of the present paper.

1.2. Introduction for universal algebraists. For universal algebraists, the title of our paper could be “The locally closed clones above the permutations”. In fact, in this article we describe the lattice of all locally closed clones on a countably infinite set that contain all permutations of this set.

Traditionally, most work on clones was done for clones on a finite domain, and only occasionally clones on infinite sets were studied. However recently, the use of methods from mathematical logic, in particular from set theory, has allowed for

the production of a considerable number of new results on clones over infinite sets. Classical results on clones, mostly on finite domains, can be found in the books [PK79], [Sze86], and [Lau06], while the recent survey paper [GP08] contains much of what is known on clones over infinite sets.

The clones on a finite domain which contain all permutations were completely determined in [HR94]; it turns out that the number of such clones is finite. If one moves on to infinite domains, there are two possible clone lattices which can be considered: The lattice of local clones, which is closer to the clone lattice on a finite domain in both size and methodology that can be employed for its study, and the lattice of all clones, the investigation of which often involves set-theoretical constructions (as in [GS05], for example).

In the full clone lattice over an infinite set, the interval of those clones which contain all permutations has been subject to investigation. For example, its dual atoms have been completely described on domains of regular cardinality ([Hei02], [Pin05a]). Moreover, its atoms were determined for all infinite domains in [MP07]. However, the number of clones above the permutations has been shown to be huge ([GSS], [Gol], [Pin05b]), and it seems unlikely that all such clones on an infinite set will ever be classified.

As our results will show, the situation is quite different for local clones which contain the permutations. These clones had not yet been subject to explicit investigation and were largely unknown except for a few implicit results, such as partial results on local clones containing all unary operations in [RS82]. It turns out that although the lattice of locally closed clones which contain the permutations is uncountable, one can obtain a reasonable understanding of it.

We remark here that the lattice of all local clones has been shown to be quite complicated; in particular, already on a countably infinite base set, the class of lattices which are completely embeddable into it properly contains the class of all algebraic lattices with a countable number of compact elements ([Pin08], [Pin]).

1.3. Introduction for computer scientists. The satisfiability problem for propositional logic is one of the most fundamental problems in computer science, and among the hardest problems in the complexity class NP. Many restrictions of this problem stay NP-hard – e.g. the problems NOT-ALL-EQUAL-3SAT and ONE-IN-3-SAT. In a groundbreaking paper in 1978, Schaefer studied a large class of natural restrictions of the satisfiability problem [Sch78]. For a fixed structure Γ over the Boolean domain $\{0, 1\}$, the computational problem is to decide whether a given sentence of the form $\exists \bar{x}. \psi_1 \wedge \dots \wedge \psi_l$ is true, where ψ_i is of the form $R(x_{i_1}, \dots, x_{i_k})$ for a k -ary relation R in Γ . This problem is now known as $\text{CSP}(\Gamma)$, the *constraint satisfaction problem* for the *constraint language* Γ . (It is easy to see that 3SAT, NOT-ALL-EQUAL-3SAT, and ONE-IN-3-SAT can be modelled as $\text{CSP}(\Gamma)$ for appropriate Boolean constraint languages Γ .) Schaefer proved that for Boolean constraint languages Γ the problem $\text{CSP}(\Gamma)$ can either be solved in polynomial time, or is NP-complete, and he gave a complete description of the tractable and the hard Boolean constraint languages.

Schaefer's theorem inspired many other complexity classification projects for computational problems that are 'parametrized' by a constraint language, and similar results have been obtained not only for constraint satisfaction, but also for other computational tasks such as *quantified constraint satisfaction*, *learning of propositional formulas*, *counting solutions*, and the *maximum solution problem*. It turns out that for all these computational problems the computational complexity does not change if we add relations to the constraint language Γ that have a primitive positive definition in Γ . This implies that the complexity of these problems only

depends on the so-called *clone* consisting of all *polymorphisms* of the constraint language, i.e., the set of those operations that preserve all relations in Γ .

The lattice of clones on a two-element set was determined already quite some time before Schaefer, namely by Post in 1941 [Pos41]. And indeed, Schaefer's classification can be formulated in terms of Post's lattice, see e.g. [BCRV04].

Our result can be viewed as an analog of Post's lattice for the logic of equality (rather than for propositional logic). In this analogy, Boolean relations correspond to relations that have a first-order definition in $(X, =)$, where X is an arbitrary infinite set. Sets of such relations have been called *equality constraint languages* [BK08a]. Hence, in this paper we study the lattice of polymorphism clones of equality constraint languages. Note that equality constraint languages are exactly those constraint languages that are preserved by all permutations of the domain; in this sense, we study those constraint languages up to primitive positive interdefinability that have the maximal possible degree of symmetry.

The constraint satisfaction problem for equality constraint languages has been investigated in [BK08a]. It was shown that $\text{CSP}(\Gamma)$ is NP-complete if the polymorphism clone of the constraint language is the smallest element in the lattice; otherwise, it is polynomial-time solvable. The *quantified constraint satisfaction problem* (*QCSP*) for equality constraint languages Γ has been investigated in [BC07].

2. BACK AND FORTH BETWEEN LOGIC AND ALGEBRA

We make the notions of the introduction precise, and fix some definitions and notation from logic and universal algebra. Moreover, we explain how to translate our classification problem from logic to algebra and vice-versa: That is, we demonstrate in what way reducts correspond to local clones.

2.1. Notions from logic. Consider the relational structure consisting of a countably infinite domain and the relation of equality only. For notational simplicity, we will assume that this domain is the set \mathbb{N} of natural numbers. By a *reduct* of $(\mathbb{N}, =)$ we mean a structure $\Gamma = (\mathbb{N}, \mathcal{R})$, where \mathcal{R} is a set of finitary relations on \mathbb{N} all of which can be defined by a first-order formula in the language of equality.

Let ϕ be a first-order formula over the relational signature τ . Then ϕ is called *primitive positive* or *pp* iff it is of the form $\exists \bar{x}. \psi_1 \wedge \cdots \wedge \psi_l$, where ψ_1, \dots, ψ_l are atomic τ -formulas. For a set \mathcal{R} of relations on \mathbb{N} , we write $\text{pp}(\mathcal{R})$ for the set of all relations on \mathbb{N} which can be defined from \mathcal{R} by pp formulas (over the signature appropriate for \mathcal{R}).

We say that a structure $\Gamma_1 = (\mathbb{N}, \mathcal{R}_1)$ is *pp-definable* from a structure $\Gamma_2 = (\mathbb{N}, \mathcal{R}_2)$ iff $\mathcal{R}_1 \subseteq \text{pp}(\mathcal{R}_2)$. We say that Γ_1, Γ_2 are *pp-interdefinable* iff they are pp-definable from each other. The structure Γ_1 is *pp-closed* iff $\mathcal{R}_1 = \text{pp}(\mathcal{R}_1)$. The notion of pp-definability imposes a quasiorder on the reducts of $(\mathbb{N}, =)$, and the relation of pp-interdefinability is an equivalence relation on this order. We are interested in the partial order which is obtained by factoring the quasiorder of reducts by this equivalence relation. That is, we consider two reducts the same iff they are pp-interdefinable, or equivalently, we consider the partial order of pp-closed reducts.

In fact, the pp-closed reducts form a complete lattice, where the meet of a family of structures $(\Gamma_i)_{i \in I}$ is the structure $(\mathbb{N}, \bigcap_{i \in I} \mathcal{R}_i)$, and the join is $(\mathbb{N}, \text{pp}(\bigcup_{i \in I} \mathcal{R}_i))$.

2.2. Notions from algebra. Fix a countably infinite base set; for convenience, we take the set \mathbb{N} of natural numbers. For all $n \geq 1$, denote by $\mathcal{O}^{(n)}$ the set $\mathbb{N}^{\mathbb{N}^n} = \{f : \mathbb{N}^n \rightarrow \mathbb{N}\}$ of n -ary operations on \mathbb{N} . Then $\mathcal{O} = \bigcup_{n \geq 1} \mathcal{O}^{(n)}$ is the set of all finitary operations on \mathbb{N} . A *clone* \mathcal{C} is a subset of \mathcal{O} satisfying the following two properties:

- \mathcal{C} contains all projections, i.e., for all $1 \leq k \leq n$ the operation $\pi_k^n \in \mathcal{O}^{(n)}$ defined by $\pi_k^n(x_1, \dots, x_n) = x_k$, and
- \mathcal{C} is closed under composition, i.e., whenever $f \in \mathcal{C}$ is n -ary and $g_1, \dots, g_n \in \mathcal{C}$ are m -ary, then the operation $f(g_1, \dots, g_n) \in \mathcal{O}^{(m)}$ defined by

$$(x_1, \dots, x_m) \mapsto f(g_1(x_1, \dots, x_m), \dots, g_n(x_1, \dots, x_m))$$

also is an element of \mathcal{C} .

Since arbitrary intersections of clones are again clones, the set of all clones on \mathbb{N} , equipped with the order of inclusion, forms a complete lattice $\text{Cl}(\mathbb{N})$. In this paper, we are interested in certain clones of $\text{Cl}(\mathbb{N})$ which satisfy an additional topological closure property: Equip \mathbb{N} with the discrete topology, and $\mathcal{O}^{(n)} = \mathbb{N}^{\mathbb{N}^n}$ with the corresponding product topology (Tychonoff topology), for every $n \geq 1$. A clone \mathcal{C} is called *locally closed* or just *local* iff each of its n -ary fragments $\mathcal{C} \cap \mathcal{O}^{(n)}$ is a closed subset of $\mathcal{O}^{(n)}$. Equivalently, a clone \mathcal{C} is local iff it satisfies the following interpolation property:

For all $n \geq 1$ and all $g \in \mathcal{O}^{(n)}$, if for all finite $A \subseteq \mathbb{N}^n$ there exists an n -ary $f \in \mathcal{C}$ which agrees with g on A , then $g \in \mathcal{C}$.

Again, taking the set of all local clones on \mathbb{N} , and ordering them according to set-theoretical inclusion, one obtains a complete lattice, which we denote by $\text{Cl}_{\text{loc}}(\mathbb{N})$: This is because arbitrary intersections of clones are clones, and because arbitrary intersections of closed sets are closed. The join of a family $(\mathcal{C}_i)_{i \in I}$ is calculated as follows: First, consider the set of all operations on \mathbb{N} which can be obtained by composing operations from $\bigcup_{i \in I} \mathcal{C}_i$; this set is a clone, but might not be locally closed. For this reason, the topological closure of this set has to be formed in addition in order to arrive at the join in $\text{Cl}_{\text{loc}}(\mathbb{N})$. For a set of operations $\mathcal{F} \subseteq \mathcal{O}$, we write $\langle \mathcal{F} \rangle$ for the smallest local clone containing \mathcal{F} ; if \mathcal{C} is a local clone and $\mathcal{C} = \langle \mathcal{F} \rangle$, then we say that \mathcal{C} is *generated* by \mathcal{F} . Observe that $\langle \mathcal{F} \rangle$ is just the intersection of all local clones containing \mathcal{F} , or equivalently the topological closure of the set of term operations that can be built from \mathcal{F} . Using this notation, the join of the family $(\mathcal{C}_i)_{i \in I}$ is simply $\langle \bigcup_{i \in I} \mathcal{C}_i \rangle$.

This paper investigates, and in some sense classifies, those local clones on \mathbb{N} which contain all permutations of \mathbb{N} . It turns out that the number of such clones is uncountable.

2.3. The translation. Let $f \in \mathcal{O}^{(n)}$ and let $\rho \subseteq \mathbb{N}^m$ be a relation. We say that f *preserves* ρ iff $f(r_1, \dots, r_n) \in \rho$ whenever $r_1, \dots, r_n \in \rho$, where $f(r_1, \dots, r_n)$ is calculated componentwise. This notion of preservation links finitary relations on \mathbb{N} to finitary operations and is the prime tool for studying the reducts of $(\mathbb{N}, =)$, and indeed of all other ω -categorical structures, up to primitive positive interdefinability.

For a set of relations \mathcal{R} , we write $\text{Pol}(\mathcal{R})$ for the set of those operations in \mathcal{O} which preserve all $\rho \in \mathcal{R}$. The operations in $\text{Pol}(\mathcal{R})$ are called *polymorphisms* of \mathcal{R} . The following is folklore in universal algebra, see e.g. [Sze86].

Proposition 1. *$\text{Pol}(\mathcal{R})$ is a local clone for all sets of relations \mathcal{R} . Moreover, every local clone is of this form.*

We have seen how to assign sets of operations to sets of relations; likewise, we can go the other way. For an operation $f \in \mathcal{O}^{(n)}$ and a relation $\rho \subseteq \mathbb{N}^m$, we say that ρ is *invariant* under f iff f preserves ρ . Given a set of operations $\mathcal{F} \subseteq \mathcal{O}$, we write $\text{Inv}(\mathcal{F})$ for the set of all relations which are invariant under all $f \in \mathcal{F}$. The following is well-known.

Proposition 2. *$\text{Inv}(\mathcal{F})$ is a pp-closed set of relations for every $\mathcal{F} \subseteq \mathcal{O}$.*

Using the Galois connection defined by the operators Pol and Inv which link operations and relations, we obtain the following well-known alternative for describing the hull operator which assigns to a set of operations the local clone which this set generates (confer [Sze86]).

Proposition 3. *Let $\mathcal{F} \subseteq \mathcal{O}$. Then $\langle \mathcal{F} \rangle = \text{Pol Inv}(\mathcal{F})$.*

The reader will now expect that our notion of closure for structures Γ , namely closure under primitive positive definitions, will coincide with the closure operator Inv Pol of our Galois connection. Although this is not true for infinite structures in general, the following theorem from [BN06] states that it holds if Γ is not too far from the finite. (In the following, we often identify structures Γ with their sets of relations, e.g., when writing $\text{Pol}(\Gamma)$.)

Theorem 4. *Let Γ be ω -categorical. Then $\text{Inv Pol}(\Gamma) = \text{pp}(\Gamma)$.*

Ordering locally closed permutation groups (just like clones, a permutation group is called locally closed iff it is a closed subset of \mathcal{O}) by set-theoretical inclusion, one obtains a complete lattice since arbitrary intersections of local permutation groups are local permutation groups again. Similarly, the first-order closed structures on \mathbb{N} form a complete lattice. Using our notation, we can state the well-known connection (see e.g. [Cam90]) between locally closed permutation groups and first-order closed reducts of ω -categorical structures as follows.

Theorem 5. *Let Γ be an ω -categorical structure. The mapping $\mathcal{G} \mapsto \text{Inv}(\mathcal{G})$ sending every locally closed permutation group \mathcal{G} containing the automorphism group $\text{Aut}(\Gamma)$ of Γ to its set of invariant relations is an antiisomorphism between the lattice of such groups and the lattice of those reducts of Γ which are closed under first-order definitions.*

Utilizing local clones, we obtain an analogous statement to Theorem 5 for reducts up to primitive positive interdefinability.

Theorem 6. *Let Γ be an ω -categorical structure. Then we have:*

- (1) *The operator Pol maps every reduct of the first-order expansion of Γ to a local clone above $\text{Aut}(\Gamma)$.*
- (2) *Two reducts Δ_1, Δ_2 are mapped to the same local clone if and only if they are equivalent with respect to primitive positive interdefinability.*
- (3) *Every local clone above $\text{Aut}(\Gamma)$ is the polymorphism clone of a reduct of Γ .*
- (4) *The mappings Pol and Inv are antiisomorphisms between the lattice of local clones above $\text{Aut}(\Gamma)$ and the lattice of those reducts of the first-order expansion of Γ which are closed under primitive positive definitions.*

Proof. (1): Let Γ' be the expansion of Γ by all first-order definable relations and let Δ be a reduct of Γ' . It is easy to see and well-known that $\text{Aut}(\Gamma) = \text{Aut}(\Gamma')$. Thus $\text{Pol}(\Delta) \supseteq \text{Aut}(\Delta) \supseteq \text{Aut}(\Gamma') = \text{Aut}(\Gamma)$. By Proposition 1, $\text{Pol}(\Gamma)$ is a local clone.

(2): If $\text{Pol}(\Delta_1) = \text{Pol}(\Delta_2)$, then $\text{pp}(\Delta_1) = \text{pp}(\Delta_2)$ by Theorem 4. On the other hand, if $\text{Pol}(\Delta_1) \neq \text{Pol}(\Delta_2)$, then $\langle \text{Pol}(\Delta_1) \rangle = \text{Pol}(\Delta_1) \neq \text{Pol}(\Delta_2) = \langle \text{Pol}(\Delta_2) \rangle$. Thus, $\text{Pol Inv Pol}(\Delta_1) \neq \text{Pol Inv Pol}(\Delta_2)$ by Proposition 3, and hence $\text{Inv Pol}(\Delta_1) \neq \text{Inv Pol}(\Delta_2)$. Theorem 4 then implies $\text{pp}(\Delta_1) \neq \text{pp}(\Delta_2)$.

(3): Given a local clone $\mathcal{C} \supseteq \text{Aut}(\Gamma)$, consider $\Delta := \text{Inv}(\mathcal{C})$. By Proposition 3 we have $\mathcal{C} = \langle \mathcal{C} \rangle = \text{Pol}(\Delta)$. It remains to show that Δ is a reduct of Γ , i.e., $\Delta \subseteq \Gamma'$, where Γ' is the first-order expansion of Γ . But this is obvious, since $\mathcal{C} \supseteq \text{Aut}(\Gamma)$ implies $\Delta = \text{Inv}(\mathcal{C}) \subseteq \text{Inv Aut}(\Gamma) = \Gamma'$.

(4): By (1), (2), and (3), Pol is a bijective mapping from the pp -closed reducts of Γ onto the local clones containing $\text{Aut}(\Gamma)$. It is obvious from its definition that

this mapping is antitone. By Theorem 4, the restriction of Inv to the local clones containing $\text{Aut}(\Gamma)$ is the inverse of Pol . \square

The above theorem tells us that classifying reducts of ω -categorical structures up to primitive positive interdefinability really amounts to understanding the lattice of local clones containing its automorphisms; in our case, since clearly all permutations are automorphisms of $(\mathbb{N}, =)$, we have to investigate local clones containing all permutations of \mathbb{N} .

3. THE RESULT

Via the Galois correspondence $\text{Inv}\text{--}\text{Pol}$, it is possible to describe the reducts of $(\mathbb{N}, =)$ either on the relational or on the operational side. We will now give the classification on the clone side in as much detail as is gainful at this point of the paper; since we still want to spare the reader the technical details, we will have to be informal at times.

An operation is called *essentially unary* iff it depends on at most one of its variables. A *unary clone* is a clone all of whose operations are essentially unary. Clearly, unary clones are just disguised monoids of transformations on \mathbb{N} . The first part of our result describes the locally closed unary clones containing the set $S_{\mathbb{N}}$ of all permutations of \mathbb{N} .

Theorem 7. *Let \mathfrak{U} be the lattice of locally closed unary clones containing $S_{\mathbb{N}}$. Then:*

- (1) \mathfrak{U} is countably infinite.
- (2) \mathfrak{U} is isomorphic to the lattice of order ideals of a certain partial order \mathbb{P} on the set of finite increasing sequences of positive natural numbers.
- (3) In particular, both \mathfrak{U} and its inverse order are algebraic (that is, isomorphic to the subalgebra lattice of an algebra), and \mathfrak{U} is a distributive lattice.
- (4) \mathfrak{U} is a well-quasi-order, that is, there exist no infinite descending chains and no infinite antichains in \mathfrak{U} . (In fact, \mathfrak{U} is a better-quasi-order.)
- (5) All elements of \mathfrak{U} are finitely generated over $S_{\mathbb{N}}$, i.e., for every $\mathcal{M} \in \mathfrak{U}$ there exists a finite set $\mathcal{F} \subseteq \mathcal{O}$ such that $\mathcal{M} = \langle \mathcal{F} \cup S_{\mathbb{N}} \rangle$.

Theorem 7 will be proven in Section 4; it is also there that the partial order \mathbb{P} is defined. We remark that \mathfrak{U} is an exceptionally well-behaved part of $\text{Cl}_{\text{loc}}(\mathbb{N})$, since $\text{Cl}_{\text{loc}}(\mathbb{N})$ is far from satisfying any statement of the theorem: $|\text{Cl}_{\text{loc}}(\mathbb{N})| = 2^{\aleph_0}$, $\text{Cl}_{\text{loc}}(\mathbb{N})$ is not algebraic (not even upper continuous), and $\text{Cl}_{\text{loc}}(\mathbb{N})$ does not satisfy any non-trivial lattice identities [Pin08].

Having described the unary clones, we proceed as follows: Consider any unary clone. Then this clone has only essentially unary operations, and therefore differs only formally from a monoid \mathcal{M} of transformations. Now it turns out that the set of all local clones \mathcal{C} which have \mathcal{M} as their unary fragment, i.e., which satisfy $\mathcal{C} \cap \mathcal{O}^{(1)} = \mathcal{M}$, forms an interval $I_{\mathcal{M}}$ of the lattice $\text{Cl}_{\text{loc}}(\mathbb{N})$; intervals of this form are called *monoidal*. The smallest element of $I_{\mathcal{M}}$ is the unary clone we started this argument with, namely the clone of those essentially unary operations which are “elements” of \mathcal{M} ; this clone is just $\langle \mathcal{M} \rangle$. The largest element of $I_{\mathcal{M}}$ is called $\text{Pol}(\mathcal{M})$ and contains all $f \in \mathcal{O}$ satisfying $f(g_1, \dots, g_n) \in \mathcal{M}$ whenever $g_1, \dots, g_n \in \mathcal{M}$. (This notation is consistent with our previous use of Pol , if one thinks of the elements of \mathcal{M} as \mathbb{N} -ary relations.) Clearly, the monoidal intervals constitute a natural partition of $\text{Cl}_{\text{loc}}(\mathbb{N})$. Our strategy for describing the local clones containing $S_{\mathbb{N}}$ is to determine the monoidal interval $I_{\mathcal{M}}$ for each locally closed monoid \mathcal{M} containing $S_{\mathbb{N}}$; confer Figure 1.




FIGURE 1. Monoidal intervals

The following theorem describes $I_{\mathcal{M}}$ for all monoids \mathcal{M} which contain an operation which is neither injective nor constant. We refer the reader to Section 5, which contains the proof of the theorem, for the definition of quasilinearity.

Theorem 8. *Let \mathcal{M} be a locally closed monoid containing $S_{\mathbb{N}}$ as well as a non-constant and non-injective operation. Then:*

- (1) *If $\mathcal{M} = \mathcal{O}^{(1)}$, then $I_{\mathcal{M}}$ is a chain of order type $\omega + 1$ with largest element \mathcal{O} .
 Its smallest element is the clone $\langle \mathcal{O}^{(1)} \rangle$ of all essentially unary operations.
 Its second smallest element is Burle's clone \mathcal{Q} of all operations which are either essentially unary or quasilinear.
 For $n \geq 3$, its n -th smallest element is the clone \mathcal{K}_n of all operations which are either essentially unary or whose range contains less than n elements.*
- (2) *If $\mathcal{M} \neq \mathcal{O}^{(1)}$, then there exists a maximal natural number $k = k(\mathcal{M}) \geq 1$ such that \mathcal{M} contains all unary operations which take at most k values.
 If $k = 1$, then $I_{\mathcal{M}}$ has only one element $\langle \mathcal{M} \rangle$.
 Otherwise, $I_{\mathcal{M}}$ is a finite chain of length $k + 1$, and:
 Its smallest element is $\langle \mathcal{M} \rangle$.
 Its second smallest element $\mathcal{Q}^{\mathcal{M}}$ consists of $\langle \mathcal{M} \rangle$ plus all quasilinear operations.
 For $3 \leq n \leq k + 1$, its n -th smallest element $\mathcal{K}_n^{\mathcal{M}}$ consists of $\langle \mathcal{M} \rangle$ plus all operations whose range is smaller than n .*

We now turn to the monoid \mathcal{I} locally generated by $S_{\mathbb{N}}$. This, as a quick check shows, consists of all injections in $\mathcal{O}^{(1)}$. Its monoidal interval is the hardest to understand, and we need a few definitions before stating the theorem describing it.

Definition 9 (The Horn clone \mathcal{H}). Let \mathcal{H} be the set of operations which are, up to fictitious variables, injective.

Definition 10 (The Bar clone \mathcal{B}). Let $f \in \mathcal{O}^{(2)}$ and let $k \geq 1$. If there exists an injection $p \in \mathcal{O}^{(2)}$ such that $f(x_1, x_2) = p(x_1, x_2)$ for all $x_1 \geq k$ and $f(x_1, x_2) = p(x_1, 0) = x_1$ for all $x_1 < k$, then we call f a k -bar function. Let \mathcal{B} be the clone generated (in the sense of 3.1) by any bar function, i.e., the smallest local clone containing that bar function and all permutations of \mathbb{N} (we will see in Section 6.2 that this definition makes sense).

Definition 11 (Richard \mathcal{R}). Let $1 \leq i \leq n$. We call an operation $f \in \mathcal{O}^{(n)}$ *injective in the i -th direction* iff $f(a) \neq f(b)$ whenever $a, b \in X^n$ and $a_i \neq b_i$. We say that $f \in \mathcal{O}^{(n)}$ is *injective in one direction* iff there exists $1 \leq i \leq n$ such that f is injective in the i -th direction. Let \mathcal{R} be the set of all operations which are injective in one direction.

Definition 12 (The odd clone \mathcal{S}). Let $f_3 \in \mathcal{O}^{(3)}$ any operation satisfying the following:

- $f_3(x, 1, 1) = 1$, $f_3(2, x, 2) = 2$, $f_3(3, 3, x) = 3$, and
- For all other arguments, the function arbitrarily takes a value that is distinct from all other function values.

We set \mathcal{S} to be the clone generated by f_3 , i.e., the smallest local clone containing f_3 and $S_{\mathbb{N}}$.


The following theorem summarizes the highlights of the monoidal interval corresponding to the monoid \mathcal{I} generated by $S_{\mathbb{N}}$; confer Figure 2. More detailed descriptions of the clones of the theorem as well as other clones in that interval can be found in Section 6.

Theorem 13. *The monoid locally generated by $S_{\mathbb{N}}$ is the monoid \mathcal{I} of injections, and:*

- (1) *The largest element of $I_{\mathcal{I}}$, $\text{Pol}(\mathcal{I})$, equals $\text{Pol}(\{\neq\})$, where \neq is the (binary) inequality relation.*
- (2) *\mathcal{H} is the unique cover of $\langle \mathcal{I} \rangle$ in $I_{\mathcal{I}}$, and all elements of $I_{\mathcal{I}}$ except $\langle \mathcal{I} \rangle$ contain \mathcal{H} . Moreover, \mathcal{H} is generated by any binary injection, and $\text{Inv}(\mathcal{H})$ consists of all relations definable by a Horn formula.*
- (3) *\mathcal{B} is the unique cover of \mathcal{H} in $I_{\mathcal{I}}$, and all elements of $I_{\mathcal{I}}$ except $\langle \mathcal{I} \rangle$ and \mathcal{H} contain \mathcal{B} .*
- (4) *\mathcal{R}, \mathcal{S} are incomparable clones in $I_{\mathcal{I}}$ and every clone in $I_{\mathcal{I}}$ is either contained in \mathcal{S} or contains \mathcal{R} .*
- (5) *The number of elements of $I_{\mathcal{I}}$ containing \mathcal{R} equals the continuum: In fact, the power set of ω , ordered by reverse inclusion, has an order embedding into the interval $[\mathcal{R}, \text{Pol}(\mathcal{I})]$. In particular, the same holds for the interval $I_{\mathcal{I}}$, as well as for the set of local clones above $S_{\mathbb{N}}$.*

The last statement of Theorem 13 is among the hardest to prove in this paper, and has strong consequences for pp classification projects, so that it deserves an own corollary.

Corollary 14. *Let Γ be any relational structure. Then the number of its pp-closed reducts is uncountable. In fact, there exists an order embedding of the power set of ω into the lattice of pp-closed reducts of Γ .*

FIGURE 2. The monoidal interval of \mathcal{I}

It is for this reason that we cannot expect to completely characterize the pp-closed reducts of any relational structure. In our case, we obtain a complete characterization of the closed monoid lattice and of all monoidal intervals except for those corresponding to the monoids \mathcal{I} and \mathcal{I}^+ (see below), where we must content ourselves with some insights on the structure of those intervals.

It remains to describe the monoidal intervals of those monoids which contain all injections, and some constant operations. Clearly, there is only one such monoid, namely the monoid \mathcal{I}^+ consisting of all constants and all injections. In general, for any set of operations $\mathcal{F} \subseteq \mathcal{O}$, write \mathcal{F}^+ for \mathcal{F} plus all constant operations, and \mathcal{F}^- for \mathcal{F} without all constant operations. It turns out that $I_{\mathcal{I}^+}$ is a complete sublattice of $I_{\mathcal{I}}$, as described in the following theorem:

Theorem 15. *Let \mathcal{I}^+ be the monoid of all injective and of all constant operations. Then:*

- (1) *If $\mathcal{C} \in I_{\mathcal{I}^+}$, then \mathcal{C}^- is a local clone in $I_{\mathcal{I}}$.*
- (2) *$\langle \mathcal{I}^+ \rangle^- = \langle \mathcal{I} \rangle$ and $(\text{Pol}(\mathcal{I}^+))^- = \mathcal{I}$.*
- (3) *The mapping from $I_{\mathcal{I}^+}$ into the subinterval $[\langle \mathcal{I} \rangle, \mathcal{I}]$ of $I_{\mathcal{I}}$ which sends every clone \mathcal{C} to \mathcal{C}^- is a complete lattice embedding which preserves the smallest and the largest element.*
- (4) *If $\mathcal{C}, \mathcal{D} \in I_{\mathcal{I}^+}$, then $\mathcal{C} \subsetneq \mathcal{D}$ iff $\mathcal{C}^- \subsetneq \mathcal{D}^-$.*
- (5) *For all $\mathcal{C} \in I_{\mathcal{I}}$ which do not contain \mathcal{R} , $\langle \mathcal{C}^+ \rangle$ is a local clone in $I_{\mathcal{I}^+}$.*

- (6) All clones \mathcal{C} in $I_{\mathcal{J}^+}$ are of this form, as $\mathcal{C} = (\mathcal{C}^-)^+$.
- (7) $\langle \mathcal{R}^+ \rangle = \mathcal{O}$.
- (8) \mathcal{H}^+ is the unique cover of $\langle \mathcal{J}^+ \rangle$ in $I_{\mathcal{J}^+}$.
- (9) \mathcal{B}^+ is the unique cover of \mathcal{H}^+ in $I_{\mathcal{J}^+}$.

We remark that the mapping that sends every $\mathcal{C} \in I_{\mathcal{J}^+}$ to \mathcal{C}^- is not surjective onto the interval $[\langle \mathcal{J} \rangle, \mathcal{J}]$; see the remark after Proposition 85.

The following sections contain the proof of our result, and of course more detailed definitions of the structures involved. Each theorem corresponds to a section: Theorems 7, 8, 13 and 15 are proven in Sections 4, 5, 6, and 7, respectively.

3.1. Additional notation and terminology. In addition to the notation introduced so far, we establish the following conventions. If $f \in \mathcal{O}$, since we are only interested in local clones containing $S_{\mathbb{N}}$, we abuse the notation $\langle \cdot \rangle$ and write $\langle f \rangle$ for the local clone generated by f together with $S_{\mathbb{N}}$. For $f, g \in \mathcal{O}$, we say that f *generates* g iff $g \in \langle f \rangle$. Similarly for $g \in \mathcal{O}$ and $\mathcal{F} \subseteq \mathcal{O}$, we say that \mathcal{F} *generates* g iff $g \in \langle \mathcal{F} \rangle$. If $f \in \mathcal{O}$, we write $\text{ran}(f)$ for its range.

For a relation R , we will usually write $\text{Pol}(R)$ instead of $\text{Pol}(\{R\})$ for the set of all operations which preserve R . If a is an n -tuple, then we refer by a_i to the i -th component of a , for all $1 \leq i \leq n$.

4. THE BASIS: MONOIDS

We prove Theorem 7 describing all unary clones containing $S_{\mathbb{N}}$. Recall that a unary clone consists only of operations depending on at most one variable, and is therefore a disguised monoid of transformations. For convenience, we therefore only deal with unary operations and monoids in this section. In particular, we adjust the meaning of certain notations for this section: For example, $\langle \mathcal{F} \rangle$ refers to the local *monoid* (rather than the local clone) generated by a set $\mathcal{F} \subseteq \mathcal{O}^{(1)}$ together with $S_{\mathbb{N}}$.

The various statements of Theorem 7 are obtained in Propositions 26, 30, 32 and Corollaries 27 and 31.

In a first lemma, we officially state what we already observed in the last section, namely that $S_{\mathbb{N}}$ locally generates all unary injections.

Lemma 16. *Let $\mathcal{M} \supseteq S_{\mathbb{N}}$ be a locally closed monoid. Then \mathcal{M} contains the monoid \mathcal{J} of all unary injective operations.*

Proof. Clearly, on any finite set every injection can be locally interpolated by a suitable permutation. \square

The following lemma implies that except for the full transformation monoid $\mathcal{O}^{(1)}$, all closed monoids above $S_{\mathbb{N}}$ consist of the injections plus some finite range operations.

Lemma 17. *Let $f \in \mathcal{O}^{(1)}$ have infinite image, and assume it is not injective. Then f generates all unary operations.*

Proof. We skip the fairly easy proof, and refer the reader to the very similar (first part of the) proof of Lemma 20. \square

We thus wish to know, given a finite range operation, which other finite range operations it generates. For that, we need the following concept.

Definition 18. Let $f \in \mathcal{O}^{(1)}$ have finite range, and write $n = |\text{ran}(f)|$. Enumerate the kernel classes of f by C_1, \dots, C_n in such a way that their sizes are increasing. The *kernel tuple* $\kappa^f \in (\omega + 1)^n$ of f is the n -tuple $(|C_1|, \dots, |C_n|)$.

Note that the last entry of a kernel tuple κ^f always equals ω since f must have at least one infinite kernel class.

Having assigned a finite sequence with positive values in $\omega + 1$ to every finite range operation, we are ready to order such sequences and give the definition of \mathbb{P} .

Definition 19.

- Let $k, n \geq 1$. For $a \in (\omega + 1)^k$ and $b \in (\omega + 1)^n$ we write $a \sqsubseteq b$ iff $k \leq n$ and the following holds: There exists a partition $\{A_1, \dots, A_k\}$ of $\{1, \dots, n\}$ into k classes such that $a_i \leq \sum_{j \in A_i} b_j$ for all $1 \leq i \leq k$.
- We write \mathbb{P}_∞ for the partial order of finite increasing sequences of non-zero elements of $\omega + 1$ ordered by \sqsubseteq .
- We write \mathbb{P} for the partial order of the finite increasing sequences of positive natural numbers (not of values in $\omega + 1$!) ordered by \sqsubseteq .

Observe that for finite increasing sequences $a, b \in (\omega + 1)^n$ of the same length we have $a \sqsubseteq b$ iff $a_i \leq b_i$ for all $1 \leq i \leq n$. The following lemma justifies our definition of \sqsubseteq .

Lemma 20. *Let $f, g \in \mathcal{O}^{(1)}$ have finite range. Then f generates g iff $\kappa^g \sqsubseteq \kappa^f$.*

Proof. Assume first that $\kappa^g \sqsubseteq \kappa^f$. Let κ^f have length n and κ^g have length $k \leq n$. If $k = n$, then $\kappa_i^g \leq \kappa_i^f$ for all $1 \leq i \leq n$. It is then not hard to see that for any finite set $A \subseteq \mathbb{N}$, there exist permutations α, β such that $\beta \circ f \circ \alpha$ agrees with g on A . If $k < n$, then let A_1, \dots, A_k be the partition provided by the definition of \sqsubseteq . Enumerate the kernel classes of f by C_1, \dots, C_n and in such a way that C_j contains κ_j^f elements, for all $1 \leq j \leq n$. Now take any $i, j \in \{1, \dots, n\}$ which are distinct but equivalent with respect to the partition A_1, \dots, A_k . By composing f with a permutation, we may assume that f maps the classes C_i and C_j into the class C_n , and all other classes into classes $\neq C_n$ in such a way that no two classes are mapped into the same class. Then $t_{n-1} := f \circ f$ is a function with $n - 1$ values in its range, and $\kappa^g \sqsubseteq \kappa^{t_{n-1}}$. Proceeding like this, we arrive after $n - k$ steps at an operation t_k which takes k values and which satisfies $\kappa_i^g \leq \kappa_i^{t_k}$ for all $1 \leq i \leq k$. Thus we are back in the case $k = n$, and the proof of this direction of the lemma is finished.

For the other direction, assume that f generates g . Let k, n be as before. Since the local clone generated by f is the topological closure of the set of term operations generated by f , we have that for every $q \in \mathbb{N}$, there exists a term t_q consisting of permutations and f which agrees with g on the finite set $\{0, \dots, q\}$. We can write each t_q as $t_q = s_q \circ f \circ \alpha_q$, where s_q consists of permutations and f , and α_q is a permutation. Thus in every term t_q , certain classes of f are joined by the application of s_q (and shifted by α_q , which we do not care about for the moment); since there are only finitely many possibilities of joining classes of f , there is one constellation which appears for infinitely many q . Since $q \leq q'$ implies that $t_{q'}$ agrees with g on $\{0, \dots, q\}$, by replacing terms we may assume that the same classes are joined for all $q \in \mathbb{N}$. Naturally, this partition of classes induces a partition A_1, \dots, A_w on $\{1, \dots, n\}$ via κ^f . If $p, r \leq q$ are equivalent modulo the kernel of g , then the kernel classes of f containing $\alpha_q(p)$ and $\alpha_q(r)$, respectively, are equivalent with respect to the partition A_1, \dots, A_w . On the other hand, if $p, r \leq q$ are not in the same class of g , then $\alpha_q(p)$ and $\alpha_q(r)$ will not lie in equivalent f -classes. Thus, taking q large enough so that $\{0, \dots, q\}$ meets all kernel classes of g , we can assign to every g -class (with, say, index $1 \leq u \leq k$) an equivalence class $A_e \in \{A_1, \dots, A_w\}$ in an injective way; in particular, $k \leq w$. For infinitely many q , this assignment is the same; again, by replacing terms where necessary, we may assume it is always the same. Since for large enough q arbitrarily large parts of the kernel classes of g are

hit, we must have $\kappa_u^g \leq \sum_{j \in A_e} \kappa_j^f$ for all $1 \leq u \leq k$, where A_e is the class assigned to u . By joining some classes A_i , we can obtain $w = k$ without changing the latter fact. \square

Lemma 21. *Let $\mathcal{F} \subseteq \mathcal{O}^{(1)}$. If there is no finite bound to the sizes of the ranges of the finite range operations in \mathcal{F} , then \mathcal{F} generates $\mathcal{O}^{(1)}$.*

Proof. Let f be any finite range operation, and let $A \subseteq \mathbb{N}$ be finite. Then there exists a finite range function g which agrees with f on A and whose kernel sequence κ^g has only one entry equal to ω . Now there exists a finite range operation h in \mathcal{F} with $\kappa^g \sqsubseteq \kappa^h$, so g is generated by \mathcal{F} . This proves that f is generated by \mathcal{F} , and hence \mathcal{F} generates all finite range operations. Clearly, any operation in $\mathcal{O}^{(1)}$ can be interpolated on any finite set by a finite range operation, which implies our assertion. \square

The following is a consequence of Lemma 20; it says that if finitely many finite range operations join forces, the joint generating power is not more than the sum of the generating powers of the single operations.

Proposition 22. *Let $\mathcal{F} \subseteq \mathcal{O}^{(1)}$ be finite. Then $\langle \mathcal{F} \rangle = \bigcup \{ \langle f \rangle : f \in \mathcal{F} \}$.*

Proof. The non-trivial direction is to show $\langle \mathcal{F} \rangle \subseteq \bigcup \{ \langle f \rangle : f \in \mathcal{F} \}$. If t is any term made of operations in \mathcal{F} , then it is of the form $r \circ f$, where r is a term and $f \in \mathcal{F}$. Clearly, $\kappa^t \sqsubseteq \kappa^f$, implying $t \in \langle f \rangle$. Thus, $\bigcup \{ \langle f \rangle : f \in \mathcal{F} \}$ contains all terms that can be built from \mathcal{F} . This implies that the union is a monoid. Being a finite union of (topologically) closed sets $\bigcup \{ \langle f \rangle : f \in \mathcal{F} \}$ is itself closed, and hence contains even $\langle \mathcal{F} \rangle$. \square

We now assign ideals of \mathbb{P} to local monoids containing $S_{\mathbb{N}}$.

Definition 23. For a local monoid $\mathcal{M} \supseteq S_{\mathbb{N}}$, we set

$$\text{Id}_{\infty}(\mathcal{M}) = \{ \kappa^f : f \in \mathcal{M}, f \text{ has finite range} \} \subseteq \mathbb{P}_{\infty}.$$

For a sequence $s \in \mathbb{P}$, write $s * \infty \in \mathbb{P}_{\infty}$ for the sequence obtained by gluing ω to the end of s . Now set

$$\text{Id}(\mathcal{M}) = \{ s \in \mathbb{P} : s * \infty \in \text{Id}_{\infty}(\mathcal{M}) \}.$$

By Lemma 20, $\text{Id}(\mathcal{M})$ is always an ideal of \mathbb{P} . Conversely, we show in the following how to get closed monoids from ideals of \mathbb{P} .

Definition 24. Let $k \geq 1$ and let $(s^n)_{n \in \omega}$ be an ascending sequence of k -tuples in \mathbb{P}_{∞} . We write $\lim_n(s^n)$ for the smallest (according to \sqsubseteq) k -tuple s in \mathbb{P}_{∞} satisfying $s^n \sqsubseteq s$ for all $n \in \omega$.

For an ideal $I \subsetneq \mathbb{P}$, we let $\text{Mon}(I)$ contain all operations in $\mathcal{O}^{(1)}$ which are either injective, or which have finite range and whose kernel sequence is a limit of an ascending sequence of k -tuples of the form $s * \infty$, where $s \in I$.

Lemma 25. *Let $I \subsetneq \mathbb{P}$ be an ideal. Then $\text{Mon}(I)$ is a local monoid containing $S_{\mathbb{N}}$.*

Proof. Observe that if $f \in \text{Mon}(I)$ is a finite range function, and if $g \in \mathcal{O}^{(1)}$ is a finite range function such that $\kappa^g \sqsubseteq \kappa^f$, then $\text{Mon}(I)$ contains also g . For, let the range of f have n elements, let the range of g have $k \leq n$ elements, and let A_1, \dots, A_k be the partition of $\{1, \dots, n\}$ provided by the definition of \sqsubseteq . A quick check shows that we may assume $A_k = \{n\}$. Let $(s^i)_{i \in \omega}$ be the sequence of $(n-1)$ -tuples in I such that $\lim_i(s^i * \infty) = \kappa^f$. Set $t_j^i := \sum_{r \in A_j} s_r^i$, for all $i \in \omega$ and all $1 \leq j \leq k-1$. Since $t^i \sqsubseteq s^i$, all t^i are elements of I , provided they are actually increasing tuples. Fixing i , define inductively $w_j^i := \max\{t_r^i : r \leq j\}$, for

all $1 \leq j \leq k-1$. It is not hard to see that the w^i are in I , too, and that κ^g is the limit of the increasing sequence $(w^i * \infty)_{i \in \omega}$, so $g \in \text{Mon}(I)$.

Using Lemma 20 and Proposition 22, one now readily derives from this that $\text{Mon}(I)$ is indeed a monoid. It remains to show that $\text{Mon}(I)$ is local. Let $f \in \mathcal{O}^{(1)}$, and assume it can be interpolated on all finite sets by operations from $\text{Mon}(I)$; assume also that it is not injective. If f has infinite range, then $\text{Mon}(I)$ must contain non-injections of arbitrarily large finite range, which in turn implies that I contains tuples of arbitrary length. The definition of \mathbb{P} then shows that $I = \mathbb{P}$, a contradiction. Thus, all non-injections in the local closure of $\text{Mon}(I)$ have finite range. Assume again that f is such a non-injection, and assume κ^f has length k . Fix for every set $\{0, \dots, n\}$ an operation $f_n \in \text{Mon}(I)$ which agrees with f on this set. From some n on, the f_n will have to take at least k values, so we take the liberty of assuming that all f_n have this property. An easy manipulation of the f_n using the fact that $g \in \text{Mon}(I)$ whenever $\kappa^g \sqsubseteq \kappa^{f_n}$ allows us to assume that every f_n takes exactly k values and that $\kappa^{f_n} \sqsubseteq \kappa^f$ for all $n \in \omega$. By thinning out the sequence, we may also assume that the kernel sequences of the f_n are increasing with respect to \sqsubseteq . We then have $\lim_n(\kappa^{f_n}) = \kappa^f$. Now replace each k -tuple κ^{f_n} by a $(k-1)$ -tuple $s^n \in \mathbb{P}$ with $s^n * \infty \sqsubseteq \kappa^{f_n}$ in such a way that $\lim_n(s^n * \infty) = \lim_n(\kappa^{f_n})$ and that the sequence $(s^n)_{n \in \omega}$ is still ascending. Clearly, $s^n \in I$ for all $n \in \omega$, and so $\kappa^f = \lim_n s^n$ implies $f \in \text{Mon}(I)$. Hence, $\text{Mon}(I)$ is locally closed. \square

Proposition 26. *The mapping $\sigma : \mathcal{M} \mapsto \text{Id}(\mathcal{M})$ is an isomorphism from the lattice \mathfrak{U} of locally closed monoids that contain $S_{\mathbb{N}}$ onto the lattice of ideals of \mathbb{P} .*

Proof. That $\text{Id}(\mathcal{M})$ is an ideal of \mathbb{P} for all monoids \mathcal{M} follows directly from Lemma 20. From Lemma 21 we know that $\text{Id}(\mathcal{M}) = \mathbb{P}$ iff $\mathcal{M} = \mathcal{O}^{(1)}$, and obviously $\text{Id}(\mathcal{M}) = \emptyset$ iff \mathcal{M} contains only injections. We have seen in Lemma 25 that for any proper ideal I of \mathbb{P} , $\text{Mon}(I)$ is a local monoid, and a straightforward verification shows $\text{Id}(\text{Mon}(I)) = I$, thus σ is onto. Also, an easy check using Lemma 20 shows that $\text{Mon}(\text{Id}(\mathcal{M})) = \mathcal{M}$ for every local monoid which contains $S_{\mathbb{N}}$ and all of whose non-injections have finite range, so σ is injective. It is obvious that both σ and σ^{-1} are order-preserving. \square

Recall that a lattice is *algebraic* iff it is isomorphic to the subalgebra lattice of an algebra.

Corollary 27. *The lattice \mathfrak{U} of local monoids above $S_{\mathbb{N}}$ is distributive. Moreover, \mathfrak{U} and its dual order are algebraic.*

Proof. By the preceding proposition, \mathfrak{U} is the lattice of ideals of a partial order. The assertions then follow from [CD73, p. 83]. \square

Definition 28. A partial order is called a *well-quasi-order* iff there are no infinite descending chains and no infinite antichains in it.

We call a sequence $(a^n)_{n \in \omega}$ in a partial order with order relation \leq *bad* iff for no $i < j \in \omega$ we have $a^i \leq a^j$.

A standard application of the infinite Ramsey's theorem shows that a partial order is well-quasi-ordered iff it contains no bad sequence (confer e.g. [Die05]).

Lemma 29. *The set of finite sequences with values in $\omega + 1$ ordered by \sqsubseteq is a well-quasi-order. In particular, its suborders \mathbb{P}_{∞} and \mathbb{P} are well-quasi-orders.*

Proof. Assume that $(a^n)_{n \in \omega}$ were a bad sequence of such finite sequences. For every $n \in \omega$, let w_n be the number of occurrences of ω in the tuple a^n . If the sequence $(w_n)_{n \in \omega}$ were unbounded, then we could find $n \in \omega$ such that w_n is greater than the length of a^0 , implying $a^0 \sqsubseteq a^n$, a contradiction. Thus, we can thin out the

sequence in such a way that all w_n are equal. Let b^n be the tuple obtained from a^n by leaving away the w_n components equal to ω , for all $n \in \omega$. Clearly, the b^n form a bad sequence of tuples with values in ω . Now if the sequence of lengths of the b^n were unbounded, then we could find some $i \in \omega$ such that the length of b^i is greater than the sum of all components of b^0 , hence $b^0 \sqsubseteq b^i$, a contradiction. Thin out the sequence $(b^n)_{n \in \omega}$ in such a way that all tuples have the same length k . Now for all $1 \leq j \leq k$, we thin out our sequence so that the sequence consisting of the j -th component of the b^n is increasing; we can do this since ω is well-ordered. The remaining sequence of b^n is ascending, a contradiction. \square

In general, the ideal lattice of a given well-quasi-order need not be a well-quasi-order. Certain well-quasi-orders which satisfy a certain strong combinatorial property and which are called *better-quasi-orders*, however, do have the property that their ideal lattice is well-quasi-ordered. Although giving the definition of a better-quasi-order (see e.g. [Mil85]) would be out of scope of the present paper, we remark that it follows from the basic theory of better-quasi-orders that our well-quasi-order \mathbb{P} is in fact a better-quasi-order (as Kruskal states in [Kru72]: “All naturally known well-quasi-ordered sets which are known are better-quasi-ordered.”). Therefore, the lattice of ideals of \mathbb{P} is a well-quasi-order (and, in fact, even a better-quasi-order as well). In order to spare the reader the pain of reading the definition of a better-quasi-order, we prove the following

Proposition 30. *The lattice of ideals of \mathbb{P} is a well-quasi-order.*

Proof. Suppose that $(I_n)_{n \in \omega}$ were a bad sequence of ideals. By taking away the first element of the sequence, we may assume that $I_n \neq \mathbb{P}$ for all $n \in \omega$. Consider for all $n \in \omega$ the set J_n consisting of I_n plus all limits of ascending chains of tuples of the same length in I_n (as in Definition 24): So $J_n \subseteq \mathbb{P}_\infty$ for all $n \in \omega$. In J_n , every chain is bounded from above: If C is a chain in J_n , then there is some $k \in \omega$ such that all tuples in C have length at most k ; for otherwise, I_n contains sequences of arbitrary length, implying $I_n = \mathbb{P}$ contrary to our assumption. But if the elements of C all have length at most k , then C is bounded by construction of J_n (i.e., adding the limits of ascending chains).

Applying Zorn’s lemma, we get that every element of I_n is below some maximal tuple of J_n . By construction of J_n , every tuple in \mathbb{P} which is below a maximal tuple of J_n also is an element of I_n . The maximal elements of J_n form an antichain with respect to \sqsubseteq . By Lemma 29, the set of sequences in $\omega + 1$, equipped with the order \sqsubseteq , is a well-quasi-order. In particular, the antichain of maximal elements of J_n is finite. For every $n > 0$, there exists a tuple in $I_0 \setminus I_n$. Thus, there is a maximal tuple of J_0 which is not in J_n . We can thin out our sequence of I_n so that this witnessing maximal tuple is the same for all $n > 0$; denote it by s^0 . Now we do the same for J_1 and all $n > 1$, obtaining a maximal tuple $s^1 \in J_1$ which is not in any J_n with $n > 1$. We continue inductively in this fashion, obtaining a sequence $(s^n)_{n \in \omega}$. By construction, this sequence is a bad sequence in the order of finite $(\omega + 1)$ -valued sequences with \sqsubseteq , a contradiction. \square

Corollary 31. *The lattice of local monoids above $S_{\mathbb{N}}$ is well-quasi-ordered.*

Proposition 32. *Every local monoid \mathcal{M} above $S_{\mathbb{N}}$ is finitely generated over $S_{\mathbb{N}}$, i.e., there exists a finite $\mathcal{F} \subseteq \mathcal{O}^{(1)}$ such that $\mathcal{M} = \langle \mathcal{F} \rangle$. Moreover, the number of such monoids is countable.*

Proof. If $\mathcal{M} = \mathcal{O}^{(1)}$, then it is generated by any non-injective operation with infinite range, by Lemma 17. Assume henceforth that \mathcal{M} contains only injections and finite range operations. Set J to consist of all kernel sequences of operations in \mathcal{M} . By local closure and what is by now a standard argument, one sees that every chain

in J has an upper bound in J (confer e.g. the proof of Lemma 25). Hence, Zorn's lemma implies that every element of J is below a maximal element of J . The maximal elements of J form an antichain with respect to \sqsubseteq , and therefore are finite in number by Lemma 29. Pick for each maximal tuple one corresponding operation in \mathcal{M} . The set \mathcal{F} of operations thus chosen is as desired, by Lemma 20.

The above argument shows that every local monoid containing $S_{\mathbb{N}}$ is determined by a finite set of finite sequences with values in $\omega+1$. There are only countably many possibilities for such finite sets, so the number of such monoids is countable. \square

5. CLONES WITH ESSENTIAL FINITE RANGE OPERATIONS

Having understood the structure of the lattice of local monoids containing $S_{\mathbb{N}}$, we move on to describe the monoidal interval corresponding to each such monoid. In this section, we will prove Theorem 8, which deals with those monoids which contain a non-constant and non-injective operation.

Definition 33. An operation f on a set Y is called *essential* iff it is not essentially unary, i.e., it depends on at least two of its variables.

We will see that except for \mathcal{O} , all clones whose unary fragment contains a non-constant and non-injective operation contain only essential operations with finite range, and none with infinite range; hence the title of this section.

It turns out that the monoidal intervals of the monoids under consideration here are all chains which can be described nicely. This is essentially a consequence of a theorem from [HR94] for clones on finite sets, and the power of local closure. In order to state that theorem, we need the following definition.

Definition 34. Let Y be any set. We call an n -ary operation f on Y *quasilinear* iff there exist functions $\phi_0 : \{0, 1\} \rightarrow Y$ and $\phi_1, \dots, \phi_n : Y \rightarrow \{0, 1\}$ such that $f(x_1, \dots, x_n) = \phi_0(\phi_1(x_1) \dot{+} \dots \dot{+} \phi_n(x_n))$, where $\dot{+}$ denotes the sum modulo 2.

Theorem 35 ([HR94]). *Let Y be a finite set of at least three elements and let f be an essential operation on Y . Set $k := |\text{ran}(f)|$. Then, writing S_Y for the set of all permutations on Y , we have:*

- *If $k \geq 3$, then f together with S_Y generate all operations on Y which take at most k values.*
- *If $k = 2$ and f is not quasilinear, then f together with S_Y generate all operations on Y which take at most two values.*
- *If $k = 2$, f is quasilinear, and $|Y|$ is odd, then f together with S_Y generate all quasilinear operations on Y .*

Observe that there is no such thing as local interpolation on finite Y , so “generates” in the theorem refers to the term closure.

The following lemma is the infinite local version of this theorem.

Lemma 36. *Let $f \in \mathcal{O}^{(n)}$ be essential, and assume that $|\text{ran}(f)| = k$, where $2 \leq k < \omega$.*

- (1) *If $k \geq 3$, then f generates all operations which take at most k values.*
- (2) *If $k = 2$ and f is not quasilinear, then f generates all functions which take at most k values.*
- (3) *If $k = 2$ and f is quasilinear, then f generates all quasilinear operations.*

Proof. (1): Let any operation g on $|\text{rang}| \leq k$ be given. It suffices to show that for every $n \in \omega$, f generates an operation h which agrees with g on $\{0, \dots, n\}$. Choosing n large enough, we may assume that the ranges of both f and g are contained in $\{0, \dots, n\}$. Also, again by making n larger, we may assume that the restriction f' of f to $\{0, \dots, n\}$ is essential and takes k values. Then f' is an

essential operation on $\{0, \dots, n\}$ which takes k values, and hence generates all such functions by Theorem 35. In particular, f' generates a function h' which agrees with g on $\{0, \dots, n\}$. The permutations which appear in the term which represents h' can be extended to X by the identity; occurrences of f' can be replaced by f . The resulting term h is a function on X which still agrees with g on $\{0, \dots, n\}$.

(2): Again, let g and n be given. Enlarge n as before, if necessary, so that the restriction of f to $\{0, \dots, n\}$ is not quasilinear. Now we argue as in the preceding proof.

(3): The proof works as before; n only has to be chosen odd in order to allow application of Theorem 35. \square

With the preceding lemma, we see that it is quite easy to understand what happens when we add an essential finite range operation to a monoid. We will now show that to the monoids relevant for this section, we in fact cannot add an essential infinite range operation without generating \mathcal{O} . To establish this, we distinguish between those operations which preserve the binary inequality relation \neq , and those which do not. The latter case can be eliminated right away:

Proposition 37. *Let f be an essential operation with infinite image. Then f preserves \neq , or it generates all operations.*

The proposition will follow from the following lemma.

Lemma 38. *Let $f \in \mathcal{O}^{(n)}$ have infinite image, and assume it does not preserve \neq . Then f generates a unary non-injective function that has infinite range.*

Before proving the lemma we show how the proposition follows from it:

Proof of Proposition 37. By Lemma 38, f generates a unary non-injective operation with infinite range; Lemma 17 then implies that f generates all unary operations. Let $k \geq 1$ be arbitrary. We can find a unary finite range operation g_k such that $g_k(f)$ is essential and takes at least k values. By Lemma 36, this implies that f generates all operations which take not more than k values. Since k was arbitrary and by local interpolation, this implies that f generates \mathcal{O} . \square

Proof of Lemma 38. We only have to prove something if f is essential, for f is itself a unary non-injective operation with infinite range otherwise. We also assume f to depend on all of its variables.

Case 1: There exist injective functions $g_1, \dots, g_n \in \mathcal{O}^{(1)}$ such that $f(g_1, \dots, g_n)$ is injective. Choose $S \subseteq X$ infinite such that $X \setminus S$ and $X \setminus (g_1[S] \cup \dots \cup g_n[S])$ are infinite. Since f does not preserve \neq , there exist $c, d \in X^n$ such that $c_i \neq d_i$ for all $1 \leq i \leq n$ and such that $f(c) = f(d)$. We may assume that those values are not in $S \cup g_1[S] \cup \dots \cup g_n[S]$. For all $1 \leq i \leq n$, set $g'_i = g_i$ on S and $g'_i(c_1) = c_i$ and $g'_i(d_1) = d_i$. Write $T = S \cup \{c_1, d_1\}$. On T , $F = f(g'_1, \dots, g'_n)$ is a function with infinite range which is not injective. Furthermore, the g'_i can be extended to permutations on X since they are injective and have co-infinite domains and ranges. This completes the first case.

Case 2: $f(g_1, \dots, g_n)$ is not injective for all injective $g_1, \dots, g_n \in \mathcal{O}^{(1)}$. Since f is not constant, there exist injections $u_1, \dots, u_n \in \mathcal{O}^{(1)}$ such that $f(u_1, \dots, u_n)$ is not constant. By assumption for this case we have that $f(u_1, \dots, u_n)$ is not injective. Thus, $f(u_1, \dots, u_n)$ generates a function $r \in \mathcal{O}^{(1)}$ which is constant except for one argument, at which it takes another value. To see this, just observe that for such an r , the kernel sequence κ^r satisfies $\kappa^r \sqsubseteq \kappa^h$ for any non-constant $h \in \mathcal{O}^{(1)}$, and apply Lemma 20. Consider arbitrary $g_1, \dots, g_n \in \mathcal{O}^{(1)}$ such that $f(g_1, \dots, g_n)$ is injective. Such functions exist since f has infinite range. Pick an infinite $S \subseteq X$ on which each g_i is either injective or constant. Observe that it is impossible that all

those restrictions of g_i are constant. Say the restriction of g_i to S is constant with value c_i for $1 \leq i \leq k$, where $1 \leq k < n$.

Since f depends on its first variable, there exist $c_{k+1}, \dots, c_n \in X$ and $a_1, \dots, a_k \in X$ such that $f(a_1, \dots, a_k, c_{k+1}, \dots, c_n) \neq f(c)$. For otherwise the value of f is determined by the values of the arguments x_{k+1}, \dots, x_n , contradicting that f depends on x_1 . Choose $d_{k+1}, \dots, d_n \in X$ such that $f(c_1, \dots, c_k, d_{k+1}, \dots, d_n)$ is distinct from both $f(a_1, \dots, a_k, c_{k+1}, \dots, c_n)$ and $f(c)$; this is possible as $F(x_{k+1}, \dots, x_n) = f(c_1, \dots, c_k, x_{k+1}, \dots, x_n)$ takes infinitely many values. Set $u_i(0) = a_i$ and $u_i(x) = c_i$ if $x \neq 0$, for $1 \leq i \leq k$. Let moreover $v_i(1) = d_i$, and $v_i(x) = c_i$ if $x \neq 1$, for all $k+1 \leq i \leq n$. All u_i and v_i are generated by the two-valued function r and hence by f . Set $F(x) = f(u_1, \dots, u_k, v_{k+1}, \dots, v_n)(x)$. We have that $F(0) = f(a_1, \dots, a_k, c_{k+1}, \dots, c_n)$ and $F(1) = f(c_1, \dots, c_k, d_{k+1}, \dots, d_n)$ and $F(2) = f(c)$ are pairwise distinct; also, F has finite range. Hence, f generates a unary function F that takes exactly three values.

Pick $a, b \in X^n$ such that $a_i \neq b_i$ for all $1 \leq i \leq n$ and such that $f(a) = f(b)$. For all $1 \leq i \leq k$, let $h_i \in \mathcal{O}^{(1)}$ satisfy $h_i(0) = a_i$, $h_i(1) = b_i$, and $h_i(x) = c_i$ for all $x \notin \{0, 1\}$. Clearly, h_i is generated by the unary function F with three values, and hence also by f . For all $k+1 \leq i \leq n$, let α_i be any permutation that maps 0 to a_i and 1 to b_i . Now we set $G(x) = f(h_1(x), \dots, h_k(x), \alpha_{k+1}(x), \dots, \alpha_n(x))$ and have that $G(0) = f(a) = f(b) = G(1)$ and G takes infinitely many values. \square

We are thus left with essential infinite range operations which do preserve the inequality relation. The crucial theorem here has been shown in [BK08a].

Theorem 39. *Every essential operation $f \in \text{Pol}(\neq)$ generates a binary injective operation.*

Able to produce a binary injection, we use this operation in order to enlarge ranges of unary operations:

Lemma 40. *Let $f \in \mathcal{O}^{(2)}$ be injective, and let $g \in \mathcal{O}$ be a non-constant function with finite range. Then f and g together generate \mathcal{O} .*

Proof. Write $k := |\text{ran}(g)|$. Either g is itself (essentially) unary, or it generates a non-constant unary operation which takes k values by Lemma 36. Pick such a unary operation g_0 , and set $k_0 := k$. The operation $t(x, y) := f(g(x), g(y))$ takes $k_1 := (k_0)^2$ values. Hence, f and g together also generate a unary operation g_1 which takes k_1 values. Continuing in this fashion, and since $k_0 = k > 1$, we obtain unary operations of all finite ranges, and hence also $\mathcal{O}^{(1)}$ by Lemma 20. But clearly, for every operation $h \in \mathcal{O}^{(2)}$, there exists a unary $h' \in \mathcal{O}^{(1)}$ such that $h(x, y) = h'(f(x, y))$. Thus, since h was arbitrary, f and g generate $\mathcal{O}^{(2)}$, and in turn also \mathcal{O} , as it is well-known and easy to see that $\langle \mathcal{O}^{(2)} \rangle = \mathcal{O}$. \square

We are thus ready to prove Theorem 8.

Proof of Theorem 8. We are given a monoid \mathcal{M} which contains an operation which is neither constant nor injective. It follows from Lemma 40 that no clone properly contained in \mathcal{O} and containing \mathcal{M} can have a binary injection. Therefore, by Theorem 39, it cannot have any essential operation which preserves the inequality relation. Nor can it contain an essential operation with infinite range which does not preserve the inequality, by Proposition 37. Thus, it cannot contain any essential operation with infinite range. We now distinguish the two cases of the theorem:

(1): If $\mathcal{M} = \mathcal{O}^{(1)}$, then the only clones above $\mathcal{O}^{(1)}$ can be the ones mentioned in the theorem, by Lemma 36. That these sets of operations are actually clones is a straightforward verification and left to the reader.

(2): If $\mathcal{M} \neq \mathcal{O}^{(1)}$, then there is a largest natural number k such that \mathcal{M} contains all

unary operations which take at most k values; this follows from Lemma 21. If $k = 1$, then no clone having \mathcal{M} as its unary fragment can have an essential operation, since this operation would generate all quasilinear operations by Lemma 36, and hence all unary operations with at most two values. Consider thus the case $k > 1$. Again by Lemma 36, the only clones in $I_{\mathcal{M}}$ can be $\langle \mathcal{M} \rangle$, $\mathcal{Q}^{\mathcal{M}}$ and the $\mathcal{K}_n^{\mathcal{M}}$, where $n \leq k + 1$. It is easy to verify that these are indeed clones. \square

6. CLONES WITH ESSENTIAL INFINITE RANGE OPERATIONS

This section deals with the monoidal interval of $I_{\mathcal{J}}$, the monoid of all injections, and contains the proof of Theorem 13. Before we start with an outline of this section, let us observe that as a straightforward verification shows, the largest element of this interval, $\text{Pol}(\mathcal{J})$, equals $\text{Pol}(\neq)$ (this is statement (1) of Theorem 13). In particular, all operations of clones in this interval have infinite range; hence the title of this section.

In Subsection 6.1 we prove that \mathcal{H} is the unique cover above $\langle \mathcal{J} \rangle$ in $I_{\mathcal{J}}$, and that $\text{Inv}(\mathcal{H})$ is the set of all relations definable by a Horn formula (Theorem 13, part (2)).

In Subsection 6.2 we prove that \mathcal{B} is the unique cover of \mathcal{H} in $I_{\mathcal{J}}$ (Theorem 13, part (3)).

In Subsection 6.3 we present an infinite strictly decreasing chain of clones that contain \mathcal{B} , and are contained in \mathcal{J} . To this end, we also give relational descriptions of \mathcal{B} and of \mathcal{J} .

In Subsection 6.4 we show that \mathcal{R} and \mathcal{J} are incomparable, and that every clone in $I_{\mathcal{J}}$ is either contained in \mathcal{J} or contains \mathcal{R} (Theorem 13, part (4)).

In Subsection 6.5 we show that the power set of ω embeds into the intervals of those clones in $I_{\mathcal{J}}$ that contain \mathcal{R} ; in particular, the size of this interval equals the continuum (Theorem 13, part (5)).

6.1. The Horn clone. In this subsection we show that $\langle \mathcal{J} \rangle$ has the cover \mathcal{H} in the monoidal interval $I_{\mathcal{J}}$.

Recall from Definition 9 that \mathcal{H} is the clone consisting of all operations f that are essentially injective, i.e., all operations f that are the composition $i(p_1, \dots, p_n)$ of an injective operation i with projections p_1, \dots, p_n (it is straightforward to verify that this set of operations indeed forms a locally closed clone). The clone \mathcal{H} is also called the *Horn clone*; the reason for this name will be given in Proposition 43.

Definition 41. Let Φ be a quantifier-free first-order formula where all atomic subformulas are of the form $x = y$. Then Φ is called *Horn* if Φ is in conjunctive normal form (henceforth abbreviated CNF) and each clause in Φ contains at most one positive literal.

Definition 42. Let $\phi(x_1, \dots, x_n)$ be a formula in CNF. We call ϕ *reduced* iff it is not logically equivalent to any of its subformulas, i.e., there is no formula ψ obtained from ϕ by deleting literals or clauses such that $\phi(a)$ iff $\psi(a)$ for all $a \in \mathbb{N}^n$.

Clearly, for every formula ϕ in CNF there exists a logically equivalent reduced formula.

In the following proposition, the equivalence of (1) and (7) proves item (2) of Theorem 13, stating that the Horn clone is the unique cover of $\langle \mathcal{J} \rangle$ in $I_{\mathcal{J}}$. Moreover, items (5) and (6) provide finite relational generating systems of $\text{Inv}(\mathcal{H})$, and (2) provides a finite operational generating system (in fact: a continuum of such systems) of \mathcal{H} . Recall that a formula is *Horn* iff it is in conjunctive normal form and each of its clauses contains at most one positive literal. Item (4) gives us a syntactical description of the formulas defining relations in $\text{Inv}(\mathcal{H})$.

Proposition 43. *For all relations R with a first-order definition in $(\mathbb{N}, =)$ the following are equivalent.*

- (1) R is preserved by an essential operation that preserves \neq .
- (2) R is preserved by a binary injective operation.
- (3) Every reduced definition of R is Horn.
- (4) R has a Horn definition.
- (5) R has a pp definition in (\mathbb{N}, \neq, I) where

$$I := \{(a, b, c, d) \in \mathbb{N}^4 \mid a = b \rightarrow c = d\}$$

- (6) R has a pp-definition in (\mathbb{N}, N) where

$$N := \{(a, b, c, d) \in \mathbb{N}^4 \mid a = b \neq c = d \vee |\{a, b, c, d\}| = 4\}$$

- (7) R is preserved by \mathcal{H} .

Proof. The implication from (1) to (2) is Theorem 39.

For the implication from (2) to (3), suppose that Φ is a reduced formula that defines R but is not Horn. Then there exists a clause ψ of Φ which contains two equalities $x_i = x_j$ and $x_k = x_l$. Construct ψ' from ψ by removing the equation $x_i = x_j$, and ψ'' by removing $x_k = x_l$. Since Φ is reduced, there exist $a, b \in \mathbb{N}^n$ such that $\phi(a)$ but not $\phi'(a)$, and $\phi(b)$ but not $\phi''(b)$. Clearly, $a_i = a_j$, $a_k \neq a_l$, $b_i \neq b_j$, and $b_k = b_l$. Set $c = f(a, b)$, where f is the binary injection preserving R . Then $c_i \neq c_j$, $c_k \neq c_l$, and in fact $\phi(c)$ does not hold. Hence R is not preserved by f , a contradiction.

The implication from (3) to (4) is trivial.

For the implication from (4) to (5), let Φ be a Horn formula. It suffices to show that all clauses ψ of Φ have a pp definition in (\mathbb{N}, I, \neq) . If ψ is of the form $(u_1 = v_1 \wedge \dots \wedge u_l = v_l) \rightarrow u = v$, consider the following pp formula.

$$(1) \quad \exists w_1, \dots, w_{l+1}. I(w_1, w_l, u, v) \wedge \bigwedge_{1 \leq i \leq l} I(u_i, v_i, w_i, w_{i+1})$$

Assume that $u_i = v_i$ for all $1 \leq i \leq l$. In this case, the pp formula implies that $w_1 = w_2 = \dots = w_{l+1}$, and hence also implies that $u = v$. Now, if $u_i \neq v_i$ for some $1 \leq i \leq l$, then for all choices of values for the other free variables the formula can be satisfied by setting w_1, \dots, w_{l+1} to values that are distinct from all other values. Hence, the formula is a pp definition of ψ in (\mathbb{N}, I) . If ψ does not contain a positive literal, consider the formula $\exists u, v. (\alpha \vee u = v) \wedge u \neq v$, which is equivalent to ψ (we assume that u and v are fresh variables). We have seen above that the term in brackets is equivalent to a pp formula. It is then straightforward to rewrite the whole expression as a pp formula.

For the implication from (5) to (6), it suffices to show that I and \neq have pp definitions in (\mathbb{N}, N) . For \neq , this is obvious. To express $a = b \rightarrow c = d$, consider the pp formula

$$\exists u, v, w. N(a, b, u, v) \wedge N(a, b, v, w) \wedge N(u, w, c, d).$$

If $a = b$, the variables u, v, w must denote the same value, and hence the formula implies $c = d$. If $a \neq b$, then for all choices of values for c and d it is possible to select values for u, v, w that satisfy the formula. Hence, the above formula is a pp definition of I in (\mathbb{N}, N) .

It is straightforward to verify the implication from (6) to (7), because N is preserved by projections and by injective operations.

The implication from (7) to (1) is immediate, since every at least binary injective operation is essential and has infinite image. \square

6.2. The bar clone. We show that the bar clone \mathcal{B} (confer Definition 10) is the smallest clone below $\text{Pol}(\neq)$ that strictly contains \mathcal{H} , thus proving item (3) of Theorem 13.

A smallest non-empty k -ary relation that is preserved by all permutations of \mathbb{N} is called an *orbit of k -tuples*. It is clear that every relation that is preserved by all permutations is the union of a finite number of orbits of k -tuples. The following lemma will be useful here and in the following subsections, and a proof can be found in [BK08b].

Lemma 44. *Let R be a k -ary relation that consists of m orbits of k -tuples. Then every operation f that violates R generates an m -ary operation that violates R .*

The relation N has been introduced in Proposition 43; we recall it for the convenience of the reader: N is the four-ary relation $\{(a, b, c, d) \in \mathbb{N}^4 \mid a = b \neq c = d \vee |\{a, b, c, d\}| = 4\}$.

Lemma 45. *Let $f \in \text{Pol}(\neq) \setminus \mathcal{H}$. Then f generates a binary operation that violates N .*

Proof. The operation f must violate the relation N . Indeed, otherwise we would have $f \in \text{Pol}(\text{Inv}(\mathcal{H})) = \mathcal{H}$, since all relations in $\text{Inv}(\mathcal{H})$ have a primitive positive definition in (\mathbb{N}, N) by Proposition 43 – a contradiction. Since N consists of two orbits of tuples, Lemma 44 implies that f generates a binary operation that violates N . \square

Clearly, the binary operation from Lemma 45 is essential and not injective.

Recall from Definition 10 that a binary function f is called a k -bar function iff there exists an injection $p \in \mathcal{O}^{(2)}$ such that $f(x_1, x_2) = p(x_1, x_2)$ for all $x_1 \geq k$ and $f(x_1, x_2) = p(x_1, 0) = x_1$ for all $x_1 < k$; recall also that \mathcal{B} is the clone generated by any bar function. We show next that \mathcal{B} is well-defined.

Clearly, for fixed $k \geq 1$, the k -bar functions generate each other. It is also easy to see that if $1 \leq n \leq k$, then any k -bar function generates all n -bar functions.

Lemma 46. *Let $k \geq 1$, and let $f \in \mathcal{O}^{(2)}$ be a k -bar function and $h \in \mathcal{O}^{(2)}$ be a $(k+1)$ -bar function. Then f generates h .*

Proof. Let f be a k -bar function, and g be a variant of the 1-bar which satisfies $g(k+1, x_2) = k+1$ for all $x_2 \in \mathbb{N}$, and which is injective otherwise. Clearly, g is generated by f . Set $h(x_1, x_2) = f(x_1, g(x_1, x_2))$. Let $j \leq k$. Then $h(j, x_2) = j$ for all $x_2 \in \mathbb{N}$ since $f(j, z) = j$ for all $z \in \mathbb{N}$. Also, $h(k+1, x_2) = f(k+1, k+1)$ for all $x_2 \in \mathbb{N}$. Moreover, if $a > k+1$ and $b, c, d \in \mathbb{N}$ and $(a, b) \neq (c, d)$, then it is easy to see that $h(a, b) \neq h(c, d)$. Hence, h is essentially a $(k+1)$ -bar function, the only difference to a $(k+1)$ -bar function being the value of $h(k+1, x_2)$, which can be undone by composing any permutation which swaps $k+1$ and $h(k+1, 0)$ with h . \square

The next lemma characterizes the binary operations in \mathcal{B} . Recall the definition of “injective in one direction” from Definition 11.

Lemma 47. *Let $f \in \mathcal{O}^{(2)}$. Then $f \in \mathcal{B}$ iff f is injective in one direction and for all $c \in \mathbb{N}$, each of the unary operations $F(x) = f(x, c)$ and $F'(x) = f(c, x)$ is either constant or injective.*

Proof. Assume that f is of that form, and let $n \geq 2$ be given. We can say without loss of generality that f is injective in the first direction. Write $\{0, \dots, n\}$ as a disjoint union $A \cup B$ in such a way that $f(c, x_2)$ is injective for all $c \in A$, and constant for all $c \in B$. If $|B| = j$, then there exist injections α, β and a j -bar function g such that $\alpha(g(\beta(x_1), x_2))$ agrees with f on $\{0, \dots, n\}$. Hence, $f \in \mathcal{B}$.

We use induction over terms to show the other direction. The statement is true for all bar-functions, including the projections. Assume it holds for $f, g_1, g_2 \in \mathcal{B}$, and consider $t = f(g_1, g_2)$. Then t is injective in one direction (this is easy to see, and also follows from Lemma 65 in Subsection 6.4). Let $c \in \mathbb{N}$ be arbitrary and consider $T(x) = t(x, c)$. By induction hypothesis, we have that $G_1(x) = g_1(x, c)$ and $G_2(x) = g_2(x, c)$ are constant or injective. If both G_1 and G_2 are injective, then so is T , since f is injective in one direction. If on the other hand both G_1 and G_2 are constant, then T is constant as well. Assume therefore without loss of generality that G_1 is injective and G_2 is constant with value d . Now if $F(x) = f(x, d)$ is constant, then so is T , and if F is injective, then the same holds for T again, and we are done. \square

Lemma 48. *Let $f \in \text{Pol}(\neq)^{(2)}$ be essential and non-injective, and let $k \geq 2$. Then f generates a binary function g which is not injective but injective on $k = \{0, \dots, k-1\}$.*

Proof. Without loss of generality we assume $f(c, d) = f(c', d)$ for some $d \in \mathbb{N}$ and distinct $c, c' \in \mathbb{N}$. Pick any binary injection h . By Lemma 43, h is generated by f . Set $f'(x_1, x_2) = h(f(x_1, x_2), x_2)$. Then f' satisfies $f'(c, d) = f'(c', d)$, and is injective in the 2-nd direction. Therefore, replacing f by f' , we may assume that f is injective in the 2-nd direction. We use induction over k to prove the assertion of the lemma.

For the induction beginning, let $k = 2$. We may assume that there exist $s > k$ and distinct $t, t' > k$ with $f(t, s) = f(t', s)$. Indeed, otherwise f is injective on the infinite square of pairs from $\{k+1, \dots\}$, in which case we get the assertion by applying permutations. Since f is essential, we may assume $f(0, 0) \neq f(1, 0)$. We have that $f(0, 1) \notin \{f(0, 0), f(1, 0)\}$ and $f(1, 1) \notin \{f(0, 0), f(1, 0)\}$ as f is injective in the 2-nd direction. Therefore, the only thing in our way to an injection on $\{0, 1\}$ is that possibly $f(0, 1) = f(1, 1)$. In that case, we let a permutation α map $f(0, 1)$ to 0, $f(t, s)$ to s , and set $f'(x_1, x_2) = f(x_1, \alpha(f(x_1, x_2)))$. It is easy to check that f' is injective on 2 and satisfies $f(t, s) = f(t', s)$.

For the induction step, let f be injective on k , where $k \geq 2$. We generate a function g that is injective on $k+1$ and satisfies $g(t, s) = g(t', s)$ for some $s, t, t' \in \mathbb{N}$ with $t \neq t'$. Again, we pick a binary injection h and set $f'(x_1, x_2) = h(f(x_1, x_2), x_2)$. Clearly, f' is still injective on k . Moreover, f' satisfies $f'(c, d) = f'(c', d)$, and is injective in the 2-nd direction. Therefore, replacing f by f' , we may henceforth assume that f is injective in the 2-nd direction.

As in the induction beginning, we may assume that there exist $s > k$ and distinct $t, t' > k$ with $f(t, s) = f(t', s)$. Since f is injective in the 2-nd direction we have $f(t, s) \notin f[(k+1)^2]$.

If $f(k, n) = f(i, n)$ for some $i < k$ and $n \leq k$, then we do the following: Choose a permutation α that exchanges k with some $j < k$ for which $j \neq i$, and which is the identity otherwise. Let β be a permutation that maps $\{f(0, n), \dots, f(k, n)\}$ into k , and $f(t, s)$ to s . Now set $f'(x_1, x_2) = f(\alpha(x_1), \beta(f(x_1, x_2)))$. We have:

- All points $u, v \in \mathbb{N}^2$ that had different values under f still have different values under f' ; in particular, f' is injective in the 2-nd direction, and injective on k .
- $f'(i, n) \neq f'(k, n)$.
- $f'(t, s) = f'(t', s)$.

If we repeat this procedure for all $n \leq k$ for which there exists $i < k$ with $f(k, n) = f(i, n)$, then in the end $f(j, n) \neq f(k, n)$ for all $j < k$ and all $n \leq k$. Now the only non-injectivity that could be left on $k+1$ could be that $f(i, k) = f(j, k)$ for some distinct $i, j < k$. One gets rid of this setting a last time $g(x_1, x_2) =$

$f(x_1, \beta(f(x_1, x_2)))$, where β is a permutation that maps $\{f(0, k), \dots, f(k, k)\}$ into k , and $f(t, s)$ to s . \square

Lemma 49. *Let $f \in \text{Pol}(\neq)^{(2)}$ be essential and non-injective. Then it generates a 1-bar function.*

Proof. Fix a 1-bar function b . We show that for arbitrary $k \geq 2$, f generates a function which agrees with b on k . Define a finite sequence of natural numbers $(m_j)_{j \leq k}$ by setting $m_0 = k$, and $m_{j+1} = (m_j)^2$. By Lemma 48 and the use of permutations, we can produce a function g which is injective on $\{1, \dots, m_k - 1\} \times m_k$ and satisfies $g(0, c) = g(0, c') = c$ for some distinct $c, c' \in \mathbb{N}$. Moreover, we may assume that $g[\{1, \dots, m_{k-1} - 1\} \times m_{k-1}] \subseteq m_k$. We first want to produce a function g' which is injective on $\{1, \dots, m_{k-1} - 1\} \times m_{k-1}$ and satisfies $g'(0, 0) = g'(0, c) = g'(0, c') = c$. There is nothing to show if $g(0, 0) = c$. Otherwise, we may assume $g(0, 0) = c'$, and set $g'(x_1, x_2) = g(x_1, g(x_1, x_2))$; it is easy to check that g' has the desired properties. Replace g by g' . By repeating this procedure we can produce a function g' which is injective on $\{1, \dots, m_{k-2} - 1\} \times m_{k-2}$ and satisfies $g'(0, 1) = g'(0, 0) = g'(0, c) = g'(0, c') = c$. Doing the same k times, we get g' which is injective on $\{1, \dots, m_0 - 1\} \times m_0$ and satisfies $g'(0, k-1) = \dots = g'(0, 0)$. Clearly, there exists a permutation α such that $\alpha(g')$ agrees with b on k^2 . \square

The following summarizes the results about \mathcal{B} obtained so far. Relational descriptions of \mathcal{B} will be given in Proposition 56. Observe that the equivalence between (1) and (5) is exactly statement (3) of Theorem 13.

Proposition 50. *Let R be a relation with a first-order definition in $(\mathbb{N}, =)$. Then the following are equivalent.*

- (1) R is preserved by an operation from $\text{Pol}(\neq) \setminus \mathcal{H}$, i.e., $\text{Pol}(R) \cap \text{Pol}(\neq) \not\subseteq \mathcal{H}$.
- (2) R is preserved by a non-injective operation that depends on all its arguments and preserves \neq .
- (3) R is preserved by an operation from $\text{Pol}(\neq)^{(2)}$ that violates N .
- (4) R is preserved by a 1-bar function.
- (5) R is preserved by \mathcal{B} , i.e., $\mathcal{B} \subseteq \text{Pol}(R)$.

Proof. Given an operation from $\text{Pol}(\neq) \setminus \mathcal{H}$, we obtain an operation as described in (2) by leaving away all fictitious arguments, so (1) implies (2). An operation that is non-injective and depends on all its arguments is not generated by an injective operation, and therefore Lemma 45 shows the implication from (1) to (2). Any binary operation that violates N is necessarily essential and non-injective, and hence the conditions of Lemma 49 are satisfied, which shows the implication from (2) to (3). The equivalence of (3) and (4) is Lemma 46, and the implication from (3) to (1) is immediate. The equivalence of (4) and (5) is by definition. \square

6.3. From the bar clone to the odd clone. In this subsection, we explore those clones in the monoidal interval $I_{\mathcal{J}}$ which contain the bar clone \mathcal{B} , but which are not contained in \mathcal{R} (the clone of operations which are injective in one direction, cf. Definition 11). It will be necessary to first give a relational description of \mathcal{B} .

6.3.1. The bar clone, relationally. In the following it will be convenient to work with first-order formulas whose atomic formulas are of the form $x = y$, *true*, or *false*. The *graph* of a quantifier-free formula Φ with such atomic formulas is the graph where the vertices are the variables of Φ , and where two vertices x, y are adjacent iff Φ contains the sub-formula $x = y$ or the sub-formula $x \neq y$. We recall standard terminology. If a formula Φ is in conjunctive normal form, the conjuncts in Φ are called *clauses*, and the disjuncts in a clause are called *literals*. Hence, literals

are formulas that are either atomic, in which case they are also called *positive*, or the negation of an atomic formula, in which case they are also called *negative*.

Definition 51. Let Φ be a quantifier-free first-order formula where all atomic formulas are of the form $x = y$ or *false*. Then Φ is called

- *Horn* iff Φ is in conjunctive normal form and each clause in Φ contains at most one positive literal;
- *connected Horn* iff Φ is Horn and the graph of a clause ψ in Φ has at most two connected components when ψ has no positive literals, and is connected when ψ has positive literals.
- *extended Horn* iff Φ is a conjunction of formulas of the form $(\phi_1 \wedge \dots \wedge \phi_l) \rightarrow (\psi_1 \wedge \dots \wedge \psi_k)$ for $l \geq 1, k \geq 1$ where $\phi_1, \dots, \phi_l, \psi_1, \dots, \psi_k$ are atomic formulas (and hence of the form $x = y$ or *false*).
- *connected extended Horn* iff Φ is extended Horn, and if a conjunct ψ of Φ is connected whenever the right-hand side of the implication in ψ contains a literal of the form $x = y$ and no literal *false*.

Clearly, every connected Horn formula can be written as a connected extended Horn formula. But there are connected extended Horn formulas that are not equivalent to any connected Horn formula; we will even see that there are connected extended Horn formulas that do not have a pp definition by connected Horn formulas (Lemmas 59 and 60).

The conjuncts in an extended Horn formula are also called *extended Horn clauses*, and the literals on the left-hand side (right-hand side) of the implication of an extended Horn clause are called the *negative literals* (*positive literals*, respectively) of the extended Horn clause. Note that extended Horn formulas can always be translated into (standard) Horn formulas: if $(\phi_1 \wedge \dots \wedge \phi_k) \rightarrow (\psi_1 \wedge \dots \wedge \psi_k)$ is an extended Horn clause, then we can replace this clause by the conjunction of Horn clauses $(\neg\phi_1 \vee \dots \vee \neg\phi_k \vee \psi_1) \wedge \dots \wedge (\neg\phi_1 \vee \dots \vee \neg\phi_k \vee \psi_k)$.

It will be convenient to say that a formula is *preserved* by an operation (or a set of operations) iff the relation that is defined by the formula is preserved by the operation (or by the set of operations, respectively). We want to give a syntactic description of the formulas that are preserved by \mathcal{B} . We show that every such formula is equivalent to a connected extended Horn formula. In fact, we show the stronger result that every *expanded Horn formula* preserved by \mathcal{B} is itself a connected extended Horn formula.

Definition 52. A extended Horn formula Φ is called *expanded Horn* iff it contains all connected extended Horn clauses on the same set of variables as Φ that are implied by Φ , and for every disconnected clause $\psi_1 \rightarrow \psi_2$ in Φ where ψ_2 does not contain *false* and for all variables x, y in Φ

- adding an atomic formula of the form $x = y$ to ψ_1 , or
- adding an atomic formula of the form $x = y$ to ψ_2 , or
- setting ψ_2 to *false*

results in a formula that is not equivalent to Φ .

Lemma 53. Every extended Horn formula is equivalent to an expanded Horn formula.

Proof. Let Φ be any given extended Horn formula. We construct the expanded Horn formula that is equivalent to Φ , as follows. We first add to Φ all connected extended Horn clauses that are implied by Φ .

Then, if Φ contains an extended Horn clause $\psi_1 \rightarrow \psi_2$ that is not connected and where ψ_2 does not contain *false*, and Φ is equivalent to a formula Φ' obtained from Φ by adding an atomic formula of the form $x = y$ to ψ_1 or to ψ_2 for variables x, y

in Φ , or by setting ψ_2 to *false*, then we replace Φ by Φ' . If Φ does not contain such a clause (which will always happen after a finite number of steps, because there is only a finite number of formulas of the form $x = y$ for variables x, y in Φ), then Φ is expanded Horn, and clearly equivalent to the formula we started with. \square

Lemma 54. *If Φ is an expanded Horn formula, and Φ defines a relation R that is preserved by \mathcal{B} , then Φ is connected extended Horn.*

Proof. Suppose for contradiction that $\Phi(x_1, \dots, x_n)$ contains an extended Horn clause ψ of the form $\psi_1 \rightarrow \psi_2$ whose graph G contains at least two connected components, and where ψ_2 does not contain *false* and contains at least one positive literal $x_i = x_j$. Let x_k be a variable from a connected component of G that does not contain x_i and x_j . Let C_1, \dots, C_p be the components of $G' := G - H$, where H is the graph of ψ_2 . Observe that x_i and x_j are in distinct components: Otherwise, if we replace the clause ψ by $(\psi_1 \wedge x_i = x_k) \rightarrow \psi_2$, we obtain an equivalent formula. Because ψ is disconnected, this is in contradiction to the assumption that Φ is expanded. Assume for the sake of notation that C_1 is the component of x_i , and C_2 the one of x_j .

We claim that there is a tuple s for which $\Phi(s)$ is true, for which $s_i \neq s_j$, and for which $s_l = s_m$ whenever x_l, x_m are both in C_1 or are both in C_2 . Otherwise, Φ implies $(\bigwedge_{(y,z) \in C_1 \cup C_2} y = z) \rightarrow x_i = x_j$, which is a connected expanded Horn formula and therefore contained in Φ . Then Φ must be equivalent to the formula Φ' obtained from Φ by replacing the disconnected clause ψ by $(\psi_1 \wedge x_i = x_j) \rightarrow \psi_2$, again in contradiction to the assumption that Φ is expanded. To show the equivalence, it clearly suffices to prove that Φ' implies Φ . So suppose that s satisfies Φ' and ψ_1 . The clause $(\bigwedge_{(y,z) \in C_1 \cup C_2} y = z) \rightarrow x_i = x_j$, which is contained in Φ and in Φ' , shows that $s_i = s_j$. But then the premise of the new extended Horn clause of Φ' is satisfied as well, and therefore s satisfies ψ_2 , which is what we had to show.

Next, we claim that there is a tuple t with $t_k \neq t_i$ satisfying Φ and where $t_l = t_m$ if x_l, x_m are from the same component C_q for some $1 \leq q \leq p$. Otherwise, Φ implies $\psi_1 \rightarrow x_k = x_i$. But then the formula Φ' obtained from Φ by replacing ψ by $\psi_1 \rightarrow (\psi_2 \wedge x_k = x_i)$ is clearly equivalent to Φ . This again contradicts the assumption that Φ is expanded.

Let f be a binary operation such that $x \mapsto f(x, v)$ is constant for all entries v of t except for $v = t_i$, and which is injective otherwise. Clearly, modulo permutations acting on its arguments, f is an n -bar operation, and hence $f \in \mathcal{B}$. Now, consider the tuple $r := f(s, t)$. Note that $t_i = t_j$, because t satisfies Φ and t satisfies the premise of the clause ψ that is contained in Φ ; since the conclusion of ψ contains $x_i = x_j$ we have that $t_i = t_j$. It is straightforward to verify that r does not satisfy Φ , because r satisfies ψ_1 , but $r_i \neq r_j$, and therefore ψ is violated. \square

The following relations play an important role in the relational description of \mathcal{B} and the clones above \mathcal{B} .

Definition 55. For $n \geq 2$, let R_n , \underline{R}_n , and R_n^\neq be the relations defined by

$$\begin{aligned} R_n(x_1, y_1, \dots, x_n, y_n) &\equiv \bigvee_{1 \leq i \leq n} x_i \neq y_i \\ \underline{R}_n(x_1, y_1, \dots, x_n, y_n) &\equiv R_n(x_1, y_1, \dots, x_n, y_n) \vee x_1 = y_1 = \dots = x_n = y_n \\ R_n^\neq(x_1, y_1, \dots, x_n, y_n) &\equiv R_n(x_1, y_1, \dots, x_n, y_n) \wedge \bigwedge_{1 \leq i \neq j \leq n} (x_i \neq y_j \wedge x_i \neq x_j \wedge y_i \neq y_j). \end{aligned}$$

Note that all three relations can be defined by connected extended Horn formulas. Observe also that the expressive power of the relations in each of these sequences increases with increasing n : For example, $\exists u, v. (u = v) \wedge R_{n+1}(x_1, y_1, \dots, x_n, y_n, u, v)$ is a pp definition of R_n from R_{n+1} .

Proposition 56. *Let R be a relation with a first-order definition in $(\mathbb{N}, =)$. Then the following are equivalent.*

- (1) R is preserved by \mathcal{B} , i.e., $\mathcal{B} \subseteq \text{Pol}(R)$.
- (2) Every expanded Horn formula that defines R is connected extended Horn.
- (3) R can be defined by a connected extended Horn formula.
- (4) There exists an n such that R has a pp-definition in $(\mathbb{N}, \underline{R}_n, \neq)$.

Proof. The implication from (1) to (2) is shown in Lemma 54. For the implication from (2) to (3), recall that \mathcal{B} contains \mathcal{H} and hence the relation R has a Horn definition. By Lemma 53, every Horn formula has an equivalent expanded Horn formula, and by assumption this formula is a connected extended Horn definition of R .

To show that (3) implies (4), let R be a relation with a connected extended Horn definition Φ . We first show that every extended Horn clause $\psi_1 \rightarrow \psi_2$ from Φ can be pp-defined in $(\mathbb{N}, \underline{R}_n)$, for a sufficiently large n , if ψ_2 is non-empty and does not contain *false*.

Let G be the graph of ψ_1 where we add isolated vertices for each variable that appears in ψ_2 but not in ψ_1 ; in other words, let G be the graph of $\psi_1 \rightarrow \psi_2$ where we remove an edge between x and y if $x = y$ is a literal of ψ_2 and not of ψ_1 . Let C_1, \dots, C_c be the components of the graph G . We claim that $\psi_1 \rightarrow \psi_2$ can be pp-defined by \underline{R}_n where $n = |C_1| + \dots + |C_c|$. A pp definition is obtained from the formula

$$\underline{R}_n(x_1, y_1, \dots, x_n, y_n) \wedge \bigwedge_{(1 \leq i < c)} \bigwedge_{(1 < j \leq |C_{i+1}|)} y_{s(i)+1} = x_{s(i)+1} = x_{s(i)+j}$$

by existentially quantifying x_1, \dots, x_n , where $s(i) := |C_1| + \dots + |C_i|$. In this pp definition, the variables $y_{s(i)+1}, \dots, y_{s(i)+|C_{i+1}|}$ are the variables in the component C_{i+1} from $\psi_1 \rightarrow \psi_2$.

To see that this is correct, consider a satisfying assignment to the variables of $\psi_1 \rightarrow \psi_2$. Set for all i the variables $x_{s(i)+1}, \dots, x_{s(i)+|C_{i+1}|}$ to the same value as $y_{s(i)+1}$. Suppose first that there exist two variables y_k, y_l from the same component C_i which are mapped to different values. Then $x_k = x_l$ cannot be equal to both y_k and y_l , and hence either $x_k \neq y_k$ or $x_l \neq y_l$. Hence, the conjunct $\underline{R}_n(x_1, y_1, \dots, x_n, y_n)$ is satisfied, and the assignment satisfies the given pp formula. Suppose now that otherwise all variables which have the same component are mapped to the same value. Since $\psi_1 \rightarrow \psi_2$ is connected Horn, this implies by transitivity of equality that all variables that appear in $\psi_1 \rightarrow \psi_2$ must have the same value. But this assignment clearly satisfies the given pp formula as well, so we are done with one direction of the correctness proof.

For the other direction, consider a satisfying assignment for the pp formula. If all variables are set to the same value, then $\psi_1 \rightarrow \psi_2$ will be satisfied by this assignment as well, because ψ_2 does not contain *false*. Otherwise, $x_k \neq y_k$ for some $1 \leq k \leq n$ because of the conjunct $\underline{R}_n(x_1, y_1, \dots, x_n, y_n)$. Let C_{i+1} be the component of y_k in G . Then $y_{s(i)+1} = x_{s(i)+1} = x_k \neq y_k$. But since C_{i+1} is a connected component, this implies that one of the equalities in ψ_1 must be falsified, showing that $\psi_1 \rightarrow \psi_2$ is satisfied by the assignment.

Now, suppose that Φ contains a clause $\psi_1 \rightarrow \psi_2$ where ψ_2 contains *false*; i.e., the clause is logically equivalent to a disjunction of literals of the form $x \neq y$. We proceed just as in the case above, but use the relation $\underline{R}_{n+1}(x_1, y_1, \dots, x_{n+1}, y_{n+1})$

instead of $\underline{R}_n(x_1, y_1, \dots, x_n, y_n)$ with two new existentially quantified variables x_{n+1} and y_{n+1} at the two additional arguments, and we also add the conjunct $x_{n+1} \neq y_{n+1}$ to the formula. We leave the verification that the resulting formula is equivalent to $\psi_1 \rightarrow \psi_2$ to the reader. As we have seen, any clause from a connected Horn formula can be pp-defined in $(\mathbb{N}, \underline{R}_{n+1}, \neq)$, and thus (3) implies (4).

Finally, (4) implies (1): It is easy to see that for all $n \geq 2$ the relations \underline{R}_n and the relation \neq are preserved by 1-Bar operations. \square

6.3.2. Above the bar clone. In this section we show that the interval between \mathcal{B} and $\text{Pol}(\neq)$ contains an infinite strictly decreasing sequence of clones whose intersection is \mathcal{B} .

Definition 57. For all $k \geq 2$, fix an operation $f_k \in \mathcal{O}^{(k)}$ satisfying the following:

$$\begin{aligned} f_k(x, 1, 1, \dots, 1, 1) &= 1 \\ f_k(2, x, 2, \dots, 2, 2) &= 2 \\ &\dots \\ f_k(k, k, \dots, k, x) &= k \end{aligned}$$

Moreover, for all other arguments, the function arbitrarily takes a value that is distinct from all other function values.

Observe that f_3 appeared already in Definition 12. We use the f_k to define another sequence of operations.

Definition 58. For all $k \geq 3$, fix an operation $g_k \in \mathcal{O}^{(k+1)}$ which satisfies

$$g_k(0, x_1, \dots, x_k) = f_k(x_1, \dots, x_k)$$

and which takes distinct values for all other arguments.

Observe that each of the operations g_k (and similarly f_k) generates \mathcal{B} : Obviously, g_k depends on all arguments, is non-injective, and preserves \neq . By Proposition 50, all relations that are preserved by g_k are also preserved by \mathcal{B} , and hence g_k generates \mathcal{B} .

We now show that the operations g_k , for increasing k , generate smaller and smaller clones.

Lemma 59. For $k \geq 3$, the operations g_k and f_k preserve \underline{R}_{k-1} .

Proof. We show the lemma for g_k ; the proof for f_k is similar and a bit easier, and left to the reader. Let t^0, \dots, t^k be $(2k-2)$ -tuples that all satisfy \underline{R}_{k-1} . Let $t := g_k(t^0, \dots, t^k)$, and suppose that $t_1 = t_2, \dots, t_{2k-3} = t_{2k-2}$. By the definition of g_k , this implies that $t^0 = (0, \dots, 0)$, and that for all $1 \leq i \leq k-1$ the tuples $(t_{2i-1}^1, \dots, t_{2i-1}^k)$ and $(t_{2i}^1, \dots, t_{2i}^k)$ have equal entries except for at most one position, call it $p(i)$. Since the function p takes at most $k-1$ values, there must be a $j \in \{1, \dots, k\}$ which is not in the range of p . For this j we thus have $t_1^j = t_2^j, \dots, t_{2k-3}^j = t_{2k-2}^j$. Because $t^j \in \underline{R}_{k-1}$, this implies that $t_1^j = \dots = t_{2k-2}^j$. Hence, the tuple $t = g_k(t^0, \dots, t^k)$ is constant as well; therefore, g_k preserves \underline{R}_{k-1} . \square

Lemma 60. For $k \geq 3$, the operation f_k preserves neither \underline{R}_k nor R_k . The operation g_k does not preserve \underline{R}_k .

Proof. If we apply f_k to the $2k$ -tuples

$$\begin{aligned} t^1 &= (0, 1, 2, 2, 3, 3, \dots, k, k) \\ t^2 &= (1, 1, 0, 2, 3, 3, \dots, k, k) \\ &\dots \\ t^k &= (1, 1, 2, 2, 3, 3, \dots, 0, k) \end{aligned}$$

that all satisfy \underline{R}_k (and all satisfy R_k), then we obtain $(1, 1, 2, 2, 3, 3, \dots, k, k)$, which does not satisfy \underline{R}_k (and does not satisfy R_k). For g_k , we define $t^0 = (0, \dots, 0)$, and apply g_k to (t^0, t^1, \dots, t^k) . \square

We have thus seen that $\text{Pol}(\underline{R}_3, \neq), \text{Pol}(\underline{R}_4, \neq), \dots$ is a strictly decreasing chain of clones (it is decreasing by the remark after the definition of these relations, and strictly so by the preceding lemma). By Proposition 56 its intersection equals \mathcal{B} .

6.3.3. The odd clone. A clone of special interest (see Theorems 13 and 15) is the clone \mathcal{S} generated by f_3 , which will be called the *odd clone* (cf. Definition 12). We will now describe the relations invariant under \mathcal{S} by providing a (finite) generating system of these relations, as well as a syntactic characterization of the formulas defining such relations.

Definition 61. Let ODD_3 be the ternary relation

$$\{(a, b, c) \in \mathbb{N}^3 \mid a = b = c \vee |\{a, b, c\}| = 3\}.$$

Proposition 62. Let R be a relation with a first-order definition in $(\mathbb{N}, =)$. Then the following are equivalent.

- (1) R can be defined by a connected Horn formula.
- (2) R has a pp definition by \underline{R}_2 and \neq .
- (3) R has a pp definition by ODD_3 and \neq .
- (4) $R \in \text{Inv}(\mathcal{S})$, i.e., R is preserved by f_3 .
- (5) $R \in \text{Inv}(\mathcal{B})$, but there is no definition of R_3^\neq in $(\mathbb{N}, R, \underline{R}_2, \neq)$.

Before proving the proposition, we need two proof-theoretical lemmas observing restrictions on what can be derived from connected extended Horn formulas.

Lemma 63. Suppose Φ is a connected extended Horn formula that does not imply $a = b$, but $\Phi \wedge x = y$ implies $a = b$. Then $\Phi \wedge x = y$ also implies $x = y = a = b$.

Proof. We rewrite Φ as a conjunction Ψ of (standard) Horn clauses as described above. When for an extended clause $\phi \rightarrow (\psi_1 \wedge \dots \wedge \psi_k)$ for some $i \leq k$ we have $\psi_i = \text{false}$, then we first remove all positive literals except ψ_i from the extended clause, which leads to an equivalent formula.

It is well-known that the restricted form of resolution for predicate logic that requires that at least one of the parent clauses in a resolution step has to be positive, also known as *P-resolution*, is complete (see e.g. [Sch89], p. 102). When we apply resolution on a formula, we assume that there are always clauses $u \neq v \vee v \neq w \vee u = w$, $u \neq v \vee v = u$, and $u = u$ for universally quantified variables u, v, w (that axiomatize that the symbol $=$ for equality is an equivalence relation), and we also assume that all other variables in the formula are in fact constant symbols (i.e., they are existential variables before skolemization).

Since *P-resolution* is complete, we have in particular that $\Psi \wedge x = y$ implies $a = b$ if and only if $a = b$ can be derived by *P-resolution* from $\Psi \wedge x = y$. We show that $\Psi \wedge x = y$ implies $x = y = a = b$ for all formulas $a = b$ that are not implied by Ψ but implied by $\Psi \wedge x = y$. Suppose for contradiction that there is a literal $a = b$ that can be derived by *P-resolution* from $\Psi \wedge x = y$ but cannot be derived from Ψ , and where $\Psi \wedge x = y$ does not imply $x = y = a = b$. Select a literal $a = b$ with these properties whose *P-derivation* from $\Psi \wedge x = y$ is shortest possible.

The literal $a = b$ has either been derived from a clause $a_1 \neq b_1 \vee \dots \vee a_k \neq b_k \vee a = b$ in Ψ by resolving with the literals $a_1 = b_1, \dots, a_k = b_k$, or by resolving with one of the three clauses having universally quantified variables (after instantiating the universal variables). In the second case, suppose that the universal clause was $u \neq v \vee v \neq w \vee u = v$, and $a = b$ has been derived by resolving with $a = c$ and $c = b$, for some c (for the other two universal clauses, the argument is analogous). Then

at least one of the literals $a = c$ and $c = b$ cannot be derived from Ψ alone, because otherwise $a = b$ can be derived from Ψ alone, in contradiction to our assumptions. Suppose the literal that cannot be derived from Ψ alone is $a = c$ (the other case is analogous). By the choice of $a = b$, the formula $\Psi \wedge x = y$ implies $x = y = a = c$. At the same time, we have that $\Psi \wedge x = y$ implies $a = b$, and hence Ψ implies $x = y = a = b$, a contradiction.

In the first case, at least one of the literals in the clause $a_1 \neq b_1 \vee \dots \vee a_k \neq b_k \vee a = b$, say $a_i = b_i$, cannot be derived from Ψ alone, because otherwise $a = b$ can be derived from Ψ alone, in contradiction to our assumptions. By the choice of $a = b$, $\Psi \wedge x = y$ implies $x = y = a_i = b_i$. Let ϕ be the extended Horn clause in Φ that corresponds to the Horn clause ψ in Ψ (i.e., ψ is among the clauses that have replaced the clause ϕ when creating Ψ from Φ).

Since ϕ is a connected extended Horn clause, there is a sequence of variables x_1, x_2, \dots, x_s of ϕ with $x_1 = a$ and $x_s = a_i$ such that $x_i = x_{i+1}$ for $1 \leq i < s$ is either a sub-formula of the right-hand side or of the left-hand side of the implication ϕ . By construction of Ψ , for each literal $x_i = x_{i+1}$ on the right-hand side there is the clause $a_1 \neq b_1 \vee \dots \vee a_k \neq b_k \vee x_i = x_{i+1}$ in Ψ . But then it is clear that $\Psi \wedge a_1 = b_1 \wedge \dots \wedge a_k = b_k$ implies that $x_1 = \dots = x_s$, and hence also that $a_i = b_i = a = b$. Therefore, $\Psi \wedge x = y$ implies $x = y = a = b$, a contradiction. \square

The following is a straightforward consequence.

Lemma 64. *Let Φ be a connected extended Horn formula and let ψ be a conjunction of equalities $x_1 = y_1 \wedge \dots \wedge x_n = y_n$. If $\Phi \wedge \psi$ implies $x_0 = y_0$, but $\Phi \wedge \psi'$ does not imply $x_0 = y_0$ for any subconjunction ψ' of ψ , then $\Phi \wedge \psi$ also implies $x_0 = y_0 = x_1 = y_1 = \dots = x_n = y_n$.*

Proof. We apply Lemma 63 to $\Phi' := \Phi \wedge x_2 = y_2 \wedge \dots \wedge x_n = y_n$, $x_1 = y_1$, and $x_0 = y_0$ to derive that $\Phi' \wedge x_1 = y_1$ implies $x_0 = y_0 = x_1 = y_1$. We can argue analogously for $x_i = y_i$ instead of $x_1 = y_1$ for all $1 \leq i \leq n$, and the statement follows. \square

Proof of Proposition 62. To show that (1) implies (2), suppose that Φ is a connected Horn definition of R . We show that each clause ψ of Φ has a pp definition in $(\mathbb{N}, \underline{R}_2, \neq)$.

First, consider the case that ψ contains a positive literal $x_0 = y_0$. The graph of ψ without the edge between x_0 and y_0 has at most two connected components. In the case that the graph is still connected even after removing the edge between x_0 and y_0 , the clause ψ is a tautology and can be ignored. Otherwise, let x_0, \dots, x_l be the variables of the connected component of x_0 , and let y_0, \dots, y_k be the variables of the connected component of y_0 .

Consider the formula ψ' defined by

$$\begin{aligned} \exists u_1, \dots, u_l, v_1, \dots, v_k. \underline{R}_2(x_0, x_1, u_1, u_1) \wedge \underline{R}_2(u_1, x_2, u_2, u_2) \wedge \dots \wedge \underline{R}_2(u_{l-1}, x_l, u_l, u_l) \\ \wedge \underline{R}_2(y_0, y_1, v_1, v_1) \wedge \underline{R}_2(v_1, y_2, v_2, v_2) \wedge \dots \wedge \underline{R}_2(v_{k-1}, y_k, v_k, v_k) \wedge u_l = v_k. \end{aligned}$$

We show that a tuple satisfies ψ if and only if it satisfies ψ' . If the tuple satisfies no negative literal of ψ , then the only situation where ψ' is true is $x_0 = u_1 = \dots = u_l = v_k = \dots = v_1 = y_0$; hence in this case, ψ' holds if and only if the tuple satisfies $x_0 = y_0$, which in turn is the case if and only if ψ holds. If on the other hand the tuple satisfies some negative literal of ψ , then either there is $i < l$ such that $x_i \neq x_{i+1}$ or $i < k$ such that $y_i \neq y_{i+1}$; assume wlog the first is the case. Then we set (u_1, \dots, u_i) to (x_1, \dots, x_i) and (v_1, \dots, v_k) to (y_1, \dots, y_k) , u_l to v_k , and set u_{i+1}, \dots, u_{l-1} to arbitrary pairwise disjoint values that are also distinct from all the values taken by any other variable of the clause. It is easy to check that this

assignment of variables is a satisfying assignment for the existentially quantified variables in the formula ψ' . Hence, ψ' is a pp definition of ψ .

If ψ does not contain a positive literal, then the graph of ψ consists of at most two connected components, because ϕ is connected Horn. Assume first it does have two components. Let x_0, \dots, x_l be the variables in the first and y_0, \dots, y_k be the variables in the second connected component. As above, it is easy to verify that the following formula is equivalent to ψ .

$$\begin{aligned} \exists u_1, \dots, u_l, v_1, \dots, v_k, z. & \underline{R}_2(x_0, x_1, u_1, u_1) \wedge \underline{R}_2(u_1, x_2, u_2, u_2) \wedge \dots \wedge \underline{R}_2(u_{l-1}, x_l, u_l, u_l) \\ & \wedge \underline{R}_2(y_0, y_1, v_1, v_1) \wedge \underline{R}_2(v_1, y_2, v_2, v_2) \wedge \dots \wedge \underline{R}_2(v_{k-1}, y_k, v_k, v_k) \\ & \wedge \underline{R}_2(x_l, u_l, z, z) \wedge \underline{R}_2(y_k, v_k, z, z) \wedge u_l \neq v_k. \end{aligned}$$

If the graph of ψ has only one component, then the following formula is equivalent to ψ :

$$\exists u_1, \dots, u_l \underline{R}_2(x_0, x_1, u_1, u_1) \wedge \underline{R}_2(u_1, x_2, u_2, u_2) \wedge \dots \wedge \underline{R}_2(u_{l-1}, x_l, u_l, u_l) \wedge u_1 \neq u_l.$$

For the implication from (2) to (3), note that $\underline{R}_2(x_1, x_2, x_3, x_4)$ can be pp defined by

$$\begin{aligned} \exists y_1, y_2, y_3, y_4, y_5. & \text{ODD}_3(x_1, x_2, y_1) \wedge \text{ODD}_3(x_1, x_2, y_2) \\ & \wedge \text{ODD}_3(y_1, y_2, y_3) \wedge \text{ODD}_3(y_3, y_4, y_5) \\ & \wedge \text{ODD}_3(x_3, x_4, y_4) \wedge \text{ODD}_3(x_3, x_4, y_5) \end{aligned}$$

We check that this formula is indeed equivalent to $\underline{R}_2(x_1, x_2, x_3, x_4)$: Let t be a 4-tuple that satisfies $\underline{R}_2(x_1, x_2, x_3, x_4)$; we assign values to the y_i such that the formula above is satisfied. If $t_1 = t_2$ then set y_1, y_2, y_3 to t_1 . If $t_3 = t_4$ then set y_3, y_4, y_5 to t_3 . Note that if both $t_1 = t_2$ and $t_3 = t_4$ then our choice for y_3 is well-defined, since $\underline{R}_2(x_1, x_2, x_3, x_4)$ then implies that $t_1 = t_2 = t_3 = t_4$. If $t_1 \neq t_2$ then set y_1, y_2 to distinct values that are also distinct from all entries of t . Similarly, if $t_3 \neq t_4$, then set y_4, y_5 to distinct values that are also distinct from all entries of t . If both $t_1 \neq t_2$ and $t_3 \neq t_4$, then set y_3 to any value that is distinct from the values for y_1, y_2, y_4, y_5 . It is straightforward to verify that these assignments satisfy all conjunct in the given formula. Conversely, if a 4-tuple t satisfies the formula, and $t_1 = t_2$ and $t_3 = t_4$, then $t_1 = t_2 = t_3 = t_4$, and hence t satisfies $\underline{R}_2(x_1, x_2, x_3, x_4)$.

For the implication from (3) to (4) it suffices to verify that f_3 preserves ODD_3 . Every 3-tuple t from ODD_3 either satisfies $t_1 = t_2 = t_3$ or $t_1 \neq t_2 \neq t_3 \neq t_1$. To violate ODD_3 , we would have to find three such tuples a, b, c such that $d := f_3(a, b, c)$ has two equal entries that are distinct from a third entry, say wlog. $d_1 = d_2 \neq d_3$. The only two ways to have $d_1 = d_2$ are when $d_1 = d_2$ is from $\{0, 1, 2\}$, or when this is not the case and the tuples (a_1, b_1, c_1) and (a_2, b_2, c_2) are equal. Consider the first case, and assume wlog. $d_1 = d_2 = 0$. Then we know that $b_1 = c_1 = 0$ and $b_2 = c_2 = 0$. But then, since b and c are in ODD_3 , we have $b_3 = c_3 = 0$ as well, and hence $d_1 = d_2 = d_3 = 0$, a contradiction. The second case is similar and left to the reader. Hence, f_3 preserves ODD_3 .

For the implication from (4) to (5), containment in $\text{Inv}(\mathcal{B})$ follows from the fact that f_3 generates all operations in \mathcal{B} (see also the remark after Definition 58). By Lemma 59, f_3 preserves \underline{R}_2 and f_3 also clearly preserves \neq , so it suffices to show that f_3 violates R_3^\neq . If we apply f_3 to the 6-tuples $t_1 = (0, 1, 2, 2, 3, 3)$, $t_2 = (0, 0, 1, 2, 3, 3)$, $t_3 = (0, 0, 2, 2, 1, 3)$, which all satisfy R_3^\neq , we obtain $(0, 0, 2, 2, 3, 3)$, which does not satisfy R_3^\neq .

Finally, we show the contraposition of the implication from (5) to (1): if R cannot be defined by a connected Horn formula, but is preserved by \mathcal{B} , then there is a pp-definition of R_n^\neq in $(\mathbb{N}, R, \underline{R}_2, \neq)$ for some $n \geq 3$. Because R_3^\neq clearly has a pp-definition by R_n^\neq for all $n \geq 3$ this will show the statement.

Let $\Phi(v_1, \dots, v_s)$ be an expanded Horn formula that defines R . By Proposition 56, we know that Φ is a connected extended Horn formula. We first claim that if for all extended Horn clauses $\psi_1 \rightarrow \psi_2$ in Φ the graph of ψ_1 has at most two connected components, then Φ has a connected Horn definition, contradicting our assumption on Φ . To see this, we show that every such clause $\psi_1 \rightarrow \psi_2$ has a connected Horn definition. If ψ_2 contains *false* or is empty then there is nothing to show; assume thus that this is not the case. If the graph of ψ_1 has exactly two components C_1 and C_2 , we know that ψ_2 contains a conjunct $x = y$ where $x \in C_1$ and $y \in C_2$, because Φ is connected extended Horn. Then $\psi_1 \rightarrow \psi_2$ is equivalent to the conjunction of the connected Horn clauses $\psi_1 \rightarrow x = y$ and $(\psi_1 \wedge x = y) \rightarrow x' = y'$ for every other conjunct $x' = y'$ of ψ_2 . The case that the graph of ψ_1 is connected is easier and left to the reader.

Therefore, Φ must contain a clause $\psi_1 \rightarrow \psi_2$ such that ψ_1 consists of $n \geq 3$ connected components C_1, \dots, C_n . We show that $R_n^\neq(x_1, y_1, \dots, x_n, y_n)$ has the following pp-definition by R , \underline{R}_2 and \neq .

$$\begin{aligned} \exists v_1, \dots, v_s. R(v_1, \dots, v_s) \wedge \bigwedge_{1 \leq i \leq n} \bigwedge_{v_k \in C_i} \underline{R}_2(x_i, y_i, v_k, v_k) \\ \wedge \bigwedge_{1 \leq i \neq j \leq n} (x_i \neq y_j \wedge x_i \neq x_j \wedge y_i \neq y_j) \end{aligned} \quad (*)$$

To prove the equivalence, assume that $t \notin R_n^\neq$. We want to show that t does not satisfy $(*)$. If t violates the second line of the formula $(*)$, we are done. Otherwise, it must hold that $x_1 = y_1, \dots, x_n = y_n$. Suppose for contradiction that there are elements v_1, \dots, v_s as specified in $(*)$. Then for all $1 \leq i \leq n$ the sub-formula $\bigwedge_{v_k \in C_i} \underline{R}_2(x_i, y_i, v_k, v_k)$ implies that all variables v_k from C_i must have the same value as $x_i = y_i$. Then the clause $\psi_1 \rightarrow \psi_2$ in Φ implies that $x_1 = y_1 = \dots = x_n = y_n$ (because $\psi_1 \rightarrow \psi_2$ is connected), which contradicts the second line of formula $(*)$.

Now assume that $t \in R_n^\neq$. We want to show that t satisfies $(*)$. If t satisfies $x_i \neq y_i$ for all $1 \leq i \leq n$, then any tuple from R gives an assignment to the existentially quantified variables in $(*)$ that shows that t satisfies $(*)$. For the sake of notation, let us thus assume that there is a k with $1 \leq k < n$ such that t satisfies $x_i = y_i$ for $1 \leq i \leq k$, and $x_i \neq y_i$ for all $k < i \leq n$. We claim that there exists a tuple $r \in R$ such that for all $i, j \leq k$ and all $v_l \in C_i, v_m \in C_j$ it holds that $r_l = r_m$ if and only if $i = j$. Otherwise, the formula

$$\Phi(v_1, \dots, v_s) \wedge \bigwedge_{i \leq k} \bigwedge_{v_l, v_m \in C_i} v_l = v_m \wedge \bigwedge_{v_l \in C_i, v_m \in C_j, 1 \leq i \neq j \leq k} v_l \neq v_m$$

is unsatisfiable. Since P-resolution (also see the proof of Lemma 63) is complete, there is a P-resolution proof of unsatisfiability. This proof uses at most once a clause of the form $\{v_p \neq v_q\}$ from this formula, since it resolves only with positive literals and since therefore every resolution step using $\{v_p \neq v_q\}$ must be the last one, deriving *false*. Therefore, $\Phi(v_1, \dots, v_s) \wedge \bigwedge_{i \leq k} \bigwedge_{v_l, v_m \in C_i} v_l = v_m$ implies $v_p = v_q$. Let ϕ be the smallest sub-formula of $\bigwedge_{i \leq k} \bigwedge_{v_l, v_m \in C_i} v_l = v_m$ such that $\Phi \wedge \phi$ implies $v_p = v_q$. If ϕ is empty, then the clause $v_p = v_q$ is contained in Φ because Φ is expanded; but then, adding $v_p \neq v_q$ to ψ in the clause $\psi_1 \rightarrow \psi_2$ results in an equivalent formula, a contradiction to Φ being expanded. If ϕ is not empty, we can apply Lemma 64 and obtain that $\Phi \wedge \phi$ implies that all variables u_1, \dots, u_l that appear in ϕ have the same value as the value of $v_p = v_q$. Because Φ is expanded, it has to contain the connected extended Horn clause $\phi \rightarrow u_1 = \dots = u_l = v_p = v_q$;

again, this shows that adding $v_p \neq v_q$ to ψ_1 in the clause $\psi_1 \rightarrow \psi_2$ results in an equivalent formula, a contradiction.

Hence, we have shown the existence of a tuple $r \in R$ with the properties as claimed above. The tuple r gives witnesses for the existentially quantified variables in $(*)$ that show that t satisfies $(*)$. \square

6.4. Richard. In this section we show that \mathcal{R} and \mathcal{S} are incomparable clones, and that every clone in $I_{\mathcal{S}}$ is either contained in \mathcal{S} or contains \mathcal{R} . We start with the fundamental observation that the set of operations \mathcal{R} (see Definition 11) is indeed a clone.

Lemma 65. *\mathcal{R} is a clone.*

Proof. Clearly, \mathcal{R} contains the projections. Let $f \in \mathcal{R}^{(n)}$ and $g_1, \dots, g_n \in \mathcal{R}^{(m)}$. There exist $1 \leq i \leq n$ and $1 \leq j \leq m$ such that f is injective in the i -th direction and g_i is injective in the j -th direction. It is readily verified that $f(g_1, \dots, g_n)$ is injective in the j -th direction. \square

Definition 66. Let ϕ be a quantifier-free first-order formula where all atomic subformulas are of the form $x = y$. Then ϕ is called *negative* if ϕ is in conjunctive normal form and each clause in ϕ either consists of a single positive literal, or does not contain positive literals.

The following theorem has been stated in equivalent form in [BC07, Lemma 8.6]. For the convenience of the reader, and for completeness, we also present a simplified proof here.

Theorem 67. *Let R be a relation having a reduced Horn definition that is not negative. Then ODD_3 has a pp definition from R and \neq .*

Proof. Let $\Phi(x_1, \dots, x_n)$ be a reduced Horn definition of R which is not negative; assume moreover that among such definitions of R , Φ is one with the minimal possible number of occurrences of variables. Since R is not negative, there must be a clause ψ in Φ of the form $x_{i_1} = x_{j_1} \vee x_{i_2} \neq x_{j_2} \vee \dots \vee x_{i_k} \neq x_{j_k}$ for $k \geq 2$. Since Φ is reduced, it contains a tuple $a \in \mathbb{N}^n$ with $a_{i_l} = a_{j_l}$ for all $1 \leq l \leq k$: Otherwise, Φ would imply $x_{i_1} \neq x_{j_1} \vee x_{i_2} \neq x_{j_2} \vee \dots \vee x_{i_k} \neq x_{j_k}$, and we would obtain an equivalent formula by removing the literal $x_{i_1} = x_{j_1}$ from ψ . Moreover, R contains a tuple $b \in \mathbb{N}^n$ with $b_{i_1} \neq b_{j_1}$, $b_{i_2} \neq b_{j_2}$, and $b_{i_l} = b_{j_l}$ for all $3 \leq l \leq k$: Otherwise, we could remove the literal $x_{i_2} \neq x_{j_2}$ from ψ .

Now consider the formula

$$R(x_1, \dots, x_n) \wedge \bigwedge_{l \geq 3}^k x_{i_l} = x_{j_l} \quad (*)$$

Note that because R is reduced, we have $|\{i_1, j_1, i_2, j_2\}| \geq 3$.

We first look at the case where $|\{i_1, j_1, i_2, j_2\}| = 3$. By appropriately renaming the variables of Φ , we may assume that $i_1 = i_2 = 1$, and that $j_1 = 2$ and $j_2 = 3$. Let $R'(x_i, x_{j_1}, x_{j_2})$ be the ternary relation defined by existentially quantifying all variables in $(*)$ except for x_i, x_{j_1}, x_{j_2} . Because $a \in R$ satisfies $a_{i_l} = a_{j_l}$ for all $3 \leq l \leq k$, it satisfies $(*)$. Since $a_{j_1} = a_{i_1} = a_{i_2} = a_{j_2}$, the relation R' contains a triple with three equal entries. Moreover, because $b \in R$ satisfies $(*)$ too, the relation R' contains a tuple (b_1, b_2, b_3) with $b_1 \neq b_2$ and $b_1 \neq b_3$.

Assume first that R' contains a triple with three distinct entries. We know that R' does not contain any tuple of the form (x, x, y) for $x \neq y$, because every n -tuple a with $a_{i_1} = a_{j_1}$ and $a_{i_2} \neq a_{j_2}$ violates $\psi \wedge \bigwedge_{l=3}^k x_{i_l} = x_{j_l}$. Hence, the relation defined by $R'(x, y, z) \wedge R'(y, z, x) \wedge R'(z, x, y)$ only contains the tuples with three

distinct entries and the tuples with three equal entries, and we have indeed found a pp-definition of the relation ODD_3 in (\mathbb{N}, R) .

Striving for a contradiction, we suppose now that $R'(x_i, x_{j_1}, x_{j_2})$ does not contain a triple with three distinct entries. Then it is impossible that there is a tuple $r \in R'$ where $r_2 \neq r_3$: This is because having a Horn definition, R is preserved by a binary injection, by Proposition 43, and so is R' . But the tuple obtained by applying a binary injective operation to r and (b_1, b_2, b_3) has three distinct entries, which contradicts our assumption for this case. Thus the clause $x_{i_3} \neq x_{j_3} \vee \dots \vee x_{i_k} \neq x_{j_k} \vee x_{j_1} = x_{j_2}$ is entailed by Φ . Hence, we could replace the clause ψ in Φ by this shorter clause and obtain a formula that is equivalent to Φ , because conversely, $x_{j_1} = x_{j_2}$ implies $x_{i_1} = x_{j_1} \vee x_{i_1} \neq x_{j_2}$ (recall that $i_1 = i_2$). This contradicts the choice of Φ .

Now take the case where $|\{i_1, j_1, i_2, j_2\}| = 4$; say wlog $i_1 = 1, j_1 = 2, i_2 = 3, j_2 = 4$. Let $R'(x_{i_1}, x_{j_1}, x_{i_2}, x_{j_2})$ be the relation defined by $(*)$ by existentially quantifying all variables except for $x_{i_1}, x_{j_1}, x_{i_2}, x_{j_2}$. As before we observe that $a' := (a_{i_1}, a_{j_1}, a_{i_2}, a_{j_2}) \in R'$ and $b' := (b_{i_1}, b_{j_1}, b_{i_2}, b_{j_2}) \in R'$. If $a_{i_1} = a_{j_1} \neq a_{i_2} = a_{j_2}$, then the tuple obtained by applying a binary injective operation to a' and b' has four distinct entries; this tuple is an element of R , by Proposition 43. It is easy to verify that then the formula

$$R'(x_1, x_2, x_3, x_4) \wedge R'(x_3, x_4, x_1, x_2) \wedge x_1 \neq x_3 \wedge x_1 \neq x_4 \wedge x_2 \neq x_3 \wedge x_2 \neq x_4 \quad (**)$$

is a pp definition of the relation $T(x_1, x_2, x_3, x_4)$ given by

$$T := \{(x_1, x_2, x_3, x_4) \mid x_1 = x_2 \neq x_3 = x_4 \vee |\{x_1, x_2, x_3, x_4\}| = 4\}.$$

The relation $\text{ODD}_3(x_1, x_2, x_3)$ has the following pp-definition.

$$\exists y \exists z. T(x_1, x_2, y, z) \wedge T(x_2, x_3, y, z) \wedge T(x_1, x_3, y, z).$$

Now suppose that $a_{j_1} = a_{i_2}$ and hence a' is the tuple with four equal entries. Assume for a moment that there exist $p \in \{1, 2\}$ and $q \in \{3, 4\}$ such that $R'(x_1, \dots, x_4)$ implies $x_p = x_q$. Then the relation obtained from $R'(x_1, \dots, x_4)$ by existentially quantifying x_q is such that we can proceed as in the case where $|\{i_1, j_1, i_2, j_2\}| = 3$. Thus, we may henceforth assume that $x_p = x_q$ is not entailed for any $p \in \{1, 2\}$ and $q \in \{3, 4\}$. For every such pair (p, q) , fix a tuple in R' witnessing that $x_p \neq x_q$ can be satisfied. Now applying a 5-ary injection to these four witnessing tuples and b' , we obtain a tuple all of whose entries are distinct. This tuple is in R' as R' is preserved by \mathcal{H} , by Proposition 43.

Consider the formula $R''(x_1, x_2, x_3, x_4)$ defined by the following formula.

$$\bigwedge_{\pi \in S_{\{1, 2, 3, 4\}}} R'(x_{\pi(1)}, x_{\pi(2)}, x_{\pi(3)}, x_{\pi(4)}),$$

where $S_{\{1, 2, 3, 4\}}$ denotes the symmetric group on $\{1, 2, 3, 4\}$. We claim that R'' defines ODD_3 . Note that tuples of the form (x, x, x, y) and (x, x, y, z) , for distinct x, y, z , are not contained in R' since every n -tuple d with $d_{i_1} = d_{j_1}$ and $d_{i_2} \neq d_{j_2}$ violates $\psi \wedge \bigwedge_{l=3}^k x_{i_l} = x_{j_l}$. Therefore, the relation R'' does not contain tuples with two values where one value appears three times, or tuples with three values where one value appears twice. But certainly R'' does contain all tuples with four equal values, and all tuples with four pairwise distinct values. If some tuple where two values occur twice is excluded from R' , then R'' does not contain any tuple with two values. Hence in this case, $\exists u. R''(x_1, x_2, x_3, u)$ is a pp definition of $\text{ODD}_3(x_1, x_2, x_3)$, and we are done. Otherwise, R'' contains all tuples where two values occur twice. Similarly as before it is easy to verify that then the expression $(**)$ above where R' has been replaced by R'' is a pp definition of $T(x_1, x_2, x_3, x_4)$. We have already

seen that ODD_3 has a pp definition from T , so it has a pp definition in (\mathbb{N}, R, \neq) as well. \square

The following proposition describes the clone \mathcal{R} : Item (1) shows that \mathcal{R} is finitely generated (namely, by any binary operation which violates ODD_3 but is injective in one direction); and (2) and (3) provide a syntactical description of the formulas defining relations in $\text{Inv}(\mathcal{R})$.

Proposition 68. *Let R be a relation with a first-order definition in $(\mathbb{N}, =)$. Then the following are equivalent.*

- (1) *R is preserved by a binary operation from $\text{Pol}(\neq)$ that violates ODD_3 .*
- (2) *Every reduced definition of R is negative.*
- (3) *R has a negative definition.*
- (4) *R is preserved by \mathcal{R} .*

Proof. To see that (1) implies (2), fix any reduced definition of R . Since R is preserved by a binary operation from $\text{Pol}(\neq)$, Proposition 43 implies that this definition is Horn. As there is a binary operation preserving R and \neq which violates ODD_3 , there is no pp definition of ODD_3 from R and \neq . Hence, our reduced definition must be negative, by Theorem 67.

The implication from (2) to (3) is trivial.

To show that (3) implies (4), fix a negative definition Φ of R . Let $f \in \mathcal{R}^{(n)}$ and $a_1, \dots, a_n \in R$ be arbitrary. Assume without loss of generality that f is injective in the 1-st direction. Then $b = f(a_1, \dots, a_n)$ satisfies all inequalities that were satisfied by a_1 , which is enough to satisfy Φ as Φ is negative. Hence, f preserves R .

The implication from (4) to (1) is witnessed by any binary operation which is not in \mathcal{R} but injective in one direction, as is easily verified. \square

The following proposition establishes statement (4) of Theorem 13.

Proposition 69. *The clones \mathcal{R} and \mathcal{S} form a maximal antichain in the interval $[\mathcal{B}, \text{Pol}(\neq)]$. In fact, every clone in this interval is either contained in \mathcal{S} or contains \mathcal{R} . In particular, \mathcal{R} is a cover of $\mathcal{R} \cap \mathcal{S}$.*

Proof. First, we have to show that \mathcal{R} and \mathcal{S} are indeed incomparable. By Proposition 68 it is clear that \mathcal{R} is not a subset of \mathcal{S} , because it contains an operation that violates ODD_3 . To show that \mathcal{S} is not a subset of \mathcal{R} , it suffices to show that the operation f_3 from Definition 12 is not injective in any direction. Indeed, this holds as $f_3(x, 1, 1) = 1$ and $f_3(2, y, 2) = 2$ and $f_3(3, 3, z) = 3$ for all $x, y, z \in \mathbb{N}$. Therefore, $f \notin \mathcal{R}$ and hence \mathcal{R} and \mathcal{S} form an antichain.

Now, let \mathcal{C} be any clone from the interval $[\mathcal{B}, \text{Pol}(\neq)]$ which is not contained in \mathcal{S} . Then it contains an operation in $\text{Pol}(\neq)$ that does not preserve ODD_3 . Since ODD_3 has two orbits, \mathcal{C} contains a binary operation with these properties, by Lemma 44. Thus, Proposition 68 implies that \mathcal{C} is above \mathcal{R} .

Hence, the antichain that consists of \mathcal{R} and \mathcal{S} is maximal, and every clone of the interval not contained in \mathcal{S} does contain \mathcal{R} . \square

6.5. Richard's many friends. We now prove that there exists an antitone and injective mapping from the power set of ω into the interval $[\mathcal{R}, \text{Pol}(\neq)]$; in particular, this interval has cardinality 2^{\aleph_0} . This proves the last statement (5) of Theorem 13.

Definition 70. For all $n \geq 3$, write

$$\delta_n := x_1 \neq y_1 \vee \dots \vee x_n \neq y_n.$$

For all $A \subseteq \{1, \dots, n\}$ with $1 < |A| < n$, writing $A = \{j_1, \dots, j_r\}$ with $j_1 < j_2 < \dots < j_r$, we set

$$\kappa_A := y_{j_1} \neq x_{j_2} \vee y_{j_2} \neq x_{j_3} \vee \dots \vee y_{j_r} \neq x_{j_1}.$$

Set

$$\gamma_n := \delta_n \wedge \bigwedge_{A \subseteq \{1, \dots, n\}, 1 < |A| < n} \kappa_A.$$

Observe that the γ_n are all negative formulas; hence, by Proposition 68, polymorphism clones of sets of relations defined by such γ_n will always contain \mathcal{R} . In the rest of this section, we will prove that no fixed γ_n can be defined from the others by primitive positive definitions. Therefore, distinct sets of such formulas define distinct clones above \mathcal{R} , which is what we want.

Definition 71. We now enumerate tuples $c \in \mathbb{N}^{2n}$ as $c = (c^{1,x}, c^{1,y}, \dots, c^{n,x}, c^{n,y})$. If $\phi(x^1, y^1, \dots, x^n, y^n)$ is a formula, then we say that c satisfies ϕ iff $\phi(c^{1,x}, c^{1,y}, \dots, c^{n,x}, c^{n,y})$ holds; this is to say that we insert $c^{i,u}$ for u^i , where $u \in \{x, y\}$ and $1 \leq i \leq n$. Moreover, if $c_1, \dots, c_m \in \mathbb{N}^{2n}$, we denote the “ (i, x) -th column tuple” $(c_1^{i,x}, \dots, c_m^{i,x})$ by $c^{i,x}$.

Definition 72. For all $n \geq 3$, we set C_n to be the $2n$ -ary relation defined by γ_n .

The operations of the following definition will separate the polymorphism clones of the C_n .

Definition 73. For all $n \geq 3$, the *Hubie-violator* H_n is an operation defined as follows: Enumerate $C_n \cap \{1, \dots, n+1\}^{2n}$, that is, those elements of C_n which have only entries in $\{1, \dots, n+1\}$, by c_1, \dots, c_m . Note that $m > 0$ since C_n always contains the tuple where $x_i = 0$ and $y_i = 1$, for all $i \leq n$. Now we define $H_n \in \mathcal{O}^{(m)}$ by $H_n(c^{j,x}) = H_n(c^{j,y}) = j$, for all $1 \leq j \leq n$; in other words, we set $H_n(c_1, \dots, c_m) = (1, 1, \dots, n, n)$. For every other input tuple in \mathbb{N}^m , H_n takes a unique value distinct from all other values in its range.

Lemma 74. Let $n \geq 3$. Then the Hubie-violator H_n violates C_n .

Proof. By definition, $H_n(c_1, \dots, c_m) = (1, 1, \dots, n, n)$, a tuple which does not satisfy δ_n and which is therefore not an element of C_n . □

Lemma 75. Let $3 \leq k < n$. Then the Hubie-violator H_n preserves C_k .

Proof. Write m for the arity of H_n . Let $t_1, \dots, t_m \in C_k$ and set $s = H_n(t_1, \dots, t_m)$. We must show $s \in C_k$.

We first check that s satisfies the clauses κ_A . Suppose it does not, and pick any $A \subseteq \{1, \dots, k\}$ with $1 < |A| < k$ witnessing this. We may assume wlog that $A = \{1, \dots, p\}$, where $1 < p < k$, and write $s = (a_p, a_1, a_1, a_2, a_2, a_3, \dots, a_{p-1}, a_{p-1}, a_p, ?, \dots)$. Set $Q := \{a_j : 1 \leq j \leq p\}$, and $Q_1 := Q \cap \{1, \dots, n\}$, and $Q_2 = Q \setminus Q_1$.

If Q_1 were empty, then since for every value in Q_2 there is only a unique tuple that H_n sends to that value, we have that the column $t^{i,y}$ equals the column $t^{i+1,x}$, for all $1 \leq i \leq p$ (where we set $p+1 := 1$). Thus in that case, no row t_j would satisfy κ_A , a contradiction.

Assuming that Q_2 were empty, we now show that there exists $1 \leq j \leq m$ such that the row t_j does not satisfy $\kappa_{\{1,2\}}$. Since $H_n(t^{1,x}) = a_p$, we have $t^{1,x} = c^{a_p,u}$, for some $u \in \{x, y\}$. Similarly, $t^{1,y} = c^{a_1,w}$, $t^{2,x} = c^{a_1,q}$, and $t^{2,y} = c^{a_2,v}$, where $w, q, v \in \{x, y\}$. Consider first the case where $a_p \neq a_1$ and $a_2 \neq a_1$ (a_p may well equal a_2). Pick $d \in C_n \cap \{1, \dots, n+1\}^{2n}$ such that $d^{a_1,x} = d^{a_1,y} = a_1$,

$d^{a_p,u} = a_2$, and $d^{a_2,v} = a_2$. This is without doubt possible by defining all other entries to be very unequal. Because in the definition of the Hubie-violator all tuples of $C_n \cap \{1, \dots, n+1\}^{2n}$ appear, there is $1 \leq j \leq m$ with $d = c_j$. We have: $t_j^{1,x} = c_j^{a_p,u} = a_2$, $t_j^{1,y} = c_j^{a_1,w} = a_1$, $t_j^{2,x} = c_j^{a_1,q} = a_1$, and $t_j^{2,y} = c_j^{a_2,v} = a_2$. Hence, $t_j = (a_2, a_1, a_1, a_2, ?, \dots)$ does not satisfy $\kappa_{\{1,2\}}$, a contradiction. The case where $a_p = a_1$ or $a_2 = a_1$ is even easier: Say wlog $a_p = a_1$, and pick $c_j \in C_n$ such that $c_j^{a_1,x} = c_j^{a_1,y} = a_1$, and $c_j^{a_2,v} = a_1$. We have: $t_j^{1,x} = c_j^{a_1,u} = a_1$, $t_j^{1,y} = c_j^{a_1,w} = a_1$, $t_j^{2,x} = c_j^{a_1,q} = a_1$, and $t_j^{2,y} = c_j^{a_2,v} = a_1$. Hence, in this case $t_j = (a_1, a_1, a_1, a_1, ?, \dots)$ does not satisfy $\kappa_{\{1,2\}}$ either, a contradiction.

Finally, consider the case where neither Q_1 nor Q_2 are empty. Assume wlog that $s^{1,x} = a_p \in Q_1$ and $s^{1,y} = a_1 \in Q_2$. We have $t^{1,x} = c^{a_p,v}$, for some $v \in \{x, y\}$. Let $r \leq p$ be minimal with the property that $s^{r,y} \in Q_1$; write $s^{r,y} = a_i$. Then $t^{r,y} = c^{a_i,u}$ for some $u \in \{x, y\}$. Pick $c_j \in C_n$ such that $c_j^{a_i,u} = c_j^{a_p,v} = a_1$. We then have that t_j does not satisfy $\kappa_{\{1,\dots,r\}}$: Indeed $t_j^{1,x} = a_1 = t_j^{r,y}$, and $t_j^{h,y} = t_j^{h+1,x}$ for all $1 \leq h < r$ as columns since $a_h \in Q_2$. A contradiction.

It remains to show that s satisfies δ_k . Suppose it does not; we claim that under this assumption, there is $1 \leq j \leq m$ such that t_j does not satisfy δ_k either. Set Q to consist of all values that appear in s , $Q_1 := Q \cap \{1, \dots, n\}$, and $Q_2 := Q \setminus Q_1$. Write p for the cardinality of Q_1 ; wlog we may assume $Q_1 = \{1, \dots, p\}$. Observe that since s has length $2k$ and since it does not satisfy δ_k , we have $p \leq k$. Consider the $2n$ -tuple $b = (1, 1, 2, 2, \dots, p, p, p+1, n+1, p+2, n+1, \dots, n, n+1)$. Using $p \leq k < n$, one readily checks that $b \in C_n$, hence $b = c_j$ for some $1 \leq j \leq m$. Thus, since c_j appears in the j -th row of the definition of the Hubie-violator H_n , we have that for all $d \in \mathbb{N}^m$ and all $1 \leq i \leq p$, if $H_n(d) = i$, then $d_j = i$. In particular, this holds for the vectors $t^{r,x}$ and $t^{r,y}$ for which $H_n(t^{r,x}) = H_n(t^{r,y}) \in Q_1$, and we have $t_j^{r,x} = t_j^{r,y}$ for those columns. But for the other columns $t^{r,x}, t^{r,y}$ whose (equal) image under H_n is in Q_2 , we have $t_j^{r,x} = t_j^{r,y}$ anyway since values outside $\{1, \dots, n\}$ are taken by at most one tuple. Summarizing these two observations, we have that $t_j^{r,x} = t_j^{r,y}$ for all $1 \leq r \leq k$, so t_j does not satisfy δ_k . A contradiction. \square

We thus know that for all $n \geq 3$, C_n has no pp-definition from $\{C_3, \dots, C_{n-1}\}$. For otherwise, the polymorphisms of $\{C_3, \dots, C_{n-1}\}$ would be contained in those of C_n , which is not the case by the preceding two lemmas. It remains to show that C_n cannot be defined from the C_k with $k > n$ either.

Lemma 76. *Let $3 \leq k < n$. Then the Hubie-violator H_k preserves C_n .*

Proof. Denote the arity of H_k by m . As before, let $t_1, \dots, t_m \in C_n$ and set $s := H_k(t_1, \dots, t_m)$. To see that s satisfies the clauses κ_A , one proceeds as in the preceding proof, with k and n exchanged. What is more difficult here is to check that s satisfies δ_n . Suppose it does not, and define Q, Q_1, Q_2 , as well as $p = |Q_1|$, as before (only with k and n switched). Now we distinguish two cases, $p < k$ and $p = k$.

In the first case, we argue like at the end of the proof of the preceding lemma: Consider the tuple $b = (1, 1, 2, 2, \dots, p, p, p+1, k+1, p+2, k+1, \dots, k, k+1)$. One readily checks that $b \in C_k$, so $b = c_j$ for some $1 \leq j \leq m$. Thus, since c_j appears in the j -th line of the definition of H_k , we have that for all $d \in \mathbb{N}^m$ and all $1 \leq i \leq p$, if $H_k(d) = i$, then $d_j = i$. In particular, this holds for those vectors $t^{r,x}$ and $t^{r,y}$ which are sent by H_k to a value in Q_1 . Since $t^{r,x} = t^{r,y}$ for the other columns (i.e., those with value in Q_2), this implies that $t_j^{r,x} = t_j^{r,y}$ for all $1 \leq r \leq n$, so t_j does not satisfy δ_n , a contradiction.

To finish the proof, consider the case where $p = k$. Observe first that for every $1 \leq i \leq k$, there exists exactly one $1 \leq r \leq n$ such that $s^{r,x} = s^{r,y} = i$. There is

at least one such r as $p = k$; suppose there were distinct r, r' with this property. Then $t^{r,x}, t^{r,y}, t^{r',x}, t^{r',y}$ are all elements of $\{c^{i,x}, c^{i,y}\}$. Pick any tuple $c_j \in C_k$ with $c^{i,x} = c^{i,y}$. We have that $t_j^{r,x}, t_j^{r,y}, t_j^{r',x}, t_j^{r',y}$ are all equal, hence t_j does not satisfy $\kappa_{\{r,r'\}}$, a contradiction. We may therefore wlog assume that s looks like this: $s = (1, 1, \dots, k, k, ?, \dots)$, where all entries starting from the question mark are elements of Q_2 . If we had $t^{i,x} = t^{i,y}$ for some $1 \leq i \leq k$, then we could derive a contradiction just like in the case $p < k$, by taking i out of Q_1 . So assume this is not the case. Then the set of vectors $\{t^{1,x}, t^{1,y}, \dots, t^{k,x}, t^{k,y}\}$ equals $\{c^{1,x}, c^{1,y}, \dots, c^{k,x}, c^{k,y}\}$. We prove that there exists $1 \leq j \leq m$ such that t_j does not satisfy $\kappa_{\{1, \dots, k\}}$. Let σ be the permutation on the indices of the tuple $(c^{1,x}, c^{1,y}, \dots, c^{k,x}, c^{k,y})$ which sends this tuple to $(t^{1,x}, t^{1,y}, \dots, t^{k,x}, t^{k,y})$. Observe that σ only switches $c^{j,x}$ and $c^{j,y}$, if necessary. Now consider the tuple $b := (k, 1, 1, 2, 2, \dots, k-1, k-1, k) \in C_k$, and apply σ^{-1} to the indices of this tuple, obtaining a new tuple $\sigma^{-1}(b)$. One readily checks $\sigma^{-1}(b) \in C_k$, so $\sigma^{-1}(b) = c_j$ for some $1 \leq j \leq m$. Thus the tuple $(t_j^{1,x}, t_j^{1,y}, \dots, t_j^{k,x}, t_j^{k,y})$, which is obtained by applying σ to $c_j = \sigma^{-1}(b)$, equals b . But b does not satisfy $\kappa_{\{1, \dots, k\}}$, so $t_j \notin C_n$, a contradiction. \square

We end this section by the following proposition, which completes the proof of Theorem 13 and of the corollary after.

Proposition 77. *The mapping σ from the power set of ω into $\text{Cl}_{\text{loc}}(\mathbb{N})$ which sends every $A \subseteq \omega$ to the local clone $\text{Pol}(\{C_{n+3} : n \in A\})$ is injective and antitone (with respect to inclusion, on both sides). In particular, the number of local clones containing $S_{\mathbb{N}}$ equals the continuum.*

Proof. Since Pol is antitone, it follows that the same holds for σ . Now let $A, B \subseteq \omega$ be unequal, say wlog that there is $n \in A \setminus B$. Then the Hubie-violator H_{n+3} violates C_{n+3} but preserves all C_{k+3} , where $k \in B$. Thus $H_{n+3} \in \sigma(B) \setminus \sigma(A)$, and σ is indeed injective. \square

7. CLONES WITH ESSENTIAL INFINITE RANGE OPERATIONS PLUS CONSTANTS

It remains to investigate one last monoidal interval, namely the one corresponding to the monoid \mathcal{S}^+ consisting of all unary operations which are either injective or constant. We will thus be concerned with the proof of Theorem 15 in this section. The crucial ingredients to the proof will be Proposition 69 and the following lemma; the definition of \mathcal{S} was given in Definition 12.

Lemma 78. $\langle \mathcal{S}^+ \rangle = \mathcal{S}^+$.

Proof. We have to show that \mathcal{S}^+ is a local clone. Let $f \in \langle \mathcal{S}^+ \rangle$ be an arbitrary non-constant operation. We claim that $f \in \mathcal{S}$. To see this, let A be any finite subset of \mathbb{N}^n , where n is the arity of f . Extend A to a finite set $B \subseteq \mathbb{N}^n$ such that f is non-constant on B . Since $f \in \langle \mathcal{S}^+ \rangle$, and since \mathcal{S} is by definition generated by f_3 , there exists a term t using f_3 and the constants which agrees with f on B . Since t is non-constant, it can as well be written without constants, using f_3 only: This is because inserting a constant as an argument of f_3 gives us an injection or essentially a bar operation, both of which are generated by f_3 without using constants. Thus, f can be interpolated on B , and hence also on A , by a term in \mathcal{S} . This proves $f \in \mathcal{S}$. \square

Recall that \mathcal{R} consists of all operations which are injective in one direction (Definition 11). Whereas \mathcal{S}^+ is a local clone, we cannot add constants to \mathcal{R} without generating all operations:

Proposition 79. $\langle \mathcal{R}^+ \rangle = \mathcal{O}$.

Proof. Let $f(x_1, \dots, x_n) \in \mathcal{O}$ be arbitrary, and let $A \subseteq \mathbb{N}^n$ be finite. Define $g(x_1, \dots, x_{n+1})$ as follows: If $x_{n+1} = 0$ and $(x_1, \dots, x_n) \in A$, then g returns $f(x_1, \dots, x_n)$. For every other input tuple, g returns a unique value distinct from all other values in its range. Clearly, g is injective in the $(n+1)$ -st direction, so $g \in \mathcal{R}$. Moreover, $g(x_1, \dots, x_n, 0)$ agrees with f on A , implying $f \in \langle \mathcal{R}^+ \rangle$ as A was arbitrary. \square

Note that the preceding proposition proves item (7) of Theorem 15. We have found the largest element of our monoidal interval $I_{\mathcal{J}^+}$:

Lemma 80. $\text{Pol}(\mathcal{J}^+) = \mathcal{J}^+$.

Proof. Clearly, the unary fragment of \mathcal{J}^+ equals \mathcal{J}^+ , hence \mathcal{J}^+ is an element of the corresponding monoidal interval $I_{\mathcal{J}^+}$. Suppose there was a local clone \mathcal{C} in this interval which properly contains \mathcal{J}^+ . Take any $f \in \mathcal{C} \setminus \mathcal{J}^+$. Clearly, f must be essential. Since the only non-injections in the unary fragment of \mathcal{C} are constant, Lemma 36 implies that f must have infinite range. By Proposition 37, f preserves \neq . Thus, $f \in \text{Pol}(\neq) \setminus \mathcal{J}$, so f generates \mathcal{R} by Proposition 69. Hence $\mathcal{C} = \mathcal{O}$ by the preceding proposition. \square

The following proposition proves statement (1) of Theorem 15.

Proposition 81. *Let $\mathcal{C} \in I_{\mathcal{J}^+}$. Then $\mathcal{C}^- = \mathcal{C} \cap \mathcal{J}$. In particular, \mathcal{C}^- is a local clone in $I_{\mathcal{J}}$.*

Proof. By Lemma 80, \mathcal{C} is a subclone of \mathcal{J}^+ ; hence intersecting \mathcal{C} with \mathcal{J} just means taking away the constants, which proves our assertion. Since the intersection of local clones is a local clone, so is \mathcal{C}^- . Clearly, the unary fragment of \mathcal{C}^- equals \mathcal{J} , hence \mathcal{C}^- is an element of $I_{\mathcal{J}}$. \square

Observe that the second statement of the theorem is trivial. We turn to the proof of item (3).

Proposition 82. *The mapping from $I_{\mathcal{J}^+}$ into the subinterval $[\langle \mathcal{J} \rangle, \mathcal{J}]$ of $I_{\mathcal{J}}$ which sends every clone \mathcal{C} to \mathcal{C}^- is a complete lattice embedding which preserves the smallest and the largest element.*

Proof. It is clear that this mapping preserves the order and by the second statement of Theorem 15, it sends the smallest (largest) element of $I_{\mathcal{J}^+}$ to the smallest (largest) element of $[\langle \mathcal{J} \rangle, \mathcal{J}]$. It is also obvious that the mapping is injective. Let $\mathcal{C}, \mathcal{D} \in I_{\mathcal{J}^+}$. Clearly, $(\mathcal{C} \cap \mathcal{D})^- = \mathcal{C}^- \cap \mathcal{D}^-$, hence the mapping preserves finite meets; larger meets work the same way. Since $(\mathcal{C} \vee \mathcal{D})^-$ contains both \mathcal{C}^- and \mathcal{D}^- , it contains also their join, so $(\mathcal{C} \vee \mathcal{D})^- \supseteq \mathcal{C}^- \vee \mathcal{D}^-$. For the other inclusion, let f be any non-constant operation in the (not necessarily local) clone generated by $\mathcal{C} \cup \mathcal{D}$. Then f can be written as a term over $\mathcal{C} \cup \mathcal{D}$. We can assume that in this term, thinking of it as a tree, the constants appear only as leaves. We may also assume that except for leaves, no subtree of this tree is constant. That just means that the nodes above the leaves are operation symbols in $\mathcal{C}^- \cup \mathcal{D}^-$, some of whose variables are set to constant values. But these operations in $\mathcal{C}^- \cup \mathcal{D}^-$ with some variables set to constants are again elements of $\mathcal{C}^- \cup \mathcal{D}^-$, as they are non-constant by assumption and since \mathcal{C} and \mathcal{D} are clones containing all constants. Summarizing, we can write f as a term over $\mathcal{C}^- \cup \mathcal{D}^-$ without the use of constants. Hence, $f \in \mathcal{C}^- \vee \mathcal{D}^-$. This proves $(\mathcal{C} \vee \mathcal{D})^- \subseteq \mathcal{C}^- \vee \mathcal{D}^-$, and hence our mapping preserves finite joins. Infinite joins work the same way. \square

Observe that item (4) of Theorem 15 is clear from the definitions. We prove (5).

Proposition 83. *For all $\mathcal{C} \in I_{\mathcal{J}}$ which do not contain \mathcal{R} , $\langle \mathcal{C}^+ \rangle$ is a local clone in $I_{\mathcal{J}^+}$.*

Proof. By Proposition 69, \mathcal{C} is contained in \mathcal{S} . Hence, $\langle \mathcal{C}^+ \rangle$ is contained in \mathcal{S}^+ by Lemma 78, which equals $\text{Pol}(\mathcal{S}^+)$ by Lemma 80. Therefore $\langle \mathcal{C}^+ \rangle$ is indeed a local clone in $I_{\mathcal{S}^+}$. \square

Item (6) is trivial from the definitions. We finish the proof of the theorem by the following proposition, which restates items (8) and (9):

Proposition 84. \mathcal{H}^+ is the unique cover of $\langle \mathcal{S}^+ \rangle$ and \mathcal{B}^+ is the unique cover of \mathcal{H}^+ in $I_{\mathcal{S}^+}$.

Proof. Let \mathcal{C} be in $I_{\mathcal{S}^+}$, and assume it has an essential operation f . Then $f \in \mathcal{S}$ by Lemma 80, thus f generates \mathcal{H} by Proposition 43. It is easy to see that \mathcal{H}^+ is a local clone, which proves the first assertion. Similarly, if \mathcal{C} is in $I_{\mathcal{S}^+}$ and properly contains \mathcal{H}^+ , then any operation witnessing this generates \mathcal{B} , by Proposition 50. Again, it is straightforward to check that \mathcal{B}^+ is a local clone. \square

We want to close this section with the remark that the mapping which sends every clone \mathcal{C} in the interval $[\langle \mathcal{S} \rangle, \mathcal{S}]$ of $I_{\mathcal{S}}$ to $\langle \mathcal{C}^+ \rangle$ is not injective. In particular, it collapses all clones between $\mathcal{R} \cap \mathcal{S}$ and \mathcal{S} .

Proposition 85. $\langle (\mathcal{R} \cap \mathcal{S})^+ \rangle = \mathcal{S}^+$

Proof. The operation g_3 (Definition 58) is injective in the first argument, and hence contained in \mathcal{R} . By Lemma 59, g_3 preserves \underline{R}_2 , and, being an element of \mathcal{R} , it also preserves \neq . Proposition 62 then implies that g_3 is contained in \mathcal{S} : Otherwise, it would violate a relation preserved by \mathcal{S} ; that relation would have a pp definition from \underline{R}_2 and \neq , a contradiction. Therefore, the operation g_3 is in $\mathcal{R} \cap \mathcal{S}$.

The operation given by $g_3(0, x_1, x_2, x_3)$ equals by definition the operation f_3 (Definition 12 or 57), which by definition generates \mathcal{S} . Hence, $\langle (\mathcal{R} \cap \mathcal{S})^+ \rangle$ contains \mathcal{S}^+ . The converse containment is trivial. \square

Observe that this proposition implies that the complete lattice embedding that sends every $\mathcal{C} \in I_{\mathcal{S}^+}$ to \mathcal{C}^- is not surjective onto the interval $[\langle \mathcal{S} \rangle, \mathcal{S}]$: If there existed a clone $\mathcal{C} \in I_{\mathcal{S}^+}$ such that $\mathcal{C}^- = \mathcal{R} \cap \mathcal{S}$, then $\mathcal{C} = (\mathcal{R} \cap \mathcal{S})^+$ since $\mathcal{C} = (\mathcal{C}^-)^+$. Thus $(\mathcal{R} \cap \mathcal{S})^+$ would be a clone and hence equal to \mathcal{S}^+ , implying $\mathcal{R} \cap \mathcal{S} = \mathcal{S}$, a contradiction. In fact the same argument shows that a clone $\mathcal{D} \in [\langle \mathcal{S} \rangle, \mathcal{S}]$ is in the range of this embedding iff \mathcal{D}^+ is a clone.

8. OPEN PROBLEMS

We have to leave the complete description of the monoidal intervals corresponding to \mathcal{S} and \mathcal{S}^+ open; we do not know if a reasonable description of these intervals is possible at all. Two particular problems seem most important to us at this point:

- Can one effectively decide whether for a given sequence R_0, R_1, \dots, R_n of relations that are first-order definable in $(\mathbb{N}, =)$ there is a pp-definition of R_0 in the structure $(\mathbb{N}, R_1, \dots, R_n)$?
- What is the cardinality of the monoidal interval corresponding to \mathcal{S}^+ ?

Acknowledgements. We thank the referee for his helpful suggestions.

REFERENCES

- [BC07] M. Bodirsky and H. Chen. Quantified equality constraints. In *LICS'07*, pages 203–212, 2007.
- [BCRV04] E. Böhrer, N. Creignou, S. Reith, and H. Vollmer. Playing with boolean blocks, part ii: constraint satisfaction problems. *ACM SIGACT-Newsletter*, 35(1):22–35, 2004.
- [BK08a] M. Bodirsky and J. Kára. The complexity of equality constraint languages. *Theory of Computing Systems*, 3(2):136–158, 2008. A conference version of the paper appeared in the proceedings of the International Computer Science Symposium in Russia (CSR'06).

- [BK08b] M. Bodirsky and J. Kára. The complexity of temporal constraint satisfaction problems. In *Proceedings of STOC'08*, pages 29–38, 2008.
- [BN06] M. Bodirsky and J. Nešetřil. Constraint satisfaction with countable homogeneous templates. *J. Logic Comput.*, 16(3):359–373, 2006.
- [Cam76] P. J. Cameron. Transitivity of permutation groups on unordered sets. *Math. Z.*, 148:127–139, 1976.
- [Cam90] P. J. Cameron. *Oligomorphic permutation groups*, volume 152 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 1990.
- [CD73] P. Crawley and R. P. Dilworth. *Algebraic theory of lattices*. Prentice-Hall, 1973.
- [Die05] R. Diestel. *Graph Theory, 3rd edition*. Springer-Verlag, New York, 2005.
- [Gol] M. Goldstern. Analytic clones. Preprint available from <http://arxiv.org/math.RA/0404214>.
- [GP08] M. Goldstern and M. Pinsker. A survey of clones on infinite sets. *Algebra Universalis*, 59:365–403, 2008.
- [GS05] M. Goldstern and S. Shelah. Clones from creatures. *Trans. Amer. Math. Soc.*, 357(9):3525–3551, 2005.
- [GSS] M. Goldstern, S. Sági, and S. Shelah. Very many clones above the unary clone. Preprint.
- [Hei02] L. Heindorf. The maximal clones on countable sets that include all permutations. *Algebra Universalis*, 48:209–222, 2002.
- [HR94] L. Haddad and I. G. Rosenberg. Finite clones containing all permutations. *Canad. J. Math.*, 46(5):951–970, 1994.
- [JZ08] M. Junker and M. Ziegler. The 116 reducts of $(\mathbb{Q}, <, a)$. *Journal of Symbolic Logic*, 73(3):861–884, 2008.
- [Kru72] J. B. Kruskal. The theory of well-quasi-ordering: A frequently discovered concept. *Journal of Combinatorial Theory (A)*, 13:297–305, 1972.
- [Lau06] D. Lau. *Function algebras on finite sets*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2006.
- [MP07] H. Machida and M. Pinsker. The minimal clones above the permutations. *Semigroup Forum*, 75:181–211, 2007.
- [Mil85] E. C. Milner. Basic wqo and bqo theory. In *Graphs and Orders* (I. Rival, ed.), Reidel, Boston, Mass., pages 487–502, 1985.
- [Pin] M. Pinsker. More sublattices of the lattice of local clones. Preprint available from <http://dmg.tuwien.ac.at/pinsker/>.
- [Pin05a] M. Pinsker. Maximal clones on uncountable sets that include all permutations. *Algebra Universalis*, 54(2):129–148, 2005.
- [Pin05b] M. Pinsker. The number of unary clones containing the permutations on an infinite set. *Acta Sci. Math.*, 71:461–467, 2005.
- [Pin08] M. Pinsker. Sublattices of the lattice of local clones. In *Proceedings of the ROGICS'08 conference*, pages 80–87, 2008.
- [PK79] R. Pöschel and L. Kalužnin. *Funktionen- und Relationenalgebren*. VEB Deutscher Verlag der Wissenschaften, 1979.
- [Pos41] E. L. Post. The two-valued iterative systems of mathematical logic. *Annals of Math. Studies*, 5, 1941.
- [RS82] I. G. Rosenberg and D. Schweigert. Locally maximal clones. *Elektron. Informationsverarb. Kybernet.*, 18(7-8):389–401, 1982.
- [Sch78] T. J. Schaefer. The complexity of satisfiability problems. In *Proceedings of STOC'78*, pages 216–226, 1978.
- [Sch89] U. Schöning. *Logic for Computer Scientists*. Springer, 1989.
- [Sze86] Á. Szendrei. *Clones in universal algebra*. Les Presses de L'Université de Montréal, 1986.
- [Tho91] S. Thomas. Reducts of the random graph. *Journal of Symbolic Logic*, 56(1):176–181, 1991.

LABORATOIRE D'INFORMATIQUE (LIX), CNRS UMR 7161, ÉCOLE POLYTECHNIQUE, 91128
PALAISEAU, FRANCE

E-mail address: `bodirsky@lix.polytechnique.fr`

URL: `http://www.lix.polytechnique.fr/~bodirsky/`

DEPARTAMENT DE TECNOLOGIES DE LA INFORMACIÓ I LES COMUNICACIONS, UNIVERSITAT POM-
PEU FABRA, BARCELONA, SPAIN

E-mail address: `hubie.chen@upf.edu`

URL: `http://www.tecn.upf.es/~hchen/`

LABORATOIRE DE MATHÉMATIQUES NICOLAS ORESME, CNRS UMR 6139, UNIVERSITÉ DE
CAEN, 14032 CAEN CEDEX, FRANCE

E-mail address: `marula@gmx.at`

URL: `http://dmg.tuwien.ac.at/pinsker/`