

The Power of Primitive Positive Definitions with Polynomially Many Variables

Victor Lagerkvist¹ and Magnus Wahlström²

¹ Department of Computer and Information Science, Linköping University, Sweden
victor.lagerkvist@liu.se (Corresponding author)

² Department of Computer Science, Royal Holloway, University of London, Great Britain
Magnus.Wahlstrom@rhul.ac.uk

Abstract. Two well-studied closure operators for relations are based on existentially quantified conjunctive formulas, *primitive positive* (p.p.) definitions, and primitive positive formulas without existential quantification, *quantifier-free primitive positive definitions* (q.f.p.p.) definitions. Sets of relations closed under p.p. definitions are known as *co-clones* and sets of relations closed under q.f.p.p. definitions as *weak partial co-clones*. The latter do however have limited expressivity, and the corresponding lattice of strong partial clones is of uncountably infinite cardinality even for the Boolean domain. Hence, it is reasonable to consider the expressiveness of p.p. definitions where only a small number of existentially quantified variables are allowed. In this paper we consider p.p. definitions allowing only polynomially many existentially quantified variables, and say that a co-clone closed under such definitions is *polynomially closed*, and otherwise *superpolynomially closed*. We investigate properties of polynomially closed co-clones and prove that if the corresponding clone contains a k -ary near-unanimity operation for $k \geq 3$ then the co-clone is polynomially closed, and if the clone does not contain a k -edge operation for any $k \geq 2$, then the co-clone is superpolynomially closed. For the Boolean domain we strengthen these results and prove a complete dichotomy theorem separating polynomially closed co-clones from superpolynomially closed co-clones. Using these results, we then proceed to investigate properties of strong partial clones corresponding to superpolynomially closed co-clones. We prove that if Γ is a finite set of relations over an arbitrary finite domain such that the clone corresponding to Γ is essentially unary, then the strong partial clone corresponding to Γ is of infinite order and cannot be generated by a finite set of partial functions.

Keywords: Clone theory, partial clone theory, universal algebra, primitive positive definitions, constraint satisfaction problems

1 Introduction

A finite or infinite set of relations Γ over a finite domain is known as a *constraint language*. Given a constraint language Γ , a natural question to ask is which other relations R can be expressed by first order formulas over Γ , or, equivalently, what is the smallest set of relations that contains Γ and is closed under such definitions. In practice one often considers restricted first order formulas, and two common restrictions are *primitive positive definitions* (p.p. definitions), where one is allowed to use existential quantification, conjunction and equality constraints, and *quantifier-free primitive positive definitions* (q.f.p.p. definitions) where only conjunction and equality constraints are allowed. A *relational clone*, or a *co-clone*, is a set of relations closed under p.p. definitions. Any set of relations which generates a given co-clone using p.p. definitions is called a *base* of the co-clone. Similarly, a set of relations closed under q.f.p.p. definitions is referred to as a *weak partial co-clone*, or a *weak system*. Both co-clones and weak partial co-clones have interesting applications in theoretical computer science, and in particular, for the study of the computational complexity of problems parameterized by constraint languages. One noteworthy example is the *constraint satisfaction problem* over a constraint language Γ ($\text{CSP}(\Gamma)$), which is the problem of determining whether a conjunctive formula over Γ has a model. The use of algebraic techniques to study the complexity of $\text{CSP}(\Gamma)$ is usually referred to as the *algebraic approach* and was first pioneered by Jeavons [15]. The success of the algebraic approach can be mainly attributed to the fact that for

every constraint language Γ one can associate a particular function algebra, known as the *clone* of Γ , or the *polymorphisms* of Γ . The exact details of this relationship between a constraint language and its polymorphisms will be covered later, but for the moment we merely state that this relationship provides a powerful method for analyzing the complexity of constraint satisfaction and related problems. Using this relationship Jeavons proved that the complexity of $\text{CSP}(\Gamma)$ up to polynomial-time reductions is determined by the polymorphisms of Γ [15]. Since then, this result has been extended and used in numerous applications, cf. the excellent survey by Creignou et al. [10] for a broad introduction to this topic.

There is also a similar relationship between weak partial co-clones and sets of partial functions closed under composition, containing all total and partial projection functions, *strong partial clones*. Again, we omit the exact definitions of these concepts for the moment, and just state that for every weak partial co-clone there exists a set of partial functions, *partial polymorphisms*, which completely characterizes this set. With this relationship Jonsson et al. proved that the partial polymorphisms of a constraint language Γ determines the complexity of $\text{CSP}(\Gamma)$ up to $O(c^n)$ complexity [18], where n denotes the number of variables in a given $\text{CSP}(\Gamma)$ instance. This result was used to give lower bounds for all NP-complete Boolean $\text{CSP}(\Gamma)$ problems. Similar results were given in Jonsson et al. [19] but in the context of Boolean optimization problems. Hence, strong partial clones and weak partial co-clones lead to interesting applications when comparing and relating computational problems vis-à-vis $O(c^n)$ time complexity.

Unfortunately, the seemingly subtle steps from p.p. definitions to q.f.p.p. definitions, and from total to partial functions, makes reasoning much more complex. One of the reasons is that, unlike Post's lattice of Boolean clones [24], the lattice of strong partial clones is of uncountably infinite cardinality even for the Boolean domain [1]. Given this fact it is reasonable to consider the expressive power of closure operators which lie between q.f.p.p. definitions and p.p. definitions. To find logical formulas of such intermediate complexity we in this article restrict the number of existentially quantified variables occurring in formulas, and are therefore interested in which n -ary relations that can be p.p. defined with $1, 2, \dots, p(n)$ existentially quantified variables, for some reasonably slowly growing function p . In the sequel we assume that p is a polynomial function. If $p(n)$ variables are sufficient to define every n -ary relation R in a co-clone then we say that the co-clone is *polynomially closed*. We remark that if $p(n) \leq 2$ then the resulting set of definable relations over some language Γ closely corresponds to the closure operator considered in Nordh and Zanuttini [23].

The first contribution of this article is a complete classification of the polynomially closed Boolean co-clones (in Section 3). Our proofs are based on comparing the least expressive base of the co-clone with the most expressive base of the co-clone, in order to obtain an upper bound of p . These languages are known as the weak base and plain base, respectively, and were introduced by Schnoor and Schnoor [28], and Creignou et al. [9]. We first give a general result and provide a sufficient condition for a co-clone over any finite domain to be polynomially closed: a co-clone X is polynomially closed if the clone corresponding to X contains a *k-ary near-unanimity* function for some $k \geq 3$. We then complete this classification for the Boolean domain and in addition prove that a Boolean co-clone X is polynomially closed if the polymorphisms of X can be represented by affine functions, or if X is of infinite order (i.e., that X does not have a finite base). To handle the last case we extend Schnoor and Schnoor's result [28] for constructing weak bases and give a condition for the existence of weak bases for co-clones of infinite order. In Section 5 we then proceed with the problem of determining whether a co-clone is superpolynomially closed. We first prove that if the number of n -ary relations in a co-clone is sufficiently large, then, for any finite base of the co-clone, there exists relations which cannot be p.p. defined using a polynomial number of existentially quantified variables. By a result of Berman et al. [3] we then obtain a sufficient condition for verifying whether a co-clone over any finite domain is superpolynomially closed. We remark that for the Boolean domain, a co-clone of finite order is polynomially closed if and only if the corresponding clone contains a *k-edge function* for some $k \geq 2$, or, equivalently, if the clone has *few subpowers* [3]. Interestingly, this does not hold for co-clones of infinite order, which suggests a quantitative difference between our notion and that of Berman et al.

The second contribution of this article (in Section 6) is an investigation of the structure of the partial polymorphisms of finite constraint languages corresponding to superpolynomially closed co-clones. Before we can present this result, we need a few additional preliminaries. Given a constraint language Γ , say that the

set of partial polymorphisms of Γ is of *finite order* if there exists a finite set of partial functions F which can generate this set, using the standard notion of functional composition, and of *infinite order* otherwise. The set F is in this case called a *base* of the set of partial polymorphisms of Γ . Assume e.g. that $R_{1/3} = \{(0, 0, 1), (0, 1, 0), (1, 0, 0)\}$, and observe that $\text{CSP}(\{R_{1/3}\})$ is an alternative formulation of the well-known NP-complete problem 1-IN-3-SAT. It is easy to verify that the co-clone of $\{R_{1/3}\}$ is the set of all Boolean relations, and from the results in Section 5 we know that this set is superpolynomially closed. Given the fact that the partial polymorphisms of a constraint language has a close relationship with the worst-case time complexity of the corresponding CSP problem [18], obtaining a finite base of the set of partial polymorphisms of $R_{1/3}$ would likely increase our understanding of the time complexity of 1-IN-3-SAT. We prove that such a finite base *cannot exist* (irregardless of any complexity theoretical assumptions). In fact, we prove something stronger: let Γ be a finite constraint language over an arbitrary finite domain. If the co-clone of Γ is superpolynomially closed, and if the polymorphisms of Γ are essentially unary, then the set of partial polymorphisms of Γ is of infinite order. This result can be seen as a continuation of the research by Haddad and Börner [7] who gave a condition for checking whether a strong partial clone is infinitely generated, but our result also has many practical consequences for the applicability of partial clone theory to the study of the computational complexity of NP-hard CSP problems. Also, it is worth noting that the fact that even though a given strong partial clone is of infinite order, it might still be possible to give a reasonably simple characterization of its functions. This problem was investigated in Lagerkvist et al. [22] by considering stronger notions of closure than functional composition.

2 Preliminaries and Notation

If Γ is a constraint language we let $\Gamma^{(n)}$ be defined as $\{R \mid R \in \Gamma, \text{ar}(R) \leq n\}$, where $\text{ar}(R)$ is the arity of the relation R . Given a finite domain D , we let Rel_D be the set of all finitary relations over D , OP_D be the set of all functions over D , and we let Eq_D denote the equality relation $\{(x, x) \mid x \in D\}$ over D . For a tuple $t = (x_1, \dots, x_i, \dots, x_n) \in D^n$ we let $t[i] = x_i$. An *n-ary projection function* over D is a function π_i^n which for some $i \in \{1, \dots, n\}$ satisfies $\pi_i^n(x_1, \dots, x_i, \dots, x_n) = x_i$ for all $(x_1, \dots, x_i, \dots, x_n) \in D^n$. We typically represent relations and constraint languages by their defining logical formulas, and write $R(x_1, \dots, x_n) \equiv \phi$, where ϕ is a logical first-order formula, to denote the n -ary relation $R = \{(f(x_1), \dots, f(x_n)) \mid f \text{ is a model of } \phi\}$. As a convenience we often write \bar{x} instead of $\neg x$.

2.1 Clones, Co-Clones and Galois Connection

Let Γ be a constraint language over a finite domain D . If f is a function over D it is said to be a *polymorphism* of Γ , or that Γ is *invariant* under f , if, for every relation $R \in \Gamma$, $f(t_1, \dots, t_n) \in R$ for all $t_1, \dots, t_n \in R$. Here $f(t_1, \dots, t_n)$ denotes the $\text{ar}(R)$ -ary tuple obtained by the component-wise application of f to t_1, \dots, t_n , i.e., $f(t_1, \dots, t_n) = (f(t_1[1], \dots, t_n[1]), \dots, f(t_1[\text{ar}(R)], \dots, t_n[\text{ar}(R)]))$. If F is a set of functions over D and Γ a set of relations we let $\text{Pol}_D(\Gamma)$ denote the set of all polymorphisms over D of Γ , and $\text{Inv}_D(F)$ denote the set of all relations over D that are invariant under F . If the domain is clear from the context we simply write $\text{Inv}(F)$ and $\text{Pol}(\Gamma)$, respectively.

Sets of the form $\text{Pol}(\Gamma)$ are usually referred to as *clones*, and, as can be verified, are composition-closed sets of functions containing all projection functions. That is, if $f \in \text{Pol}(\Gamma)$ is an n -ary function and $g_1, \dots, g_n \in \text{Pol}(\Gamma)$ are m -ary functions, then $\text{Pol}(\Gamma)$ also contains the m -ary function

$$(f \circ (g_1, \dots, g_n))(x_1, \dots, x_m) = f(g_1(x_1, \dots, x_m), \dots, g_n(x_1, \dots, x_m)).$$

Dually, sets of the form $\text{Inv}(F)$ are referred to as *co-clones*, and are sets of relations closed under *primitive positive definitions* (p.p. definitions), i.e., whenever $\Gamma \subseteq \text{Inv}(F)$ then $\text{Inv}(F)$ also contains all n -ary relations R of the form $R(x_1, \dots, x_n) \equiv \exists y_1, \dots, y_{n'} . R_1(\mathbf{x}_1) \wedge \dots \wedge R_m(\mathbf{x}_m)$, where each $R_i \in \Gamma \cup \{\text{Eq}\}$ and each \mathbf{x}_i is an $\text{ar}(R_i)$ -ary tuple of variables over $x_1, \dots, x_n, y_1, \dots, y_{n'}$. Let $[F] = \text{Pol}(\text{Inv}(F))$ and $\langle \Gamma \rangle = \text{Inv}(\text{Pol}(\Gamma))$, and note that $[F]$ is then the smallest clone containing F , while $\langle \Gamma \rangle$ is the smallest co-clone containing Γ . The sets




Fig. 1. The lattice of Boolean co-clones. The co-clones of finite order which are polynomially closed are coloured in grey. The co-clones of infinite order that are polynomially closed are coloured in white. The superpolynomially closed co-clones are coloured in dark grey.

F and Γ are said to be *bases* of $[F]$ and $\langle \Gamma \rangle$, respectively, and a clone or a co-clone is said to be of *finite order* if it has a finite base, and is said to be of *infinite order* otherwise. We have the following *Galois connection* between $\text{Inv}(\cdot)$ and $\text{Pol}(\cdot)$.

Theorem 1 ([4, 5, 12]). *Let Γ and Δ be two constraint languages. Then $\Gamma \subseteq \langle \Delta \rangle$ if and only if $\text{Pol}(\Delta) \subseteq \text{Pol}(\Gamma)$.*

If D is a finite domain it is well known that the set of all clones over D form a lattice structure when ordered by set inclusion, where the meet-operator \sqcap is defined as $X \sqcap Y = X \cap Y$ and the join-operator \sqcup as $X \sqcup Y = [X \cup Y]$. Similarly, the set of all co-clones over D also form a lattice structure under set inclusion, where $X \sqcap Y = X \cap Y$ and $X \sqcup Y = \langle X \cup Y \rangle$. For the Boolean domain, all clones have been completely determined, and the lattice of Boolean clones is typically referred to as *Post's lattice* due to Post's original classification [24]. See Table 1 for a complete list of Boolean clones and their bases [6]. As a shorthand we let **BF** denote the set of all Boolean functions and **BR** the set of all Boolean relations. Due to the Galois connection in Theorem 1, each clone C in Table 1 uniquely determines a co-clone $\text{Inv}(C)$, and it is not difficult to see that the lattice of Boolean co-clones is dually isomorphic to the lattice of Boolean clones. See Figure 1 for a visualization of the Boolean co-clone lattice. In this figure, each node IC is an abbreviation of $\text{Inv}(C)$, where C is a clone from Table 1.

In this article we in addition need more restricted closure operators. Say that an n -ary relation R has a *quantifier-free primitive positive* (q.f.p.p.) definition in a constraint language Γ if $R(x_1, \dots, x_n) \equiv R_1(\mathbf{x}_1) \wedge \dots \wedge R_k(\mathbf{x}_k)$, where each $R_i \in \Gamma \cup \{\text{Eq}\}$ and each \mathbf{x}_i is an $\text{ar}(R_i)$ -ary tuple of variables over x_1, \dots, x_n . Hence, q.f.p.p. definitions are more restricted than p.p. definitions since we do not allow existential quantification. We also need an alternative notion of polymorphisms. An n -ary *partial function* over a finite domain is a map $f : X \mapsto D$, where $X \subseteq D^n$. In other words X is the set of arguments for which the function is defined. We let $\text{domain}(f) = X$, and say that f is *undefined* for all $(x_1, \dots, x_n) \in D^n \setminus \text{domain}(f)$. Composition of partial functions is defined in an analogous manner to the case of total functions, but the resulting function is only defined for those arguments where all involved functions are defined. Say that an n -ary partial function f is a *partial polymorphism* of a constraint language Γ , or that Γ is *invariant* under f , if, for every $R \in \Gamma$, $f(t_1, \dots, t_n) \in R$ for all $t_1, \dots, t_n \in R$ such that $f(t_1, \dots, t_n)$ is defined. Let $\text{pPol}_D(\Gamma)$ denote the set of all partial polymorphisms over D of a constraint language Γ and $\text{Inv}_D(F)$ the set of all relations over D invariant under the set of partial functions F . Sets of the form $\text{pPol}(\Gamma)$ are known as *strong partial clones*, and are composition-closed sets of partial functions, containing all projection functions, and closed under subfunctions. The latter means that whenever $f \in \text{pPol}(\Gamma)$ then $\text{pPol}(\Gamma)$ also contains all functions g such that $\text{domain}(g) \subseteq \text{domain}(f)$ and such that $g(x_1, \dots, x_n) = f(x_1, \dots, x_n)$ for all $(x_1, \dots, x_n) \in \text{domain}(g)$. Sets of the form $\text{Inv}(\text{pPol}(\Gamma))$ are known as *weak partial co-clones*, or, *weak systems*, and are sets of relations closed under q.f.p.p. definability. Given a set of partial functions F and a constraint language Γ let $[F]_s = \text{pPol}(\text{Inv}(F))$ and $\langle \Gamma \rangle_{\overline{\exists}} = \text{Inv}(\text{pPol}(\Gamma))$. As is easily verified $[F]_s$ is the smallest strong partial clone containing F and $\langle \Gamma \rangle_{\overline{\exists}}$ the smallest weak partial co-clone containing Γ . Observe that this implies that $\langle \Gamma \rangle_{\overline{\exists}}$ is the smallest set of relations which is closed under q.f.p.p. definitions over Γ . Similar to the definitions for clones and co-clones the sets F and Γ are said to be *bases* of $[F]_s$ and $\langle \Gamma \rangle_{\overline{\exists}}$, and we say that a strong partial clone or a weak partial co-clone is of *finite order* if it has a finite base, and is of *infinite order* otherwise. The relationship between strong partial clones and weak partial co-clones is given by the following Galois connection.

Theorem 2 ([12, 26]). *Let Γ and Δ be two constraint languages. Then $\Gamma \subseteq \langle \Delta \rangle_{\overline{\exists}}$ if and only if $\text{pPol}(\Delta) \subseteq \text{pPol}(\Gamma)$.*

2.2 Weak and Plain Bases of Co-Clones

The structure of the lattice of strong partial clones is largely undetermined, since it is of uncountably infinite cardinality for every non-trivial finite domain [1]. Due to the Galois connection in Theorem 2, this also implies that the dually isomorphic lattice of weak partial co-clones is of uncountably infinite cardinality. Despite this, it is possible to describe parts of this lattice by considering a particular kind of sublattice.

Table 1. List of all Boolean clones with definitions and bases, where $\text{id}(x) = x$ and $h_n(x_1, \dots, x_{n+1}) = \bigvee_{i=1}^{n+1} x_1 \cdots x_{i-1} x_{i+1} \cdots x_{n+1}$, $\text{dual}(f)(a_1, \dots, a_n) = 1 - f(\bar{a}_1, \dots, \bar{a}_n)$.

Clone	Definition	Base
BF	All Boolean functions	$\{x \wedge y, \neg x\}$
R ₀	$\{f \mid f \text{ is 0-reproducing}\}$	$\{x \wedge y, x \oplus y\}$
R ₁	$\{f \mid f \text{ is 1-reproducing}\}$	$\{x \vee y, x \oplus y \oplus 1\}$
R ₂	$R_0 \cap R_1$	$\{x \vee y, x \wedge (y \oplus z \oplus 1)\}$
M	$\{f \mid f \text{ is monotonic}\}$	$\{x \vee y, x \wedge y, 0, 1\}$
M ₁	$M \cap R_1$	$\{x \vee y, x \wedge y, 1\}$
M ₀	$M \cap R_0$	$\{x \vee y, x \wedge y, 0\}$
M ₂	$M \cap R_2$	$\{x \vee y, x \wedge y\}$
S ₀ ⁿ	$\{f \mid f \text{ is 0-separating of degree } n\}$	$\{x \rightarrow y, \text{dual}(h_n)\}$
S ₀	$\{f \mid f \text{ is 0-separating}\}$	$\{x \rightarrow y\}$
S ₁ ⁿ	$\{f \mid f \text{ is 1-separating of degree } n\}$	$\{x \wedge \neg y, h_n\}$
S ₁	$\{f \mid f \text{ is 1-separating}\}$	$\{x \wedge \neg y\}$
S ₀₂ ⁿ	$S_0^n \cap R_2$	$\{x \vee (y \wedge \neg z), \text{dual}(h_n)\}$
S ₀₂	$S_0 \cap R_2$	$\{x \vee (y \wedge \neg z)\}$
S ₀₁ ⁿ	$S_0^n \cap M$	$\{\text{dual}(h_n), 1\}$
S ₀₁	$S_0 \cap M$	$\{x \vee (y \wedge z), 1\}$
S ₀₀ ⁿ	$S_0^n \cap R_2 \cap M$	$\{x \vee (y \wedge z), \text{dual}(h_n)\}$
S ₀₀	$S_0 \cap R_2 \cap M$	$\{x \vee (y \wedge z)\}$
S ₁₂ ⁿ	$S_1^n \cap R_2$	$\{x \wedge (y \vee \neg z), h_n\}$
S ₁₂	$S_1 \cap R_2$	$\{x \wedge (y \vee \neg z)\}$
S ₁₁ ⁿ	$S_1^n \cap M$	$\{h_n, 0\}$
S ₁₁	$S_1 \cap M$	$\{x \wedge (y \vee z), 0\}$
S ₁₀ ⁿ	$S_1^n \cap R_2 \cap M$	$\{x \wedge (y \vee z), h_n\}$
S ₁₀	$S_1 \cap R_2 \cap M$	$\{x \wedge (y \vee z)\}$
D	$\{f \mid f \text{ is self-dual}\}$	$\{(x \wedge \neg y) \vee (x \wedge \neg z) \vee (\neg y \wedge \neg z)\}$
D ₁	$D \cap R_2$	$\{(x \wedge y) \vee (x \wedge \neg z) \vee (y \wedge \neg z)\}$
D ₂	$D \cap M$	$\{h_2\}$
L	$\{f \mid f \text{ is affine}\}$	$\{x \oplus y, 1\}$
L ₀	$L \cap R_0$	$\{x \oplus y\}$
L ₁	$L \cap R_1$	$\{x \oplus y \oplus 1\}$
L ₂	$L \cap R_2$	$\{x \oplus y \oplus z\}$
L ₃	$L \cap D$	$\{x \oplus y \oplus z \oplus 1\}$
V	$\{f \mid f \text{ is a disjunction or constant}\}$	$\{x \vee y, 0, 1\}$
V ₀	$V \cap R_0$	$\{x \vee y, 0\}$
V ₁	$V \cap R_1$	$\{x \vee y, 1\}$
V ₂	$V \cap R_2$	$\{x \vee y\}$
E	$\{f \mid f \text{ is a conjunction or constant}\}$	$\{x \wedge y, 0, 1\}$
E ₀	$E \cap R_0$	$\{x \wedge y, 0\}$
E ₁	$E \cap R_1$	$\{x \wedge y, 1\}$
E ₂	$E \cap R_2$	$\{x \wedge y\}$
N	$\{f \mid f \text{ depends on at most one variable}\}$	$\{\neg x, 0, 1\}$
N ₂	$N \cap R_2$	$\{\neg x\}$
I	$\{f \mid f \text{ is a projection or a constant}\}$	$\{\text{id}, 0, 1\}$
I ₀	$I \cap R_0$	$\{\text{id}, 0\}$
I ₁	$I \cap R_1$	$\{\text{id}, 1\}$
I ₂	$I \cap R_2$	$\{\text{id}\}$

Definition 3. Let C be a clone over a finite domain D . The interval of C , $\mathcal{I}(C)$, is the set $\mathcal{I}(C) = \{\text{pPol}(\Delta) \mid \Delta \subseteq \text{Rel}_D, \text{Pol}(\Delta) = C\}$.

Hence, the interval $\mathcal{I}(C)$ of a clone C is simply the set of all strong partial clones where the total component equals C . Even though $\mathcal{I}(C)$ can still be of uncountably infinite cardinality [29], it is known that there always exists a largest and smallest element [28]. A constraint language Γ_w such that $\text{pPol}(\Gamma_w) \in \mathcal{I}(C)$ satisfying $\text{pPol}(\Gamma_w) \supseteq \text{pPol}(\Delta)$ for any $\text{pPol}(\Delta) \in \mathcal{I}(C)$ is called a *weak base* of $\text{Inv}(C)$ [28]. Note that if a co-clone is of finite order then a weak base can always be given as a single relation. As one can verify with the Galois connection, from the functional point of view, a weak base Γ_w results in the largest element $\text{pPol}(\Gamma_w)$ in $\mathcal{I}(C)$, but from the relational point of view, the weak base has the least expressive power with respect to q.f.p.p. definability. Hence, we have the following theorem.

Theorem 4 ([28]). Let Γ_w be a weak base of a co-clone $\text{Inv}(C)$. Then $\Gamma_w \subseteq \langle \Gamma \rangle_{\exists}$ for any base Γ of $\text{Inv}(C)$.

Dually, a constraint language Γ_p such that $\Gamma_p \in \mathcal{I}(C)$ and satisfying $\text{pPol}(\Gamma_p) \subseteq \text{pPol}(\Gamma)$ for any $\text{pPol}(\Gamma) \in \mathcal{I}(C)$ is called a *plain base* of $\text{Inv}(C)$ [9]. Again, using the Galois connection, we see that $\text{pPol}(\Gamma_p)$ is the smallest element in $\mathcal{I}(C)$ but that Γ_p is the most expressive language with respect to q.f.p.p. definability. Hence, we have the following theorem.

Theorem 5. Let Γ_p be a plain base of a co-clone $\text{Inv}(C)$. Then $R \in \langle \Gamma_p \rangle_{\exists}$ for any $R \in \text{Inv}(C)$.

It is not difficult to verify that $\text{Inv}(C)$ is a plain base of itself since $[C]_s = \text{pPol}(\text{Inv}(C))$ and $\langle \text{Inv}(C) \rangle_{\exists} = \text{Inv}(C)$. However, Creignou et al. [9] gave a much more systematic and highly regular description of plain bases for Boolean co-clones. These bases can be found in Table 2, and we remark that every such plain base Γ_p in addition fulfils the condition that $R \in \langle \Gamma^{(n)} \rangle_{\exists}$ for each n -ary $R \in \langle \Gamma_p \rangle$. Hence, we have the following theorem.

Theorem 6. Let Γ_p be the plain base from Table 2 for some Boolean co-clone $\text{Inv}(C)$. Then $R \in \langle \Gamma_p^{(n)} \rangle_{\exists}$ for any n -ary $R \in \text{Inv}(C)$.

For weak bases, Schnoor and Schnoor [28] gave a systematic procedure for obtaining weak bases, which was later refined in Lagerkvist [20] in order to get a complete list of weak bases for all Boolean co-clones of finite order. These relations can be found in Table 2. We give a short description of some of the involved relations: for a full description, see Lagerkvist [20, 21] and Creignou et al. [9]. We write F and T for the constant relations $\{(0)\}$ and $\{(1)\}$; $\text{OR}^n(x_1, \dots, x_n)$ for the disjunction $x_1 \vee \dots \vee x_n$, $\text{NAND}^n(x_1, \dots, x_n)$ for the relation $\overline{x_1} \vee \dots \vee \overline{x_n}$; and define the $(n + m)$ -ary relation $\text{Compl}_{m,n}$ as

$$\text{Compl}_{m,n}(x_1, \dots, x_{m+n}) \equiv (x_1 \vee \dots \vee x_n \vee \overline{x_{n+1}} \vee \dots \vee \overline{x_{n+m}}) \wedge (\overline{x_1} \vee \dots \vee \overline{x_n} \vee x_{n+1} \vee \dots \vee x_{n+m}).$$

2.3 The Constraint Satisfaction Problem

The *constraint satisfaction problem* over a constraint language Γ ($\text{CSP}(\Gamma)$) is the following computational decision problem.

INSTANCE: A set V of variables and a set C of constraint applications $R(v_1, \dots, v_k)$ where $R \in \Gamma$, $k = \text{ar}(R)$, and $v_1, \dots, v_k \in V$.

QUESTION: Is there a function $f : V \rightarrow D$ such that $(f(v_1), \dots, f(v_k)) \in R$ for each $R(v_1, \dots, v_k)$ in C ?

The $\text{CSP}(\Gamma)$ problem is in general NP-complete and can be used to model many classical NP-complete problems such as the k -colorability problem and the k -clique problem [15]. Jeavons et al. proved that the complexity of $\text{CSP}(\Gamma)$, up to polynomial-time reductions, is determined by $\text{Pol}(\Gamma)$ [17]. With this result Schaefer's dichotomy theorem for the Boolean satisfiability problem [27] can be formulated in a particularly simple way: for Boolean constraint languages Γ , $\text{CSP}(\Gamma)$ is NP-complete if and only if $\text{Pol}(\Gamma) \subseteq [\neg x]$. Consulting Table 1 we see that this furthermore holds if and only if $\text{Pol}(\Gamma) \in \{\mathbf{I}_2, \mathbf{N}_2\}$.

Table 2. Weak and plain bases of all Boolean co-clones of finite order.

Co-clone	Weak base	Plain base
Inv(BF)	$\text{Eq}(x_1, x_2)$	$\{\text{Eq}(x_1, x_2)\}$
Inv(R ₀)	$F(c_0)$	$\{F(c_0)\}$
Inv(R ₁)	$T(c_1)$	$\{T(c_1)\}$
Inv(R ₂)	$F(c_0) \wedge T(c_1)$	$\{F(c_0), T(c_1)\}$
Inv(M)	$(x_1 \rightarrow x_2)$	$\{(x_1 \rightarrow x_2)\}$
Inv(M ₀)	$(x_1 \rightarrow x_2) \wedge F(c_0)$	$\{(x_1 \rightarrow x_2), F(c_0)\}$
Inv(M ₁)	$(x_1 \rightarrow x_2) \wedge T(c_1)$	$\{(x_1 \rightarrow x_2), T(c_1)\}$
Inv(M ₂)	$(x_1 \rightarrow x_2) \wedge F(c_0) \wedge T(c_1)$	$\{(x_1 \rightarrow x_2), F(c_0), T(c_1)\}$
Inv(S ₀ ⁿ), $n \geq 2$	$\text{OR}^n(x_1, \dots, x_n) \wedge T(c_1)$	$\{\text{OR}^n(x_1, \dots, x_n)\}$
Inv(S ₀₂ ⁿ), $n \geq 2$	$\text{OR}^n(x_1, \dots, x_n) \wedge F(c_0) \wedge T(c_1)$	$\{\text{OR}^n(x_1, \dots, x_n), F(c_0)\}$
Inv(S ₀₁ ⁿ), $n \geq 2$	$\text{OR}^n(x_1, \dots, x_n) \wedge (x \rightarrow x_1 \cdots x_n) \wedge T(c_1)$	$\{\text{OR}^n(x_1, \dots, x_n), (x_1 \rightarrow x_2)\}$
Inv(S ₀₀ ⁿ), $n \geq 2$	$\text{OR}^n(x_1, \dots, x_n) \wedge (x \rightarrow x_1 \cdots x_n) \wedge F(c_0) \wedge T(c_1)$	$\{\text{OR}^n(x_1, \dots, x_n), (x_1 \rightarrow x_2), F(c_0)\}$
Inv(S ₁ ⁿ), $n \geq 2$	$\text{NAND}^n(x_1, \dots, x_n) \wedge F(c_0)$	$\{\text{NAND}^n(x_1, \dots, x_n)\}$
Inv(S ₁₂ ⁿ), $n \geq 2$	$\text{NAND}^n(x_1, \dots, x_n) \wedge F(c_0) \wedge T(c_1)$	$\{\text{NAND}^n(x_1, \dots, x_n), T(c_1)\}$
Inv(S ₁₁ ⁿ), $n \geq 2$	$\text{NAND}^n(x_1, \dots, x_n) \wedge \neg(x \rightarrow x_1 \cdots x_n) \wedge F(c_0)$	$\{\text{NAND}^n(x_1, \dots, x_n), (x_1 \rightarrow x_2)\}$
Inv(S ₁₀ ⁿ), $n \geq 2$	$\text{NAND}^n(x_1, \dots, x_n) \wedge \neg(x \rightarrow x_1 \cdots x_n) \wedge F(c_0) \wedge T(c_1)$	$\{\text{NAND}^n(x_1, \dots, x_n), (x_1 \rightarrow x_2), T(c_1)\}$
Inv(D)	$(x_1 \oplus x_2 = 1)$	$\{(x_1 \oplus x_2 = 1)\}$
Inv(D ₁)	$(x_1 \oplus x_2 = 1) \wedge F(c_0) \wedge T(c_1)$	$\{(x_1 \oplus x_2 = 1)\} \cup \{F(c_0), T(c_1)\}$
Inv(D ₂)	$\text{OR}_{2 \neq}^2(x_1, x_2, x_3, x_4) \wedge F(c_0) \wedge T(c_1)$	$\{F(c_0), T(c_1), (x_1 \vee x_2), (\neg x_1 \vee x_2), (\neg x_1 \vee \neg x_2)\}$
Inv(L)	$\text{EVEN}^4(x_1, x_2, x_3, x_4)$	$\{(x_1 \oplus \dots \oplus x_k = 0) \mid k \text{ even}\}$
Inv(L ₀)	$\text{EVEN}^3(x_1, x_2, x_3) \wedge F(c_0)$	$\{(x_1 \oplus \dots \oplus x_k = 0) \mid k \in \mathbb{N}\}$
Inv(L ₁)	$\text{ODD}^3(x_1, x_2, x_3) \wedge T(c_1)$	$\{(x_1 \oplus \dots \oplus x_k = c) \mid k \in \mathbb{N}, c = k \bmod 2\}$
Inv(L ₂)	$\text{EVEN}_{3 \neq}^3(x_1, \dots, x_6) \wedge F(c_0) \wedge T(c_1)$	$\{(x_1 \oplus \dots \oplus x_k = c) \mid k \in \mathbb{N}, c \in \{0, 1\}\}$
Inv(L ₃)	$\text{EVEN}_{4 \neq}^4(x_1, \dots, x_8)$	$\{(x_1 \oplus \dots \oplus x_k = c) \mid k \text{ even}, c \in \{0, 1\}\}$
Inv(V)	$(\bar{x}_1 \leftrightarrow \bar{x}_2 \bar{x}_3) \wedge (\bar{x}_2 \vee \bar{x}_3 \rightarrow \bar{x}_4)$	$\{(x_1 \vee \dots \vee x_k \vee \neg x) \mid k \geq 1\}$
Inv(V ₀)	$(\bar{x}_1 \leftrightarrow \bar{x}_2 \bar{x}_3) \wedge F(c_0)$	$\{(x_1 \vee \dots \vee x_k \vee \neg x) \mid k \in \mathbb{N}\}$
Inv(V ₁)	$(\bar{x}_1 \leftrightarrow \bar{x}_2 \bar{x}_3) \wedge (\bar{x}_2 \vee \bar{x}_3 \rightarrow \bar{x}_4) \wedge T(c_1)$	$\{\text{OR}^n(x_1, \dots, x_n) \mid n \in \mathbb{N}\} \cup \{(x_1 \vee \dots \vee x_k \vee \neg x) \mid k \geq 1\}$
Inv(V ₂)	$(\bar{x}_1 \leftrightarrow \bar{x}_2 \bar{x}_3) \wedge F(c_0) \wedge T(c_1)$	$\{\text{OR}^n(x_1, \dots, x_n) \mid n \in \mathbb{N}\} \cup \{(x_1 \vee \dots \vee x_k \vee \neg x) \mid k \in \mathbb{N}\}$
Inv(E)	$(x_1 \leftrightarrow x_2 x_3) \wedge (x_2 \vee x_3 \rightarrow x_4)$	$\{(\neg x_1 \vee \dots \vee \neg x_k \vee x) \mid k \geq 1\}$
Inv(E ₀)	$(x_1 \leftrightarrow x_2 x_3) \wedge (x_2 \vee x_3 \rightarrow x_4) \wedge F(c_0)$	$\{(\neg x_1 \vee \dots \vee \neg x_k \vee x) \mid k \in \mathbb{N}\}$
Inv(E ₁)	$(x_1 \leftrightarrow x_2 x_3) \wedge T(c_1)$	$\{\text{NAND}^n(x_1, \dots, x_n) \mid n \in \mathbb{N}\} \cup \{(\neg x_1 \vee \dots \vee \neg x_k \vee x) \mid k \geq 1\}$
Inv(E ₂)	$(x_1 \leftrightarrow x_2 x_3) \wedge F(c_0) \wedge T(c_1)$	$\{\text{NAND}^n(x_1, \dots, x_n) \mid n \in \mathbb{N}\} \cup \{(\neg x_1 \vee \dots \vee \neg x_k \vee x) \mid k \in \mathbb{N}\}$
Inv(N)	$\text{EVEN}^4(x_1, x_2, x_3, x_4) \wedge x_1 x_4 \leftrightarrow x_2 x_3$	$\{\text{Compl}_{m,n} \mid m, n \geq 1\}$
Inv(N ₂)	$\text{EVEN}_{4 \neq}^4(x_1, \dots, x_8) \wedge x_1 x_4 \leftrightarrow x_2 x_3$	$\{\text{Compl}_{m,n} \mid m, n \in \mathbb{N}\}$
Inv(I)	$(x_1 \leftrightarrow x_2 x_3) \wedge (\bar{x}_4 \leftrightarrow \bar{x}_2 \bar{x}_3)$	$\{(x_1 \vee \dots \vee x_m \vee \neg y_1 \vee \dots \vee \neg y_n) \mid m, n \geq 1\}$
Inv(I ₀)	$(\bar{x}_1 \vee \bar{x}_2) \wedge (\bar{x}_1 \bar{x}_2 \leftrightarrow \bar{x}_3) \wedge F(c_0)$	$\{(x_1 \vee \dots \vee x_m \vee \neg y_1 \vee \dots \vee \neg y_n) \mid m \in \mathbb{N}, n \geq 1\}$
Inv(I ₁)	$(x_1 \vee x_2) \wedge (x_1 x_2 \leftrightarrow x_3) \wedge T(c_1)$	$\{(x_1 \vee \dots \vee x_m \vee \neg y_1 \vee \dots \vee \neg y_n) \mid m \geq 1, n \in \mathbb{N}\}$
BR	$R_{3 \neq}^{1/3}(x_1, \dots, x_6) \wedge F(c_0) \wedge T(c_1)$	$\{(x_1 \vee \dots \vee x_m \vee \neg y_1 \vee \dots \vee \neg y_n) \mid m, n \in \mathbb{N}\}$

Similarly, it has been shown that $\text{pPol}(\Gamma)$ determines the complexity of $\text{CSP}(\Gamma)$ up to $O(c^n)$ time complexity [18], where n denotes the number of variables in a given $\text{CSP}(\Gamma)$ instance. Hence, a better understanding of the partial polymorphisms of a constraint language Γ could lead to a better understanding of the worst-case time complexity of $\text{CSP}(\Gamma)$. However, as we will see in Section 6, obtaining simple characterizations of strong partial clones $\text{pPol}(\Gamma)$ is likely very difficult for many natural choices of constraint languages such that $\text{CSP}(\Gamma)$ is NP-complete.

3 Polynomially Closed Co-Clones of Finite Order

In this section we formally introduce the notion of a polynomially closed co-clone. Intuitively, the notion means that for any base of the co-clone, a polynomial amount of variables is sufficient to p.p. define any relation in the co-clone.

Definition 7. *Let $\text{Inv}(\mathbf{C})$ be a co-clone over a finite domain. We say that $\text{Inv}(\mathbf{C})$ is polynomially closed if there exists a polynomial p such that for all bases Γ of $\text{Inv}(\mathbf{C})$ and all n -ary $R \in \text{Inv}(\mathbf{C})$ it holds that R can be p.p. defined in Γ with at most $p(n)$ existentially quantified variables.*

Observe that $\text{Inv}(\mathbf{C})$ in this definition is allowed to be of infinite order. In this section, however, we restrict our focus to co-clones of finite order, while we in Section 4 investigate co-clones of infinite order. If a co-clone is not polynomially closed then we say that it is *superpolynomially closed*. We remark that if a co-clone is polynomially closed then the corresponding CSP problem is globally tractable [14]. As we will see later in Section 5, the opposite is not necessarily true for superpolynomially closed co-clones, however, since there exists superpolynomially closed co-clones where the corresponding CSP problem is trivially tractable.

We now turn to the problem of determining whether a co-clone is polynomially closed. First observe that to prove that a co-clone is polynomially closed it is sufficient to prove that there exists some polynomial p such that the weak base of the co-clone can p.p. define any n -ary relation with $p(n)$ variables.

Say that a plain base Γ_p of a co-clone $\text{Inv}(\mathbf{C})$ is a *polynomial base* if there exists a polynomial p , such that every n -ary $R \in \text{Inv}(\mathbf{C})$ has a q.f.p.p. definition over $\Gamma_p^{(n)}$, with at most $p(n)$ constraints. Polynomial bases and polynomially closed co-clones are related by the following lemma, which states that a polynomial base for a co-clone implies polynomial closure, under some additional conditions.

Lemma 8. *Let $\text{Inv}(\mathbf{C})$ be a co-clone with a weak base R_w . If there exists a polynomial, plain base Γ_p of $\text{Inv}(\mathbf{C})$, and a polynomial p such that, for each $n \geq 1$, R_w can p.p. define every relation in $\Gamma_p^{(n)}$ with at most $p(n)$ existentially quantified variables, then $\text{Inv}(\mathbf{C})$ is polynomially closed.*

Proof. Let $R \in \text{Inv}(\mathbf{C})$ be an n -ary relation. By Theorem 5 and the assumption that Γ_p is a polynomial, plain base it follows that $\Gamma_p^{(n)}$ can q.f.p.p. define R using at most $g(n)$ constraints for some polynomial g . Let ϕ denote the q.f.p.p. definition of R in $\Gamma_p^{(n)}$. For every constraint C_i in ϕ we then replace C_i with its p.p. definition in $\{R_w, \text{Eq}\}$. Let the resulting formula be ϕ' . Since ϕ had $g(n)$ constraints and each constraint in ϕ' introduced at most $p(n)$ new existentially quantified variables, the total number of variables in ϕ' is $g(n) \cdot p(n)$, clearly polynomial with respect to n . Hence, $\text{Inv}(\mathbf{C})$ is polynomially closed. \square

It is not difficult to see that this condition is satisfied whenever a co-clone has a finite plain base.

Lemma 9. *If $\text{Inv}(\mathbf{C})$ has a finite plain base then $\text{Inv}(\mathbf{C})$ is polynomially closed.*

Proof. Assume that $\text{Inv}(\mathbf{C})$ has a plain base Γ_p of finite cardinality and let R_w denote a weak base of $\text{Inv}(\mathbf{C})$. Observe that Γ_p is trivially a polynomial base. Since Γ_p is finite there exists a polynomial p such that R_w can p.p. define $\Gamma_p^{(n)}$ for every $n \geq 1$ with $p(n)$ variables. To see this, simply take the number of existentially quantified variables of the relation requiring the largest number of quantified variables in the p.p. definition in Γ_p . Such a relation must exist since Γ_p is finite. The result then follows from Lemma 8. \square

For the Boolean domain one can simply consult Table 2 to see which co-clones have finite plain bases. It is however possible to give a general characterization of the co-clones admitting finite plain bases. Let D be an arbitrary finite domain. A k -ary near-unanimity (NU) operation on D for $k \geq 3$ is an operation $f : D^k \rightarrow D$ such that $f(x, y, \dots, y) = f(y, x, \dots, y) = \dots = f(y, y, \dots, x) = y$ for all $x, y \in D$ (we may note that this is a special case of a k -edge operation, used in Section 5).

Theorem 10. *Let $\text{Inv}(\mathbf{C})$ be a co-clone over a finite domain D such that $\text{Pol}(\text{Inv}(\mathbf{C}))$ contains a k -ary NU operation for some $k \geq 3$. Then $\text{Inv}(\mathbf{C})$ is polynomially closed.*

Proof. We recall some definitions from Jeavons et al. [16]. Let $R \subseteq D^n$ be a relation and $I = \{i_1, \dots, i_d\}$ a set of indices, $1 \leq i_1 < \dots < i_d \leq n$. The *projection of R onto I* is the relation $\pi_I(R) = \{(t[i_1], \dots, t[i_d]) \mid t \in R\}$. A relation $R \subseteq D^n$ over D is *r -decomposable* if it is equivalent to the conjunction of all its projections of arity at most r , i.e., for every $t \in (D^n \setminus R)$ there is a set $I = \{i_1, \dots, i_d\}$ as above such that $(t[i_1], \dots, t[i_d]) \notin \pi_I(R)$. It is known that any relation preserved by a k -ary, $k \geq 3$, NU operation is $(k-1)$ -decomposable [16].

Now let $R \in \text{Inv}(\mathbf{C})$, of arity n . Observe that $\pi_I(R)$ can be defined using existential quantification over R , hence $\pi_I(R) \in \langle \{R\} \rangle$ for every set of indices I . Also note that $\bigwedge_{I:|I|<k} \pi_I(R)$ is a q.f.p.p. definition of R . Hence the set of all relations $R' \in \text{Inv}(\mathbf{C})$ of arity at most $k-1$ is a plain base of $\text{Inv}(\mathbf{C})$. Clearly, this is a finite set (since $|D|$ is finite). Thus $\text{Inv}(\mathbf{C})$ is polynomially closed by Lemma 9.

Observe, however, that Lemma 9 or Theorem 10 are not applicable for $\text{Inv}(\mathbf{L})$, $\text{Inv}(\mathbf{L}_0)$, $\text{Inv}(\mathbf{L}_1)$, $\text{Inv}(\mathbf{L}_3)$ and $\text{Inv}(\mathbf{L}_2)$ since they do not admit finite plain bases. Fortunately, it is rather straightforward to prove that these co-clones admit polynomial bases, since the included relations can be viewed as linear equations over the field $\text{GF}(2)$.

Lemma 11. *$\text{Inv}(\mathbf{L}), \text{Inv}(\mathbf{L}_0), \text{Inv}(\mathbf{L}_1), \text{Inv}(\mathbf{L}_3)$ and $\text{Inv}(\mathbf{L}_2)$ have polynomial, plain bases.*

Proof. We only consider $\text{Inv}(\mathbf{L}_2)$ since the other cases follow through similar arguments. Every n -ary relation $R \in \text{Inv}(\mathbf{L}_2)$ can according to Theorem 6 be expressed by a $\Gamma_p^{(n)}$ formula ϕ with m constraints, where Γ_p is the plain base of $\text{Inv}(\mathbf{L}_2)$ in Table 2. Thus every constraint C_i in ϕ is of the form $(x_{i_1} \oplus \dots \oplus x_{i_n}) = c_i$, where $c_i \in \{0, 1\}$. Create an $m \times (n+1)$ -matrix M such that each entry $r_{i,j}$, $1 \leq j \leq n$, is equal to 1 if the variable x_j is included in the constraint C_i , and 0 otherwise. The entry $r_{i,n+1}$ is equal to the constant c_i in C_i . Then it is not hard to verify that if the row r_{i+1} is linearly dependent on r_1, \dots, r_i then C_1, \dots, C_i entails C_{i+1} in any satisfying assignment. Hence we only need to keep the rows that are linearly independent, which gives the bound $\min(n+1, m)$ on the number of constraints. \square

Lemma 12. *$\text{Inv}(\mathbf{L}), \text{Inv}(\mathbf{L}_0), \text{Inv}(\mathbf{L}_1), \text{Inv}(\mathbf{L}_3)$ and $\text{Inv}(\mathbf{L}_2)$ are polynomially closed.*

Proof. We only present the proof of $\text{Inv}(\mathbf{L}_2)$ since the other co-clones follow through entirely analogous arguments. Let Γ_p and R_w be the plain and weak base of $\text{Inv}(\mathbf{L}_2)$ from Table 2, respectively. Since $\text{Inv}(\mathbf{L}_2)$ has a polynomial base by Lemma 11 all we need to prove is that there exists a polynomial p such that R_w can p.p. define $\Gamma_p^{(n)}$ using at most $p(n)$ existentially quantified variables. We first and most crucially show that $\Gamma_p^{(n)}$ can p.p. define $\Gamma_p^{(n+1)}$ with only one extra variable, for every $n \geq 3$, with the definition $(x_1 \oplus \dots \oplus x_{n+1} = c) \equiv \exists x.(x_1 \oplus \dots \oplus x_{n-1} \oplus x = c) \wedge (x_n \oplus x_{n+1} \oplus x = 0)$. In addition to one quantified variable this requires one extra $\Gamma_p^{(3)}$ -constraint. Hence if $3 \leq n \leq n'$ then $\Gamma_p^{(n)}$ can p.p. define every relation in $\Gamma_p^{(n')}$ with $O(n' - n)$ variables and $n' - n$ additional $\Gamma_p^{(3)}$ -constraints. By this it first follows that $\Gamma_p^{(3)}$ can p.p. define any relation in $\Gamma_p^{(n)}$ with at most $n - 3$ variables and $n - 2$ constraints. The weak base R_w can then p.p. define $\Gamma_p^{(3)}$ with a fixed number of variables since the arity of each relation is bounded, for example we have that $(x_1 \oplus x_2 \oplus x_3 = 0) \equiv \exists y_1, y_2, y_3, c_0, c_1. R_w(x_1, x_2, x_3, y_1, y_2, y_3, c_0, c_1)$ and $(x_1 \oplus x_2 \oplus x_3 = 1) \equiv \exists y_1, y_2, y_3, c_0, c_1. R_w(y_1, y_2, y_3, x_1, x_2, x_3, c_0, c_1)$. Put together this implies that R_w can p.p. define any $\Gamma_p^{(n)}$ with $O(n)$ existentially quantified variables, and by Lemma 8 that $\text{Inv}(\mathbf{L}_2)$ is polynomially closed. \square

Combining all results so far in this section, we obtain the following characterization of the polynomially closed Boolean co-clones of finite order.

Theorem 13. *If $\text{Inv}(C) \subseteq \text{Inv}(X)$ for some $X \in \{\text{L}_2, \text{D}_2\} \cup \{\text{S}_{00}^n, \text{S}_{10}^n \mid n \geq 2\}$ then $\text{Inv}(C)$ is polynomially closed.*

4 Polynomially closed co-clones of infinite order

So far we have only been concerned with polynomially closed co-clones of finite order. For co-clones of infinite order, we cannot use any of the machinery introduced in Section 3. In particular, Lemma 9 breaks down since there by definition cannot exist a finite plain base of a co-clone of infinite order. In this section we give a general result to obtain weak bases of co-clones of infinite order, and leverage this result to show that the eight Boolean co-clones of infinite order in Figure 1 are polynomially closed.

Theorem 14. *Let $\text{Inv}(C)$ be a co-clone of infinite order over a finite domain and let $\text{Inv}(C_1), \text{Inv}(C_2), \dots$ be an infinite chain of co-clones of finite order such that $C = \bigcup_{i=1}^{\infty} \text{Inv}(C_i)$. Let R_{C_i} denote the weak base of each $\text{Inv}(C_i)$. Assume that $R_{C_i} \in \langle R_{C_{i+1}} \rangle_{\neq}$ for each $i \geq 1$. Then the weak base of $\text{Inv}(C)$ is the language $\Gamma_C = \{R_{C_i} \mid i \geq 1\}$.*

Proof. First observe that each $\text{Inv}(C_i)$ does indeed have a finite weak base since by assumption they are of finite order. To prove that Γ_C is a weak base of $\text{Inv}(C)$ we must prove that it is a base of $\text{Inv}(C)$ and that $\Gamma_C \subseteq \langle \Gamma \rangle_{\neq}$ for each base Γ of $\text{Inv}(C)$. It is easy to see that Γ_C is a base of $\text{Inv}(C)$ since $\text{Inv}(C_i) = \langle \{R_{C_i}\} \rangle \subseteq \langle \Gamma_C \rangle$ for every $i \geq 1$ for some $R_{C_i} \in \Gamma_C$. Let Γ be a constraint language such that $\langle \Gamma \rangle = \text{Inv}(C)$. Observe that Γ must be infinite, and that there for every $R \in \Gamma$ exists some m such that $R \in \text{Inv}(C_m)$, since $\langle \Gamma \rangle = \text{Inv}(C) = \bigcup_{i=1}^{\infty} \text{Inv}(C_i)$.

We must prove that $\Gamma_C \subseteq \langle \Gamma \rangle_{\neq}$. Let $R \in \Gamma_C$ be an n -ary relation. Then there exists an m such that R is the weak base of $\text{Inv}(C_m)$. We prove that there exists $\Delta \subseteq \Gamma$ such that $\langle \Delta \rangle = \text{Inv}(C_{m'})$ for some $m' \geq m$, since this implies that $R \in \langle \Delta \rangle_{\neq} \subseteq \langle \Gamma \rangle_{\neq}$, by the original assumption. Assume for contradiction that no such set Δ exists. But this implies that there exists some $k < m$ such that $\langle \Gamma \rangle = \text{Inv}(C_k)$, which is clearly impossible since C_k is of finite order. Hence, there exists $\Delta \subseteq \Gamma$ such that $R \in \langle \Delta \rangle_{\neq}$. Since R was chosen arbitrarily, this in turn implies that $\Gamma_C \subseteq \langle \Gamma \rangle_{\neq}$, and that Γ_C is a weak base of $\text{Inv}(C)$. \square

Table 3. Weak bases of all Boolean co-clones of infinite order.

Co-clone	Weak base
$\text{Inv}(\text{S}_0)$	$\{\text{OR}(x_1, \dots, x_n) \wedge \text{T}(c_1) \mid n \geq 2\}$
$\text{Inv}(\text{S}_{02})$	$\{\text{OR}(x_1, \dots, x_n) \wedge \text{F}(c_0) \wedge \text{T}(c_1) \mid n \geq 2\}$
$\text{Inv}(\text{S}_{01})$	$\{\text{OR}(x_1, \dots, x_n) \wedge (x \rightarrow x_1 \cdots x_n) \wedge \text{T}(c_1) \mid n \geq 2\}$
$\text{Inv}(\text{S}_{00})$	$\{\text{OR}(x_1, \dots, x_n) \wedge (x \rightarrow x_1 \cdots x_n) \wedge \text{F}(c_0) \wedge \text{T}(c_1) \mid n \geq 2\}$
$\text{Inv}(\text{S}_1)$	$\{\text{NAND}(x_1, \dots, x_n) \wedge \text{F}(c_0) \mid n \geq 2\}$
$\text{Inv}(\text{S}_{12})$	$\{\text{NAND}(x_1, \dots, x_n) \wedge \text{F}(c_0) \wedge \text{T}(c_1) \mid n \geq 2\}$
$\text{Inv}(\text{S}_{11})$	$\{\text{NAND}(x_1, \dots, x_n) \wedge \neg(x \rightarrow x_1 \cdots x_n) \wedge \text{T}(c_1) \mid n \geq 2\}$
$\text{Inv}(\text{S}_{10})$	$\{\text{NAND}(x_1, \dots, x_n) \wedge \neg(x \rightarrow x_1 \cdots x_n) \wedge \text{F}(c_0) \wedge \text{T}(c_1) \mid n \geq 2\}$

We remark that since $R_{C_i} \in \langle R_{C_{i+1}} \rangle_{\neq}$ for every $i \geq 1$ we can in fact remove any finite number of relations from the weak base Γ_C and still obtain a weak base of $\text{Inv}(C)$. According to Theorem 14 all that is needed to obtain weak bases for the eight co-clones of infinite order in the Boolean co-clone lattice, is to show that the condition $R_{C_i} \in \langle R_{C_{i+1}} \rangle_{\neq}$ is satisfied for every $C_i \in \{\text{S}_0^i, \text{S}_{02}^i, \text{S}_{01}^i, \text{S}_{00}^i, \text{S}_1^i, \text{S}_{12}^i, \text{S}_{11}^i, \text{S}_{10}^i\}$. We only consider the case $\text{Inv}(\text{S}_{00}^n)$ since the remaining proofs are entirely analogous. Hence, we need to show that the weak base of $\text{Inv}(\text{S}_{00}^n)$ can q.f.p.p. define the weak base of $\text{Inv}(\text{S}_{00}^{n-1})$ for each $n \geq 3$. For $n \geq 2$ let $R_{\text{S}_{00}}^n(x_1, \dots, x_n, x, c_0, c_1) \equiv \text{OR}(x_1, \dots, x_n) \wedge (x \rightarrow x_1 \cdots x_n) \wedge \text{F}(c_0) \wedge \text{T}(c_1)$. Then we can q.f.p.p. define $R_{\text{S}_{00}}^n$ with $R_{\text{S}_{00}}^{n+1}$ using the definition

$$R_{\text{S}_{00}}^n(x_1, \dots, x_n, x, c_0, c_1) \equiv R_{\text{S}_{00}}^{n+1}(x_1, \dots, x_n, x_n, x, c_0, c_1).$$

Hence, we obtain the following corollary, summarized in Table 3.

Corollary 15. *The following statements hold.*

- $\{\text{OR}(x_1, \dots, x_n) \wedge \text{T}(c_1) \mid n \geq 2\}$ is a weak base of $\text{Inv}(\mathbf{S}_0)$,
- $\{\text{OR}(x_1, \dots, x_n) \wedge \text{F}(c_0) \wedge \text{T}(c_1) \mid n \geq 2\}$ is a weak base of $\text{Inv}(\mathbf{S}_{02})$,
- $\{\text{OR}(x_1, \dots, x_n) \wedge (x \rightarrow x_1 \cdots x_n) \wedge \text{T}(c_1) \mid n \geq 2\}$ is a weak base of $\text{Inv}(\mathbf{S}_{01})$,
- $\{\text{OR}(x_1, \dots, x_n) \wedge (x \rightarrow x_1 \cdots x_n) \wedge \text{F}(c_0) \wedge \text{T}(c_1) \mid n \geq 2\}$ is a weak base of $\text{Inv}(\mathbf{S}_{00})$,
- $\{\text{NAND}(x_1, \dots, x_n) \wedge \text{F}(c_0) \mid n \geq 2\}$ is a weak base of $\text{Inv}(\mathbf{S}_1)$,
- $\{\text{NAND}(x_1, \dots, x_n) \wedge \text{F}(c_0) \wedge \text{T}(c_1) \mid n \geq 2\}$ is a weak base of $\text{Inv}(\mathbf{S}_{12})$,
- $\{\text{NAND}(x_1, \dots, x_n) \wedge (x \rightarrow x_1 \cdots x_n) \wedge \text{T}(c_1) \mid n \geq 2\}$ is a weak base of $\text{Inv}(\mathbf{S}_{11})$,
- $\{\text{NAND}(x_1, \dots, x_n) \wedge (x \rightarrow x_1 \cdots x_n) \wedge \text{F}(c_0) \wedge \text{T}(c_1) \mid n \geq 2\}$ is a weak base of $\text{Inv}(\mathbf{S}_{10})$.

We are now in a position to prove that all Boolean co-clones of infinite order are polynomially closed. For the proof we use the fact that relations in $\text{Inv}(\mathbf{S}_0)$, $\text{Inv}(\mathbf{S}_1)$, $\text{Inv}(\mathbf{S}_{01})$ and $\text{Inv}(\mathbf{S}_{00})$ can be expressed by a particularly simple form of Boolean formula. Before we can formally state this result we need a few additional preliminaries. If $\varphi = C_1 \wedge \dots \wedge C_m$ is a Boolean formula with m clauses we say that C_i is a *prime implicate* of φ if φ does not entail any proper subclause of C_i . A formula φ is said to be *prime* if all of its clauses are prime implicates. Obviously any finite Boolean relation is representable by a prime formula. If R is an n -ary Boolean relation we can therefore prove that $R \in \langle \Gamma \rangle_{\exists}$ by showing that $R(x_1, \dots, x_n)$ can be expressed as a conjunction $\varphi_1(\mathbf{y}_1) \wedge \dots \wedge \varphi_k(\mathbf{y}_k)$, where each \mathbf{y}_i is a tuple of variables over x_1, \dots, x_n , and each φ_i is a prime formula representation of a relation in Γ . This is advantageous since relations in $\text{Inv}(\mathbf{S}_0^n)$, $\text{Inv}(\mathbf{S}_{02}^n)$, $\text{Inv}(\mathbf{S}_{01}^n)$, $\text{Inv}(\mathbf{S}_{00}^n)$, $\text{Inv}(\mathbf{S}_1^n)$, $\text{Inv}(\mathbf{S}_{12}^n)$, $\text{Inv}(\mathbf{S}_{11}^n)$ and $\text{Inv}(\mathbf{S}_{10}^n)$ are representable by prime *implicative hitting set-bounded* (IHSB) formulas [9]. We let IHSB_+^n be the set of formulas of the form $(x_1 \vee \dots \vee x_m), 1 \leq m \leq n, (\neg x_1), (\neg x_1 \vee x_2)$, and dually for IHSB_-^n .

Theorem 16. $\text{Inv}(\mathbf{S}_0)$, $\text{Inv}(\mathbf{S}_{02})$, $\text{Inv}(\mathbf{S}_{01})$, $\text{Inv}(\mathbf{S}_{00})$, $\text{Inv}(\mathbf{S}_1)$, $\text{Inv}(\mathbf{S}_{12})$, $\text{Inv}(\mathbf{S}_{11})$ and $\text{Inv}(\mathbf{S}_{10})$ are polynomially closed.

Proof. We only consider $\text{Inv}(\mathbf{S}_{00})$ since the other cases follow through similar arguments. Let $\Gamma_{\mathbf{S}_{00}} = \{R_{\mathbf{S}_{00}}^i \mid i \geq 2\}$ denote the weak base of $\text{Inv}(\mathbf{S}_{00})$ from Corollary 15. We must prove that there exists a polynomial p such that $\Gamma_{\mathbf{S}_{00}}$ can p.p. define any n -ary $R \in \text{Inv}(\mathbf{S}_{00})$ using at most $p(n)$ existentially quantified variables. Since $R \in \text{Inv}(\mathbf{S}_{00})$ it is easily seen that there exists some $n' \geq 2$ such that $R \in \text{Inv}(\mathbf{S}_{00}^{n'})$.

Hence, R can be written as a prime $\text{IHSB}_+^{n'}$ formula ϕ over x_1, \dots, x_n [9], and we need to show that it is possible to p.p. define this formula without requiring more than a polynomial number of existentially quantified variables. There are a few different cases to consider depending on the clauses of ϕ . Let c_0 and c_1 be two fresh variables distinct from x_1, \dots, x_n . First, we implement every clause in ϕ of the form $(x_{i_1} \vee \dots \vee x_{i_j})$ for some $j \leq n'$ with the constraint $R_{\mathbf{S}_{00}}^j(x_{i_1}, \dots, x_{i_j}, c_0, c_0, c_1)$. Second, we implement every clause of the form $(\neg x_i)$ as $R_{\mathbf{S}_{00}}^2(c_1, c_1, c_1, x_i, c_1)$. Third, we implement every clause of the form $(\neg x_{i_1} \vee x_{i_2})$ as $R_{\mathbf{S}_{00}}^2(x_{i_2}, c_1, x_{i_1}, c_0, c_1)$. Let ϕ' be the $\Gamma_{\mathbf{S}_{00}}$ -formula resulting from replacing every clause in ϕ in the above manner. We see that $R(x_1, \dots, x_n) \equiv \exists c_0 \exists c_1. \phi'$, and since we in total only require 2 existentially quantified variables, it follows that $\text{Inv}(\mathbf{S}_{00})$ is polynomially closed. \square

5 Superpolynomially closed co-clones

From Section 3 and Section 4 we now have straightforward, necessary conditions for verifying whether a given co-clone is polynomially closed. We now turn to the problem of determining whether a co-clone is not polynomially closed, i.e., superpolynomially closed. We show that this question is related to counting the number of n -ary relations in a co-clone — a problem that has attracted significant attention in universal algebra and conceptual learning problems [3, 14]. Before we can present this result, we for every finite domain D , introduce a particular constraint language Γ_D , which will turn out to be a plain base of Rel_D . The language Γ_D is defined as

$$\Gamma_D = \{R \mid n \geq 1, t \in D^n, R = D^n \setminus \{t\}\}.$$

In other words each n -ary relation in Γ_D contains all n -ary tuples over D except one. Observe that $\Gamma_{\{0,1\}}$ is equivalent to the plain base of BR in Table 2.

Lemma 17. *For any finite domain D the language Γ_D is a plain base of Rel_D .*

Proof. We must prove that $\langle \Gamma_D \rangle_{\exists} = \text{Rel}_D$, i.e. that we can q.f.p.p. define all relations over D . Hence, let $R \in \text{Rel}_D$ be an n -ary relation. For every $t \in D^n \setminus R$ we let $R_t \in \Gamma_D$ denote the unique relation satisfying $R_t = D^n \setminus \{t\}$. Hence, a constraint of the form $R_t(x_1, \dots, x_n)$ implies that x_1, \dots, x_n can take any value except for $t[1], \dots, t[n]$. With this observation it is then easy to see that we can define R with the q.f.p.p. definition

$$R(x_1, \dots, x_n) \equiv R_{t_1}(x_1, \dots, x_n) \wedge \dots \wedge R_{t_m}(x_1, \dots, x_n),$$

where $\{t_1, \dots, t_m\} = D^n \setminus R$. □

Also observe that $\langle \Gamma_D^{(n)} \rangle_{\exists} \subseteq \langle \Gamma_D^{(n+1)} \rangle_{\exists}$ for each $n \geq 1$. To see this, simply note that there for every n -ary relation R exists a $(n+1)$ -ary relation R' defined as $R'(x_1, \dots, x_n, x_{n+1}) \equiv R(x_1, \dots, x_n) \wedge \text{Eq}(x_{n+1}, x_{n+1})$, which is equivalent with respect to q.f.p.p. definitions. We will now prove that if a co-clone $\text{Inv}(\mathbf{C})$ contains a sufficiently large number of n -ary relations, then for every polynomial p there will exist some n -ary relation in $\text{Inv}(\mathbf{C})$ that Γ_D cannot p.p. define using only $p(n)$ existentially quantified variables. To make this counting argument more precise we, given a constraint language Γ , first let $\Gamma^{=n} = \{R \mid R \in \Gamma, \text{ar}(R) = n\}$, and then define the function s_Γ as

$$s_\Gamma(n) = \log_2(|\{R \mid R \in \Gamma, \text{ar}(R) = n\}|).$$

With this notation we see that $s_{\text{Inv}(\mathbf{C})}(n)$ denotes the exponent of the number of n -ary relations in the co-clone $\text{Inv}(\mathbf{C})$, and obtain the following lemma.

Lemma 18. *Let $\text{Inv}(\mathbf{C})$ be a co-clone of finite order over a finite domain D . If $\text{Inv}(\mathbf{C})$ is polynomially closed, then $s_{\text{Inv}(\mathbf{C})}(n) \leq g(n)$ for some polynomial g .*

Proof. Let Γ be a finite base of $\text{Inv}(\mathbf{C})$ and let R be the relation with the highest arity k in Γ . We make a few observations before the proof: first, $\langle \Gamma \rangle_{\exists} \subseteq \langle \Gamma_D^{(=k)} \rangle_{\exists}$; second, if some $R' \notin \langle \Gamma_D^{(=k)} \rangle_{\exists}$ then $R' \notin \Gamma$. This also implies that if Γ can p.p. define some n -ary relation R with $p(n)$ existentially quantified then the same is true for $\Gamma_D^{(=k)}$. By contraposition this also implies that if $\Gamma_D^{(=k)}$ cannot p.p. define some n -ary relation R with $p(n)$ variables then neither can Γ . It is not too difficult to see that the number of q.f.p.p. definitions with $\Gamma_D^{(=k)}$ over n variables is bounded by $2^{|D|^k n^k}$, since (1) $\Gamma_D^{(=k)}$ contains $|D|^k$ relations and (2) for each relation in $\Gamma_D^{(=k)}$ one can form at most n^k distinct constraints. Since $\text{Inv}(\mathbf{C})$ is polynomially closed, we are allowed to introduce at most $p(n)$ existentially quantified variables to define any n -ary relation, hence, the number of definable relations is at most $2^{|D|^k (p(n)+n)^k}$, which implies that $s_{\text{Inv}(\mathbf{C})}(n) \leq |D|^k (p(n)+n)^k$ and that there exists a polynomial g such that $s_{\text{Inv}(\mathbf{C})}(n) \leq g(n)$. □

Since the number of n -ary relations over a finite domain D is $2^{|D|^n}$ it immediately follows that Rel_D is superpolynomially closed. To handle the other cases where it is not so obvious how to count the number of n -ary relations we utilize a result from Berman et al. [3]. Before we can present their result, we need a few additional preliminaries. If Γ is a constraint language over D the algebra $\mathbf{A}_\Gamma = (D, \text{Pol}(\Gamma))$ is said to have *few subpowers* if $s_{\langle \Gamma \rangle}(n) \in O(n^k)$ for some polynomial $k \geq 1$, and to have *many subpowers* if $c^n \in O(s_{\langle \Gamma \rangle}(n))$ for some real number $c > 1$. A k -edge operation over D , $k \geq 2$, is a $(k+1)$ -ary operation f satisfying the k identities

- $f(x, x, y, y, y, \dots, y, y) = y$,
- $f(x, y, x, y, y, \dots, y, y) = y$,
- $f(y, y, y, x, y, \dots, y, y) = y$,
- $f(y, y, y, y, x, \dots, y, y) = y$,
- \vdots
- $f(y, y, y, y, y, \dots, x, y) = y$,
- $f(y, y, y, y, y, \dots, y, x) = y$.

We then have the following useful theorem from Berman et al. [3], restated in terms of our terminology of clones and co-clones

Theorem 19. [3] *Let Γ be a constraint language over a finite domain D . If $\text{Pol}(\Gamma)$ does not contain a k -edge operation for any $k \geq 2$ then (1) the algebra $(D, \text{Pol}(\Gamma))$ has many subpowers and (2) $s_{\langle \Gamma \rangle}(n) \notin O(n^l)$ for any $l \geq 0$.*

Hence, if $(D, \text{Pol}(\Gamma))$ has many subpowers, then, intuitively, $\langle \Gamma \rangle$ contains too many relations for it to be polynomially closed. Combining Lemma 18 and Theorem 19 we obtain the following classification of the superpolynomially closed co-clones.

Theorem 20. *Let $\text{Inv}(C)$ be a co-clone of finite order over a finite domain D . If C does not contain a k -edge operation for any $k \geq 2$ then $\text{Inv}(C)$ is superpolynomially closed.*

With the help of Table 1 one can verify that any Boolean co-clone of finite order above or equal to $\text{Inv}(\mathbf{V})$, $\text{Inv}(\mathbf{E})$, or $\text{Inv}(\mathbf{N})$ in Figure 1, fulfil this property.

Theorem 21. *If $\text{Inv}(C) \supseteq \text{Inv}(X)$ for some $\text{Inv}(X) \in \{\text{Inv}(\mathbf{V}), \text{Inv}(\mathbf{E}), \text{Inv}(\mathbf{N})\}$ then $\text{Inv}(C)$ is superpolynomially closed.*

Due to the close relationship between a polynomially closed co-clone and the existence of a polynomial, plain base, one might suspect that superpolynomially closed co-clones are unlikely to admit such polynomial bases. This can in fact be proven by a straightforward counting argument, using the bounds from Theorem 19 on the number of n -ary relations in these co-clones.

Theorem 22. *Let $\text{Inv}(C)$ be a superpolynomially closed co-clone over a finite domain D such that there exists a plain base Γ_p of $\text{Inv}(C)$ satisfying $|\Gamma_p^{(n)}| \leq 2^{p(n)}$ for some polynomial p . Then Γ_p is not a polynomial, plain base of $\text{Inv}(C)$.*

Proof. Assume that $\text{Inv}(C)$ has a polynomial base with respect to a polynomial c . We show the theorem with a counting argument, using the results of Section 5. First, recall from Lemma 18 that $s_{\text{Inv}(C)}(n)$ cannot be bounded by a polynomial function since $\text{Inv}(C)$ is superpolynomially closed. In other words it cannot hold that $|\{R \in \text{Inv}(C) \mid \text{ar}(R) = n\}| \leq 2^{p(n)}$ for some polynomial p .

Now observe that for each $R \in \Gamma_p^{(n)}$, there are at most n^n different possible constraints one can form with R ; thus the number of different possible constraints overall is bounded by $|\Gamma_p^{(n)}| \cdot n^n$. The number of possible formulas with at most $c(n)$ constraints is then bounded by $(|\Gamma_p^{(n)}| \cdot n^n)^{c(n)} \leq (2^{p(n)} \cdot n^n)^{c(n)} \leq 2^{q(n)}$ for a polynomial $q(n)$, which implies that $s_{\text{Inv}(C)}(n) \leq 2^{q(n)}$, contradicting the original assumption. \square

Using Table 2 we see that each Boolean plain base Γ_p contains at most polynomially many n -ary relations. Hence, we obtain the following theorem for Boolean co-clones.

Theorem 23. *If $\text{Inv}(C) \supseteq \text{Inv}(X)$ for some $\text{Inv}(X) \in \{\text{Inv}(\mathbf{V}), \text{Inv}(\mathbf{E}), \text{Inv}(\mathbf{N})\}$ then the plain base of $\text{Inv}(C)$ in Table 2 is not a polynomial, plain base.*

Thus a Boolean co-clone of finite order has a polynomial, plain base in Table 2 if and only if it is polynomially closed. In conjunction, the results of Section 3 and Section 5 therefore imply the following corollary.

Corollary 24. *Let $\langle \Gamma \rangle$ be a Boolean co-clone of finite order. Then the following statements are equivalent.*

- $\langle \Gamma \rangle$ is polynomially closed.
- $\langle \Gamma \rangle$ has a polynomial, plain base in Table 2.
- The algebra $(\{0, 1\}, \text{Pol}(\Gamma))$ has few subpowers.
- There exists a polynomial p such that the number of n -ary relations in $\langle \Gamma \rangle$ is not larger than $2^{p(n)}$.

For arbitrary finite domains our result do not form a sharp dichotomy. Combining Theorem 10 and Theorem 20, we however obtain the following corollary.

Corollary 25. *Let $\langle \Gamma \rangle$ be a co-clone of finite order over a finite domain. Then the following statements hold.*

- *If $\text{Pol}(\Gamma)$ does not contain a k -edge operation for any $k \geq 2$ then $\langle \Gamma \rangle$ is superpolynomially closed.*
- *If $\text{Pol}(\Gamma)$ contains a k -ary near-unanimity operation for some $k \geq 3$ then $\langle \Gamma \rangle$ is polynomially closed.*

For co-clones of infinite order this situation differs drastically, as evident in Section 4, since even in the Boolean domain it can be the case that a co-clone of infinite order is polynomially closed even if the corresponding algebra has many subpowers.

6 Strong Partial Clones of Finite and Infinite Order

So far we have been interested in obtaining conditions for separating polynomially closed co-clones from superpolynomially closed co-clones, and obtained a complete dichotomy theorem for the Boolean domain. Since we for polynomially closed co-clones can define all relations in the co-clone with a comparably few number of existentially quantified variables, one might conjecture that a strong partial clone $\text{pPol}(\Gamma)$ has a more complex structure if $\langle \Gamma \rangle$ is superpolynomially closed. To make this intuition a bit more precise, given a co-clone $\text{Inv}(C)$ and a base Γ of $\text{Inv}(C)$, we are interested in determining when $\text{pPol}(\Gamma)$ is of infinite order and when it is of finite order. Hence, we make the following definition (recall from Section 2.2 that $\mathcal{I}(C)$ denotes the interval of all strong partial clones where the total component equals C).

Definition 26. *Let C be a clone over a finite domain. We say that $\mathcal{I}(C)$ is finitely generated if every $\text{pPol}(\Delta) \in \mathcal{I}(C)$ is of finite order and that $\mathcal{I}(C)$ is infinitely generated if $\text{pPol}(\Delta)$ is of infinite order for every $\text{pPol}(\Delta) \in \mathcal{I}(C)$.*

A few basic observations are in place. First, determining whether a partial clone is of finite or infinite order is a problem that has attracted significant attention in the literature, see e.g. [7, 8, 13]. However, observe that the authors in this case study partial clones that are not necessarily strong, and that a partial clone P might be of infinite order even though the smallest strong partial clone containing P is of finite order. Second, if C is a clone of finite order then $\text{pPol}(\text{Inv}(C))$ is of finite order. This implies that as long as C is of finite order, $\mathcal{I}(C)$ will contain at least one strong partial clone of finite order. Hence, in general, intervals of the form $\mathcal{I}(C)$ may contain strong partial clones of both finite and infinite order. If we restrict ourself to strong partial clones $\text{pPol}(\Gamma)$ where Γ is finite, this phenomenon is not as likely to occur, however. We thus make the following definition as well.

Definition 27. *Let C be a clone over a finite domain D . The finite interval of C , $\mathcal{I}_{\text{fin}}(C)$, is the set $\mathcal{I}_{\text{fin}}(C) = \{\text{pPol}(\Delta) \mid \Delta \subseteq \text{Rel}_D, \Delta \text{ is finite, } C = \text{Pol}(\Delta)\}$.*

In Section 6.1 we prove that the existence of finitely generated intervals is related to the question whether the cardinality of the interval is finite or uncountably infinite, and give examples of polynomially closed co-clones over arbitrary finite domains resulting in finitely generated intervals. Since not much is currently known about the lattice of strong partial clones over arbitrary finite domains, these results are necessarily inconclusive, and we cannot yet hope to provide a complete classification of finitely generated intervals. In Section 6.2 we study the opposite question: given a superpolynomially closed co-clone $\text{Inv}(C)$, is $\mathcal{I}_{\text{fin}}(C)$ infinitely generated? We give a general result and prove that $\mathcal{I}_{\text{fin}}(C)$ is infinitely generated whenever C consists of essentially unary functions, i.e., if $C = \{e_1, \dots, e_k\}$ for some unary functions e_1, \dots, e_k . The results are summarized in Figure 2.

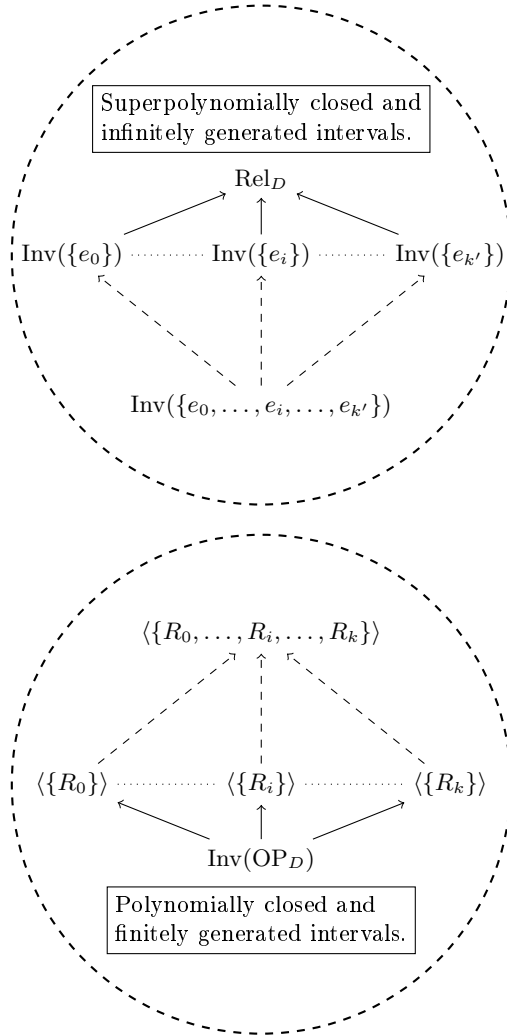


Fig. 2. An illustration of some fragments of the the structure of $\mathcal{I}(\Gamma)$ for Γ over an arbitrary finite domain $D = \{0, \dots, i, \dots, k\}$. For $a \in D$ let R_a denote the relation $\{(a)\}$. Let $e_0, \dots, e_{k'}$ be an enumeration of the unary functions over D which are not projections. A directed arrow from node A to B means that $A \subset B$. A dashed arrow from node A to B means that there exists A' such that $A \subset A' \subset B$. Some inclusions have been omitted.

6.1 Strong Partial Clones of Finite Order

We first remark that if $\mathcal{I}(C)$ is finitely generated then the cardinality of $\mathcal{I}(C)$ is at most countably infinite. Hence, we get the following proposition.

Theorem 28. *Let C be clone such that $\mathcal{I}(C)$ is of uncountably infinite cardinality. Then $\mathcal{I}(C)$ is not finitely generated.*

On the other hand, if $\mathcal{I}(C)$ is finite and C is of finite order, it is not too difficult to see that $\mathcal{I}(C)$ must be finitely generated.

Lemma 29. *Let C be a clone of finite order over D such that $\mathcal{I}(C)$ is finite. Then $\mathcal{I}(C)$ is finitely generated.*

Proof. Let F denote an arbitrary finite base of C . Then $[F]_s$ is the least element in $\mathcal{I}(C)$. Assume, for contradiction, that there exists a strong partial clone $C' \in \mathcal{I}(C)$ of infinite order. Obviously $C' \supset [F]_s$. Let $f \in C' \setminus [F]_s$. Then it is easy to see that $C' \supset [F \cup \{f\}]_s \supset [F]_s$ since by assumption C' is of infinite order and cannot be generated by $F \cup \{f\}$. This procedure can be repeated arbitrarily many times, which contradicts the assumption that $\mathcal{I}(C)$ was finite. \square

Hence, the question of whether an interval is finitely generated or not is tightly connected to whether the interval is finite. In the Boolean domain it has been proven that $\mathcal{I}(\text{Pol}(\Gamma))$ is finite if $\langle \Gamma \rangle$ is a subset of $\text{Inv}(\mathbf{M}_2)$ or $\text{Inv}(\mathbf{D}_1)$, and is of uncountably infinite cardinality otherwise [29]. Hence, we obtain the following proposition.

Proposition 30. *Let Γ be a Boolean constraint language. Then $\mathcal{I}(\text{Pol}(\Gamma))$ is finitely generated if and only if $\langle \Gamma \rangle \subseteq \text{Inv}(\mathbf{X})$ for $\mathbf{X} \in \{\mathbf{M}_2, \mathbf{D}_1\}$.*

In Schölzel [29] it is conjectured that intervals of the form $\mathcal{I}(\text{Pol}(\Gamma))$ are either finite or uncountably infinite for arbitrary finite domains. Such a dichotomy theorem would therefore also answer the question which intervals are finitely generated and which are not. We remark that such a dichotomy theorem is likely very difficult to obtain, since not much is known of the structure of the lattice of strong partial clones over arbitrary finite domains. We give an exemplary case of a simple kind of constraint language where the intervals of strong partial clones is always finite.

Given a finite domain $D = \{0, \dots, k\}$ let $R_i, i \in D$, denote the unary, constant relation $\{(i)\}$. Say that a co-clone $\text{Inv}(C)$ over D is *essentially constant* if there exists a set $\Gamma \subseteq \{R_0, \dots, R_k\}$ such that $\langle \Gamma \rangle = \text{Inv}(C)$. In other words $\text{Inv}(C)$ can be generated from a finite set of constant relations.

Theorem 31. *Let $\text{Inv}(C)$ be an essentially constant co-clone over some finite domain D . Then $\mathcal{I}(C)$ is finitely generated.*

Proof. Since $\text{Inv}(C)$ is essentially constant there exists $\Gamma \subseteq \{R_0, \dots, R_k\}$ such that $\langle \Gamma \rangle = \text{Inv}(C)$. It is known both that $\text{Pol}(\{R_0, \dots, R_k\})$, the clone consisting of all idempotent functions over D , is finitely generated [25] and that there exists a finite number of (strong) partial clones containing $\text{Pol}(\{R_0, \dots, R_k\})$ [11]. From this it easily follows that $\text{Pol}(\Gamma)$ is of finite order and that $\mathcal{I}(C)$ is finite. By applying Lemma 29 it follows that $\mathcal{I}(C)$ is finitely generated. \square

The reader might well ask why we do not attempt to prove a more general result than Theorem 31. The reason is that, currently, not much is known about the structure of finitely generated intervals of strong partial clones over arbitrary finite domains. For instance, it is not even known whether $\text{pPol}(\{u_1, \dots, u_n\})$, where each $u_i \subseteq D$, is of finite order. Moreover, it is known that the intersection of two strong partial clones of finite order can be of infinite order [8], which suggests that this problem is more difficult than one might believe at a first glance.

6.2 Strong Partial Clones of Infinite Order

We now turn to the problem of determining whether an interval $\mathcal{I}_{\text{fin}}(C)$ is infinitely generated. We show that $\mathcal{I}_{\text{fin}}(C)$ is always infinitely generated if C is an essentially unary clone over an arbitrary finite domain D . For finite Boolean constraint languages Γ this implies that if $\langle \Gamma \rangle \supseteq \text{Inv}(\mathbf{N}_2)$, i.e. $\text{CSP}(\Gamma)$ is NP-hard assuming $\text{P} \neq \text{NP}$, then $\text{pPol}(\Gamma)$ is of infinite order. For the proofs, we first need the following construction of a universal hash family, due to Alon et al. [2]. Given a natural number k we let $[k] = \{1, \dots, k\}$.

Theorem 32 (Section 4 of [2]). *For any k and n , there is a family H of $2^{O(k)} \log n$ functions $h_i : [n] \mapsto [k]$ such that for every $S \subset [n]$ of size k there is a function in H that is injective on S .*

Note that the bound $O(k)$ has no hidden dependency on n . Hence, if k is a constant, then $2^{O(k)} \log n \in O(\log(n))$. The purpose of a universal hash family in this paper is to, given an n -ary relation R , create an n' -ary relation R' using the universal hash family such that $\text{pPol}(R) \subseteq \text{pPol}(R')$, and such that $n' = O(n)$. In the following definition we exploit the fact that any n -ary relation R can be viewed as an $|R| \times n$ matrix where each row corresponds to a tuple in R .

Definition 33. *Let R be a relation over D , $|R| = m$, let $r \geq 1$ and let H be the universal hash family from $[m]$ to $[r]$. The closure of R under H , $H(R)$, is the relation defined as follows.*

1. let M be the matrix corresponding to R ,
2. let $g_1, \dots, g_{|D|^r}$ be an enumeration of all functions $g : [r] \mapsto D$,
3. for every $h_i \in H$ and every g_j add the column $y_{i,j}$ to M which in row $x \in [m]$ takes the value $g_j(h_i(x))$,
4. let $H(R)$ be the relation corresponding to M .

Say that a relation R over D is n -saturated if for every $t_1, \dots, t_{n'} \in R$, $n' \leq n$, for every $(x_1, \dots, x_{n'}) \in D^{n'}$ there exists an i such that $(t_1[i], \dots, t_{n'}[i]) = (x_1, \dots, x_{n'})$.

Lemma 34. *Let R be a relation with m tuples and let $r \geq 1$. Let H be the universal hash family from $[m]$ to $[r]$. Then $H(R)$ is r -saturated.*

Proof. Let $t_1, \dots, t_q \in H(R)$, $q \leq r$, let M be the matrix corresponding to $H(R)$. For every $(x_1, \dots, x_q) \in D^q$ we must prove that there exists some j such that $(x_1, \dots, x_q) = (t_1[j], \dots, t_q[j])$. Let $P = (p_1, \dots, p_q) \in [m]^q$ be the row indices of t_1, \dots, t_q , i.e., $t_i = M[p_i, \cdot]$ for each $i \in [q]$. Since H is a universal hash family, there is some $h \in H$ which is injective on P . Let $g : D^q \mapsto D$ be the function satisfying $(g(h(p_1)), \dots, g(h(p_q))) = (x_1, \dots, x_q)$. Due to the construction of $H(R)$ this implies that the column in M corresponding to h and g will enumerate (x_1, \dots, x_q) . Hence, there is a j such that $(x_1, \dots, x_q) = (t_1[j], \dots, t_q[j])$. \square

If R is a relation and Γ a constraint language we let $\text{Pol}(\Gamma)(R)$ denote the closure of R under $\text{Pol}(\Gamma)$. Formally this relation can be defined as $\text{Pol}(\Gamma)(R) = \bigcap_{R' \in \langle \Gamma \rangle, R \subseteq R'} R'$.

Lemma 35. *Let $\text{Pol}(\Gamma)$ be an essentially unary clone. If $\text{pPol}(\Gamma)$ is of finite order, then Γ can p.p. define all n -ary relations $R \in \langle \Gamma \rangle$ with at most $O(n)$ existentially quantified variables.*

Proof. Let R be an n -ary relation in $\langle \Gamma \rangle$, and let $m \leq |D|^n$ be the number of tuples in R . Let S be a finite base of $\text{pPol}(\Gamma)$, let r be the largest arity of any function in S , and let H be the r -universal hash family from $[m]$ to $[r]$ of Theorem 32. Let $R' = H(R)$. By the construction of $H(R)$ in Definition 33 it follows that $\text{ar}(R') = n + |D|^r |H| = n + |D|^r 2^{O(r)} \log m = |D|^r 2^{O(r)} O(n)$. To see that the last equality holds simply note that $\log(m) \leq \log(|D|^n) = O(n)$. Moreover, since r is a constant, it also holds that $\text{ar}(R') = O(n)$. Let $p = \text{ar}(R')$, and let $R'' = \text{Pol}(\Gamma)(R')$, i.e. R' closed under all polymorphisms of Γ . Note that $\text{Pol}(\Gamma) \subseteq \text{Pol}(R'')$. Note that

$$R(x_1, \dots, x_n) \equiv \exists x_{n+1}, \dots, x_p R''(x_1, \dots, x_n, x_{n+1}, \dots, x_p),$$

or, put in other words, R'' can p.p. define R with at most $O(n)$ existentially quantified variables. To see that this holds, simply note that $\{(x_1, \dots, x_n) \mid (x_1, \dots, x_n, x_{n+1}, \dots, x_p) \in R''\} = R$.

It remains to prove that $\text{pPol}(\Gamma) \subseteq \text{pPol}(R'')$, since this, due to the Galois connection in Theorem 2, implies that $\langle R'' \rangle_{\bar{\exists}} \subseteq \langle \Gamma \rangle_{\bar{\exists}}$ and that Γ can p.p. define R using at most $O(n)$ existentially quantified variables. Hence, let $f \in S$ be a q -ary, $q \leq r$, function. If $f \notin \text{pPol}(R'')$ then there exists $t_1, \dots, t_q \in R''$ such that $f(t_1, \dots, t_q) \notin R''$. We may assume that all t_1, \dots, t_q are distinct, as otherwise the application of f is equivalent to the application of some q' -ary partial polymorphism f' on distinct rows, where q' is the number of distinct rows represented in (t_1, \dots, t_q) [22].

Our strategy is now, using Lemma 34, to prove that we can define a total function h using the partial function f such that h does not preserve R'' . However, this also implies that $h \notin \text{Pol}(\Gamma)$, which is a contradiction since $f \in \text{pPol}(\Gamma)$. Before this proof we make one observation: for every $t \in R''$ there exists $t' \in R'$ and a unary function $h \in \text{Pol}(\Gamma)$ such that $h(t') = t$. Hence, for the tuples t_1, \dots, t_q there exists $t'_1, \dots, t'_q \in R'$ and $h_1, \dots, h_q \in \text{Pol}(\Gamma)$ such that $h_i(t'_i) = t_i$. We now define the q -ary function h as

$$h(x_1, \dots, x_q) = f(h_1(\pi_1^q(x_1, \dots, x_q)), \dots, h_q(\pi_q^q(x_1, \dots, x_q))).$$

Obviously, $h \in \text{pPol}(\Gamma)$ since it is a composition of f, h_1, \dots, h_q , and projection functions. This in turn implies that

$$h(t'_1, \dots, t'_q) = f(h_1(t'_1), \dots, h_q(t'_q)) = f(t_1, \dots, t_q) \notin R'',$$

but since $t'_1, \dots, t'_q \in R'$, R' is r -saturated and $q \leq r$, h must be a total polymorphism, i.e. $h \in \text{Pol}(\Gamma) \subseteq \text{Pol}(R'')$. This is a contradiction since $h(t'_1, \dots, t'_q) \notin R''$. Hence, $f \in \text{pPol}(R'')$. \square

With the help of this Lemma we can now prove that $\text{pPol}(\Gamma)$ is of infinite order whenever Γ is finite and $\text{Pol}(\Gamma)$ is essentially unary.

Theorem 36. *Let C be an essentially unary clone over a finite domain D . Then $\mathcal{I}_{\text{fin}}(C)$ is infinitely generated.*

Proof. Let Γ be a finite constraint language such that $\text{Pol}(\Gamma) = C$. Assume that $\text{pPol}(\Gamma)$ can be finitely generated. By Lemma 35 we then have that Γ can p.p. define all n -ary relations in $\text{Inv}(C)$ with $O(n)$ existentially quantified variables. However, this is a contradiction since $\langle \Gamma \rangle$ is superpolynomially closed by Theorem 20. To see this simply note that C cannot contain a k -edge operation for any $k \geq 2$ since a k -edge operation by definition is not essentially unary. This fact together with Lemma 18 results in a contradiction. Hence, $\text{pPol}(\Gamma)$ cannot be of finite order. \square

This theorem has a number of interesting applications. First, recall from Section 2.3 that for Boolean constraint languages Γ , $\text{CSP}(\Gamma)$ is NP-complete if and only if $\text{Pol}(\Gamma) \subseteq [\neg x]$. Hence, assuming $\text{P} \neq \text{NP}$, $\text{pPol}(\Gamma)$ is of infinite order whenever Γ is finite and $\text{CSP}(\Gamma)$ is NP-complete. This implies that describing partial polymorphisms of finite constraint languages resulting in NP-hard CSP problems is a very difficult problem. For some illustrative usages of this theorem, let $R_{1/k} = \{(x_1, \dots, x_k) \mid x_1, \dots, x_k \in \{0, 1\}, \sum_{i=1}^k x_i = 1\}$, and let Γ_{SAT}^k and Γ_{NAE}^k be the restrictions of the plain bases of BR and $\text{Inv}(\mathbf{N}_2)$, respectively, restricted to relations of arity at most k . It is easy to see that $\text{CSP}(\{R_{1/k}\})$, $\text{CSP}(\Gamma_{\text{NAE}}^k)$, and $\text{CSP}(\Gamma_{\text{SAT}}^k)$ can be seen as alternative formulations of the well-known NP-complete problems 1-IN- k -SAT, NOT-ALL-EQUAL- k -SAT, and k -SAT, respectively. Since all these languages are finite we obtain the following corollary to Theorem 36.

Corollary 37. *Let $k \geq 3$. Then $\text{pPol}(\Gamma_{\text{SAT}}^k)$, $\text{pPol}(\Gamma_{\text{NAE}}^k)$, and $\text{pPol}(R_{1/k})$ are of infinite order.*

It is worth noting that a complete dichotomy theorem for $\text{CSP}(\Gamma)$ for constraint languages Γ defined over arbitrary finite domains is not yet known. However, if $\text{Pol}(\Gamma)$ is essentially unary and every $f \in \text{Pol}(\Gamma)$ is injective, then $\text{CSP}(\Gamma)$ is NP-complete [17]. Hence, Theorem 36 also extends to many non-Boolean cases where $\text{CSP}(\Gamma)$ is NP-complete.

7 Concluding Remarks and Open Research Questions

We have studied the question of whether a polynomial amount of variables is sufficient to define any relation in a given co-clone, have provided a complete dichotomy theorem for the Boolean co-clones where this is possible, and have also given several general results for arbitrary finite domains. In the process, we have also extended the concept of a weak base from Schnoor and Schnoor [28] and have given weak bases of all Boolean co-clones of infinite order. Using these notions we have then studied the question of whether a given strong partial clone is of finite or infinite order, and proven that the latter holds for a large variety of well-studied constraint languages. We now discuss some possibilities of future research.

Polynomially closed co-clones and few subpowers

From the results in Section 3 and Section 5 we see that the question whether a co-clone of finite order is polynomially closed is related to the question whether the corresponding algebra has few subpowers. For the Boolean domain, these two notions exactly coincide, and it would be interesting to see whether this holds in the more general setting of arbitrary finite domains, possibly using some of the machinery developed in Berman et al. [3].

Partial polymorphisms and superpolynomially closed co-clones

Theorem 36 states that a $\text{pPol}(\Gamma)$ is always of infinite order whenever Γ is finite and $\text{Pol}(\Gamma)$ is essentially unary. It would be interesting to try to extend this theorem to the case when $\langle \Gamma \rangle$ is an arbitrary superpolynomially closed co-clone, and a possible starting point is to investigate the case when $\text{Pol}(\Gamma)$ can be generated from a finite set of binary functions. However, this appears to be far from straightforward, and even in the apparently simple case when $\text{Pol}(\Gamma) = [x_1 \wedge x_2]$, the proof strategy in Lemma 35, based on constructing a universal hash family, breaks down.

Partial Polymorphisms of finite Boolean constraint languages

In the light of Theorem 36, describing the partial polymorphisms of *any* finite Boolean constraint language Γ such that $\text{CSP}(\Gamma)$ is NP-complete is a challenging problem since $\text{pPol}(\Gamma)$ is of infinite order. Nevertheless, recent research shows that this problem can be circumvented by considering stronger closure operators than functional composition [22]. Using this approach it would be interesting to attempt to give a general characterization of the partial polymorphisms of the constraint languages in the “bottom” of BR, e.g., all constraint languages Γ such that $\langle \Gamma \rangle = \text{BR}$ and $\langle \Gamma \rangle_{\overline{\mathcal{F}}} \subseteq \langle R_{1/3} \rangle_{\overline{\mathcal{F}}}$.

Acknowledgements

The authors are grateful toward Peter Jonsson, Karsten Schölzel and Bruno Zanuttini, for helpful comments and suggestions.

References

1. V. B. Alekseev and A. A. Voronenko. On some closed classes in partial two-valued logic. *Discrete Mathematics and Applications*, 4(5):401–419, 1994.
2. N. Alon, R. Yuster, and U. Zwick. Color-coding. *Journal of the ACM*, 42(4):844–856, July 1995.
3. J. Berman, P. Marković, R. McKenzie, M. Valeriote, and R. Willard. Varieties with few subalgebras of powers. *Transactions of the American Mathematical Society*, 362(3):1445 – 1473, 2006.
4. V. G. Bodnarchuk, L. A. Kaluzhnin, V. N. Kotov, and B. A. Romov. Galois theory for Post algebras. I. *Cybernetics*, 5:243–252, 1969.

5. V. G. Bodnarchuk, L. A. Kaluzhnin, V. N. Kotov, and B. A. Romov. Galois theory for Post algebras. II. *Cybernetics*, 5:531–539, 1969.
6. E. Böhler, N. Creignou, S. Reith, and H. Vollmer. Playing with Boolean blocks, part I: Post’s lattice with applications to complexity theory. *ACM SIGACT-Newsletter*, 34(4):38–52, 2003.
7. F. Börner and L. Haddad. Maximal partial clones with no finite basis. *Algebra Universalis*, 40(4):453–476, 1998.
8. M. Couceiro and L. Haddad. Finitely generated maximal partial clones and their intersections. In *Proceedings of the 40th IEEE International Symposium on Multiple-Valued Logic, (ISMVL-2010)*, pages 122–127. IEEE Computer Society, 2010.
9. N. Creignou, P. Kolaitis, and B. Zanuttini. Structure identification of Boolean relations and plain bases for co-clones. *Journal of Computer and System Sciences*, 74(7):1103–1115, November 2008.
10. N. Creignou, P. G. Kolaitis, and H. Vollmer, editors. *Complexity of Constraints - An Overview of Current Research Themes*, volume 5250 of *Lecture Notes in Computer Science*. Springer, 2008.
11. J. Fugère and L. Haddad. On partial clones containing all idempotent partial operations. In *Proceedings of the 28th IEEE International Symposium on Multiple-Valued Logic (ISMVL-98)*, pages 369–373, 1998.
12. D. Geiger. Closed systems of functions and predicates. *Pacific Journal of Mathematics*, 27(1):95–100, 1968.
13. L. Haddad and J. Fugère. Maximal chains of partial clones containing all idempotent partial functions. In *Proceedings of the 29th IEEE International Symposium on Multiple-Valued Logic (ISMVL-99)*, pages 80–84. IEEE Computer Society, 1999.
14. P. Idziak, P. Marković, R. McKenzie, M. Valeriote, and R. Willard. Tractability and learnability arising from algebras with few subpowers. *SIAM Journal on Computing*, 39(7):3023–3037, June 2010.
15. P. Jeavons. On the algebraic structure of combinatorial problems. *Theoretical Computer Science*, 200:185–204, 1998.
16. P. Jeavons, D. Cohen, and M. C. Cooper. Constraints, consistency and closure. *Artificial Intelligence*, 101(1–2):251 – 265, 1998.
17. P. Jeavons, D. Cohen, and M. Gyssens. Closure properties of constraints. *Journal of the ACM*, 44(4):527–548, July 1997.
18. P. Jonsson, V. Lagerkvist, G. Nordh, and B. Zanuttini. Complexity of SAT problems, clone theory and the exponential time hypothesis. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA-2013)*, pages 1264–1277, 2013.
19. P. Jonsson, V. Lagerkvist, J. Schmidt, and H. Uppman. Relating the time complexity of optimization problems in light of the exponential-time hypothesis. In *Proceedings of the 39th International Symposium on Mathematical Foundations of Computer Science (MFCS-14)*, pages 408–419, Berlin, Heidelberg, 2014. Springer-Verlag.
20. V. Lagerkvist. Weak bases of Boolean co-clones. *Information Processing Letters*, 114(9):462–468, 2014.
21. V. Lagerkvist. *Strong Partial Clones and the Complexity of Constraint Satisfaction Problems: Limitations and Applications*. PhD thesis, Linköping University, The Institute of Technology, 2016.
22. V. Lagerkvist, M. Wahlström, and B. Zanuttini. Bounded bases of strong partial clones. In *Proceedings of the 45th International Symposium on Multiple-Valued Logic (ISMVL-2015)*, pages 189–194, 2015.
23. G. Nordh and B. Zanuttini. Frozen Boolean partial co-clones. In *Proceedings of the 39th International Symposium on Multiple-Valued Logic (ISMVL-2009)*, pages 120 –125, 2009.
24. E. Post. The two-valued iterative systems of mathematical logic. *Annals of Mathematical Studies*, 5:1–122, 1941.
25. R. W. Quackenbush. On the composition of idempotent functions. *Algebra Universalis*, 1(1):7–12, 1971.
26. B.A. Romov. The algebras of partial functions and their invariants. *Cybernetics*, 17(2):157–167, 1981.
27. T. Schaefer. The complexity of satisfiability problems. In *Proceedings of the 10th Annual ACM Symposium on Theory Of Computing (STOC-78)*, pages 216–226. ACM Press, 1978.
28. H. Schnoor and I. Schnoor. Partial polymorphisms and constraint satisfaction problems. In N. Creignou, P. G. Kolaitis, and H. Vollmer, editors, *Complexity of Constraints*, volume 5250 of *Lecture Notes in Computer Science*, pages 229–254. Springer Berlin Heidelberg, 2008.
29. K. Schölzel. Dichotomy on intervals of strong partial boolean clones. *Algebra Universalis*, 73(3-4):347–368, 2015.