Galois correspondence for counting quantifiers

ANDREI A. BULATOV[1]*, AMIR HEDAYATY[1]†

Simon Fraser University

We introduce a new type of closure operator on the set of relations, max-implementation, and its weaker analog, max-quantification. Then we show that approximation preserving reductions between counting constraint satisfaction problems (CSPs) are preserved by these two types of closure operators. Together with some previous results this means that the approximation complexity of counting CSPs is determined by partial co-clones of relations that are additionally closed under these new types of closure operators. Galois correspondences of various kind have proved to be quite helpful in the study of the complexity of the CSP. While we were unable to identify a Galois correspondence for partial co-clones closed under max-implementation and max-quantification, we obtain such results for slightly different type of closure operators, k-existential quantification. This type of quantifiers are known as counting quantifiers in model theory, and often used to enhance first order logic languages. We characterize partial co-clones of relations closed under k-existential quantification as sets of relations invariant under a set of partial functions that satisfy the condition of k-subset surjectivity.

This is an extended version of [*A.Bulatov and A.Hedayaty. Counting quantifiers, subset surjective functions, and counting CSPs. ISMVL 2012, p.331–336.*]

Key words: counting constraint satisfaction problem, approximation, co-clones, Galois correspondence

1 INTRODUCTION

Clones of functions and clones of relations in their various incarnations have proved to be an immensely powerful tool in the study of the complexity of different versions of the Constraint Satisfaction Problem (CSP, for short). In a CSP the aim is to find an assignment of values to a given set of variables, subject to constraints on the values that can be assigned simultaneously to certain specified subsets of variables. A CSP can also be expressed as the problem of deciding whether a given conjunctive formula has a model. In the counting version of the CSP the goal is to find the number of satisfying assignments, and in the quantified version we need to verify if a first order sentence, whose quantifier-free part is conjunctive, is true in a given model.

The general CSP is NP-complete [22]. However, many practical and theoretical problems can be expressed in terms of CSPs using constraints of a certain restricted form. One of the most widely used way to restrict a constraint satisfaction problem is to specify the set of allowed constraints, which is usually a collection of relations on a finite set. The key result is that this set of relations can usually be assumed to be a co-clone of relations of a certain kind. More precisely, a generic statement asserts that if a relation R belongs to the co-clone generated by a set Γ of relations then the CSP over $\Gamma \cup \{R\}$ is polynomial time reducible to the CSP over Γ. Then we can use the appropriate Galois connection to transfer the question about sets of relations to a question about certain classes of functions.

* email: abulatov@cs.sfu.ca
† email: aha49@cs.sfu.ca

For the classical decision CSP such a result was obtained by Jeavons et al. [21], who proved that intersection of relations (that is, conjunction of the corresponding predicates) and projections (that is, existential quantification) give rise to polynomial time reducibility of CSPs. Therefore in the study of the complexity of the CSP it suffices to focus on co-clones. Using the result of Geiger [17] or the one of Bodnarchuk et al. [2] one can instead consider clones of functions. A similar result is true for the counting CSP as shown by Bulatov and Dalmau [8]. In the case of quantified CSP, Börner et al. proved [3] that conjunction, existential quantification, and also universal quantification give rise to a polynomial time reduction between quantified problems. The appropriate class of functions is then the class of surjective functions. Along with the usual counting CSP, a version, in which one is required to approximate the number of solutions, has also been considered. The standard polynomial time reduction between problems is not suitable for approximation complexity. In this case, therefore, another type of reductions, approximation preserving, or, AP-reductions, is used. The first author proved in [7] that conjunction of predicates gives rise to an AP-reduction between approximation counting CSPs. By the Galois connection established by Fleischner and Rosenberg [16], the approximation complexity of a counting CSP is a property of a clone of partial functions.

In most cases establishing the connection between clones of functions and reductions between CSPs has led to a major success in the study of the CSP. For the decision problem, a number of very strong results have been proved using methods of universal algebra [9, 4, 5, 1, 19]. For the exact counting CSP a complete complexity classification of such problems has been obtained [6]. Substantial progress has been also made in the case of quantified CSP [13].

Compared to the results cited above the progress made in the approximation counting CSP is modest. Perhaps, one reason for this is that clones of partial functions are much less studied, and much more diverse than clones of total functions. In this paper we attempt to overcome to some extent the difficulties arising from this weakness of partial clones.

In the first part of the paper we introduce new types of quantification and show that such quantifications, we call them max-implementation and max-quantification, give rise to AP-reductions between approximation counting CSPs. Intuitively, applying the max-quantifier to a relation $R(x_1, \dots, x_n, y)$ results in the relation $\exists^1_{\max} y R(x_1, \dots, x_n, y)$ that contains those tuples $(a_1, \dots, a_n)$ that have a maximal number of extensions $(a_1, \dots, a_n, b)$ such that $R(a_1, \dots, a_n, b)$ is satisfied. Max-implementation, $\exists_{\max}$, is a similar construction, but applied to a group of variables. Sets of relations closed with respect to this new type of quantification will be called max-co-clones. Thus we strengthen the closure operator on sets of relation hoping that the sets of functions corresponding to the new type of Galois connection are easier to study. We were unable, however, to describe a Galois connection for sets closed under max-implementation and max-quantification. Instead, we consider a somewhat close type of quantifiers, k-existential quantifiers. Quantifiers of this type are known as counting quantifiers in model theory, and often used to enhance first order logic languages (see, e.g. [20, 15]). Counting quantifiers are similar to max-quantifiers, although do not capture them completely. We call sets of relations closed under conjunctions and k-existential quantification k-existential co-clones. On the functional side, an n-ary (partial) function on a set D is said to be k-subset surjective if it is surjective on any collection of k-element subsets. More precisely, for any k-element subsets $A_1, \dots, A_n \subseteq D$ the set $f(A_1, \dots, A_n)$ contains at least k elements. The second result of the paper asserts that k-existential co-clones are exactly the sets of relations invariant with respect to a set of k-subset surjective (partial) functions.

2 PRELIMINARIES

By $[n]$ we denote the set $\{1, \dots, n\}$. For a set D, by D^n we denote the set of all *n-tuples* of elements of D. An n-ary relation is any set $R \subseteq D^n$. The number n is called the *arity* of R and denoted $\mathsf{ar}(R)$. Tuples will be denoted in boldface, say, $\mathbf{a}$, and their entries will be denoted by $\mathbf{a}[1], \dots, \mathbf{a}[n]$. For $I = (i_1, \dots, i_k) \subseteq [n]$ by $\mathrm{pr}_I \mathbf{a}$ we denote the tuple $(\mathbf{a}[i_1], \dots, \mathbf{a}[i_k])$, and we use $\mathrm{pr}_I R$ to denote $\{\mathrm{pr}_I \mathbf{a} \mid \mathbf{a} \in R\}$. We will also need predicates corresponding to relations. To simplify the notation we use the same symbol for a relation and the corresponding predicate, for instance, for an n-ary relation R the corresponding predicate $R(x_1, \dots, x_n)$ is given by $R(\mathbf{a}[1], \dots, \mathbf{a}[n]) = 1$ if and only if $\mathbf{a} \in R$. Relations and predicates are used interchangeably.

For a set Γ of relations over a set D, the set $\langle\langle\Gamma\rangle\rangle$ includes all relations that can be expressed (as a predicate) using (a) relations from Γ, together with the binary equality relation $=_D$ on D, (b) conjunctions, and (c) existential quantification. This set is called the *co-clone generated by* Γ.

Partial co-clone $\langle\Gamma\rangle$ *generated by* Γ is obtained in a similar way by disallowing existential quantification. $\langle\Gamma\rangle$ includes all relations that can be expressed using (a) relations from Γ, together with $=_D$, and (b) conjunctions,

If $\Gamma = \langle\Gamma\rangle$ or $\Gamma = \langle\langle\Gamma\rangle\rangle$, the set Γ is said to be a *partial co-clone*, or a *co-clone*, respectively.

Sometimes there is no need to apply even conjunction to produce a new relation. For instance, $Q(x, y) = R(x, y, y)$ defines a binary relation from a ternary one. Therefore it is often convenient, especially for technical purposes, to group manipulations with variables of a relation into a separate category. More formally, for a relation $R(x_1, \dots, x_n)$ and a mapping $\pi\colon \{x_1, \dots, x_n\} \to V$, where V is some set of variables, πR denotes the relation $R(\pi(x_1), \dots, \pi(x_n))$. We will understand by (partial) co-clones sets of relations closed under manipulation with variables, conjunction, and existential quantification (respectively, closed under manipulation with variables and conjunction).

Co-clones and partial co-clones can often be conveniently and concisely represented through functions and partial functions, respectively.

Let R be a (k-ary) relation on a set D, and $f\colon D^n \to D$ an n-ary function on the same set. Function f *preserves* R, or is a *polymorphism* of R, if for any n tuples $\mathbf{a}_1, \dots, \mathbf{a}_n \in R$ the tuple $f(\mathbf{a}_1, \dots, \mathbf{a}_n)$ obtained by component-wise application of f also belongs to R. Relation R in this case is said to be *invariant* with respect to f. The set of all functions that preserve every relation from a set of relations Γ is denoted by $\mathsf{Pol}(\Gamma)$, the set of all relations invariant with respect to a set of functions C is denoted by $\mathsf{Inv}(C)$.

Operators Inv and Pol form a Galois connection between sets of functions and sets of relations. Sets of the form $\mathsf{Inv}(C)$ are precisely co-clones; on the functional side there is another type of closed sets.

A set of functions is said to be a *clone* of functions if it is closed under superpositions and contains all the *projection* functions, that is, functions of the form $f(x_1, \dots, x_n) = x_i$. Sets of functions of the form $\mathsf{Pol}(\Gamma)$ are exactly clones of functions [23] .

The study of the counting CSP also makes use of another Galois connection, a connection between partial co-clones and sets of *partial functions*. An n-ary partial function f on a set D is just a partial mapping $f\colon D^n \to D$. As in the case of total functions, a partial function f *preserves* relation R, if for any n tuples $\mathbf{a}_1, \dots, \mathbf{a}_n \in R$ the tuple $f(\mathbf{a}_1, \dots, \mathbf{a}_n)$ obtained by component-wise application of f is either undefined or belongs to R. The set of all partial functions that preserve every relation from a set Γ of relations is denoted by $\mathsf{pPol}(\Gamma)$.

The set of all tuples from D^n on which f is defined is called the *domain* of f and denoted by $\mathsf{Dom}(f)$. A set of functions is said to be *down-closed* if along with a function f it contains any function f' such that $\mathsf{Dom}(f') \subseteq \mathsf{Dom}(f)$ and $f'(a_1, \dots, a_n) = f(a_1, \dots, a_n)$ for every tuple $(a_1, \dots, a_n) \in \mathsf{Dom}(f')$. A down-closed set of functions, containing all projections and closed under superpositions is called a *partial clone*. Fleischner and Rosenberg [16] proved that partial clones are exactly the sets of the form $\mathsf{pPol}(\Gamma)$ for a certain Γ, and that the partial co-clones are precisely the sets $\mathsf{Inv}(C)$ for collections C of partial functions.

3 APPROXIMATE COUNTING AND MAX-IMPLEMENTATION

Let D be a set, and let Γ be a finite set of relations over D. An instance of the counting Constraint Satisfaction Problem, $\#\mathrm{CSP}(\Gamma)$, is a pair $\mathcal{P} = (V, \mathcal{C})$ where V is a set of *variables*, and $\mathcal{C}$ is a set of *constraints*. Every constraint is a pair $\langle \mathbf{s}, R\rangle$, in which R is a member of Γ, and $\mathbf{s}$ is a tuple of variables from V of length $\mathsf{ar}(R)$ (possibly with repetitions). A *solution* to $\mathcal{P}$ is a mapping $\varphi : V \to D$ such that $\varphi(\mathbf{s}) \in R$ for every constraint $\langle \mathbf{s}, R\rangle \in \mathcal{C}$. The objective in $\#\mathrm{CSP}(\Gamma)$ is to find the number $\#\mathcal{P}$ of solutions to a given instance $\mathcal{P}$.

We are interested in the complexity of this problem depending on the set Γ. The complexity of the exact counting problem (when we are required to find the exact number of solutions) is settled in [6] by showing that for any finite D and any set Γ of relations over D the problem is polynomial time solvable or is complete in a natural complexity class $\#P$. One of the key steps in that line of research is the following result: For a relation R and a set of relations Γ over

D, if R belongs to the co-clone generated by Γ, then $\#\mathrm{CSP}(\Gamma \cup \{R\})$ is polynomial time reducible to $\#\mathrm{CSP}(\Gamma)$. This results emphasizes the importance of co-clones in the study of constraint problems.

The situation is different when we are concerned about approximating the number of solutions. We will need some notation and terminology. Let A be a counting problem. An algorithm Alg is said to be an *approximation algorithm* for A with relative error ε (which may depend on the size of the input) if it is polynomial time and for any instance $\mathcal{P}$ of A it outputs a certain number $\mathsf{Alg}(\mathcal{P})$ such that $\mathsf{Alg}(\mathcal{P}) = 0$ if $\mathcal{P}$ has no solution and

$$\frac{|\#\mathcal{P} - \mathsf{Alg}(\mathcal{P})|}{\#\mathcal{P}} < \varepsilon$$

otherwise, where $\#\mathcal{P}$ denotes the exact number of solutions to $\mathcal{P}$.

The following framework is viewed as one of the most realistic models of efficient computations. A *fully polynomial approximation scheme* (FPAS, for short) for a problem A is an algorithm Alg such that: It takes as input an instance $\mathcal{P}$ of A and a real number $\varepsilon > 0$, the relative error of Alg on the input $(\mathcal{P}, \varepsilon)$ is less than ε, and Alg is polynomial time in the size of $\mathcal{P}$ and $\log(\frac{1}{\varepsilon})$.

To determine the approximation complexity of problems approximation preserving reductions are used. Suppose A and B are two counting problems whose complexity (of approximation) we want to compare. An *approximation preserving reduction* or *AP-reduction* from A to B is an algorithm Alg, using B as an oracle, that takes as input a pair $(\mathcal{P}, \varepsilon)$ where $\mathcal{P}$ is an instance of A and $0 < \varepsilon < 1$, and satisfies the following three conditions: (i) every oracle call made by Alg is of the form $(\mathcal{P}', \delta)$, where $\mathcal{P}'$ is an instance of B, and $0 < \delta < 1$ is an error bound such that $\log\left(\frac{1}{\delta}\right)$ is bounded by a polynomial in the size of $\mathcal{P}$ and $\log\left(\frac{1}{\varepsilon}\right)$; (ii) the algorithm Alg meets the specifications for being an FPAS for A whenever the oracle meets the specification for being an FPAS for B; and (iii) the running time of Alg is polynomial in the size of $\mathcal{P}$ and $\log(\frac{1}{\varepsilon})$. If an approximation preserving reduction from A to B exists we denote it by $A \leq_{\mathrm{AP}} B$, and say that A is *AP-reducible to* B.

Similar to co-clones and polynomial time reductions, partial co-clones can be shown to be preserved by AP-reductions.

Theorem 1 ([7]) *Let R be a relation and Γ be a set of relations over a finite set such that R belongs to $\langle\Gamma\rangle$. Then $\#\mathrm{CSP}(\Gamma \cup \{R\})$ is AP-reducible to $\#\mathrm{CSP}(\Gamma)$.*

This result however has two significant setbacks. First, partial co-clones are not studied to the same extent as regular co-clones, and, due to greater diversity, are not believed to be ever studied to a comparable level. Second, it does not use the full power of AP-reductions, and therefore leaves significant space for improvements. In the rest of this section we try to improve upon the second issue.

Definition 2 *Let Γ be a set of relations on a set D, and let R be an n-ary relation on D. Let $\mathcal{P}$ be an instance of $\#\mathrm{CSP}(\Gamma)$ over the set of variables consisting of $V = V_x \cup V_y$, where $V_x = \{x_1, x_2, \cdots, x_n\}$ and $V_y = \{y_1, y_2, \cdots, y_q\}$. For any assignment $\varphi : V_x \to D$, let $\#\varphi$ be the number of assignments $\psi : V_y \to D$ such that $\varphi \cup \psi$ satisfy $\mathcal{P}$. Let M be the maximum value of $\#\varphi$ among all assignments φ of V_x. The instance $\mathcal{P}$ is said to be a* max-implementation *of R if a tuple $(\varphi(x_1), \cdots, \varphi(x_n))$ is in R if and only if $\#\varphi = M$.*

Theorem 3 *If there is max-implementation of R by Γ, then $\#\mathrm{CSP}(\Gamma \cup \{R\}) \leq_{AP} \#\mathrm{CSP}(\Gamma)$.*

Proof: Let $\mathcal{P} = (V = V_x \cup V_y, \mathcal{C})$ be a max-implementation of R by Γ, and let M be the maximal number of extensions of assignments of V_x to solutions of $\mathcal{P}$. For any instance $\mathcal{P}_1 = (V_1, \mathcal{C}_1)$ of $\#\mathrm{CSP}(\Gamma \cup \{R\})$ we construct an instance $\mathcal{P}_2 = (V_2, \mathcal{C}_2)$ of $\#\mathrm{CSP}(\Gamma)$ as follows.

- Choose a sufficiently large integer m (to be determined later).
- Let $C_1, \ldots, C_\ell \in \mathcal{C}_1$ be the constraints from $\mathcal{P}_1$ involving R, $C_i = \langle \mathbf{s}_i, R\rangle$. Set $V_2 = V_1 \cup \bigcup_{i=1}^{\ell}(V_1^i \cup \ldots \cup V_m^i)$, where each V_j^i is a fresh copy of V_y.

- Let $\mathcal{C}$ be the set of constraints of $\mathcal{P}$. Set $\mathcal{C}_2 = (\mathcal{C}_1 - \{C_1, \dots, C_\ell\}) \cup \bigcup_{i=1}^{\ell}(\mathcal{C}_1^i \cup \dots \cup \mathcal{C}_m^i)$, where each $\mathcal{C}_j^i$ is a copy of $\mathcal{C}$ defined as follows. For each $\langle \mathbf{s}, Q\rangle \in \mathcal{C}$ we include $\langle \mathbf{s}_j^i, Q\rangle$ into $\mathcal{C}_j^i$, where $\mathbf{s}_j^i$ is obtained from $\mathbf{s}$ replacing every variable from V_y with its copy from V_j^i.

Now, as is easily seen, every solution of $\mathcal{P}_1$ can be extended to a solution of $\mathcal{P}_2$ in $M^{\ell m}$ ways. Observe that sometimes the restriction of a solution ψ of $\mathcal{P}_2$ to V_1 is not a solution of $\mathcal{P}_1$. Indeed, it may happen that although ψ satisfies every copy $\mathcal{C}_j^i$ of $\mathcal{P}$, its restriction to $\mathbf{s}_j^i$ does not belong to R, simply because this restriction does not have sufficiently many extensions to solutions of $\mathcal{P}$. However, any assignment to V_1 that is not a solution to $\mathcal{P}_1$ can be extended to a solution of $\mathcal{P}_2$ in at most $(M-1)^m \cdot M^{(\ell-1)m}$ ways. Hence,

$$M^{\ell m} \cdot \#\mathcal{P}_1 \le \#\mathcal{P}_2 \le M^{\ell m} \cdot \#\mathcal{P}_1 + |V_1|^{|D|} \cdot (M-1)^m \cdot M^{(\ell-1)m}.$$

Then we output $\frac{\#\mathcal{P}_2}{M^{\ell m}}$.

Let $|V_1| = k$ and $|D| = d$. Given a desired relative error ε we have to find m such that

$$\frac{\frac{\#\mathcal{P}_2}{M^{\ell m}} - \#\mathcal{P}_1}{\#\mathcal{P}_1} < \varepsilon.$$

A straightforward computation shows that any

$$m > \frac{d\log k - \log\varepsilon}{\log M - \log(M-1)}$$

achieves the goal. □

Max-implementation can be used as another closure operator on the set of relations. Let $R(x_1, \dots, x_n, y_1, \dots, y_m)$ be a relation on a set D. By $\exists_{\max}(y_1, \dots, y_m)R(x_1, \dots, x_n, y_1, \dots, y_m)$ we denote the relation $Q(x_1, \dots, x_n)$ on the same set given by the rule: $\mathbf{a} \in Q$ if and only if there are M tuples $\mathbf{b} \in D^m$ such that $(\mathbf{a}, \mathbf{b}) \in R$, where M is the maximal number of elements in the set $\{\mathbf{b} \mid (\mathbf{a}, \mathbf{b}) \in Q\}$ over all $\mathbf{a} \in D^n$. A set Γ of relations over D is said to be a *max-co-clone* if it contains the equality relations, and closed under conjunctions and max-implementations. The smallest max-co-clone containing a set Γ of relations is called the *max-co-clone generated by* Γ and denoted $\langle\Gamma\rangle_{\max}$.

Lemma 4 *Let Γ be a set of relations and $R \in \langle\Gamma\rangle_{\max}$, $R \neq \varnothing$. Then there is a max-implementation of R by Γ.*

Proof: Suppose $R \in \langle\Gamma\rangle_{\max}$. We need to show that R can be represented as $R(x_1, \dots, x_n) = \exists_{\max}(y_1, \dots, y_m)$ $\Phi(x_1, \dots, x_n, y_1, \dots, y_m)$, where Φ is quantifier free. To this end it suffices to prove three equalities:

1. if $R(x_1, \dots, x_n) = \exists_{\max}(y_1, \dots, y_m)\Phi(x_1, \dots, x_n, y_1, \dots, y_m)$ and π is a transformation of the set $\{x_1, \dots, x_n\}$ then $(\pi R)(x_1, \dots, x_n) = \exists_{\max}(y_1, \dots, y_m)\Phi(\pi(x_1), \dots, \pi(x_n), y_1, \dots, y_m)$;
2. if $R(x_1, \dots, x_n) = \exists_{\max}(y_1, \dots, y_m)\Phi_1(x_1, \dots, x_n, y_1, \dots, y_m) \wedge \exists_{\max}(z_1, \dots, z_r)\Phi_2(x_1, \dots, x_n, z_1, \dots, z_r)$, then $R(x_1, \dots, x_n) = \exists_{\max}(y_1, \dots, y_m, z_1, \dots, z_r)(\Phi_1(x_1, \dots, x_n, y_1, \dots, y_m) \wedge \Phi_2(x_1, \dots, x_n, z_1, \dots, z_r))$;
3. if $R(x_1, \dots, x_n) = \exists_{\max}(y_1, \dots, y_m)\exists_{\max}(z_1, \dots, z_r)\Phi(x_1, \dots, x_n, y_1, \dots, y_m, z_1, \dots, z_r)$, then there is a quantifier free formula Ψ such that $R(x_1, \dots, x_n) = \exists_{\max}(u_1, \dots, u_s)\Psi(x_1, \dots, x_n, u_1, \dots, u_s)$.

(1) follows straightforwardly from definitions.

(2) $\mathbf{a} \in R$ if and only if it has the maximal number of extensions in both Φ_1 and Φ_2. Without loss of generality, sets $\{y_1, \dots, y_m\}$ and $\{z_1, \dots, z_r\}$ are disjoint. Let a tuple $\mathbf{a} \in R$ have M_1 extensions in Φ_1 and M_2 extensions in Φ_2. Then it has M_1M_2 extensions in $\Phi_1 \wedge \Phi_2$. On the other hand, let $\mathbf{a} \notin R$. Let also it have M_1' extensions in Φ_1 and M_2' extensions in Φ_2, and either $M_1' < M_1$ or $M_2' < M_2$. Since such tuple has $M_1'M_2' < M_1M_2$ extensions, it does not belong to the relation defined by $R(x_1, \dots, x_n) = \exists_{\max}(y_1, \dots, y_m, z_1, \dots, z_r)(\Phi_1(x_1, \dots, x_n, y_1, \dots, y_m) \wedge \Phi_2(x_1, \dots, x_n, z_1, \dots, z_r))$ as well.

(3) Observe first that $R(x_1, \ldots, x_n)$ does not necessarily equal

$$\exists_{\max}(y_1, \ldots, y_m, z_1, \ldots, z_r)\Phi(x_1, \ldots, x_n, y_1, \ldots, y_m, z_1, \ldots, z_r).$$

Indeed, let Φ' denote the formula

$$Q(x_1, \ldots, x_n, y_1, \ldots, y_m) = \exists(z_1, \ldots, z_r)\Phi(x_1, \ldots, x_n, y_1, \ldots, y_m, z_1, \ldots, z_r).$$

Then it is possible that although every extension of a tuple $\mathbf{a}$ to $(\mathbf{a}, \mathbf{b}) \in Q$ has very few extensions to a tuple from Φ, and so $\mathbf{a} \notin R$, the number of extensions $\mathbf{b}$ is large so that combined $\mathbf{a}$ has enough extensions to tuples from Φ.

To avoid this we make sure that extensions to tuples from Q cannot make up for extensions to Φ. Let M be the maximal number of extensions $\mathbf{b}$ of tuple $\mathbf{a}$ such that $(\mathbf{a}, \mathbf{b}) \in Q$, and N the maximal number of extensions $\mathbf{c}$ of $(\mathbf{a}, \mathbf{b}) \in Q$ to $(\mathbf{a}, \mathbf{b}, \mathbf{c}) \in \Phi$. Let also L be the maximal number of extensions $\mathbf{b}$ of $\mathbf{a} \in R$; it is possible that $L < M$. Set

$$c = \max\left(1, \left\lceil \log \frac{L}{M} / \log \frac{N-1}{N} \right\rceil\right).$$

We show that $R(x_1, \ldots, x_n) = \exists_{\max}(u_1, \ldots, u_s)\Psi(x_1, \ldots, x_n, u_1, \ldots, u_s)$, where $\{u_1, \ldots, u_s\} = \{y_1, \ldots, y_m, z_1^1, \ldots, z_r^1, \ldots, z_1^c, \ldots, z_r^c\}$, and

$$\Psi(x_1, \ldots, x_n, u_1, \ldots, u_s) = \bigwedge_{s=1}^{c} \Phi(x_1, \ldots, x_n, y_1, \ldots, y_m, z_1^s, \ldots, z_r^s).$$

If a tuple $\mathbf{a}$ belongs to R it is extendable in L ways to a tuple from Q, and then every such extended tuple $(\mathbf{a}, \mathbf{b})$ is extendable in N ways to a tuple from Φ. Therefore $\mathbf{a}$ has LN^c extensions to a tuple from Ψ. On the other hand, if $\mathbf{a} \notin R$, then it can be extended in at most M ways to a tuple $(\mathbf{a}, \mathbf{b}) \in Q$, then this tuple is extendable in at most $N-1$ ways to a tuple from Φ. Thus $\mathbf{a} \notin R$ has

$$M(N-1)^c = LN^c \cdot \frac{M}{L}\left(\frac{N-1}{N}\right)^c < LN^c$$

extensions. □

The next natural step would be to find a class of functions and a closure operator on the set of functions that give rise to a Galois connection capturing max-co-clones.

Problem 1 *Find a class $\mathcal{F}$ of (partial) functions and a closure operator $[\cdot]$ on this class such that for any set of relations Γ and any set $C \subseteq \mathcal{F}$ it holds that $\langle\Gamma\rangle_{\max} = \mathsf{Inv}(\mathcal{F} \cap \mathsf{pPol}(\Gamma))$, and $[C] = \mathcal{F} \cap \mathsf{pPol}\,\mathsf{Inv}(C)$.*

In all the cases previously studied the projection (or quantification) type operators on relations can be reduced to quantifying away a single variable. Therefore a meaningful relaxation of max-co-clones restricts the use of max-implementation to one auxiliary variable. Let Φ be a formula with free variables $x_1, \ldots, x_n$ and y over set D and some predicate symbols. Then $a_1, \ldots, a_n$ satisfy $\Psi(x_1, \ldots, x_n) = \exists^1_{\max} y\Phi(x_1, \ldots, x_n, y)$ if and only if the number of $b \in D$ such that $\Phi(a_1, \ldots, a_n, b)$ is true is maximal among all tuples $(c_1, \ldots, c_n) \in D^n$. The quantifier $\exists^1_{\max}$ will be called *max-quantifier*. A set of relations Γ over D is said to be a *max-existential co-clone* if it contains the equality relation, and closed under conjunctions and max-existential quantification. The smallest max-existential co-clone containing a set of relations Γ is called the *max-existential co-clone generated by* Γ and denoted $\langle\Gamma\rangle^1_{\max}$.

However, max-implementations seem to inherently involve a number of variables, rather than a single variable. In the end of this paper we mention a description of Boolean max-co-clones. This description can be used to show that max-implementations are provably more powerful than max-quantification (see below). In the Boolean case every max-quantification is equivalent to either existential quantification, or universal quantification. Sets of relations on $\{0, 1\}$ closed under these two types of quantifications are well known: these are sets of invariant relations of sets of surjective functions [3]. However, not all of them are max-co-clones.

Problem 2 *Find a class $\mathcal{F}$ of (partial) functions and a closure operator $[\cdot]$ on this class such that for any set of relations Γ and any set of functions $C \subseteq \mathcal{F}$ it holds that $\langle\Gamma\rangle^1_{\max} = \mathsf{Inv}(\mathcal{F} \cap \mathsf{pPol}(\Gamma))$, and $[C] = \mathcal{F} \cap \mathsf{pPol}\,\mathsf{Inv}(C)$.*

In the next section we consider certain constructions approximating max-existential co-clones.

4 K-EXISTENTIAL AND MAX-EXISTENTIAL CO-CLONES

In order to approach max-quantification we consider counting quantifiers that have been used in model theory to increase the power of first order logic [20, 15].

Let Φ be a formula with free variables $x_1, \dots, x_n$ and y over set D and some predicate symbols. Then $a_1, \dots, a_n$ satisfy $\Psi(x_1, \dots, x_n) = \exists_k y \Phi(x_1, \dots, x_n, y)$ if and only if $\Phi(a_1, \dots, a_n, b)$ is true for at least k values $b \in D$. The quantifier $\exists_k$ will be called *k-existential quantifier*. It is easy to see that 1-existential quantifier is just the regular existential quantifier, and the $|D|$-existential quantifier is equivalent to the universal quantifier on set D.

We now introduce several types of co-clones depending on what kind of k-existential quantifiers are allowed. A set of relations Γ over set D is said to be a *k-existential partial co-clone* if it contains the equality relation $=_D$, and closed under manipulations with variables, conjunction, and k-existential quantification. The smallest k-existential partial co-clone containing a set of relations Γ is called the *k-existential partial co-clone generated by* Γ and denoted $\langle\Gamma\rangle_k$. In a similar way we can define sets of relations closed under several counting quantifiers. Let $K \subseteq \mathbb{N}$. A set of relations Γ over set D is said to be a *K-existential partial co-clone* if it contains the equality relation $=_D$, and closed under manipulations with variables, conjunction, and k-existential quantification for $k \in K$. Clearly, if Γ is a set of relations on an m-element set, we may assume $K \subseteq [m]$. If $1 \in K$, set Γ is closed under existential quantification, and so it is called a *K-existential co-clone*. If, in addition, $K = \{1, k\}$, Γ is called *k-existential co-clone*. The set Γ is said to be a *counting co-clone** if it is an $\mathbb{N}$-existential partial co-clone, that is, if it contains $=_D$, and closed under conjunctions and k-existential quantification for all $k \geq 1$. The smallest K-existential partial co-clone (K-existential co-clone, k-existential co-clone, counting co-clone) containing Γ is called the *K-existential partial co-clone* (*K-existential co-clone*, *k-existential co-clone*, *counting co-clone*) *generated by* Γ and denoted $\langle\Gamma\rangle_K$ ($\langle\langle\Gamma\rangle\rangle_K$, $\langle\langle\Gamma\rangle\rangle_k$, $\langle\langle\Gamma\rangle\rangle_\infty$, respectively).

We observe some simple properties of counting quantifiers.

Lemma 5 *Let $\Phi(x_1, \dots, x_n, y_1, \dots, y_m)$ and $\Psi(x_1, \dots, x_n, z_1, \dots, z_\ell)$ be conjunctive quantifier free formulas. Then*

$$\begin{aligned}&\exists_{s_1} y_1 \dots \exists_{s_m} y_m \exists_{t_1} z_1 \dots \exists_{t_\ell} z_\ell \, (\Phi(x_1, \dots, x_n, y_1, \dots, y_m) \wedge \Psi(x_1, \dots, x_n, z_1, \dots, z_\ell)) \\ &= \quad (\exists_{s_1} y_1 \dots \exists_{s_m} y_m \, (\Phi(x_1, \dots, x_n, y_1, \dots, y_m)) \wedge (\exists_{t_1} z_1 \dots \exists_{t_\ell} z_\ell \, \Psi(x_1, \dots, x_n, z_1, \dots, z_\ell)),\end{aligned}$$

for any $s_1, \dots, s_m, t_1, \dots, t_\ell \in \mathbb{N}$, with $y_1, \dots, y_m, z_1, \dots, z_\ell \notin \{x_1, \dots, x_n\}$ and $\{y_1, \dots, y_m\} \cap \{z_1, \dots, z_\ell\} = \varnothing$.

Corollary 6 *Let Γ be a set of relations on a set D, $K \subseteq \mathbb{N}$, and $R(x_1, \dots, x_n) \in \langle\Gamma\rangle_K$. Then there is a conjunctive quantifier free formula $\Phi(x_1, \dots, x_n, y_1, \dots, y_m)$ using relations from Γ and the equality relation such that $R(x_1, \dots, x_n) = \exists_{s_1} y_1 \dots \exists_{s_m} y_m \, \Phi(x_1, \dots, x_n, y_1, \dots, y_m)$.*

The following observation summarizes some relationship between the constructions introduced.

Observation 7 *For a set of relations Γ on D, $|D| = m$, the following hold.*

- *Γ is a 1-existential (partial) co-clone if and only if it is a co-clone.*
- *Γ is a (partial) m-existential clone if and only if it is a (partial) co-clone closed under universal quantification.*
- *if Γ is a counting co-clone then it is a max-existential co-clone.*
- *if Γ is a max-existential co-clone then it is a partial m-existential co-clone.*

* 'Counting' in this term comes from counting quantifiers and has nothing to do with counting constraint satisfaction.

In all other cases the introduced versions of co-clones are incomparable.

Example 8 Fix a natural number m and let D be a set with $\frac{m(m+1)}{2}$ elements. Consider an equivalence relation R_m on D with classes $D_1, \dots, D_m$ such that $|D_i| = i$. Then the co-clone generated by R_m corresponds to one of the Rosenberg's maximal clones [24], and so the structure of relations from this co-clone is well understood. For any n-ary relation $Q \in \langle\langle R_m \rangle\rangle$ there is a partition $I_1, \dots, I_k$ of $[n]$ such that a tuple $\mathbf{a}$ belongs to Q if and only if for each $j \in [k]$ and every $i, i' \in I_j$ the entries $\mathbf{a}[i], \mathbf{a}[i']$ are R_m-related. This also means that $\langle R_m \rangle = \langle\langle R_m \rangle\rangle$.

Applying k-existential and max-existential quantifiers one can easily find the k-existential, counting, and max-existential clones generated by R_m:

1. $\langle R_m \rangle_k = \langle\langle R_m \rangle\rangle_k$ is the set of relations Q: There is a partition $I_1, \dots, I_t$ of $[\mathsf{ar}(Q)]$ and $J \subseteq [t]$ such that a tuple $\mathbf{a}$ belongs to Q if and only if for each $j \in [t]$ and every $i, i' \in I_j$ the entries $\mathbf{a}[i], \mathbf{a}[i']$ are R_m-related and $\mathbf{a}[i] \in D_k \cup \ldots \cup D_m$ for $i \in I_j$, $j \in J$.
2. $\langle\langle R_m \rangle\rangle_\infty$ is the set of relations Q: There is a partition $I_1, \dots, I_t$ of $[\mathsf{ar}(Q)]$ and a function $\varphi : [t] \to [m]$ such that a tuple $\mathbf{a}$ belongs to Q if and only if for each $j \in [t]$ and every $i, i' \in I_j$ the entries $\mathbf{a}[i], \mathbf{a}[i']$ are R_m-related and $\mathbf{a}[i] \in D_{\varphi(j)} \cup \ldots \cup D_m$ for $i \in I_j$, $j \in J$.
3. $\langle R_m \rangle_{\max} = \langle R_m \rangle^1_{\max}$ is the set of relations Q: There is a partition $I_1, \dots, I_t$ of $[\mathsf{ar}(Q)]$ and $J \subseteq [t]$ such that a tuple $\mathbf{a}$ belongs to Q if and only if for each $j \in [t]$ and every $i, i' \in I_j$ the entries $\mathbf{a}[i], \mathbf{a}[i']$ are R_m-related and $\mathbf{a}[i] \in D_m$ for $i \in I_j$, $j \in J$.

A set Γ such that $\langle \Gamma \rangle_k \neq \langle\langle \Gamma \rangle\rangle_k$ can be easily found among usual weak co-clones. For instance, for any weak co-clone Γ that is not a co-clone we have $\langle \Gamma \rangle_1 \neq \langle\langle \Gamma \rangle\rangle_1$. Such a weak co-clone can be found in, say, [18].

In the example given we have $\langle R_m \rangle^1_{\max} = \langle R_m \rangle_m$. However, since $\langle R_{m-1} \rangle_m = \langle R_{m-1} \rangle$, we have $\langle R_{m-1} \rangle^1_{\max} \neq \langle R_{m-1} \rangle_m$. An example distinguishing between $\langle \Gamma \rangle_{\max}$ and $\langle \Gamma \rangle^1_{\max}$ has been mentioned above.

We give a sketchy proof of (1) here, the remaining results are similar. Let $Q(x_1, \dots, x_n)$ satisfies the conditions in (1) for a partition $I_1, \dots, I_t$ of $[n]$ and $J \subseteq [t]$. Without loss of generality assume $J = [s]$, $s \leq t$. Choose variables $y_1, \dots, y_s \notin \{x_1, \dots, x_n\}$ and consider relation $S(x_1, \dots, x_n, y_1, \dots, y_s)$ given by: $\mathbf{a} \in S$ if and only if $(\mathbf{a}[i], \mathbf{a}[j]) \in R_m$ for any $i, j \in I_\ell$ and any $\ell \in [t]$ and $(\mathbf{a}[i], \mathbf{a}[n + \ell]) \in R_m$ for any $i \in I_\ell$ where $\ell \in J$. Clearly, $S \in \langle R_m \rangle = \langle\langle R_m \rangle\rangle$. Now, as it is easy to see,

$$Q(x_1, \dots, x_n) = \exists_k y_1 \dots \exists_k y_s \, S(x_1, \dots, x_n, y_1, \dots, y_s).$$

In order to show that every relation from $\langle\langle R_m \rangle\rangle_k$ satisfies these conditions, it suffices to prove that the set of relations Γ satisfying them is closed under manipulations with variables, conjunction, existential quantification, and k-existential quantification. The first three operations are easy, since Γ is a co-clone generated by R_m and unary relation $D' = D_k \cup \ldots \cup D_m$. Let $Q(x_1, \dots, x_n) \in \Gamma$ and $S(x_1, \dots, x_{n-1}) = \exists_k x_n \, Q(x_1, \dots, x_n)$. Let also $I_1, \dots, I_t$ and $J \subseteq [t]$ be the partition and the set from conditions (1). We may assume $n \in I_t$. Then if $t \in J$ then $S(x_1, \dots, x_{n-1}) = \exists x_n \, Q(x_1, \dots, x_n)$. Otherwise $\mathbf{a} \in S$ if and only if (a) for any $i, j \in I_\ell$, $\ell < t$, we have $(\mathbf{a}[i], \mathbf{a}[j]) \in R_m$, (b) for any $i, j \in I'_t = I_t - \{n\}$, we have $(\mathbf{a}[i], \mathbf{a}[j]) \in R_m$, and (c) $\mathbf{a}[i] \in D'$, whenever $i \in I'_t \cup \bigcup_{s \in J} I_s$. Therefore $S \in \langle\langle R_m \rangle\rangle_k$.

In some fairly general cases usual co-clones and max-co-clones coincide.

Recall that an (n-ary) relation is said to be *affine* if it can be represented as the set of solutions to a system of linear equations over a finite field $GF(p^m)$, m prime. The next lemma follows from basic linear algebra, as sets of extensions of tuples are cosets of the same vector subspace.

Lemma 9 *Let R be an (n-ary) affine relation. Then for any $I \subseteq [n]$ any two tuples $\mathbf{a}, \mathbf{b} \in \mathrm{pr}_I R$ have the same number of extensions to tuples from R.*

Proof: Let R be the set of solutions of a system of linear equations $A \cdot \mathbf{x} = \mathbf{c}$, where A is a $\ell \times n$-matrix over $GF(m)$, $\mathbf{x} = (x_1, \ldots, x_n)^\top$, and $\mathbf{c} \in \{0,1\}^\ell$. Without loss of generality $I = [k]$. Then A can be represented as $A = [A_1 \mid A_2]$, where A_1 is a $\ell \times k$-matrix and A_2 is a $\ell \times (n-k)$-matrix; $\mathbf{x}$ can be represented as $\mathbf{x} = (\mathbf{x}^1, \mathbf{x}^2)^\top$, where $\mathbf{x}^1 = (x_1, \ldots, x_k)$, $\mathbf{x}^2 = (x_{k+1}, \ldots, x_n)$. Fix $\mathbf{a} \in \mathrm{pr}_{[k]} R$ and set $\mathbf{c}_\mathbf{a} = \mathbf{c} \oplus (A_1 \cdot \mathbf{a})$. The set of extensions of $\mathbf{a}$ is the set of solutions of the system $A_2 \cdot \mathbf{x}^2 = \mathbf{c}_\mathbf{a}$. Clearly, the number of solutions of this system does not depend on $\mathbf{a}$, provided the system is consistent. □

Corollary 10 *Let Γ be a set of affine relations over a finite field $GF(m)$. Then $\langle\langle\Gamma\rangle\rangle = \langle\Gamma\rangle_{\max}$.*

Proof: Lemma 9 implies that for any (n-ary) affine relation R and any set $J = \{i_1, \ldots, i_k\} \subseteq [n]$ the max-implementation $\exists_{\max}(x_{i_1}, \ldots, x_{i_k})$ is equivalent to a sequence of ordinary existential quantifiers $\exists x_{i_1} \ldots \exists x_{i_k}$. □

This result can be generalized using certain notions from universal algebra (see [12] for definitions and details). A ternary operation $f(x,y,z)$ on a set A is said to be *Mal'tsev* if it satisfies the identities $f(x,y,y) = f(y,y,x) = x$. An algebra is said to be Mal'tsev if it has a Mal'tsev term operation. An algebra $\mathbb{A}$ is called *congruence uniform* if for any congruence α of any subalgebra $\mathbb{A}$ the blocks of α have the same cardinality.

Proposition 11 *Let Γ be a set of relations invariant under operations of a Mal'tsev congruence uniform algebra. Then $\langle\langle\Gamma\rangle\rangle = \langle\Gamma\rangle_{\max}$.*

Proof: Denote the algebra mentioned in the proposition by $\mathbb{A}$. Every relation from Γ, as well as, every relation from $\langle\langle\Gamma\rangle\rangle$ is invariant with respect to the Mal'tsev term operation of $\mathbb{A}$. A folklore result (in the form we need it here it is observed in, for example, [8]). Let R be an (n-ary) relation. An (n-ary) relation R is said to be *rectangular* if, for any partition of $[n]$ into subsets I, J, and any $\mathbf{a}, \mathbf{b} \in \mathrm{pr}_I R$, $\mathbf{c}, \mathbf{d} \in \mathrm{pr}_J R$, if $(\mathbf{a}, \mathbf{c}), (\mathbf{a}, \mathbf{d}), (\mathbf{b}, \mathbf{c}) \in R$, then $(\mathbf{b}, \mathbf{d}) \in R$ (here $(\mathbf{a}, \mathbf{c})$ denotes the tuple $\mathbf{e}$ such that $\mathbf{e}[i] = \mathbf{a}[i]$ if $i \in I$ and $\mathbf{e}[i] = \mathbf{c}[i]$ if $i \in J$). The rectangularity of R implies that there are partitions $A_1, \ldots, A_s$ of $\mathrm{pr}_I R$ and $B_1, \ldots, B_s$ of $\mathrm{pr}_J R$ such that

$$R = \bigcup_{i=1}^{s} (A_i \times B_i).$$

Observe that if we prove that for any relation R invariant under operations of a Mal'tsev congruence uniform algebra and any partition I, J of the set of its coordinate positions, that $|B_1| = \ldots = |B_s|$, then we obtain a statement similar to Lemma 9, and the result follows in the same way as in the proof of Corollary 10.

We prove by induction on $|J|$. If $|J| = 1$ the required property follows from the fact that $B_1, \ldots, B_s$ are blocks of a certain congruence of $\mathbb{A}$. Suppose that the property is true for any J with $|J| = m$. Without loss of generality assume $n \in J$. Consider relation $\mathrm{pr}_{[n-1]} R$ and the partition of $[n-1]$ into $I, J - \{n\}$. The corresponding partitions of $\mathrm{pr}_I R$ and $\mathrm{pr}_{J-\{n\}} R$ are $A_1, \ldots, A_s$ and $B'_1, \ldots, B'_s$ where $B'_i = \mathrm{pr}_{J-\{n\}} B_i$. By induction hypothesis $|B'_1| = \ldots = |B'_s| = r$. Now take another partition of $[n]$, namely, $[n-1], n$. We again have some partition $C_1, \ldots, C_t$ of $\mathrm{pr}_n R$ such that $|C_1| = \ldots = |C_t| = q$. Therefore every tuple $\mathbf{a} \in \mathrm{pr}_{[n-1]} R$ can be extended to a tuple from R in q ways, and every tuple from $\mathrm{pr}_I R$ can be extended to a tuple from $\mathrm{pr}_{[n-1]} R$ in r ways. Hence, $|B_1| = \ldots = |B_s| = rq$. □

5 GALOIS CORRESPONDENCE

Let D be a finite set. A (partial) function $f\colon D^n \to D$ is said to be *k-subset surjective* if for any k-element subsets $A_1, \ldots, A_n \subseteq D$ the image $f(A_1, \ldots, A_n)$ has cardinality at least k. A (partial) function that is k-subset surjective

for each k, $1 \le k \le |D|$ is said to be *subset surjective*. The set of all arity n k-subset surjective partial functions [arity n k-subset surjective functions, subset surjective functions] on D will be denoted by $P_D^{k,(n)}$ [resp., $F_D^{k,(n)}$, $F_D^{(n)}$]; furthermore, $P_D^k = \bigcup_{n\ge 0} P_D^{k,(n)}$, $F_D^k = \bigcup_{n\ge 0} F_D^{k,(n)}$, $F_D = \bigcup_{n\ge 0} F_D^{(n)}$. Any total function is 1-subset surjective, while $|D|$-subset surjective partial functions are exactly the surjective partial functions. Observe that this definition can be strengthened by allowing the sets A_i, $i \in [n]$, to have at least k elements.

Lemma 12 *If an n-ary function f is k-subset surjective, then for any subsets $A_1, \dots, A_n \subseteq D$ with $|A_i| \ge k$, $i \in [n]$, the image $f(A_1, \dots, A_n)$ has cardinality at least k.*

Proof: Choose any $B_i \subseteq A_i$, $i \in [n]$, and set $B = f(B_1, \dots, B_n)$. As f is k-subset surjective, $|B| \ge k$. Finally, $B \subseteq f(A_1, \dots, A_n)$, and the result follows. □

The conditions of being k-subset surjective for different k are incomparable, as the following example shows.

Example 13 Let $D = \{0, \dots, k-1\}$ be a k-element set and $1 < m \le k$. Then the following function f is not m-subset surjective, but is ℓ-subset surjective for any $\ell \in [k]$, $\ell \ne m$. Function f is binary and is given by its operation table:

$$\begin{pmatrix}
0 & 0 & \cdots & 0 & 1 & m & \cdots & k-1 \\
1 & 1 & \cdots & 1 & 2 & m & \cdots & k-1 \\
\vdots & \vdots & & \vdots & \vdots & \vdots & & \vdots \\
m-3 & m-3 & \cdots & m-3 & m-2 & m & \cdots & k-1 \\
m-2 & m-2 & \cdots & m-2 & 0 & m & \cdots & k-1 \\
0 & 1 & \cdots & m-2 & 0 & m & \cdots & k-1 \\
0 & 1 & \cdots & m-2 & m-1 & m & \cdots & k-1 \\
\vdots & \vdots & & \vdots & \vdots & \vdots & & \vdots \\
0 & 1 & \cdots & m-2 & m-1 & m & \cdots & k-1
\end{pmatrix}.$$

Clearly, f is not m-subset surjective, because $f(B, B) = \{0, \dots, m-2\}$ for $B = \{0, \dots, m-1\}$. Also, as it is a total function, f is 1-subset surjective. Take $\ell \in [k]$, $\ell > 1$, and $B_1, B_2 \subseteq \{0, \dots, k-1\}$ with $|B_1| = |B_2| = \ell$. If there is $a \in B_1$ with $i \ge m$ then $f(a, b_1) \ne f(a, b_2)$ whenever $b_1 \ne b_2$. This means that $|f(B_1, B_2)| \ge \ell$ in this case, and, in particular, f is ℓ-subset surjective for any $\ell > m$. So, suppose $\ell < m$ and $B_1 \subseteq \{0, \dots, m-1\}$. If $B_1 \subseteq \{0, \dots, m-2\}$ then take $b \in B_2 \cap \{0, \dots, m-1\}$ and observe that $f(a_1, b) \ne f(a_2, b)$ for any $a_1, a_2 \in \{0, \dots, m-1\}$, $a_1 \ne a_2$. Thus, $|f(B_1, \{b\})| = \ell$. Suppose $m-1 \in B_1$. If $m-1 \notin B_2$, then $|f(m-1, B_2)| = \ell$; assume $m-1 \in B_2$. As is easily seen, $B_1 \cap \{0, \dots, m-2\} \subseteq f(B_1, B_2)$. There is $a \in \{0, \dots, m-2\}$ such that $a \notin B_1$ but $a - 1 \pmod{m-1} \in B_1$. Then $a \in f(B_1, B_2)$, since $a = f(a-1, m-1)$. Thus, $|f(B_1, B_2)| \ge \ell$.

The notion of invariance for k-subset surjective functions is the standard one for partial functions and relations. As usual, if C is a set of (k-) subset surjective (partial) functions, $\mathsf{Inv}(C)$ denotes the set of relations invariant with respect to every function from C. For a set Γ of relations, $\mathrm{m}(k)\text{–}\mathsf{Pol}(\Gamma)$ and $\mathrm{m}(k)\text{–}\mathsf{pPol}(\Gamma)$ denote the set of all k-subset surjective functions and partial functions, respectively, preserving every relation from Γ. For a set $K \subseteq \mathbb{N}$ by $\mathrm{m}(K)\text{–}\mathsf{Pol}(\Gamma)$ and $\mathrm{m}(K)\text{–}\mathsf{pPol}(\Gamma)$ we denote the set of all functions and, respectively, partial functions preserving every relation from Γ that are k-subset surjective for each $k \in K$. Thus, in particular,

$$\mathrm{m}(K)\text{–}\mathsf{Pol}(\Gamma) = \bigcap_{k\in K} \mathrm{m}(k)\text{–}\mathsf{Pol}(\Gamma), \quad \text{and} \quad \mathrm{m}(K)\text{–}\mathsf{pPol}(\Gamma) = \bigcap_{k\in K} \mathrm{m}(k)\text{–}\mathsf{pPol}(\Gamma).$$

By $\mathrm{m}\text{–}\mathsf{Pol}(\Gamma)$ we denote the analogous set of subset surjective functions.

The operator Inv on one side and the operators $\mathrm{m}(k)\text{–}\mathsf{pPol}(\Gamma)$, $\mathrm{m}(k)\text{–}\mathsf{Pol}(\Gamma)$, $\mathrm{m}(K)\text{–}\mathsf{Pol}(\Gamma)$, $\mathrm{m}\text{–}\mathsf{pPol}(\Gamma)$, $\mathrm{m}\text{–}\mathsf{Pol}(\Gamma)$ on the other side form Galois correspondences in the standard fashion. We characterize closed sets of relations that give rise from this correspondence.

Lemma 14 *Let $R(x_1,\dots,x_\ell,y)$ be a relation on D, and let $Q(x_1,\dots,x_\ell)=\exists_k y R(x_1,\dots,x_\ell,y)$. Then if a k-subset surjective (partial) function f preserves R, it also preserves Q.*

Proof: Suppose f is n-ary. Take $\mathbf{a}_1,\dots,\mathbf{a}_n\in Q$. Since each of them is put into Q by k-existential quantification, it has at least k extensions to a tuple from R. Let $B_1,\dots,B_n\subseteq D$ be such that $|B_i|\geq k$ and $(\mathbf{a}_i,b)\in R$ for $b\in B_i$ and $i\in[n]$. Let also $\mathbf{b}=f(\mathbf{a}_1,\dots,\mathbf{a}_n)$. For any $b\in B=f(B_1,\dots,B_n)$ the tuple $(\mathbf{b},b)$ belongs to R. As f is k-subset surjective, $|B|\geq k$, hence, $\mathbf{b}\in Q$. □

Theorem 15 *Let Γ be a set of relations on a set D and $K\subseteq\mathbb{N}$. Then* $\mathsf{Inv}(\mathrm{m}(K)\text{−}\mathsf{pPol}(\Gamma))=\langle\Gamma\rangle_K$.

Proof: We will assume that $K=\{k_1,\dots,k_s\}\subseteq\{1,\dots,|D|\}$. Indeed, if $k\geq|D|$ then $\exists_k xR$ is empty for any relation on D. The equality relation, $=_D$, is invariant with respect to any partial function on D. Let f be a k-subset surjective functions. It is straightforward to verify that manipulations of variables of a predicate invariant under f and the conjunction of any two predicates invariant under f result in predicates invariant under f, again, since it is true for any partial function. By Lemma 14 applying k-quantification to a predicate invariant under f gives a predicate invariant under f, again because it is true for any partial function. Hence, $\langle\Gamma\rangle_K\subseteq\mathsf{Inv}\,(\mathrm{m}(K)\text{−}\mathsf{pPol}(\Gamma))$. Moreover, it follows that $\mathsf{Inv}\,(\mathrm{m}(K)\text{−}\mathsf{pPol}(\Gamma))=\mathsf{Inv}\,(\mathrm{m}(K)\text{−}\mathsf{pPol}(\langle\Gamma\rangle_K))$.

To establish the reverse inclusion, take an ℓ-ary relation $R\in\mathsf{Inv}(\mathrm{m}(K)\text{−}\mathsf{pPol}(\Gamma))$. We need to show that $R\in\langle\Gamma\rangle_k$. Define a relation Q as follows. Let $R=\{\mathbf{a}_1,\dots,\mathbf{a}_t\}$. For each $k\in K$ we consider sequences $(B_1,\dots,B_t)$ of k-element subsets of D. Let also $(B_1^{k1},\dots,B_t^{k1}),\dots,(B_1^{kr_k},\dots,B_t^{kr_k})$ be a list of all such sequences. Let S_k^j be the relation

$$\underbrace{B_j^{k1}\times\ldots\times B_j^{k1}}_{k\text{ times}}\times\ldots\times\underbrace{B_j^{kr_k}\times\ldots\times B_j^{kr_k}}_{k\text{ times}},$$

and $S^j=S^j_{k_1}\times\ldots\times S^j_{k_s}$. Then Q is the union of relations $\mathbf{a}_j\times S^j$, for all $j\in[t]$. We show that there is $S\in\langle\Gamma\rangle_k$ such that $Q\subseteq S$ and $\mathrm{pr}_{[\ell]}S=R$. Then applying k-quantifications, $k\in K$, to all coordinates of S except for the first ℓ we infer that $R\in\langle\Gamma\rangle_K$.

Set $M=\sum_{k\in K}kr_k$ and $M_j=\sum_{i=1}^{j}k_ir_{k_i}$; by N_K, $k\in K$, we denote the set $\{M_j+1,\dots,M_{j+1}\}$. Let us consider the relation $S=\bigcap\{Q'\in\langle\Gamma\rangle_K\mid Q\subseteq Q'\}$. Since $\langle\Gamma\rangle_K$ is closed under conjunctions and contains the total relation $D^{\ell+M}$, we have $S\in\langle\Gamma\rangle_K$ and $Q\subseteq S$.

Now choose any tuple $\mathbf{b}=(b_1,\dots,b_\ell,d_1,\dots,d_M)\in S$. There are sets $C_1,\dots,C_M$ such that $|C_i|=k_j$, $i\in[M]$, whenever $i\in N_j$, for any $t\in[r_j]$, $C_{M_{j-1}+k_j(t-1)+1}=\ldots=C_{M_{j-1}+k_jt}$, $d_i\in C_i$, and for any $d_i'\in C_i$, $i\in[M]$, the tuple $(b_1,\dots,b_\ell,d_1',\dots,d_M')\in S$. Indeed, otherwise we can apply a sequence of k-quantifications for $k\in K$ to obtain an ℓ-ary relation S' containing R, but not $(b_1,\dots,b_\ell)$. Then , $(S'\times D^{\ell+M})\cap Q$ belongs to $\langle\Gamma\rangle_K$, but is smaller than S. Therefore we can choose $\mathbf{b}$ such that for any $j\in[s]$ and any $t\in[r_j]$ all the values $d_{M_{j-1}+k_j(t-1)+1},\dots,d_{M_{j-1}+k_jt}$ are distinct, and $\{d_{M_{j-1}+k_j(t-1)+1},\dots,d_{M_{j-1}+k_jt}\}=C_{M_{j-1}+k_jt}$.

Since $\langle\Gamma\rangle_K$ is closed under conjunctions, by the Fleischer and Rosenberg result [16] it satisfies $\langle\Gamma\rangle_K=\mathsf{Inv}(\mathsf{pPol}(\langle\Gamma\rangle_K))$. Moreover, by the proof of Theorem 2 of [16] S is the set of all tuples of the form $f(\mathbf{c}_1,\dots,\mathbf{c}_n)$ for $n\geq1$, $\mathbf{c}_1,\dots,\mathbf{c}_n\in Q$, and $f\in\mathsf{pPol}(\langle\Gamma\rangle_K)$. Therefore there exist $n\geq1$, $\mathbf{c}_1,\dots,\mathbf{c}_n\in Q$, and $f\in\mathsf{pPol}(\langle\Gamma\rangle_K)$ such that $\mathbf{b}=f(\mathbf{c}_1,\dots,\mathbf{c}_n)$. Let $\mathrm{pr}_{[\ell]}\mathbf{c}_q=\mathbf{a}_{i_q}$. For any selection $E_1,\dots,E_n$ of k_j-element subsets of D, $j\in[s]$, there is $t\in[r_{k_j}]$ such that $E_q=B_{i_q}^{k_jt}$ for $q\in[n]$. By the choice of $\mathbf{b}$ the range of f on $E_1\times\ldots\times E_n=B_{i_1}^{k_jt}\times\ldots\times B_{i_n}^{k_jt}$ contains $C_{M_{j-1}+k_jt}$. Hence f is k_j-subset surjective for any $k_j\in K$, and so $f\in\mathrm{m}(K)\text{−}\mathsf{pPol}(\Gamma)$, as it is equal to $\mathrm{m}(K)\text{−}\mathsf{pPol}(\langle\Gamma\rangle_k)$. Therefore R is invariant under f, and so $(b_1,\dots,b_\ell)\in R$. Relation S satisfies the required conditions, which completes the proof. □

Corollary 16 *There is a Galois correspondence between K-existential partial co-clones on one side and partial clones generated by K-surjective partial functions on the other side.*

More precisely, for any set Γ of relations on D, any $K \subseteq \{1, \ldots, |D|\}$, and any set C of K-surjective partial functions on D,

- $\mathsf{Inv}(C)$ *is a K-existential partial co-clone;*
- $\mathsf{pPol}(\langle\Gamma\rangle_K)$ *is a partial co-clone generated by the set* $\mathrm{m}(K)\text{−}\mathsf{pPol}(\langle\Gamma\rangle_K)$ *of K-surjective partial functions;*
- $\mathsf{Inv}(\mathrm{m}(K)\text{−}\mathsf{pPol}(\Gamma)) = \langle\Gamma\rangle_K$*;*
- $\mathrm{m}(K)\text{−}\mathsf{pPol}(\mathsf{Inv}(C))$ *is the set of K-surjective functions from the partial clone generated by C.*

Corollary 17 *Let Γ be a set of relations on a set D.*

(a) $\mathsf{Inv}(\mathrm{m}(k)\text{−}\mathsf{pPol}(\Gamma)) = \langle\Gamma\rangle_k$*;*

(b) $\mathsf{Inv}(\mathrm{m}(k)\text{−}\mathsf{Pol}(\Gamma)) = \langle\langle\Gamma\rangle\rangle_k$*;*

(c) $\mathsf{Inv}(\mathrm{m}\text{−}\mathsf{Pol}(\Gamma)) = \langle\langle\Gamma\rangle\rangle_\infty$*;*

REFERENCES

[1] L. Barto and M. Kozik. Constraint satisfaction problems of bounded width. In *FOCS*, pages 595–603, 2009.

[2] V.G. Bodnarchuk, L.A. Kaluzhnin, V.N. Kotov, and B.A. Romov. Galois theory for Post algebras. I. *Kibernetika*, 3:1–10, 1969.

[3] F. Börner, A. Bulatov, H. Chen, P. Jeavons, and Andrei A. Krokhin. The complexity of constraint satisfaction games and QCSP. *Inf. Comput.*, 207(9):923–944, 2009.

[4] A. Bulatov. Tractable conservative constraint satisfaction problems. In *LICS* 2003, pages 321–330 .

[5] A. Bulatov. A dichotomy theorem for constraint satisfaction problems on a 3-element set. *J. ACM*, 53(1):66–120, 2006.

[6] A. Bulatov. The complexity of the counting constraint satisfaction problem. In *ICALP (1)*, pages 646–661, 2008.

[7] A. Bulatov. Counting problems and clones of functions. In *ISMVL*, pages 1–6, 2009.

[8] A. Bulatov and V. Dalmau. Towards a dichotomy theorem for the counting constraint satisfaction problem. *Inf. Comput.*, 205(5):651–678, 2007.

[9] A. Bulatov, P. Jeavons, and A.A. Krokhin. Classifying the complexity of constraints using finite algebras. *SIAM J. Comput.*, 34(3):720–742, 2005.

[10] A. Bulatov and A. Hedayaty. Counting quantifiers, subset surjective functions, and counting CSPs. In *ISMVL*, pages 331-336, 2012.

[11] A. Bulatov and A. Hedayaty. Galois correspondence for counting quantifiers. *CoRR*, abs/1210.3344, 2012.

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

[13] H. Chen. The complexity of quantified constraint satisfaction: Collapsibility, sink algebras, and the three-element case. *SIAM J. Comput.*, 37(5):1674–1701, 2008.

[14] A. Dawar and D. Richerby. The power of counting logics on restricted classes of finite structures. In *CSL*, pages 84–98, 2007.

[15] K. Etessami. Counting quantifiers, successor relations, and logarithmic space. *J. Comput. Syst. Sci.*, 54(3):400–411, 1997

[16] I. Fleischner and I.G. Rosenberg. The Galois connection between partial operations and relations. *Pacific J. Math.*, 79:93–97, 1978.

[17] D. Geiger. Closed systems of function and predicates. *Pacific Journal of Mathematics*, pages 95–100, 1968.

[18] L. Haddad and I.G. Rosenberg. Partial clones containing all permutations. *Bull. Austral. Math. Soc.*, 52:263–278, 1995.

[19] P.M. Idziak, P. Markovic, R. McKenzie, M. Valeriote, and R. Willard. Tractability and learnability arising from algebras with few subpowers. *SIAM J. Comput.*, 39(7):3023–3037, 2010.

[20] N. Immerman and E. Lander. Describing graphs: a first-order approach to graph canonization. In *Complexity Theory Retrospective*, Alan Selman, ed., Springer-Verlag, 1990, pages 59-81.

[21] P.G. Jeavons, D.A. Cohen, and M. Gyssens. Closure properties of constraints. *Journal of the ACM*, 44:527–548, 1997.

[22] U. Montanari. Networks of constraints: Fundamental properties and applications to picture processing. *Inf. Sci.*, 7:95–132, 1974.

[23] R. Pöschel and L.A. Kalužnin. *Funktionen- und Relationenalgebren*. DVW, Berlin, 1979.

[24] I.G. Rosenberg. Über die funktionale Vollständikeit in dem mehrwertigen Logiken. *Rozpravy Čs. Akademie Věd. Ser. Math. Nat. Sci.*, 80:3–93, 1970.

[25] T. Schaefer. The complexity of satisfiability problems. In *STOC*, pages 216–226, 1978.