A NEW TRACTABLE CLASS OF CONSTRAINT SATISFACTION PROBLEMS

VÍCTOR DALMAU

Departament de Tecnologia, Universitat Pompeu Fabra
Estació de França, Passeig de la circumval.lació, Barcelona 08003, Spain
victor.dalmau@upf.edu

Abstract. In this paper we consider constraint satisfaction problems where the set of constraint relations is fixed. Feder and Vardi (1998) identified three families of constraint satisfaction problems containing all known polynomially solvable problems. We introduce a new class of problems called *para-primal* problems, incomparable with the families identified by Feder and Vardi (1998) and we prove that any constraint problem in this class is decidable in polynomial time. As an application of this result we prove a complete classification for the complexity of constraint satisfaction problems under the assumption that the basis contains all the permutation relations. In the proofs, we make an intensive use of algebraic results from clone theory about the structure of para-primal and homogeneous algebras.

Keywords: Constraint satisfaction problem, complexity, para-primal algebra.
AMS Classification: 08A70.

1. Introduction

The constraint satisfaction problem provides a framework in which it is possible to express, in a natural way, many combinatorial problems encountered in artificial intelligence and elsewhere. A constraint satisfaction problem is represented by a set of variables, a domain of values for each variable, and a set of constraints between variables. The aim of a constraint satisfaction problem is then to find an assignment of values to the variables that satisfies the constraints.

Since solving a general constraint satisfaction problem is known to be NP-complete [10, 26], it arises naturally the important question of identifying restrictions to the general problem which are sufficient to ensure tractability. Such restrictions may either involve the scope in the constraints, in other words, which variables may be constrained with other variables [14, 16, 17, 30, 31], or they may involve the relation in the constraints, in other words, which combinations of values are allowed for variables which are mutually constrained [11, 21, 24, 25, 30, 35, 36].

In this paper, we focus on the second approach. More precisely we are interested in characterizing the complexity of solving constraint satisfaction problems in which every constraint relation belongs to a fixed set called *basis*. This broad class includes problems such as K-SAT or GRAPH K-COLORABILITY.

Jeavons, Cohen and Gyssens in [21, 22] introduced a novel framework for studying this class of constraint satisfaction problems. The approach started in [21, 22] relies on the fact that the tractability of constraint satisfaction problems using certain relations depends on some algebraic condition (that of closure) of the set of relations used to build the problem.

This approach has lead to the identification of several tractable families of problems or basis. The exhaustive list of tractable families known so far in the framework of closure functions is: (**1**) near-unanimity problems [20], (**2**) coset generating problems [15, 12], including the class of affine problems [23], and (**3**) width 1 problems [13] including some important class of problems such as constant problems and ACI problems [23]. These three classes were identified originally by Feder and Vardi [15], using a different approach. In [15] this three classes were called respectively bounded strict width, subgroup and width 1 problems.

In this paper we increase the previous list with a new family of tractable basis that we call *para-primal* basis, incomparable with the previous three classes.

This family of basis is tightly related with a family of algebras that is well known in Universal algebra: para-primal algebras. In this paper we make an intensive use of algebraic results. So far, the link between constraint satisfaction problem with clones in universal algebra has been used mostly to establish an unifying point of view. Nevertheless, a genuine application of results in clone theory to constraint satisfaction problem is rather rare. In fact, with the exception of the use of Post's classification to give an alternative and simpler proof of Schaefer's dichotomy theorem in [23], the few results in clone theory used in the constraint literature can be easily reformulated and proved in constraint terms. The results of this paper constitute a genuine application of clone theory to constraint satisfaction problems and a clear evidence of the power of the algebraic-based approach in the study of constraint satisfaction problems. On the other hand, due precisely to the profusion of algebraic concepts and results, some familiarity with algebraic notions would be advisable to follow some sections of the paper.

The definition of para-primal basis is rather technical and, in fact, a bit involved. Nevertheless, it is possible to give a natural interpretation of para-primal basis in terms of "simulated relations" which might provide some intuitive grasp of the concept. We illustrate the concept of "simulated relation" by means of an example. Assume that our basis, or constraint language contains only 3-clauses. It is well-known that it is possible to express any clause of arbitrary size by a conjunction of 3-clauses by adding some auxiliary variables. We say that clauses of size, say 4, can be "simulated" by clauses of size 3. That is, although 4-clauses are not part of the basis and, henceforth, cannot be included directly into our formula, it is possible to built a subformula or *gadget* that expresses any 4-clause (and in fact every possible clause). In general, given a constraint language, the class of relations that can be "simulated" characterizes faithfully the "expressiveness power" of the constraint language. The notion of "relation simulated" has been more or less implicitly used in most of the constraint literature and formally defined in [23]. More precisely, for a fixed set of relations Γ, the relations "simulated" by Γ is defined as the smallest set containing Γ, and closed under cartesian product, identification and projection (see Section 4). Para-primal basis can be entirely described in terms of the binary relations simulated. More precisely, we show that the binary relations "simulated"by para-primal basis are either: (1) a direct product of two subsets of the domain, or (2) a bijection between two subsets of the domain. In fact, under some additional technical assumptions, the previous property characterizes exactly para-primal basis. Thus, tractability of para-primal algebras can be restated in the following way: for any arbitrary basis Γ, if the binary relations "simulated" by Γ are "simple enough" then its associated constraint satisfaction problem is efficiently solvable.

An alternative interpretation is obtained considering it as an extension of the subgroup problems introduced by Feder and Vardi [15]. In a subgroup problem every relation in the basis is a coset of some subgroup of some group power for a fixed group. Feder and Vardi proposed a way to extend the previous definition to include relations others than cosets. More precisely, they introduced a generalization of the concept of coset, called nearsubgroup, and they proved that for several families of groups, this extended problem is still efficiently solvable. Para-primal basis can be regarded as a different extension of subgroup problems. Nevertheless, in para-primal algebras the approach relies in combining different subgroup problems rather than extending a unique subgroup problem. In a para-primal problem, different subgroup problems are combined in a non-trivial way but such that the interaction between them is bounded and the problem can be solved by decomposing the problem in subproblems corresponding to every one of the subgroup problems.

As an application of the tractability of para-primal algebras we prove a complete classification for constraint satisfaction problems under the assumption that the basis contains all the permutations. This result generalizes a known previous result due to Cooper et al. [11] in several ways.

As discussed in Section 8 the state of affairs has changed since this paper was originally written. In particular, the list of tractable cases of CSP has grown since then and, henceforth, the list of tractable cases described in this introduction is no longer exhaustive.

2. Constraint Satisfaction Problems

A constraint satisfaction problem is a natural way to express simultaneous requirements for values of variables. More precisely,

Definition 1. An instance of a *constraint satisfaction problem* consists of:

- a finite set of variables, V;
- a finite domain of values, D;
- a finite set of constraints $\{C_1, \ldots, C_q\}$; each constraint C_i $(1 \leq i \leq q)$ is a pair (s_i, R_i) where:
  - s_i is a tuple of variables of length k_i, called the *constraint scope*; and
  - R_i is an k_i-ary relation over D, called the *constraint relation*.

For each constraint (s_i, R_i), the tuples in R_i indicate the allowed combinations of simultaneous values for the variables in s_i. The length of s_i, and of the tuples in R_i, is called the *arity* of the constraint.

A *solution* to a constraint satisfaction problem instance is a function from the variables to the domain such that the image of each constraint scope is an element of the corresponding constraint relation. Deciding whether or not a given problem instance has a solution is NP-complete in general, even when the constraints are restricted to binary constraints [26] or the domain of the problem has size 2 [10]. However, by imposing restrictions on the constraint interconnections (see [14, 16, 17, 30, 31]), or the form of the constraints (see [11, 21, 24, 25, 30, 35, 36]), it is possible to obtain restricted versions of the problem that are tractable. As we have said, the aim underlying this research is to determine all the possible restrictions on the form of the constraints that ensure tractability. The central object of study is the set of problems defined from some fixed set (or basis) of relations, and the associated complexity of the constraint problems in that set.

Definition 2. For any set of relations Γ, CSP(Γ) is defined to be the class of decision problems with:

- Instance: A constraint satisfaction problem instance $\mathcal{P}$, in which all constraint relations are elements of Γ.
- Question: Does $\mathcal{P}$ have a solution?

3. Some Algebraic Terminology

First, we introduce some concepts that will be used intensively in this paper (See [29] for further details).

Definition 3. *An* algebra $\mathbf{A}$ *is an ordered pair* $(A; F)$ *such that* A *is a nonempty set, called the* universe, *and* $F = \langle f_i : i \in I \rangle$ *where* f_i *is a finitary operation on* A *for each* $i \in I$. A *is called the* universe *of* $\mathbf{A}$, f_i *is referred* *to as a* fundamental *or* basic operation *of* $\mathbf{A}$ *for each* $i \in I$*, and* I *is called the* index set *or the* set of operations symbols *of* $(A; F)$.

Endowing algebras with indexed systems of operations rather than a mere set of operations has many advantages. Nevertheless, sometimes, it is most acceptable to display algebras following the customary viewpoint. For example, $(A; \cdot, 1, ^{-1})$ is a more comfortable way to define a group than $(A; \langle f_i : i \in \{\cdot, 1, ^{-1}\}\rangle)$.

If ϱ is an operation symbol of $\mathbf{A}$, then we use $\varrho^{\mathbf{A}}$ to stand for the fundamental operation of $\mathbf{A}$ indexed by ϱ.

Given an algebra $\mathbf{A}$ with index set I, there is a function Rank called the *rank function* from I to the set of natural numbers defined by:

$$\text{Rank}(\varrho) \text{ is the rank of } \varrho^{\mathbf{A}} \text{ for all } \varrho \in I$$

Algebras $\mathbf{A}$ and $\mathbf{B}$ are said to be *similar* if and only if they have the same rank function.

There are some standard ways to assemble new algebras from those already at hand. The chief tools we will use are the formation of subuniverses and the formation of direct powers.

Let f be an m-ary operation over A, and let B be a subset of A. We say that B is closed under f (or f preserves B) if and only if

$$\forall a_1, \ldots, a_m \in B, \quad f(a_1, \ldots, a_m) \in B.$$

Definition 4. *Let* $\mathbf{A} = (A; F)$ *be an algebra. A subset* B *of* A *which is closed under every operation in* F *is called a* subuniverse *of* $(A; F)$*. The algebra* $\mathbf{B}$ *is said to be a* subalgebra *of* $\mathbf{A}$ *if and only if* $\mathbf{A}$ *and* $\mathbf{B}$ *are similar, the universe* B *of* $\mathbf{B}$ *is a subuniverse of* $\mathbf{A}$*, and* $\varrho^{\mathbf{B}}$ *is the restriction to* B *of* $\varrho^{\mathbf{A}}$*, for each operation symbol* ϱ *of* $\mathbf{A}$.

Now consider similar algebras $\mathbf{A}$ and $\mathbf{B}$ and let ϱ be an operation symbol of rank r. A function h from A into B is said to *respect the interpretation of* ϱ if and only if

$$h(\varrho^{\mathbf{A}}(a_1, \ldots, a_r)) = \varrho^{\mathbf{B}}(h(a_1), \ldots, h(a_r))$$

for all $a_1, \ldots, a_r \in A$.

Definition 5. *Let* $\mathbf{A}$ *and* $\mathbf{B}$ *be similar algebras. A function* h *from* A *into* B *is called a* homomorphism *from* $\mathbf{A}$ *into* $\mathbf{B}$ *if and only if* h *respects the interpretation of every operation symbol of* $\mathbf{A}$*. If* h *is bijective then we say that* h *is an* isomorphism*. In this case we say that* $\mathbf{A}$ *and* $\mathbf{B}$ *are isomorphic and we denote it by* $\mathbf{A} \sim \mathbf{B}$*. If* h *is bijective and* $\mathbf{A} = \mathbf{B}$ *then* h *is called an* automorphism.

Let f be a k-ary operation on A and let n be a positive integer. Then operation f^n is a k-ary operation on A^n defined by:

$$f^n(x_1, \ldots, x_k) = \langle f(x_1[1], \ldots, x_k[1]), \ldots, f(x_1[n], \ldots, x_k[n])\rangle$$

Let $\mathbf{A} = (A; F)$ be an algebra and n any positive integer. The *n-th direct power* of $(A; F)$ is the algebra, denoted by $\mathbf{A}^n$, with the same similarity type, with universe A^n such that for every operation symbol ϱ of $\mathbf{A}$, $\varrho^{(\mathbf{A}^n)} = (\varrho^{\mathbf{A}})^n$

The formation of direct powers is the natural way to extend an algebra to tuples of elements. From now on, due to the correspondence between an algebra and its direct powers we will do some notation abuse and for instance we will say that an n-ary relation R over A is closed under a k-ary operation $f : A^k \to A$, meaning that R is closed under f^n. This concept deserves to be stated as a definition.

Definition 6. *Let f be an m-adic operation over A and let R be a n-ary relation over A. We say that f* preserves *R (R* is closed under *f, or f* is a polymorphism of *R) if*

$$\forall t_1, t_2, \ldots, t_m \in R, \quad f^n(t_1, t_2, \ldots, t_m) \in R.$$

4. Relationship with Constraint Satisfaction Problems

The notion of preserving operation is the key to study the complexity of CSP(Γ). Jeavons [19] showed that the complexity of constraint satisfaction problems is characterized by the set of polymorphisms of the basis.

Let $\mathcal{F}^A$ be the set of all finitary operations on A. Let Γ be a set of relations over A. We denote by $\mathrm{Pol}(\Gamma) = \bigcap_{R\in\Gamma}\{f \in \mathcal{F}^A : f \text{ preserves } R\}$ the set of all operations preserving all the relations in Γ.

Theorem 1. [19] *Let Γ_1 and Γ_2 be finite sets of relations over A such that* $\mathrm{Pol}(\Gamma_1) \subseteq \mathrm{Pol}(\Gamma_2)$*. If* CSP($\Gamma_1$) *is polynomially solvable then so is* CSP(Γ_2)*.*

It is easy to see (see [34] for example) that the property of being closed under some operation is preserved by some operations on relations, such as cartesian product, identification and projection.

Lemma 1. *Let R_1 be an n-ary relation over a domain A and let R_2 be a m-ary relation over a domain A which are closed under f, for some operation f. The following relations are also closed under f.*

- *The* Cartesian Product, *$R_1 \times R_2$ defined to be the $(n+m)$-ary relation*

$$R_1 \times R_2 = \{\langle t[1], \ldots, t[n+m]\rangle \quad | \quad \langle t[1], \ldots, t[n]\rangle \in R_1 \ \wedge \ \langle t[n+1], \ldots, t[n+m]\rangle \in R_2\}$$

- *The* identification *$\sigma_{i=j}(R)$ $(1 \le i, j \le n)$ defined to be the n-ary relation*

$$\sigma_{i=j}(R) = \{t \in R \quad | \quad t[i] = t[j]\}$$

- *The* projection *$\mathrm{pr}_I(R)$ where $I = \{i_1, \ldots, i_k\}$ $i_1 < \cdots < i_k$ is a set of indices chosen from $\{1, 2, \ldots, n\}$, defined to be the k-ary relation*

$$\mathrm{pr}_{\{i_1,\ldots,i_k\}}(R) = \{\langle t[i_1], \ldots, t[i_k]\rangle \quad | \quad t \in R\}.$$

On the other hand, it is also easy to see (see [34] for example) that the property of being closed under some operation is preserved by some operations on functions.

Lemma 2. *Let $g : A^m \to A$ be a m-ary function and let $f_1, f_2, \dots, f_m : A^n \to A$ be n-ary functions. If $g, f_1, f_2, \dots, f_m$ preserve R, for some relation R, then the following functions preserve also R:*

- *The* composition $g(f_1, f_2, \dots, f_m)$ *defined to be the n-ary operation*

$$g(f_1, f_2, \dots, f_m)(x_1, \dots, x_n) = g(f_1(x_1, \dots, x_n), \dots, f_m(x_1, \dots, x_n))$$

- *For every $j \leq k$, the* projection $proj_{j,k}$ *defined to be the k-ary operation*

$$proj_{j,k}(x_1, x_2, \dots, x_k) = x_j$$

Any set of operations closed under composition and containing all the projection operations is called *Clone* (See [34] for more information about clones).

There is a smallest clone which is the intersection of all clones; we shall denote it $\mathcal{I}_A$ ($\mathcal{I}$ if the domain A is clear). It is easy to see that $\mathcal{I}_A$ contains exactly the projections over A. If F is any set of operations over A, there is a smallest clone containing all of the operations in F; this is the clone generated by F, and we denote it $\langle F \rangle$. If $F = \{f_1, \dots, f_k\}$ is a finite set, we may write $\langle f_1, \dots, f_k \rangle$ for $\langle F \rangle$, and refer to $f_1, \dots, f_k$ as generators of $\langle F \rangle$.

By virtue of Theorem 1, the complexity of the constraint satisfaction problem is characterized by the set of polymorphisms, which constitutes a clone. Thus we define

Definition 7. We say that a clone F over A is tractable if for every finite set Γ of relations over A such that $F \subseteq \mathrm{Pol}(\Gamma)$, $\mathrm{CSP}(\Gamma)$ is solvable in polynomial time. An operation f is tractable if $\langle f \rangle$ is tractable.

The relationship between constraint satisfaction problems and clones was first pointed out by Jeavons et al. [23].

The introduction of clones of operations in order to study constraint satisfaction problems started a research project consisting in determining tractable clones. Up to the present moment, three main families of tractable clones have been identified: (**1**) clones containing a near-unanimity operation [20], (**2**) clones containing a coset generating operation [12], and (**3**) clones "containing a set function" [13] including important subclasses such as ACI operations [23] and constant operations [23]. All these three families were previously identified in [15] using a different approach not based on algebraic conditions.

4.0.1. *Near-Unanimity operations.* An operation $f : A^k \to K$ is called a *near-unanimity* operation if for all $a, b \in A$,

$$f(a, b, \dots, b) = f(b, a, \dots, b) = \dots = f(b, b, \dots, a) = b.$$

In [20], it is shown that every clone containing a near-unanimity operation is tractable.

Theorem 2. [20] *For every near-unanimity operation f, $\langle f \rangle$ is tractable.*

This broad family of problems contains, as a particular case, 2-SAT and the class of 0/1/all constraints [25, 11].

4.0.2. *Coset generating operations.* An operation $f : A^3 \to A$ is a coset generating operation if for all $a, b, c, d \in A$:

- $f(a, a, b) = f(b, a, a) = b.$
- $f(f(a, b, c), c, d) = f(a, b, d)$
- $f(d, c, f(c, b, a)) = f(d, b, a)$

It is easy to verify (see [12] for example) that the previous definition is equivalent to the existence of some group $(A; \cdot, ^{-1}, 1)$ such that $f(a, b, c) = a \cdot b^{-1} \cdot c$. Using this fact, we can show the tractability of clones containing a coset generating operation using a result due to Feder and Vardi [15].

Theorem 3. [15] *For every coset generating operation f, $\langle f \rangle$ is tractable.*

If $(A; \cdot, ^{-1}, 1)$ is *Abelian*, we have a particular kind of coset generating operation called an affine operation. Affine operations have been intensively studied in universal algebra (see [34] for more results about affine operations). Tractability of clones containing an affine operation was shown originally in [23].

4.0.3. *Set functions.* An operation $f : \mathcal{P}(A)/\emptyset \to A$ is said to be a *set function.* Associated to every set function f we have a *family of functions associated to f* $\{f_i : i \geq 1\}$ where for every $i \geq 1$, $f_i : A^i \to A$ is given by

$$f_i(x_1, \ldots, x_i) = f(\{x_1, \ldots, x_i\}).$$

Any clone containing a family of functions associated to a set function is tractable. Informally, we say that a clone "contains" a set function f if it contains the family of functions associated to f.

Theorem 4. [13] *For every set function f, $\langle \{f_i : i \geq 1\} \rangle$ is tractable.*

This class of problems contains important subclasses such as ACI operations [23] and constant operations [23].

We are interested in defining tractability over algebras. To do it we use the following standard definition which relates a given algebra with the clone generated by its basic set of operations.

Definition 8. The *clone of term operations* of an algebra $\mathbf{A} = (A; F)$, denoted by $T(\mathbf{A})$, is the smallest clone containing the basic operations of $\mathbf{A}$, that is $T(\mathbf{A}) = \langle F \rangle$. A *term operation* of $\mathbf{A}$ is any operation in $T(\mathbf{A})$. Two algebras with a common universe are called *term equivalent* iff they have the same term operations.

Thus, we define

Definition 9. An algebra is tractable if the clone of its term operations is tractable.

Note: in a recent paper [7], Bulatov et al. distinguish two sorts of tractability for algebras: global tractability and local tractability. The notion of tractability used in this paper corresponds to the local tractability in the sense of [7].

An operation f on A is *idempotent* iff it satisfies the identity $f(x, \ldots, x) = x$. A clone on A is called idempotent iff it consists of idempotent operations, while an algebra $(A; F)$ is idempotent iff its basic operations (or equivalently, its term operations) are idempotent. Let $\mathbf{A_1} = (A; F_1)$ be an algebra. An algebra $\mathbf{A_2} = (A; F_2)$ is said to be a *reduct* of $\mathbf{A_1}$ iff $T(\mathbf{A_2}) \subseteq T(\mathbf{A_1})$. The full idempotent reduct of $\mathbf{A_1}$ is the algebra on A whose operations are the idempotent term operations of $\mathbf{A_1}$.

5. Para-Primal Algebras

See [34] for extensive results about para-primal algebras.

For arbitrary mapping $g : A^n \to A$, we define a subset $g^\square$ of A^{n+1} by

$$g^\square = \{\langle x_1, \ldots, x_n, g(x_1, \ldots, x_n)\rangle : x_1, \ldots, x_n \in A\}.$$

Let B a subset of A^n $(n \geq 1)$. The symbol $B \preceq B_1 \times \cdots \times B_n$ will be used to designate that $\mathrm{pr}_{\{i\}} B = B_i$ for all $1 \leq i \leq n$. The number $\max\{|B_i| : 1 \leq i \leq n\}$ will be called the *size* of B. Let B and B' subsets of A^n, we denote by $B \approx B'$ the property that B and B' coincide up to the order of their components. The subset B will be termed *directly indecomposable* iff $B \not\approx (\mathrm{pr}_I B) \times (\mathrm{pr}_J B)$ holds for all nontrivial partitions $\{I, J\}$ of $\{1, \ldots, n\}$.

Furthermore, B is said to be *reduced* iff B is directly indecomposable and no projection $\mathrm{pr}_{\{i,j\}} B$ $(1 \leq i < j \leq n)$ of B is of the form $\pi^\square$ for a bijection $\pi : B_i \to B_j$.

An isomorphism between two subalgebras of an algebra $\mathbf{A} = (A; F)$ is called an *internal isomorphism* of $\mathbf{A}$. Notice that a bijection $\pi : B \to C$ with $B, C \subseteq A$ is an internal isomorphism of $\mathbf{A} = (A; F)$ if and only if $\pi^\square$ is a subuniverse of $\mathbf{A}^2$.

Definition 10. Let A be a finite set and $B \subseteq A^k$ $(k \geq 1)$. A nonvoid subset I of $\{1, \ldots, k\}$ is said to be *B-minimal* iff it is minimal in the sense that the projection $B \to \mathrm{pr}_I B$ is one-to-one.

Definition 11. An algebra $\mathbf{A}$ is called *para-primal* iff it is finite and for every natural number k, every subuniverse B of $\mathbf{A}^k$, and every B-minimal set I, we have $\mathrm{pr}_I B = \prod_{i \in I} B_i$.

As an immediate consequence of the definition we have that if $\mathbf{A}$ is para-primal, then $\mathbf{A}^2$ has no reduced subuniverses. This is precisely the property of para-primal algebras that we will use.

Actually, for finite idempotent algebras the previous condition characterizes exactly para-primality.

Theorem 5. *([34], Corollary 4.13) A finite idempotent algebra $\mathbf{A}$ is para-primal if and only if $\mathbf{A}^2$ has no reduced subuniverses.*

In the rest of the paper we will assume idempotency. This assumption is fully justified because of the following result.

Lemma 3. *(*[34]*,Lemma 4.8) A finite algebra is para-primal if and only if its full idempotent reduct is para-primal.*

In consequence:

Lemma 4. *If all idempotent para-primal algebras are tractable, then every para-primal algebra is tractable.*

Proof. Let $\mathbf{A}$ be any para-primal algebra and let F be its full idempotent reduct. Recall that $\mathbf{A}$ is tractable iff $T(\mathbf{A})$ is tractable. By definition, $F \subseteq T(\mathbf{A})$. F is para-primal and idempotent and consequently tractable. Thus, by virtue of Theorem 1, $T(\mathbf{A})$ is tractable. ∎

Thus, tractability of idempotent para-primal algebras implies tractability of para-primal algebras in general. The same result could be obtained as a consequence of the classification of the minimal para-primal algebras given later.

The class of idempotent para-primal algebras has an interesting interpretation: Let Γ be a finite set of relations over A. Associate to every problem in CSP(Γ) with variables $x_1, \ldots, x_n$, the relation given by $\{\langle t(x_1), \ldots, t(x_n)\rangle : t \text{ is a solution of } \mathcal{P}\}$. The set $\mathcal{S}$ containing all the relations associated to all the problems in CSP(Γ) is a good representative of the expressiveness power of the basis Γ. Informally, we can say that Γ can "simulate" every relation in $\mathcal{S}$. It is possible to see that the set of subuniverses of $(A; \mathrm{Pol}(\Gamma))^n$ for all $n \geq 1$ is, under some technical assumptions [23], identical to $\mathcal{S}$. Thus, the previous theorem says for every basis Γ, the algebra $(A; \mathrm{Pol}(\Gamma))$ is para-primal if and only if the class of binary relations "simulated" by Γ is very simple. More precisely, every such relation B is, either, a cartesian product $B_1 \times B_2$ or $\pi^{\square}$ for some bijection $\pi : B_1 \to B_2$.

5.1. **Some Examples.** We present some examples of para-primal algebras. Para-primality of the following algebras can be shown using the general description of para-primal algebras presented later, but it can be proved also by showing that they do not exhibit reduced subuniverses of arity 2.

5.1.1. *The Discriminator Operation.* For a set A, the *discriminator* t on A is defined by

$$t(a, b, c) = \begin{cases} c & \text{if } a = b \\ a & \text{otherwise} \end{cases}$$

We show later that algebra $(A; t)$ is para-primal. Actually, it is known that finite algebras containing the discriminator operator t among their term operations constitute an important subclass of para-primal algebras, called *quasi-primal* algebras. Quasi-primal algebras satisfy the following alternative characterization that we will use later.

Lemma 5. *(*[34]*, Theorem 4.2) An algebra* $\mathbf{A}$ *is quasi-primal algebra iff* $\mathbf{A}$ *is finite and, for every* $k \geq 0$, $\mathbf{A}^k$ *has no reduced subuniverses.*

The term operation $d(x,y,z) = t(x,t(x,y,z),z)$ is the following majority operation:

$$d(a,b,c) = \begin{cases} a & \text{if } a = b \\ c & \text{otherwise} \end{cases}$$

Operation d is called *dual discriminator* and it has been intensively studied (See [34]).

5.1.2. *Affine operations on Universes of Prime Order.* Recall that a 3-ary operation f on a universe A such that $f(a,b,c) = a - b + c$ where $(A;+,-)$ is an abelian group is called an *affine* operation. If A is of prime order, then $(A;f)$ is para-primal. Actually the original idea in the definition of para-primal algebras was to extend quasi-primal algebras in order to contain groups of prime order [8]. It is known [34] that the subuniverses of $(A;f)^n$ are precisely the subsets of A^n which can be expressed as a system of counting functions modulo $|A|$.

5.1.3. *Switching Operation.* The simplest example of a para-primal algebra not contained in the so far tractable cases is the *switching operation* s on A, defined by:

$$s(a,b,c) = \begin{cases} c & \text{if } \; a = b \\ b & \text{if } \; a = c \\ a & \text{otherwise} \end{cases}$$

Algebra $(A;s)$ is para-primal. Notice that restricted to a boolean domain $A = \{0,1\}$, the switching operation $s(a,b,c)$ is equivalent to the affine operation $(a \oplus b \oplus c)$. We will see later that for larger domains $|A| = k \geq 3$, algebra $(A;s)$ is not affine. Actually, the algebra $(A;s)$ is a pure representative of this new tractable class in the sense, that it cannot be included in any of the classes known to be tractable.

5.2. **Tractability of Para-Primal Algebras.** In this section we will prove that idempotent para-primal algebras are tractable.

Theorem 6. *Idempotent para-primal algebras are tractable.*

The proof is splitted in two parts. First, we prove that subuniverses of product powers of para-primal algebras satisfy a certain decomposability property. Then, we derive an polynomial-time algorithm that uses this decomposability property.

5.2.1. *Structure of the Subuniverses.* The key notion in the following sections is that of *isomorphic-decomposable.* First, we introduce some notation.

Let $\mathbf{A} = (A;F)$ be an algebra. Consider the isomorphy relation $\sim$ over the subalgebras of $\mathbf{A}$. Clearly, relation $\sim$ is an equivalence relation and, in consequence, it determines a partition of the subalgebras of $\mathbf{A}$.

Let $B \preceq B_1 \times \cdots \times B_n$ be an arbitrary subuniverse of $\mathbf{A}^n$. Thus, the isomorphy relation determines a partition $I_1, \ldots, I_m$ of $\{1, \ldots, n\}$ in the following way: let $\mathbf{B_i} = (B_i;F)$ $(1 \leq i \leq n)$ be the subalgebra of $\mathbf{A}$ with universe B_i. Then $1 \leq i, j \leq n$ belong to the same class iff $\mathbf{B_i} \sim \mathbf{B_j}$. We will say that $I_1, \dots, I_m$ is the *isomorphic-decomposition associated to* B *with respect to* $\mathbf{A}$.

Subuniverse B is said to be *isomorphic-decomposable* if for all $t \in A^n$,

$$\mathrm{pr}_{I_1} t \in \mathrm{pr}_{I_1} B, \dots, \mathrm{pr}_{I_m} t \in \mathrm{pr}_{I_m} B \Longrightarrow t \in B.$$

For simplicity, we will assume that I_1 contains the first $|I_1|$ indices in $\{1, \dots, n\}$ and so on. Thus, B is isomorphic-decomposable iff $B = \mathrm{pr}_{I_1} B \times \cdots \times \mathrm{pr}_{I_m} B_m$.

In this section we will show that if $\mathbf{A}$ is an idempotent para-primal algebra then every subuniverse B of $\mathbf{A}^n$ is isomorphic-decomposable. We will proceed by contradiction. First we will show that if B is non-isomorphic-decomposable then there is a projection $\mathrm{pr}_I B$ which is non-isomorphic-decomposable and reduced. Later, we will use some previous results about the structure of para-primal algebras to show that such a subuniverse cannot exist.

Lemma 6. *For any algebra* $\mathbf{A}$ *and for any non-isomorphic-decomposable subuniverse* B *of* $\mathbf{A}^n$ *there exists some set of indices* $I \subseteq \{1, \dots, n\}$ *such that* $\mathrm{pr}_I B$ *is reduced and non-isomorphic-decomposable.*

Proof. Recall that a relation C is reduced if: (1) C is directly indecomposable and (2) no projection $\mathrm{pr}_{i,j} C$ is of the form $\pi^{\square}$ for some bijection $\pi : C_i \to C_j$. The structure of the proof is the following. The underlying idea consists in to prove that, given a non-isomorphic-decomposable subuniverse C that violates one of the previous conditions, it is possible to find a proper subset of indices $\mathrm{pr}_I C$ that is still non-isomorphic-decomposable. Since every subuniverse of arity 1 is isomorphic-decomposable, we must eventually find a reduced non-isomorphic-decomposable projection $\mathrm{pr}_I C$.

We do it in two stages corresponding to the two conditions in the definition of reducedness. In the first stage we show that if C is non-isomorphic-decomposable and directly decomposable there exists a proper projection $\mathrm{pr}_I C$ which is still non-isomorphic-decomposable.

In the second stage we find that if a pair of indices i, j violate the second condition, then we can obtain a projection I that does not contain one of the indices and such that $\mathrm{pr}_I C$ is non-isomorphic-decomposable.

First stage: Let $B \preceq B_1 \times \cdots \times B_n$ be a non-isomorphic-decomposable subuniverse of $\mathbf{A}^n$. If B is directly decomposable, then there exists some proper subset of indices $I \subset \{1, \dots, n\}$ such that $\mathrm{pr}_I B$ is non-isomorphic-decomposable.

Proof: Let I, J be any partition of $\{1, \dots, n\}$ such that $B \approx \mathrm{pr}_I B \times \mathrm{pr}_J B$. Clearly, the set B is isomorphic-decomposable whenever both $\mathrm{pr}_I B$ and $\mathrm{pr}_J B$ are isomorphic-decomposable.

Second stage: Let $B \preceq B_1 \times \cdots \times B_n$ be a non-isomorphic-decomposable subuniverse of $\mathbf{A}^n$, and let i, j two indices such that $\mathrm{pr}_{\{i,j\}} B$ is $\pi^{\square}$ for some bijection $\pi : B_i \to B_j$. Then $B' = \mathrm{pr}_{\{1,\dots,j-1,j+1,\dots,n\}} B$ is non-isomorphic-decomposable.

Proof: Let $I_1, \dots, I_m$ be the isomorphic-decomposition associated to B with respect to $\mathbf{A}$. Clearly i, j belong to the same class, say I_k. By the non-isomorphic decomposability of B, there exists $t \notin B$ such that for all $1 \leq l \leq m$, $\mathrm{pr}_{I_l} t \in \mathrm{pr}_{I_l} B$.

Let $t' = \mathrm{pr}_{\{1,\dots,j-1,j+1,\dots,n\}} t$. We will see that assignment t' proves the non-isomorphic-decomposability of B'. Clearly, $I_1, \dots, I_k/\{j\}, \dots, I_m$ is the isomorphic-decomposition associated to relation B'. By construction we have $\mathrm{pr}_{I_l/\{j\}} t \in \mathrm{pr}_{I_l/\{j\}} B'$. Furthermore, $t' \notin B'$, since $\mathrm{pr}_{\{i,j\}} B = \pi^{\square}$, and, in consequence, t is the unique valid extension of t'. Thus, t' certificates the non-isomorphic-decomposability of B' and the proof is done.

The main result is proved by induction on the arity n of B. If $(n = 1)$, then B is trivially isomorphic-decomposable. Consider the case $(n > 1)$. Let $B \preceq B_1 \times \cdots \times B_n$ be any non-isomorphic-decomposable subuniverse of $\mathbf{A}^n$. If B is reduced we are done. Otherwise, we should consider two possibilities: (1) B is directly decomposable, in this case there exists some set of indices I with $|I| < n$ such that B_I is non-isomorphic-decomposable. Otherwise (2) there exists two indices $1 \leq i, j \leq n$ such that $\mathrm{pr}_{\{i,j\}} B$ is $\pi^{\square}$ for some bijection $\pi : B_i \to B_j$. Thus, $\mathrm{pr}_{\{1,\dots,j-1,j+1,\dots,n\}} B$ is non-isomorphic-decomposable.

In consequence, in both cases there exists some set of indices I with $|I| < n$ such that $\mathrm{pr}_I B$ is non-isomorphic-decomposable. Then, there exists some subset $I' \subseteq I$, such that $\mathrm{pr}_{I'} B$ is non-isomorphic-decomposable and reduced. ∎

Now, we will apply some known results to prove that if an algebra $\mathbf{A}$ is para-primal, then $\mathbf{A}^n$ has no reduced non-isomorphic-decomposable subuniverses. By contradiction, we have the following result: Every subuniverse of $\mathbf{A}^n$ is isomorphic-decomposable.

We study two cases: quasi-primal and non-quasi-primal algebras. The first case is easy, since if $\mathbf{A}$ is a quasi-primal algebra then, by Lemma 5, $\mathbf{A}^n$ does not have any reduced subuniverse at all. If $\mathbf{A}$ is not quasi-primal, then the following result is applicable.

Theorem 7. *([34], Theorem 4.3) Let $\mathbf{A} = (A; F)$ be a finite idempotent algebra. Suppose $\mathbf{A}$ is not quasi-primal, and let $B \preceq B_1 \times \cdots \times B_n$ be a reduced subuniverse of $\mathbf{A}^n$ for some $n \geq 2$. Then either*

a.- *$\mathbf{A}^2$ has a reduced subuniverse of the same size as B, or*
b.- *$\mathbf{B_i} = (B_i; F)$ $(1 \leq i \leq n)$ are all isomorphic affine subalgebras.*

See [34] for the definition of affine algebras. The only property of affine algebras that we will use in this paper is that affine algebras contain an affine operation among their term operations.

As a consequence of the Theorem 7 we have that, if $\mathbf{A}$ is para-primal, condition (b) holds (recall that if $\mathbf{A}$ is para-primal then $\mathbf{A}^2$ has no reduced subuniverses) and, in consequence, every reduced subuniverse $B \preceq B_1 \times \cdots \times B_n$ is isomorphic-decomposable, since $\mathbf{B_i}$ are isomorphic algebras.

Thus, in both cases (quasi-primal and non-quasi-primal algebras) we can imply that there exists non-isomorphic-decomposable reduced subuniverses of $\mathbf{A}^n$, and, by virtue of Lemma 6, every subuniverse of $\mathbf{A}^n$ is isomorphic-decomposable.

Thus, we have proved the following result:

Corollary 1. *Let $\mathbf{A}$ be a para-primal algebra and let B be a subuniverse of $\mathbf{A}^n$. Then B is isomorphic-decomposable.*

Furthermore, from Theorem 7 we can infer some interesting properties of $\mathrm{pr}_{I_l} B$ for each class I_l of the isomorphic-decomposition. The following theorem summarizes the results.

Theorem 8. *Let A be an idempotent para-primal algebra and let B be a subuniverse of $\mathbf{A}^n$, then B is isomorphic decomposable. Furthermore, let $I_1, \ldots, I_m$ be the isomorphic-decomposition associated to B in $\mathbf{A}$. Then for every $1 \leq i \leq m$, $\mathrm{pr}_{I_l} B$ is preserved by an affine operation or by the discriminator operation t.*

Proof. We say that an isomorphy class of $\mathbf{A}$ is affine if so is any of its representatives. Let I_l $(1 \leq l \leq m)$ be any of the equivalence classes given by the isomorphy decomposition in $B \preceq B_1 \times \cdots \times B_n$. Assume that I_l does not correspond to an affine isomorphy class. We will prove that $\mathrm{pr}_{I_l} B$ is preserved by the discriminator operation t (and consequently is preserved also by the dual discriminator operator d): Let I be any set of indices containing exactly one of the representatives of the relation: $i \equiv j$ $(i, j \in I_l)$ iff $\mathrm{pr}_{\{i,j\}} B$ is of the form $\pi^\square$ for a bijection $\pi : B_i \to B_j$. We will study two cases: $\mathrm{pr}_I B \neq \prod_{i \in I} B_i$ or $\mathrm{pr}_I B = \prod_{i \in I} B_i$.

If $\mathrm{pr}_I B \neq \prod_{i \in I} B_i$, then we will show that there exists some $J \subseteq I$ such that $\mathrm{pr}_J B$ is reduced.

First notice that the set of indices I contains exactly one of the representatives of the relation $\equiv$. In consequence, for any two indices $i \neq j \in I$, $\mathrm{pr}_{i,j} B \neq \pi^\square$ for any bijection $\pi : B_i \to B_j$. Thus, the second condition of reducedness holds for every subset of I.

Let $J \subseteq I$ be a minimal subset of I satisfying $\mathrm{pr}_J B \neq \prod_{j \in J} B_j$. We show by contradiction that pr_J is directly indecomposable. Consider any partition J_1, J_2 of J such that $\mathrm{pr}_J B \approx \mathrm{pr}_{J_1} B \times \mathrm{pr}_{J_2} B$. Then, it must be the case that $\mathrm{pr}_{J_1} B \neq \prod_{j \in J_1} B_j$ or $\mathrm{pr}_{J_2} B \neq \prod_{j \in J_2} B_j$ in contradiction with the fact that J is minimal.

In consequence, by Theorem 7, $\mathbf{A}$ is quasi-primal and therefore contains t among its term operations.

Otherwise, $\mathrm{pr}_I B = \prod_{i \in I} B_i$. It can be easily verified that, in this case, discriminator t preserves $\mathrm{pr}_{I_l} B$. To see it, notice that for any tuple t over

A, $t \in B$ iff for all $1 \leq l \leq m$ $\mathrm{pr}_{I_l} t \in \mathrm{pr}_{I_l} B$. Let $t_1, t_2, t_3 \in B$ any arbitrary tuples in B.

We will show that for all $1 \leq l \leq m$, $t(\mathrm{pr}_{I_l} t_1, \mathrm{pr}_{I_l} t_2, \mathrm{pr}_{I_l} t_3) \in \mathrm{pr}_{I_l} B$. Since I_l is one of the equivalence classes of the relation $\equiv$ we have that $\mathrm{pr}_{I_l} t_1 \neq \mathrm{pr}_{I_l} t_2$ implies that $t_1[i] \neq t_2[i]$, for all $i \in I_l$. In this case $t(\mathrm{pr}_{I_l} t_1, \mathrm{pr}_{I_l} t_2, \mathrm{pr}_{I_l} t_3) = \mathrm{pr}_{I_l} t_1 \in \mathrm{pr}_{I_l} B$. Otherwise, we have that $\mathrm{pr}_{I_l} t_1 = \mathrm{pr}_{I_l} t_2$ and, in consequence, $t(\mathrm{pr}_{I_l} t_1, \mathrm{pr}_{I_l} t_2, \mathrm{pr}_{I_l} t_3) = \mathrm{pr}_{I_l} t_3 \in \mathrm{pr}_{I_l} B$.

5.2.2. *Algorithm.* The underlying idea of the algorithm is that using the decomposability property shown in the previous section, it is possible, after an appropriate transformation of the problem, to divide the problem in subproblems. Every subproblem falls in one of the following already known tractable classes: affine problems and dual discriminator problems.

Let $\mathcal{P}$ be a constraint instance in $\mathrm{CSP}(\Gamma)$ over the variables $x_1, \ldots, x_n$ and with domain A. Consider any constraint $C = \langle s, R \rangle$ in $\mathcal{P}$ and any variable x. We will denote the set of values enforced by C in x by C_x. The set C_x is given by $C_x = \bigcap_{i \in I} \mathrm{pr}_{\{i\}} R$, where I is the set of indices containing the positions of x in s (I is empty if x is not contained in s). Furthermore, we can associate to x the set of values D_x given by the intersection of all the values enforced by all the constraints in $\mathcal{P}$, that is $D_x = A \cap \bigcap_{C \in \mathcal{P}} C_x$. Replace every constraint $\langle (x_{i_1} \ldots, x_{i_k}), R \rangle$ by the new constraint $\langle (x_{i_1} \ldots, x_{i_k}), R \cap (D_{x_{i_1}} \times \cdots \times D_{x_{i_k}}) \rangle$ obtaining a new problem $\mathcal{P}_1$ with identical set of solutions.

It is important to notice that every constraint relation occurring in $\mathcal{P}$ belongs to Γ and that D_{x_i} have been obtained from the constraint relations using exclusively cartesian products, identification projections, and intersections, that is, the operations that, according to Lemma 1, preserve the polymorphisms of the relations. In consequence, for all $1 \leq i \leq n$ D_{x_i} is preserved by all the operations preserved by Γ. In consequence, every constraint relation in the new problem $\mathcal{P}_1$ is also preserved by all the polymorphisms of Γ.

Iterate the previous step while some progress is made, ie., some constraint is modified. We obtain a sequence of problem instances: $\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_l$. For each $1 \leq i \leq l$, $\mathcal{P}_i$ has the same solutions than $\mathcal{P}$, and every constraint relation occurring in $\mathcal{P}_i$ is preserved by all the operation in $\mathrm{Pol}(\Gamma)$.

Furthermore, instance $\mathcal{P}_l$ satisfies the following important property: Let $C = \langle (x_{i_1}, \ldots, x_{i_k}), R \rangle$ be any constraint in $\mathcal{P}_l$, then $R \preceq D_{x_{i_1}} \times \cdots \times D_{x_{i_k}}$

Consider the isomorphy relation over $\mathbf{D_{x_1}}, \ldots, \mathbf{D_{x_n}}$, as defined above. Consider any constraint $C = \langle s, R \rangle$ in $\mathcal{P}_l$. If $\mathrm{Pol}(\Gamma)$ is para-primal then the constraint relation R is isomorphic-decomposable. In consequence, it is possible to replace the constraint C by the new constraints $\langle s_1, R_1 \rangle, \ldots, \langle s_m, R_m \rangle$ obtained by dividing R in its components. Each component corresponds to a class with respect to the isomorphy relation. We obtain a new problem instance $\mathcal{Q}$ with the same set of solutions and satisfying the following property: for every constraint $C = \langle (x_{i_1}, \ldots, x_{i_k}), R \rangle$ in $\mathcal{Q}$ we have that $\mathbf{D_{x_{i_1}}}, \ldots, \mathbf{D_{x_{i_k}}}$ are isomorphic.

Since the scope of every constraint C only involves variables belonging to the same isomorphy class, we can split problem $\mathcal{Q}$ in subproblems $\mathcal{Q}_1, \ldots, \mathcal{Q}_r$. Each problem $\mathcal{Q}_i$ $(1 \leq i \leq r)$ contains the variables of only one isomorphy class in $\mathbf{D}_{\mathbf{x}_1}, \ldots, \mathbf{D}_{\mathbf{x}_\mathbf{m}}$.

Consider a subproblem $\mathcal{Q}_i$ $(1 \leq i \leq r)$. If problem $\mathcal{Q}_i$ corresponds to an affine isomorphy class then every constraint relation occurring in $\mathcal{Q}_i$ is preserved by the corresponding affine operation and, consequently, $\mathcal{Q}_i$ is polynomial-time solvable (Theorem 3). Otherwise, by Theorem 8, every constraint relation in $\mathcal{Q}_i$ is preserved by the discriminator operation t and, consequently, by the dual discriminator d. In this case, $\mathcal{Q}_i$ is polynomial-time solvable since d is a near-unanimity operation and Theorem 2 applicable. ∎

6. The Class of Para-Primal Problems

In the previous section we shown the tractability of para-primal algebras. In order to complete the study of para-primal algebras we must consider the class of basis (sets of relations) associated to para-primal algebras.

Definition 12. A finite set of relations Γ over A is called *para-primal* iff $(A; \mathrm{Pol}(\Gamma))$ is para-primal.

Thus we have,

Theorem 9. *For every para-primal set of relations Γ, the class of problems* $\mathrm{CSP}(\Gamma)$ *is solvable in polynomial-time.*

We must show that the class of para-primal basis satisfies two important properties: (1) It has to be decidable whether or not a given basis is para-primal and (2) the class of para-primal basis must contain some representative not contained in the classes already known to be tractable.

6.1. **Para-primal problems are decidable.** The decidability of para-primal problems follows from a general description of para-primal algebras, due to Quackenbush [32].

The structure of para-primal algebras has been intensively studied [8, 9, 32] (see [34] for an overview). It is known that para-primal algebras constitute an order filter in the clone lattice. For a finite universe A, up to term equivalence, there are only finitely many minimal para-primal algebras which were identified by Quackenbush [32].

Let A be a finite set, and $A_1, \ldots, A_k$ be a (possibly empty) family of subsets of A with $|A_i \cap A_j| \leq 1$ $(1 \leq i < j \leq k)$ and $|A_i| = q_i$, a prime power, for all $1 \leq i \leq k$. We say that $A_1, \ldots, A_k$ is a *decomposition* of A.

Let us fix a generating element c_i of the multiplicative group of each finite galois field $\mathrm{GF}(q_i)$ so that $c_i = c_j$ whenever $q_i = q_j$ $(1 \leq i, j \leq k)$. Let us fix, for each $1 \leq i \leq k$, a group $(A; +_i, -_i, 0_i)$. Consider the vector space ${}_{\mathrm{GF}(q_i)}\underline{A_i} = (A_i; +_i, -_i, 0_i, \mathrm{GF}(q_i))$ for every $1 \leq i \leq k$, and define a ternary operation p on A by,

$$p(a,b,c) = \begin{cases} a -_i b +_i c & \text{if } a,b,c \in A_i \text{ for some } 1 \leq i \leq k \\ t(a,b,c) & \text{otherwise} \end{cases}$$

and a binary operation g on A by

$$g(a,b) = \begin{cases} c_i a +_i (1-c_i)b & \text{if } a,b \in A_i \text{ for some } 1 \leq i \leq k \\ a & \text{otherwise} \end{cases}$$

We say that p and g is a pair of operations associated to the decomposition $A_1, \dots, A_k$ of A.

The algebra $(A; p, g)$ is exactly the minimal para-primal algebra with non-singleton affine subuniverses $A_1, \dots, A_k$.

Theorem 10. *([34], Theorem 4.15) For every decomposition $A_1, \dots, A_k$ and for every pair of operations p and g associated to $A_1, \dots, A_k$, the algebra $(A; p, g)$ is para-primal, and has no proper para-primal reduct in which the nonsingleton affine subuniverses are exactly $A_1, \dots, A_k$. Moreover, every para-primal algebra $(A; F)$ with nonsingleton affine subuniverses $A_1, \dots, A_k$ has a reduct of the form $(A; p, g)$.*

In consequence, a basis Γ is para-primal iff there exists some decomposition $A_1, \dots, A_k$, such that the operations p, and g associated to $A_1, \dots, A_k$ belong to $\mathrm{Pol}(\Gamma)$. Since we have a finite number of decompositions we have the following result.

Theorem 11. *The problem of deciding whether a finite basis is para-primal is decidable.*

Using the general description of para-primal algebras we can show that the examples in Section 5.1 are para-primal algebras. Actually, we show an stronger result: every algebra introduced in Section 5.1 is a minimal para-primal algebra.

- Let A be a finite set and let t the discriminator on A. We show that algebra $(A; t)$ is para-primal. More precisely, we will show that the pair of functions p and g associated to the empty decomposition ($k = 0$) are term operations of $(A; t)$. To see it, notice that if $k = 0$ then g is the first projection and $p = t$. Thus, we are done. Furthermore, by the second identify we have also that t is a term operation of $(A; p, g)$ and, in consequence, we have that algebras $(A; t)$ and $(A; p, g)$ are term equivalent.
- Let $f(x, y, z) = x - y + z$ be an affine operation on A of prime size q. Consider the decomposition of A given by $A_1 = A$. Let $\mathbf{1}$ be the identity element of the multiplicative group of $\mathrm{GF}(q)$ and let $c = \mathbf{1} \cdot n \quad (n \geq 2)$ be a generating element of the multiplicative group of $\mathrm{GF}(q)$. Consider the vector space ${}_{\mathrm{GF}(q)}\underline{A} = (A; +, -, \mathrm{GF}(q))$ and define the functions p and g as above. We have that

$$g(a,b) = \overbrace{f \dots f}^{n-1}(a,b,a),b,a) \dots),b,a)$$

and $p(a,b,c) = f(a,b,c)$. Thus p and g are term operations of $(A;f)$. Since $p = f$, we have also that $(A;f)$ and $(A;p,g)$ are term equivalent.

- Finally, let s be the switching operation on A. Consider the decomposition $A_1, \ldots, A_{\binom{|A|}{2}}$ of A containing all the 2-element subsets of A. Let p and g be the pair of operations associated to the previous decomposition. Clearly g is the first projection and $p(a,b,c) = s(a,b,c)$. In consequence $(A;s)$ and $(A;p,g)$ are term equivalent.

6.2. **Para-primal Problems Contain New Tractable Problems.** Let us place the class of para-primal problems with respect to the classes already known to be tractable. Recall that, so far, there are three main families of tractable classes [15], in our terminology: near-unanimity, set function, and coset generating problems.

The class of para-primal algebras contains some problems included in those three families. For example, among the operations presented in Section 5.1, the discriminator operation is a near-unanimity operation and affine operations are particular cases of coset generating operations.

Nevertheless, some algebras cannot be included in any of the three families. Let A be a finite set with at least three elements and let s be the switching operation on A. We will show that $T(A;s)$ does not contain any of the clones known to be tractable. Let $k = |A|$ be the cardinality of A.

First, we will show that $T(A;s)$ does not contain any coset generating operation. To see it, notice that $(A;s)$ has subuniverses of size $k-1$ which does not divide k for $k \geq 3$, in contradiction with the fact that the cardinality of any subuniverse of $\mathbf{A}^n$ for an algebra $\mathbf{A}$ containing a coset generating operation among its term operations should divide $|A|^n$ [12].

At this point, we would like to add an important remark: the definition of coset generating operations (and affine operations) suffers from some domain dependence. An example: consider the following relation

$$R = \{t \in \{0,1\}^3 : t[1] \oplus t[2] \oplus t[3] = 0\}$$

Clearly, operation $f(x,y,z) = x \oplus y \oplus z$ preserves R. Thus, we could say that R is affine, meaning that R is preserved by an affine operation. Nevertheless, if the domain of our problem is $A = \{0,1,2\}$ then there not exists any affine operation and (any coset generating operation) over A that preserves R. In consequence, we must extend the definition of coset generating problems in order to avoid the above-mentioned problem. We will say that a basis Γ on A is coset generating if there exists some subset B of A and a coset generating operation f on B that preserves Γ.

The previous proof is adapted in a straightforward way to this extended definition. To do it, just consider a basis $\Gamma = \{R_1, R_2\}$ containing two different subsets of cardinality $k-1$. By an analogous reasoning to the previous case, it is easy to show that Γ is not a coset-generating basis.

Finally, we show that no near-unanimity operation nor a "set function" belongs to the clone of term operations of $(A; s)$. To do it we use the fact that the switching operation, restricted to a two-element domain $\{0, 1\}$ is equivalent to the affine operation $f(x, y, z) = x \oplus z \oplus z$. Later, we show that $T(\{0, 1\}; f)$ does not contain any near-unanimity operation nor a "set function". In consequence, $T(A; s)$ does not contain any near-unanimity operation (set function), since the restriction of a near-unanimity operation (set function) to any subuniverse, and in particular to a 2-element subuniverse, is still a near-unanimity operation (set function).

To verify that $T(\{0, 1\}; f)$ does not contain any near-unanimity operation consider the following relation preserved by f.

$$R_3 = \{(x, y, z) : x \oplus y \oplus z = 0\}$$

We prove by contradiction that R_3 is not preserved by any near-unanimity operation. Let $n : \{0, 1\}^k \to \{0, 1\}$ be any near-unanimity operation of arity $k \geq 3$ that preserves R_3. If $k = 3$ we get easily the following contradiction:

$$n(\langle 1, 1, 0\rangle, \langle 1, 0, 1\rangle, \langle 0, 1, 1\rangle) = \langle 1, 1, 1\rangle \notin R_3$$

Assume $k \geq 4$. Since $n(0, 1, \ldots, 1) = 1$ and $n(0, \ldots, 0, 1) = 0$, there exists some value $1 \leq n \leq k - 2$ such that

$$n(\overbrace{0, \ldots, 0}^{n}, 1, \ldots, 1) = 1 \text{ and } n(\overbrace{0, \ldots, 0}^{n+1}, 1, \ldots, 1) = 0$$

Then we have the following contradiction:

$$n(\overbrace{\langle 0, 0, 0\rangle, \ldots, \langle 0, 0, 0\rangle}^{n}, \langle 1, 0, 1\rangle, \langle 1, 1, 0\rangle, \ldots, \langle 1, 1, 0\rangle) = \langle 1, 0, 0\rangle \notin R_3$$

Finally, we prove that the clone of term operations of $(\{0, 1\}; f)$ does not contain the family associated to any set function. Consider the following relation preserved by f:

$$R_4 = \{\langle x, y\rangle : x \oplus y = 1\}$$

Consider now any set function g over $\{0, 1\}$ and let $x = g(\{0, 1\})$. Clearly g_2 does not preserve R_4, since

$$g_2(\langle 0, 1\rangle, \langle 1, 0\rangle) = \langle x, x\rangle \notin R_4$$

for any value for x.

Therefore, the class of para-primal problems constitutes a new tractable class.

As a direct corollary of the previous proof we obtain that the basis $\{R_1, R_2, R_3, R_4\}$, where R_1, R_2, R_3, R_4 are defined as above is not contained in any of the classes known to be tractable: near-unanimity, coset-generating and set function problems.

The class of problems associated to para-primal algebras can be regarded as a combination of some affine algebras, which are particular cases of subgroup problems. Feder and Vardi [15] introduced another way to extend the class of coset-generating problems. The approach introduced in [15] is based in adding some non-subgroup constraints rather that in combining different subgroup problems.

7. An Application: Dichotomy under the Availability of Permutation Relations

A dichotomic classification of CSP(Γ) is a major open question. In fact, the mere existence of such a dichotomy is an interesting theoretical question [15]. It is known that restricting the problem in several ways we have a dichotomy [33, 18, 11]. In this section we extend a known dichotomy classification of constraint satisfaction problems due to Cooper et al. [11]. The problem addressed in [11] is to classify the computational complexity of constraint satisfaction problems whenever the basis: (1) contains only binary relations, (2) includes all the unary relations, and (3) includes all the *permutation relations*, i.e., relations of the form $\pi^{\square}$ where $\pi : A \to A$ is a permutation. In [11], it is proved that the only family of tractable basis under the above-mentioned restriction is the class of the so-called 0/1/all relations which in our terminology can be defined as the binary relations preserved by the dual discriminator operation d.

In this section we use the tractability of switching operation and some known results about the classification of homogeneous algebras to extend the previous dichotomy by dropping restrictions (1) and (2). Let A be a finite set. A *permutation relation* on A is any relation of the form $\pi^{\square}$ where $\pi : A \to A$ is a permutation. We will study the problems given by basis of relations containing all the permutation relations on the domain of the basis. Let Γ be a finite set of relations on A. Then

$$\mathrm{CSP}^*(\Gamma) = \mathrm{CSP}(\Gamma \cup \{R : R \text{ permutation relation on } A\})$$

7.0.1. *Homogeneous Algebras.* The key observation of this section is that for every permutation $\pi : A \to A$, $\pi^{\square}$ is a subuniverse of $\mathbf{A}$ if and only if π is an automorphism of $\mathbf{A}$. Thus, we can restrict to algebras where every permutation is an automorphism.

Definition 13. An algebra $\mathbf{A} = (A; F)$ is called *homogeneous* iff every permutation of A is an automorphism of $\mathbf{A}$. An operation f on A is homogeneous if $(A; f)$ is homogeneous.

Example 1. *Let A be a finite set, $|A| = n$. The following operations on A are homogeneous:*

- *the switching operation s;*
- *the discriminator operation t;*
- *the dual discriminator operation d;*
- *the k-ary* near-projection l_k $(3 \leq k \leq n)$ *defined by*

$$l_k(a_1, \ldots, a_k) = \begin{cases} a_1 & \text{if } |\{a_1, \ldots, a_k\}| < k \\ a_k & \text{otherwise}; \end{cases}$$

- *the $(n-1)$-ary operation r_n defined by*

$$r_n(a_1,\ldots,a_{n-1}) = \begin{cases} a_1 & \text{if } |\{a_1,\ldots,a_{n-1}\}| < n-1 \\ a_n \text{ with } \{a_n\} = A/\{a_1,\ldots,a_{n-1}\} & \text{otherwise}; \end{cases}$$

  for $n=2$, r_2 is the unique nonidentity permutation on A, while for $n=3$, $r_3(x,y) = 2x+2y$ for arbitrary group $(A;+)$;
- *the $(n-1)$-ary operation d_n defined by*

$$d_n(a_1,\ldots,a_{n-1}) = \begin{cases} d(a_1,a_2,a_3) & \text{if } |\{a_1,\ldots,a_{n-1}\}| < n-1 \\ a_n \text{ with } \{a_n\} = A/\{a_1,\ldots,a_{n-1}\} & \text{otherwise}; \end{cases}$$

  provided $n \geq 4$.
- *the operation $x+y+z$ if $(A;+)$ is a 4-element group of exponent 2.*

Homogeneous algebras were first investigated by Marczewski [28] and were finally completely classified by Marchenkov [27]. Here we give a description of Marchenkov's classification (See [34] for a detailed presentation of homogeneous algebras). There exists some "exceptional" homogeneous algebras on small sets ($|A| \leq 5$). Thus, they are presented separately. The lattice of clones of homogeneous algebras on a set A of $n(\geq 5)$ elements is depicted in Figure 1. The lattice of clones of homogeneous algebras on a 4-element, 3-element, and 2-element set are depicted in Figure 2. The notation is as follows:

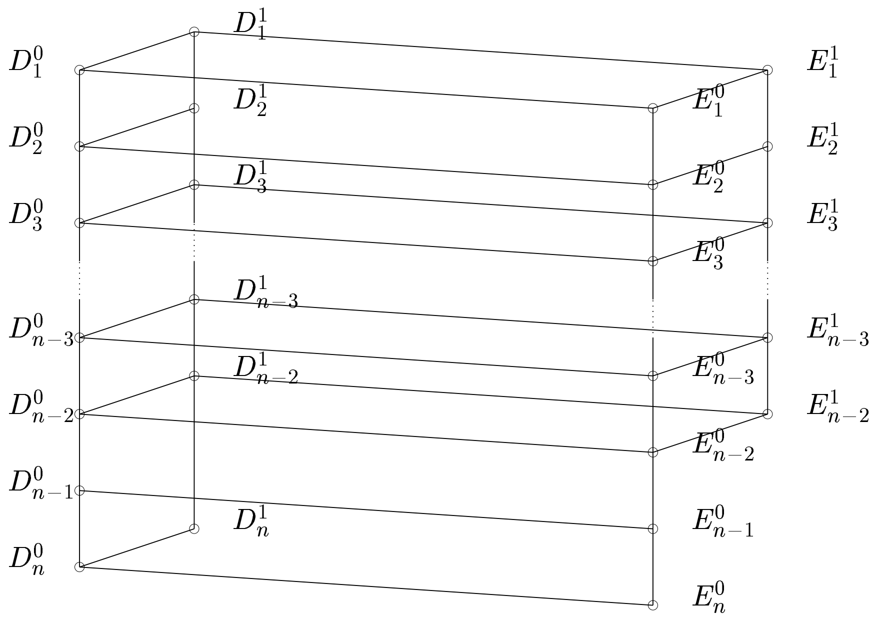

FIGURE 1. Lattice of homogeneous algebras with size $n \geq 5$

- $D_1^0 = \langle t \rangle$; $D_n^0 = \langle d \rangle$;
- $D_i^0 = \langle d, l_{i+1} \rangle$ for $2 \leq i \leq n-1$;
- $D_1^1 = \langle t, r_n \rangle$; $D_n^1 = \langle d_n \rangle$;

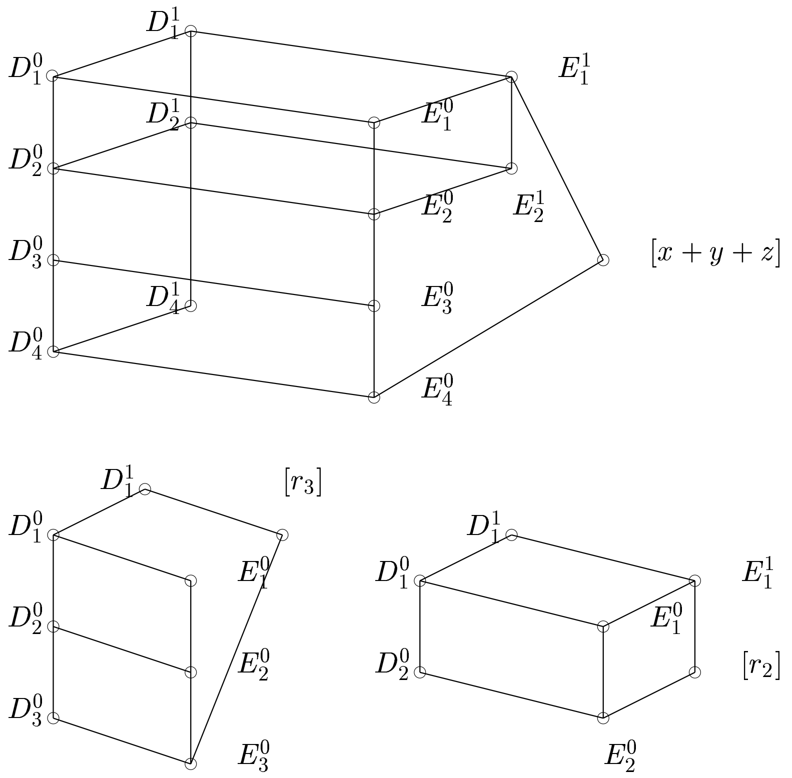

FIGURE 2. Lattice of homogeneous algebras with size ≤ 4

- $D_i^1 = \langle d, l_{i+1}, r_n \rangle$ for $2 \leq i \leq n-2$;
- $E_1^0 = \langle s \rangle$; $E_n^0 = \mathcal{I}$
- $E_i^0 = \langle l_{i+1} \rangle$ for $2 \leq i \leq n-1$;
- $E_1^1 = \langle s, r_n \rangle$; $E_{n-2}^1 = \langle r_n \rangle$;
- $E_i^1 = \langle l_{i+1}, r_n \rangle$ for $2 \leq i \leq n-3$;

Since the classification of homogeneous algebras is complete we can study the tractability of every element in the classification individually. Let us consider first some tractable clones:

- The dual discriminator d is a near-unanimity operation and, hence, tractable. In consequence, D_i^j $(1 \leq i,\ \ j \in \{0,1\})$ is tractable, since d is contained in them.
- The switching operation s is tractable and therefore E_1^j $j \in \{0,1\}$ are tractable.

Now consider the remaining clones. Fix two arbitrary values $a, b \in A$. Notice that operations l_k $(3 \leq k \leq n)$ and r_n $(n \geq 4)$ preserve all the relations over any 2-element subdomain $\{a, b\}$, and in particular the NOT-ALL-EQUAL with set of tuples

$$\begin{array}{ll} \{ & \langle a,a,b\rangle \\ & \langle a,b,a\rangle \\ & \langle b,a,a\rangle \\ & \langle b,b,a\rangle \\ & \langle b,a,b\rangle \\ & \langle b,b,a\rangle \quad \} \end{array}$$

which is known to be NP-complete [33]. In consequence, E_i^0 $(2 \leq i \leq n)$ and E_i^1 $(2 \leq i \leq n-2)$ are not tractable.

Thus, homogeneous algebras with $|A| = n \geq 5$ are completely classified, being the tractable clones the ones containing the switching operation s and the dual discriminator d.

For $n = 4$ we have to study the special case of the clone generated by $(x + y + z)$ where $(A; +)$ is a group of exponent 2. Operation $(x + y + z)$ is affine and, in consequence, the clone is tractable. Thus, for $n = 4$ the only tractable classes are the clones containing s, d or $(x + y + z)$.

For $n = 3$ we have to consider the particular case $\langle r_3\rangle$. Clearly, operation r_3 is tractable since $r_3(r_3(x, y), r_3(y, x)) = 2(2(x + y) + 2(y + z)) = 4x + 8y + 4z = x + 2y + z = x - y + z$. Thus, the tractable classes for $n = 3$ are the clones containing s, d, and r_3.

For $n = 2$, we have to consider operation r_2, which is the non-identity permutation. Clearly r_2 preserves the NOT-ALL-EQUAL relation and therefore clone $\langle r_2\rangle$ is not tractable. Summarizing, the tractable classes for $n = 2$ are the clones containing d and s.

The results in this section can be summarized in the following theorem.

Theorem 12. *Let* Γ *be a set of relations on* A. *Then* $\mathrm{CSP}^*(\Gamma)$ *is polynomial-time solvable iff* $\mathrm{Pol}(\Gamma)$ *contains the dual discriminator* d*; or the switching operation* s*; or an affine operation. More precisely,* $\mathrm{CSP}^*(\Gamma)$ *is polynomial-time solvable if and only if*

- $\{s, d\} \cap \mathrm{Pol}(\Gamma) \neq \emptyset$ *if* $|A| \notin \{3, 4\}$.
- $\{s, d, x + y + z\} \cap \mathrm{Pol}(\Gamma) \neq \emptyset$*, where* $(A; +)$ *is a group of exponent 2, if* $|A| = 4$.
- $\{s, d, r_3\} \cap \mathrm{Pol}(\Gamma) \neq \emptyset$*, if* $|A| = 3$.

Otherwise, $\mathrm{CSP}^*(\Gamma)$ *is NP-complete.*

■

8. Final Remarks

A considerable amount of progress has been done in the identification tractable cases of CSP since this work was submitted. In fact, new tractable classes of the CSP have been found by means of the algebraic approach used in this paper, such that the picture of the "islands of tractability" presented in the introduction is no longer complete. Up to the author's knowledge, any exhaustive list of tractable classes should include also the class of conservative commutative grupoids [6], or even more generally, the class of 2-semilattices [1] which generalize the class of ACI operations. Another different extension of ACI operations is the class of all block groups introduced in [5]. Some additional tractable cases are obtained as a byproduct in the complete classification of the 3-element domain in [2] and of conservative CSPs in [4]. Lately, in [3] the family of all Malt'sev operations is shown to be tractable. Para-primal algebras are also Malt'sev algebras so that the results in [3] extend the tractability of para-primal algebras presented in this paper. Nevertheless, the results presented in this paper were found independently and, in fact, much earlier than those in [3].

References

[1] A. Bulatov. Combinatorial Problems Raised from Semilattices. manuscript.

[2] A. Bulatov. A dichotomy theorem for constraints on a three-element set. In *Proceedings of the 43rd IEEE Symposium on Foundations of Computer Science (FOCS'02)*, pages 649–658, 2002.

[3] A. Bulatov. Malt'sev Constraints Are Tractable. Technical Report PRG-RR-02-05, Computing Laboratory, Oxford University, 2002.

[4] A. Bulatov. Tractable Conservative Constraint Satisfaction Problems. In *Proceedings of 18th IEEE Symposium on Logig in Computer Science (LICS,03)*, pages 321–330, 2003.

[5] A. Bulatov, P. Jeavons, and M. Volkov. Finite Semigroups Imposing Tractable Constraints. In *Proceedings of the School on Algorithmic Aspects of the Theory of Semigroups and its Applications*, pages 313–329, 2002.

[6] A. Bulatov and P. G. Jeavons. Tractable Constraints Closed under a Binary Operation. Technical Report PRG-TR-12-00, Computing Laboratory, Oxford University, 2000.

[7] A. Bulatov, A. Krokhin, and P. Jeavons. Constraint Satisfaction Problems and Finite Algebras. Technical Report PRG-TR-4-99, Computing Laboratory, Oxford University, 1999.

[8] D.M. Clark and P.H. Krauss. Para Primal Algebras. *Algebra Universalis*, 6:165–192, 1976.

[9] D.M Clark and P.H. Krauss. Plain Para Primal Algebras. *Algebra Universalis*, 11:365–388, 1980.

[10] S.A. Cook. The Complexity of Theorem-Proving Procedures. In *3rd Annual ACM Symposium on Theory of Computing STOC'71*, pages 151–158, 1971.

[11] M.C. Cooper, D.A. Cohen, and P.G. Jeavons. Characterizing Tractable Constraints. *Artificial Intelligence*, 65:347–361, 1994.

[12] V. Dalmau and P. Jeavons. Learnability of Quantified Formulas. In *4th European Conference on Computational Learning Theory Eurocolt'99*, volume 1572 of *Lecture Notes in Artificial Intelligence*, pages 63–78, Berlin/New York, 1999. Springer-Verlag.

[13] V. Dalmau and J. Pearson. Set Functions and Width 1. In *5th International Conference on Principles and Practice of Constraint Programming, CP'99*, volume 1713 of *Lecture Notes in Computer Science*, pages 159–173, Berlin/New York, 1999. Springer-Verlag.

[14] R. Dechter and J. Pearl. Network-based Heuristics for Constraint Satisfaction Problems. *Artificial Intelligence*, 34(1):1–38, 1988.

[15] T. Feder and M.Y. Vardi. The Computational Structure of Monotone Monadic SNP and Contraint Satisfaction: A Study through Datalog and Group Theory. *SIAM J. Computing*, 28(1):57–104, 1998.

[16] E.C. Freuder. A Sufficient Condition for Backtrack-bounded Search. *Journal of the ACM*, 32:755–761, 1985.

[17] M. Gyssens, P. Jeavons, and D. Cohen. Decomposing Constraint Satisfaction Problems using Database Techniques. *Artificial Intelligence*, 66(1):57–89, 1994.

[18] P. Hell and J. Nešetřil. On the Complexity of H-coloring. *J. Comb. Theory, Series B*, 48:92–110, 1990.

[19] P. Jeavons. On the Algebraic Structure of Combinatorial Problems. *Theoretical Computer Science*, 200:185–204, 1998.

[20] P. Jeavons, D. Cohen, and M.C. Cooper. Constraints, Consistency and Closure. *Artificial Intelligence*, 101:251–265, 1998.

[21] P. Jeavons, D. Cohen, and M. Gyssens. A Unifying Framework for Tractable Constraints. In *1st International Conference on Principles and Practice of Constraint Programming, CP'95, Cassis (France), September 1995*, volume 976 of *Lecture Notes in Computer Science*, pages 276–291, Berlin/New York, 1995. Springer-Verlag.

[22] P. Jeavons, D. Cohen, and M. Gyssens. A Test for Tractability. In *2nd International Conference on Principles and Practice of Constraint Programming CP'96*, volume 1118 of *Lecture Notes in Computer Science*, pages 267–281, Berlin/New York, August 1996. Springer-Verlag.

[23] P. Jeavons, D. Cohen, and M. Gyssens. Closure Properties of Constraints. *Journal of the ACM*, 44(4):527–548, July 1997.

[24] P. Jeavons and M. Cooper. Tractable Constraints on Ordered Domains. *Artificial Intelligence*, 79:327–339, 1996.

[25] L. Kirousis. Fast Parallel Constraint Satisfaction. *Artificial Intelligence*, 64:147–160, 1993.

[26] A. K. Mackworth. Consistency in networks of relations. *Artificial Intelligence*, 8:99–118, 1977.

[27] S.S. Marchenkov. Homogeneous Algebras. *Problemy Kibernetiki*, 39:85–106, 1982.

[28] E. Marczewski. Homogeneous Algebras and Homogeneous Operations. *Fund. Math.*, 56:81–103, 1964.

[29] R.N. McKenzie, G.F. McNulty, and W.F. Taylor. *Algebras, Lattices and Varieties*, volume 1. Wadsworth and Brooks, 1987.

[30] U. Montanari. Networks of Constraints: Fundamental Properties and Applications to Picture Processing. *Information Sciences*, 7:95–132, 1974.

[31] U. Montanari and F. Rossi. Constraint Relaxation may be Perfect. *Artificial Intelligence*, 48:143–170, 1991.

[32] R.W. Quackenbush. Minimal Para Primal Algebras. In *Contributions to General Algebra 2, Proc. Klagenfurt Conf. 1982*, pages 291–304, 1983.

[33] T.J. Schaefer. The Complexity of Satisfiability Problems. In *10th Annual ACM Symposium on Theory of Computing*, pages 216–226, 1978.

[34] A. Szendrei. *Clones in Universal Algebra*, volume 99 of *Seminaires de Mathématiques Supéreiores*. University of Montreal, 1986.

[35] P. van Beek and R. Dechter. On the Minimality and Decomposability of Row-convex Constraint Networks. *Journal of the ACM*, 42:543–561, 1995.

[36] P. van Hentenryck, Y. Deville, and C-M. Teng. A Generic Arc-consistency Algorithm and its Specializations. *Artificial Intelligence*, 1992.