Electronic Colloquium on Computational Complexity, Report No. 59 (2005)

ECCC

Tractable Clones of Polynomials over Semigroups

Víctor Dalmau*[1], Ricard Gavaldà**[2], Pascal Tesson***[3], and Denis Thérien†[4]

[1] Departament de Tecnologia, Universitat Pompeu Fabra
victor.dalmau@upf.edu
[2] Department of Software (LSI), Universitat Politècnica de Catalunya
gavalda@lsi.upc.edu
[3] Département d'Informatique et de Génie Logiciel, Université Laval
pascal.tesson@ift.ulaval.ca
[4] School of Computer Science, McGill University
denis@cs.mcgill.ca

Abstract. It is well known that coset-generating relations lead to tractable constraint satisfaction problems. These are precisely the relations closed under the operation $xy^{-1}z$ where the multiplication is taken in some finite group. Bulatov et al. have on the other hand shown that any clone containing the multiplication of some "block-group" (a particular class of semigroups) also yields a tractable CSP. We consider more systematically the tractability of $\mathrm{CSP}(\Gamma)$ when Γ is a set of relations closed under operations that are expressible as polynomials over a finite semigroup. In particular, we unite the two results above by showing that if S is a block-group of exponent ω and Γ is a set of relations over S preserved by the operation defined by the polynomial $f(x,y,z) = xy^{\omega-1}z$ over S, then $\mathrm{CSP}(\Gamma)$ is tractable. We show one application of this result by reproving an upper bound by Klíma et al. on the complexity of solving systems of equations over certain block-groups.

We show that if S is a commutative semigroup and $\mathcal{C}$ is an idempotent clone consisting of polynomials over S, then $\mathcal{C}$ is tractable iff it contains the polynomial $xy^{\omega-1}z$. If S is a nilpotent group, we show that a clone of polynomials over S is tractable iff it contains a Malt'sev operation, and conjecture that this holds for all groups.

* Research partially supported by the MCyT under grants TIC 2002-04470-C03 and TIC 2002-04019-C03, the EU PASCAL Network of Excellence, IST-2002-506778, and the MODNET Marie Curie Research Training Network, MRTN-CT-2004-512234.

** Research partially supported by the EU PASCAL Network of Excellence, IST-2002-506778. Partly done while visiting McGill University, supported in part by the AIRE and ACI programs of the DURSI of the Generalitat de Catalunya.

*** Part of this research took place while the author was at the University of Tübingen, supported by the von Humboldt Foundation.

† Research supported in part by NSERC, FQRNT and the von Humboldt Foundation.

ISSN 1433-8092

1 Introduction

Constraint satisfaction problems (CSPs) provide a natural way to study in a unified framework a number of combinatorial problems arising in various areas of computer science. An instance of CSP consists of a list of variables, a domain and a set of constraints relating the variables and we ask whether the variables can be assigned domain values such that all constraints are satisfied.

In general, the CSP problem is NP-complete and one thus tries to identify tractable (i.e., polynomial-time solvable) restrictions of the problem. In particular, a lot of attention has been paid to the case where the relations available to construct constraints lie in a fixed finite set Γ of relations over a finite domain. It is conjectured that for any such Γ the problem $\mathrm{CSP}(\Gamma)$ is always either tractable or NP-complete [10]. An algebraic approach has been particularly successful in making progress on this question [12, 11]: it was shown that the tractability of $\mathrm{CSP}(\Gamma)$ depends on the algebraic properties of the set of operations under which all relations of Γ are closed. This has lead to the identification of very broad classes of sets of relations (often called *islands of tractability*) for which $\mathrm{CSP}(\Gamma)$ is known to have polynomial-time algorithms [2, 7, 8, 10, 5] and has validated the conjecture mentioned above for domains of size two [16] and three [1].

Feder and Vardi [10] have shown that if a set Γ of relations over a finite group G is coset-generating (i.e. every R in Γ is a coset of a power of the group), then $\mathrm{CSP}(\Gamma)$ is tractable. Equivalently, Γ is coset-generating iff it is closed under the ternary operation $t(x, y, z) = xy^{-1}z$ where multiplication is taken in the group. Another island of tractability uncovered by Bulatov, Jeavons, and Volkov [5] states that $\mathrm{CSP}(\Gamma)$ is tractable if Γ is closed under the multiplication of a particular type of semigroup called a block-group. This result generalizes a previous result of Jeavons, Cohen, and Gyssens [12] where multiplication was taken in a semilattice.

In light of these two results, we consider more systematically classes Γ of relations whose closure properties can, as above, be expressed using polynomials over a semigroup. Our long term objective is to classify all corresponding problems $\mathrm{CSP}(\Gamma)$ as either tractable or NP-complete. Our first result gives a new sufficient condition for the tractability of $\mathrm{CSP}(\Gamma)$. We show that if S is a block-group of exponent ω and Γ is a set of relations over S that are preserved under the ternary operation $f(x, y, z) = xy^{\omega-1}z$ then $\mathrm{CSP}(\Gamma)$ is tractable. This result generalizes both of the results [10, 5] just mentioned. We show that our theorem can be applied to reprove an upper bound of [13] on the complexity of solving systems of equations over a certain subclass of block-groups.

Next we consider necessary conditions for tractability. These are expressed in terms of *clones,* or sets of such closure operations – definitions are given in Section 2. For technical reasons, we restrict ourselves to *idempotent* clones, which are known to still determine the complexity of $\mathrm{CSP}(\Gamma)$. We show:

- If S is a commutative semigroup the sufficient condition given by the previous theorem is also necessary: if $\mathcal{C}$ is a nontrivial idempotent clone of polynomials over S, then $\mathcal{C}$ is tractable iff it contains the operation $xy^{\omega-1}z$.

– If S is a nilpotent group and $\mathcal{C}$ is a nontrivial clone of polynomials over S then $\mathcal{C}$ is tractable iff it contains a Malt'sev operation. This is a type of operations of which $xy^{\omega-1}z$ is a prime example and that is known to imply tractability [2, 3]. We conjecture that this in fact holds for any finite group.

The paper is organized as follows. Section 2 presents the required notions of semigroup theory as well as an introduction to constraint satisfaction problems and the algebraic approach to their study. In Section 3 we prove the new sufficient condition for the tractability of CSP(Γ) and give one application of the result. In Section 4 we consider sets of relations Γ that are closed solely under operations that can be described by polynomials over a semigroup and investigate necessary and sufficient conditions for the tractability of CSP(Γ).

2 Preliminaries and Background

2.1 Finite Semigroups

A semigroup is a set S with a binary associative operation that we denote multiplicatively as $\cdot_S$ or $\cdot$ when no ambiguity exists. An element $s \in S$ is said to be idempotent if it is its own square, i.e. $s^2 = s$. In this paper we are solely concerned with finite semigroups, and in that case there exists a minimal integer ω such that for all $s \in S$ the element s^ω is idempotent. We call ω the *exponent* of the semigroup. If S is a group then s^ω is the identity element of the group since it is the only idempotent element.

A class of finite semigroups $\mathbf{V}$ is a *pseudo-variety* if it is closed under finite direct products and formation of subsemigroups and homomorphic images. Some of the pseudo-varieties that we will use are:

– **SL**, the pseudo-variety of finite semilattices, i.e. of commutative semigroups in which every element is idempotent;
– **Ab**, the pseudo-variety of finite Abelian groups;
– **BG**, the pseudo-variety of *block-groups*, or semigroups that satisfy the identity $(x^\omega y^\omega)^\omega = (y^\omega x^\omega)^\omega$.

Our main theorem concerns block-groups, an important class in the theory of finite semigroups that admits a number of interesting characterizations [15]. We state some of their relevant properties. Most useful to us will be the following: the finite semigroup S is *not* a block-group iff it contains two distinct idempotents e, f such that $ef = e$ and $fe = f$ or such that $ef = f$ and $fe = e$. It can easily be deduced that semilattices and groups are special cases of block-groups, but not all block-groups are in the pseudo-variety generated by semilattices and groups.

Let $E_S : \{s^\omega : s \in S\}$ be the set of idempotents of S. If $S \in \mathbf{BG}$, the subsemigroup of S generated by E_S satisfies $(xy)^\omega = (yx)^\omega$. This can be used to show that if two sequences $e_1, \ldots, e_n$ and $f_1, \ldots, f_m$ of idempotents of S satisfy $\{e_1, \ldots, e_n\} = \{f_1, \ldots, f_m\}$ (as sets), then $(e_1 \ldots e_n)^\omega = (f_1 \ldots f_m)^\omega$.

For any semigroup S and any idempotent $e \in S$, the set of elements s such that $es = se = s$ forms a subgroup G_e with identity element e. For all $s \in G_e$

we have $s^\omega = e$ and thus $s^{\omega+1} = s$. We will say that $s \in S$ is a *subgroup element* if it lies in some G_e. We will say that a semigroup is a *union of groups* if all its elements are subgroup elements.

2.2 CSPs and Universal Algebra

Let D be a finite domain and Γ be a finite set of relations over D. In the sequel, D and Γ will always denote respectively a finite domain and a finite set of relations over that domain. The constraint satisfaction problem over Γ, denoted $\mathrm{CSP}(\Gamma)$ is the following decision problem. The input consists of a list of variables $x_1, \ldots, x_n$ and constraints that are pairs (S_i, R_i) where R_i is a k_i-ary relation in Γ and S_i, the scope of the constraint, is an ordered list of k_i variables. We ask whether the variables can be assigned values in D such that every constraint is satisfied. It is conjectured that for any Γ the problem $\mathrm{CSP}(\Gamma)$ is either tractable or NP-complete [10]. Over the last ten years, a lot of ground was covered towards establishing this conjecture using an algebraic approach pioneered by [12] that considers the closure properties of Γ, as we next explain formally.

An *operation* f on D is simply a function $f : D^t \to D$. We naturally extend f so that it takes as inputs t k-tuples $\overline{a_1}, \ldots, \overline{a_t}$ of values in D by defining

$$f(\overline{a_1}, \ldots, \overline{a_t}) = (f(a_{11}, \ldots, a_{t1}), \ldots, f(a_{1k}, \ldots, a_{tk})).$$

We say that a k-ary relation R over D is *closed under* f if for any t k-tuples of R, say $\overline{a_1}, \ldots, \overline{a_t}$ we also have $f(\overline{a_1}, \ldots, \overline{a_t}) \in R$.

By extension we say that Γ is closed under f if every relation of Γ is closed under f, and denote as $\mathrm{Pol}(\Gamma)$ the set of all such finitary operations f (the notation is due to the fact that every such f is called a *polymorphism* of Γ in universal algebra). The fundamental link to the complexity of CSPs is the following theorem.

Theorem 1 ([11]). *If Γ_1, Γ_2 are sets of relations over D such that $\mathrm{Pol}(\Gamma_1) \subseteq \mathrm{Pol}(\Gamma_2)$ then $\mathrm{CSP}(\Gamma_2)$ is polynomial-time reducible to $\mathrm{CSP}(\Gamma_1)$.*

The following is a crucial property of all the sets of the form $\mathrm{Pol}(\Gamma)$.

Lemma 2 (see e.g. [12]). *For any set of relations Γ over D: (1) $\mathrm{Pol}(\Gamma)$ contains all the projection functions $\pi_{i,n}(x_1, \ldots, x_n) = x_i$. (2) If g is a k-ary operation in $\mathrm{Pol}(\Gamma)$ and $f_1, \ldots, f_k$ are t-ary operations in $\mathrm{Pol}(\Gamma)$, then their composition*

$$g(f_1, \ldots, f_k)(x_1, \ldots, x_t) = g(f_1(x_1, \ldots, x_t), \ldots, f_k(x_1, \ldots, x_t))$$

is also in $\mathrm{Pol}(\Gamma)$.

Note that from (1) and (2) it follows that $\mathrm{Pol}(\Gamma)$ is also closed under identification of variables, since this can be obtained by composition with projections.

In universal algebra lingo, a set of operations containing all the projections and closed under composition is called a *clone*. For a set of operations F, we denote by $\langle F \rangle$ the clone generated by F, i.e. the smallest clone containing F.

Using the connection between CSP(Γ) and Pol(Γ) given by Theorem 1, we say that a clone $\mathcal{C}$ is *tractable* if CSP(Γ) is tractable for every Γ such that $\mathcal{C} \subseteq \text{Pol}(\Gamma)$. On the other hand, we say that $\mathcal{C}$ is NP-complete if there exists a set of relations Γ such that $\mathcal{C} \subseteq \text{Pol}(\Gamma)$ and CSP(Γ) is NP-complete. Note that we will shamelessly assume $P \neq NP$.

We can thus view the task of resolving the CSP conjecture as that of proving that any clone is either tractable or NP-complete. An important simplification is known [6]: in order to obtain such a classification it suffices in fact to consider clones in which every operation f is *idempotent*, i.e. satisfies $f(x, \ldots, x) = x$. We call these the *idempotent clones*.

The first half of this task is to identify tractable clones and many such "islands of tractability" have already been identified in this way. For example, a ternary operation $M(x, y, z)$ is said to be *Malt'sev* if it satisfies $M(x, x, y) = y$ and $M(x, y, y) = x$: Bulatov showed that any clone containing a Malt'sev operation is tractable [2, 3]. This very general result covers an important special case first identified as tractable by [10]: Suppose that the domain D is a finite group and that any k-ary relation of Γ is a coset of a subgroup of D^k. We then say that Γ is *coset generating,* and it can be verified from the definition of a coset that Γ is closed under the operation $M(x, y, z) = x \cdot y^{-1} \cdot z$ (where multiplication and inverse are those of the group D). This operation is Malt'sev since $M(x, x, y) = xx^{-1}y = y$ and $M(x, y, y) = xy^{-1}y = x$ as required.

Bulatov et al. [5] considered the tractability of clones generated by a semigroup, i.e. generated by the binary operation $x \cdot_S y$ for some semigroup S. They showed that the clone $\langle \cdot_S \rangle$ is tractable if S is a block-group and NP-complete otherwise. This result extends another well-known result stating that any Γ closed under the multiplication in a semilattice is tractable [12]. In both cases, it is shown that such CSPs are solved by an arc-consistency algorithm.

For the remainder of this paper, we focus on clones whose operations can all be described by expressions over a finite semigroup S. Formally, a *polynomial* P over the semigroup S is simply a finite sequence $P = x_{i_1} \cdots x_{i_m}$ of (possibly repeating) variables. A polynomial containing k distinct variables naturally defines a k-ary function, but in order to express the projections with such polynomials we allow for unused variables and e.g. represent the projection $\pi_{i,n}(x_1, \ldots, x_n)$ by the polynomial x_i. We say that a clone is a *clone of polynomials* if every operation in the clone can be represented in this way. Note that the composition of polynomials is again a polynomial so that the clone generated by a set of polynomials is indeed a clone of polynomials.

3 A Polynomial that Guarantees Tractability

Our main goal in this section is to prove the following sufficient condition for the tractability of a clone.

Theorem 3. *If S is a block-group and $\mathcal{C}$ is a clone containing the polynomial $xy^{\omega-1}z$ then $\mathcal{C}$ is tractable.*

Note that when S is a group, this condition is equivalent to saying that every Γ such that $\mathcal{C} \subseteq \mathrm{Pol}(\Gamma)$ is coset-generating. Also if $\mathcal{C}$ is generated by $\cdot_S$ for some block-group S then in particular it must contain the polynomial $xy^{\omega-1}z$, which can be obtained from xy by composition. Hence this result generalizes the results of [10, 5] mentioned before.

Proof (Theorem 3). Let $\mathcal{P}$ be an instance of $\mathrm{CSP}(\Gamma)$, with $\mathcal{C} \subseteq \mathrm{Pol}(\Gamma)$. If $\mathcal{P}$ has any solution then it has one in which every variable x_i has a value a_i that is a subgroup element. Indeed, by the closure properties of Γ, if $\overline{a}$ is a solution, then $\overline{a}\overline{a}^{\omega-1}\overline{a} = \overline{a}^{\omega+1}$ also is and every element of the latter is a subgroup element.

We will give a polynomial-time algorithm to solve $\mathcal{P}$, which will work in two stages. In the first stage, we will assign to every variable x_i some subgroup G_{e_i} such that if $\mathcal{P}$ is satisfiable then it is satisfiable by an assignment that sets each x_i to a value in G_{e_i}. We will do this by using an arc-consistency procedure. In the second stage we will reduce the $\mathrm{CSP}(\Gamma)$ problem to an instance of $\mathrm{CSP}(\Lambda)$, where Λ is a coset-generating set of relations over the direct product of the subgroups G_e, and then solve this CSP with the algorithm of [10].

We begin by enforcing arc-consistency for $\mathcal{P}$. To every variable x_i, we associate a set of possible values $V_i \subseteq S$. We find the largest V_i s.t. for any constraint of $\mathcal{P}$, say $(\{x_{i_1}, \ldots, x_{i_r}\}, R)$ and for any value $a_{i_j} \in V_{i_j}$ there exist $a_{i_k} \in V_{i_k}$ s.t. $(a_{i_1}, \ldots, a_{i_r}) \in R$. It is well known that this can be done in polynomial time by initializing each V_i to S and gradually removing values that violate the above requirement. Also, if any V_i becomes empty we know that $\mathcal{P}$ has no solution.

If $V_1, \ldots, V_n$ are the sets produced by the arc-consistency algorithm, we define e_i to be the idempotent $(\prod_{a \in V_i} a^\omega)^\omega$. Recall that since S is a block group, the value of a product of the form $(s_1^\omega \ldots s_t^\omega)^\omega$ depends solely on the set $\{s_1, \ldots, s_t\}$ and our definition of e_i is thus sound.

Lemma 4. *If $\mathcal{P}$ has a solution then it has one in which each variable x_i is assigned a value a_i that lies in the subgroup G_{e_i} where the e_i are the idempotents obtained through the arc-consistency algorithm.*

Proof. Let $\overline{b} \in S^n$ be any solution to $\mathcal{P}$. We claim that $\overline{a} = \overline{e}\overline{b}^{\omega+1}$ is then a solution of $\mathcal{P}$ satisfying $\overline{a}^\omega = \overline{e}$.

Since Γ is closed under the operation $xy^{\omega-1}z$, it is also closed under $x^\omega z$ (by identifying x and y) and under $x^\omega y^\omega z$ (by substituting $y^\omega z$ for z in the previous polynomial). By iterating this procedure, we get that for any n, Γ is closed under the polynomial $F(x_1, \ldots, x_{n+1}) = (x_1^\omega \ldots x_n^\omega)^\omega x_{n+1}^{\omega+1}$.

Consider any constraint of $\mathcal{P}$, e.g. $(\{x_{i_1}, \ldots, x_{i_k}\}, R)$. By assumption, we have $\overline{c} = (b_{i_1}, \ldots, b_{i_k}) \in R$. For a tuple $\overline{t}$, we denote as $\overline{t}[j]$ the jth coordinate of $\overline{t}$. Let $\{\overline{t}_1, \ldots, \overline{t}_m\}$ be the tuples of R such that each $\overline{t}_k[j]$ lies in V_{i_j}. Since we have enforced arc-consistency, we have in fact $V_{i_j} = \{\overline{t}_k[j] : k = 1 \ldots m\}$ and we can thus deduce that $e_{i_j} = (\prod_{k=1\ldots m} \overline{t}_k[j])^\omega$ for each j. By the closure properties of Γ we also know that $F(\overline{t}_1, \ldots, \overline{t}_m, \overline{c})$ is in R. Hence,

$$F(\overline{t}_1, \ldots, \overline{t}_m, \overline{c})[j] = (\overline{t}_1[j]^\omega \ldots \overline{t}_n[j]^\omega)^\omega b_{i_j}^{\omega+1} = e_{i_j} b_{i_j}^{\omega+1}$$

Therefore $\overline{e}\overline{b}^{\omega+1}$ is indeed a solution to $\mathcal{P}$. The same methods allow one to further verify that $(\overline{e}\overline{b}^{\omega+1})^{\omega} = \overline{e}$. □

Thus, in polynomial time, we can associate to each variable x_i a subgroup G_{e_i} such that if CSP(Γ) has any solution then it has one where x_i is assigned a value in G_{e_i}. Let G be the direct product $\prod G_e$ where the product is taken over the n distinct idempotents e of S. We will identify any element lying in one of the n canonical subgroups G_e of G with the corresponding element in the subgroup G_e of S. For any k-ary relation $R \in \Gamma$ and any k-idempotents $e_{i_1}, \ldots, e_{i_k}$ (not necessarily distinct) we define the relation $R_{e_{i_1},\ldots,e_{i_k}} \subseteq G^k$ as consisting of tuples $(a_1, \ldots, a_k)$ such that[1]

1. a_j lies in the subgroup G_{e_j} of G;
2. $(a_1, \ldots, a_k) \in R$ when we view the a_j's as elements of S.

The crucial observation is that each $R_{e_{i_1},\ldots,e_{i_k}}$ is coset-generating, i.e. closed under the operation $xy^{\omega-1}z$. Indeed, if $\overline{a}, \overline{b}, \overline{c} \in R_{e_{i_1},\ldots,e_{i_k}}$ then certainly the jth component of $\overline{a}\overline{b}^{\omega-1}\overline{c}$ also lies in the subgroup G_{e_j}. Furthermore, the second condition is satisfied since R is closed under $xy^{\omega-1}z$.

Let Λ consist of relations over G of the form $R_{e_{i_1},\ldots,e_{i_k}}$ for some $R \in \Gamma$. Given an instance of CSP(Γ) in which every variable has been restricted to lie in some particular subgroup we can naturally construct an instance of CSP(Λ) that will be satisfiable iff the instance of CSP(Γ) can be satisfied. Since Λ is coset-generating, we can solve CSP(Λ) in polynomial time. □

As we mentioned earlier, the "island of tractability" uncovered by this theorem subsumes the tractability results for coset-generating relations of [10] and for clones generated by a block-group [5]. We give an application of this theorem to a problem studied in [13]: for a finite semigroup S, let EQN$^*_S$ denote the problem of determining whether a system of equations over S has a solution. Note that by introducing dummy variables we can assume that a system of equations over S consists only of equations of the form $xy = z$ or $x = y$ where x, y, z are variables or constants. We can thus think of the problem EQN$^*_S$ as CSP(Γ_S) where Γ_S is the set of relations definable by such an equation over S.

Theorem 5 ([13]). *Let* EQN$^*_S$ *be the problem of testing whether a system of equations over the semigroup S has a solution. If S is in* **SL** $\vee$ **Ab** *(the pseudovariety generated by* **SL** *and* **Ab***) then* EQN$^*_S$ *lies in P.*

Proof. Any semigroup in **SL** $\vee$ **Ab** is commutative and satisfies $x^{\omega+1} = x$ since both semilattices and Abelian groups have these properties. Consider an equation over S of the form $x_1x_2 = x_3$. If (a_1, a_2, a_3), (b_1, b_2, b_3), and (c_1, c_2, c_3) are solutions of this equation then we have by commutativity

$$a_1b_1^{\omega-1}c_1a_2b_2^{\omega-1}c_2 = a_1a_2(b_1b_2)^{\omega-1}c_1c_2 = a_3b_3^{\omega-1}c_3.$$

[1] Alternatively, we could view this relation as multi-sorted in the sense of [4].

Similarly, if we consider an equation in which a constant appears, e.g. $sx_1 = x_2$ then since $s = s^{\omega+1}$ we get

$$sa_1 b_1^{\omega-1} c_1 = sa_1 (sb_1)^{\omega-1} sc_1 = a_2 b_2^{\omega-1} c_2.$$

Thus Γ_S is closed under the polynomial $xy^{\omega-1}z$ over the block-group S and EQN_S^* is tractable by Theorem 3. □

One cannot directly infer the tractability of EQN_S^* for $S \in \mathbf{SL} \vee \mathbf{Ab}$ by simply using the tractability of Malt'sev operations or the tractability of clones generated by a block group so the result of Theorem 3 seems required in this case. It is worth noting that if S is a finite monoid then EQN_S^* is NP-complete when S is not in $\mathbf{SL} \vee \mathbf{Ab}$ [13]. An alternative proof of this latter fact was given in [14] using an elegant universal algebra argument.

4 Tractable Clones of Polynomials

We have just shown that for a clone of polynomials over a block-group S to be tractable, a sufficient condition is that it contains the operation $xy^{\omega-1}z$. In this section, we consider necessary conditions for tractability. Our goal is to eventually be able to classify all clones of polynomials over any S as either tractable or NP-complete. We first note that this question is only of real interest if S is a block-group.

Theorem 6. *If S is not a block-group, any clone of polynomials over S is NP-complete.*

Proof. Since S is not a block group, there exist idempotents $e, f \in S$ such that $ef = e$ and $fe = f$ or $ef = f$ and $fe = e$. Suppose w.l.o.g. that the former occurs and let $P = x_{i_1} \dots x_{i_m}$ be a polynomial in the clone. If we set all variables to one of e, f this polynomial will always return the value of x_{i_1} since $ef = ee = e$ and $fe = ff = f$. Hence, as first observed in [5], this operation preserves any relation over the subdomain $\{e, f\}$ and since there are NP-complete CSPs over a binary domain (e.g. 3SAT), the clone is NP-complete. □

We also provide a condition on clones of polynomials that guarantees NP-completeness over any semigroup. We need an extra semigroup-theoretic notion: the *subgroup exponent* η of the semigroup S is the least common multiple of the exponents of the subgroups in S. When S is a union of groups, we have $\eta = \omega$ but in general we can only say that η is a divisor of ω. We say that an operation (polynomial) $x_{i_1}^{n_1} \cdots x_{i_r}^{n_r}$ is a *d-factor* if $d > 1$, d is a divisor of η (possibly η itself), and $|\{1 \leq i \leq r : d | n_i\}| = r - 1$, that is, if every n_i but one is divided by d. We say that the clone $\mathcal{C}$ is a d-factor if every operation in $\mathcal{C}$ is a d-factor.

Theorem 7. *If $\mathcal{C}$ is a d-factor for some d then $\mathcal{C}$ is NP-complete.*

Proof. We have that $\eta = da$ for some $1 \leq a < \omega$. Since η is the least common multiple of the exponents of all subgroup elements of S, there exists some subgroup element $s \in S$ with exponent η' (i.e. $s^{\eta'}$ is idempotent) such that $lcm(\eta', a) = ka$ for some $k > 1$. Notice that k must divide d. Notice also that for every $1 \leq l < k$, η' does not divide la. Consequently s^a has exponent k.

Consider the subgroup A of S generated by s^a, that is $A = \{s^a, s^{2a}, \ldots, s^{ka}\}$. Clearly $|A| \geq 2$. Let $f(x_1, \ldots, x_q) = x_{i_1}^{n_1} \cdots x_{i_r}^{n_r}$ be any operation in $\mathcal{C}$. First, note that A forms a subuniverse for f: if $a_1, \ldots, a_q \in A$, $f(a_1, \ldots, a_q)$ is in A too. Furthermore, the polynomial $x_{i_1}^{n_1} \cdots x_{i_r}^{n_r}$ is a d-factor so there exists some $1 \leq j \leq r$ such that d divides n_i for every $1 \leq i \leq r$ with $i \neq j$. If again $a_1, \ldots, a_q \in A$, each a_i is $a_i = s^{l_i a}$ for some l_i. So we have $a_{t_i}^{n_i} = s^{l_{t_i} a n_i} = s^{m_i \eta}$, for some m_i, and then $f(a_1, \ldots, a_q) = s^{m\eta + l_{t_j} a n_j}$ for some $m \geq 0$, a value that depends only on x_j. Therefore, if g is is the restriction g of f to A, we have $g(x_1, \ldots, x_q) = s^{m\eta} x_{t_j}^{n_j}$. It is easy to see that $s^{m\eta} x^{n_j}$ is one-to-one over A. Let $s^{k_1 a}$ and $s^{k_2 a}$ be two different elements in A and let $s^{m\eta + k_1 a n_j}$, $s^{m\eta + k_2 a n_j}$ be their corresponding images. In order to be identical, ka has to divide $(k_1 - k_2) a n_j$, or equivalently k has to divide $(k_1 - k_2) n_j$. First notice that since $s^{k_1 a}$ and $s^{k_2 a}$ are different, k cannot divide $k_1 - k_2$. Furthermore, we shall show that $gcd(k, n_j) = 1$. Let p be any common divisor to k and n_j. Since p divides k, it also divides d and then it also divides n_i for every i such that $1 \leq i \neq j \leq r$. If p also divides n_j then it must divide $\eta + 1$, in contradiction with the fact that p divides η.

Summarizing we have shown that there exists a set, namely A, of cardinality at least 2, that is a subuniverse of every operation in $\mathcal{C}$. Furthermore, for every operation $f(x_1, \ldots, x_r)$ in $\mathcal{C}$, the restriction $f_{|A}(x_1, \ldots, x_r)$ of f to A is equivalent to $g(x_j)$ for some one-to-one function $g : A \to A$. It is well known that this implies that $\mathcal{C}$ is NP-complete. □

In the next two subsections we will see that in the cases of commutative semigroups and nilpotent groups, an idempotent clone of polynomials is tractable iff it is not a d-factor.

4.1 The Commutative Case

As we mentioned in Section 2, in order to understand the tractability of clones, it suffices to consider idempotent ones. We concentrate on idempotent clones from now on; this allows us to consider only unions of groups.

Lemma 8. *Let S be a semigroup and $\mathcal{C}$ a nontrivial, idempotent clone over S. Then S is a union of groups and for every polynomial $x_{i_1}^{\alpha_1} \ldots x_{i_m}^{\alpha_m}$ in $\mathcal{C}$, $\sum \alpha_j$ is congruent with 1 modulo ω.*

Proof. Take any operation in $\mathcal{C}$ other than a projection, and suppose it is defined by the polynomial $x_{i_1}^{\alpha_1} \ldots x_{i_m}^{\alpha_m}$. This operation is idempotent if $\mathcal{C}$ is, so the semigroup S must satisfy $s^{\sum \alpha_j} = s$ for all s. Consequently, S must be a union of groups. Furthermore $\sum \alpha_j$ is congruent with 1 modulo ω for otherwise we contradict the minimality of the exponent ω. □

So in the proofs of the rest of the section we will implicitly assume that S is a union of groups and that the polynomials in $\mathcal{C}$ satisfy the condition above.

Theorem 9. *Let S be a commutative semigroup and let $\mathcal{C}$ be a nontrivial idempotent clone of polynomials over S. If $\mathcal{C}$ is not a d-factor for any d, then it contains $xy^{\omega-1}z$.*

Proof. By Lemma 8 S is a union of groups, and is commutative, so we can assume that for every polynomial $x_1^{n_1}\ldots,x_r^{n_r}$ all x_i's are different. Furthermore, since $x^{\omega+1}=x$, every two polynomials $x_1^{n_1}\cdots x_r^{n_r}$ and $x_1^{n'_1}\cdots x_r^{n'_r}$ such that for every $1\leq i\leq r$, $n_i\equiv n'_i \mod \omega$, denote the same operation. We sometimes allow negative indices n_i in an expression, meaning by that any positive integer of the form $n_i+n\omega$. Note also that since we are dealing with unions of groups, the subgroup exponent of S is simply its exponent ω.

We need the following auxiliary lemma.

Lemma 10. *Let $x_1^{n_1}\cdots x_r^{n_r}$ be any operation in $\mathcal{C}$, let $1\leq i\neq j\leq r$, let $a=\gcd(n_in_j,\omega)$, and let $r\geq 1$. Then $x^{ra}y^{\omega-2ra+1}z^{ra}$ belongs to $\mathcal{C}$.*

Proof. First notice that we can identify and rename variables in the expression $x_1^{n_1}\cdots x_r^{m_r}$ to obtain in $\mathcal{C}$ the expression

$$y^{n_1}\cdots y^{n_{i-1}}x^{n_i}y^{n_{i+1}}\cdots y^{n_{j-1}}z^{n_j}y^{n_{j+1}}\cdots y^{n_r},$$

which is equivalent to $x^{n_i}y^{\omega-n_i-n_j+1}z^{n_j}$. By a further replacement, the expression $(y^{n_i}y^{\omega-n_i-n_j+1}x^{n_j})^{n_i}y^{\omega-n_i-n_j+1}(z^{n_i}y^{\omega-n_i-n_j+1}y^{n_j})^{n_j}$ also belongs to $\mathcal{C}$. Setting $c=n_in_j$, this can be rewritten as as $x^cy^{\omega-2c+1}z^c$. Now we will show that for every $m\geq 1$ the expression

$$x_1^{c^m}\cdots x_{2^{m-1}}^{c^m}y^{\omega-2^mc^m+1}z_1^{c^m}\cdots z_{2^{m-1}}^{c^m}$$

belongs to $\mathcal{C}$. We will show it by induction on m. The case $m=1$ has already been proven. Let us assume that the statement holds for m. Then by identification and composition we construct the expression

$$\begin{aligned}(x_1^{c^m}\cdots x_{2^{m-1}}^{c^m}\quad & y^{\omega-2^mc^m+1}x_{2^{m-1}+1}^{c^m}\cdots x_{2^m}^{c^m})^cy^{\omega-2c+1}\\ &\cdot(z_1^{c^m}\cdots z_{2^{m-1}}^{c^m}y^{\omega-2^mc^m+1}z_{2^{m-1}+1}^{c^m}\cdots z_{2^m}^{c^m})^c\\ &=x_1^{c^{m+1}}\cdots x_{2^m}^{c^{m+1}}y^{\omega-2^{m+1}c^{m+1}+1}z_1^{c^{m+1}}\cdots z_{2^m}^{c^{m+1}}.\end{aligned}$$

We are now almost done. There exists some $l\geq 1$ such that $c^{2l}\equiv c^l \mod\omega$. Since $\gcd(c^l,\omega)=\gcd(c,\omega)=a$ we have that for every $r\geq 1$, there exists some integers α,β such that $\alpha c^l+\beta\omega=ra$. We can also assume that $\alpha\geq 0$. Fix some n such that $2^{nl-1}\geq\alpha$. By setting $m=nl$ we can infer that the expression

$$x_1^{c^{nl}}\cdots x_{2^{nl-1}}^{c^{nl}}y^{\omega-2^{nl}c^{nl}+1}z_1^{c^{nl}}\cdots z_{2^{nl-1}}^{c^{nl}}$$

belongs to $\mathcal{C}$, and consequently that the expression

$$\overbrace{x^{c^{nl}}\cdots x^{c^{nl}}}^{\alpha}y^{c^{nl}}\cdots y^{c^{nl}}y^{\omega-2^{nl}c^{nl}+1}y^{c^{nl}}\cdots y^{c^{nl}}\overbrace{z^{c^{nl}}\cdots z^{c^{nl}}}^{\alpha},$$

obtained by identification and renaming of variables also belongs to $\mathcal{C}$. This becomes $x^b y^{\omega-2b+1} z^b$ if we set $b = \alpha c^{nl}$. Finally notice that

$$\alpha c^{nl} \equiv \alpha c^l \equiv ra \mod \omega.$$

□

We now continue the proof of Theorem 9. Assume that $\mathcal{C}$ is not a d-factor for every divisor $d > 1$ of ω. We will show that $xy^{\omega-1}z$ is in $\mathcal{C}$. Let $p_1, \ldots, p_k$ be the set of prime divisors of ω strictly larger than 1. We shall show by induction that for every $1 \leq l \leq k$, there exists some a such that $\gcd(a, p_1 \times \cdots \times p_l) = 1$ and such that $x^a y^{\omega-2a+1} z^a$ belongs to $\mathcal{C}$ (*). Notice that the statement follows from $l = k$ and Lemma 10 since in this case $\gcd(a, \omega) = \gcd(a \times a, \omega) = 1$. The case $l = 1$ is easy. Since $\mathcal{C}$ is nontrivial it contains an operation $x_1^{n_1} \cdots x_r^{n_r}$ and some $i \neq j$ such that p_1 does not divide n_i and does not divide n_j. Consequently p_1 does not divide $a = \gcd(\omega, n_i \times n_j)$. By Lemma 10, $x^a y^{\omega-2a+1} z^a$ belongs to $\mathcal{C}$. Assume now that statement (*) holds for $l < k$. We shall show that it also holds for $l+1$. By induction hypothesis there exists some a with $\gcd(a, p_1 \times \cdots \times p_l) = 1$ such that $x^a y^{\omega-2a+1} z^a$ belongs to $\mathcal{C}$. Also, by a reasoning analogous to the case $l = 1$ we can infer that since $\mathcal{C}$ is not a p_{l+1}-factor we have that there exists some b not divided by p_{l+1} such that $x^b y^{\omega-2b+1} z^b$. We can also assume that $p_1 \times \cdots \times p_l$ divides b.

Let us consider two cases: if p_{l+1} does not divide a then we are done since we have that $\gcd(a, p_1 \times \cdots \times p_{l+1}) = 1$. Otherwise, we proceed as follows. Notice that the operation $x^a (x^b y^{\omega-2b+1} z^b)^{\omega-2a+1} z^a$ belongs to $\mathcal{C}$ since it is obtained by composition and identification of variables. Notice also that if we set $c = a + b(\omega - 2a + 1)$ the previous expression is equivalent to $x^c y^{\omega-2c+1} z^c$. It is easy to see that none of $p_1, \ldots, p_{l+1}$ divides c and consequently $\gcd(c, p_1 \times \cdots \times p_{l+1}) = 1$. To see this, note that for every $1 \leq l' \leq l$, we have that $p_{l'}$ divides b (and consequently $b(\omega - 2a + 1)$) but not a. Consequently $p_{l'}$ cannot divide its sum. Similarly since p_{l+1} divides a, it cannot divide $\omega - 2a + 1$ (since otherwise it would divide $\omega + 1$ in contradiction with the fact that it divides ω). Thus, p_{l+1} does not divide $b(\omega - 2a + 1)$, so it cannot divide c. □

We obtain the following corollary from Theorems 7 and 9, plus the fact that the operation $xy^{\omega-1}z$ implies tractability.

Corollary 11. *Let S be a commutative semigroup and $\mathcal{C}$ a nontrivial, idempotent clone of polynomials over S. Then the following are equivalent:*

- *$\mathcal{C}$ is tractable.*
- *$\mathcal{C}$ is not a d-factor, for any d.*
- *$\mathcal{C}$ contains $xy^{\omega-1}z$.*

If S is a semilattice, we can easily obtain a stronger result.

Theorem 12. *Every non-trivial clone of polynomials over a semilattice S is tractable.*

Proof. Since S is idempotent and commutative, any polynomial over S can be rewritten as simply $x_1 \dots x_k$ and if the clone is non-trivial, it must contain such a polynomial with $k \geq 2$. By identification of variables, the clone contains the polynomial xy and is tractable. □

4.2 The Group Case

We now turn our attention to tractable clones of polynomials over groups. First, we do not think that Corollary 11 can be extended to groups. We believe in particular that the "near subgroup" operation of Feder [9] can in some cases be represented as a polynomial over a group. In any case, the closure function that defines near-subgroup problems is a Malt'sev operation and we conjecture:

Conjecture 13. Let $\mathcal{C}$ be an idempotent clone of polynomials over a group. Then $\mathcal{C}$ is tractable iff it contains a Malt'sev operation.

By Theorem 11, the conjecture is true for Abelian groups because, as noted in Section 2 the operation $xy^{\omega-1}z$ over a finite group is Malt'sev. Next we will show that our conjecture holds for nilpotent groups which, in many ways, form one of the simplest class of non-Abelian groups. A group is said to be *nilpotent* if it is a direct product of p-groups. An alternative description will be more useful for our purposes: Let G be a group. For every $g, h \in G$, the commutator $[g, h]$ of g and h is the element $g^{-1}h^{-1}gh$. Note that $gh = hg[g, h]$, so if g and h commute $[g, h]$ is the identity. An element of G is *central* if it commutes with every element in G; the set of central elements of G is an Abelian subgroup.

For two subgroups G_1 and G_2 of G, $[G_1, G_2]$ is the subgroup generated by all commutators $[g, h]$ with $g \in G_1$ and $h \in G_2$. Define the lower central series of G by $G_0 = G$, and $G_{i+1} = [G_i, G]$. Elements in G_i are called *commutators of weight* $i + 1$ of G. We say that G is *nilpotent class* k if G_k is trivial; note that if G_k is trivial all elements in G_{k-1} are central in G. A group is *nilpotent* if it is nilpotent class k for some k.

We define commutator polynomials analogously. For two polynomials p_1, p_2, $[p_1, p_2]$ is the polynomial $p_1^{\omega-1}p_2^{\omega-1}p_1p_2$. The only commutator polynomial of weight 1 is the empty polynomial. A commutator polynomial of weight 2 is $[x, y]$ for two variables x, y. A commutator polynomial of weight k is $[p, x]$, where x is a variable and p a commutator polynomial of weight $k - 1$. Commutator polynomials of weight $k+1$ or more are the identity in a nilpotent class k group.

Theorem 14. *If $\mathcal{C}$ is an idempotent clone of polynomials over a nilpotent group and not a d-factor for any d, then $\mathcal{C}$ contains a Malt'sev operation.*

Proof. Any polynomial p defining an operation over G also defines an operation over any subgroup H of G since the value of p lies in H when all variables are themselves set to values in H. We will say that a polynomial (or a clone) is interpreted over a subgroup H to mean that the variables of the polynomial take values in H. Note that if $\mathcal{C}$ is not a d-factor interpreted over G it is not a d-factor interpreted over H. Suppose H is an Abelian subgroup and suppose we

can show that $\mathcal{C}$ interpreted over H contains the polynomial x^2y. One cannot conclude that $\mathcal{C}$ interpreted over G contains the polynomial x^2y itself for it may be that it contains, say, xyx that, interpreted over the Abelian subgroup H is indeed the same as x^2y.

Fix a nilpotent class K group G. It is more convenient for the proof to redefine the indices of the central series of G as follows: $G_K = G$ and $G_{k-1} = [G_k, G]$, so that G_1 is the set of central elements of G and G_k is nilpotent of class k. It can be shown that any polynomial p in n variables over a nilpotent group of class k can be rewritten in the following normal form

$$p = \prod_{i=1}^{n} x_i^{\alpha_i} \prod_{i,j \le n} [x_i, x_j]^{\alpha_{ij}} \cdots \prod_{i_1,\ldots,i_k \le n} [\ldots[[x_{i_1}, x_{i_2}]x_{i_3}]\ldots x_{i_k}]^{\alpha_{i_1\ldots i_k}}.$$

In other words, p can be rewritten as a product of distinct commutator polynomials raised to some power, where the "lightest" commutators appear first. Note that if we interpret p over a nilpotent subgroup H of class $k < K$, it is equivalent to the polynomial obtained by deleting all occurrences of commutator polynomials of weight larger than k, since they are the identity over H.

We prove the following for all k, by induction: There is a sequence Q of commutator polynomials of weight at most k over the variables x, y, z such that the operation $xy^{-1}zQ$ is Malt'sev over G and belongs to $\mathcal{C}$ when interpreted over G_k. The theorem follows from this statement for $k = K$.

For $k = 1$, Q is be the empty sequence. As $\mathcal{C}$ interpreted over G_1 is not a d-factor, and G_1 is Abelian, by Corollary 11 $\mathcal{C}$ contains $xy^{-1}z$ when interpreted over G_1.

Inductively, let Q be a sequence of commutator polynomials of weight at most $k-1$ such that $xy^{-1}zQ$ is Malt'sev over G and contained in $\mathcal{C}$ when interpreted over G_{k-1}. Then, $\mathcal{C}$ interpreted over G_k contains a polynomial of the form

$$P(x, y, z) = xy^{-1}z\, Q \cdot \odot_i C_i^{\alpha_i}$$

for some set of exponents α_i, where $\odot_i$ denotes concatenation over all commutator polynomials C_i of weight k on the variables x, y, z.

Let R be obtained by identifying x and y in P. If $\{D_j\}$ is the set of commutator polynomials of weight k in x and z, this identification maps each C_i to some D_j. Every commutator D_j is central in G_k, so we can group all its occurrences, and since $xy^{-1}zQ$ is Malt'sev,

$$R(x, z) = P(x, x, z) = z \cdot \odot_i D_i^{\beta_i}$$

for appropriate exponents β_i. Similarly, let S be obtained by identifying z and y in P, so we have, for appropriate exponents γ_i,

$$S(x, z) = P(x, z, z) = x \cdot \odot_i D_i^{\gamma_i}.$$

Now we move to the domain G_k and remain there until further notice. Let $R^{(2)}$ be obtained by replacing z with $R(x, z)$ in R, that is,

$$\begin{aligned} R^{(2)}(x, z) &= R(x, R(x, z)) = R(x, z) \cdot \odot_i (D_i(x, R(x, z)))^{\beta_i} \\ &= z \cdot \odot_i (D_i(x, z))^{\beta_i} \cdot \odot_i (D_i(x, R(x, z)))^{\beta_i} \end{aligned}$$

Now observe that if c is central, then we have $[u, vc] = u^{-1}(vc)^{-1}u(vc) = u^{-1}v^{-1}uv = [u, v]$ for any u, v. Every D_j has weight k, so it is central in G_k. Thus $D_i(x, R(x, z)) = D_i(x, z)$ over G_k and

$$R^{(2)}(x, z) = z \cdot \odot_i (D_i(x, z))^{2\beta_i}.$$

Iterating this process $\omega - 1$ times, we deduce that the polynomial

$$R^{(\omega-1)}(x, z) = z \cdot \odot_i (D_i(x, z))^{(\omega-1)\beta_i}$$

is in $\mathcal{C}$ when interpreted over G_k. Similarly, let $S^{(2)}$ be obtained by replacing x with S in S. By the same argument as before, we see that the polynomial

$$S^{(\omega-1)}(x, z) = x \cdot \odot_i (D_i(x, z))^{(\omega-1)\gamma_i}$$

is in $\mathcal{C}$ when interpreted over G_k.

Now, build T by replacing x with $R^{(\omega-1)}(z, y)$ and z with $S^{(\omega-1)}(y, x)$ in P. Observe that the two previous replacements have no effect on Q and D_i other than permuting x and z, since any commutator of weight k placed in an argument of a commutator polynomial of weight larger than 1 can simply be deleted. Using again that the D_i are central in G_k we have

$$T(x, y, z) = zy^{-1}xQ(z, y, x) \cdot \odot_i D_i^{(\omega-1)\beta_i}(z, y) \cdot \odot_i D_i^{(\omega-1)\gamma_i}(y, x) \cdot \odot_i C_i^{\alpha_i}(z, y, x).$$

While still in G_k, reorder the D_i and C_i in the following way: Let $\phi(i)$ be such that $C_i(x, x, z)$ is $D_{\phi(i)}(x, z)$, and similarly let $\varphi(i)$ be such that $C_i(x, z, z)$ is $D_{\varphi(i)}(x, z)$. By the definition of β_i, γ_i, ϕ, and φ, there are just the right number of each of the D_i in T so that T can be equivalently written in G_k as

$$T(x, y, z) = zy^{-1}xQ(z, y, x) \cdot \odot_i (C_i(z, y, x)(D_{\phi(i)}(y, x))^{\omega-1}(D_{\varphi(i)}(z, y))^{\omega-1})^{\alpha_i}$$

So this polynomial $T(x, y, z)$ belongs to $\mathcal{C}$ when interpreted over G_k, although it need not be in $\mathcal{C}$ when interpreted in G. Still, it is easy to see that T is Malt'sev in all G: When $x = y$, all the commutators on x and y vanish so we have

$$\begin{aligned} T(x, x, z) &= zx^{-1}xQ(z, x, x) \cdot \odot_i (C_i(z, x, x) \cdot (D_{\varphi(i)}(z, x))^{\omega-1})^{\alpha_i} \\ &= z \cdot \odot_i (D_{\varphi(i)}(z, x) \cdot (D_{\varphi(i)}(z, x))^{\omega-1})^{\alpha_i} = z, \end{aligned}$$

and when $y = z$, all the commutators on z and y vanish and

$$\begin{aligned} T(x, z, z) &= zz^{-1}xQ(z, z, x) \cdot \odot_i (C_i(z, z, x) \cdot (D_{\phi(i)}(z, x))^{\omega-1})^{\alpha_i} \\ &= x \cdot \odot_i (D_{\phi(i)}(z, x) \cdot (D_{\phi(i)}(z, x))^{\omega-1})^{\alpha_i} = x. \end{aligned}$$

This concludes the induction step and the proof of the theorem. □

Now from Theorems 7 and 14, plus the tractability of Malt'sev operations, we obtain the following corollary.

Corollary 15. *Let G be a nilpotent group and $\mathcal{C}$ a nontrivial, idempotent clone of polynomials over G. Then the following are equivalent:*

– $\mathcal{C}$ *is tractable.*
– $\mathcal{C}$ *is not a d-factor, for any d.*
– $\mathcal{C}$ *contains a Malt'sev operation.*

It is tempting to conjecture in light of Corollaries 11 and 15 that if S is a block-group then an idempotent clone of polynomials over S is tractable iff it is not a d-factor. It would also be interesting to identify the largest class of finite groups for which the presence of $xy^{-1}z$ is necessary and sufficient for tractability.

References

1. A. Bulatov. A dichotomy theorem for constraints on a three-element set. In *Proc. of 43rd Foundations of Comp. Sci. (FOCS'02)*, pages 649–658, 2002.
2. A. Bulatov. Malt'sev constrains are tractable. *Electronic Colloquium on Computational Complexity (ECCC)*, 2002.
3. A. Bulatov and V. Dalmau. A simple algorithm for Malt'sev constraints. Submitted, 2005.
4. A. Bulatov and P. Jeavons. An algebraic approach to multi-sorted constraints. In *Principles and Practice of Constraint Programming (CP'03)*, pages 183–193, 2003.
5. A. Bulatov, P. Jeavons, and M. Volkov. Finite semigroups imposing tractable constraints. In G. Gomez, P. Silva, and J-E.Pin, editors, *Semigroups, Algorithms, Automata and Languages*, pages 313–329. WSP, 2002.
6. A. Bulatov, A. Krokhin, and P. Jeavons. Constraint satisfaction problems and finite algebras. In *Proc. 27th Int. Col. on Automata, Languages and Programming (ICALP'00)*, pages 272–282, 2000.
7. V. Dalmau. A new tractable class of constraint satisfaction problems. In *6th Int. Symp on Artificial Intelligence and Mathematics*, 2000.
8. V. Dalmau and J. Pearson. Closure functions and width 1 problems. In *Principles and Practice of Constraint Programming—CP'99*, pages 159–173, 1999.
9. T. Feder. Constraint satisfaction on finite groups with near subgroups. *Electronic Colloquium on Computational Complexity (ECCC)*, 2005.
10. T. Feder and M. Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory. *SIAM J. on Computing*, 28(1):57–104, 1998.
11. P. Jeavons. On the algebraic structure of combinatorial problems. *Theoretical Computer Science*, 200(1-2):185–204, 1998.
12. P. Jeavons, D. Cohen, and M. Gyssens. Closure properties of constraints. *J. ACM*, 44(4):527–548, 1997.
13. O. Klíma, P. Tesson, and D. Thérien. Dichotomies in the complexity of solving systems of equations over finite semigroups. *Theory of Computing Systems*, 2005. To appear.
14. B. Larose and L. Zádori. Taylor terms, constraint satisfaction and the complexity of polynomial equations over finite algebras. Submitted for publication, 2004.
15. J.-É. Pin. *PG = BG*, a success story. In J. Fountain, editor, *NATO Advanced Study Institute Semigroups, Formal Languages and Groups*, pages 33–47. Kluwer Academic Publishers, 1995.
16. T. J. Schaefer. The complexity of satisfiability problems. In *Proc. 10^{th} ACM STOC*, pages 216–226, 1978.

ECCC ISSN 1433-8092
http://www.eccc.uni-trier.de/eccc
ftp://ftp.eccc.uni-trier.de/pub/eccc
ftpmail@ftp.eccc.uni-trier.de, subject 'help eccc'