

Computing Science
REPRESENTING AND SOLVING
FINITE-DOMAIN CONSTRAINT PROBLEMS
USING SYSTEMS OF POLYNOMIALS

Chris Jefferson, Peter Jeavons
Computing Laboratory, University of Oxford, Oxford
{Chris.Jefferson,Peter.Jeavons}@comlab.ox.ac.uk

Martin J. Green
Department of Computer Science,
Royal Holloway, University of London, Egham, UK
M.J.Green@cs.rhul.ac.uk

M.R.C. van Dongen
Department of Computer Science, University College Cork, Ireland
Dongen@cs.ucc.ie

CS-RR-07-07



Oxford University Computing Laboratory
Wolfson Building, Parks Road, Oxford OX1 3QD

Abstract

In this paper we investigate the use of a system of multivariate polynomials to represent the restrictions imposed by a collection of constraints. The advantage of using polynomials to represent constraints is that it allows many different forms of constraints to be treated in a uniform way. Systems of polynomials have been widely studied, and a number of general techniques have been developed, including algorithms that generate an equivalent system with certain desirable properties, called a Gröbner Basis. General algorithms to compute a Gröbner Basis have doubly exponential complexity, but we observe that for the systems we use to describe constraint problems over finite domains a Gröbner Basis can be computed more efficiently than this. We also describe a family of algorithms, related to the calculation of Gröbner Bases, that can be used on any system of polynomials to compute an equivalent system in polynomial time. We show that these modified algorithms can simulate the effect of the local-consistency algorithms used in constraint programming. Finally, we describe how these algorithms can be used in a similar way to local consistency techniques to solve certain broad classes of constraint problems in polynomial time.

1 Introduction

The constraint programming paradigm [17] involves modelling a real-world problem as a set of *variables* together with a set of *constraints*; the constraints restrict the allowed combinations of values that can be simultaneously assigned to the variables. As well as capturing many practical computational problems, this very general paradigm includes a number of special cases corresponding to problems with particular forms of constraints, such as INTEGER PROGRAMMING and SATISFIABILITY. The study of efficient ways to represent and solve such constraint problems is currently a very active area of research within artificial intelligence [17].

In a constraint programming system each variable has to be assigned a value from some specified (often finite) domain, and each constraint describes the set of allowed combinations of values for some subset of the variables. A constraint may sometimes be specified *explicitly*, by listing the allowed (or disallowed) combinations of values in a table. More often, a constraint is specified *implicitly* by using some combination of available predicates, equations, inequalities and logical connectives.

In this paper we explore the use of *systems of polynomials* to specify constraints. We will say that a system of polynomials *allows* a particular combination of values for a set of variables if the simultaneous assignment of those values to the variables makes all of the polynomials in the system evaluate to zero. Such an assignment is called a *solution* to the system of polynomials.

The use of systems of polynomials has been considered a number of times in the constraints literature [2, 13, 14, 16], but is typically used to represent constraints on continuous variables. Here we focus on the use of polynomials to represent finite-domain constraints. The advantage of representing such constraints using polynomials is that they can then be treated in a uniform way along with continuous constraints, allowing the development of very general constraint-solving techniques. Another advantage is that systems of polynomials can be processed by standard computer algebra packages such as

Mathematica and REDUCE, so our approach helps to unify constraint programming with other forms of mathematical programming. The approach of using polynomials to represent finite-domain constraints has previously been explored by Clegg, Edmonds and Impagliazzo, but only for the special case of SATISFIABILITY problems [6].

Polynomials provide a very flexible and generic representation for problems from a wide variety of areas, and they have been intensively studied. In particular, Buchberger [4] showed that for any system of polynomials it is possible to compute an equivalent system with a number of desirable properties, which he called a *Gröbner Basis*. Given a Gröbner Basis, it is possible to obtain the solutions to a system of polynomials very easily (or determine that it has no solutions). A Gröbner Basis provides a convenient representation for the whole set of solutions which can be used to answer a wide range of questions about them, such as the correlations between individual variables, or the total number of solutions.

In general, the complexity of computing a Gröbner Basis for a system of polynomials is doubly exponential in the size of the system, except in certain special cases. However, we show in Section 5 that for the systems we use to represent constraints over finite domains this complexity is only singly exponential, and so is comparable with other general search techniques. In Section 6 we discuss a modified form of the algorithm for computing a Gröbner Basis which runs in *polynomial* time, but is incomplete, in the sense that it computes a system which is not necessarily a Gröbner Basis. The system computed by our modified algorithm has the same solutions as the original system and represents a set of equivalent constraints which is *locally consistent*. We show in Section 7 that this modified form of the algorithm can be used to achieve the same benefits as the local-consistency algorithms which are widely used in constraint processing. Finally, in Section 8 we discuss the use of adaptive consistency techniques for systems of polynomials.

2 Polynomials, Varieties and Ideals

In this section we introduce the basic concepts related to polynomials which are used throughout the paper. For a more detailed introduction see, for example, [8].

Definition 2.1. *Let \mathbb{K} be any field and $\{x_1, \dots, x_n\}$ a set of variables. A **polynomial** $p(x_1, \dots, x_n)$ over \mathbb{K} with variables x_1, \dots, x_n is a finite sum of **terms** of the form $\kappa x_1^{\alpha_1} \dots x_n^{\alpha_n}$, where $\kappa \in \mathbb{K}$ and $\alpha_i \in \{0, 1, 2, \dots\}$, $i = 1, \dots, n$. A single product $x_1^{\alpha_1} \dots x_n^{\alpha_n}$ is called a **monomial**² over $\{x_1, \dots, x_n\}$.*

The set of all polynomials over \mathbb{K} with variables x_1, \dots, x_n , together with the standard addition and multiplication operations, forms a ring, which is denoted $\mathbb{K}[x_1, \dots, x_n]$. A **system of polynomials** is a set of polynomials which all belong to the same polynomial ring; this set may be finite or infinite.

²Unfortunately, the words “term” and “monomial” are used by some authors the opposite way round! Here we follow the terminology of [8].

Definition 2.2. Let P be a system of polynomials from $\mathbb{K}[x_1, \dots, x_n]$. The **variety** of P , denoted $V(P)$, is the set of all solutions to P , that is,

$$V(P) = \{ \langle \kappa_1, \dots, \kappa_n \rangle \in \mathbb{K}^n : \forall p \in P, p(\kappa_1, \dots, \kappa_n) = 0 \}.$$

Definition 2.3. An **ideal** I over the polynomial ring $\mathbb{K}[x_1, \dots, x_n]$ is a system of polynomials from $\mathbb{K}[x_1, \dots, x_n]$ which satisfies the following conditions:

1. For all $f_1, f_2 \in I$ we have that $f_1 + f_2 \in I$;
2. For all $u \in \mathbb{K}[x_1, \dots, x_n]$ and all $f \in I$ we have that $uf \in I$.

Definition 2.4. Let $\{p_1, \dots, p_m\}$ be a finite system of polynomials from $\mathbb{K}[x_1, \dots, x_n]$. The **ideal generated by** $\{p_1, \dots, p_m\}$, denoted $I(\{p_1, \dots, p_m\})$, is given by

$$I(\{p_1, \dots, p_m\}) = \left\{ \sum_{i=1}^m u_i p_i : u_i \in \mathbb{K}[x_1, \dots, x_n] \right\}.$$

Given a finite system of polynomials $P = \{p_1, \dots, p_m\}$, from $\mathbb{K}[x_1, \dots, x_n]$, it is easy to see that the ideal generated by P has exactly the same set of solutions as the original system, that is, $V(I(P)) = V(P)$.

Example 2.5. In this example, assume the polynomials are drawn from $\mathbb{C}[x, y, z]$, where \mathbb{C} is the field of complex numbers.

The ideal $I(\{x + y, y + z\})$ has variety $\{ \langle a, -a, a \rangle : a \in \mathbb{C} \}$.

The ideal $I(\{x(x - 1), y + x - 3, z - x\})$ has variety $\{ \langle 0, 3, 0 \rangle, \langle 1, 2, 1 \rangle \}$. □

There are two standard ways to combine ideals to obtain new ideals:

Definition 2.6. The **sum** of two ideals I and J , denoted by $I + J$, is defined by $I + J = \{f + g : f \in I, g \in J\}$. The **product** of two ideals I and J , denoted by $I \times J$, is defined by $I \times J = \{fg : f \in I, g \in J\}$.

These two standard ways of combining two ideals have a straightforward effect on the corresponding varieties, as the next result indicates.

Lemma 2.7 ([8]). For any pair of ideals I and J , we have:

1. $V(I + J) = V(I) \cap V(J)$;
2. $V(I \times J) = V(I) \cup V(J)$.

Furthermore, given two finite systems of polynomials, they can be combined in straightforward ways to obtain *generating sets* for the sum and product of their corresponding ideals, as the next result indicates.

Lemma 2.8 ([8]). For any pair of ideals $I = I(\{p_1, \dots, p_m\})$ and $J = I(\{q_1, \dots, q_l\})$, we have:

1. $I + J = I(\{p_1, \dots, p_m, q_1, \dots, q_l\})$;
2. $I \times J = I(\{p_i q_j : i \in \{1, \dots, m\}, j \in \{1, \dots, l\}\})$.

Combining Lemmas 2.7 and 2.8, if we have two finite systems of polynomials, each with an associated set of solutions, then it is easy to obtain a new system of polynomials whose solutions are precisely the union or intersection of these two sets of solutions. For example, in the case of the intersection, we simply concatenate the two systems of polynomials.

3 Representing Constraints by Systems of Polynomials

In this section we show how various forms of constraint can be mapped to a system of polynomials whose set of solutions is precisely the set of combinations of values allowed by the constraint. If we do this for each constraint in a problem, then by the observation above, we can simply concatenate the systems of polynomials we obtain to create a system which represents the entire problem.

We shall assume for simplicity that each variable in the constraints we are trying to represent has a specified **domain** of possible values, and that these values are represented by natural numbers. Hence each constraint allows some set of tuples of natural numbers.

We shall represent a constraint over some subset of the variables x_1, \dots, x_n by a system of polynomials from $\mathbb{C}[x_1, \dots, x_n]$. The solutions to this system will be precisely the assignments (of natural numbers) that satisfy the constraint.

First note that we can restrict the domain of possible values for each individual variable to a specified finite set by including in our system a polynomial of the following form for each variable.

Definition 3.1. *The **domain polynomial** of a variable x with domain D is the polynomial $\prod_{j \in D} (x - j)$.*

In this paper we shall assume that all the variables have *finite* domains and we will always include a domain polynomial for each variable in the system used to represent a constraint (even though they are sometimes redundant). This has the advantage³ that it makes our solution methods complete, and provides a simple bound on the complexity (see Sections 4 and 5). Note that the system of polynomials consisting of just the domain polynomials for the variables x_1, \dots, x_r represents the rather trivial constraint which allows *all possible* combinations of values (from the domains) for these variables.

Example 3.2. [Constant constraints] Probably the simplest non-trivial form of constraint on a collection of variables $\langle x_1, \dots, x_r \rangle$ is a **constant constraint** which allows only a single assignment, say $\langle \kappa_1, \dots, \kappa_r \rangle$, assigning κ_i to x_i for $i = 1, \dots, r$. This constraint is easily represented by the system of polynomials $\{x_1 - \kappa_1, x_2 - \kappa_2, \dots, x_r - \kappa_r\}$, which has just one solution corresponding to the single allowed assignment. \square

³It has the further technical advantage that it ensures that the ideals generated by our systems of polynomials are *radical* [3, Lemma 8.19].

Example 3.3. [Disjunctive constraints] Some forms of constraint, such as the propositional **clauses** used in the SATISFIABILITY problem, are specified as a disjunction of simpler constraints. Using Lemmas 2.7 and 2.8, it follows that any such constraint can be represented by the system of polynomials consisting of the set of all products of all the polynomials in the systems representing the disjuncts. For example, consider the polynomial $\prod_{i \in \{1, \dots, r\}} (x_i - \kappa_i)$. This polynomial is satisfied if, for some $i \in \{1, \dots, r\}$, we have that x_i is assigned κ_i . In other words, at least one of the x_i 's is assigned the associated κ_i . This type of polynomial can therefore be used to represent a clause in an instance of SATISFIABILITY, by setting each κ_i to either 0 or 1, depending on whether the corresponding literal in the clause is negative or positive. For example, the clause $\neg x \vee y \vee \neg z$ is represented by the polynomial $x(y - 1)z$, which is equal to $xyz - xz$. \square

Example 3.4. [Linear Equations] Any linear equation involving variables x_1, \dots, x_r can be written in the form $p(x_1, \dots, x_r) = 0$, where p is a special form of polynomial in which $\sum \alpha_i \leq 1$ for all terms $\kappa x_1^{\alpha_1} \dots x_r^{\alpha_r}$. This means that any constraint which is expressed by one or more linear equations is naturally represented by a system of polynomials. This includes many of the constraints encountered in INTEGER PROGRAMMING problems. \square

Example 3.5. [Inequality] Consider the constraint “ $x \leq y$ ” over the variables x and y which each have the domain $\{1, 2, \dots, d\}$. The set of solutions to this constraint can be represented by the system of polynomials $P = \{p_i(x, y) : i = 1, \dots, d\}$, where

$$p_i(x, y) = \prod_{1 \leq j \leq i} (x - j) \prod_{i+1 \leq j \leq d} (y - j).$$

\square

Example 3.6. [Table constraints] Any constraint which specifies an explicit list of allowed combinations can be seen as a disjunction of constant constraints, so in principle we can use the methods described in Examples 3.2 and 3.3 to build a suitable system of polynomials to represent such constraints. However, in practice this method has a major limitation, in that it may produce systems of huge size.

Fortunately, there is a more sophisticated algorithm which computes a system of polynomials with a specified finite set of solutions in polynomial time in the size of that set [5]. In other words, there is a known algorithm which can be used to find a system of polynomials to represent any constraint specified by an explicit list of allowed tuples in polynomial time. \square

The examples above illustrate how constraints specified in various ways can be mapped to systems of polynomials whose solutions are precisely the combinations of values allowed by the constraints.

4 Gröbner Bases

Once we have constructed a system of polynomials whose solutions are precisely the assignments that satisfy a given collection of constraints, there are a number of questions

we may wish to ask, such as: Does this system of polynomials have a solution? Does x_3 take the value 7 in all solutions? In this section we examine ways in which a system of polynomials can be manipulated to obtain a new system with the same set of solutions which makes answering such questions easier.

Note that any two systems of polynomials which generate the same ideal have the same sets of solutions, since the variety of a system of polynomials is equal to the variety of the ideal that it generates. It turns out that a system of polynomials is particularly convenient to work with if it is a *Gröbner Basis* for the ideal that it generates. To define this property, we first have to define an ordering for the monomials in our polynomials.

Definition 4.1. A **monomial ordering**, \prec , on a set of variables $\{x_1, \dots, x_n\}$ is a total ordering on monomials over $\{x_1, \dots, x_n\}$ which satisfies the conditions:

1. For all monomials μ_1, μ_2 and μ_3 we have $\mu_1 \prec \mu_2 \Rightarrow \mu_1 \times \mu_3 \prec \mu_2 \times \mu_3$;
2. Every non-empty set of monomials has a smallest element with respect to \prec .

Example 4.2. One important monomial ordering is the **lexicographic** ordering, \prec_{lex} , which is defined by

$$x_1^{\alpha_1} \cdots x_n^{\alpha_n} \prec_{\text{lex}} x_1^{\beta_1} \cdots x_n^{\beta_n} \iff \langle \alpha_1, \dots, \alpha_n \rangle < \langle \beta_1, \dots, \beta_n \rangle,$$

where the ordering on tuples of natural numbers is the usual lexicographic ordering. \square

Example 4.3. Another important monomial ordering is the **total degree lexicographic** ordering, $\prec_{\text{d-lex}}$, which is defined by

$$x_1^{\alpha_1} \cdots x_n^{\alpha_n} \prec_{\text{d-lex}} x_1^{\beta_1} \cdots x_n^{\beta_n} \iff \left(\sum \alpha_i < \sum \beta_i \right) \vee \left(\sum \alpha_i = \sum \beta_i \right) \wedge \langle \alpha_1, \dots, \alpha_n \rangle < \langle \beta_1, \dots, \beta_n \rangle$$

where the ordering on tuples of natural numbers is the usual lexicographic ordering. \square

Once we have ordered the monomials, we can identify a distinguished term in every polynomial, in the following way.

Definition 4.4. Given a monomial ordering \prec on the set of variables $\{x_1, \dots, x_n\}$ and a non-zero polynomial $u = \sum_i \kappa_i \mu_i$ from $\mathbb{K}[x_1, \dots, x_n]$, where each $\kappa_i \in \mathbb{K}$ and each μ_i is a monomial over $\{x_1, \dots, x_n\}$, the **leading monomial** (LM) of u is the maximum of the μ_i under the monomial ordering \prec . If the leading monomial is μ_j (for some j), then the **leading term** (LT) of u is $\kappa_j \mu_j$.

Example 4.5. Consider the polynomial $4x^3 + 2y^2 + xy$ in the polynomial ring $\mathbb{C}[x, y]$. Under the lexicographic ordering, with $x \prec y$, the leading monomial is y^2 whilst the leading term is $2y^2$. Under total degree lexicographic ordering, with $x \prec y$, the leading monomial is x^3 whilst the leading term is $4x^3$. \square

Using the notion of leading term we can now define a form of division algorithm for multi-variate polynomials.

Definition 4.6. For any system of polynomials $P = \{p_1, \dots, p_m\}$ from $\mathbb{K}[x_1, \dots, x_n]$, and any polynomial $u \in \mathbb{K}[x_1, \dots, x_n]$, a **remainder** of u on division by P , denoted $u|_P$, is obtained by repeatedly performing the following “reduction” rule until it cannot be further applied:

Choose any $i \in \{1, \dots, m\}$ such that $\text{LT}(p_i)$ divides some term τ of u and replace u with $u - \frac{\tau}{\text{LT}(p_i)}p_i$.

Unfortunately, in the general case the remainder of a given polynomial on division by a given system of polynomials is not uniquely defined. This is because the result of performing the division algorithm in Definition 4.6 is not unique: it can produce different results depending on the order in which the polynomials in P are considered. Furthermore, in general, even when a polynomial f is in the ideal $I(P)$, it is not always the case that $f|_P = 0$. For example, x cannot be reduced by either polynomial in the system $\{x^2 + x, x^2\}$, although it can clearly be written as the first polynomial minus the second, and hence belongs to the ideal generated by this system.

Definition 4.7 ([4, 3]). A system of polynomials G from $\mathbb{K}[x_1, \dots, x_n]$ is called a **Gröbner Basis** (with respect to a given monomial ordering) if it satisfies the property that $u|_G$ is uniquely defined for any $u \in \mathbb{K}[x_1, \dots, x_n]$.

One very useful property of a Gröbner Basis is that it can be used to determine whether a given polynomial belongs to the ideal that the basis generates, as the next result shows.

Theorem 4.8 ([8]). If G is a Gröbner Basis, then $f|_G = 0$ if and only if $f \in I(G)$.

The next example demonstrates the use of a Gröbner Basis to provide a convenient representation for a table constraint.

Example 4.9. The following relation over the set $D = \{1, 2, 3, 4\}$ appeared in [12, Example 13] where it was used as an example of a relation having a large amount of symmetry:

$$R = \{ \langle 1, 2, 2, 3, 4 \rangle, \langle 1, 2, 2, 4, 3 \rangle, \langle 1, 2, 3, 1, 2 \rangle, \\ \langle 2, 4, 2, 3, 4 \rangle, \langle 2, 4, 2, 4, 3 \rangle, \langle 2, 4, 3, 1, 2 \rangle \}.$$

If we consider this relation as a constraint on the variables $\langle a, b, c, d, e \rangle$, then it can be represented by the following system of polynomials:

$$G = \{ e^3 - 9e^2 + 26e - 24, b^2 - 6b + 8, \\ d + 2e^2 - 13e + 17, 2c - e^2 + 7e - 16, 2a - b \}.$$

This system of polynomials is a Gröbner Basis (with respect to the lexicographic monomial order where $e \prec b \prec d \prec c \prec a$). The two uni-variate polynomials in the set G correspond to unary constraints: they indicate that the second component of any solution satisfying the constraint has to be in $\{2, 4\}$ and the last component has to be

in $\{2, 3, 4\}$. The remaining polynomials in the set G are bi-variate: they indicate relationships which were implicit in R and which are revealed by the Gröbner Basis. The fact that these polynomials are *linear* in the variables d, c and a indicates that for any $\langle b, e \rangle \in \{2, 4\} \times \{2, 3, 4\}$ there is *exactly one* $\langle a, b, c, d, e \rangle \in R$ and vice versa. In particular, if $\langle a, b, c, d, e \rangle \in R$ then $d = -2e^2 + 13e - 17$, $c = (e^2 - 7e + 16)/2$, and $a = b/2$. \square

One remarkable result established by Buchberger is that given any finite system of polynomials, and any fixed monomial ordering, it is possible to compute a new system of polynomials which generates the same ideal and is a Gröbner Basis [4, 8]. One way to do this is to use Algorithm 1, which is known as Buchberger's Algorithm [4, 8]⁴. The algorithm makes use of a construct called the *S-polynomial* of two given polynomials, which is defined as follows.

Definition 4.10. Given two monomials $x_1^{\alpha_1} \dots x_n^{\alpha_n}$ and $x_1^{\beta_1} \dots x_n^{\beta_n}$, their **lowest common multiple** (LCM) is $x_1^{\max(\alpha_1, \beta_1)} \dots x_n^{\max(\alpha_n, \beta_n)}$.

Given two polynomials u_1 and u_2 and a monomial ordering, the **S-polynomial** of u_1 and u_2 , denoted $\text{S-Pol}(u_1, u_2)$, is

$$\frac{\text{LCM}(\text{LM}(u_1), \text{LM}(u_2))}{\text{LT}(u_1)} u_1 - \frac{\text{LCM}(\text{LM}(u_1), \text{LM}(u_2))}{\text{LT}(u_2)} u_2 .$$

Example 4.11. Let $p = x^3 + z^3$ and $q = y^3 + z^3$, and take a total degree lexicographic ordering with $z \prec y \prec x$. Then we have $\text{S-Pol}(p, q) = y^3 z^3 - x^3 z^3$.

Note that although the leading terms of the two polynomials are eliminated in the calculation of their S-polynomial, the result may have a leading monomial which is higher in the monomial ordering than either of these original leading terms, as this example shows. \square

Example 4.12. Consider a set Q of domain polynomials corresponding to distinct variables. Note that each domain polynomial in Q is uni-variate, and the leading terms of any two domain polynomials for distinct variables do not share a variable. This implies that the S-polynomial of any two domain polynomials in Q is reduced by those two polynomials to zero (see Lemma 3.3.1 of [1]). Hence, applying Buchberger's Algorithm to Q will not change the set. It follows that any set of domain polynomials corresponding to any collection of distinct variables is a Gröbner Basis with respect to any monomial ordering. \square

Example 4.13. Recall from Example 3.5 that the constraint " $x \leq y$ " over the variables x and y which each have the domain $\{1, 2, \dots, d\}$ can be represented by the system of polynomials $P = \{p_i(x, y) : i = 1, \dots, d\}$, where

$$p_i(x, y) = \prod_{1 \leq j \leq i} (x - j) \prod_{i+1 \leq j \leq d} (y - j).$$

⁴More sophisticated and efficient refinements of this algorithm have been developed, and are generally used in practice, but this simple version is sufficient to illustrate the technique and allow a straightforward complexity analysis.

Algorithm 1 Buchberger's Algorithm

```
1: GIVEN:  $P$ : Set of initial polynomials
2: RETURNS:  $G$ : Gröbner Basis
3: BEGIN:  $G := P$ ,  $T := \{\{g_1, g_2\} : \{g_1, g_2\} \subseteq G\}$ 
4: while  $T \neq \emptyset$  do
5:   Select  $\{g_1, g_2\}$  from  $T$  and set  $T := T \setminus \{\{g_1, g_2\}\}$ 
6:    $h := \text{S-Pol}(g_1, g_2)$ 
7:    $h_0 := h|_G$ 
8:   if  $h_0 \neq 0$  then
9:      $T := T \cup \{\{g, h_0\} : g \in G\}$ 
10:     $G := G \cup \{h_0\}$ 
11:   end if
12: end while
```

Whatever monomial ordering we choose, P is a Gröbner Basis. To see this, it is sufficient to show that $\text{S-Pol}(p_a, p_b)|_P = 0$ for all $a, b \in \{1, \dots, d\}$. Without loss of generality assume $a < b$. Then the greatest common divisor of p_a and p_b is:

$$\prod_{1 \leq j \leq a} (x - j) \prod_{b+1 \leq j \leq d} (y - j)$$

After dividing both polynomials by this, the first will become a polynomial only in y , and the second a polynomial only in x . The result then follows by Lemma 3.3.1 of [1]. \square

Once we have computed a Gröbner Basis, it is then much easier to obtain the solutions to the original set of polynomials (or determine that there are no solutions). In particular, we can decide whether the system has a solution using the following well-known result (proved in [3, Theorem 7.40]).

Theorem 4.14 (Hilbert's Nullstellensatz). *Let \mathbb{K} be an algebraically closed field and I be an ideal in $\mathbb{K}[x_1, \dots, x_n]$. The variety $V(I)$ is empty if and only if I contains the polynomial 1.*

Hence to decide whether a system of polynomials has any solutions we can simply calculate a Gröbner Basis (with respect to any monomial ordering) and then check whether it contains a polynomial that is a non-zero constant. If so, the system has no solutions. Otherwise, at least one solution exists.

Example 4.15. A collection of benchmark instances of the SATIFIABILITY problem are provided at www.satlib.org. These can easily be converted to systems of polynomials and then passed to a computer algebra package, such as Mathematica, to obtain a Groebner Basis.

Some of these benchmark instances contain random 3-clauses. For example, the in-

stance “uf20-01” contains 91 random 3-clauses over 20 variables, as follows:

$$\begin{aligned} &\{x_4 \vee \neg x_{18} \vee x_{19}, x_3 \vee \neg x_5 \vee x_{18}, \neg x_5 \vee \neg x_8 \vee \neg x_{15}, \\ &\quad \vee x_7 \vee \neg x_{16} \vee \neg x_{20}, \neg x_7 \vee x_{10} \vee \neg x_{13}, \neg x_9 \vee \neg x_{12} \vee x_{17}, \dots, \\ &\quad x_9 \vee x_{17} \vee \neg x_{19}, \neg x_2 \vee x_{12} \vee x_{17}, x_4 \vee \neg x_5 \vee \neg x_{16}\}. \end{aligned}$$

A straightforward mapping to polynomials as described in Example 3.3 gives the following system of 91 polynomials:

$$\begin{aligned} &\{(1 - x_4)x_{18}(1 - x_{19}), (1 - x_3)x_5(1 - x_{18}), x_5x_8x_{15}, \\ &\quad (1 - x_7)x_{16}x_{20}, x_7(1 - x_{10})x_{13}, x_9x_{12}(1 - x_{17}), \dots, \\ &\quad (1 - x_9)(1 - x_{17})x_{19}, x_2(1 - x_{12})(1 - x_{17}), (1 - x_4)x_5x_{16}\}. \end{aligned}$$

If we combine this system with domain polynomials $x_i(x_i - 1)$, restricting each variable x_i to the Boolean values 0 or 1, then the in-built `GroebnerBasis` command of Mathematica calculates the following Gröbner Basis for the combined system in around 2 seconds⁵:

$$\begin{aligned} &\{x_{20} - 1, x_{19}^2 - x_{19}, x_{18} - x_{19}, x_{17} - 1, x_{16}, x_{15} - 1, x_{14} - 1, x_{13}x_{19}, \\ &\quad x_{13}^2 - x_{13}, x_{12}, x_{11} - x_{19}, x_{10}x_{19} - x_{19}, x_{13}x_{10} - x_{10} + x_{19}, x_{10}^2 - x_{10}, x_9x_{19} - x_{19}, \\ &\quad x_{13}x_9 - x_9 - x_{13} + 1, x_9x_{10} - x_{19}, x_9^2 - x_9, x_8x_{19} - x_{19}, x_{13}x_8 - x_8 + x_{19}, x_8x_{10} - x_8, \\ &\quad x_8x_9 - x_{19}, x_8^2 - x_8, x_7, x_6x_{19}, x_{13}x_6 - x_6 - x_{13} - x_{19} + 1, x_{10}x_6 - x_6 - x_{10} + 1, x_6x_9 - x_9 + x_{19}, \\ &\quad x_6x_8, x_6^2 - x_6, x_5, x_4x_{19} - x_{19}, x_{13}x_4 - x_4 + x_{19}, x_4x_{10} - x_{10}, x_4x_9 - x_{19}, x_4x_8 - x_8, \\ &\quad x_6x_4 - x_4 - x_6 + 1, x_4^2 - x_4, x_3 - x_{19}, x_2 - x_{19}, x_1 + x_{19} - 1\}. \end{aligned}$$

By Theorem 4.14, this indicates that this instance is soluble. Furthermore, many of the polynomials in this basis are linear, indicating that the corresponding variables have fixed values in all solutions. For example, variable x_{20} must be assigned the value 1 in all solutions. \square

Once we have determined that a system has a solution, we can build such a solution by constructing progressively larger satisfying partial assignments which we can verify will extend to a complete solution.

One way to do this is to generate a Gröbner Basis G for the system just once. We can then check whether any particular partial assignment s extends to a solution by checking whether the polynomial p_s which forbids precisely that assignment is in the ideal generated by G . To do this we simply have to compute $p_s|_G$. If $p_s|_G = 0$, then p_s is in the ideal, and hence the assignment s cannot be part of a solution, otherwise p_s is not in the ideal, and hence s can be extended to a solution.

An alternative method is to add polynomials representing the partial assignment (recall Example 3.2) to the system and regenerate the Gröbner Basis. We then use Hilbert’s Nullstellensatz to determine the satisfiability of the new system. Since we can construct solutions in this way without backtracking, all solutions to a problem with n variables, each with d possible values, can be found with at most $n \times d$ iterations.

⁵All timings reported in this paper were obtained using Mathematica 5.2 running under Windows XP on a Pentium 4 processor running at 3.2GHz.

5 Complexity

For a general system of polynomials P , the worst-case complexity of Buchberger's Algorithm is doubly exponential in the size of P [3, Page 513], and hence it can only be used with confidence for very small systems.

However, the systems of polynomials we are considering have the special property that they contain a domain polynomial for each variable, limiting the possible values for that variable to a fixed finite set. We will now show that in this special case the complexity of Buchberger's Algorithm is only singly exponential, and hence essentially no worse than other general solution methods for finite domain constraint satisfaction problems, such as backtrack search.

Lemma 5.1. *Let P be a system of polynomials over x_1, \dots, x_n which contains a domain polynomial for each variable of degree at most d . The total number of new polynomials added to P by Buchberger's Algorithm is at most d^n , each of which contains at most d^n monomials.*

Proof. Note that in Algorithm 1 each polynomial h is reduced by P before being added to the set G . If any term of h contains any variable x whose exponent is larger than d , then it will be reduced by the domain polynomial to give a lower power of x . Moreover, each new polynomial added to G must have a different leading monomial from all current members of G , or else it will be further reduced before being added. The number of distinct monomials over x_1, \dots, x_n , where each power is at most $d - 1$ is d^n . Further, each of these monomials has at most d^n monomials with this property which precede it in any monomial ordering, so each polynomial added contains at most d^n monomials. \square

Lemma 5.1 shows that the number of times that we can generate a new non-zero polynomial (line 8 of Algorithm 1) is bounded by d^n . Each of these adds at most $|P| + d^n$ elements to T (line 9). Therefore, the central loop (lines 4-12) executes at most $|P|^2 + (|P| + d^n)d^n < (|P| + d^n)^2$ times. Calculating the S-polynomials and reduced polynomials (lines 6 and 7) are both polynomial in the *size* of their input (also polynomial-size in $|P| + d^n$). Therefore, Algorithm 1 runs in polynomial time in $|P| + d^n$.

On the other hand, our next example demonstrates that a Gröbner Basis representing a constraint problem can be very large, even for some very simple problems.

Example 5.2. Consider the constraint problem with variables x_1, \dots, x_n (each with domain $\{1, \dots, d\}$) and a line of inequality constraints $x_1 \leq x_2 \leq \dots \leq x_n$.

We can construct a system of polynomials, P , whose solutions are the solutions to this constraint problem simply by combining the Gröbner Bases for the inequality constraint on each pair of adjacent variables, as given in Example 3.5. However, the resulting system of dn polynomials is not a Gröbner Basis for the ideal generated by P , as we shall now demonstrate.

For any set $J = \{i_1, i_2, \dots, i_d\}$ such that $1 \leq i_1 \leq i_2 \leq \dots \leq i_d \leq n$, we define the polynomial p_J as follows: $p_J(x_{i_1}, \dots, x_{i_d}) = (x_{i_1} - 1)(x_{i_2} - 2) \cdots (x_{i_d} - d)$. Note that any solution to the constraint problem we are considering is a solution to $p_J = 0$, so $p_J \in I(P)$. But this means that this ideal contains polynomials which have as leading

monomials all possible monomials of total degree d . It is known [8] that the leading monomials of the polynomials in any Gröbner Basis generate all the leading monomials of all polynomials in the ideal. Hence P is not a Gröbner Basis. However, the much larger set of polynomials, containing all $\binom{n+d-1}{n-1}$ polynomials of the form p_J , can be shown to be a Gröbner Basis with respect to any monomial ordering.

For example, when $n = 10$ and $d = 5$ this set of polynomials contains $\binom{14}{9} = 2002$ polynomials. It takes more than 7 minutes to calculate this Gröbner Basis using Mathematica. \square

Example 5.2 shows that the size of the Gröbner Basis grows rapidly with n and d in this case, whatever monomial ordering is used, and can be expensive to compute, even though the underlying constraint problem is very straightforward. This is because the Gröbner Basis encodes information about *all* of the relationships that can be deduced between all of the variables in the problem. In the next section we consider modifying Buchberger's Algorithm to obtain a system of polynomials which encodes only *local* relationships among the variables.

6 A Weaker Form of Basis

In order to obtain a polynomial-time variant of Buchberger's Algorithm, we will now describe a modified form of the basic algorithm which will generally be incomplete, in the sense that it generates a system of polynomials which may not always be a Gröbner Basis, but can still provide useful information, as we will establish below. This approach has been previously studied by Clegg et al. [6] for the special case of constraint problems over Boolean domains (that is, for instances of the SATISFIABILITY problem). In this section we will extend their approach to a more general setting by defining a new form of basis for a polynomial system.

We first define a link between Gröbner Bases and a certain *proof system*.

Definition 6.1. *Given a system of polynomials P from $\mathbb{K}[x_1, \dots, x_n]$ we define a **derivation-proof** from P of a polynomial f to be a sequence of polynomials $\langle f_1, \dots, f_b \rangle$ such that $f_b = f$, and for $i = 1, \dots, b$, we have:*

$$\begin{aligned} \exists p \in P \cup \{f_1, \dots, f_{i-1}\}, \exists u \in \mathbb{K}[x_1, \dots, x_n], f_i = pu \quad \vee \\ \exists f_j, f_k \in \{f_1, \dots, f_{i-1}\}, f_i = f_j + f_k \end{aligned}$$

If there exists a derivation-proof of f from P , then we denote this by $P \vdash f$.

Notice that a polynomial f has a derivation-proof from some system of polynomials P if and only if f is in the ideal generated by P . Hence, by Theorem 4.8, given a Gröbner Basis G , and a polynomial f , $f|_G = 0$ if and only if $G \vdash f$.

In order to obtain more efficient algorithms for processing polynomials we will now restrict the notion of a derivation-proof, in order to define a weaker form of proof system, and a correspondingly weaker notion of basis. We do this by restricting the polynomials that can be used in a derivation.

Definition 6.2. A derivation-proof $\langle f_1, \dots, f_b \rangle$ is called a ϕ -**proof** if the polynomial derived at every step satisfies a specified property ϕ (that is, $\phi(f_i)$ holds, for $i = 1, \dots, b$). If there exists a ϕ -proof of f from P then we write $P \vdash_\phi f$.

We will say that a system of polynomials G is a ϕ -**basis** if, for any polynomial f satisfying ϕ , $f|_G = 0$ if and only if $G \vdash_\phi f$.

Our next theorem will establish a sufficient condition for a system of polynomials to be a ϕ -basis, for a broad class of properties ϕ . The proof is similar in style to the usual proof that ‘‘Buchberger’s S-pair criterion’’ is a sufficient condition to establish a Gröbner Basis (see, for example, [8, Page 83, Theorem 6]). First we need some technical definitions, and a standard lemma.

Definition 6.3. Given a system of polynomials P , and a property ϕ , we say that a polynomial f has a ϕ -**representation** over P if f can be written as $\sum_i p_i h_i$ where each $p_i \in P$, and each $p_i h_i$ satisfies ϕ .

Given a system of polynomials P , and a monomial ordering \prec , we say that a polynomial f is **semi-reducible** over P with respect to \prec if f can be written as $\sum_i p_i h_i$ where each $p_i \in P$, and each $p_i h_i$ satisfies $\text{LM}(p_i h_i) \preceq \text{LM}(f)$.

Lemma 6.4 ([8, Page 81, Lemma 5]). Suppose we have a linear sum of polynomials $\sum_{i=1}^s c_i f_i$, where $c_i \in \mathbb{K}$ and $\text{LM}(f_i) = \mu$ for all i .

If $\text{LM}(\sum_{i=1}^s c_i f_i) \prec \mu$, then $\sum_{i=1}^s c_i f_i$ is a linear sum, with coefficients in \mathbb{K} , of the S-polynomials $\text{S-Pol}(f_j, f_k)$ for $1 \leq j, k \leq s$. Furthermore, $\text{LM}(\text{S-Pol}(f_j, f_k)) \prec \mu$ for each j, k .

Theorem 6.5. Let \prec be a monomial ordering, and let ϕ be a property of polynomials such that if $\phi(f)$ holds, and $\text{LM}(g) \preceq \text{LM}(f)$, then $\phi(g)$ also holds.

Let Q be a system of polynomials such that, for all $q_1, q_2 \in Q$, if $\text{S-Pol}(q_1, q_2)$ satisfies ϕ , then $\text{S-Pol}(q_1, q_2)$ is semi-reducible over Q with respect to \prec .

Then for any polynomial p satisfying ϕ the following are equivalent:

1. $Q \vdash_\phi p$;
2. p is ϕ -representable over Q ;
3. p is semi-reducible over Q with respect to \prec ;
4. $p|_Q = 0$

Proof. We first show that (2) \Leftrightarrow (3).

To establish that (2) \Rightarrow (3), assume for contradiction that there exists some polynomial p which is ϕ -representable over Q but not semi-reducible over Q with respect to \prec . Since p is ϕ -representable, we know that p can be expressed as $\sum h_i q_i$, where each $q_i \in Q$ and each $h_i q_i$ satisfies ϕ . Since p is not semi-reducible, we know that in any such representation $\text{LM}(p) \prec \max_i \{\text{LM}(h_i q_i)\}$. Choose a representation for which $\max_i \{\text{LM}(h_i q_i)\}$ is minimal (according to \prec) and call this monomial μ_p .

If we set $I = \{i : \text{LM}(h_i q_i) = \mu_p\}$, then we obtain:

$$p = \sum_{i \in I} \text{LT}(h_i) q_i + \sum_{i \in I} (h_i - \text{LT}(h_i)) q_i + \sum_{i \notin I} h_i q_i$$

where all of the polynomials in the first summation have leading monomial μ_p , and all polynomials in the other summations have a leading monomial which is strictly smaller (according to \prec).

By Lemma 6.4, we can express the first summation as a linear sum of S-polynomials of pairs of $\text{LT}(h_i) q_i$ terms. We can then rewrite these S-polynomials, as follows, where we write L_{jk} for $\text{LCM}(\text{LT}(q_j), \text{LT}(q_k))$:

$$\begin{aligned} \text{S-Pol}(\text{LT}(h_j) q_j, \text{LT}(h_k) q_k) &= \frac{\mu_p}{\text{LT}(h_j) \text{LT}(q_j)} \text{LT}(h_j) q_j - \frac{\mu_p}{\text{LT}(h_k) \text{LT}(q_k)} \text{LT}(h_k) q_k \\ &= \frac{\mu_p}{\text{LT}(q_j)} q_j - \frac{\mu_p}{\text{LT}(q_k)} q_k \\ &= \frac{\mu_p}{L_{jk}} \left(\frac{L_{jk}}{\text{LT}(q_j)} q_j - \frac{L_{jk}}{\text{LT}(q_k)} q_k \right) \\ &= \frac{\mu_p}{L_{jk}} \text{S-Pol}(q_j, q_k). \end{aligned}$$

By Lemma 6.4, each $\text{S-Pol}(\text{LT}(h_j) q_j, \text{LT}(h_k) q_k)$ has leading monomial strictly smaller than μ_p , and hence so does each $\text{S-Pol}(q_j, q_k)$. By our choice of property ϕ , this means that each $\text{S-Pol}(q_j, q_k)$ satisfies ϕ . By our assumption about Q , this means that each $\text{S-Pol}(q_j, q_k)$ is semi-reducible over Q with respect to \prec . Hence we can express each $\text{S-Pol}(\text{LT}(h_j) q_j, \text{LT}(h_k) q_k)$ as a sum of polynomials from Q with polynomial coefficients where each term in the summation has leading monomial strictly smaller than μ_p . This contradicts the minimality of μ_p and hence establishes the result that (2) \Rightarrow (3).

The converse is straightforward: by the choice of property ϕ , it follows immediately from Definition 6.3 that any polynomial which is semi-reducible and satisfies ϕ is also ϕ -representable. Hence (3) \Rightarrow (2).

Next we show that (1) \Leftrightarrow (2). Let p be any polynomial such that $Q \vdash_\phi p$. We will show that p is ϕ -representable over Q by induction on the length of the ϕ -proof of p from Q (see Definition 6.1). For the base case, if $p = qu$, for some $q \in Q$, and p satisfies ϕ , then p is clearly ϕ -representable over Q .

Now assume that p has a ϕ -proof of length i , and the result holds for all shorter proofs.

If $p = f_j u$ for some f_j with a shorter proof, then we know by the induction hypothesis that f_j is ϕ -representable over Q , and hence semi-reducible over Q with respect to \prec , by the argument above. This means that $f_j = \sum h_i q_i$ where $\text{LM}(h_i q_i) \preceq \text{LM}(f_j)$. Hence $p = \sum h_i q_i u$ where $\text{LM}(h_i q_i u) \preceq \text{LM}(f_j u) = \text{LM}(p)$. Hence in this case p is semi-reducible over Q with respect to \prec , and hence ϕ -representable, by the argument above.

If $p = f_j + f_k$ for some f_j, f_k with shorter proofs, then we know by the induction hypothesis that f_j and f_k are ϕ -representable over Q . Adding these representations gives a ϕ -representation for p over Q .

Hence, in all cases p is ϕ -representable, so we have shown that (1) \Rightarrow (2).

The converse is immediate from Definition 6.3 and Definition 6.1, since any polynomial which is ϕ -representable over a set Q clearly has a ϕ -proof from Q .

Finally, we show that (3) \Leftrightarrow (4). Assume for contradiction that there exists a polynomial p satisfying ϕ which is semi-reducible over Q with respect to \prec , but with $p|_Q \neq 0$. Chose such a polynomial whose leading monomial is as small as possible (according to \prec). Since p is semi-reducible, we know that p can be expressed as $\sum h_i q_i$, where each q_i in Q and each $h_i q_i$ satisfies $\text{LM}(h_i q_i) \preceq \text{LM}(p)$. In this summation there must exist at least one $h_i q_i$ such that $\text{LM}(h_i q_i) = \text{LM}(p)$, and for this value of i we have $\text{LT}(q_i) | \text{LT}(p)$. Now consider the polynomial $p' = p - \frac{\text{LT}(p)}{\text{LT}(q_i)} q_i$.

By the results above, $Q \vdash_\phi p$. Hence $Q \vdash_\phi p'$, so p' is also semi-reducible over Q with respect to \prec , by the results above. Since the leading monomial of p' is strictly smaller than the leading monomial of p , we know by the choice of p that $p'|_Q = 0$. However this implies that $p|_Q = 0$, which contradicts the choice of p . Hence we have shown that (3) \Rightarrow (4).

The converse follows immediately from Definition 4.6, since any polynomial which reduces to 0 over Q must be semi-reducible over Q . \square

We have shown that the conditions described in Theorem 6.5 are sufficient to ensure that the system of polynomials Q is a ϕ -basis. Now consider a modified form of the original Buchberger Algorithm, as shown in Algorithm 1, which only considers S-polynomials which satisfy the property ϕ . In other words, the algorithm only performs lines 7-11 when the S-polynomial h satisfies ϕ . Such a modified form of the algorithm will be called a ϕ -**truncated** Buchberger Algorithm. When ϕ is a property of the form specified in Theorem 6.5, the system of polynomials computed by a ϕ -truncated Buchberger Algorithm will satisfy the conditions of Theorem 6.5, and hence will be a ϕ -basis.

Even though a ϕ -basis computed in this way is not guaranteed to be a Gröbner Basis, it does generate the same ideal as the original system of polynomials, and can reveal significant information about this ideal, and hence the corresponding set of solutions. Moreover, we will show below that for some classes of problems computing a ϕ -basis in this way (for an appropriate choice of ϕ) is sufficient to decide whether any solutions exist.

Moreover, for some choices of the property ϕ , the time complexity of the ϕ -truncated Buchberger Algorithm is much lower than the time complexity of the complete Buchberger Algorithm.

The first specific property ϕ that we consider is based on the **total degree** of the monomials occurring in a polynomial. We will say that a polynomial satisfies the property [degree $\leq m$] if, for each of its monomials $x_1^{\alpha_1} \cdots x_n^{\alpha_n}$, the sum of the exponents $\sum_{i=1}^n \alpha_i$ is at most m . We first need to ensure that the monomial ordering we are using has the property that if μ_1 satisfies the [degree $\leq m$] property, and $\mu_2 \prec \mu_1$, then μ_2 also satisfies the [degree $\leq m$] property. Such an ordering is called a **graded** monomial ordering. Many such orderings exist, including the standard total degree lexicographic ordering [3, 8] (see Example 4.3).

Lemma 6.6. *Let P be a system of polynomials from $\mathbb{K}[x_1, \dots, x_n]$. The total number of new polynomials added to P by the $[\text{degree} \leq m]$ -truncated Buchberger Algorithm using any graded monomial ordering is at most $\binom{n+m}{m}$, each of which contains fewer than $\binom{n+m}{m}$, monomials.*

Proof. Similar to the proof of Lemma 5.1, except that the number of distinct monomials over x_1, \dots, x_n satisfying $[\text{degree} \leq m]$ is $\binom{n+m}{m}$, and each of these therefore has fewer than $\binom{n+m}{m}$, monomials which precede it in any graded monomial ordering. \square

By a similar argument to the argument given for the full algorithm above, this result implies that the truncated algorithm is polynomial-time in $|P|$ and $\binom{n+m}{m}$. The value of $\binom{n+m}{m}$ is $O(n^m)$, and therefore polynomial in n for fixed m .

7 Local Consistency

In this section we show that the local consistency algorithms used in constraint programming can be expressed by ϕ -proofs on the corresponding systems of polynomials, for a suitable choice of property ϕ . By relating this property to the property $[\text{degree} \leq m]$ considered in the previous section, we show that a local consistency algorithm can be simulated by a truncated Buchberger Algorithm, and hence any information obtained by enforcing local consistency can be obtained by performing a truncated Buchberger Algorithm.

A constraint problem is said to be **strong k -consistent** if any consistent assignment to $k - 1$ or fewer variables can be extended to a consistent assignment to any additional variable. Strong k -consistency can be enforced in polynomial-time (for any fixed k) [7]. One method of achieving this is given in Algorithm 7.1.

Algorithm 7.1 (Enforcing strong $(k+1)$ -consistency).

1. *For each set of k or fewer variables, define a new constraint, equal to the relational join of the projections of all constraints onto those variables.*
2. *Replace each constraint by the relational join of that constraint with the projections of all other constraints onto its variables.*
3. *Repeat (2) until no changes are made.*

It is known that, for several broad classes of constraint problems, the existence of a solution can be decided by enforcing strong k -consistency. These include problems with bounded tree-width [11] and problems where the constraints are binary and *max-closed* [15].

We now observe that if our original constraints each involve at most k variables, then the operations described in Algorithm 7.1 only ever involve constraints on k or fewer variables. These operations can be simulated using operations on polynomials representing the constraints in such a way that the polynomials computed never involve more than k variables. (For example, projection of a constraint onto some subset of

variables can be simulated using the Extension Theorem [8, Page 115] on some Gröbner Basis of the system which is generated using a lexicographic monomial ordering.)

We will say that a polynomial satisfies the property $[\#\text{vars} \leq k]$ if it involves at most k variables. Recall that if some system of polynomials over at most k variables has no solution then, by Hilbert's Nullstellensatz, any Gröbner Basis of this system must contain a constant polynomial. This implies that there is a $[\#\text{vars} \leq k]$ -proof of 1 from these polynomials.

Now consider any constraint problem represented by a system of polynomials P , where each polynomial in P satisfies $[\#\text{vars} \leq k]$. Our observations above imply that, if applying strong $(k + 1)$ -consistency to this constraint problem results in an empty constraint, then there must be a $[\#\text{vars} \leq k]$ -proof of 1 from P .

Unfortunately, the property $[\#\text{vars} \leq k]$ does not satisfy the conditions of Theorem 6.5 with any standard monomial ordering, and it does not seem to be possible to verify directly whether a polynomial has a $[\#\text{vars} \leq k]$ -proof of 1 using a $[\#\text{vars} \leq k]$ -truncated Buchberger Algorithm. However, the next two results show that the existence of a $[\#\text{vars} \leq k]$ -proof of 1 implies the existence of a $[\text{degree} \leq k \times (d - 1) + 1]$ -proof of 1, where d is the maximum domain size.

Lemma 7.2. *Let Q be a set of domain polynomials for x_1, \dots, x_n , with degree at most d . For any polynomial u from $\mathbb{K}[x_1, \dots, x_n]$, if u satisfies $[\#\text{vars} \leq k]$, then $u|_Q$ is uniquely defined, and satisfies both $[\#\text{vars} \leq k]$ and $[\text{degree} \leq k \times (d - 1)]$.*

Proof. By Example 4.12, any set of domain polynomials forms a Gröbner Basis, and so we can reduce any polynomial by this set and obtain a unique result.

If a monomial in u is reduced by the domain polynomial for variable x , it must have contained x and therefore the resulting monomials introduced cannot contain any more variables than the monomial being reduced. This means that reduction by the domain polynomials preserves the property $[\#\text{vars} \leq k]$.

After all possible reductions have been performed, no monomial can contain a power of any variable higher than $d - 1$, else it could be further reduced. As no monomial can contain more than k variables, this means all monomials must satisfy $[\text{degree} \leq k \times (d - 1)]$. \square

Theorem 7.3. *Let p_1, \dots, p_m, u be polynomials from $\mathbb{K}[x_1, \dots, x_n]$, and let Q be a set of domain polynomials for x_1, \dots, x_n , of degree at most d . If each p_i satisfies $[\#\text{vars} \leq k]$, and $\{p_1, \dots, p_m\} \vdash_{[\#\text{vars} \leq k]} u$, then*

$$\{p_1|_Q, \dots, p_m|_Q\} \cup Q \vdash_{[\text{degree} \leq k \times (d-1) + 1]} u|_Q .$$

Proof. We first note that $d \leq k \times (d - 1) + 1$, so each domain polynomial in Q satisfies $[\text{degree} \leq k \times (d - 1) + 1]$. Moreover, by Lemma 7.2, the polynomials $p_1|_Q, \dots, p_m|_Q$ and $u|_Q$ also all satisfy $[\text{degree} \leq k \times (d - 1) + 1]$.

Assume that the derivation-proof of u from $\{p_1, \dots, p_m\}$ is given by the list of polynomials $\langle f_1, \dots, f_b \rangle$. By induction on the length of this proof, b , we shall show that we can obtain a $[\text{degree} \leq k \times (d - 1) + 1]$ -derivation of $u|_Q$ from $\{p_1|_Q, \dots, p_m|_Q\} \cup Q$.

If f_b is derived by adding two earlier polynomials f_j and f_k , then it is straightforward to show that $f_b|_Q = f_j|_Q + f_k|_Q$, so by the induction hypothesis we are done.

If f_b is derived by multiplying some polynomial p from the set $\{p_1, \dots, p_m\} \cup \{f_1, \dots, f_{b-1}\}$ by some arbitrary polynomial h from $\mathbb{K}[x_1, \dots, x_n]$, then, without loss of generality, we may assume that h is a polynomial consisting of a single variable, say x . There are then two cases to consider:

- $(px)|_Q = (p|_Q)x$; this means that $f_b|_Q = (p|_Q)x$, so by the induction hypothesis, we are done.
- $(px)|_Q \neq (p|_Q)x$; this can only happen if $(p|_Q)x$ contains some monomials where the exponent of x is equal to the degree of q_i , where q_i is the domain polynomial for x . In this case we can subtract appropriate multiples of q_i to obtain $(px)|_Q$, and hence, by the induction hypothesis, obtain a (possibly longer) [degree $\leq k \times (d - 1) + 1$]-proof of $f_i|_Q$ from $\{p_1|_Q, \dots, p_m|_Q\} \cup Q$.

□

It follows from Theorem 7.3 that we can check for the existence of a [$\#\text{vars} \leq k$]-proof indirectly, using a [degree $\leq m$]-truncated Buchberger Algorithm, for $m = k \times (d - 1) + 1$. Hence, for a fixed domain size d , any constraint problem which can be decided by strong k -consistency can be decided in polynomial-time by a truncated Buchberger Algorithm.

8 Adaptive Consistency

Another local consistency method which is commonly used in constraint programming is to vary the level of consistency which is enforced during the solution process depending on the local structure of the constraints. This technique is known as *adaptive consistency* [10], and is often implemented using a general algorithmic framework called *bucket elimination* [9].

The bucket elimination algorithm for a constraint problem proceeds as follows. It first orders the variables of the problem, and then partitions the constraints into separate collections, known as *buckets*. Each bucket is associated with a particular variable. The bucket associated with variable x contains all the constraints involving the variable x which do not involve any variables occurring higher in the ordering. In other words, each constraint is allocated to the bucket which is associated with whichever of the variables of that constraint occurs highest in the ordering. Buckets are then processed in order, according to the chosen variable ordering on their associated variables, from highest to lowest. When the bucket associated with a variable x is processed, an “elimination procedure” is performed over all the constraints in that bucket, yielding a new constraint that does not involve the variable x . This new constraint specifies the “effect” of the variable x on the remainder of the problem. In other words, it allows just those combinations of values that are allowed by the constraints in the bucket along with *some* value for x . The new constraint is then placed in the appropriate (lower) bucket and the processing continues until all buckets are processed, or some constraint is generated which allows no solutions, in which case the problem is declared to be insoluble.

```

bucketeliminate[polys_, vars_] := Module[{buckets = {}, newpolys = polys},
  Scan[{buckets = Join[buckets, bucketassign[newpolys, vars]];
    newpolys = GroebnerBasis[Cases[buckets, {{#}, poly_] -> poly], vars, {#}];
    If[newpolys == {1},
      Print["Unsatisfiable at variable ", #]; Return[False];] &,
    vars]]

bucketassign[polys_, vars_] :=
  Map[{Select[vars, Function[var, MemberQ[Variables[#], var]], 1], #} &, polys]

```

Figure 1: Bucket elimination for a system of polynomials implemented in Mathematica

If constraints are represented by systems of polynomials, in the way that we are proposing here, then bucket elimination can be implemented very simply using Gröbner Basis calculations. For each variable x , the bucket associated with x is assigned a subset of the polynomials in the given system, consisting of all polynomials over x (and possibly other variables lower in the ordering). If we calculate a Gröbner Basis for these polynomials, with respect to an appropriate *elimination ordering* [1], then the new constraint we require is represented by the subset of polynomials in this Gröbner Basis which do not involve x . These new polynomials can then be allocated to lower buckets and processing continues. An implementation of this algorithm as a Mathematica function is shown in Figure 1, where the arguments are the set of polynomials, and the ordered list of variables (highest first). A subsidiary function is also shown in the figure, which associates each polynomial in the given set with the name of a variable indicating the bucket to which that polynomial is allocated.

The complexity of the bucket elimination algorithm is exponential in a structural parameter of the given constraint problem, called the *induced width* [9]. This value depends on the way in which the constraints overlap and the chosen variable ordering. For problems with (known) bounded tree-width it is possible to find (in polynomial time) a variable ordering such that the induced width is also bounded, so all such problems can be solved in polynomial time.

Example 8.1. Recall the simple constraint problem discussed in Example 5.2, with variables x_1, \dots, x_n (each with domain $\{1, \dots, d\}$) and a line of inequality constraints: $x_1 \leq x_2 \leq \dots \leq x_n$. It was shown in Example 5.2 that any Gröbner Basis for this problem contains $\binom{n+d}{d}$ polynomials, and the time required to compute such a Basis rises rapidly with n and d .

However, if the variables are ordered along the line, the induced width of this problem is 1, and the bucket elimination algorithm never has to deal with polynomials over more than 2 variables. This means that the time required to decide whether a solution exists increases only linearly with the number of variables in the problem.

For example, when $n = 10$ and $d = 5$ the Mathematica implementation shown in Figure 1 verifies that the corresponding constraint problem is soluble in 0.015 seconds. \square

9 Conclusion

We have shown how polynomials provide a powerful general language for expressing many different forms of constraint problems.

Once a problem has been expressed with a corresponding system of polynomials, we have described a standard technique to find a new system of polynomials with the same set of solutions, called a Gröbner Basis. A Gröbner Basis provides a convenient representation of all the solutions to a system of polynomials and hence of all solutions to the original constraint problem. It can be computed using a standard computer algebra package, and can be used to answer many different questions about the solutions in a straightforward way.

We have also shown that we may truncate the standard Gröbner Basis algorithm by defining a property which all polynomials added to the system must satisfy. If this property is suitably specified, then the algorithm generates a useful system of polynomials with the same solutions in polynomial time, albeit not normally a Gröbner Basis. We have demonstrated that we can use this truncated algorithm to achieve a kind of local consistency which can simulate the k -consistency algorithms commonly used in constraint programming.

Finally we have shown that adaptive consistency techniques for constraint problems can be implemented very easily using Gröbner Basis techniques.

This paper presents our initial findings on using Gröbner Basis techniques for solving constraint problems. Obvious directions for future research are to determine for which classes of constraint problems these techniques can efficiently find solutions, and to refine the techniques so that they can be implemented more efficiently.

References

- [1] W. Adams and P. Loustau. *An Introduction to Gröbner Bases*. Graduate Studies in Mathematics. American Mathematical Society, 1994.
- [2] A. Aiba, K. Sakai, Y. Sato, D. Hawley, and R. Hasegawa. Constraint logic programming language CAL. In *Proceedings of the International Conference on Fifth Generation Computer Systems*, pages 263–76. Ohmsha Publishers, 1988.
- [3] T. Becker and V. Weispfenning. *Gröbner Bases A Computational Approach to Commutative Algebra*. Graduate Texts in Mathematics. Springer, 1993.
- [4] B. Buchberger. Gröbner bases: An algorithmic method in polynomial ideal theory. In R. Bose, editor, *Multidimensional Systems Theory, Mathematics and Its Applications*, chapter 6, pages 184–232. D. Reidel Publishing Company, 1985.
- [5] B. Buchberger and M. Moeller. The construction of multivariate polynomials with preassigned zeros. In *Computer Algebra, EUROCAM 82*, number 144 in LNCS, pages 24–31, 1982.

- [6] M. Clegg, J. Edmonds, and R. Impagliazzo. Using the Groebner basis algorithm to find proofs of unsatisfiability. In *STOC '96: Proceedings of the 28th ACM Symposium on Theory of Computing*, pages 174–183. ACM Press, 1996.
- [7] M. Cooper. An optimal k-consistency algorithm. *Artificial Intelligence*, 41:89–95, 1989.
- [8] D. Cox, J. Little, and J. O’Shea. *Ideals, Varieties, and Algorithms* An Introduction to Computational Algebraic Geometry and Commutative Algebra. Springer, 1996.
- [9] R. Dechter. Bucket elimination: A unifying framework for reasoning. *Artificial Intelligence*, 113(1-2):41–85, 1999.
- [10] R. Dechter. *Constraint Processing*. Morgan Kaufmann, 2003.
- [11] E. Freuder. A sufficient condition for backtrack-bounded search. *Journal of the ACM*, 32:755–761, 1985.
- [12] R. Gault and P. Jeavons. Implementing a test for tractability. *Journal of Constraints*, 9:139–160, 2004.
- [13] L. Granvilliers, E. Monfroy, and F. Benhamou. Symbolic-interval cooperation in constraint programming. In *Proceedings of ISSAC 2001*, pages 150–166, 2001.
- [14] H. Hong. RISC-CLP(Real): Constraint logic programming over real numbers. In B. F. and A. Colmerauer, editors, *Constraint Logic Programming: Selected Research*. MIT Press, 1993.
- [15] P. Jeavons and M. Cooper. Tractable constraints on ordered domains. *Artificial Intelligence*, 79(2):327–339, 1995.
- [16] E. Monfroy. The constraint solver collaboration language of BALI. In *In Proceedings of JCSLP 1998*, pages 349–350, 1998.
- [17] F. Rossi, P. van Beek, and T. Walsh, editors. *The Handbook of Constraint Programming*. Elsevier, 2006.