

Closure Properties of Constraints*

Peter Jeavons, David Cohen
Department of Computer Science
Royal Holloway, University of London, UK

Marc Gyssens
Department WNI
University of Limburg
B-3590 Diepenbeek, Belgium

May 14, 1997

Abstract

Many combinatorial search problems can be expressed as ‘constraint satisfaction problems’, and this class of problems is known to be NP-complete in general. In this paper we investigate the subclasses which arise from restricting the possible constraint types. We first show that any set of constraints which does not give rise to an NP-complete class of problems must satisfy a certain type of algebraic closure condition. We then investigate all the different possible forms of this algebraic closure property, and establish which of these are sufficient to ensure tractability. As examples, we show that all known classes of tractable constraints over finite domains can be characterised by such an algebraic closure property. Finally, we describe a simple computational procedure which can be used to determine the closure properties of a given set of constraints. This procedure involves solving a particular constraint satisfaction problem, which we call an ‘indicator problem’.

Keywords: Constraint satisfaction problem, complexity, NP-completeness, indicator problem

1 Introduction

Solving a constraint satisfaction problem is known to be NP-complete [20]. However, many of the problems which arise in practice have special properties which allow them to be solved efficiently. The question of identifying restrictions to the general problem which are sufficient to ensure tractability is important from both a practical and a theoretical viewpoint, and has been extensively studied.

Such restrictions may either involve the *structure* of the constraints, in other words, which variables may be constrained by which other variables, or they may involve the

*Earlier versions of parts of this paper were presented at the International Conferences on Constraint Programming in 1995 and 1996. (See references [14] and [15].)

nature of the constraints, in other words, which combinations of values are allowed for variables which are mutually constrained. Examples of the first approach can be found in [7, 9, 12, 22, 23] and examples of the second approach can be found in [4, 14, 16, 8, 17, 22, 28, 29].

In this paper we take the second approach, and investigate those classes of constraints which only give rise to tractable problems whatever way they are combined. A number of distinct classes of constraints with this property have previously been identified and shown to be maximal [4, 14, 16].

In this paper we establish that any class of constraints which does not give rise to NP-complete problems must satisfy a certain algebraic closure condition, and hence this algebraic property is a *necessary* condition for a class of constraints to be tractable (assuming that P is not equal to NP). We also show that many forms of this algebraic closure property are *sufficient* to ensure tractability.

As an example of the wide applicability of these results, we show that all known examples of tractable constraint classes over finite domains can be characterized by an algebraic condition of this kind, even though some of them were originally defined in very different ways.

Finally, we describe a simple computational procedure to determine the algebraic closure properties of a given set of constraints. The test involves calculating the solutions to a fixed constraint satisfaction problem involving constraints from the given set.

The work described in this paper represents a generalization of earlier results concerning tractable subproblems of the GENERALIZED SATISFIABILITY problem. Schaefer [26] identified all possible tractable classes of constraints for this problem, which corresponds to the special case of the constraint satisfaction problem in which the variables are Boolean. The tractable classes described in [26] are special cases of the tractable classes of general constraints described below, and they are given as examples.

A number of tractable constraint classes have also been identified by Feder and Vardi in [8]. They define a notion of ‘width’ for constraint satisfaction problems in terms of the logic programming language Datalog, and show that problems with bounded width are solvable in polynomial time. It is stated in [8] that the problem of determining whether a fixed collection of constraints gives rise to problems of bounded width is not known to be decidable. However, it is shown that a more restricted property, called ‘bounded strict width’, is decidable, and in fact corresponds to an algebraic closure property of the form described here. Other tractable constraint classes are also shown to be characterised by a closure property of this type. This paper builds on the work of Feder and Vardi by examining the general question of the link between algebraic closure properties and tractability of constraints, and establishing necessary and sufficient conditions for these closure properties.

A different approach to identifying tractable constraints is taken in [28], where it is shown that a property of constraints, referred to as ‘row-convexity’, together with path-consistency, is sufficient to ensure tractability in binary constraint satisfaction problems. It should be noted, however, that because of the additional requirement for path-consistency, row-convex constraints do *not* constitute a tractable class in the sense defined in this paper. In fact, the class of problems which contain only row-convex constraints is NP-complete [4].

The paper is organised as follows. In Section 2 we give the basic definitions, and in Section 3 we define what we mean by an algebraic closure property for a set of relations

and examine the possible forms of such a closure property. In Section 4 we identify which of these forms are *necessary* conditions for tractability, and in Section 5 we identify which of them are *sufficient* for tractability. In Section 6 we describe a computational method to determine the closure properties satisfied by a set of relations. Finally, we summarise the results presented and draw some conclusions.

2 Definitions

2.1 The constraint satisfaction problem

Notation 2.1 For any set D , and any natural number n , we denote the set of all n -tuples of elements of D by D^n . For any tuple $t \in D^n$, and any i in the range 1 to n , we denote the value in the i th coordinate position of t by $t[i]$. The tuple t may then be written in the form $\langle t[1], t[2], \dots, t[n] \rangle$.

A subset of D^n is called an n -ary relation over D .

We now define the (finite) constraint satisfaction problem which has been widely studied in the Artificial Intelligence community [18, 20, 22].

Definition 2.2 An instance of a constraint satisfaction problem consists of

- a finite set of variables, V ;
- a finite domain of values, D ;
- a set of constraints $\{C_1, C_2, \dots, C_q\}$.

Each constraint C_i is a pair (S_i, R_i) , where S_i is a list of variables of length m_i , called the constraint scope, and R_i is an m_i -ary relation over D , called the constraint relation. (The tuples of R_i indicate the allowed combinations of simultaneous values for the variables in S_i .)

The length of the tuples in the constraint relation of a given constraint will be called the *arity* of that constraint. In particular, unary constraints specify the allowed values for a single variable, and binary constraints specify the allowed combinations of values for a pair of variables. A *solution* to a constraint satisfaction problem is a function from the variables to the domain such that the image of each constraint scope is an element of the corresponding constraint relation.

Deciding whether or not a given problem instance has a solution is NP-complete in general [20] even when the constraints are restricted to binary constraints. In this paper we shall consider how restricting the allowed constraint relations to some fixed subset of all the possible relations affects the complexity of this decision problem. We therefore make the following definition.

Definition 2.3 For any set of relations, Γ , $\text{CSP}(\Gamma)$ is defined to be the decision problem with

INSTANCE: An instance, P , of a constraint satisfaction problem, in which all constraint relations are elements of Γ .

QUESTION: Does P have a solution?

If there exists an algorithm which solves every problem instance in $\text{CSP}(\Gamma)$ in polynomial time, then we shall say that $\text{CSP}(\Gamma)$ is a *tractable* problem, and Γ is a *tractable* set of relations.

Example 2.4 The binary disequality relation over a set D , denoted \neq_D , is defined as

$$\neq_D = \{\langle d_1, d_2 \rangle \in D^2 \mid d_1 \neq d_2\}.$$

Note that $\text{CSP}(\{\neq_D\})$ corresponds precisely to the GRAPH $|D|$ -COLORABILITY problem [11]. This problem is tractable when $|D| \leq 2$ and NP-complete when $|D| \geq 3$. \square

Example 2.5 Consider the ternary relation δ over the set $D = \{0, 1\}$ which is defined by

$$\delta = \{\langle 0, 0, 1 \rangle, \langle 0, 1, 0 \rangle, \langle 1, 0, 0 \rangle, \langle 0, 1, 1 \rangle, \langle 1, 0, 1 \rangle, \langle 1, 1, 0 \rangle\}.$$

The problem $\text{CSP}(\{\delta\})$ corresponds precisely to the NOT-ALL-EQUAL SATISFIABILITY problem [11], which is NP-complete [26]. \square

Example 2.6 We now describe three relations which will be used as examples of constraint relations throughout the paper.

Each of these relations is a set of tuples of elements from the domain $D = \{0, 1, 2\}$, as defined below:

$$\begin{aligned} R_1 = \{ & \langle 0, 0 \rangle, & R_2 = \{ & \langle 0, 1, 2 \rangle, \\ & \langle 1, 2 \rangle, & & \langle 1, 2, 0 \rangle, \\ & \langle 0, 1 \rangle, & & \langle 2, 0, 1 \rangle \} \\ & \langle 2, 1 \rangle \} \end{aligned}$$

$$\begin{aligned} R_3 = \{ & \langle 2, 2 \rangle, \\ & \langle 2, 1 \rangle, \\ & \langle 2, 0 \rangle, \\ & \langle 1, 2 \rangle, \\ & \langle 1, 1 \rangle, \\ & \langle 1, 0 \rangle, \\ & \langle 0, 2 \rangle, \\ & \langle 0, 1 \rangle \} \end{aligned}$$

The problem $\text{CSP}(\{R_1, R_2, R_3\})$ consists of all constraint satisfaction problem instances in which the constraint relations are all chosen from the set $\{R_1, R_2, R_3\}$.

The complexity of $\text{CSP}(\Gamma)$ for arbitrary subsets Γ of $\{R_1, R_2, R_3\}$ will be determined using the techniques developed later in this paper (see Example 6.6). \square

2.2 Operations on relations

In Section 4 we shall examine conditions on a set of relations Γ which allow known NP-complete problems to be reduced to $\text{CSP}(\Gamma)$. The reductions will be described using standard operations from relational algebra [1], which are described in this section.

Definition 2.7 We define the following operations on relations.

- Let R be an n -ary relation over a domain D and let S be an m -ary relation over D . The Cartesian product $R \times S$ is defined to be the $(n + m)$ -ary relation

$$R \times S = \{\langle t[1], t[2], \dots, t[n + m] \rangle \mid (\langle t[1], t[2], \dots, t[n] \rangle \in R) \wedge (\langle t[n + 1], t[n + 2], \dots, t[n + m] \rangle \in S)\}.$$

- Let R be an n -ary relation over a domain D . Let $1 \leq i, j \leq n$. The equality selection $\sigma_{i=j}(R)$ is defined to be the n -ary relation

$$\sigma_{i=j}(R) = \{t \in R \mid t[i] = t[j]\}.$$

- Let R be an n -ary relation over a domain D . Let i_1, \dots, i_m be a subsequence of $1, \dots, n$. The projection $\pi_{i_1, \dots, i_m}(R)$ is defined to be the m -ary relation

$$\pi_{i_1, \dots, i_m}(R) = \{\langle t[i_1], \dots, t[i_m] \rangle \mid t \in R\}.$$

It is well-known that the combined effect of two constraints in a constraint satisfaction problem can be obtained by performing a relational *join* operation [1] on the two constraints [12]. The next result is a simple consequence of the definition of the relational join operation.

Lemma 2.8 *Any relational join of relations R and S can be calculated by performing a sequence of Cartesian product, equality selection, and projection operations on R and S .*

In view of this result, it will be convenient to use the following notation.

Notation 2.9 *The set of all relations which can be obtained from a given set of relations, Γ , using some sequence of Cartesian product, equality selection, and projection operations will be denoted Γ^+ .*

Note that Γ^+ contains exactly those relations which can be obtained as ‘derived’ relations in a constraint satisfaction problem instance with constraint relations chosen from Γ [2].

3 Closure operations

We shall establish below that significant information about $\text{CSP}(\Gamma)$ can be determined from algebraic properties of the set of relations Γ . In order to describe these algebraic properties we need to consider arbitrary operations on D , in other words, arbitrary functions from D^k to D , for arbitrary values of k .

For the results below we shall be particularly interested in certain special kinds of operations. We therefore make the following definition:

Definition 3.1 *Let \otimes be a k -ary operation from D^k to D .*

- If \otimes is such that, for all $d \in D$, $\otimes(d, d, \dots, d) = d$ then \otimes is said to be idempotent.
- If there exists an index $i \in \{1, 2, \dots, k\}$ such that for all $\langle d_1, d_2, \dots, d_n \rangle \in D^k$ we have $\otimes(d_1, d_2, \dots, d_n) = f(d_i)$, where f is a non-constant unary operation on D , then \otimes is called essentially unary. (Note that f is required to be non-constant, so constant operations are not essentially unary.)

If f is the identity operation, then \otimes is called a projection.

- If $k \geq 3$ and there exists an index $i \in \{1, 2, \dots, k\}$ such that for all $d_1, d_2, \dots, d_k \in D$ with $|\{d_1, d_2, \dots, d_k\}| < k$ we have $\otimes(d_1, d_2, \dots, d_k) = d_i$, but \otimes is not a projection, then \otimes is called a semiprojection [25, 27].
- If $k = 3$ and for all $d, d' \in D$ we have $\otimes(d, d, d') = \otimes(d, d', d) = \otimes(d', d, d) = d$, then \otimes is called a majority operation.
- If $k = 3$ and for all $d_1, d_2, d_3 \in D$ we have $\otimes(d_1, d_2, d_3) = d_1 - d_2 + d_3$, where $+$ and $-$ are binary operations on D such that $\langle D, +, - \rangle$ is an Abelian group [27], then \otimes is called an affine operation.

Any operation on D can be extended to an operation on tuples over D by applying the operation in each coordinate position separately (i.e., pointwise). Hence, any operation defined on the domain of a relation can be used to define an operation on the tuples in that relation, as follows:

Definition 3.2 Let $\otimes : D^k \rightarrow D$ be a k -ary operation on D and let R be an n -ary relation over D .

For any collection of k tuples, $t_1, t_2, \dots, t_k \in R$, (not necessarily all distinct) the n -tuple $\otimes(t_1, t_2, \dots, t_k)$ is defined as follows:

$$\otimes(t_1, t_2, \dots, t_k) = \langle \otimes(t_1[1], t_2[1], \dots, t_k[1]), \otimes(t_1[2], t_2[2], \dots, t_k[2]), \dots, \otimes(t_1[n], t_2[n], \dots, t_k[n]) \rangle.$$

Finally, we define $\otimes(R)$ to be the n -ary relation

$$\{\otimes(t_1, \dots, t_k) \mid t_1, \dots, t_k \in R\}.$$

Using this definition, we now define the following closure property of relations.

Definition 3.3 Let \otimes be a k -ary operation on D , and let R be an n -ary relation over D . The relation R is closed under \otimes if $\otimes(R) \subseteq R$.

Example 3.4 Let Δ be the ternary majority operation defined as follows:

$$\Delta(x, y, z) = \begin{cases} y & \text{if } y = z; \\ x & \text{otherwise.} \end{cases}$$

The relation R_2 defined in Example 2.6 is closed under Δ , since applying the Δ operation to any 3 elements of R_2 yields an element of R_2 . For example,

$$\Delta(\langle 0, 1, 2 \rangle, \langle 1, 2, 0 \rangle, \langle 1, 2, 0 \rangle) = \langle 1, 2, 0 \rangle \in R_2.$$

The relation R_3 defined in Example 2.6 is *not* closed under Δ , since applying the Δ operation to the last 3 elements of R_3 yields a tuple which is not an element of R_3 :

$$\Delta(\langle 1, 0 \rangle, \langle 0, 2 \rangle, \langle 0, 1 \rangle) = \langle 0, 0 \rangle \notin R_3.$$

□

For any set of relations Γ , and any operation \otimes , if every $R \in \Gamma$ is closed under \otimes , then we shall say that Γ is closed under \otimes . The next lemma indicates that the property of being closed under some operation is preserved by all possible projection, equality selection, and product operations on relations, as defined in Section 2.2.

Lemma 3.5 *For any set of relations Γ , and any operation \otimes , if Γ is closed under \otimes , then Γ^+ is closed under \otimes .*

Proof: Follows immediately from the definitions. ■

Notation 3.6 *For any set of relations Γ with domain D , the set of all operations on D under which Γ is closed will be denoted Γ^\triangleright .*

The set of closure operations, Γ^\triangleright , can be used to obtain a great deal of information about the problem $\text{CSP}(\Gamma)$, as we shall demonstrate in the next two Sections.

As a first example of this, we shall show that the operations in Γ^\triangleright can be used obtain a reduction from one problem to another.

Proposition 3.7 *For any set of finite relations Γ , and any $\otimes \in \Gamma^\triangleright$, there is a polynomial-time reduction from $\text{CSP}(\Gamma)$ to $\text{CSP}(\otimes(\Gamma))$, where $\otimes(\Gamma) = \{\otimes(R) \mid R \in \Gamma\}$.*

Proof: Let \mathcal{P} be any problem instance in $\text{CSP}(\Gamma)$ and consider the instance \mathcal{P}' obtained by replacing each constraint relation R_i of \mathcal{P} by the relation $\otimes(R_i)$. It is clear that \mathcal{P}' can be obtained from \mathcal{P} in polynomial-time. It follows from Definition 3.3 that \mathcal{P}' has a solution if and only if \mathcal{P} has a solution. ■

It follows from this result that if Γ^\triangleright contains a non-injective unary operation, then $\text{CSP}(\Gamma)$ can be reduced to a problem over a smaller domain. One way to view this is that the presence of a non-injective unary operation in Γ^\triangleright indicates that constraints with relations chosen from Γ allow a form of global ‘substitutability’, similar to the notion defined by Freuder in [10].

If Γ^\triangleright does *not* contain any non-injective unary operations, then we shall say that Γ is *reduced*. The next theorem uses a general result from universal algebra [25, 27] to show that the possible choices for Γ^\triangleright are quite limited.

Theorem 3.8 *For any reduced set of relations Γ , on a finite set, either*

1. Γ^\triangleright contains essentially unary operations only, or
2. Γ^\triangleright contains an operation which is
 - (a) a constant operation; or
 - (b) a majority operation; or
 - (c) an idempotent binary operation (which is not a projection); or
 - (d) an affine operation; or
 - (e) a semiprojection.

Proof: The set of operations Γ^\triangleright contains all projections and is closed under composition, hence it constitutes a ‘clone’ [3, 27]. It was shown in [25] that any non-trivial clone on a finite set must contain a minimal clone, and that any minimal clone contains either

1. a non-identity unary operation; or
2. a constant operation; or
3. a majority operation; or
4. an idempotent binary operation (which is not a projection); or
5. an affine operation; or
6. a semiprojection.

Furthermore, if Γ is reduced, and Γ^\triangleright contains *any* operations which are *not* essentially unary, then it is straightforward to show, by considering such an operation of the smallest possible arity, that Γ^\triangleright contains an operation in one of the last five of these classes [27, 19]. ■

In the next two Sections we shall examine each of these possibilities in turn, in order to establish what can be said about the complexity of $\text{CSP}(\Gamma)$ in the various cases.

4 A necessary condition for tractability

In this Section we will show that any set of relations which is only closed under essentially unary operations will give rise to a class of constraint satisfaction problems which is NP-complete.

Theorem 4.1 *For any finite set of relations, Γ , over a finite set D , if Γ^\triangleright contains essentially unary operations only then $\text{CSP}(\Gamma)$ is NP-complete.*

Proof: When $|D| \leq 2$, then we may assume without loss of generality that $D \subseteq \{0, 1\}$, where 0 corresponds to the Boolean value **false** and 1 corresponds to the Boolean value **true**. It follows that the problem $\text{CSP}(\Gamma)$ corresponds to the GENERALISED SATISFIABILITY problem over the set of Boolean relations Γ , as defined in [26] (see also [11]).

It was established in [26] that this problem is NP-complete unless one of the following conditions holds:

1. Every relation in Γ contains the tuple $(0, 0, \dots, 0)$;
2. Every relation in Γ contains the tuple $(1, 1, \dots, 1)$;
3. Every relation in Γ is definable by a formula in conjunctive normal form in which each conjunct has at most one negated variable;
4. Every relation in Γ is definable by a formula in conjunctive normal form in which each conjunct has at most one unnegated variable;
5. Every relation in Γ is definable by a formula in conjunctive normal form in which each conjunct contains at most 2 literals;
6. Every relation in Γ is the set of solutions of a system of linear equations over the finite field $\text{GF}(2)$.

It is straightforward to show that in each of these cases Γ is closed under some operation which is not essentially unary (see [13] for details). Hence the result holds when $|D| = 2$.

For larger values of $|D|$ we proceed by induction. Assume that $|D| \geq 3$ and the result holds for all smaller values of $|D|$. Let $m = |D|(|D| - 1)$ and let $n = |D|^m$. Let M be an m by n matrix over D in which the columns consist of all possible m -tuples over D (in some order). Let R_0 be the relation consisting of all the tuples occurring as rows of M . The only condition we place on the choice of order for the columns of M is that $\pi_{1,2}(R_0) = \neq_D$, where \neq_D is the binary disequality relation over D , as defined in Example 2.4.

We now construct a relation \hat{R}_0 which is the ‘closest approximation’ to R_0 that we can obtain from the relations in Γ and the domain D using the Cartesian product, equality selection and projection operations:

$$\hat{R}_0 = \bigcap \{R \in (\Gamma \cup D^1)^+ \mid R_0 \subseteq R\}.$$

Since this is a finite intersection, and intersection is a special case of join, we have from Lemma 2.8 that $\hat{R}_0 \in (\Gamma \cup D^1)^+$. In other words, the relation \hat{R}_0 can be obtained as a derived constraint relation in some problem belonging to $\text{CSP}(\Gamma)$.

There are now two cases to consider:

1. If there exists some tuple $t_0 \in \hat{R}_0$ with $t_0[1] = t_0[2]$, then we will construct, using t_0 , an appropriate operation under which Γ is closed.

Define the function $\otimes : D^m \rightarrow D$ by setting $\otimes(d_1, d_2, \dots, d_m) = t_0[j]$ where j is the unique column of M corresponding to the m -tuple $\langle d_1, d_2, \dots, d_m \rangle$. We will show that Γ is closed under \otimes .

Choose any $R \in \Gamma$, and let p be the arity of R . We are required to show that R is closed under \otimes . Consider any sequence t_1, t_2, \dots, t_m of tuples of R (not necessarily distinct), and, for $i = 1, 2, \dots, p$, let c_i be the m -tuple $\langle t_1[i], t_2[i], \dots, t_m[i] \rangle$. For each pair of indices, i, j , such that $c_i = c_j$, apply the equality selection $\sigma_{i=j}$ to R , to obtain a new relation R' .

Now choose a maximal set of indices, $I = \{i_1, i_2, \dots, i_s\}$, such that the corresponding c_i are all distinct, and construct the relation $R'' = \pi_I(R') \times D^{n-|I|}$. Finally, permute the coordinate positions of R'' (by a sequence of Cartesian product, equality selection, and projection operations), such that $R'' \supseteq R_0$ (this is always possible, by the construction of R_0 and R''). Since $R'' \in (\Gamma \cup D^1)^+$, we know that t_0 is a tuple of R'' , by the definition of R_0 . Hence the appropriate projection of t_0 is an element of R , and R is closed under \otimes .

If \otimes is not essentially unary, then we have the result. Otherwise, let $f : D \rightarrow D$ be the corresponding unary operation, and set

$$\begin{aligned} f(D) &= \{f(d) \mid d \in D\}; \\ f(\Gamma) &= \{\{\langle f(d_1), f(d_2), \dots, f(d_r) \rangle \mid \langle d_1, d_2, \dots, d_r \rangle \in C\} \mid C \in \Gamma\}. \end{aligned}$$

By the choice of t_0 , f cannot be injective, so $|f(D)| < |D|$. By the inductive hypothesis, we know that either $\text{CSP}(f(\Gamma))$ is NP-complete (in which case $\text{CSP}(\Gamma)$ must also be NP-complete) or else $f(\Gamma)$ is closed under some operation \oplus which is not essentially unary (in which case Γ is closed under the operation $\oplus f$, which is also not essentially unary). Hence, the result follows by induction in this case.

2. Alternatively, if \hat{R}_0 contains no tuple t such that $t[1] = t[2]$, then $\pi_{1,2}(\hat{R}_0) = \neq_D$, so $\neq_D \in (\Gamma \cup D^1)^+$. But this implies that $\text{CSP}(\{\neq_D\})$ is reducible to $\text{CSP}(\Gamma)$, since every occurrence of the constraint relation \neq_D can be replaced with an equivalent collection of constraints with relations chosen from Γ . However, it was pointed out in Example 2.4 that $\text{CSP}(\{\neq_D\})$ corresponds to the GRAPH $|D|$ -COLORABILITY problem [11], which is NP-complete when $|D| \geq 3$. Hence, this implies that $\text{CSP}(\Gamma)$ is NP-complete, and the result holds in this case also. ■

Combining Theorem 4.1 with Theorem 3.8 gives the following necessary condition for tractability.

Corollary 4.2 *Assuming that P is not equal to NP , any tractable set of reduced relations must be closed under either a constant operation, or a majority operation, or an idempotent binary operation, or an affine operation, or a semiprojection.*

Note that the arity of a semiprojection is at most $|D|$, so for any finite set D there are only finitely many operations matching the given criteria, which means that there is a finite procedure to check whether this necessary condition is satisfied (see Corollary 6.5, below).

5 Sufficient conditions for tractability

We have shown in the previous section that when Γ is a tractable set of relations, then Γ^\triangleright must contain an operation from a limited range of types. We now consider each of these possibilities in turn, to determine whether or not they are sufficient to ensure tractability.

5.1 Constant operations

Closure under a constant operation is easily shown to be a sufficient condition for tractability.

Proposition 5.1 *For any set of relations Γ , if Γ is closed under a constant operation, then $\text{CSP}(\Gamma)$ is solvable in polynomial time.*

Proof: If every relation in Γ is closed under some constant operation \otimes , with constant value d , then every non-empty relation in Γ must contain the tuple $\langle d, d, \dots, d \rangle$. Hence, in this case, the decision problem for any constraint satisfaction problem instance \mathcal{P} in $\text{CSP}(\Gamma)$ is clearly trivial to solve, since \mathcal{P} either contains an empty constraint, in which case it does not have a solution, or else \mathcal{P} allows the solution in which every variable is assigned the value d . ■

The class of sets of relations closed under some constant operation is a rather trivial tractable class. It is referred to in [14] as **Class 0**.

Example 5.2 Let \top denote the unary operation on the domain $D = \{0, 1, 2\}$ which returns the constant value 1. The constraint R_3 defined in Example 2.6 is closed under \top , since applying the \top operation to any element of R_3 yields the tuple $\langle 1, 1 \rangle$, which is

an element of R_3 . The constraint R_2 defined in Example 2.6 is *not* closed under \top , since applying the \top operation to any element of R_2 yields the tuple $\langle 1, 1, 1 \rangle$, which is not an element of R_2 . In fact, R_2 is clearly not closed under any constant operation. \square

Example 5.3 When $D = \{\mathbf{true}, \mathbf{false}\}$, there are only two possible constant operations on D .

The first two tractable subproblems of the GENERALISED SATISFIABILITY problem identified by Schaefer in [26] correspond to the tractable classes of relations characterised by closure under these two constant operations. \square

5.2 Majority operations

We will now show that closure under a majority operation is a sufficient condition for tractability.

We first establish that when a relation R is closed under a majority operation, any constraint involving R can be decomposed into binary constraints.

Proposition 5.4 *Let R be a relation of arity n which is closed under a majority operation, and let C be any constraint $C = (S, R)$ constraining the variables in S with relation R .*

For any problem \mathcal{P} with constraint C , the problem \mathcal{P}' which is obtained by replacing C by the set of constraints

$$\{((S[i], S[j]), \pi_{i,j}(R)) \mid 1 \leq i < j \leq n\}$$

has exactly the same solutions as \mathcal{P} .

Proof: It is clear that any solution to \mathcal{P} is a solution to \mathcal{P}' , since \mathcal{P}' is obtained by taking binary projections of a constraint from \mathcal{P} .

Now let σ be any solution to \mathcal{P}' , and set $t = \langle \sigma(S[1]), \sigma(S[2]), \dots, \sigma(S[n]) \rangle$. We shall prove, by induction on n , that $t \in R$, thereby establishing that σ is a solution to \mathcal{P} .

For $n < 3$ the result holds trivially, so assume that $n \geq 3$, and that the result holds for all smaller values. Let $I = \{1, 2, \dots, n\}$ be the set of indices of positions in S and choose $i_1, i_2, i_3 \in I$. By Proposition 3.5 and the inductive hypothesis, applied to $\pi_{I \setminus \{i_j\}}(R)$, there is some $t_j \in R$ which agrees with t at all positions except i_j , for $j = 1, 2, 3$. Since R is closed under a majority operation, applying this operation to t_1, t_2, t_3 gives $t \in R$. \blacksquare

Example 5.5 Recall the relation R_2 defined in Example 2.6.

It was shown in Example 3.4 that R_2 is closed under the operation Δ . Since this operation is a majority operation, we know by Proposition 5.4 that any constraint with relation R_2 can be decomposed into a collection of binary constraints with the following relations:

- $\pi_{1,2}(R_2) = \{\langle 0, 1 \rangle, \langle 1, 2 \rangle, \langle 2, 0 \rangle\}$
- $\pi_{1,3}(R_2) = \{\langle 1, 0 \rangle, \langle 2, 1 \rangle, \langle 0, 2 \rangle\}$
- $\pi_{2,3}(R_2) = \{\langle 0, 1 \rangle, \langle 1, 2 \rangle, \langle 2, 0 \rangle\}$

□

It is, of course, not always the case that a constraint can be replaced by a collection of binary constraints on the same variables. In many cases the binary projections of the constraint relation allow extra solutions, as the following example demonstrates.

Example 5.6 Recall the relation δ on domain $D = \{0, 1\}$ defined in Example 2.5. The binary projections of δ are as follows:

- $\pi_{1,2}(\delta) = \{\langle 0, 0 \rangle, \langle 0, 1 \rangle, \langle 1, 0 \rangle, \langle 1, 1 \rangle\}$
- $\pi_{1,3}(\delta) = \{\langle 0, 0 \rangle, \langle 0, 1 \rangle, \langle 1, 0 \rangle, \langle 1, 1 \rangle\}$
- $\pi_{2,3}(\delta) = \{\langle 0, 0 \rangle, \langle 0, 1 \rangle, \langle 1, 0 \rangle, \langle 1, 1 \rangle\}$

The join of these binary projections contains the tuples $\langle 0, 0, 0 \rangle$ and $\langle 1, 1, 1 \rangle$, which are not elements of δ . It clearly follows that δ cannot be replaced by any set of binary constraints on the same variables. □

Theorem 5.7 *Let Γ be any set of relations over a finite domain, D .*

If Γ is closed under a majority operation, then $\text{CSP}(\Gamma)$ is solvable in polynomial time.

Proof: For any problem instance \mathcal{P} in $\text{CSP}(\Gamma)$ we can impose strong $(|D| + 1)$ -consistency [5] in polynomial time to obtain a new instance \mathcal{P}' with the same solutions. All of the constraints in \mathcal{P}' are elements of Γ^+ , and so they are closed under a majority operation, by Proposition 3.5. Hence, all of the constraints of \mathcal{P}' are decomposable into binary constraints by Proposition 5.4. Hence, by Corollary 3.2 of [5], \mathcal{P}' is solvable in polynomial time. ■

Example 5.8 When $D = \{\text{true}, \text{false}\}$, there is only one possible majority operation on D , (which is equal to the Δ operation defined in Example 3.4). It is easily shown that all possible binary Boolean relations are closed under Δ . Hence, it follows from Proposition 5.4 that the Boolean relations of arbitrary arity which are closed under this majority operation are precisely the relations which are definable by a formula in conjunctive normal form in which each conjunct contains at most 2 literals. Hence, if a set of Boolean relations Γ is closed under a majority operation, then $\text{CSP}(\Gamma)$ is equivalent to the 2-SATISFIABILITY problem (2-SAT) [24], which is well-known to be a tractable subproblem of the SATISFIABILITY problem [26]. □

Recall the class of tractable constraints identified independently in [4] and [17], and referred to as *0/1/all constraints* or *implicational constraints*. (This class of tractable constraints is referred to as **Class I** in [14].) It was shown in [14] that these constraints are in fact precisely the relations closed under the majority operation Δ defined in Example 3.4. This result is rather unexpected, in view of the fact that 0/1/all constraints were originally defined purely in terms of their syntactic structure [4].

However, we remark here that the class of tractable sets of relations defined by closure under some majority operation is a true generalization of the class containing all sets of 0/1/all constraints. In other words, there exist tractable sets of relations which are closed under some majority operation but are not closed under the Δ operation, as the following example demonstrates.

Example 5.9 Let μ be the ternary majority operation on $D = \{0, 1, 2\}$ which returns the median value of its three arguments (in the standard ordering of D).

Recall the relation R_3 defined in Example 2.6. It is easy to show that R_3 is closed under μ , since applying the μ operation to any 3 elements of R_3 yields an element of R_3 . For example,

$$\mu(\langle 1, 0 \rangle, \langle 0, 2 \rangle, \langle 0, 1 \rangle) = \langle 0, 1 \rangle \in R_3.$$

Hence, by Theorem 5.7, $\text{CSP}(\{R_3\})$ is tractable.

However, it was shown in Example 3.4 that R_3 is not closed under Δ , and hence R_3 is not a 0/1/all constraint. \square

5.3 Binary operations

We first show that closure under an *arbitrary* idempotent binary operation is *not* in general sufficient to ensure tractability.

Lemma 5.10 *There exists a set of relations Γ closed under an idempotent binary operation (which is not a projection) such that $\text{CSP}(\Gamma)$ is NP-complete.*

Proof: Consider the binary operation \square on the set $D = \{0, 1, 2, 3\}$ which is defined by the following table:

\square	0	1	2	3
0	0	1	0	1
1	0	1	0	1
2	2	3	2	3
3	2	3	2	3

This operation is idempotent but it is not a projection (in fact, it is an example of a form of binary operation known as a ‘rectangular band’ [21].)

Now consider the functions $b_1 : D \rightarrow \{0, 1\}$ and $b_2 : D \rightarrow \{0, 1\}$ which return the first and second bit in the binary expression for the numerical value of each element of D .

Using these functions, we define ternary relations R_1 and R_2 over D , as follows:

$$R_1 = \{\langle d, d', d'' \rangle \in D^3 \mid (b_1(d) \neq b_1(d')) \vee (b_1(d') \neq b_1(d'')) \vee (b_1(d'') \neq b_1(d))\}$$

$$R_2 = \{\langle d, d', d'' \rangle \in D^3 \mid (b_2(d) \neq b_2(d')) \vee (b_2(d') \neq b_2(d'')) \vee (b_2(d'') \neq b_2(d))\}$$

Finally, we define $R = R_1 \cap R_2$.

It is easily shown that R is closed under \square , since applying the operation \square to any 2 elements of R yields an element of R .

However, it can also be shown that the NOT-ALL-EQUAL SATISFIABILITY problem [26], which is known to be NP-complete, is reducible in polynomial time to $\text{CSP}(\{R\})$. Hence, $\text{CSP}(\{R\})$ is NP-complete, and the result follows. \blacksquare

We now describe some additional conditions which may be imposed on binary operations. It will be shown below that closure under any binary operation satisfying these additional conditions is a sufficient condition for tractability.

Definition 5.11 *Let $\square : D^2 \rightarrow D$ be an idempotent binary operation on the set D such that, for all $d_1, d_2, d_3 \in D$,*

- $\sqcap(\sqcap(d_1, d_2), d_3) = \sqcap(d_1, \sqcap(d_2, d_3))$; and *(Associativity)*
- $\sqcap(d_1, d_2) = \sqcap(d_2, d_1)$. *(Commutativity)*

Then \sqcap is said to be an ACI operation.

We will make use of the following result about ACI operations, which is well-known from elementary algebra [3, 21].

Lemma 5.12 *Let \sqcap be an ACI operation on the set D . The binary relation R on D defined by*

$$R(d_1, d_2) \iff \sqcap(d_1, d_2) = d_2$$

is a partial order on D in which any two elements d_1, d_2 have a least upper bound given by $\sqcap(d_1, d_2)$.

It follows from Lemma 5.12 that any (finite) non-empty set $D' \subseteq D$ which is \sqcap -closed contains a least upper bound with respect to the partial order R . This upper bound will be denoted $\sqcap(D')$.

Using Lemma 5.12, we now show that relations which are closed under some arbitrary ACI operation form a tractable class.

Theorem 5.13 *For any set of relations Γ over a finite domain D , if Γ is closed under some ACI operation, then $\text{CSP}(\Gamma)$ is solvable in polynomial time.*

Proof: Let Γ be a set of relations closed under the ACI operation \sqcap , and let \mathcal{P} be any problem instance in $\text{CSP}(\Gamma)$. First enforce pairwise consistency to obtain a new instance \mathcal{P}' with the same set of solutions which is pairwise consistent. Such a \mathcal{P}' can be obtained by forming the join of every pair of constraints in \mathcal{P} , replacing these constraints with the (possibly smaller) constraints obtained by projecting down to the original scopes, and then repeating this process until there are no further changes in the constraints. The time complexity of this procedure is polynomial in the size of \mathcal{P} , and the resulting \mathcal{P}' is a member of $\text{CSP}(\Gamma^+)$. Hence, all the constraint relations in \mathcal{P}' are closed under \sqcap , by Proposition 3.5.

Now let $D(v)$ denote the set of values allowed for variable v by the constraints of \mathcal{P}' . Since $D(v)$ equals the projection of some \sqcap -closed constraint onto v , it must be \sqcap -closed, by Proposition 3.5. There are two cases to consider:

1. If any of the sets $D(v)$ is empty then \mathcal{P}' has no solutions, so the decision problem is trivial.
2. On the other hand, if all of these sets are non-empty, then we claim that assigning the value $\sqcap(D(v))$ to each variable v gives a solution to \mathcal{P}' , so the decision problem is again trivial. To establish this claim, consider any constraint $C = (S, R)$ in \mathcal{P}' , with relation R of arity n , and scope S . For each $i \in \{1, 2, \dots, n\}$, there must be some tuple $t_i \in R$ such that $t_i[i] = \{\sqcap(D(S[i]))\}$, by the definition of $D(S[i])$. Now consider the tuple $t = \sqcap(t_1, \sqcap(t_2, \dots, \sqcap(t_{r-1}, t_r)) \dots)$. We know that $t \in R$, since R is closed under \sqcap . Furthermore, for each i , $t[i] = \sqcap(D(S[i]))$, because $\sqcap(D(S[i]))$ is an upper bound of $D(S[i])$, so

$$\sqcap(t, \sqcap(D(S[i]))) = \sqcap(D(S[i]))$$

for all $d \in D(S[i])$. Hence the constraint C allows the assignment of $\sqcap(D(v))$ to each variable v in S . Since C was arbitrary, we have shown that this assignment is a solution to \mathcal{P}' , and hence a solution to \mathcal{P} . ■

Example 5.14 When $D = \{\mathbf{true}, \mathbf{false}\}$, there are only two idempotent binary operations on D (which are not projections), corresponding to the logical AND operation and the logical OR operation. These two operations are both ACI operations, and they correspond to the two possible orderings of D .

It is well-known [6, 16, 24] that a Boolean relation is closed under AND if and only if it can be defined by a Horn sentence (that is, a conjunction of clauses each of which contains at most one unnegated literal). Hence, if a set of Boolean relations Γ is closed under AND, then $\text{CSP}(\Gamma)$ is equivalent to the Horn clause satisfiability problem, HORNSAT [24], which is a tractable subproblem of the SATISFIABILITY problem [26].

Similarly, a Boolean relation is closed under OR if and only if it can be defined by a conjunction of clauses each of which contains at most one negated literal, and this class of relations also gives rise to a tractable subproblem of the SATISFIABILITY problem [26]. □

Example 5.15 Let D be a finite subset of the natural numbers. The operation $MAX : D^2 \rightarrow D$ which returns the larger of any pair of numbers is an ACI operation. The following types of arithmetic constraints (amongst many others) are closed under this operation:

- $aX \neq b$
- $aX = bY + c$
- $aX \leq bY + c$
- $aX \geq bY + c$
- $a_1X + a_2Y + \dots + a_rZ \geq c$
- $aXY \geq c$
- $(a_1X \geq b_1) \vee (a_2Y \geq b_2) \vee (a_3Z \leq b_3)$

where upper-case letters represent variables and lower-case letters represent positive constants. Hence, by Theorem 5.13 it is possible to determine efficiently whether any collection of constraints of these types has a solution. These constraints include (and extend) the ‘basic’ arithmetic constraints allowed by the well-known constraint programming language, CHIP [29]. □

The class of tractable constraints first identified in [16], and referred to as *max-closed constraints*, are in fact relations closed under an \sqcap operation with the additional property that the partial order R , defined in Lemma 5.12, is a total ordering of D . Hence, a set of constraints is max-closed if and only if the constraint relations are closed under some

specialized ACI operation of this kind (see, for example, Example 5.15). This class of tractable constraints is referred to as **Class II** in [14].

However, we remark here that the class of tractable sets of relations defined by closure under some ACI operation is a true generalization of the class containing all sets of max-closed constraints. In other words, there exist tractable relations which are closed under some ACI operation but are not closed under the maximum operation associated with any (total) ordering of the domain. (An example of such a relation is the relation R_1 defined in Example 2.6, see Example 6.4 below.)

5.4 Affine operations

We will now show that closure under an affine operation is a sufficient condition for tractability.

This result was established in [14] using elementary methods, for the special case when the domain D contains a prime number, p , of elements. It was shown in [14] that, in this special case, constraints which are closed under an affine operation correspond precisely to constraints which may be expressed as conjunctions of *linear equations* modulo p . (This class of tractable constraints is referred to as **Class III** in [14].)

We now generalise this result to arbitrary finite domain sizes by making use of a result stated by Feder and Vardi in [8].

Theorem 5.16 ([8]) *For any finite group G , and any set Γ of cosets of subgroups of direct products of G , $\text{CSP}(\Gamma)$ is solvable in polynomial time.*

Corollary 5.17 *For any set of relations Γ , if Γ is closed under an affine operation, then $\text{CSP}(\Gamma)$ is solvable in polynomial time.*

Proof: By Definitions 3.1 and 3.2, any relation R which is closed under an affine operation is a subset of a direct product of some Abelian group, with the property that for all $t_1, t_2, t_3 \in R$, $t_1 - t_2 + t_3 \in R$. However, this is equivalent to saying that R is a coset of a subgroup of this direct product group [21], so we may apply Theorem 5.16 to Γ to obtain the result. ■

Example 5.18 Let ∇ be the affine operation on $D = \{0, 1, 2\}$ which is defined by $\nabla(d_1, d_2, d_3) = d_1 - d_2 + d_3$, where addition and subtraction are both modulo 3.

The relation R_2 defined in Example 2.6 is closed under ∇ , since applying the ∇ operation to any 3 elements of R_2 yields an element of R_2 . For example,

$$\nabla(\langle 0, 1, 2 \rangle, \langle 1, 2, 0 \rangle, \langle 2, 0, 1 \rangle) = \langle 1, 2, 0 \rangle \in R_2.$$

Since $|D|$ is prime, the results of [14] indicate that R_2 must be the set of solutions to some system of linear equations over the integers modulo 3. In fact, we have

$$R_2 = \{\langle x_1, x_2, x_3 \rangle \mid (x_1 - x_2 \equiv 2 \pmod{3}) \wedge (x_2 - x_3 \equiv 2 \pmod{3})\}.$$

□

Example 5.19 Let G be the Abelian group $\langle D, +, - \rangle$, where $D = \{0, 1, 2, 3\}$ and the $+$ operation is defined by the following table:

$+$	0	1	2	3
0	0	1	2	3
1	1	0	3	2
2	2	3	0	1
3	3	2	1	0

Now let ∇ be the affine operation on $D = \{0, 1, 2, 3\}$ which is defined by $\nabla(d_1, d_2, d_3) = d_1 - d_2 + d_3$, where addition and subtraction are as defined in G .

Any relation R over D which is a coset of a subgroup of a direct product of G will be closed under ∇ , and hence $\text{CSP}(\{R\})$ will be tractable by Corollary 5.17. One example of such a relation is the following:

$$R = \{\langle 0, 1, 2 \rangle, \langle 1, 1, 2 \rangle, \langle 2, 2, 0 \rangle, \langle 3, 2, 0 \rangle\}$$

It is easily seen that in this case R is *not* the set of solutions to any system of linear equations over a field. \square

Example 5.20 When $D = \{\mathbf{true}, \mathbf{false}\}$, there is only one possible Abelian group structure over D and hence only one possible affine operation on D .

If a set of Boolean relations Γ is closed under this affine operation, then $\text{CSP}(\Gamma)$ is equivalent to the problem of solving a set of simultaneous linear equations over the integers modulo 2. This corresponds to the final tractable subproblem of the GENERALISED SATISFIABILITY problem identified by Schaefer in [26]. \square

5.5 Semiprojections

We now show that closure under a semiprojection operation is *not* in general a sufficient condition for tractability. In fact we shall establish a much stronger result, which shows that even being closed under *all* semiprojections is not sufficient to ensure tractability.

Lemma 5.21 *For any finite set D , with $|D| \geq 3$, there exists a set of relations Γ over D , such that Γ is closed under all semiprojections on D , and $\text{CSP}(\Gamma)$ is NP-complete.*

Proof: Let D be a finite set with $|D| \geq 3$ and let d_1, d_2 be elements of D . Consider the relation $R = \{\langle d_1, d_1, d_2 \rangle, \langle d_1, d_2, d_1 \rangle, \langle d_2, d_1, d_1 \rangle, \langle d_2, d_2, d_1 \rangle, \langle d_2, d_1, d_2 \rangle, \langle d_1, d_2, d_2 \rangle\}$. This relation is closed under all semiprojections on D , since any 3 elements of R contain at most two distinct values in each coordinate position.

However, if we identify d_1 with the Boolean value **true** and d_2 with the Boolean value **false**, then it is easy to see that $\text{CSP}(\{R\})$ is isomorphic to the NOT-ALL-EQUAL SATISFIABILITY problem [11], which is NP-complete [26] (see Example 2.5). \blacksquare

It is currently unknown whether there are tractable sets of relations closed under some combination of semiprojections, unary operations and binary operations which are not included in any of the tractable classes listed above.

However, when $|D| = 2$ the situation is very simple, as the next example shows.

Example 5.22 When $D = \{\mathbf{true}, \mathbf{false}\}$, there are no semiprojections on D , so there are no subproblems of the SATISFIABILITY problem which are characterised by a closure operation of this form. \square

6 Calculating closure operations

For any set of relations Γ , over a set D , the operations under which Γ is closed are simply mappings from D^k to D , for some k , which satisfy certain constraints, as described in Definition 3.3. In this Section we show that it is possible to identify these operations by solving a single constraint satisfaction problem in $\text{CSP}(\Gamma)$. In fact, we shall show that these closure operations are precisely the solutions to a constraint satisfaction problem of the following form.

Definition 6.1 *Let Γ be a set of relations over a finite domain D .*

For any natural number $m > 0$, the indicator problem for Γ of order m is defined to be the constraint satisfaction problem $\mathcal{IP}(\Gamma, m)$ with

- *Set of variables D^m ;*
- *Domain of values D ;*
- *Set of constraints $\{C_1, C_2, \dots, C_q\}$, such that for each $R \in \Gamma$, and for each sequence t_1, t_2, \dots, t_m of tuples from R , there is a constraint $C_i = (S_i, R)$ with $S_i = (v_1, v_2, \dots, v_n)$ where n is the arity of R and $v_j = \langle t_1[j], t_2[j], \dots, t_m[j] \rangle$.*

Example 6.2 Consider the relation R_1 over $D = \{0, 1, 2\}$, defined in Example 2.6.

The indicator problem for $\{R_1\}$ of order 1, $\mathcal{IP}(\{R_1\}, 1)$, has 3 variables and 4 constraints. The set of variables is

$$\{\langle 0 \rangle, \langle 1 \rangle, \langle 2 \rangle\},$$

and the set of constraints is

$$\{ ((\langle 0 \rangle, \langle 0 \rangle), R_1), \\ ((\langle 0 \rangle, \langle 1 \rangle), R_1), \\ ((\langle 1 \rangle, \langle 2 \rangle), R_1), \\ ((\langle 2 \rangle, \langle 1 \rangle), R_1) \}.$$

The indicator problem for $\{R_1\}$ of order 2, $\mathcal{IP}(\{R_1\}, 2)$, has 9 variables and 16 constraints. The set of variables is

$$\{\langle 0, 0 \rangle, \langle 0, 1 \rangle, \langle 0, 2 \rangle, \langle 1, 0 \rangle, \langle 1, 1 \rangle, \langle 1, 2 \rangle, \langle 2, 0 \rangle, \langle 2, 1 \rangle, \langle 2, 2 \rangle\},$$

and the set of constraints is

$$\{ ((\langle 0, 0 \rangle, \langle 0, 0 \rangle), R_1), ((\langle 0, 0 \rangle, \langle 0, 1 \rangle), R_1), \\ ((\langle 0, 0 \rangle, \langle 1, 0 \rangle), R_1), ((\langle 0, 0 \rangle, \langle 1, 1 \rangle), R_1), \\ ((\langle 0, 1 \rangle, \langle 0, 2 \rangle), R_1), ((\langle 0, 1 \rangle, \langle 1, 2 \rangle), R_1), \\ ((\langle 0, 2 \rangle, \langle 0, 1 \rangle), R_1), ((\langle 0, 2 \rangle, \langle 1, 1 \rangle), R_1), \\ ((\langle 1, 0 \rangle, \langle 2, 0 \rangle), R_1), ((\langle 1, 0 \rangle, \langle 2, 1 \rangle), R_1), \\ ((\langle 1, 1 \rangle, \langle 2, 2 \rangle), R_1), ((\langle 1, 2 \rangle, \langle 2, 1 \rangle), R_1), \\ ((\langle 2, 0 \rangle, \langle 1, 0 \rangle), R_1), ((\langle 2, 0 \rangle, \langle 1, 1 \rangle), R_1), \\ ((\langle 2, 1 \rangle, \langle 1, 2 \rangle), R_1), ((\langle 2, 2 \rangle, \langle 1, 1 \rangle), R_1) \}.$$

Further illustrative examples of indicator problems are given in [15]. □

Solutions to the indicator problem for Γ of order m are functions from D^m to D , or in other words, m -ary operations on D . We now show that they are precisely the m -ary operations under which Γ is closed.

Theorem 6.3 *For any set of relations Γ over domain D , the set of solutions to $\mathcal{IP}(\Gamma, m)$ is equal to the set of m -ary operations under which Γ is closed.*

Proof: By Definition 3.3, we know that Γ is closed under the m -ary operation \otimes if and only if $\otimes(t_1, t_2, \dots, t_m) \in R$ for each possible choice of $R \in \Gamma$ and $t_1, t_2, \dots, t_m \in D$ (not necessarily all distinct). But this is equivalent to saying that \otimes satisfies all the constraints in $\mathcal{IP}(\Gamma, m)$, so the result follows. ■

Example 6.4 Consider the relation R_1 over $D = \{0, 1, 2\}$, defined in Example 2.6.

The indicator problem for $\{R_1\}$ of order 1, defined in Example 6.2, has 2 solutions, which may be expressed in tabular form as follows:

	Variables		
	$\langle 0 \rangle$	$\langle 1 \rangle$	$\langle 2 \rangle$
Solution 1	0	0	0
Solution 2	0	1	2

One of these solutions is a constant operation, so $\text{CSP}(\{R_1\})$ is tractable, by Proposition 5.1. In fact, any problem in $\text{CSP}(\{R_1\})$ has the solution which assigns the value 0 to each variable, so the complexity of $\text{CSP}(\{R_1\})$ is trivial.

The indicator problem for $\{R_1\}$ of order 2, defined in Example 6.2, has 4 solutions, which may be expressed in tabular form as follows:

	Variables								
	$\langle 0, 0 \rangle$	$\langle 0, 1 \rangle$	$\langle 0, 2 \rangle$	$\langle 1, 0 \rangle$	$\langle 1, 1 \rangle$	$\langle 1, 2 \rangle$	$\langle 2, 0 \rangle$	$\langle 2, 1 \rangle$	$\langle 2, 2 \rangle$
Solution 1	0	0	0	0	0	0	0	0	0
Solution 2	0	1	2	0	1	2	0	1	2
Solution 3	0	0	0	1	1	1	2	2	2
Solution 4	0	0	0	0	1	0	0	0	2

The first of these solutions is a constant operation, and the second and third are essentially unary operations. However, the fourth solution shown in the table is more interesting. It is easily checked that this operation is an associative, commutative, idempotent (ACI) binary operation, so we have a second proof that $\text{CSP}(\{R_1\})$ is tractable, by Theorem 5.13. Furthermore, this result shows that R_1 can be combined with any other relations (of any arity) which are also closed under this ACI operation to obtain larger tractable problem classes. □

Corollary 6.5 *For any set of relations Γ over a domain D , with $|D| \geq 3$, if all solutions to $\mathcal{IP}(\Gamma, |D|)$ are essentially unary, then $\text{CSP}(\Gamma)$ is NP-complete.*

Proof: Follows from Theorem 3.8, Theorem 4.1, and Theorem 6.3. ■

Example 6.6 Recall the relations R_1 , R_2 and R_3 defined in Example 2.6. It has been shown in Examples 6.4, 5.18, and 5.2 that a set containing any *one* of these relations on its own is tractable.

For any set Γ containing more than one of these relations, it can be shown, using Corollary 6.5, that $\text{CSP}(\Gamma)$ is NP-complete. \square

In the special case when $|D| = 2$ we obtain an even stronger result.

Corollary 6.7 *For any set of relations Γ over a domain D , with $|D| = 2$, if all solutions to $\mathcal{IP}(\Gamma, 3)$ are essentially unary then $\text{CSP}(\Gamma)$ is NP-complete, otherwise it is polynomial.*

Proof: It has been shown in Examples 5.3, 5.8, 5.14, 5.20, and 5.22 that when $|D| = 2$ all possible closure operations of the restricted types specified in Corollary 4.2 are sufficient to ensure tractability. \blacksquare

This result demonstrates that solving the indicator problem of order 3 provides a simple and complete test for tractability of any set of relations over a domain with 2 elements. This answers a question posed by Schaefer in 1978 [26] concerning the existence of an efficient test for tractability in the GENERALISED SATISFIABILITY problem. Note that carrying out the test requires finding the solutions to a constraint satisfaction problem with just 8 Boolean variables.

7 Conclusion

In this paper we have shown how the algebraic properties of relations can be used to distinguish between sets of relations which give rise to tractable constraint satisfaction problems and those which give rise to NP-complete problems. Furthermore, we have proposed a method for determining the operations under which a set of relations is closed by solving a particular form of constraint satisfaction problem, which we have called an indicator problem.

For problems where the domain contains just two elements these results provide a necessary and sufficient condition for tractability (assuming that P is not equal to NP), and an efficient test to distinguish the tractable sets of relations.

For problems with larger domains we have described algebraic closure properties which are a necessary condition for tractability. We have also shown that in many cases these closure properties are sufficient to ensure tractability.

In particular, we have shown that closure under any constant operation, any majority operation, any ACI operation, or any affine operation, is a sufficient condition for tractability. It can be shown using the results of [13] that for any operation of one of these types, the set, Γ , containing all relations which are closed under that operation is a *maximal* set of tractable relations. In other words, the addition of *any* other relation which is not closed under the same operation changes $\text{CSP}(\Gamma)$ from a tractable problem into an NP-complete problem. Hence, the tractable classes defined in this way are as large as possible.

We are now investigating the application of these results to particular problem types, such as temporal problems involving subsets of the interval algebra. We are also attempting to determine how the presence of particular algebraic closure properties in the constraints can be used to derive appropriate efficient algorithms for tractable problems.

Acknowledgments

This research was partially supported by EPSRC research grant GR/L09936 and by the ‘British-Flemish Academic Research Collaboration Programme’ of the Belgian National Fund for Scientific Research and the British Council. We are also grateful to Martin Cooper for many helpful discussions and for suggesting the proof of Theorem 5.7.

References

- [1] Codd, E.F., “A relational model of data for large shared databanks”, *Communications of the ACM* 13 (1970), pp. 377–387.
- [2] Cohen D., Gyssens, M., Jeavons, P., “Derivation of constraints and database relations”, In *Proceedings 2nd International Conference on Constraint Programming—CP ’96* (Boston, August 1996), *Lecture Notes in Computer Science*, 1118 Springer-Verlag, Berlin/New York, pp. 134–148.
- [3] Cohn, P.M., *Universal Algebra*, Harper & Row, New York (1965).
- [4] Cooper, M.C., Cohen, D.A., Jeavons, P.G., “Characterizing tractable constraints”, *Artificial Intelligence* 65, (1994), pp. 347–361.
- [5] Dechter, R., “From local to global consistency”, *Artificial Intelligence* 55 (1992), pp. 87–107.
- [6] Dechter, R. & Pearl, J., “Structure identification in relational data”, *Artificial Intelligence* 58 (1992) pp. 237–270.
- [7] Dechter, R. & Pearl J. “Network-based heuristics for constraint-satisfaction problems”, *Artificial Intelligence* 34 (1988), pp. 1–38.
- [8] Feder, T. & Vardi, M.Y., “Monotone monadic SNP and constraint satisfaction”, In *Proceedings of 25th ACM Symposium on the Theory of Computing (STOC)* (1993), pp. 612–622.
- [9] Freuder, E.C., “A sufficient condition for backtrack-bounded search”, *Journal of the ACM* 32 (1985) pp. 755–761.
- [10] Freuder, E.C., “Eliminating interchangeable values in constraint satisfaction problems”, *Proceedings of AAAI-91*, (1991) pp. 227–233.
- [11] Garey, M.R. & Johnson, D.S., *Computers and intractability: a guide to the theory of NP-completeness*, Freeman, San Francisco, California, (1979).
- [12] Gyssens, M., Jeavons, P., Cohen, D., “Decomposing constraint satisfaction problems using database techniques”, *Artificial Intelligence* 66, (1994), pp. 57–89.
- [13] Jeavons, P.G., “An algebraic characterization of tractable constraints”, Technical Report CSD-TR-95-05, Royal Holloway, University of London (1995).

- [14] Jeavons, P., Cohen D., Gyssens, M., “A unifying framework for tractable constraints”, In *Proceedings 1st International Conference on Principles and Practice of Constraint Programming—CP '95* (Cassis, France, September 1995), *Lecture Notes in Computer Science*, 976, Springer-Verlag, Berlin/New York, (1995), pp. 276–291.
- [15] Jeavons, P., Cohen D., Gyssens, M., “A test for tractability”, In *Proceedings 2nd International Conference on Constraint Programming—CP '96* (Boston, August 1996), *Lecture Notes in Computer Science*, 1118 Springer-Verlag, Berlin/New York, pp. 267–281.
- [16] Jeavons, P.G. & Cooper, M.C., “Tractable constraints on ordered domains”, *Artificial Intelligence* 79 (1996), pp. 327–339.
- [17] Kirousis, L., “Fast parallel constraint satisfaction”, *Artificial Intelligence* 64, (1993), pp. 147–160..
- [18] Ladkin, P.B. & Maddux, R.D., “On binary constraint problems”, *Journal of the ACM* 41 (1994), pp. 435–469.
- [19] Machida, H. & Rosenberg, I.G., “Classifying essentially minimal clones”, *Proceeding 14th Internat. Symp. on Multiple-Valued Logic*, (Winnipeg, 1984), (1984) pp. 4–7.
- [20] Mackworth, A.K. “Consistency in networks of relations”, *Artificial Intelligence* 8 (1977) pp. 99–118.
- [21] McKenzie, R.N., McNulty, G.F., Taylor, W.F., *Algebras, lattices and varieties. Volume I*, Wadsworth and Brooks, California (1987).
- [22] Montanari, U., “Networks of constraints: fundamental properties and applications to picture processing”, *Information Sciences* 7 (1974), pp. 95–132.
- [23] Montanari, U. & Rossi, F., “Constraint relaxation may be perfect”, *Artificial Intelligence* 48 (1991), pp. 143–170.
- [24] Papadimitriou, C.H., *Computational Complexity*, Addison-Wesley, (1994).
- [25] Rosenberg, I.G., “Minimal clones I: the five types”, in *Lectures in Universal Algebra (Proc. Conf. Szeged 1983)*, Colloq. Math. Soc. Janos Bolyai 43, North-Holland, (1986), pp. 405–427.
- [26] Schaefer, T.J., “The complexity of satisfiability problems”, *Proc 10th ACM Symposium on Theory of Computing (STOC)* , (1978) pp. 216–226.
- [27] Szendrei, A., *Clones in Universal Algebra*, Seminaires de Mathematiques Superieures 99, University of Montreal, (1986).
- [28] van Beek, P., “On the Minimality and Decomposability of Row-Convex Constraint Networks”, *Proceedings of the Tenth National Conference on Artificial Intelligence, AAAI-92*, MIT Press, (1992) pp. 447–452.
- [29] Van Hentenryck, P., Deville, Y., Teng, C-M., “A generic arc-consistency algorithm and its specializations”, *Artificial Intelligence* 57 (1992), pp. 291–321.