

Symmetry Definitions for Constraint Satisfaction Problems

David Cohen¹, Peter Jeavons², Christopher Jefferson³, Karen E. Petrie³, and
Barbara M. Smith⁴

¹ Department of Computer Science, Royal Holloway, University of London, UK
D.Cohen@rhul.ac.uk,

² Computing Laboratory, University of Oxford, UK
peter.jeavons@comlab.ox.ac.uk,

³ School of Computer Science, University of St Andrews, UK
{caj,kep}@dcs.st-and.ac.uk,

⁴ Cork Constraint Computation Centre, University College Cork, Ireland
b.smith@4c.ucc.ie

Abstract. We review the many different definitions of symmetry for constraint satisfaction problems (CSPs) that have appeared in the literature, and show that a symmetry can be defined in two fundamentally different ways: as an operation preserving the solutions of a CSP instance, or else as an operation preserving the constraints. We refer to these as *solution symmetries* and *constraint symmetries*. We define a constraint symmetry more precisely as an automorphism of a hypergraph associated with a CSP instance, the microstructure complement. We show that the solution symmetries of a CSP instance can also be obtained as the automorphisms of a related hypergraph, the *k-ary nogood hypergraph* and give examples to show that some instances have many more solution symmetries than constraint symmetries. Finally, we discuss the practical implications of these different notions of symmetry.

1 Introduction

The issue of *symmetry* is now widely recognised as of fundamental importance in constraint satisfaction problems (CSPs). It seems self-evident that in order to deal with symmetry we should first agree what we mean by symmetry. Surprisingly, this appears not to be true: researchers in this area have defined symmetry in fundamentally different ways, whilst often still identifying the same collection of symmetries in a given problem and dealing with them in the same way.

In this paper, we first survey the various symmetry definitions that have appeared in the literature. We show that the existing definitions reflect two distinct views of symmetry: that symmetry is a property of the solutions, i.e. that any mapping that preserves the solutions is a symmetry; or that symmetry preserves the constraints, and therefore as a consequence also preserves the solutions. We propose two new definitions of *solution symmetry* and *constraint symmetry* to capture these two distinct views, and show that they are indeed different: although any constraint symmetry is also a solution symmetry, there can be many

solution symmetries that are not constraint symmetries. We discuss the relationship between the symmetry groups identified by these definitions and show that each is the automorphism group of a hypergraph, derived from either the solutions or the constraints of the CSP. We illustrate these ideas by discussing how they apply to a well-studied example problem, the n -queens problem. We point out that problems with very few solutions must have many solution symmetries, even if they have no constraint symmetries. Finally, we discuss how these definitions of symmetry may be used in practice, by enlarging the symmetry group used in searching for solutions to the CSP.

2 A Brief Survey of Symmetry Definitions

There have been many papers in recent years on symmetry in constraint satisfaction and related problems, not all of which give a clear definition of symmetry. In this section, we review the variety of definitions that have been used.

We first fix our terminology by defining a CSP instance as follows.

Definition 1. A CSP instance is a triple $\langle V, D, C \rangle$ where:

- V is a set of variables;
- D is a universal domain, specifying the possible values for those variables;
- C is a set of constraints. Each constraint $c \in C$ is a pair $c = \langle \sigma, \rho \rangle$ where σ is a list of variables from V , called the constraint scope, and ρ is a $|\sigma|$ -ary relation over D , called the constraint relation.

An assignment of values to variables is a set $\{\langle v_1, a_1 \rangle, \langle v_2, a_2 \rangle, \dots, \langle v_k, a_k \rangle\}$ where $\{v_1, v_2, \dots, v_k\} \subseteq V$ and $a_i \in D, \forall 1 \leq i \leq k$. Note that the constraint relation of a constraint c is intended to specify the assignments that are allowed by that constraint.

A solution to the CSP instance $\langle V, D, C \rangle$ is a mapping from V into D whose restriction to each constraint scope is in the corresponding constraint relation, i.e. is allowed by the constraint.

We will call a CSP k -ary if the maximum arity of any of its constraints is k . A 2-ary CSP will be called a binary CSP.

There are two basic types of definition for symmetry in a CSP instance: those that define symmetry as a property of the set of solutions, and those that define symmetry as a property that can be identified in the statement of the problem, without solving it. We shall refer to these informally in this section as *solution symmetry* and *problem symmetry* or *constraint symmetry*. In Section 3 we will define them formally and use these definitions to show how the two types of symmetry are related.

An example of an early definition of *solution symmetry* is given by Brown, Finkelstein & Purdom [6], who define a symmetry as a permutation of the problem variables that leaves invariant the set of solutions. Backofen and Will [2] similarly define a symmetry as a bijective function on the set of solutions of

a CSP: they allow a symmetry to be specified by its effect on the individual assignments of values to variables.

A number of papers have defined *problem symmetry* in propositional calculus. Aguirre [1] and Crawford, Ginsberg, Luks & Roy [7] each define symmetry similarly: if S is a set of clauses in CNF, then a permutation π of the variables in those clauses is a symmetry of S if $\pi(S) = S$. The expression $\pi(S)$ denotes the result of applying the permutation π to the clauses in S . If this permutation simply re-orders the literals in individual clauses, and reorders the clauses, then it leaves S effectively unchanged, and so in this case $\pi(S) = S$ and π is a symmetry. Benhamou and Sais [4] use a slightly more general definition, in which a symmetry is a permutation defined on the set of *literals* that preserves the set of clauses. For example, given two variables x and y , x may be mapped to $\neg y$.

In CSPs, some authors have similarly defined a symmetry as a mapping that leaves the constraints unchanged, but have often restricted the allowed mappings to those that affect only the variables or only the values. Note that a constraint may be specified *extensionally* by listing its allowed tuples, or *intensionally* by giving an expression such as $x < y$ from which the allowed tuples could be determined. Permuting the variables in a constraint will in general change it: for example, the constraint $x + y = z$ is not the same as the constraint $x + z = y$. Puget [18] defines the notion of a *symmetrical constraint*, that is, a constraint which is unaffected by the order of the variables. For example, the binary inequality constraint, \neq , is symmetrical. He defines a symmetry of a CSP as a permutation of the variables which maps the set of constraints into an equivalent set: any constraint is either unchanged by the permutation or is an instance of a symmetrical constraint and is mapped onto a constraint on the same set of variables.

A similar idea was introduced by Roy and Pachaet [21]. They define the notion of *intensional permutability*: two variables are intensionally permutable if: they have the same domain; any constraint affecting either of them affects both; for any constraint affecting these two variables, interchanging them in the expression defining the constraint does not change it. (The constraint is assumed to be defined intensionally, hence the name.) For example, in a linear constraint, any two variables with the same coefficient are intensionally permutable (with respect to that constraint, and assuming that they have the same domain).

Both Puget [18] and Roy and Pachaet [21] restrict their definitions of symmetries to mappings that permute the *variables* of the problem only. Meseguer and Torras [16] define symmetries that act on both the variables and the values of a CSP. They define a symmetry on a CSP with n variables as a collection of $n + 1$ bijective mappings $\Theta = \{\theta, \theta_1, \dots, \theta_n\}$. The mapping θ is a bijection on the set of variables $\{x_1, x_2, \dots, x_n\}$; each θ_i is a bijection from $D(x_i)$ to $D(\theta(x_i))$ (where $D(x_i)$ is the domain D restricted to the acceptable values for x_i by unary constraints). These mappings will also transform each constraint. The set Θ is called a symmetry if it does not change the set of constraints C , as a whole.

Meseguer and Torras's definition allows both variable symmetries (that permute only the variables) and value symmetries (that permute only the values)

as special cases, and hence is more general than many earlier definitions. However, it does not allow mappings in which variable-value pairs involving the same variable (say $\langle x_i, a_1 \rangle$ and $\langle x_i, a_2 \rangle$) can be mapped to variable-value pairs involving different variables (say $\langle x_j, a_j \rangle$ and $\langle x_k, a_k \rangle$, where $x_j \neq x_k$). For example, Meseguer and Torras consider the n -queens problem, in the commonly-used CSP formulation in which the variables correspond to the rows of the chessboard and the values to the columns. They show that the reflections in the horizontal and vertical axes and the rotation of the chessboard through 180° are symmetries of the corresponding CSP according to their definition, but the other four chessboard symmetries (reflection in the diagonals, rotation through 90° and 270°) are not. This example will be considered in more detail in Section 4 below.

Finally, we consider the notion of *interchangeability*, as defined by Freuder [10]. This is a form of solution symmetry: two values a, b for a variable v are *fully interchangeable* if every solution to the CSP containing the assignment $\langle v, a \rangle$ remains a solution when b is substituted for a , and vice versa. As Freuder notes, in general identifying fully interchangeable values requires finding all solutions to the CSP. He therefore defines local forms of interchangeability that can be identified by inspecting the problem. Neighbourhood interchangeability, for example, is a form of constraint symmetry: two values a, b for a variable v are said to be neighbourhood interchangeable if for every constraint c whose scope includes v , the set of assignments that satisfy c which contain the pair $\langle v, a \rangle$ still satisfy c when this is replaced by $\langle v, b \rangle$, and vice versa.

Benhamou [3] extends the ideas of interchangeability slightly and distinguishes between *semantic* and *syntactic* symmetry in CSPs, corresponding to our notions of solution symmetry and constraint symmetry, respectively. He defines two kinds of semantic symmetry: two values a_i and b_i for a CSP variable v_i are *symmetric for satisfiability* if the following property holds: there is a solution which assigns the value a_i to v_i if and only if there is a solution which assigns the value b_i to v_i . The values are *symmetric for all solutions* if: each solution containing the value a_i can be mapped to a solution containing the value b_i . (The latter property implies the former.) Identifying semantic symmetries requires solving the CSP to find all solutions, and examining them. The notion of *syntactic symmetry* in [3] is defined as follows. Let $P = \langle V, D, C \rangle$ be a binary CSP instance, whose constraint relations are all members of some set R . A permutation π of D is a syntactic symmetry if $\forall r_{ij} \in R$, we have $(d_i, d_j) \in r_{ij} \implies (\pi(d_i), \pi(d_j)) \in r_{ij}$. In other words, the permutation π does not change any constraint relation of P , considered as a set of tuples.

From this brief survey of existing symmetry definitions, it can be seen that they differ both in what aspect of the CSP they act on (only the values, only the variables, or variable-value pairs) and in what they preserve (the constraints or the set of solutions). It should be noted that it has become standard in the symmetry breaking methods that act during search (e.g. [2, 8, 9, 13]), as opposed to adding constraints to the CSP, to describe symmetries by their action on variable-value pairs. Hence, almost all the definitions described in this section are more restrictive than these systems allow.

Under all the definitions, symmetries map solutions to solutions and non-solutions to non-solutions; the definitions disagree over whether this is a defining property, so that any bijective mapping of the right kind that preserves the solutions must be a symmetry, or whether it is simply a consequence of leaving the constraints unchanged. In the next section we will show that this distinction is critical: the choice we make can seriously affect the symmetries that we find.

3 Constraint Symmetries and Solution Symmetries

In this section we will give two definitions of symmetries for constraint satisfaction problems that are sufficiently general to encompass all the types of symmetry allowed by the definitions given in the last section.

Note that the essential feature that allows any bijective mapping on a set of objects to be called a symmetry is that it leaves some property of those objects unchanged. It follows from this that the identity mapping will always be a symmetry, and the inverse of any symmetry will also be a symmetry. Furthermore, given two symmetries we can combine them (by composing the mappings) to obtain another symmetry. Hence, the set of symmetries forms a *group*.

The particular group of symmetries that we obtain depends on exactly what property it is that we choose to be preserved. Our first definition uses the property of being a solution.

Definition 2. *For any CSP instance $P = \langle V, D, C \rangle$, a solution symmetry of P is a permutation of the set $V \times D$ that preserves the set of solutions to P .*

In other words, a solution symmetry is a bijective mapping defined on the set of possible variable-value pairs of a CSP, that maps solutions to solutions. Note that this general definition allows variable and value symmetries as special cases.

To state our definition of *constraint symmetries* we first describe a mathematical structure associated with any CSP instance. For a binary CSP instance, the details of the constraints can be captured in a graph, the *microstructure* [10, 14] of the instance.

Definition 3. *For any binary CSP instance $P = \langle V, D, C \rangle$, the microstructure of P is a graph with set of vertices $V \times D$ where each edge corresponds either to an assignment allowed by a specific constraint, or to an assignment allowed because there is no constraint between the associated variables.*

Hence, the vertices of the microstructure correspond to variable-value pairs of the CSP. For our purposes, it is more convenient to deal with the complement of this graph. The *microstructure complement* has the same set of vertices as the microstructure, but with edges joining all pairs of vertices which are *disallowed* by some constraint, or else are incompatible assignments for the same variable. In other words, two vertices $\langle v_1, a_1 \rangle$ and $\langle v_2, a_2 \rangle$ in the microstructure complement are connected by an edge if and only if:

- the vertices v_1 and v_2 are in the scope of some constraint, but the assignment of a_1 to v_1 and a_2 to v_2 is disallowed by that constraint; or

- $v_1 = v_2$ and $a_1 \neq a_2$.

Recall that any set of vertices of a graph which does not contain an edge is called an *independent set*. An immediate consequence of the definition of the microstructure complement is that a solution to a CSP instance P is precisely an independent set of size $|V|$ in its microstructure complement.

The definition extends naturally to the non-binary case. Here the microstructure complement is a *hypergraph* whose set of vertices is again the set of all variable-value pairs. In this case, a set of vertices E is a hyperedge of the microstructure complement if it represents an assignment *disallowed* by a constraint, or else consists of a pair of incompatible assignments for the same variable. In other words, a set of vertices $\{\langle v_1, a_1 \rangle, \langle v_2, a_2 \rangle, \dots, \langle v_k, a_k \rangle\}$ is a hyperedge if and only if:

- $\{v_1, v_2, \dots, v_k\}$ is the set of variables in the scope of some constraint, but the constraint disallows the assignment $\{\langle v_1, a_1 \rangle, \langle v_2, a_2 \rangle, \dots, \langle v_k, a_k \rangle\}$; or
- $k = 2$, $v_1 = v_2$ and $a_1 \neq a_2$.

Example 1. The following system of linear equations over the integers modulo 2 (that is, where $1 + 1 = 0$)

$$x + y + z = 0$$

$$w + y = 1$$

$$w + z = 0$$

can be modelled as a CSP instance $P = \langle V, D, C \rangle$, with $V = \{w, x, y, z\}$, $D = \{0, 1\}$ and $C = \{c_1, c_2, c_3\}$, where c_1, c_2, c_3 correspond to the three equations.

The microstructure complement of P is shown in Figure 1. It has eight vertices: $\langle w, 0 \rangle, \langle w, 1 \rangle, \langle x, 0 \rangle, \langle x, 1 \rangle, \langle y, 0 \rangle, \langle y, 1 \rangle, \langle z, 0 \rangle, \langle z, 1 \rangle$, and twelve hyperedges. The equation $x + y + z = 0$ disallows the assignment $\{\langle x, 0 \rangle, \langle y, 0 \rangle, \langle z, 1 \rangle\}$ and three other assignments. Hence, the microstructure complement has four ternary hyperedges arising from this constraint, including $\{\langle x, 0 \rangle, \langle y, 0 \rangle, \langle z, 1 \rangle\}$. Each binary constraint also gives two binary hyperedges. In total therefore the microstructure complement has four ternary and four binary hyperedges arising from disallowed tuples. Finally, there are four binary hyperedges (one per variable) corresponding to pairs of different values for the same variable; for example, the hyperedge $\{\langle y, 0 \rangle, \langle y, 1 \rangle\}$.

We are now in a position to define a constraint symmetry. Recall that an *automorphism* of a graph or hypergraph is a bijective mapping of the vertices that preserves the edges (and hence also preserves the non-edges).

Definition 4. For any CSP instance $P = \langle V, D, C \rangle$, a constraint symmetry is an automorphism of the microstructure complement of P (or, equivalently, of the microstructure).




Fig. 1. The microstructure complement of the CSP instance P defined in Example 1.

The microstructure complement is related to the direct encoding of a CSP as a SAT instance [22]. The direct encoding has a variable for each variable-value pair in the original CSP; a clause for each pair of values for each variable, forbidding both values being assigned at the same time; and a clause for each tuple of variable-value pairs not allowed by a constraint (as well as other clauses ensuring that a value is chosen for every variable). A constraint symmetry as defined here is therefore equivalent to a permutation of the variables in the SAT encoding that does not change the set of clauses, and so is related to the definition of symmetry in SAT given by Crawford *et al.* [7].

Example 2. We consider the constraint symmetries of the CSP defined in Example 1, whose microstructure complement is shown in Figure 1. The automorphisms of this graph are:

- $(\langle w, 0 \rangle \langle w, 1 \rangle) (\langle y, 0 \rangle \langle y, 1 \rangle) (\langle z, 0 \rangle \langle z, 1 \rangle)$;
- $(\langle w, 0 \rangle \langle w, 1 \rangle) (\langle y, 0 \rangle \langle z, 1 \rangle) (\langle y, 1 \rangle \langle z, 0 \rangle)$;
- $(\langle w, 0 \rangle \langle w, 1 \rangle) (\langle y, 0 \rangle \langle z, 0 \rangle) (\langle y, 1 \rangle \langle z, 1 \rangle)$;

and the identity mapping. (The permutations of the vertices are written in cyclic form: the first swaps the vertices $\langle w, 0 \rangle$ and $\langle w, 1 \rangle$ while simultaneously swapping $\langle y, 0 \rangle$ and $\langle y, 1 \rangle$ and swapping $\langle z, 0 \rangle$ and $\langle z, 1 \rangle$, but leaves $\langle x, 0 \rangle$ and $\langle x, 1 \rangle$ unchanged.) Hence, these four mappings are the constraint symmetry group of the CSP.

This example also shows that there can be more solution symmetries than constraint symmetries. The CSP has only two solutions, $\{\langle w, 0 \rangle, \langle x, 1 \rangle, \langle y, 1 \rangle, \langle z, 0 \rangle\}$ and $\{\langle w, 1 \rangle, \langle x, 1 \rangle, \langle y, 0 \rangle, \langle z, 1 \rangle\}$. The permutation $(\langle w, 0 \rangle \langle z, 0 \rangle \langle y, 1 \rangle)$, which maps $\langle w, 0 \rangle$ to $\langle z, 0 \rangle$, $\langle z, 0 \rangle$ to $\langle y, 1 \rangle$, $\langle y, 1 \rangle$ to $\langle w, 0 \rangle$ and leaves all other variable-value pairs unchanged, is a solution symmetry. This mapping transforms each solution into itself, but clearly is not a constraint symmetry.

Although Definition 2 and Definition 4 appear to be very different, we now show that there are some simple relationships between solution symmetries and constraint symmetries.

Theorem 1. *The group of constraint symmetries of a CSP instance P is a subgroup of the group of solution symmetries of P .*

Proof. Let P be a CSP instance and let π be any automorphism of the microstructure complement of P . We will show that π maps solutions to solutions, and hence is a solution symmetry of P .

Let s be any solution of P , and let W be the corresponding set of vertices in the microstructure complement of P . By the construction of the microstructure complement, W is an independent set of size $|V|$. Since π is an automorphism, we know that $\pi(W)$ is also an independent set of size $|V|$. Hence $\pi(W)$ is not disallowed by any of the constraints of P , and so is a solution. \square

Next we show that the group of all solution symmetries of an instance P is also the automorphism group of a certain hypergraph. We first define a *nogood*.

Definition 5. *For any CSP instance P , a k -ary nogood is an assignment to k variables of P that cannot be extended to a solution of P .*

The k -nogood hypergraph of P is a hypergraph whose set of vertices is $V \times D$ and whose set of edges is the set of all m -ary nogoods for all $m \leq k$.

The k -nogood hypergraph of a CSP instance has the same vertices as the microstructure complement. For a k -ary CSP (one whose constraints have maximum arity k), the k -ary nogood hypergraph contains every hyperedge of the microstructure complement, and possibly some others. The additional hyperedges represent partial assignments of up to k variables that are allowed by the constraints, but do not appear in any solution because they cannot be extended to a full assignment satisfying all the constraints. Note that although enforcing $(k+1)$ -consistency in a CSP instance may add new k -ary nogoods, it will not in general find all such nogoods; except in special cases, we identify the k -ary nogoods from the set of solutions.

Example 3. Consider again the CSP instance P defined in Example 1, with solutions $\{\langle w, 0 \rangle, \langle x, 1 \rangle, \langle y, 1 \rangle, \langle z, 0 \rangle\}$ and $\{\langle w, 1 \rangle, \langle x, 1 \rangle, \langle y, 0 \rangle, \langle z, 1 \rangle\}$.

This instance has a large number of 3-ary nogoods, and the 3-nogood hypergraph of P has a large number of hyperedges, in addition to those in the microstructure complement. These include the hyperedge $\{\langle x, 0 \rangle, \langle y, 0 \rangle, \langle z, 0 \rangle\}$, for example. This assignment is allowed by the 3-ary constraint on the variables x, y, z , but cannot be extended to a complete solution of P . Many of the additional hyperedges do not correspond to the scope of any constraint: for example, the hyperedge $\{\langle w, 0 \rangle, \langle x, 1 \rangle, \langle y, 0 \rangle\}$.

Theorem 2. *For any k -ary CSP instance P , the group of all solution symmetries of P is equal to the automorphism group of the k -nogood hypergraph of P .*

Proof. Let F be the k -nogood hypergraph of P and let π be any automorphism of F . We will show that π preserves solutions, and hence is a solution symmetry.

Let s be any solution of P , and let W be the corresponding set of vertices in F . By the construction of this hypergraph, W is an independent set of size $|V|$.

Since π is an automorphism of F , we know that $\pi(W)$ is also an independent set of size $|V|$. Hence $\pi(W)$ is not disallowed by any of the constraints of P , and is a solution.

Conversely, let π be a solution symmetry of P . We will show that π maps every set of k or fewer vertices of F which is not a hyperedge to another non-hyperedge, and hence π is an automorphism of this hypergraph.

Let E be any set of k or fewer vertices in F which is not a hyperedge. Since every nogood of P of size k or less is a hyperedge of the k -nogood hypergraph, it follows that E can be extended to at least one solution of P .

Hence we may suppose that E is part of some solution s . Now, s is mapped to the solution $\pi(s)$ by the solution symmetry π . Every k -ary projection of this solution, including the image $\pi(E)$ of E , is a non-hyperedge in F , and so we are done. \square

Theorem 2 shows that to obtain the solution symmetries of a CSP instance it is sufficient to consider the automorphisms of the hypergraph obtained by adding all the nogoods of arity k or less to the microstructure complement. We will show in the next section that in some cases there are hypergraphs obtained by adding a smaller number of edges to the microstructure complement which already have all solution symmetries as automorphisms. However, the next result shows that there are cases where it is in fact necessary to add all nogoods of arity k or less to the microstructure complement in order to obtain a hypergraph with all solution symmetries as automorphisms.

Proposition 1. *For some k -ary CSP instances P , the k -nogood hypergraph is the only hypergraph containing the microstructure complement of P whose automorphisms are exactly the solution symmetries.*

Proof. Consider a CSP instance P , with constraints of every arity less than or equal to k , which has no solutions. Let H be the microstructure complement of P .

Since P has no solutions, every permutation of the vertices of H is a solution symmetry. For each positive integer $m \leq k$, there is at least one m -tuple disallowed by a constraint, so H has at least one m -ary hyperedge. Since every permutation of the vertices of H is a solution symmetry, applying the solution symmetry group to H will give all m -sets of vertices as hyperedges. Hence the only hypergraph containing H whose automorphisms include all solution symmetries is the hypergraph with all m -sets of vertices as edges, for all $m \leq k$, which is equal to the k -nogood hypergraph. \square

Hence to obtain all solution symmetries of a k -ary CSP instance it is sometimes necessary to consider all m -ary nogoods, for all $m \leq k$. On the other hand, Theorem 2 shows that we do not need to consider nogoods of any size *larger* than k . (In fact the same proof shows that adding all nogoods of size l , for any l larger than k , to the k -nogood hypergraph does not change its automorphism group.) In particular, this means that to obtain all the solution symmetries of a binary CSP instance we need only consider the binary and unary nogoods.

Theorem 1 and Theorem 2 help to clarify the relationship between solution symmetries and constraint symmetries. One reason that it is important to distinguish these two kinds of symmetries carefully is that, in general, there can be many more solution symmetries than constraint symmetries for a given CSP instance, as we will show in the next section.

4 Case Study: Symmetry in n -queens

In this section we will illustrate the relationship between solution symmetries and constraint symmetries by examining the n -queens problem. This problem is useful for discussing symmetry because the common CSP formulation has several different types of symmetry, some of which are beyond the scope of some earlier definitions.

The standard formulation of the n -queens problem as a CSP has n variables corresponding to the rows of the chessboard, say r_1, r_2, \dots, r_n . The domain of values corresponds to the columns of the chessboard, say $D = \{1, 2, \dots, n\}$. The constraints can be expressed as follows, to give a binary CSP:

- for all $i, j, 1 \leq i < j \leq n, r_i \neq r_j$;
- for all $i, j, 1 \leq i < j \leq n, |r_i - r_j| \neq |i - j|$.

Considered as a geometric object, a chessboard has eight symmetries: reflections in the horizontal and vertical axes and the two diagonals, rotations through 90° , 180° and 270° , and the identity.

Recall, however, that Meseguer & Torras [16] did not allow the full set of geometric symmetries of the chessboard as symmetries of the usual CSP formulation of n -queens. This is because the formulation introduces an asymmetry between rows and columns, so that some of the geometric symmetries do not leave the constraints syntactically unchanged. In particular, the rotational symmetries through 90° and 270° map assignments forbidden by some of the constraints to assignments that are mutually incompatible because they assign two values to the same variable. For example, the forbidden pair consisting of $\langle r_1, 1 \rangle$ and $\langle r_2, 1 \rangle$ is mapped by the rotation through 90° to the incompatible pair consisting of $\langle r_1, n \rangle$ and $\langle r_1, n - 1 \rangle$.

The microstructure complement restores the symmetry between rows and columns, by treating in the same way both of these reasons for a pair of assignments to be disallowed. Hence each geometric symmetry of the chessboard gives rise to a constraint symmetry of the n -queens problem for any n , according to our definition of constraint symmetry (Definition 4).

Clearly, the set of *solutions* to an instance of the n -queens problem is invariant under each of the eight geometric symmetries of the chessboard. Hence each of these geometric symmetries is a solution symmetry of the n -queens problem for any n , according to our definition of solution symmetry (Definition 2). However, there can be many other solution symmetries for instances of this problem, as we will now show.

The 3-queens problem has no solutions; like any other CSP with no solution, any permutation of the possible variable-value pairs is a solution symmetry. This is confirmed by Theorem 2: the binary nogood hypergraph is the complete graph with nine vertices, and any permutation of the vertices is an automorphism.

The 4-queens problem has two solutions, shown in Figure 2. In this case, it is easier to consider the *complement* of the binary nogood hypergraph, in which each edge represents a pair of variable-value assignments that is allowed by the solutions. Figure 2 also shows this graph, drawn so that the position of each vertex, representing a variable-value pair, and hence a square of the chessboard, corresponds to the position on a chessboard of that square. Each solution is represented as a 4-clique in this graph, rather than as an independent set of size 4 in the binary nogood hypergraph. The automorphisms of this graph are that: the vertices within either clique can be permuted; the vertices in one clique can be swapped with those in the other; and the eight isolated vertices (representing unary nogoods) can be permuted; and we can also compose these permutations. This gives a total of $4! \times 4! \times 2 \times 8!$ automorphisms, or 46,448,640. Since the automorphisms of a graph are the same as the automorphisms of its complement, these 46,448,640 automorphisms are the solution symmetries of 4-queens, by Theorem 2.




Fig. 2. The solutions of the 4-queens problem (left) and the complement of the binary nogood graph (right).

The 5-queens problem has ten solutions, shown in Figure 3. These solutions are divided into two equivalence classes by the geometric symmetries of the chessboard; they transform any solution into another solution from the same equivalence class. (The first eight solutions shown form one class, the last two the second class.)

Every square of the chessboard has a queen on it in at least one of these ten solutions, so that every value in the original domains is part of at least one solution and there are no unary nogoods. However, there are some *pairs* of squares, where two queens can be placed consistently with the original constraints, but which are not allowed in any solution.

Some of these new binary nogoods are shown in Figure 4; the others are their symmetric equivalents under the geometric symmetries of the chessboard. In this case, the 40 additional binary nogoods can be derived by making the problem path consistent, without finding all solutions to the problem. A binary CSP instance is path consistent if for every possible assignment $\{\langle v_i, a_i \rangle, \langle v_j, a_j \rangle\}$




Fig. 3. The ten solutions of the 5-queens problem.

satisfying the constraint between v_i and v_j and for every third variable v_k , there is a value a_k such that the assignments $\{\langle v_i, a_i \rangle, \langle v_k, a_k \rangle\}$ and $\{\langle v_k, a_k \rangle, \langle v_j, a_j \rangle\}$ satisfy the constraints between v_i and v_k and between v_k and v_j , respectively. For example in the 5-queens problem, a pair of queens in row 1, column 1 and row 3, column 4 together attack every square in row 2, and so this pair of assignments cannot be part of any solution.

If these additional binary nogoods are added to the microstructure complement, then we obtain the 2-ary nogood graph for the 5-queens problem. By using the software tool NAUTY [15], we can find the automorphism group of this graph: it has 28,800 elements, so the 5-queens problem has a total of 28,800




Fig. 4. New binary nogoods derived from the solutions of the 5-queens problem.

solution symmetries, by Theorem 2. (We have also used NAUTY to confirm that the microstructure complement has just the eight geometric symmetries.)

In this case, unlike the 4-queens problem, it is difficult to develop an intuitive understanding of the additional solution symmetries. Some are easy to see: for instance, the rows of the board can be cyclically permuted (row 1 \rightarrow row 2 \rightarrow row 3 \rightarrow row 4 \rightarrow row 5 \rightarrow row 1), as can the columns. The subgroup consisting of these permutations together with the geometric symmetries and all combinations makes all 10 solutions symmetrically equivalent, i.e. there is just one equivalence class, rather than two. However, this subgroup is still much smaller than the full solution symmetry group.⁵

In the 6-queens problem, the additional binary nogoods can again be derived by path consistency. For example, queens in row 1, column 2 and row 3, column 5 together attack all the squares in row 2. There are also unary nogoods, since the problem has only four solutions. Adding all these nogoods to the microstructure complement yields a very large symmetry group, as shown in Table 1.

For $n \geq 7$, path consistency does not give any new nogoods, because two queens can together attack at most six squares on any row. Even so, there can be binary nogoods in addition to those in the original constraints for $n = 7, 8, 9$; if the constraints allow queens to be placed on a pair of squares, but no solution has both squares occupied, we have a new binary nogood. For example, in the 7-queens problem there are no solutions with queens in row 1, column 1 and in row 2, column 7. Again, in the 8-queens problem, there are many binary nogoods to add; for example, in the set of 12 non-equivalent solutions (with respect to the geometric symmetries), only two have a queen in a corner square. Since a queen in a corner only directly attacks two squares in any other row, there are many squares where a queen does not directly conflict with a queen in a corner square, but where the combination with a queen in a corner square does not appear in any solution. For instance, there are no solutions (including symmetric equivalents) with a queen in row 1, column 1, and a queen in row 2, columns 2, 3 or 8.

Thereafter, we have examined the sets of solutions up to $n = 14$ and found no further additional binary nogoods. Table 1 shows that for $n \geq 7$, the solution symmetries appear to be just the geometric symmetries: in spite of the additional binary nogoods for $n = 7, 8, 9$, the binary nogood graph for these instances has the same automorphism group as the microstructure complement. This demonstrates that in spite of Proposition 1, the minimal hypergraph containing the microstructure complement whose automorphisms are the solution symmetries can sometimes be smaller than the k -ary nogood hypergraph. Note that the number of solutions to the n -queens problem increases rapidly with n : intuitively,

⁵ In an attempt to understand the solution symmetries of 5-queens in terms of simple transformations of the chessboard, we used NAUTY and GAP [11] to find a small number of generators of the group, including one or more of the geometric symmetries. The group can be generated by two permutations of the variable-value pairs, together with the rotation through 90° . However, these two permutations are still far from meaningful in terms of the chessboard.

Table 1. The number of additional binary nogoods derived from the sets of solutions to the n -queens problem, and the number of solution symmetries

n	Additional binary nogoods	Solution symmetries
3	8	$9! = 362,880$
4	32	$4! \times 4! \times 2 \times 8! = 46,448,640$
5	40	28,800
6	280	3,089,428,805,320,704,000,000
7	72	8
8	236	8
9	40	8
10	0	8

it becomes more difficult for a solution symmetry to preserve them all, so that eventually the solution symmetries are just the constraint symmetries.

A principal reason for identifying symmetry in CSPs is to reduce search effort by not exploring assignments that are symmetrically equivalent to assignments considered elsewhere in the search. Clearly, if the solution symmetry group is larger than the constraint symmetry group, there will potentially be a greater search reduction from using the solution symmetries, if they can somehow be identified in advance. In some cases, as in the 5-queens problem, establishing some level of consistency in the problem to find new nogoods (of arity $\leq k$), and adding these to the microstructure complement, will give an automorphism group that is nearer to the solution symmetry group, if not equal to it.

When finding all solutions, the aim in symmetry breaking is to find just one solution from each symmetry equivalence class; in the 5-queens problem, the solutions fall into two equivalence classes when using the constraint symmetries and only one when using the solution symmetries. Once the search is completed, it is possible to reconstruct the full set of ten solutions, if required, from either the single solution found using the solution symmetries, or the two solutions found using the constraint symmetries.

However, there seems no straightforward way of going from the single representative solution found in the first case to two representatives that are distinct under the constraint symmetries, except by generating the full set of solutions and eliminating those that are equivalent under the constraint symmetries. Although feasible in this case, when there are only ten solutions altogether, it would not often be practicable to generate all possible solutions. Hence, if a set of non-isomorphic solutions were required, under some group of symmetries, it seems that only those symmetries should be used in the search. The appropriate symmetry group should be chosen in advance, as the choice can affect the number of solutions found.

This raises the question of how to identify the symmetries of a CSP, whether the constraint symmetries or the solution symmetries; we discuss this next.

5 Identifying Symmetry in Practice

Symmetry in CSPs is usually identified, in practice, by applying human insight: the programmer sees that some transformation would transform a hypothetical solution into another hypothetical solution. The definition of constraint symmetry given earlier can be used to confirm that candidate transformations are genuine symmetries. It is not necessary to generate the entire microstructure complement for this purpose, but only to demonstrate that each candidate mapping will map edges to edges and non-edges to non-edges in this hypergraph.

For example, in the n -queens problem, each vertex in the microstructure corresponds to a queen placed on a specific square of the chessboard, i.e. the variable-value pair $\langle r_i, j \rangle$ represents a queen placed in row i , column j . Rotation through 90° maps this assignment to $\langle r_j, n+1-i \rangle$. An edge in the microstructure complement corresponds to a pair of squares that are on the same row, column or diagonal; the rotation clearly maps rows to columns and vice versa, and also maps diagonals to diagonals, and hence maps edges to edges in the microstructure complement. For example, $\langle r_i, j \rangle$ and $\langle r_i, k \rangle$, which are in the same row, are mapped to $\langle r_j, n+1-i \rangle$ and $\langle r_k, n+1-i \rangle$, which are in the same column. Similarly, it can be shown that this rotation maps variable-value pairs that are in different rows, columns, and diagonals to pairs that remain in different rows, columns, and diagonals, and hence it maps non-edges to non-edges.

Identifying symmetry in a CSP by inspection is prone to missing some of the symmetry: a feasible method to identify symmetry automatically would be preferable. We can in theory identify all the constraint symmetries in a problem by generating the microstructure complement and finding its automorphism group. However, this may not be feasible, especially for CSPs with global constraints, where identifying the satisfying tuples may itself be intractable [5], and in any case each constraint may result in a very large number of hyperedges. It may then be possible to represent the constraints more compactly while preserving the important details; for instance, Ramani and Markov [20] propose to represent constraints by parse trees and find the automorphisms of the resulting graph. Puget [19] also shows how to represent a CSP instance compactly in a graph whose automorphisms are the constraint symmetries.

One advantage of considering the solution symmetry group is that it is well-defined for any problem instance with a given set of solutions. On the other hand, different CSP models of the same problem, with the same set of solutions but differing slightly in the way that the constraints are expressed, may have different constraint symmetry groups. The 5-queens example shows that if the problem is made path consistent, and hence new nogoods are added, the constraint symmetry group is the full solution symmetry group with 28,800 elements rather than the eight constraint symmetries of the original formulation. Conversely, it would often be possible, either deliberately or inadvertently, to write the constraints of a CSP in such a way that the symmetry of the problem being modelled is not reflected in the microstructure of the CSP.

Many authors have defined symmetry in CSPs in a similar way to our definition of solution symmetry, perhaps for this reason, but have effectively identified

constraint symmetries; we have shown that the solution symmetry group can be much larger than the constraint symmetry group. It is inescapably hard to determine the solution symmetry group: since it is the complete permutation group on the variable-value pairs if and only if the CSP has no solution, finding the solution symmetry group is equivalent to determining whether or not the CSP has a solution. However, since we have shown that the solution symmetries are the automorphisms of the k -ary nogood hypergraph, any additional nogoods that can be found, for instance by preprocessing the problem, can be added to the microstructure complement and may allow a larger group of symmetries to be found.

Methods such as those proposed here may find a potentially very large group of symmetries, but with possibly only a small number of generators. For instance, as shown earlier, the solution symmetry group of 5-queens has 28,800 elements but just three generators. Symmetry-breaking methods that combine dynamic symmetry breaking during search with computational group theory, e.g. [12], can tackle such symmetry groups.

6 Solution Symmetry without Constraint Symmetry

As already mentioned, if a CSP instance has no solutions, any permutation of the variable-value pairs is a solution symmetry, irrespective of whether the problem has any constraint symmetry. We can take this further, since it is easy to see that any CSP instance with very few solutions must have many solution symmetries.

Example 4. Suppose a CSP instance with n variables, and d values for each variable, has only one solution. Any permutation of the variable-value pairs in that solution is a solution symmetry; so is any permutation of the variable-value pairs that are not in the solution. The size of the solution symmetry group is therefore $n! \times (n(d-1))!$.

Example 5. Suppose a CSP instance with n variables, and d values for each variable, has just 2 solutions, and these solutions have k assignments in common. A solution symmetry must transform the common assignments to themselves, but then can map the other assignments in each solution either to assignments in the same solution, or to assignments in the other solution. Again, the $nd - 2n + k$ assignments that are not part of either solution can be mapped to themselves. This gives $(nd - 2n + k)! \times k! \times 2 \times (n - k)!^2$. (For example, the 4-queens problem has two solutions; $n = d = 4$, $k = 0$ and the number of solution symmetries is $8! \times 2 \times 4!^2$, as shown in Table 1.)

With larger numbers of solutions, it becomes more difficult to construct permutations of the variable-value pairs (other than the identity mapping) that can map each solution into itself without changing the other solutions, and still more difficult to map one solution into another, unless there is some inherent symmetry that leads to symmetry in the way that the solutions overlap.

Example 6. Suppose a CSP instance has the following three solutions:

$$\begin{aligned} A &: \{\langle x, a \rangle, \langle y, b \rangle, \langle z, c \rangle, \langle w, d \rangle\} \\ B &: \{\langle x, a \rangle, \langle y, e \rangle, \langle z, f \rangle, \langle w, g \rangle\} \\ C &: \{\langle x, h \rangle, \langle y, b \rangle, \langle z, c \rangle, \langle w, i \rangle\} \end{aligned}$$

it is not possible to map any solution to another while preserving the third, because of the different numbers of variable-value pairs in common. On the other hand, if solution B is instead:

$$B' : \{\langle x, a \rangle, \langle y, e \rangle, \langle z, f \rangle, \langle w, d \rangle\}$$

then the permutation $(\langle x, a \rangle \langle y, b \rangle)(\langle w, d \rangle \langle z, c \rangle)(\langle x, h \rangle \langle y, e \rangle)(\langle z, f \rangle \langle w, i \rangle)$ maps solution B' to solution C and solution A to itself, and there are many other solution symmetries too.

Hence, as the number of solutions gets larger, the existence of solution symmetries depends on the structure of the solutions, and hence on the properties of the CSP. Nevertheless, any CSP instance that has d values for each variable and fewer than d solutions must have some solution symmetries arising from permuting the assignments that do not appear in any solution.

These considerations show that to some extent the existence of solution symmetry reflects how many solutions the CSP instance has, rather than any other property. This can lead to paradoxical results, as we will show. One way of dealing with (constraint) symmetry that has been identified in a CSP instance is to add constraints to eliminate it; the resulting new CSP instance has fewer solutions than the original. Even though the new instance has fewer constraint symmetries, it may have more solution symmetries.

Example 7. Consider the problem of finding a graceful labelling of a graph [17]. A labelling f of the nodes of a graph with q edges is *graceful* if f assigns each node a unique label from $\{0, 1, \dots, q\}$ and when each edge xy is labelled with $|f(x) - f(y)|$, the edge labels are all different (and form a permutation of $\{1, 2, \dots, q\}$).

A simple CSP model has a variable for each node, and its values are the potential labels for that node. Suppose that in the CSP instance, P , constructed to find a graceful labelling of the graph in Figure 5, the nodes in one clique are represented by variables x_1, \dots, x_5 and those in the other by x_6, \dots, x_{10} , in such a way that there is an edge in the graph joining x_1 and x_6 , x_2 and x_7 , and so on. The solution shown in Figure 5 is therefore $\{\langle x_1, 0 \rangle, \langle x_2, 4 \rangle, \langle x_3, 18 \rangle, \langle x_4, 19 \rangle, \langle x_5, 25 \rangle, \langle x_6, 23 \rangle, \langle x_7, 14 \rangle, \langle x_8, 6 \rangle, \langle x_9, 3 \rangle, \langle x_{10}, 1 \rangle\}$.

The constraint symmetries of P are, first, those corresponding to the symmetry of the graph: any permutation of the variables in each 5-clique which acts




Fig. 5. A graceful labelling of $K_5 \times P_2$.

on both sets of variables in the same way leaves P unchanged; also the variables in the first clique can be interchanged with the adjacent variables in the second. In addition, there is a value symmetry: replacing every value v in every domain by its complement $q - v$ leaves P unchanged. The graph symmetries can be combined with each other and with the complement symmetry. Hence, the size of the symmetry group is $5! \times 2 \times 2$, or 480.

In this case, it is possible to add symmetry-breaking constraints to eliminate all the symmetry. A possible set of constraints to eliminate the graph symmetry is $x_1 < x_2 < x_3 < x_4 < x_5$ and $x_1 < x_6$. Eliminating the complement symmetry is more difficult, but there must be an edge labelled $q - 1$ in any graceful labelling, and the edge must join either the nodes labelled 0 and $q - 1$ or those labelled 1 and q ; these two possibilities are equivalent under the complement symmetry. Suppose we add a constraint eliminating the first possibility. The resulting CSP instance P' has no constraint symmetry.

The graph $K_5 \times P_2$ has a unique graceful labelling, apart from symmetric equivalents. Hence, the original CSP instance P has 480 solutions. P has more solution symmetries than the 480 constraint symmetries, since the labelling shown in Figure 5 and its complement use only 16 different node labels; the assignments of the remaining 10 values to each variable can be permuted without changing the solutions. However, P' has only one solution and its solution symmetry group therefore has $10! \times 250!$ elements, as shown in Example 4. We again used NAUTY to find the solution symmetry group of P ; as well as permutations of the 10 unused values for each of the 10 variables, and the original constraint symmetries, there are some other swaps of variable-value pairs that leave the solutions unchanged, giving a group with $100! \times 2^{25} \times 480$ elements. This is much smaller than the solution symmetry group of P' , but is clearly not a subgroup of it: the constraint symmetries of P are also solution symmetries of P , but not of P' .

We can see a similar effect in the n -queens problem, where eliminating the symmetry by adding symmetry-breaking constraints also results in a new CSP with more solution symmetries:

Example 8. The constraint symmetry of the n -queens problem can be eliminated by adding ordering constraints (of arity up to n) in a manner similar to that

given by Crawford *et al.* [7]. The solution symmetry group can be found for the resulting CSP by finding all the solutions and finding the automorphisms of the graph that has an n -ary hyperedge for every solution. Table 2 shows the size of the solution symmetry group.

Table 2. The size of the solution symmetry group and number of solutions for the n -queens problem with constraints to eliminate the constraint symmetry.

n	Size of solution symmetry group	Number of solutions
3	$9!$	0
4	$12! \times 4!$	1
5	$16! \times 4! \times 4! \times 2$	2
6	$30! \times 6!$	1
7	$24! \times 13824$	6
8	$24! \times 12$	12
9	$11! \times 4$	46
10	$10!$	92
11	$10!$	341
12	$11!$	1787

These results can be compared with the size of the solution symmetry group for the original problem given in Table 1. The solution symmetry group is the same as that of the original problem for $n = 3$: both problems have no solution. Otherwise, the problem with the symmetry-breaking constraints has a larger symmetry group than the original; this is particularly striking for $n \geq 7$, where the solution symmetry group of the original problem has just the 8 constraint symmetries, as shown in Table 1. The symmetry-breaking constraints result in a larger group because they prevent queens from being placed on certain squares; for instance, a queen can be placed in at most one of the corner squares, and the queen in the top row can only be placed in the left half of the board. This gives rise to several variable-value pairs (the number increasing with n) that cannot appear in any solution, any permutation of which is a solution symmetry.

Hence, these examples show that eliminating the constraint symmetry in a CSP instance, by adding symmetry-breaking constraints to give a new instance, can result in an *increase* in the number of solution symmetries. In Example 7, this is simply because the new instance has only one solution, and such problems have many solution symmetries. In Example 8, the symmetry-breaking constraints ensure that some variable-value pairs cannot be used in any solution, and so give rise to solution symmetries.

7 Incremental Symmetry Breaking

In trying to find a set of non-isomorphic solutions to a CSP, the number of solutions found may depend on the symmetry group, and hence, the symmetry group should be decided in advance. However, if we are only trying to find a single solution, and are making use of symmetry to eliminate wasted search, the results of this paper suggest a novel, *incremental* approach to using symmetry during search, which we now describe.

In this approach we maintain a set of currently known symmetries throughout the solution process. This set is initialised to the group of constraint symmetries by constructing the microstructure complement, and finding its automorphism group. The microstructure complement serves as a first approximation to the k -nogood hypergraph. Each time a nogood of arity k or less is found during the solution process it is added to our current view of the k -nogood hypergraph, together with all of its images under currently known symmetries. Adding these edges might increase the number of automorphisms of this graph, and hence increase the set of currently known symmetries. The bigger this group of symmetries gets the more information we get from each additional nogood.

We know by Theorem 2 that when we have found all k -ary nogoods then we have the whole group of solution symmetries. However, we may get the whole group of solution symmetries long before we have explicitly considered all of the nogoods, as illustrated in the following example.

Example 9. Consider the binary CSP P which has three variables $V = \{x, y, z\}$, domain $D = \{1, 2\}$ and three constraints $\{x \neq y, y \neq z, x \neq z\}$.

The microstructure complement of P has six vertices $\langle x, 1 \rangle, \langle y, 1 \rangle, \langle z, 1 \rangle, \langle x, 2 \rangle, \langle y, 2 \rangle, \langle z, 2 \rangle$ and nine edges.




Fig. 6. The microstructure complement of the CSP instance P defined in Example 9.

As shown in Figure 6, this graph has two disjoint 3-cycles, consisting of $\langle x, 1 \rangle, \langle y, 1 \rangle, \langle z, 1 \rangle$ and $\langle x, 2 \rangle, \langle y, 2 \rangle, \langle z, 2 \rangle$. These 3-cycles must be preserved by any automorphism, so it follows that the automorphism group is generated by three permutations:

- $(\langle x, 1 \rangle \langle y, 1 \rangle)(\langle x, 2 \rangle \langle y, 2 \rangle)$,
- $(\langle x, 1 \rangle \langle y, 1 \rangle \langle z, 1 \rangle)(\langle x, 2 \rangle \langle y, 2 \rangle \langle z, 2 \rangle)$, and
- $(\langle x, 1 \rangle \langle x, 2 \rangle)(\langle y, 1 \rangle \langle y, 2 \rangle)(\langle z, 1 \rangle \langle z, 2 \rangle)$.

The full automorphism group of this graph has 12 elements (it is the group $S_3 \times S_2$). Hence, P has 12 constraint symmetries.

On the other hand, as the instance has no solutions, the solution symmetry group is the full permutation group of the six vertices, with $6! = 720$ elements.

By one application of path consistency to the constraints $x \neq y$ and $y \neq z$, we immediately obtain the binary nogood $x = 1, z = 2$. If we add this nogood as an additional edge to the microstructure complement of P , together with the five images of it obtained by applying the constraint symmetries described above, we obtain a graph with 15 edges in total, representing the 15 known binary nogoods. Since this graph is the complete graph on six vertices, we have established that every permutation is a solution symmetry and also that the problem has no solutions.

As already outlined, an obvious way to use new nogoods of arity up to k to increase the size of the symmetry group is to add each nogood, with its images under the current symmetry group, to the microstructure complement and to find the new automorphism group. Adding the symmetric equivalents of each nogood ensures that this process will never lead to a decrease in the size of the symmetry group. Unfortunately, however, the order in which nogoods are added to the graph can affect the size of the resulting symmetry group, as the following example shows:

Example 10. Consider a CSP instance with five variables $V = \{a, b, c, d, e\}$, each with domain $D = \{0, 1\}$, and the constraints $\{a = b, b = c, c = d, d = e\}$. The microstructure complement is shown in Figure 7. The constraint symmetry group has four elements: apart from the identity mapping, the values 0 and 1 can be swapped for every variable; a can be swapped with e and simultaneously b swapped with d ; and the last two symmetries can be combined.⁶

Now suppose that the nogood $a = 0 \ \& \ e = 1$ is discovered. Figure 8 shows the microstructure complement after the edges $\{a, 0\} \langle e, 1 \rangle$ and $\{a, 1\} \langle e, 0 \rangle$ corresponding to this nogood and its symmetric equivalent have been added. This graph has a larger symmetry group; the variables a, b, c, d, e can now be cyclically permuted.

If the nogood $b = 0 \ \& \ e = 1$ is found next and added to the microstructure complement, with its images under the new symmetry group, the resulting graph is shown in Figure 9. The symmetry group is again enlarged, and is now the full solution symmetry group, since the complement of this graph consists of two

⁶ Note that in finding the constraint symmetries of a CSP instance, Puget [19] proposes to take the transitive closure of equality constraints, to ensure that the full symmetry between the variables is represented in the constraints. In this example, the microstructure complement would then be as shown below in Figure 9. However, the example serves to illustrate the general point.




Fig. 7. The microstructure complement of a CSP instance with $V = \{a, b, c, d, e\}$, $D = \{0, 1\}$ and constraints $\{a = b, b = c, c = d, d = e\}$.




Fig. 8. The microstructure complement after the nogood $a = 0 \ \& \ e = 1$ is found.

cliques with vertex sets $\{\langle a, 0 \rangle, \langle b, 0 \rangle, \langle c, 0 \rangle, \langle d, 0 \rangle, \langle e, 0 \rangle\}$ and $\{\langle a, 1 \rangle, \langle b, 1 \rangle, \langle c, 1 \rangle, \langle d, 1 \rangle, \langle e, 1 \rangle\}$ corresponding to the two solutions $a = b = c = d = e = 0$ and $a = b = c = d = e = 1$.

Suppose instead that the first nogood found was $b = 0 \ \& \ e = 1$; the edges $\{\langle b, 0 \rangle \langle e, 1 \rangle\}$, $\{\langle b, 1 \rangle \langle e, 0 \rangle\}$, $\{\langle d, 0 \rangle \langle a, 1 \rangle\}$ and $\{\langle d, 1 \rangle \langle a, 0 \rangle\}$ would be added to the microstructure complement, as shown in Figure 10 (left). The new graph has the same automorphism group as the original. Now, if the nogood $a = 0 \ \& \ e = 1$ is discovered, giving the microstructure complement shown on the right of Figure 10, the symmetry group is again unchanged. So by adding the same nogoods in the reverse order, a smaller symmetry group has been generated.

In order to get the graph shown in Figure 9, and hence the solution symmetry group, further nogoods would need to be discovered, e.g. $a = 0 \ \& \ c = 1$. Note that adding both nogoods $b = 0 \ \& \ e = 1$ and $a = 0 \ \& \ e = 1$ to the microstructure complement at the same time, rather than successively, along with their images in the original constraint symmetry group, would also result in the microstructure complement shown on the right of Figure 10, and no change to the symmetry group.




Fig. 9. The microstructure complement after the nogoods $a = 0$ & $e = 1$ and $b = 0$ & $e = 1$ are found, in that order.




Fig. 10. The microstructure complement (left) after the nogood $b = 0$ & $e = 1$ is found and (right) after the nogoods $b = 0$ & $e = 1$ and $a = 0$ & $e = 1$ are found, in that order.

This example shows that to get the maximum benefit from finding new nogoods and using them to enlarge the symmetry group, the nogoods must be added to the microstructure complement in the right order; finding a good order efficiently seems difficult. Hence, although enlarging the symmetry group either before or during search could potentially reduce the search effort required to solve the CSP, there are some difficulties to overcome. Future work will address these.

8 Conclusion

We have reviewed definitions of symmetry in CSPs and have proposed definitions of *constraint symmetry* and *solution symmetry* to encompass two types of definition that have been used. We have shown that there can be many more solution symmetries, i.e. permutations of the variable-value pairs that preserve the solutions, than constraint symmetries, i.e. permutations that preserve the constraints. In practice, researchers have identified constraint symmetries in CSPs

rather than solution symmetries, regardless of their definition of symmetry, because of the difficulty of identifying solution symmetries that are not also constraint symmetries without examining the set of solutions. We have pointed out that a CSP instance with very few solutions must have many solution symmetries, even if it has no constraint symmetry, and that eliminating constraint symmetry can lead to an increase in solution symmetry. For a k -ary CSP, the solution symmetries are the automorphisms of the k -ary nogood hypergraph; hence, finding new nogoods of arity up to k and adding them to the CSP can allow the constraint symmetry group to expand towards the solution symmetry group. It is not yet clear how best to do this, since we have shown that the order in which new nogoods are added can affect the resulting symmetry group. However, this approach has the potential to reduce the search effort to solve the problem, since symmetry-breaking methods avoid exploring assignments that are symmetrically equivalent to assignments explored elsewhere in the search, and hence working with a larger symmetry group may allow more of the search tree to be pruned.

Acknowledgments

We thank members of the SBDS group, especially Warwick Harvey, Ian Gent and Steve Linton for their discussions on this topic; a review of symmetry definitions by Iain McDonald was also useful. We are grateful to Brendan McKay for his help in finding automorphism groups of graphs using the software tool NAUTY. This material is based in part on works supported by the Science Foundation Ireland under Grant No. 00/PI.1/C075; the authors were also supported by SymNet, the U.K. Symmetry and Search Network. Most of this work was completed while the third author was at the Department of Computer Science, University of York, UK, and the fourth author was at Cork Constraint Computation Centre, University College, Cork, Ireland.

References

1. A. Aguirre. How to Use Symmetries in Boolean Constraint Solving. In F. Benhamou and A. Colmerauer, editors, *Constraint Logic Programming: Selected Research*, pages 287–306. MIT Press, 1992.
2. R. Backofen and S. Will. Excluding Symmetries in Constraint-Based Search. In J. Jaffar, editor, *Principles and Practice of Constraint Programming - CP'99*, LNCS 1713, pages 73–87. Springer, 1999.
3. B. Benhamou. Study of symmetry in constraint satisfaction problems. In *Proceedings of the 2nd Workshop on Principles and Practice of Constraint Programming, PPCP'94*, pages 246–254, May 1994.
4. B. Benhamou and L. Sais. Theoretical study of symmetries in propositional calculus and applications. In D. Kapur, editor, *Automated Deduction - CADE-11*, LNAI 607, pages 281–294. Springer-Verlag, 1992.
5. C. Bessière, E. Hebrard, B. Hnich, and T. Walsh. The tractability of global constraints. In M. Wallace, editor, *Principles and Practice of Constraint Programming - CP 2004*, volume LNCS 3258, pages 716–720. Springer, 2004.

6. C. A. Brown, L. Finkelstein, and J. Paul Walter "Purdum. Backtrack Searching in the Presence of Symmetry. In T. Mora, editor, *Applied Algebra, Algebraic Algorithms and Error-Correcting Codes*, LNCS 357, pages 99–110. Springer-Verlag, 1988.
7. J. Crawford, M. Ginsberg, E. Luks, and A. Roy. Symmetry-Breaking Predicates for Search Problems. In *Proceedings KR'96*, pages 149–159, Nov. 1996.
8. T. Fahle, S. Schamberger, and M. Sellmann. Symmetry Breaking. In T. Walsh, editor, *Principles and Practice of Constraint Programming - CP 2001*, LNCS 2239, pages 225–239. Springer, 2001.
9. F. Focacci and M. Milano. Global Cut Framework for Removing Symmetries. In T. Walsh, editor, *Principles and Practice of Constraint Programming - CP 2001*, LNCS 2239, pages 77–92. Springer, 2001.
10. E. C. Freuder. Eliminating Interchangeable Values in Constraint Satisfaction Problems. In *Proceedings AAAI'91*, volume 1, pages 227–233, 1991.
11. The GAP Group. *GAP – Groups, Algorithms, and Programming, Version 4.4*, 2005. (<http://www.gap-system.org>).
12. I. P. Gent, W. Harvey, T. Kelsey, and S. Linton. Generic SBDD using Computational Group Theory. In F. Rossi, editor, *Principles and Practice of Constraint Programming - CP 2003*, LNCS. Springer, 2003.
13. I. P. Gent and B. M. Smith. Symmetry Breaking During Search in Constraint Programming. In W. Horn, editor, *Proceedings ECAI'2000, the European Conference on Artificial Intelligence*, pages 599–603, 2000.
14. P. Jégou. Decomposition of Domains Based on the Micro-Structure of Finite Constraint-Satisfaction Problems. In *Proceedings AAAI'93*, pages 731–736, 1993.
15. B. McKay. Practical Graph Isomorphism. *Congressus Numerantium*, 30:45–87, 1981. (The software tool NAUTY is available for download from <http://cs.anu.edu.au/~bdm/nauty/>)
16. P. Meseguer and C. Torras. Exploiting symmetries within constraint satisfaction search. *Artificial Intelligence*, 129:133–163, 2001.
17. K. E. Petrie and B. M. Smith. Symmetry Breaking in Graceful Graphs. In F. Rossi, editor, *Principles and Practice of Constraint Programming - CP 2003*, LNCS 2833, pages 930–934. Springer, 2003.
18. J.-F. Puget. On the Satisfiability of Symmetrical Constrained Satisfaction Problems. In J. Komorowski and Z. W. Ras, editors, *Methodologies for Intelligent Systems (Proceedings of ISMIS'93)*, LNAI 689, pages 350–361. Springer-Verlag, 1993.
19. J.-F. Puget. Automatic Detection of Variable and Value Symmetries. In P. van Beek, editor, *Principles and Practice of Constraint Programming - CP 2005*, LNCS 3709, pages 475–489. Springer, 2005.
20. A. Ramani and I. L. Markov. Automatically Exploiting Symmetries in Constraint Programming. In B. Faltings, A. Petcu, F. Fages, and F. Rossi, editors, *Recent Advances in Constraints, Joint ERCIM/CoLogNet International Workshop on Constraint Solving and Constraint Logic Programming, CSCLP 2004*, LNCS 3419, pages 98–112. Springer, 2005.
21. P. Roy and F. Pachet. Using Symmetry of Global Constraints to Speed up the Resolution of Constraint Satisfaction Problems. In *Workshop on Non Binary Constraints, ECAI-98*, Aug. 1998.
22. T. Walsh. SAT v CSP. In *Proceedings CP'2000*, pages 441–456, 2000.