

Internal Symmetry

Marijn Heule¹ and Toby Walsh²

¹ *Delft University of Technology, The Netherlands*

`marijn@heule.nl`

² *NICTA and UNSW, Sydney, Australia*

`toby.walsh@nicta.com.au`

Abstract. We have been studying the internal symmetries within an individual solution of a constraint satisfaction problem [1]. Such internal symmetries can be compared with solution symmetries which map between different solutions of the same problem. We show that we can take advantage of both types of symmetry when solving constraint satisfaction solutions within two benchmark domains. By identifying internal symmetries and breaking solution symmetries, we are able to increase the size of problems which have been solved.

1 Introduction

Symmetry occurs in several different forms. For example, when finding magic squares (prob019 in CSPLib [2]), we have the solution symmetries that describe the rotation and reflection of one solution onto another. We can also have internal symmetries which describe the mappings *within* each solution [1]. Methods for identifying and dealing with solution symmetries have been extensively studied. We can, for instance, post symmetry breaking constraints to eliminate symmetric solutions [3, 4]. Internal symmetries, on the other hand, have not received as much attention. However, they can be dealt with in a similar way. We simply add constraints that limit search to those solutions with a given internal symmetry. In addition, we can limit search further by only branching on the subset of decisions that then generate a complete solution. To demonstrate the value of exploiting such internal symmetries, we report results on two benchmark domains: Van der Waerden numbers and graceful graphs. Some (but not all) of the results reported here will first appear in [1].

2 Solution symmetry

A symmetry σ is a bijection on assignments. Given a set of assignments A and a symmetry σ , we write $\sigma(A)$ for $\{\sigma(a) \mid a \in A\}$. Similarly, given a set of symmetries Σ , we write $\Sigma(A)$ for $\{\sigma(a) \mid a \in A, \sigma \in \Sigma\}$. A special type of symmetry, called *solution symmetry* is a symmetry *between* the solutions of a problem. Such a symmetry maps solutions onto solutions. A solution is a set of assignments that satisfy every constraint in the problem. More formally, a problem has the *solution symmetry* σ iff σ of any solution is itself a solution [5]. The set of solution symmetries Σ of a problem forms a group under composition.

We say that two sets of assignments A and B are in the same *symmetry class* of Σ iff there exists $\sigma \in \Sigma$ such that $\sigma(A) = B$.

Running example. *The magic squares problem is to label a n by n square so that every row, column and diagonal have the same sum (prob019 in CSPLib [2]). A normal magic square contains the integers 1 to n^2 . We model this with n^2 variables where $X_{i,j} = k$ iff the i th column and j th row is labelled with the integer k .*

“Lo Shu” is an important object in ancient Chinese mathematics. It is the smallest non-trivial normal magic square and has been known for over four thousand years:

4	9	2
3	5	7
8	1	6

(1)

The magic squares problem has a number of solution symmetries. For example, consider the symmetry σ_d that reflects a solution in the leading diagonal. This map “Lo Shu” onto a symmetric solution:

6	7	2
1	5	9
8	3	4

(2)

Any other rotation or reflection of the square maps one solution onto another. The 8 symmetries of the square (the dihedral group of order 8) are thus all solution symmetries of this problem. In fact, there are only 8 different magic square of order 3, and all are in the same symmetry class as “Lo Shu”.

To eliminate such solution symmetry, we can, for instance, post symmetry breaking constraints that rule out symmetric solutions [3, 4, 6–13]. For example, to eliminate σ_d , we simply post an inequality constraint to ensure that the top left corner is smaller than its symmetry, the bottom right corner. This selects (1) and eliminates (2).

3 Internal symmetry

Symmetries can also be found within individual solutions of a constraint satisfaction problem. We say that a solution A *contains* the internal symmetry σ (or equivalently σ is an internal symmetry *within* this solution) iff $\sigma(A) = A$.

Running example. *Consider again “Lo Shu”, the smallest normal magic square. This contains a simple internal symmetry. To see this, consider the solution symmetry σ_{inv} that inverts labels, mapping k onto n^2+1-k . This solution symmetry maps “Lo Shu” onto a different (but symmetric) solution. However, if we now*

apply the solution symmetry σ_{180} that rotates the square 180° , we map back onto the original solution:

$$\begin{array}{ccc}
 \begin{array}{|c|c|c|} \hline 4 & 9 & 2 \\ \hline 3 & 5 & 7 \\ \hline 8 & 1 & 6 \\ \hline \end{array} & \begin{array}{c} \sigma_{inv} \\ \Rightarrow \\ \Leftarrow \\ \sigma_{180} \end{array} & \begin{array}{|c|c|c|} \hline 6 & 1 & 8 \\ \hline 7 & 5 & 3 \\ \hline 2 & 9 & 4 \\ \hline \end{array}
 \end{array}$$

Consider the composition of these two symmetries: $\sigma_{inv} \circ \sigma_{180}$. This symmetry both inverts the labels in the square and rotates the square 180° . As this symmetry maps “Lo Shu” onto itself, we see that the solution “Lo Shu” contains the internal symmetry $\sigma_{inv} \circ \sigma_{180}$.

Note that a solution symmetry is a property of every solution whilst an internal symmetry is a property of just the given solution.

Running example. Consider the following magic square:

14	11	5	4
1	8	10	15
12	13	3	6
7	2	16	9

This is one of the oldest known magic squares, dating from a 10th century engraving on the Parshvanath Jain temple in Khajuraho, India. $\sigma_{inv} \circ \sigma_{180}$ is not an internal symmetry contained within this solution:

$$\begin{array}{ccc}
 \begin{array}{|c|c|c|c|} \hline 14 & 11 & 5 & 4 \\ \hline 1 & 8 & 10 & 15 \\ \hline 12 & 13 & 3 & 6 \\ \hline 7 & 2 & 16 & 9 \\ \hline \end{array} & \begin{array}{c} \Leftrightarrow \\ \sigma_{inv} \circ \sigma_{180} \end{array} & \begin{array}{|c|c|c|c|} \hline 8 & 1 & 15 & 10 \\ \hline 11 & 14 & 4 & 5 \\ \hline 2 & 7 & 9 & 16 \\ \hline 13 & 12 & 6 & 3 \\ \hline \end{array}
 \end{array}$$

However, this internal symmetry is found within other order 4 solutions. Consider Albrecht Dürer’s famous magic square:

16	3	2	13
5	10	11	8
9	6	7	12
4	15	14	1

(3)

This appears in his engraving “Melencolia I” of 1514 (as indicated by the two middle squares of the bottom row). It also plays a role in Dan Brown’s novel “The Lost Symbol”. The internal symmetry $\sigma_{inv} \circ \sigma_{180}$ is contained within (3):

$$\begin{array}{ccc}
 \begin{array}{|c|c|c|c|} \hline 16 & 3 & 2 & 13 \\ \hline 5 & 10 & 11 & 8 \\ \hline 9 & 6 & 7 & 12 \\ \hline 4 & 15 & 14 & 1 \\ \hline \end{array} & \begin{array}{c} \sigma_{inv} \\ \Rightarrow \\ \Leftarrow \\ \sigma_{180} \end{array} & \begin{array}{|c|c|c|c|} \hline 1 & 14 & 15 & 4 \\ \hline 12 & 7 & 6 & 9 \\ \hline 8 & 11 & 10 & 5 \\ \hline 13 & 2 & 3 & 16 \\ \hline \end{array}
 \end{array}$$

Thus we can conclude that $\sigma_{inv} \circ \sigma_{180}$ is an internal symmetry contained within some but not all solutions of the normal magic squares problem. In fact, 48 out of the 880 distinct normal magic squares of order 4 contain this internal symmetry. On the other hand, $\sigma_{inv} \circ \sigma_{180}$ is a solution symmetry of normal magic square problems of every size.

A solution containing an internal symmetry can often be described by a subset of assignments and one or more symmetries acting on this subset that generate a complete set of assignments. Given a set of symmetries Σ , we write Σ^* for the closure of Σ . That is, $\Sigma^0 = \Sigma$, $\Sigma^i = \{\sigma_1 \circ \sigma_2 \mid \sigma_1 \in \Sigma, \sigma_2 \in \Sigma^{i-1}\}$, $\Sigma^* = \bigcup_i \Sigma^i$. Given a solution A , we say the subset B of A and the symmetries Σ generate A iff $A = B \cup \Sigma^*(B)$. In this case, we also describe A as containing the internal symmetries Σ .

Running example. Consider the following magic square:

$$\begin{array}{|c|c|c|c|} \hline 1 & 8 & 12 & 13 \\ \hline 14 & 11 & 7 & 2 \\ \hline 15 & 10 & 6 & 3 \\ \hline 4 & 5 & 9 & 16 \\ \hline \end{array} \quad (4)$$

This contains the internal symmetry $\sigma_{inv} \circ \sigma_{180}$. Half this magic square and $\sigma_{inv} \circ \sigma_{180}$ generate the whole solution:

$$\begin{array}{|c|c|c|c|} \hline 1 & 8 & 12 & 13 \\ \hline 14 & 11 & 7 & 2 \\ \hline - & - & - & - \\ \hline - & - & - & - \\ \hline \end{array} \quad \Leftrightarrow \quad \begin{array}{|c|c|c|c|} \hline - & - & - & - \\ \hline - & - & - & - \\ \hline 15 & 10 & 6 & 3 \\ \hline 4 & 5 & 9 & 16 \\ \hline \end{array}$$

$\sigma_{inv} \circ \sigma_{180}$

In fact, (4) can be generated from just the first quadrant and two symmetries: $\sigma_{inv} \circ \sigma_{180}$ and a symmetry τ which constructs a 180° rotation of the first quadrant in the second quadrant, decrementing those squares on the leading diagonal and incrementing those on the trailing diagonal (the same symmetry constructs the third quadrant from the fourth). More precisely, τ makes the following mappings:

$$\begin{array}{|c|c|c|c|} \hline a & b & - & - \\ \hline c & d & - & - \\ \hline - & - & - & - \\ \hline - & - & - & - \\ \hline \end{array} \quad \Rightarrow \quad \begin{array}{|c|c|c|c|} \hline - & - & d+1 & c-1 \\ \hline - & - & b-1 & a+1 \\ \hline - & - & - & - \\ \hline - & - & - & - \\ \hline \end{array}$$

τ

The example hints at how we can exploit internal symmetries within solutions. We will limit search to a subset of the decision variables that generates a complete set of assignments and construct the rest of the solution using the generating symmetries.

4 Theoretical properties

We list some properties of internal symmetries that will be used to help find solutions.

4.1 Set of internal symmetries within a solution

Like solution symmetries, the internal symmetries within a solution form a group. A solution A *contains* a set of internal symmetries Σ (or equivalently Σ are internal symmetries *within* the solution) iff A contains σ for every $\sigma \in \Sigma$.

Proposition 1. *The set of internal symmetries Σ within a solution A form a group under composition.*

The proof of this proposition and all subsequent propositions are given in [1].

4.2 Symmetries within and between solutions

In general, there is no relationship between the solution symmetries of a problem and the internal symmetries within a solution of that problem. There are solution symmetries of a problem which are not internal symmetries within any solution of that problem, and vice versa. The problem $Z_1 \neq Z_2$ has the solution symmetry that swaps Z_1 with Z_2 , but no solutions of $Z_1 \neq Z_2$ contain this internal symmetry. On the other hand, the solution $Z_1 = Z_2 = 0$ of $Z_1 \leq Z_2$ contains the internal symmetry that swaps Z_1 and Z_2 , but this is not a solution symmetry of $Z_1 \leq Z_2$ (since $Z_1 = 0, Z_2 = 1$ is a solution but its symmetry is not). When all solutions of a problem contain the same internal symmetry, we can be sure that this is a solution symmetry of the problem itself.

Proposition 2. *If all solutions of a problem contain an internal symmetry then this is a solution symmetry.*

By modus tollens, it follows that if σ is not a solution symmetry of a problem then there exists at least one solution which does not contain the internal symmetry σ .

4.3 Symmetries of symmetric solutions

We next consider internal symmetries contained within symmetric solutions. In general, the symmetry of a solution contains the conjugate of any internal symmetry contained within the original solution.

Proposition 3. *If the solution A contains the internal symmetry σ and τ is any (other) symmetry then $\tau(A)$ contains the internal symmetry $\tau \circ \sigma \circ \tau^{-1}$.*

In the special case that symmetries commute, the symmetry of a solution contains the same internal symmetries as the original problem. Two symmetries σ and τ *commute* iff $\sigma \circ \tau = \tau \circ \sigma$.

Proposition 4. *If the solution A contains the internal symmetry σ and τ commutes with σ then $\tau(A)$ also contains the internal symmetry σ .*

Running example. Consider again “Lo Shu”, the smallest normal magic square. This solution contains the internal symmetry $\sigma_{inv} \circ \sigma_{180}$. This particular symmetry commutes with any rotation symmetry. For instance, consider the rotation of “Lo Shu” by 90° clockwise:

$$\begin{array}{ccc}
 \begin{array}{|c|c|c|} \hline 8 & 3 & 4 \\ \hline 1 & 5 & 9 \\ \hline 6 & 7 & 2 \\ \hline \end{array} & \begin{array}{c} \sigma_{inv} \\ \Rightarrow \\ \Leftarrow \\ \sigma_{180} \end{array} & \begin{array}{|c|c|c|} \hline 2 & 7 & 6 \\ \hline 9 & 5 & 1 \\ \hline 4 & 3 & 8 \\ \hline \end{array}
 \end{array}$$

This symmetry of “Lo Shu” also contains the internal symmetry $\sigma_{inv} \circ \sigma_{180}$.

4.4 Symmetry breaking

Finally, we consider the compatibility of eliminating symmetric solutions and focusing search on those solutions that contain particular internal symmetries. In general, the two techniques are incompatible. Symmetric breaking may eliminate all those solutions which contain a given internal symmetry.

Running example. Consider the following magic square:

$$\begin{array}{|c|c|c|c|} \hline 1 & 4 & 13 & 16 \\ \hline 14 & 15 & 2 & 3 \\ \hline 8 & 5 & 12 & 9 \\ \hline 11 & 10 & 7 & 6 \\ \hline \end{array} \tag{5}$$

This contains the internal symmetry $\sigma_v \circ \sigma_{inv}$ that inverts all values and reflects the square in the vertical axis:

$$\begin{array}{ccc}
 \begin{array}{|c|c|c|c|} \hline 1 & 4 & 13 & 16 \\ \hline 14 & 15 & 2 & 3 \\ \hline 8 & 5 & 12 & 9 \\ \hline 11 & 10 & 7 & 6 \\ \hline \end{array} & \begin{array}{c} \sigma_{inv} \\ \Rightarrow \\ \Leftarrow \\ \sigma_v \end{array} & \begin{array}{|c|c|c|c|} \hline 16 & 13 & 4 & 1 \\ \hline 3 & 2 & 15 & 14 \\ \hline 9 & 12 & 5 & 8 \\ \hline 6 & 7 & 10 & 11 \\ \hline \end{array}
 \end{array}$$

Note that this internal symmetry can only occur within magic squares of even order or of order 1.

Suppose symmetry breaking eliminates all solutions in the same symmetry class as (5) except for a symmetric solution which is a 90° clockwise rotation of (5). This solution does not contain the internal symmetry $\sigma_v \circ \sigma_{inv}$. In fact, this rotation of (5) contains the internal symmetry that inverts all values and reflects the square in the horizontal axis.

$$\begin{array}{ccc}
 \begin{array}{|c|c|c|c|} \hline 11 & 8 & 14 & 1 \\ \hline 10 & 5 & 15 & 4 \\ \hline 7 & 12 & 2 & 13 \\ \hline 6 & 9 & 3 & 16 \\ \hline \end{array} & \begin{array}{c} \Leftrightarrow \\ \sigma_v \circ \sigma_{inv} \end{array} & \begin{array}{|c|c|c|c|} \hline 16 & 3 & 9 & 6 \\ \hline 13 & 2 & 12 & 7 \\ \hline 4 & 15 & 5 & 10 \\ \hline 1 & 14 & 8 & 11 \\ \hline \end{array}
 \end{array}$$

We can identify a special case where symmetry breaking does not change any internal symmetry within solutions. Suppose symmetry breaking only eliminates symmetries which commute with the internal symmetry contained within a particular solution. In this case, whilst symmetry breaking may eliminate the given solution, it must leave a symmetric solution containing the given internal symmetry. Given a set of constraints C with solution symmetries Σ , we say that a set of symmetry breaking constraints S is *sound* iff for every solution of C there exists at least one solution of $C \cup S$ in the same symmetry class.

Proposition 5. *Given a set of constraints C with solution symmetries Σ , a sound set of symmetry breaking constraints S , and a solution A containing the internal symmetry σ , if σ commutes with every symmetry in Σ then there exists a solution of $C \cup S$ in the same symmetry class as A also containing the internal symmetry σ .*

Running example. *Consider the internal symmetry $\sigma_{inv} \circ \sigma_{180}$ contained within some (but not all) normal magic squares. This particular symmetry commutes with every rotation, reflection and inversion solution symmetry of the problem. Hence, if there is a solution with the internal symmetry $\sigma_{inv} \circ \sigma_{180}$, this remains true after breaking the rotational, reflection and inversion symmetries. However, as in the last example, there are internal symmetries contained within some solutions (like reflection in the vertical axis) which do not commute with all symmetries of the square.*

5 Detecting internal symmetries

Detection of internal symmetries in a class of problems consists of two steps: 1) selecting a good candidate solution, and 2) finding internal symmetries in that solution. As we already discussed earlier, an internal symmetry of one solution may not be an internal symmetry of another solution for the same problem. Recall in this context Albrecht Dürer’s famous magic square (3) that has internal symmetry $\sigma_{inv} \circ \sigma_{180}$, while most magic squares with the same size don’t. Before trying to detect internal symmetries, we consider a strategy to select a solution that is more likely to reveal internal symmetries.

The strategy consists of two parts: 1) determine a good size of the problem to look at and 2) filtering out solutions. First, consider the problem size. In order to exploit internal symmetries, internal symmetries observed at a certain size need also to hold for other sizes. Therefore, it makes sense to select problems of a size that have only very few (preferably only one) solutions. Yet for most problems, such as the magic squares, only the smallest problems have few solutions. Unfortunately these sizes might be too small to do some proper internal symmetry detection. However, for some problems, such as Van der Waerden numbers which we consider in Section 7, the number of solutions decreases while increasing the size. In those cases a “good” size is much larger.

After selecting the size, we filter out solutions using the following method. First, we compute the set \mathcal{S} consisting of all solutions of the selected size while

applying symmetry breaking. Let set \mathcal{V} be the set of variables that describe the solutions in \mathcal{S} . Then we repeat until $|\mathcal{S}| = 1$: Select a variable $v \in \mathcal{V}$ that is most frequently assigned the same value in the solutions of \mathcal{S} . Remove from \mathcal{S} all solutions that have v assigned to another value and $\mathcal{V} := \mathcal{V} \setminus \{v\}$.

Running example. Consider the magic squares problem of size 5×5 . After symmetry breaking³ (in this case 25 transpositions, 4 rotations, and 2 reflections), this problem has 144 solutions. In all those solutions, the first element is placed in the first (top left) entry. After fixing element 1, element 6 occurs most frequently in the same position. More specifically, it appears in 72 solutions on the last row in the fourth column. After fixing element 6 to that position, element 11 occurs in 36 of the remaining 72 solutions on second entry of the fourth row. The next elements to be placed are 16 and 21 and after fixing them, 24 solutions are left with the following pattern:

1	-	-	-	-
-	-	21	-	-
-	-	-	-	16
-	11	-	-	-
-	-	-	6	-

By repeating the procedure, we will select the following solution:

1	7	13	19	25
14	20	21	2	8
22	3	9	15	16
10	11	17	23	4
18	24	5	6	12

(6)

This solution has an interesting internal symmetry, which will be explained in the next example.

The last step is to detect internal symmetries in the selected solution. This step is the hardest to automate. As a general approach we propose the following. Construct a set of symmetries Σ that include at least all solutions symmetries. Let A be the selected solution. We then use a search program to find $s_1, s_2, \dots, s_k \in \Sigma$ that satisfy $s_1 \circ s_2 \circ \dots \circ s_k(A) = A$.

Although internal symmetries can be found using this approach – even when restricted to the solution symmetries – in our experience the most effective internal symmetries in the selected solution were often found manually.

³ For this example we use the symmetry breaking based on the score function that is defined on <http://www.grogon.com/magic/5x5pan144.php>. In short, out of each symmetry group the one is selected for which the elements in the top left corner have the smallest elements.

Running example. Consider the magic square we selected. We denote by $\sigma_{\text{row}(x,y,z)}$ a transposition of the rows such the row 1 is replaced by row x , row 2 by row y and row 3 by row z . An automated search program that only combines solution symmetries to detect internal symmetries can find the following symmetry $\sigma_{\text{row}(2,3,1)} \circ \sigma_{180} \circ \sigma_{\text{inv}} \circ \sigma_{\text{row}(3,1,2)}$:

1	7	13	19	25	$\sigma_{\text{row}(2,3,1)}$ \Rightarrow \Leftarrow $\sigma_{\text{row}(3,1,2)}$	14	20	21	2	8	σ_{180} \Rightarrow \Leftarrow σ_{inv}	12	6	5	24	18
14	20	21	2	8		22	3	9	15	16		4	23	17	11	10
22	3	9	15	16		1	7	13	19	25		25	19	13	7	1
10	11	17	23	4		10	11	17	23	4		16	15	9	3	22
18	24	5	6	12		18	24	5	6	12		8	2	21	20	14

However, this is not the most useful internal symmetry that can be observed in this magic square. A careful look at the filtering procedure above reveals that for the first five elements i that are fixed, we have $i = 1(\text{mod } 5)$. More importantly, all elements are fixed at a location that is a knight's move (two left, one up) from each other. A similar pattern is observable for the first five elements $i \in \{1, \dots, 5\}$: They are also a knight's move (two left, one down) from each other. Combining both patterns gives a construction method for odd magic squares that is very similar to the Siamese method⁴: After placing the first element, place the next element on the entry reachable by a knight's move (two left, one down). If the entry is already filled, then place the element in the entry directly right of the last placed element.

6 Exploiting internal symmetries

Once we have identified an internal symmetry which we conjecture may be contained in solutions of other (perhaps larger) instances of the problem, it is a simple matter to restrict search of a constraint solver to solutions of this form. In general, if we want to find solutions containing the internal symmetry σ , we post symmetry constraints of the form:

$$Z_i = j \Rightarrow \sigma(Z_i = j)$$

In addition, we can limit branching decisions to a subset of the decisions variables that generates a complete set of assignments. This can significantly reduce the size of the search space. Propagation of the problem and symmetry constraints may prune the search space even further.

7 Van der Waerden numbers

We illustrate the use of internal symmetries within solutions with two applications where we have been able to extend the state of the art. In the first, we

⁴ see http://en.wikipedia.org/wiki/Siamese_method

found new lower bound certificates for Van der Waerden numbers. Such numbers are an important concept in Ramsey theory. In the second application, we increased the size of graceful labellings known for a family of graphs. Graceful labelling has practical applications in areas like communication theory.

The Van der Waerden number, $W(k, l)$ is the smallest integer n such that if the integers 1 to n are colored with k colors then there are always at least l integers in arithmetic progression. For instance, $W(2, 3)$ is 9 since the two sets $\{1, 4, 5, 8\}$ and $\{2, 3, 6, 7\}$ contain no arithmetic progression of length 3, but every partitioning of the integers 1 to 9 into two sets contains an arithmetic progression of length 3 or more. The certificate that $W(2, 3) > 8$ can be represented with the following blocks:



Finding such certificates can be encoded as a constraint satisfaction problem. To test if $W(k, l) > n$, we introduce the Boolean variable $x_{i,j}$ where $i \in [0, k)$, $j \in [0, n)$ and constraints that each integer takes one color ($\bigvee_{i \in [0, k)} x_{i,j}$), and that no row of colors contains an arithmetic progression of length l ($x_{i,a} \wedge \dots \wedge x_{i,a+d(l-2)} \rightarrow \neg x_{i,a+d(l-1)}$). This problem has a number of solution symmetries. For example, we can reverse any certificate and get another symmetric certificate. We can also permute the colors and get another symmetric certificate:



Individual certificates also often contain internal symmetry. For example, the second half of the last certificate repeats the first half:



Hence, this certificate contains an internal symmetry that maps $x_{i,j}$ onto $x_{i,j+4 \pmod 8}$.

In fact, many known certificates can be generated from some simple symmetry operations on just the colors assigned to the first two or three integers. For instance, the first construction method for Van der Waerden certificates [14] made use of the observation that the largest possible certificates for the known numbers $W(k, l)$ ⁵ consist of a repetition of $l - 1$ times a base pattern. All these certificates, as well as all best lower bounds, have a base pattern of size $m = \frac{n}{l-1}$. This first method only worked for certificates for which m is prime. An improved construction method [15] generalises it for non-prime m .

An important concept in both construction methods is the primitive root⁶ of m denoted by r . Let p be the largest prime factor of m , then r is the smallest number for which:

$$r^i \pmod m \neq r^j \pmod m \quad \text{for } 1 \leq i < j < p \quad (7)$$

⁵ except for $W(3, 3)$

⁶ our use slightly differs from the conventional definition

We identified four internal symmetries:

- σ_{+m} : Apply to all elements $x_{i,j} := x_{i,j+m} \pmod n$
- σ_{+p} : Apply to all elements $x_{i,j} := x_{i,j+p} \pmod m$
- $\sigma_{\times r}$: Apply to all elements $x_{i,j} := x_{i,j \times r} \pmod m$
- $\sigma_{\times r^t}$: At least one subset maps onto itself after applying $x_{i,j} := x_{i,j \times r^t} \pmod m$ for a $t \in \{1, \dots, k\}$

Consider the largest known certificate for $W(5, 3)$ which has 170 elements. For this certificate, $m = 85, p = 17$, and $r = 3$. Below the base pattern is shown the first 85 elements. Notice that for this certificate A , $\sigma_{+p}(A)$ and $\sigma_{\times r}(A)$ are also certificates. In fact, after sorting the elements and permuting the subsets, this certificate is mapped onto itself after applying these symmetries.

18	20	24	26	33	36	38	44	65	66	74	76	79	80	5	13	17
22	30	34	35	37	41	43	50	53	55	61	82	83	6	8	11	12
23	25	28	29	39	47	51	52	54	58	60	67	70	72	78	14	15
31	32	40	42	45	46	56	64	68	69	71	75	77	84	2	4	10
19	21	27	48	49	57	59	62	63	73	81	85	1	3	7	9	16

↑ σ_{+p}

1	3	7	9	16	19	21	27	48	49	57	59	62	63	73	81	85
5	13	17	18	20	24	26	33	36	38	44	65	66	74	76	79	80
6	8	11	12	22	30	34	35	37	41	43	50	53	55	61	82	83
14	15	23	25	28	29	39	47	51	52	54	58	60	67	70	72	78
2	4	10	31	32	40	42	45	46	56	64	68	69	71	75	77	84

↓ $\sigma_{\times r}$

3	9	21	27	48	57	63	81	59	62	1	7	16	19	49	73	85
15	39	51	54	60	72	78	14	23	29	47	25	28	52	58	67	70
18	24	33	36	66	5	17	20	26	38	44	65	74	80	13	76	79
42	45	69	75	84	2	32	56	68	71	77	4	10	31	40	46	64
6	12	30	8	11	35	41	50	53	83	22	34	37	43	55	61	82

Given these symmetries, we can easily construct a complete certificate. We place the first and last elements (1 and 85) in the first subset and apply $\sigma_{\times r}$ to generate all elements in this subset. We apply σ_{+p} to partition the elements $\{1, \dots, 85\}$. Finally, we obtain a complete certificate by applying σ_{+m} . We generalised this into a construction method. To find a larger certificate $W(k, l, n)$, we test with a constraint solver for increasing $n \equiv 0 \pmod{l-1}$ whether a certificate can be obtained using the following steps:

- break solution symmetry by forcing that the first subset of the partition maps onto itself after applying $\sigma_{\times r^t}$
- choose $t \in \{1, \dots, k\}$, $q \in \{1, \dots, \frac{m}{p}\}$
- place elements q and m in the first subset
- apply the symmetries $\sigma_{\times r^t}$, $\sigma_{\times r}$, σ_{+p} , and σ_{+m} , to construct a certificate A with n' elements
- check with a constraint solver if A lacks an arithmetic progression of length l


Using this method we significantly improved some of the best known lower bounds⁷:

- $W(3, 7) > 48811$. The old bound was 43855.
- $W(4, 7) > 420217$. The old bound was 393469.

8 Graceful graphs

Our second application of internal symmetries is graceful labelling. A graph with e edges is called graceful if its vertices can be labelled with the distinct values $\{0, \dots, e\}$ in such a way that each edge gets a unique label when it is assigned the absolute difference of the vertices it connects. Graceful labelling has a wide range of applications in areas like radio astronomy, cryptography, communication networks and circuit design.

Whilst various classes of graphs are known to be graceful [16], there are others where it is not known but is conjectured that they are graceful. One such class is the class of double wheel graphs. The graph DW_n consists of two cycles of size n and a hub connected all the vertices. The largest double wheel graph that we have seen graceful labelled in the literature⁸ has size 10.



The problem of finding a graceful labelling can be specified using $2n + 1$ variables X_i with domain $\{0, \dots, e\}$. This problem has $16n^2$ solution symmetries [17]:

- Rotation of the vertices (n^2 symmetries)
- Inversion of the order of the vertices (4 symmetries)
- Swapping of the inner and outer wheel (2 symmetries)
- Inversion of the labels, $X_i := 4n - X_i$ (2 symmetries)

To identify internal symmetries, we generated all graceful labellings for DW_4 . This is the smallest double wheel graph with a graceful labelling. We observed two internal symmetries within the 44 solutions of DW_4 :

σ_{4n} : In 31 solutions, the hub had label $4n$ or 0 (σ_{inv}).

σ_{+2} : If $1 \leq X_i \leq n - 2$, then $X_{i+2} := X_i + 2$


Although we observed σ_{+2} , we restrict this internal symmetry to $1 \leq X_i \leq n - 4$ because it proved more effective.

When both symmetries are applied, the computational costs to find a graceful labelling is significantly reduced. Consider DW_{24} . To construct a graceful labelling, we first assign the hub to value 96 (applying σ_{4n}). Second, we label the first vertex of the outer wheel with 1 and label the first vertex of the inner

⁷ see www.st.ewi.tudelft.nl/sat/~waerden.php

⁸ see www.comp.leeds.ac.uk/bms/Graceful/

wheel with 2. Third, we apply symmetry σ_{+2} to label $n-1$ vertices with the labels $\{1, \dots, n-1\}$. Finally, we use a constraint solver to label the remaining vertices. Using this method we found the first known graceful labeling for DW_{24} .



The right table gives the runtime (in seconds) for our constraint solver to find graceful labellings of DW_n for the original problem (P) with and without symmetry breaking (SB) constraints [17]. The last column shows the results when we force internal symmetries within solutions. This also breaks the solution symmetries.

n	P	$P + \text{SB}$	$P + \sigma_{4n}, \sigma_{+2}$
4	0.01	0.01	0.01
8	1.88	0.78	0.1
12	82.12	25.18	0.3
16	1,608	706.35	4.97
20	21,980	8,272	17.25
24	> 36,000	> 36,000	157.87

9 Related work

Several forms of symmetry have been identified and exploited in search. For instance, Brown, Finkelstein and Purdom defined symmetry as a permutation of the variables leaving the set of solutions invariant [18]. This is a subset of the solution symmetries. For the propositional calculus, Krishnamurthy was one of the first to exploit symmetry [19]. He defined symmetry as a permutation of the variables leaving the set of clauses unchanged. Benhamou and Sais extended this to a permutation of the literals preserving the set of clauses [20].

The symmetry of individual constraints has also been considered. For example, Puget considered permutations of the variables leaving the set of constraints invariant [3]. He proved that such symmetry can be eliminated by the

addition of static constraints. Crawford *et al.* presented the first general method for constructing static constraints for eliminating solution symmetries [4]. Perhaps closest to this work is Puget's symmetry breaking method that considers symmetries which stabilize the current partial set of assignments [21]. By comparison, we consider only those symmetries which stabilize a complete set of assignments. A stabilizer maps individual assignments onto themselves, whilst an internal symmetry maps a set of assignments onto the same set (but may change every individual assignment).

Apart from symmetry, the idea of exploiting regularities in solutions of small sized problems in order to constrain large sized problems has been studied before. For instance, streamlining constraints in CP [22] and resolution tunnels in SAT [23]. In contrast to other work, internal symmetries focuses on a specific regularity: a mapping of a set of assignments onto itself.

10 Conclusions

We have been studying internal symmetries within a single solution of a constraint satisfaction problem [1]. Internal symmetries are properties of an individual solution. They can be compared with solution symmetries which are properties of all solutions of a constraint satisfaction problem. Both types of symmetry can be profitably exploited when solving constraint satisfaction problems. We illustrated the potential of doing this on two benchmark domains: Van der Waerden numbers and graceful graphs. With the first, we improved some of the best known lower bounds by around 10%. With the second, we more than doubled the size of the largest known double wheel graph with a graceful labelling from a wheel of size 10 to a wheel of size 24.

Acknowledgments

The authors are supported by the Dutch Organization for Scientific Research (NWO) under grant 617.023.611, the Australian Government's Department of Broadband, Communications and the Digital Economy and the ARC.

References

1. Heule, M., Walsh, T.: Symmetry within solutions. In: Proceedings of AAAI'10. (2010) 77–82
2. Gent, I., Walsh, T.: CSPLib: a benchmark library for constraints. Technical report, Technical report APES-09-1999 (1999).
3. Puget, J.F.: On the satisfiability of symmetrical constrained satisfaction problems. In Komorowski, J., Ras, Z., eds.: Proceedings of ISMIS'93. LNAI 689, Springer-Verlag (1993) 350–361
4. Crawford, J., Ginsberg, M., Luks, G., Roy, A.: Symmetry breaking predicates for search problems. In: Proceedings of the 5th International Conference on Knowledge Representation and Reasoning, (KR '96). (1996) 148–159

5. Cohen, D., Jeavons, P., Jefferson, C., Petrie, K., Smith, B.: Symmetry definitions for constraint satisfaction problems. *Constraints* **11**(2–3) (2006) 115–137
6. Flener, P., Frisch, A., Hnich, B., Kiziltan, Z., Miguel, I., Pearson, J., Walsh, T.: Breaking row and column symmetry in matrix models. In: 8th International Conference on Principles and Practices of Constraint Programming (CP-2002), Springer (2002)
7. Frisch, A., Hnich, B., Kiziltan, Z., Miguel, I., Walsh, T.: Global constraints for lexicographic orderings. In: 8th International Conference on Principles and Practices of Constraint Programming (CP-2002), Springer (2002)
8. Walsh, T.: General symmetry breaking constraints. In Benhamou, F., ed.: 12th International Conference on Principles and Practice of Constraint Programming (CP 2006). Volume 4204 of LNCS, Springer (2006)
9. Walsh, T.: Symmetry breaking using value precedence. In Brewka, G., Coradeschi, S., Perini, A., Traverso, P., eds.: 17th European Conference on Artificial Intelligence, IOS Press (2006) 168–172
10. Law, Y.C., Lee, J., Walsh, T., Yip, J.: Breaking symmetry of interchangeable variables and values. In: 13th International Conference on Principles and Practices of Constraint Programming (CP-2007), Springer-Verlag (2007)
11. Walsh, T.: Breaking value symmetry. In: 13th International Conference on Principles and Practices of Constraint Programming (CP-2007), Springer-Verlag (2007)
12. Katsirelos, G., Walsh, T.: Symmetries of symmetry breaking constraints. In: Proceedings of the 19th ECAI, IOS Press (2010)
13. Katsirelos, G., Narodytska, N., Walsh, T.: Static constraints for breaking row and column symmetry. In: 16th International Conference on Principles and Practices of Constraint Programming (CP-2010), Springer-Verlag (2010)
14. Rabung, J.R.: Some progression-free partitions constructed using Folkman’s method. *Canadian Mathematical Bulletin* **22**(1) (1979) 87–91
15. Herwig, P., Heule, M.J., van Lambalgen, M., van Maaren, H.: A new method to construct lower bounds for Van der Waerden numbers. *The Electronic Journal of Combinatorics* **14** (2007)
16. Gallian, J.A.: A dynamic survey of graph labelling. *The Electronic Journal of Combinatorics* **5** (1998) #DS6 Updated in 2008.
17. Petrie, K.E., Smith, B.M.: Symmetry breaking in graceful graphs. In: Principles and Practice of Constraint Programming. Volume 2833 of LNCS. (2003) 930–934
18. Brown, C., Finkelstein, L., Jr., P.P.: Backtrack searching in the presence of symmetry. In Mora, T., ed.: Applied Algebra, Algebraic Algorithms and Error-Correcting Codes, 6th International Conference. Volume 357 of LNCS. (1988) 99–110
19. Krishnamurthy, B.: Short proofs for tricky formulas. *Acta Informatica* **22**(3) (1985) 253–275
20. Benhamou, B., Sais, L.: Theoretical study of symmetries in propositional calculus and applications. In: Proceedings of 11th International Conference on Automated Deduction. Volume 607 of LNCS. (1992) 281–294
21. Puget, J.F.: Symmetry breaking using stabilizers. In Rossi, F., ed.: Proceedings of 9th International Conference on Principles and Practice of Constraint Programming (CP2003), Springer (2003)
22. Gomes, C., Sellmann, M.: Streamlined constraint reasoning. In: Principles and Practice of Constraint Programming CP 2004. Volume 3258 of Lecture Notes in Computer Science. (2004) 274–289
23. Kouril, M., Franco, J.: Resolution tunnels for improved SAT solver performance. In: Theory and Applications of Satisfiability Testing. Volume 3569 of Lecture Notes in Computer Science. (2005) 143–157