Symmetry within Solutions

Marijn Heule
TU Delft
The Netherlands
marijn@heule.nl

Toby Walsh
NICTA and UNSW
Sydney, Australia
toby.walsh@nicta.com.au

Abstract

We define the concept of an internal symmetry. This is a symmety within a solution of a constraint satisfaction problem. We compare this to solution symmetry, which is a mapping between different solutions of the same problem. We argue that we may be able to exploit both types of symmetry when finding solutions. We illustrate the potential of exploiting internal symmetries on two benchmark domains: Van der Waerden numbers and graceful graphs. By identifying internal symmetries we are able to extend the state of the art in both cases.

Introduction

Symmetry is an important feature of many combinatorial search problems. To be able to solve such problems, we often need to take account of symmetry. For example, when finding magic squares (prob019 in CSPLib (Gent and Walsh 1999)), we have the symmetries that rotate and reflect the square. Factoring such symmetry out of the search space is often critical when trying to solve large instances of a problem. Up till now, research on symmetry has mostly focused on symmetries *between* different solutions of the same problem. In this paper, we propose considering in addition the internal symmetries (that is, symmetries *within* each solution). Whilst it appears to be challenging to identify useful internal symmetries, such symmetries are easy to exploit. We simply add constraints that restrict search to those solutions with the required internal symmetry and limit branching to the subset of decisions that generate a complete solution. We will demonstrate the value of exploiting internal symmetries within solutions with experimental results on two benchmark domains: Van der Waerden numbers and graceful graphs.

Symmetry between solutions

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 simply a set of assignments that satisfy every constraint in the problem. More formally, we say that a problem has the *solution symmetry σ* iff σ of any solution is itself a solution (Cohen et al. 2006). As such mappings are associativity, and the inverse of a solution symmetry and the identity mapping are solution symmetries, 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 the sum of every row, column and diagonal are equal (prob019 in CSPLib (Gent and Walsh 1999)). A normal magic square contains the integers 1 to n^2. We model this with n^2 variables $X_{i,j}$ where $X_{i,j} = k$ iff the ith column and jth row is labelled with the integer k.*

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

$$\begin{array}{|c|c|c|} \hline 4 & 9 & 2 \\ \hline 3 & 5 & 7 \\ \hline 8 & 1 & 6 \\ \hline \end{array} \qquad (1)$$

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:

$$\begin{array}{|c|c|c|} \hline 6 & 7 & 2 \\ \hline 1 & 5 & 9 \\ \hline 8 & 3 & 4 \\ \hline \end{array} \qquad (2)$$

Any other rotation or reflection of the square maps one solution onto another. The 8 symmetries of the square 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.

Methods like posting symmetry breaking constraints have been developed to factor solution symmetry out of the search space (Puget 1993; Crawford et al. 1996; Flener et al. 2002; Frisch et al. 2002; Walsh 2006a; Walsh 2006b; Law et al. 2007; Walsh 2007). For example, we can eliminate σ_d by posting a constraint which ensures that the top left corner is smaller than its symmetry, the bottom right corner. This selects (1) and eliminates (2).

Symmetry within a solution

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 a internal symmetry *within* this solution) iff $\sigma(A) = A$.

Running example. *Consider again "Lo Shu". This contains an 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:*

4	9	2
3	5	7
8	1	6

$\begin{array}{c}\sigma_{inv}\\ \Rightarrow\\ \Leftarrow\\ \sigma_{180}\end{array}$

6	1	8
7	5	3
2	9	4

Consider the composition of these two symmetries: $\sigma_{inv} \circ \sigma_{180}$. As this symmetry maps "Lo Shu" onto itself, the solution "Lo Shu" contains the internal symmetry $\sigma_{inv} \circ \sigma_{180}$.

One significant difference between a solution symmetry and an internal symmetry is 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:*

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

(3)

$\sigma_{inv} \circ \sigma_{180}$ *is* not *an internal symmetry contained within this solution:*

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

$\begin{array}{c}\Leftrightarrow\\ \sigma_{inv} \circ \sigma_{180}\end{array}$

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

However, $\sigma_{inv} \circ \sigma_{180}$ is an internal symmetry found within the following solution:

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

(4)

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 again the solution (4) which contains the internal symmetry $\sigma_{inv} \circ \sigma_{180}$. Half this magic square and $\sigma_{inv} \circ \sigma_{180}$ generate the whole solution:*

1	8	12	13
14	11	7	2
-	-	-	-
-	-	-	-

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

-	-	-	-
-	-	-	-
15	10	6	3
4	5	9	16

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:

a	b	-	-
c	d	-	-
-	-	-	-
-	-	-	-

$\begin{array}{c}\Rightarrow\\ \tau\end{array}$

-	-	d+1	c-1
-	-	b-1	a+1
-	-	-	-
-	-	-	-

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.

Theoretical properties

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

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.*

Proof: The identity symmetry is trivially an internal symmetry. Internal symmetries are also trivially closed under composition, Finally, consider any $\sigma \in \Sigma$. As $\sigma(A) = A$, $\sigma^{-1}(\sigma(A)) = \sigma^{-1}(A)$. That is $A = \sigma^{-1}(A)$. Hence, the inverse of σ is an internal symmetry. $\square$

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.*

Proof: Consider any solution A. As all solutions of the problem contain the internal symmetry σ, $\sigma(A) = A$. Hence σ maps A onto itself, and $\sigma(A)$ is also a solution. □

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.

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 $\tau \circ \sigma \circ \tau^{-1}$.*

Proof: Consider the action of $\tau \circ \sigma \circ \tau^{-1}$ on $\tau(A)$. Now $\tau(\sigma(\tau^{-1}(\tau(A)))) = \tau(\sigma(A))$. But as A contains the internal symmetry σ, $\sigma(A) = A$. Hence $\tau(\sigma(A)) = \tau(A)$. Thus $\tau \circ \sigma \circ \tau^{-1}$ maps $\tau(A)$ onto itself. □

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.*

Proof: By Proposition 3, $\tau(A)$ contains the internal symmetry $\tau^\circ \sigma \circ \tau^{-1}$. But $\tau \circ \sigma \circ \tau^{-1} = \tau \circ \tau^{-1} \circ \sigma = \sigma$. □

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 again the solution (3). This contains the internal symmetry $\sigma_v \circ \sigma_{inv}$ that inverts all values and reflects the square in the vertical axis:*

1	4	13	16			16	13	4	1
14	15	2	3	σ_{inv}		3	2	15	14
8	5	12	9	$\Rightarrow$		9	12	5	8
11	10	7	6	$\Leftarrow$		6	7	10	11

σ_v

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 (3) except for a symmetric solution which is a $90°$ clockwise rotation of (3). This solution does not contain the internal symmetry $\sigma_v \circ \sigma_{inv}$. In fact, this rotation of (3) contains the internal symmetry that inverts all values and reflects the square in the horizontal *axis.*

11	8	14	1			16	3	9	6
10	5	15	4			13	2	12	7
7	12	2	13	$\Leftrightarrow$		4	15	5	10
6	9	3	16	$\sigma_v \circ \sigma_{inv}$		1	14	8	11

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.*

Proof: As S is sound, there exists a solution B of $C \cup S$ and $\tau \in \Sigma$ with $B = \tau(A)$. Now τ commutes with σ. Therefore by Proposition 4, B contains 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.*

Exploiting symmetries within solutions

The exploitation of internal symmetries involves two steps: finding internal symmetries, and then restricting search to solutions containing just these internal symmetries. The first step appears challenging. The definition of an internal symmetry is rather weak. There will be many uninteresting internal symmetries contained within a solution. We want to find internal symmetries that are likely to be contained within as yet unsolved instances of our problem. Although we do not yet have an efficient set of automated methods to do this, we can focus on simple symmetries (like the solution symmetries of the problem) and on small and already solved instances of a problem. This may suggest internal symmetries which might be contained in solutions of larger (perhaps open) problems.

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:

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

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.

Running example. *Consider again the problem of finding normal magic squares. We coded this problem in BProlog on a Pentium 4 3.2 GHz processor with 3GB of memory. In addition to the problem constraints, we used symmetry breaking constraints that eliminated most of the rotation, reflection and inversion solution symmetries:*

$$X_{1,1} < \min(X_{1,n}, X_{n,1}, X_{n,n}),\ X_{1,n} < X_{n,1},$$
$$X_{1,1} \leq n^2 + 1 - \max(X_{1,1}, X_{1,n}, X_{n,1}, X_{n,n}) \quad (5)$$

We also used symmetry constraints to ensure a simple internal symmetry was within the solution. Even and odd order magic squares often contain different internal symmetries so we used different symmetry constraints for even and odd n. For even n, we looked for solutions containing $\sigma_v \circ \sigma_{inv}$. Recall that this internal symmetry cannot be contained in solutions with odd n (except $n = 1$). For odd n, we looked instead for solutions containing $\sigma_{inv} \circ \sigma_{180}$. Hence, we used the following symmetry constraints for $1 \leq i, j \leq n$:

$$odd(n) \rightarrow X_{n+1-j,n+1-i} = n^2 + 1 - X_{i,j}$$
$$even(n) \rightarrow X_{n+1-i,j} = n^2 + 1 - X_{i,j} \quad (6)$$

In the following table, we report backtracks (b) and time (t) in seconds to find an order n normal magic square using the default branching heuristic, the problem constraints (P), the symmetry breaking constraints (5) and the internal symmetry constraints (6).

| n	P: b/t	$P + 5$: b/t	$P + 6$: b/t	$P + 5,6$: b/t
3	**2**/**0.00**	**1**/**0.00**	**1**/**0.00**	**1**/**0.00**
4	18/**0.00**	72/**0.00**	13/**0.00**	**12**/**0.00**
5	6656/0.13	5693/0.12	2287/0.03	**38**/**0.00**
6	$4.47 \cdot 10^9$/	$1.92 \cdot 10^8$/	**959018**/	**959018**/
	$2.43 \cdot 10^5$	6618.38	**28.08**	28.63

We see that both symmetry breaking and internal symmetry constraints speed up search. In addition, the combination of the two is usually better than either on their own.

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 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:
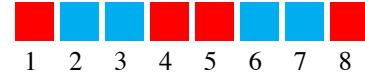
Finding such certificates can be encoded as a constraint satisfaction problem. To test if $W(k,l) > n$, we introduce the Boolean variable $x_{i,j}$ where $i \in [0,k)$, $j \in [0,n)$ and constraints that each integer takes one color ($\bigvee_{i \in [0,k)} x_{i,j}$), and that no row of colors contains an arithmetic progression of length l ($x_{i,a} \wedge \ldots \wedge x_{i,a+d(l-2)} \rightarrow \neg x_{i,a+d(l-1)}$). 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 the internal symmetry that maps $x_{i,j}$ onto $x_{i,j+4 \mod 8}$.

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 (Rabung 1979) made use of the observation that the largest possible certificates for the known numbers $W(k,l)$[1] 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 (Herwig et al. 2007) generalises it for non-prime m.

An important concept in both construction methods is the primitive root[2] of m denoted by r. Let p be the largest prime factor of m, then r is the smallest number for which:

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

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
$\quad\quad x_{i,j} := x_{i,j} \times r^t \pmod{m}$ for a $t \in \{1, \ldots, k\}$

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.

---

[1] Except for $W(3,3)$
[2] Our use slightly differs from the conventional definition

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

$$\Uparrow \sigma_{+p}$$

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

$$\Downarrow \sigma_{\times r}$$

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 $\{1, \ldots, 85\}$. 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 $t \in \{1, \ldots, k\}, q \in \{1, \ldots, \frac{m}{p}\}$
- place elements q and m in the first subset
- apply the symmetries $\sigma_{\times r^t}, \sigma_{\times r}, \sigma_{+p}$, 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 bounds[3]:

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

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 $\{0, \ldots, e\}$ 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 (Gallian 1998), there are others where it is not known but is conjectured that they are graceful. One such

[3]See www.st.ewi.tudelft.nl/sat/~waerden.php

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 literature[4] has size 10.
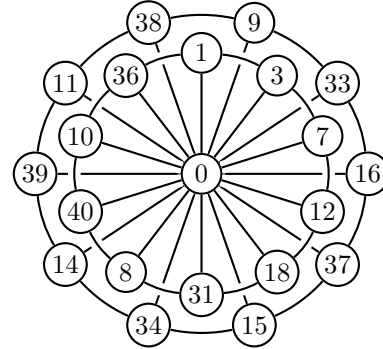
The problem of finding a graceful labelling can be specified using $2n + 1$ variables X_i with domain $\{0, \ldots, e\}$. This problem has $16n^2$ solution symmetries (Petrie and Smith 2003):

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

The following 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 (Petrie and Smith 2003). 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**

[4]See www.comp.leeds.ac.uk/bms/Graceful/
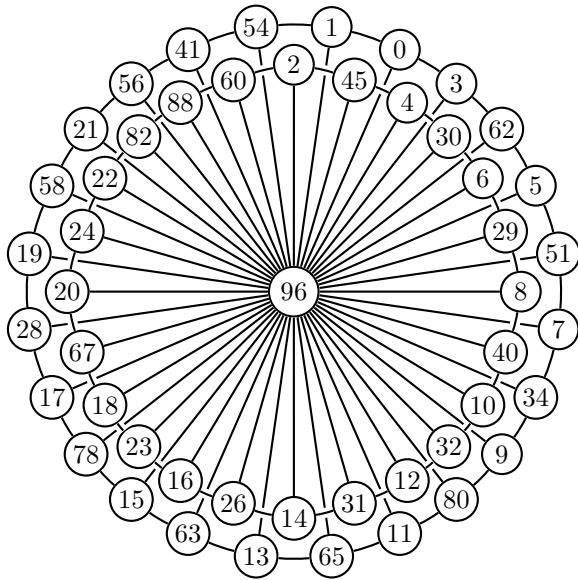
Related work

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 (Brown, Finkelstein, and Purdom. 1988). This is a subset of the solution symmetries. For the propositional calculus, Krishnamurthy was one of the first to exploit symmetry (Krishnamurthy 1985). 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 (Benhamou and Sais 1992). Perhaps closest to this work is Puget's symmetry breaking method that considers symmetries which stabilize the current partial set of assignments (Puget 2003). By comparison, we consider only those symmetries which stabilize a complete set of assignments.

Conclusions

We have defined the concept of an internal symmetry within a single solution of a constraint satisfaction problem. We compared this with the existing notion of symmetry between different solutions of the same problem. We demonstrated that we can exploit both types of symmetry when solving constraint satisfaction problems. We illustrated the potential of exploiting internal symmetry on two benchmark domains: Van der Waerden numbers and graceful graphs. By identifying internal symmetries, we were able to extend the state of the art in both cases. With Van der Waerden numbers, we improved two lower bounds by around 10%. With graceful graphs, we more than doubled the size of the largest known double wheel graph with a graceful labelling.

Acknowledgments

References

[Benhamou and Sais 1992] Benhamou, B., and Sais, L. 1992. Theoretical study of symmetries in propositional calculus and applications. In *Proc of CADE-11*, 281–294.

[Brown, Finkelstein, and Purdom. 1988] Brown, C.; Finkelstein, L.; and Purdom, P. 1988. Backtrack searching in the presence of symmetry. In *Proc. of AAECC-6*, 99–110.

[Cohen et al. 2006] Cohen, D.; Jeavons, P.; Jefferson, C.; Petrie, K.; and Smith, B. 2006. Symmetry definitions for constraint satisfaction problems. *Constraints* 11(2–3):115–137.

[Crawford et al. 1996] Crawford, J.; Ginsberg, M.; Luks, G.; and Roy, A. 1996. Symmetry breaking predicates for search problems. In *Proc. of KR '96*, 148–159.

[Flener et al. 2002] Flener, P.; Frisch, A.; Hnich, B.; Kiziltan, Z.; Miguel, I.; Pearson, J.; and Walsh, T. 2002. Breaking row and column symmetry in matrix models. In *Principles and Practice of Constraint Programming (CP 2002*, 462–476.

[Frisch et al. 2002] Frisch, A.; Hnich, B.; Kiziltan, Z.; Miguel, I.; and Walsh, T. 2002. Global Constraints for Lexicographic Orderings. In *Principles and Practice of Constraint Programming (CP 2002)*, 93–108.

[Gallian 1998] Gallian, J. A. 1998. A dynamic survey of graph labelling. *The Electronic Journal of Combinatorics* 5: DS6. Updated in 2009.

[Gent and Walsh 1999] Gent, I., and Walsh, T. 1999. CSPLib: a benchmark library for constraints. Technical report APES-09-1999.

[Herwig et al. 2007] Herwig, P.; Heule, M. J.; van Lambalgen, M.; and van Maaren, H. 2007. A new method to construct lower bounds for Van der Waerden numbers. *The Electronic Journal of Combinatorics* 14.

[Krishnamurthy 1985] Krishnamurthy, B. 1985. Short proofs for tricky formulas. *Acta Informatica* 22(3):253–275.

[Law et al. 2007] Law, Y.-C.; Lee, J.; Walsh, T.; and Yip, J. 2007. Breaking symmetry of interchangeable variables and values. In *Principles and Practice of Constraint Programming (CP 2007)*, 423–437.

[Petrie and Smith 2003] Petrie, K. E., and Smith, B. M. 2003. Symmetry breaking in graceful graphs. In *Principles and Practice of Constraint Programming*, 930–934.

[Puget 1993] Puget, J.-F. 1993. On the satisfiability of symmetrical constrained satisfaction problems. In *Proc. of IS-MIS'93*, 350–361.

[Puget 2003] Puget, J.-F. 2003. Symmetry Breaking Using Stabilizers. In *Principles and Practice of Constraint Programming (CP 2003)*, 585–599.

[Rabung 1979] Rabung, J. R. 1979. Some progression-free partitions constructed using Folkman's method. *Canadian Mathematical Bulletin* 22(1):87–91.

[Walsh 2006a] Walsh, T. 2006. General Symmetry Breaking Constraints. In *Principles and Practice of Constraint Programming (CP 2006)*, 650–664.

[Walsh 2006b] Walsh, T. 2006. Symmetry Breaking using Value Precedence. In *Proc. of ECAI-2006*, 168–172.

[Walsh 2007] Walsh, T. 2007. Breaking Value Symmetry. In *Principles and Practice of Constraint Programming (CP 2007)*, 880–887.