Groupoids and Conditional Symmetry

I.P. Gent[1], T. Kelsey[1], S.A. Linton[1], J. Pearson[2], and C.M. Roney-Dougal[3]

[1]School of Computer Science, University of St Andrews, St Andrews, UK
[2]Uppsala University, Department of Information Technology, Uppsala, Sweden
[3]School of Mathematics and Statistics, University of St Andrews, St Andrews, UK
{ipg,tom,sal}@cs.st-and.ac.uk, justin.pearson@it.uu.se,
colva@mcs.st-and.ac.uk

Abstract. We introduce groupoids – generalisations of groups in which not all pairs of elements may be multiplied, or, equivalently, categories in which all morphisms are invertible – as the appropriate algebraic structures for dealing with conditional symmetries in Constraint Satisfaction Problems (CSPs). We formally define the Full Conditional Symmetry Groupoid associated with any CSP, giving bounds for the number of elements that this groupoid can contain. We describe conditions under which a Conditional Symmetry sub-Groupoid forms a group, and, for this case, present an algorithm for breaking all conditional symmetries that arise at a search node. Our algorithm is polynomial-time when there is a corresponding algorithm for the type of group involved. We prove that our algorithm is both sound and complete – neither gaining nor losing solutions.

1 Introduction

Conditional symmetries in CSPs are parts of a problem that become interchangeable when some condition is satisfied. Typically the condition is that a subset of the variables have been given particular values. Definitions of conditional symmetry, together with initial approaches for identifying and breaking them, are given in [1,2,3,4]. A key problem is that the set of conditional symmetries of a CSP does not, in general, form a group. This motivates the research question: is there a suitable mathematical abstraction of conditional symmetry that can be used to identify, classify and break such symmeteries in an arbitrary CSP?

We describe groupoids as the class of mathematical objects that are the appropriate abstraction for conditional symmetries. Using groupoids we can describe, enumerate and analyse conditional symmetries for any CSP. Given a CSP, the Full Conditional Symmetry Groupoid, containing elements that capture both the symmetry and the condition under which it arises. This approach allows us to classify conditional symmetries in terms of the sub-groupoid(s) in which they are contained. Moreover, we can identify conditional symmetries that have properties that allow us to develop effective symmetry breaking techniques.

Groupoids are generalisations of groups, in which not all elements can be composed. A basic introduction to groupoids is given in [5]. The reason that we

C. Bessiere (Ed.): CP 2007, LNCS 4741, pp. 823–830, 2007.

have to leave groups behind when talking about conditional symmetries is that every element of a symmetry group can be applied to *all* partial assignments, and hence one cannot capture the concept of "condition".

2 Groupoids and Conditional Symmetries

We will use the following notation throughout: the set of variables will denoted with a V, individual variables $v_1, \ldots, v_n$ and sometimes w. For simplicity, we will assume that there is a single domain D for all variables (since constraints can be added to give different domains for chosen variables). A *literal* of a CSP is a pair (w, d) where $w \in V$ and $d \in D$.

Definition 1. *A* partial assignment *f is a set of literals that contains at most one literal for each variable. A partial assignment g is an* extension *of f, if $f \subseteq g$.*

We say that a partial assignment f is *complete* if $|f| = |V|$. That is, all variables are assigned values by f. A *solution* to a CSP is a complete partial assignment that violates no constraint.

Definition 2. *A* groupoid *is a set G, and a basis set B, together with a partial operation, composition, and two maps s and t from G to B that together satisfy the following:*

1. *composition, gh, of two elements, g and h in G is defined only when $t(g) = s(h)$ — this is what is meant by saying that the composition is partial;*
2. *if the products gh and hk are defined then $g(hk) = (gh)k$ is defined;*
3. *for every g in G there are left and right identity elements λ_g and ρ_g s.t. $\lambda_g g = g = g\rho_g$;*
4. *Each element has an inverse g^{-1} s.t. $gg^{-1} = \lambda_g$ and $g^{-1}g = \rho_g$.*

Often groupoids will be presented as collections of triples, $(s(g), g, t(g))$, with composition of two elements g and h can written as: $(s(g), g, t(g))(s(h), h, t(h)) = (s(g), gh, t(h))$, provided that $t(g) = s(h)$. convention of "acting from the right" so that gh means "do g then do h".

A *subgroupoid* of a group G is a subset of the elements of G that itself forms a groupoid under the same partial operation. The base set of a subgroupoid is the set of all sources and targets that occur for elements of the subgroupoid.

We now turn to formalising conditional symmetries as groupoid elements. The key problem in providing the definitions is that for composition to be defined, the target of one element has to be *equal* to the source of the next. This makes it impossible to represent sources and targets as the simple condition given by the constraint. Instead, we have a separate groupoid element for every partial assignment that satisfies a condition. We will develop the theory and show that this leads to well-defined groupoids and useful theoretical results.

We will denote the image of an object α under a map ϕ by α^{ϕ}. A *literal bijection* of a CSP is a bijection ϕ from the set of literals of the CSP to itself. The following definition begins the process of capturing the notion of a condition.

Definition 3. *A literal bijection is a* symmetry with respect to *a, where a is a subset of literals, if in its induced action on sets of literals, whenever f is a solution that contains a then f^{ϕ} is a solution and whenever f is a non-solution that contains a then f^{ϕ} is a non-solution.*

Before defining conditional symmetry groupoids, we need precise descriptions of conditions and symmetries arising from conditions. A *condition* is a predicate on literals. Any such predicate can be described by listing the sets of literals upon which it holds. If a symmetry in a CSP is present only when some condition holds, we will describe this situation by adding one generator to the conditional symmetry groupoid for each set of literals which satisfy the condition.

Definition 4. *We define the* full conditional symmetry groupoid G *of a CSP* P *to be the set of all triples as follows:*

$$G = \{(g, \pi, f) : g, f \text{ sets of literals}, \pi \text{ a symmetry with respect to } g, g^{\pi} = f\}.$$

The product (in this order) of two groupoid elements $(g, \pi, f), (h, \sigma, k)$ is defined only if $h = f$ in which case the product is $(g, \pi\sigma, k)$ and $g^{\pi\sigma}$ will equal k.

We work throughout with subgroupoids of the full conditional symmetry groupoid of the CSP. Note that if $(g, \pi, f) \in G$ and h is a set of literals that contains g, then $(h, \pi, h^{\pi}) \in G$, as π is a symmetry with respect to g only if it is a symmetry with respect to all extensions of g. We say that (h, π, h^{π}) is an *extension* of (g, π, f), and that the conditional symmetry groupoid is *closed under extensions*. We call the first entry of each triple its *precondition* and the last its *postcondition*.

Lemma 1. *With these definitions, the full conditional symmetry groupoid G of a CSP is a groupoid.*

Proof. The set of elements of G is clear, its base set is the power set of the set of literals. We must show that composition, where defined, is associative and that each element has left and right identities and inverses. Let $(g_1, \pi_1, f_1), (g_2, \pi_2, f_2)$ and $(g_3, \pi_3, f_3) \in G$, and suppose that $g_2 = f_1$ and $g_3 = f_2$. Then

$$((g_1, \pi_1, f_1)(g_2, \pi_2, f_2))(g_3, \pi_3, f_3) = (g_1, \pi_1\pi_2, g_1^{\pi_1\pi_2})(g_3, \pi_3, f_3).$$

Now, $g_2 = f_1 = g_1^{\pi_1}$, so $f_2 = g_2^{\pi_2} = g_1^{\pi_1\pi_2}$. Hence $g_3 = f_2 = g_1^{\pi_1\pi_2}$ and the product of the two groupoid elements in the previous displayed equation is defined, and is equal to $(g_1, \pi_1\pi_2\pi_3, g_1^{\pi_1\pi_2\pi_3})$. Conversely

$$\begin{aligned}(g_1, \pi_1, f_1)((g_2, \pi_2, f_2)(g_3, \pi_3, f_3)) &= (g_1, \pi_1, f_1)(g_2, \pi_2\pi_3, g_2^{\pi_2\pi_3}) \\ &= (g_1, \pi_1\pi_2\pi_3, g_1^{\pi_1\pi_2\pi_3}),\end{aligned}$$

as required. The right identity of (g, π, f) is $(f, 1, f)$, the left identity is $(g, 1, g)$, and the inverse is (f, π^{-1}, g), where by 1 we mean the identity mapping on the set of all literals.

Lemma 2. *The full conditional symmetry groupoid of a CSP has a well-defined partial action on the set of all sets of literals, which maps (non-)solutions to (non-)solutions.*

Proof. The action of an element (g, π, f) on a set h of literals is as follows. If $g \not\subseteq h$ then the condition of (g, π, f) is not satisfied, and so the action is undefined. If $g \subseteq h$ then $h^{(g,\pi,f)} := h^{\pi}$. To show that this is an action, we note that $h^{(g,\pi_1,f)(f,\pi_2,k)}$ is defined if and only if $g \subseteq h$ are sets of literals in which case we have

$$h^{(g,\pi_1,f)(f,\pi_2,k)} = h^{(g,\pi_1\pi_2,k)} = h^{\pi_1\pi_2} = (h^{(g,\pi_1,f)})^{(f,\pi_2,k)},$$

and that $h^{(g,1,f)} = h$ whenever $g \subseteq h$. Whenever h is a full assignment and $h^{(g,\pi,f)}$ is defined, then, since π is a symmetry with respect to g and $g \subseteq h$, h^{π} is a (non-)solution exactly when h is a (non-)solution.

Next we show that our notions of conditional symmetry strictly generalise the standard notions of unconditional symmetry.

Lemma 3. *The full conditional symmetry groupoid of a CSP P contains the group of all symmetries of P.*

Proof. Elements of the symmetry group of the CSP are of the form $(\emptyset, \pi, \emptyset)$.

In general, however, the full conditional symmetry groupoid of a CSP is far too large for practial computation, and we must restrict the situation somewhat.

Lemma 4. *The full conditional symmetry groupoid of a CSP P with $|V| = n$ and $|D| = d$ has order at least 2^{dn} and at most $2^{dn}(dn)!$.*

Proof. A precondition can be any subset of literals and the corresponding permutation can be any permutation of the full set of literals. The identity permutation is a symmetry with respect to any set of literals.

Of course, in practise we cannot identify *all* conditional symmetries of a CSP before solving it, so we identify as many non-identity conditional symmetries as possible, and then form the resulting groupoid.

Definition 5. *We say that a collection of conditional symmetry groupoid elements* $\{(g_i, \pi_i, f_i) \ : \ 1 \leq i \leq k\}$ generate *a conditional symmetry groupoid G if G is the smallest subgroupoid of the full conditional symmetry groupoid that contains all of the elements and is closed under extensions.*

The reason why any conditional symmetry groupoid is defined to be closed under extensions is that it must always be possible to consider conditions that are stronger than those initially given: if (v_1, α) implies some symmetry, then so does $\{(v_1, \alpha), (v_2, \beta)\}$. Later, when we discuss computing with conditional symmetry groupoids, we will describe a technique which avoids the creation of these additional elements.

Lemma 5. *The conditional symmetry groupoid G is generated by elements from $A := \{(g_i, \pi_i, f_i) \ : \ 1 \leq i \leq k\}$ if and only if each element of G is a product of elements that are extensions of elements of A and their inverses.*

Proof. We first show that the set of all elements of the conditional symmetry groupoid that are products of extensions of elements of A forms a groupoid and is closed under extensions. It is clear that it is closed under the partial product and the associativity follows from the associativity of the generalised conditional symmetry groupoid. If $(x_0, \sigma, x_k) = (x_0, \sigma_1, x_1)(x_1, \sigma_2, x_2) \cdots (x_k, \sigma_{k+1}, x_{k+1})$ is a product of extensions of elements of A, then we note that $(x, 1, x)$ is an extension of a left identity of an element of A, and similarly for $(y, 1, y)$, and that $(x_i, \sigma_i^{-1}, x_{i-1})$ and $(x_{i+1}, \sigma_{i+1}^{-1}, x_i)$ are extensions of right and left inverses of elements of A. Thus this set forms a groupoid. If (x, σ, y) is an extension of (x_0, σ, x_k) then it is clear that it is a product of elements of the form $(\overline{x_i}, \sigma_{i+1}, \overline{x_{i+1}})$, where $\overline{x_i}$ is a partial assignment that extends x_i.

Conversely, if (x, σ, y) is a product of extensions of elements of A then it must be contained in any groupoid that contains A, and so is an element of the groupoid generated by A. Thus the result follows.

3 Symmetry Breaking with Groupoids

Symmetry breaking by dominance detection for groups of symmetries uses the following principle: given the current search node g then for each previously found nogood f, and each symmetry π test the inclusion $f^\pi \subseteq g$. This can be done either by some supplied procedure, as in SBDD, or by adding new constraints that rule out all images under the symmetry group of extensions of f.

Lemma 6. *If $f \subseteq g$ are partial assignments, and (h, π, k) is an element of the full conditional symmetry groupoid, then $f^{(h,\pi,k)} \subseteq g^{(h,\pi,k)}$.*

Proof. If h is not a subset of f, then neither map is defined. Otherwise, $f^{(h,\pi,k)} = f^\pi = (f \setminus h)^\pi \cup k$ and $g^{(h,\pi,k)} = g^\pi = (g \setminus h)^\pi \cup k$. Since $(f \setminus h)^\pi \subseteq (g \setminus h)^\pi$, we have $f^\pi \subseteq g^\pi$ as required.

When implementing SBDD efficiently for groups or groupoids, the fact that groupoid actions are well-behaved with respect to extensions of partial assignments means that only failures at the top of subtrees need be stored during depth first search.

To check the dominance in SBDD in the group case we search for an element π of the symmetry group such that g is an extension of f^π. The situation with groupoids is a bit more complicated, as the action is only partial and so the following two cases have to be checked:

- Does there exist a groupoid element (h, γ, k) such that $f \subseteq h$ and $g \subseteq f^\gamma$?
- Otherwise we must search for some extension, f', of f and a groupoid element (h, γ, k) such that $f' \subseteq h$ and $g \subseteq f'^\gamma$.

Note that in case (ii) it is only worth considering extensions of f that assign values to no more variables than g, and one should only consider extensions that enable some new conditional symmetry to be used.

To avoid this two-case analysis, we *reverse* the normal process of SBDD and instead look for conditional symmetries that map the current partial assignment to a partial assignment that is an extension of the previous nogood.

Definition 6. *A partial assignment g is* dominated *by a nogood f with respect to a conditional symmetry groupoid G if there exists an element $(h, \gamma, k) \in G$ with $h \subseteq g$ and $f \subseteq g^\gamma$.*

This simplification is possible because groupoids *are* closed under inversion, and so if a map exists in one direction we can always consider its inverse.

Lemma 7. *Let G be a conditional symmetry groupoid of a CSP P and let a be a fixed partial assignment of P. The set of all elements of G of the form (a, π, a) form a group.*

Proof. Let $(a, \pi_1, a), (a, \pi_2, a) \in G$. The product $(a, \pi_1\pi_2, a)$ is defined and is of the correct form. Associativity holds because within this set of elements the product of two elements is always defined. The element $(a, 1, a) \in G$ and is the identity of this subgroup. The inverse $(a, \pi_1^{-1}, a) \in G$ is of the correct form.

Definition 7. *Let P be a CSP with conditional symmetry groupoid G generated by $X := \{(a_i, \pi_{ij}, b_i) : 1 \leq i \leq m, j \in \mathcal{J}_i\}$. Let A be a partial assignment, and let*

$$C := \bigcup_{a_i \subseteq A} a_i.$$

Then the local group at A with respect to X, *denoted $\mathcal{L}_X(A)$ is the group consisting of the permutations σ of all elements of G of the form (C, σ, C).*

It follows from Lemma 7 that the group at A is always a group, if nonempty. Since there is always the identity unconditional symmetry, the group at A is always nonempty, although it may be the trivial group $(C, 1, C)$.

Lemma 8. *Checking for dominance under the $\mathcal{L}_X(A)$ is sound.*

Proof. Since all elements of $\mathcal{L}_X(A)$ are symmetries with respect to a subset of A, they will all map (non-)solutions extending A to (non-)solutions.

We now consider a special case where checking for dominance under the group at A is also complete.

Theorem 1. *Let G be a conditional symmetry groupoid of a CSP P, generated by a set $X = \{(a_i, \pi_{ij}, a_i) \ : 1 \leq i \leq m, j \in \mathcal{J}_i\}$, where each a_i is a partial assignment. Assume that for $1 \leq i, k \leq m$ we have $a_i^{\pi_{kj}} = a_i$ for all $j \in \mathcal{J}_k$. Then checking for dominance under $\mathcal{L}_X(A)$ is complete as well as sound.*

Proof. Let $H := \mathcal{L}_X(A)$. We show that if $(g, \sigma, h) \in G$ and $g \subseteq A$ then $\sigma \in H$. It is clear that if σ is a product of π_{ij} for $a_i \subseteq A$ then $\sigma \in H$. We wish to show that if σ is a product including at least one π_{kj} for some $a_k \not\subseteq H$ then (g, σ, h) does not act on A. It is clear that no extension of (a_k, π_{kj}, a_k) can act on g as $a_k \not\subseteq A$, and hence that one cannot postmultiply any extension of (a_k, π_{kj}, a_k) to make it act on A. We show that one cannot premultiply any extension of (a_k, π_{kj}, a_k) by an element of G to make it act on A. Since $a_k^{\pi_{il}} = a_k$ for all l, if (x, σ, y) can premultiply an extension of (a_k, π_{kj}, a_k) then $a_k \subseteq y$ and so $a_k \subseteq x$ and hence $x \not\subseteq A$ and we are done.

Using this theorem, we can identify some special cases where testing dominance is possible in polynomial time.

Theorem 2. *Let G be a conditional symmetry groupoid of a CSP P, generated by a set $\{(a_i, \pi_{ij}, a_i) \ : 1 \leq i \leq m, j \in \mathcal{J}_i\}$ where each a_i is a partial assignment. Assume that for $1 \leq i, k \leq m$, whenever $a_i \cup a_k$ is a partial assignment then $a_i^{\pi_{kj}} = a_i$ for all $j \in \mathcal{J}_k$. If the actions of π_{ij} on the set of extensions of a_i all belong to a class of symmetries for which there exists a polynomial time algorithm to determine group dominance, then testing conditional symmetry dominance under the groupoid G can be done in polynomial-time at each node.*

Proof. By Theorem 1 it is both sound and complete to test each partial assignment A for dominance by previous nogoods using only the local group at A with respect to X. The local group at A with respect to X can be constructed in polynomial-time, by examining which conditions of which generators are subsets of A and then generating a group with the corresponding permutations. Hence

PREPROCESSING For each generator (a_i, π_{ij}, b_i) do

(i) If $a_i \neq b_i$ then return "inapplicable".
(ii) If π_{ij} is not a value symmetry on all literals in a_i, and on all literals involving variables not in a_i then return "inapplicable".
(ii) For each (a_k, π_{kl}, b_k) with $k \neq i$ such that $a_i \cup a_k$ is a partial assignment do
  - If $a_i^{\pi_{kl}} \neq a_k$ then return "inapplicable".

AT NODE A

1. Make a set L of the generators of G such that $a_i \subseteq A$.
2. Construct the local group $\mathcal{L}_X(A)$, which is $\langle \sigma : (a, \sigma, a) \in L \rangle$.
3. For each nogood B do
   (a) If B contains any literals with different variables from those in A, consider the next nogood.
   (b) Otherwise, by repeatedly computing point stabilisers check whether there exists a $\sigma \in \mathcal{L}_X(A)$ such that for all $(w, \alpha) \in A$ with $(w, \beta) \in B$ for some β we have $(w, \alpha)^\sigma = (w, \beta)$.
   (c) If such a σ exists, return "dominated".
4. Return "not dominated".

Fig. 1. Algorithm for sound and complete conditional symmetry breaking

if there exists a polynomial-time algorithm for testing whether A is dominated by a nogood by an element of a group of symmetries then it can now be applied.

Corollary 1. *Let G be a conditional symmetry groupoid of a CSP P, generated by a set $\{(a_i, \pi_{ij}, a_i) : 1 \leq i \leq m, j \in \mathcal{J}_i\}$ where each a_i is a partial assignment. Assume that for $1 \leq i, k \leq m$, whenever $a_i \cup a_k$ is a partial assignment then $a_i^{\pi_{kj}} = a_i$ for all $j \in \mathcal{J}_k$. If the actions of π_{ij} on the set of extensions of a_i are value symmetries then testing conditional symmetry dominance is in P.*

Proof. This is immediate from Theorem 2 and [6].

Fig. 1 shows a polynomial-time algorithm which takes as input a set of conditional symmetry groupoid generators for a CSP and determines whether the conditions of Corollary 1 apply and if so tests for groupoid dominance.

As in SBDD, we can safely backtrack when dominated, and continue to the next search node otherwise. Our approach is not restricted to value symmetries. We can obtain polynomial-time algorithms whenever a polynomial-time algorithm exists for symmetry breaking in the unconditional case.

4 Conclusions and Further Work

We have shown that careful use of the theory of groupoids allows us fully to capture the notion of conditional symmetry in CSPs.

We have identified the algebraic structure of conditional symmetries. With this we can study the whole set of conditional symmetries of a problem, rather than just a small subset given to us by a programmer, and analyse sub- and super-groupoids. We have shown the enormous numbers of conditional symmetries and complicated structure that arises when we generate them all. We have provided definitions and algorithms for conditional symmetry breaking. We have defined a notion of dominance, allowing us to give an analogue of SBDD for conditional symmetries. We have also shown that it is possible to identify useful conditional symmetry sub-groupoids that are small enough to permit effective conditional symmetries breaking.

References

1. Gent, I.P., McDonald, I., Smith, B.M.: Conditional symmetry in the all-interval series problem. In: Symcon'03 (2003)
2. Gent, I.P., McDonald, I., Miguel, I., Smith, B.M.: Approaches to conditional symmetry breaking. In: SymCon'04 (2004)
3. Zhang, Y., Freuder, E.C.: Conditional interchangeability and substitutability. In: SymCon'04 (2004)
4. Walsh, T.: General symmetry breaking constraints. In: Benhamou, F. (ed.) CP 2006. LNCS, vol. 4204, pp. 650–664. Springer, Heidelberg (2006)
5. Brown, R.: From groups to groupoids: A brief survey. Bulletins of the London Mathematical Society 19, 113–134 (1987)
6. Roney-Dougal, C.M., Gent, I.P., Kelsey, T., Linton, S.A.: Tractable symmetry breaking using restricted search trees. In: ECAI'04, pp. 211–215 (2004)