Static and Dynamic Structural Symmetry Breaking

Pierre Flener[1], Justin Pearson[1], Meinolf Sellmann[2], and Pascal Van Hentenryck[2]

[1] Dept of Information Technology, Uppsala University, Box 337, 751 05 Uppsala, Sweden
[2] Dept of Computer Science, Brown University, Box 1910, Providence, RI 02912, USA
{pierref,justin}@it.uu.se {sello,pvh}@cs.brown.edu

Abstract. We reconsider the idea of structural symmetry breaking (SSB) for constraint satisfaction problems (CSPs). We show that the dynamic dominance checks used in symmetry breaking by dominance-detection search for CSPs with piecewise variable *and* value symmetries have a static counterpart: there exists a set of constraints that can be posted at the root node and that breaks *all* these symmetries. The amount of these symmetry-breaking constraints is *linear* in the size of the problem, but they possibly remove a super-exponential number of symmetries on both values and variables. Moreover, static and dynamic structural symmetry breaking coincide for static variable and value orderings.

1 Introduction

Symmetry breaking has been the topic of intense research in recent years. Substantial progress was achieved in many directions, often exhibiting significant speedups for complex real-life problems arising, say, in configuration and network design. One of the interesting recent developments has been the design of general symmetry-breaking schemes such as symmetry breaking by dominance detection (SBDD) and symmetry breaking during search (SBDS). SBDD [1, 2] is particularly appealing as it combines low memory requirements with a number of dominance checks at each node linearly proportional to the depth of the search tree. It then became natural to study which classes of symmetries for CSPs admit polynomial-time dominance-checking algorithms. This issue was first studied in [9], where symmetry breaking for various classes of value symmetries was shown to take constant time and space (see also [7] for an elegant generalization to all value symmetries). It was revisited for CSPs with piecewise variable and value symmetry in [8], where a polynomial-time dominance-checking algorithm was given and the name 'structural symmetry breaking' (SSB) was coined. In parallel, researchers have investigated for many years (e.g., [4]) static symmetry breaking, which consists in adding constraints to the CSP in order to remove symmetries.

In this paper, after reviewing the basic concepts in Section 2, we show in Section 3 that the polynomial-time dominance-checking algorithm of [8] has a static counterpart, namely that there exists a static set of constraints for CSPs with piecewise symmetric variables and values that, when added to the CSP, results in a symmetry-free search tree. The amount of symmetry-breaking constraints is *linear* in the size of the problem, but possibly removes a super-exponential number of symmetries on both values and variables. In Section 4, we establish a clear link between static (SSSB) and dynamic structural symmetry breaking (DSSB) by showing that the obtained SSSB scheme explores the same tree as DSSB [8] whenever the variable and value orderings are fixed.

2 Basic Concepts

Definition 1 (CSP, Assignment, Solution). *A* constraint satisfaction problem (CSP) *is a triplet* $\langle V, D, C\rangle$*, where* V *denotes the set of variables,* D *denotes the set of possible values for these variables and is called their* domain*, and* $C : (V \to D) \to Bool$ *is a constraint that specifies which assignments of values to the variables are solutions. An* assignment *for a CSP* $\mathcal{P} = \langle V, D, C\rangle$ *is a function* $\alpha : V \to D$*. A* partial assignment *for a CSP* $\mathcal{P} = \langle V, D, C\rangle$ *is a function* $\alpha : W \to D$*, where* $W \subseteq V$*. The* scope *of* α*, denoted by* $scope(\alpha)$*, is* W*. A* solution *to a CSP* $\mathcal{P} = \langle V, D, C\rangle$ *is an assignment* σ *for* $\mathcal{P}$ *such that* $C(\sigma) = \mathbf{true}$*. The set of all solutions to a CSP* $\mathcal{P}$ *is denoted by* $Sol(\mathcal{P})$*.*

Definition 2 (Partition, Piecewise Bijection). *Given a set* S *and a set of sets* $P = \{P_1, \dots, P_n\}$ *such that* $S = \bigcup_i P_i$ *and the* P_i *are pairwise non-overlapping, we say that* P *is a* partition *of* S *and that each* P_i *is a* component*, and we write* $S = \sum_i P_i$*. A bijection* $b : S \to S$ *is a* piecewise bijection *over* $\sum_i P_i$ *iff* $\{b(e) \mid e \in P_i\} = P_i$*.*

Definition 3 (Piecewise Symmetric CSP). *A CSP* $\mathcal{P} = \langle \sum_k V_k, \sum_\ell D_\ell, C\rangle$ *is a* piecewise symmetric CSP *iff, for each solution* $\alpha \in Sol(\mathcal{P})$*, each piecewise bijection* τ *over* $\sum_\ell D_\ell$*, and each piecewise bijection* σ *over* $\sum_k V_k$*, we have* $\tau \circ \alpha \circ \sigma \in Sol(\mathcal{P})$*.*

Definition 4 (Dominance Detection). *Given two partial assignments* α *and* β *for a piecewise symmetric CSP* $\mathcal{P} = \langle \sum_k V_k, \sum_\ell D_\ell, C\rangle$*, we say that* α dominates β *iff there exist piecewise bijections* σ *over* $\sum_k V_k$ *and* τ *over* $\sum_\ell D_\ell$ *such that* $\alpha(v) = \tau \circ \beta \circ \sigma(v)$ *for all* $v \in scope(\alpha)$*.*

Dominance detection constitutes the core operation of symmetry breaking by dominance detection (SBDD) [1, 2], and its tractability immediately implies that we can efficiently limit ourselves to the exploration of symmetry-free search trees only. For piecewise symmetric CSPs, [8] showed that dominance detection is tractable.

3 Static SSB for Piecewise Symmetric CSPs

When we assume a total ordering of the variables $V = \{v_1, \dots, v_n\}$ and the values $D = \{d_1, \dots, d_m\}$, we can break the variable symmetries within each variable component as usual, by requiring that earlier variables take smaller or equal values. To break the value symmetries, we resort to structural abstractions, so-called *signatures*, which generalize from an exact assignment of values to variables by quantifying how often a given value is assigned to variables in each component. Let the frequency $f_h^k = |\{v_g \in V_k \mid v_g \in scope(\alpha) \;\;\&\;\; \alpha(v_g) = d_h\}|$ denote how often each value d_h is taken under partial assignment α by the variables in each variable component V_k. For a partial assignment α, we then denote by $\mathrm{sig}_\alpha(d_h) := (f_h^1, \dots, f_h^a)$ the *signature* of d_h under α. Then, for all consecutive values d_h, d_{h+1} in the same value component, we require that their signatures are lexicographically non-increasing, i.e., $\mathrm{sig}_\alpha(d_h) \geq_{lex} \mathrm{sig}_\alpha(d_{h+1})$. So the problem boils down to computing the signatures of values efficiently. Fortunately, this is an easy task when using the existing global cardinality constraint (gcc) [6]. We thus propose to add the following static set of constraints to a piecewise symmetric CSP $\langle \sum_{k=1}^{a} V_k, \sum_{\ell=1}^{b} D_\ell, C\rangle$ with $V_k = \{v_{i(k)}, \ldots, v_{i(k+1)-1}\}$ and $D_\ell = \{d_{j(\ell)}, \ldots, d_{j(\ell+1)-1}\}$:

$$\begin{array}{l}\forall\ 1 \le k \le a : \forall\ i(k) \le h < i(k+1)-1 : v_h \le v_{h+1} \\ \forall\ 1 \le k \le a : \mathrm{gcc}(v_{i(k)}, \ldots, v_{i(k+1)-1}, d_1, \ldots, d_m, f_1^k, \ldots, f_m^k) \\ \forall\ 1 \le \ell \le b : \forall\ j(\ell) \le h < j(\ell+1)-1 : (f_h^1, \ldots, f_h^a) \ge_{lex} (f_{h+1}^1, \ldots, f_{h+1}^a)\end{array}$$

where $i(k)$ denotes the index in $\{1, \ldots, n\}$ of the first variable in V of variable component V_k, with $i(a+1) = n+1$, and $j(\ell)$ denotes the index in $\{1, \ldots, m\}$ of the first value in D of value component D_ℓ, with $j(b+1) = m+1$.

Example. Consider scheduling study groups for two sets of five indistinguishable students each. There are six identical tables with four seats each. Let $\{v_1, \ldots, v_5\} + \{v_6, \ldots, v_{10}\}$ be the partitioned set of piecewise interchangeable variables, one for each student. Let the domain $\{t_1, \ldots, t_6\}$ denote the set of tables, which are fully interchangeable. The static structural symmetry-breaking constraints are:

$$\begin{array}{l}v_1 \le v_2 \le v_3 \le v_4 \le v_5,\ \ v_6 \le v_7 \le v_8 \le v_9 \le v_{10}, \\ \mathrm{gcc}(v_1, \ldots, v_5,\ t_1, \ldots, t_6,\ f_1^1, \ldots, f_6^1),\ \ \mathrm{gcc}(v_6, \ldots, v_{10},\ t_1, \ldots, t_6,\ f_1^2, \ldots, f_6^2), \\ (f_1^1, f_1^2) \ge_{lex} (f_2^1, f_2^2) \ge_{lex} \cdots \ge_{lex} (f_6^1, f_6^2)\end{array}$$

Consider the assignment $\alpha = \{v_1 \mapsto t_1, v_2 \mapsto t_1, v_3 \mapsto t_2, v_4 \mapsto t_2, v_5 \mapsto t_3\} \cup \{v_6 \mapsto t_1, v_7 \mapsto t_2, v_8 \mapsto t_3, v_9 \mapsto t_4, v_{10} \mapsto t_5\}$. Within each variable component, the $\le$ ordering constraints are satisfied. Having determined the frequencies using the gcc constraints, we observe that the $\ge_{lex}$ constraints are satisfied, because $(2,1) \ge_{lex} (2,1) \ge_{lex} (1,1) \ge_{lex} (0,1) \ge_{lex} (0,1) \ge_{lex} (0,0)$. If student 10 moves from table 5 to table 6, producing a symmetrically equivalent assignment because the tables are fully interchangeable, the $\ge_{lex}$ constraints are no longer satisfied, because $(2,1) \ge_{lex} (2,1) \ge_{lex} (1,1) \ge_{lex} (0,1) \ge_{lex} (0,0) \not\ge_{lex} (0,1)$.

Theorem 1. *For every solution α to a piecewise symmetric CSP, there exists exactly one symmetric solution that obeys the static structural symmetry-breaking constraints.*

Proof. (a) We show that there exists at least one symmetric solution that obeys all the symmetry-breaking constraints. Denote by $\tau_\alpha^\ell : \{j(\ell), \ldots, j(\ell+1)-1\} \to \{j(\ell), \ldots, j(\ell+1)-1\}$ the function that ranks the values in D_ℓ according to the signatures over some solution α, i.e., $\mathrm{sig}_\alpha(d_{\tau_\alpha^\ell(h)}) \ge \mathrm{sig}_\alpha(d_{\tau_\alpha^\ell(h+1)})$ for all $j(\ell) \le h < j(\ell+1)-1$. We obtain a symmetric solution β where we re-order the values in each D_ℓ according to τ_α^ℓ. Then, when we denote by $\sigma_\beta^k : \{i(k), \ldots, i(k+1)-1\} \to \{i(k), \ldots, i(k+1)-1\}$ the function that ranks the variables in V_k according to β, i.e., $\beta(v_{\sigma_\beta^k(h)}) \le \beta(v_{\sigma_\beta^k(h+1)})$ for all $i(k) \le h < i(k+1)-1$, we can re-order the variables in each V_k according to σ_β^k, and we get a new symmetric solution γ. Note that the re-ordering of the variables within each component has no effect on the signatures of the values, i.e., $\mathrm{sig}_\gamma(d) = \mathrm{sig}_\beta(d)$ for all $d \in D$. Thus, γ obeys all the symmetry-breaking constraints.

(b) Now assume there are two symmetric solutions α and β to the piecewise symmetric CSP that both obey all the symmetry-breaking constraints. Denote by τ^ℓ the re-ordering of the values in D_ℓ and denote by σ^k the re-ordering of the variables in V_k. Then, we denote by τ the piecewise bijection over the values based on the τ^ℓ, and by σ

the piecewise bijection over the variables based on the σ^k, such that $\alpha = \tau \circ \beta \circ \sigma$. The first thing to note is that the application of the piecewise bijection σ on the variables has no effect on the signatures of the values, i.e., $\mathrm{sig}_\beta(d) = \mathrm{sig}_{\beta\circ\sigma}(d)$ for all $d \in D$. Consequently, the total lexicographic ordering constraints on the signatures of each value d and its image $\tau(d)$ require that $\mathrm{sig}_\beta(d) = \mathrm{sig}_\beta(\tau(d)) = \mathrm{sig}_{\tau\circ\beta}(d) = \mathrm{sig}_{\tau\circ\beta\circ\sigma}(d) = \mathrm{sig}_\alpha(d)$. Thus, the signatures under α and β are identical. However, with the signatures of all the values fixed and with the ordering on the variables, there exists exactly one assignment that gives these signatures, so α and β must be identical. □

4 Static versus Dynamic SSB for Piecewise Symmetric CSPs

The advantage of a static symmetry-breaking method lies mainly in its ease of use and its moderate costs per search node. The number of constraints added is *linear* in the size of the problem, unlike the general method in [5], but they may break super-exponentially many variable and value symmetries. Constraint propagation and incrementality are inherited from the existing lex-ordering and gcc constraints. However, it is well-known that static symmetry breaking can collide with dynamic variable and value orderings, whereas dynamic methods such as SBDD do not suffer from this drawback.

Theorem 2. *Given static variable and value orderings, static (SSSB) and dynamic SSB (DSSB) explore identical search trees for piecewise symmetric CSPs.*

Proof. (a) Proof by contradiction. Assume there exists a node in the SSSB search tree that is pruned by DSSB. Without loss of generality, we may consider the first node in a depth-first search tree where this occurs. We identify this node with the assignment $\beta := \{v_1, \dots, v_t\} \to D$, and the node that dominates β is identified with the assignment $\alpha := \{v_1, \dots, v_s\} \to D$, for some $1 \le s \le t \le n$. By the definition of DSSB, we have that $\alpha(v_i) = \beta(v_i)$ for all $1 \le i < s$ (since every no-good considered by SBDD differs in exactly its last variable assignment from the current search node), and $\alpha(v_s) < \beta(v_s)$.

First consider $s = t$. Assume the dominance check between α and β is successful. Then, $\mathrm{sig}_\beta(\beta(v_s)) = \mathrm{sig}_\alpha(\alpha(v_s)) \gneqq \mathrm{sig}_\beta(\alpha(v_s))$. However, since $\alpha(v_s) < \beta(v_s)$, it must also hold that $\mathrm{sig}_\beta(\alpha(v_s)) \ge \mathrm{sig}_\beta(\beta(v_s))$. Contradiction.

Now consider $s < t$. Since the parent of β is not dominated by α, as β was chosen minimally, we know that v_t must be interchangeable with some v_p with $p \le s < t$. If we denote the component of v_t by $\{v_q, \dots, v_t, \dots, v_u\}$, we can deduce that $q \le s < t \le u$, i.e., v_s and v_t must belong to the same component. By definition of SSSB, we also know that $\alpha(v_q) \le \dots \le \alpha(v_s) < \beta(v_s) \le \dots \le \beta(v_t)$. Moreover, we know that $\beta(v_t)$ and $\alpha(v_p)$ must be interchangeable. Consequently, $\alpha(v_s)$ and $\beta(v_s)$ are also interchangeable. Now, since setting $\alpha(v_s)$ and $\beta(v_s)$ to v_s was not considered symmetric by DSSB, together with $\beta(v_s) > \alpha(v_s)$, we know that $\mathrm{sig}_\beta(\alpha(v_s)) \lneqq \mathrm{sig}_\alpha(\alpha(v_s))$. It follows that $\mathrm{sig}_\alpha(\alpha(v_s)) \gneqq \mathrm{sig}_\beta(\alpha(v_s)) \ge \mathrm{sig}_\beta(\beta(v_s))$ (1). When $\alpha(v_s)$ is matched with $\beta(v_i)$, for $q \le i \le t$, by the successful dominance check of α and β, then it must hold that $i < s$ as otherwise $\mathrm{sig}_\alpha(\alpha(v_s)) \le \mathrm{sig}_\beta(\beta(v_i)) \le \mathrm{sig}_\beta(\beta(v_s))$, which is in conflict with (1). This implies that $\beta(v_t)$ must be matched with some $\alpha(v_r)$ for $q \le r < s$ by the successful dominance check. Hence all the values in $\{\alpha(v_r), \dots, \alpha(v_s), \beta(v_s), \dots, \beta(v_t)\}$

are pairwise interchangeable. But then $\mathrm{sig}_\alpha(\alpha(v_r)) \leq \mathrm{sig}_\beta(\beta(v_t)) \leq \mathrm{sig}_\beta(\beta(v_s)) \lneqq \mathrm{sig}_\alpha(\alpha(v_s)) \leq \mathrm{sig}_\alpha(\alpha(v_r))$. Contradiction.

(b) Assume there exists a node in the DSSB search tree that is pruned by SSSB. Without loss of generality, we may consider the first node in a depth-first search tree where this occurs. We identify this node with the assignment $\beta := \{v_1, \ldots, v_t\} \to D$.

First assume a variable ordering constraint is violated, i.e., $\beta(v_j) > \beta(v_i)$ for some $1 \leq i < j \leq t$ where v_i and v_j are interchangeable. Consider $\alpha : \{v_1, \ldots, v_i\} \to D$ such that $\alpha(v_k) := \beta(v_k)$ for all $1 \leq k < i$, and $\alpha(v_i) := \beta(v_j)$. Then, due to the static variable and value orderings, α is a node that has been fully explored before β, and α dominates β, which is clear by mapping v_i to v_j. Thus, β is also pruned by DSSB.

Now assume a lex-ordering constraint on the value signatures is violated. Denote the interchangeable values by d_i and d_j, with $1 \leq i < j$. Since β was chosen minimally, when we denote the variable component that shows that $\mathrm{sig}_\beta(d_i) < \mathrm{sig}_\beta(d_j)$ by V_k, we know that $\mathrm{sig}_\beta(d_i)[\ell] = \mathrm{sig}_\beta(d_j)[\ell]$ for all $\ell < k$ and $\mathrm{sig}_\beta(d_i)[k] + 1 = \mathrm{sig}_\beta(d_j)[k]$. With $s := \max\{p \mid p < t \;\;\&\;\; \beta(v_p) = d_i\}$, we set $\alpha : \{v_1, \ldots, v_{s+1}\} \to D$ with $\alpha(v_r) := \beta(v_r)$ for all $r \leq s$ and $\alpha(v_{s+1} := d_i)$. Again, due to the static variable and value orderings, α is a node that has been fully explored before β, and α dominates β, which is clear simply by mapping d_i to d_j. Hence, β is also pruned by DSSB. □

We conclude that dynamic symmetry breaking draws its strength from its ability to accommodate dynamic variable and value orderings, but causes an unnecessary overhead when these orderings are fixed. In this case, static symmetry breaking offers a much more light-weight method that achieves exactly the same symmetry-breaking effectiveness for piecewise symmetric CSPs. Can we find general conditions under which a static symmetry-breaking method leads to symmetry-free search trees?

Acknowledgements: The authors are partly supported by grant IG2001-67 of STINT, the Swedish Foundation for Int'l Cooperation in Research and Higher Education.

References

1. T. Fahle, S. Schamberger, and M. Sellmann. Symmetry breaking. In T. Walsh, editor, *Proceedings of CP'01*, volume 2293 of *LNCS*, pages 93–107. Springer-Verlag, 2001.
2. F. Focacci and M. Milano. Global cut framework for removing symmetries. In T. Walsh, editor, *Proceedings of CP'01*, volume 2239 of *LNCS*, pages 77–92. Springer-Verlag, 2001.
3. I. P. Gent and B. M. Smith. Symmetry breaking during search in constraint programming. In *Proceedings of ECAI'00*, pages 599–603, 2000.
4. J.-F. Puget. On the satisfiability of symmetrical constrained satisfaction problems. In *Proceedings of ISMIS'93*, volume 689 of *LNAI*, pages 350–361. Springer-Verlag, 1993.
5. J.-F. Puget. An efficient way of breaking value symmetries. In *Proceedings of AAAI'06*. AAAI Press, 2006.
6. J.-C. Régin. Generalized arc-consistency for global cardinality constraint. In *Proceedings of AAAI'96*, pages 209–215. AAAI Press, 1996.
7. C. M. Roney-Dougal, I. P. Gent, T. Kelsey, and S. Linton. Tractable symmetry breaking using restricted search trees. In *Proceedings of ECAI'04*, pages 211–215. IOS Press, 2004.
8. M. Sellmann and P. Van Hentenryck. Structural symmetry breaking. In *Proceedings of IJCAI'05*. Morgan Kaufmann, 2005.
9. P. Van Hentenryck, P. Flener, J. Pearson, and M. Ågren. Tractable symmetry breaking for CSPs with interchangeable values. In *Proceedings of IJCAI'03*, Morgan Kaufmann, 2003.