

An Adaptive Model Restarts Heuristic

Nina Narodytska and Toby Walsh

NICTA and UNSW, Sydney, Australia

{Nina.Narodytska, Toby.Walsh}@nicta.com.au

Abstract. We propose an adaptive heuristic for model restarts that aligns symmetry breaking with the dynamic branching heuristic. Experiments show that this method performs very well compared to other symmetry breaking methods.

1 Introduction

Symmetry is an important but often problematic feature of constraint satisfaction problems. One way to deal with symmetry is to add constraints to eliminate symmetric solutions [1–7]. Posting static symmetry breaking constraints has both good and bad features. On the positive side, static constraints are easy to post, and a few simple constraints can eliminate most symmetry in a problem. On the negative side, static symmetry breaking constraints pick out particular solutions in each symmetry class, and this may conflict with the branching heuristic. An alternative is a dynamic approach that modifies the search method to ignore symmetric states [8–11]. Whilst this reduces the conflict with the branching heuristic, we may get less propagation. In particular there is no pruning of symmetric values deeper in the search tree. An effective method to tackle this tension is model restarts [12]. This restarts search frequently with new and different symmetry breaking constraints. The hope is that we will find symmetry breaking constraints that do not clash with the branching heuristic. The original model restarts method proposed a random choice of symmetry breaking constraints. We show here that we can improve performance with an adaptive heuristic that aligns symmetry breaking with the dynamic branching heuristic.

2 An Adaptive Model Restarts Heuristic

Our adaptive heuristic collects information about branching decisions in earlier restarts in order to build a heuristic friendly ordering of variables within the static symmetry breaking constraints. This ordering is based on variable scores. We describe three different techniques to obtain these scores. The first two reuse statistics collected by the branching heuristics. If we use the domain over weighted degree variable heuristic, then we can use the DOWD ratio to compute variable scores. $Score(X_k) = D(X_k)/(w \times deg(X_k))$ [13], where $D(X_k)$ is the domain size, $deg(X_k)$ is the number of constraints involving the variable, and w the sum of the counters associated with these constraints. We order variables in increasing order of their scores. We call this the DOWD-based heuristic. Similarly, we can use statistics associated with impact based branching heuristics to build variable scores [14]. $Score(X_k) = \sum_{v \in D(x)} (1 - impact(x, v))$ where

$impact(x, v)$ is the impact of a branching decision measured by the reduction of the search space induced when the decision was posted. We now order variables in decreasing order of their scores. We call this the IMPACT-based heuristic. Our third approach is based on the branching levels of variables. This offers some robustness to the choice of branching heuristic. For example, it could be used with some other branching heuristic than DOWD or IMPACT. If a variable is instantiated at level i then it gets a Borda type score of $n - i$, where n is the number of variables, and 0 otherwise. We order variables in decreasing order of the average Borda score over the last restart. We call this the ADAPT heuristic.

We use these three scoring heuristics within model restarts [12]. Model restarts was proposed to use a random variable ordering within symmetry breaking constraints in each restart. Frequent restarting ensures we eventually select a good representative symmetric solution that is aligned with the dynamic branching heuristic. Instead of using randomization, our adaptive heuristics build a variable ordering for symmetry breaking in each restart that is aligned with the branching heuristic. This variable ordering is a permutation of the original variables, and hence itself can be seen as a variable symmetry. As noted in [15], applying a symmetry to a (sound/complete) set of symmetry breaking constraints generates a new (sound/complete) set of symmetry breaking constraints. Thus, we can safely use this permutation to reorder the variables in the symmetry breaking constraints.

3 Experimental Results

We carried out experiments with 3 sets of commonly used benchmarks. We used Choco 2.1.2 on an Intel Core 8 CPU, 2.7 Ghz, 4Gb RAM with 1000 sec timeout. We branch with DOWD or IMPACT heuristics [13, 14].¹

The first set of benchmarks, DIMACS graph colouring problems was used in earlier studies of symmetry breaking for interchangeable values [4, 16]. Such problems are particularly suitable to a dynamic symmetry breaking labeling rule that avoids symmetric solutions (DYN) [10]. We compared four symmetry breaking methods, including DYN, the static symmetry breaking precedence constraint (PREC) [4, 16], model restarts and one modification of model restarts. We use the suffix ADAPT, DOWD and IMPACT to denote that variables are reordered in the symmetry breaking PRECEDENCE constraint based on the corresponding scores.

Model restarts constructs a random permutation of variables in the scope of the symmetry breaking PRECEDENCE constraint (MR). Our adapted method works in the following way. The search starts on the model without symmetry breaking constraints. Until the first restart, we collect statistic about the search tree. If we use the ADAPT heuristic, we store the information about variables that the solver branched on as described in Section 2. If we use DOWD or IMPACT heuristics then the solver accumulates statistics in weights and impact factors. On the first restart, we order variables based on their scores obtained from the heuristic. The scores are described in Section 2. Then we post the PRECEDENCE constraint and align variables in the scope of the constraint

¹ We would like to thank Charles Prud'homme for his help in implementing the model restarts technique.

Table 2. Graph coloring. The branching heuristic is impact based branching. The average #backtracks over ten runs with a random seed to initialize the branching heuristic. #vals is the number of values in a problem.

problem	#val	Domain over weighted degree branching						MR _{val} + DOWD ADAPT						PREC						Impact based branching						MR _{val} + IMPACT ADAPT					
		MR + DOWD		ADAPT		PREC		MR _{val} + DOWD		ADAPT		PREC		MR + IMPACT		ADAPT		PREC		MR _{val} + IMPACT		ADAPT		PREC							
		bt	vals	bt	vals	bt	vals	bt	vals	bt	vals	bt	vals	bt	vals	bt	vals	bt	vals	bt	vals	bt	vals	bt	vals						
SAT (Satisfiable instances)																															
queen8	8	909599	109529	426741	159069	115148	380833	229315	122541	99789	59922	387503	172e+06	82926	130842	130842	199e+04	78942													
DSJC125.1	5	504	7087	2551	994	928	2847	937	826	9902	40019	5977	1.08e+06	3518	6512	142994	3896														
school1	14	8037	5667	166212	6397	6669	-	2673	5609	231	7325	4	4	93	4	93															
school1	15	3299	3472	-	1192	3904	-	1987	3452	39	10283	10	10	125	10	10	127														
DSJC125.5	19	290634	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-			
DSJC125.5	20	826	2.29e+06	1e+06	1.05e+06	1.05e+06	1.05e+06	1.05e+06	1.05e+06	2324	1.89e+06	390780	14711	2.88e+06	3.1e+06	5.1e+06	0.6e+06														
DSJC250.5	35	572804	-	-	-	-	-	-	-	2324	1.89e+06	390780	14711	2.88e+06	3.1e+06	5.1e+06	0.6e+06														
DSJC250.5	36	2693	-	-	-	-	-	-	-	15455	1.34e+06	17	17	7277	25888	159019	4204														
DSJC250.5	37	8	85164	5	86	44	15	23	15	86	763464	925	9	1379	24	44															
queen9	10	8.87e+06	575929	2.45e+06	1.15e+06	763242	6.75e+06	3.75e+06	426910	589342	1.21e+06	1.02e+06	6.1e+06	279081	1.18e+06	5.7e+06	450170														
le450_15b	15	1.79e+06	11684	1.26e+06	81512	45835	111520	13011	32670	1.24e+06	-	297147	174565	317851	1.43e+06	679485	322694														
school1 nsh	14	134	2.17e+06	-	376	1423	-	483	1274	367	-	1	1	440	1	1	440														
school1 nsh	15	0	111659	0	0	0	0	0	0	2088	478604	0	0	437	0	0	305														
school1 nsh	16	0	170079	0	0	0	0	0	0	4404	-	0	0	3342	999	255628	739														
queen10	10	11	4.11e+06	-	2.66e+06	-	3.8e+07	-	1.89e+07	1.05e+07	7.62e+06	-	-	-	-	-	1.04e+07	7.1e+06	66e+07												
queen10	12	353154	10618	7400	3688	2120	1168	1427	1432	3007	6614	4640	3149	2481	2139	17143	2178														
TOTALS		15/16		13/16	14/16	13/16	13/16	15/16	16/16	16/16	11/16	15/16	14/16	15/16	16/16	16/16	16/16	16/16	16/16	16/16	15/16	14/16	15/16	16/16	16/16	16/16	16/16	16/16			
solved problems/total		16/18		10/18	7/18	15/18	14/18	8/18	17/18	17/18	17/18	17/18	8/18	7/18	8/18	17/18	9/18	8/18	17/18	8/18	17/18	7/18	8/18	17/18	9/18	8/18	17/18	17/18			
UNSAT (Unsatisfiable instances)																															
mcycle5	5	27100	47759	1.1e+06	727665	518428	1.03e+06	450168	423985	52560	30503	540891	459642	327131	500700	633030	271010														
R50_5gb	9	653	26311	150714	47840	10339	128823	12867	10848	7068	40257	174693	1.24e+06	16270	207598	1.37e+06	13847														
queen8	8	223229	337913	1.86e+07	4.87e+06	5.93e+04	1.93e+07	8.64e+06	0.6e+06	521777	425590	6.97e+06	1.5e+06	4.73e+06	7.6e+06	4.01e+06															
multisol.1	47	-	-	-	8022	83800	-	145618	494667	986138	-	-	-	101	-	-	101														
multisol.1	48	-	-	-	1634	-	-	1.06e+06	371767	-	-	-	-	101	-	-	101														
school1	13	5935	1645	-	14005	6040	-	1038	5910	211	903	-	-	125	-	-	125														
school1	13	6241	-	-	1915	-	-	1249	530532	13968	-	-	-	125	-	-	125														
fpso12.12	29	8552	771308	516916	140871	806771	148121	142850	161052	21137	806651	520652	232815	738818	557867	267675															
mcycle6g	6	7757	-	-	-	520442	-	1.29e+06	933692	5.93e+06	-	-	-	1.26e+06	-	-	563975														
DSJC125.5	12	68793	-	-	-	-	-	4.75e+06	64e+06	606155	-	-	-	-	-	-	8237														
DSJC250.5	14	760284	-	-	-	-	-	108	113	606155	-	-	-	-	-	-	8237														
le450_15b	13	3.64e+06	-	-	-	-	-	108	113	606155	-	-	-	-	-	-	8237														
school1 nsh	12	13	3458	-	521	136	321	133	43	2011	-	-	-	115	3.87e+06	-	115														
school1 nsh	13	28	330336	-	279	156	-	236	156	150	1.53e+06	-	-	1293	-	-	1293														
4-Fullins 4	7	2.96e+06	881513	2679	827	854747	2506	1019	1.83e+06	582695	300621	632518	424	1.37e+06	1.5e+06	424															
4-Fullins 4	5	11263	31267	9525	657	179	10679	554	173	118198	9011	15523	29506	1618	15523	31130	1833														
2-Fullins 4	3	1.23e+086	2.2e+07	-	-	-	-	-	-	475e+06	97e+06	4.35e+07	4.9e+07	5.6e+07	5e+07	4.42e+07	43e+07														
4-Insertions 3	10	361738	334880	1.04e+06	214	322	1.04e+06	-	334	5.37e+06	339937	1.02e+06	9.5e+07	395	974031	4.4e+07	395														
huck	12	1.15e+07	-	-	1133	1279	-	813	1450	8.01e+06	-	-	-	2422	-	-	3181														
homer	12	1.15e+07	-	-	1133	1279	-	813	1450	8.01e+06	-	-	-	2422	-	-	3181														
TOTALS		16/18		10/18	7/18	15/18	14/18	8/18	17/18	17/18	17/18	8/18	7/18	8/18	17/18	9/18	8/18	17/18	8/18	17/18	7/18	8/18	17/18	9/18	8/18	17/18	17/18	17/18			
solved problems/total		16/18		10/18	7/18	15/18	14/18	8/18	17/18	17/18	17/18	8/18	7/18	8/18	17/18	9/18	8/18	17/18	8/18	17/18	7/18	8/18	17/18	9/18	8/18	17/18	17/18	17/18			

with the obtained variable ordering. Statistics for the ADAPT heuristic are reset to zero after first and later restarts. Statistics for DOWD and IMPACT heuristics have built-in mechanisms to gradually forget previous decisions. Between the first and the second restarts we again collect search statistics. On the second and later restarts, we remove the PRECEDENCE constraint that we posted on the previous restart from the model and post a new PRECEDENCE constraint where the variables in its scope are aligned with the ordering obtained from the heuristic. We continue this procedure until we find a solution or timeout. Note that our adaptive approach can be applied to all problems where model restarts can be applied as we replace a random ordering of variables with one derived from heuristics.

MR + DOWD and MR + IMPACT use the DOWD and IMPACT heuristics, whilst MR + ADAPT uses our adaptive version of model restarts.

We also consider limiting the cost of symmetry breaking. In MR_{sh} , $MR_{sh} + ADAPT$, $MR_{sh} + DOWD$ and $MR_{sh} + IMPACT$, we shorten the PRECEDENCE constraint to the first $2m$ variables, where m is the number of values. The intuition behind this idea is based on an empirical observation that an instantiation of a relatively small number of variables in the scope of the PRECEDENCE constraint entails the constraint in most benchmarks. The value $2m$ was chosen based on statistical analysis of the benchmarks. We use a geometric restart policy with the base of 100 backtracks and a growth coefficient of 1.1. This ensure that restarts are rapid as in [12]. Tables 1–2 give average times and the number of backtracks for the DOWD and IMPACT branching heuristics over 10 runs. In addition, Table 1 shows the number of runs where a problem was solved. We also computed geometric means for these instances to reduce impact of outliers. However, as this gives the same picture of results and we have limited space, we do not include these results here. We removed instances solved by all methods in under 3 seconds and separated results for satisfiable and unsatisfiable.

Effect of the adaptive heuristic. By comparing PREC, MR and MR_{sh} with their adaptive counterparts, we see that our adaptive heuristic ADAPT dramatically improves performance on the majority of instances. For example, the adaptive heuristic helps solve 9 additional benchmarks if we compare MR and MR + ADAPT. The adaptive heuristic is especially useful on unsatisfiable instances. Note that many of these additionally solved benchmarks are easy once we remove conflict between the branching heuristic and static symmetry breaking. We observed that DOWD-based adaptive ordering also performs well. Unfortunately, the IMPACT ordering does not perform well on these benchmarks.

Effect of shortening. By comparing MR + ADAPT and $MR_{sh} + ADAPT$, as well as other models with their shortened counterparts, we see that shortening achieves much better performance. However, it slightly increases the number of backtracks in some cases. Shortening does not increase significantly the number of solved instances, or change substantially the search tree. However, it improves the efficiency of search. Overall, $MR_{sh} + ADAPT$ gives the best performance over all benchmarks among all symmetry breaking methods using the DOWD and IMPACT branching heuristic.

Our second and third case studies consider classes of problems on which model restarts has been shown to outperform other static and dynamic symmetry breaking methods [12]. We ran experiments with the “signature” based static symmetry breaking

constraints proposed for variable and value interchangeability in [11] and denoted here as GCC-based. We decomposed the GCC constraint into AMONG constraints so we can have access to the cardinality variables. Following [12] we only order partitions within the symmetry breaking constraints. We compute a score for all variables in each partition with respect to the used heuristic and sort the partitions according to these scores. Again, the main advantage of our approach is that instead of random ordering of partitions in model restarts we align them with branching heuristics.

We generated 20 problems of each size and averaged statistics over these problems. We report time to find an optimum solution and prove optimality. Note that all results are shown on instances that are solved by all techniques for at least 10 generated problems.

As in [4, 12], we tested on graph colouring and Concert Hall scheduling problems. In [12], the model restarts technique was shown to outperform other symmetry breaking methods on these benchmarks. Hence, we only compare our adaptive strategy with the simple static symmetry breaking constraints and the highly effective model restarts technique (GCC-based +MR). As previously, we biased the ordering of variables in the simple static symmetry breaking constraint to put large partitions first. Figure 1 (left part) shows the results for uniform and biased graph colouring problems with $q = 0.5$ using IMPACT branching heuristic. The results confirm that model restarts is better than static symmetry breaking. Our adaptive ordering of partitions significantly improves performance of model restarts. In particular, the ADAPT heuristic is more robust compared to the IMPACT heuristic.

For the Concert Hall Problem, we generated problems as in [4]. As it is important to put large partitions first, we assumed that any partition with size greater than 4 is a large partition (the maximum partition size is 8 in this setup). The number of halls is 12 or 14. Figure 1 (right part) shows the results for 14 halls. As can be seen from the graphs, using an adaptive heuristic to order partitions improves model restarts significantly using both DOWD or IMPACT branching heuristics. Moreover, ADAPT shows the best performance across all instances.

4 Other Related Work

Crawford *et al.* proposed a general method to break symmetry statically using lex-leader constraints [17]. Most static symmetry breaking constraints (including the PRECEDENCE constraints used here) can be derived from such constraints. Efficient algorithms have been developed to propagate many static symmetry breaking constraints (e.g. [21–24]). Lex-leader constraints pick out the lexicographically smallest solution in each symmetry class. However, this may conflict with the branching heuristic. A number of dynamic methods have been proposed to deal with this conflict. For example, *SBDS* posts lex-leader constraints dynamically during search [8]. Another dynamic method for breaking symmetry is *SBDD* [9]. This checks if a node of the search tree is symmetric to some previously explored node. *GAPLex* is a hybrid method that combines together static and dynamic symmetry breaking [25]. However, it is limited to dynamically posting lex-leader constraints, and to searching with a fixed variable ordering (which can be a considerable burden). *Dynamic Lex* is another hybrid method that dynamically posts static symmetry breaking constraints during search which works with dynamic variable




Fig. 1. Mean times for uniform (top left) and biased (bottom left) graph colouring benchmarks with $q=0.5$ using IMPACT based branching. Mean times for Concert Hall Problem benchmarks with 14 halls using DOWD (top right) and IMPACT based branching (bottom right).

ordering heuristics [26]. This method adds lex-leader constraints during search that are compatible with the current partial assignment. Hence the first solution found is not removed by symmetry breaking. However, unlike here, the method assumes that values are tried in a fixed order.

5 Conclusions

Static symmetry breaking constraints are often an easy and effective way to deal with symmetry in a constraint or optimisation problem. However, there can be a conflict between static symmetry breaking constraints and branching heuristics. To reduce this conflict, we propose a simple adaptive heuristic for model restarts. This orders variables within symmetry breaking constraints to align with the dynamic branching heuristic. Experimental results suggest that it is a very promising alternative between purely static and purely dynamic symmetry breaking methods. In particular, the results show that the proposed ADAPT heuristic works well across all benchmarks and two state-of-the-art branching heuristics. Our adaptive method thus appears to be more robust compared to the original model restarts algorithm.

References

1. Puget, J.F.: On the satisfiability of symmetrical constrained satisfaction problems. In: Komorowski, J., Raś, Z.W. (eds.) ISMIS 1993. LNCS (LNAI), vol. 689, pp. 350–361. Springer, Heidelberg (1993)
2. Shlyakhter, I.: Generating effective symmetry-breaking predicates for search problems. In: Proc. of Workshop on Theory and Applications of Satisfiability Testing, SAT 2001 (2001)

3. Flener, P., Frisch, A.M., Hnich, B., Kiziltan, Z., Miguel, I., Pearson, J., Walsh, T.: Breaking row and column symmetries in matrix models. In: Van Hentenryck, P. (ed.) CP 2002. LNCS, vol. 2470, pp. 462–477. Springer, Heidelberg (2002)
4. Law, Y., Lee, J.: Symmetry Breaking Constraints for Value Symmetries in Constraint Satisfaction. *Constraints* 11(2-3), 221–267 (2006)
5. Walsh, T.: Symmetry breaking using value precedence. In: Proc. of ECAI 2006, pp. 168–172 (2006)
6. Walsh, T.: General Symmetry Breaking Constraints. In: Benhamou, F. (ed.) CP 2006. LNCS, vol. 4204, pp. 650–664. Springer, Heidelberg (2006)
7. Walsh, T.: Breaking value symmetry. In: Proc. of the 23rd National Conf. on AI, pp. 1585–1588. AAAI (2008)
8. Gent, I., Smith, B.: Symmetry breaking in constraint programming. In: Proc. of ECAI 2000, pp. 599–603 (2000)
9. Fahle, T., Schamberger, S., Sellmann, M.: Symmetry breaking. In: Walsh, T. (ed.) CP 2001. LNCS, vol. 2239, pp. 93–107. Springer, Heidelberg (2001)
10. Hentenryck, P.V., Agren, M., Flener, P., Pearson, J.: Tractable symmetry breaking for CSPs with interchangeable values. In: Proc. of the 18th IJCAI (2003)
11. Flener, P., Pearson, J., Sellmann, M., Van Hentenryck, P.: Static and dynamic structural symmetry breaking. In: Benhamou, F. (ed.) CP 2006. LNCS, vol. 4204, pp. 695–699. Springer, Heidelberg (2006)
12. Heller, D., Panda, A., Sellmann, M., Yip, J.: Model restarts for structural symmetry breaking. In: Stuckey, P.J. (ed.) CP 2008. LNCS, vol. 5202, pp. 539–544. Springer, Heidelberg (2008)
13. ChocoTeam: Documentation. CHOCO is a java library for constraint satisfaction problems (CSP) and constraint programming (CP), <http://choco.mines-nantes.fr/>
14. Refalo, P.: Impact-based search strategies for constraint programming. In: Wallace, M. (ed.) CP 2004. LNCS, vol. 3258, pp. 557–571. Springer, Heidelberg (2004)
15. Katsirelos, G., Walsh, T.: Symmetries of symmetry breaking constraints. In: Proc. of ECAI 2010 (2010)
16. Walsh, T.: Breaking value symmetry. In: Bessière, C. (ed.) CP 2007. LNCS, vol. 4741, pp. 880–887. Springer, Heidelberg (2007)
17. Crawford, J., Luks, G., Ginsberg, M., Roy, A.: Symmetry breaking predicates for search problems. In: Proc. of the 5th Int. Conf. on Knowledge Representation and Reasoning (KR 1996), pp. 148–159 (1996)
18. Luby, M., Sinclair, A., Zuckerman, D.: Optimal speedup of Las Vegas algorithms. *Information Processing Letters* 47, 173–180 (1993)
19. Boussemart, F., Hemery, F., Lecoutre, C., Sais, L.: Boosting systematic search by weighting constraints. In: Proc. of the 16th ECAI 2004, pp. 146–150 (2004)
20. Gomes, C., Selman, B., Crato, N.: Heavy-tailed distributions in combinatorial search. In: Smolka, G. (ed.) CP 1997. LNCS, vol. 1330, pp. 121–135. Springer, Heidelberg (1997)
21. Frisch, A.M., Hnich, B., Kiziltan, Z., Miguel, I., Walsh, T.: Global constraints for lexicographic orderings. In: Van Hentenryck, P. (ed.) CP 2002. LNCS, vol. 2470, pp. 93–108. Springer, Heidelberg (2002)
22. Frisch, A., Hnich, B., Kiziltan, Z., Miguel, I., Walsh, T.: Propagation algorithms for lexicographic ordering constraints. *Artificial Intelligence* 170(10), 803–908 (2006)
23. Law, Y.C., Lee, J.H.M., Walsh, T., Yip, J.Y.K.: Breaking symmetry of interchangeable variables and values. In: Bessière, C. (ed.) CP 2007. LNCS, vol. 4741, pp. 423–437. Springer, Heidelberg (2007)

24. Katsirelos, G., Narodytska, N., Walsh, T.: Combining symmetry breaking and global constraints. In: Oddi, A., Fages, F., Rossi, F. (eds.) CSCLP 2008. LNCS, vol. 5655, pp. 84–98. Springer, Heidelberg (2009)
25. Jefferson, C., Kelsey, T., Linton, S., Petrie, K.: GAPLex: Generalised static symmetry breaking. In: Proc. of 6th Int. Workshop on Symmetry in Constraint Satisfaction Problems, Sym-Con 2006 (2006)
26. Puget, J.-F.: Symmetry breaking using stabilizers. In: Rossi, F. (ed.) CP 2003. LNCS, vol. 2833, pp. 585–599. Springer, Heidelberg (2003)