

A Portfolio Approach for Enforcing Minimality in a Tree Decomposition*

Daniel J. Geschwender^{1,2,†}, Robert J. Woodward^{1,2}, Berthe Y. Choueiry^{1,2,‡},
and Stephen D. Scott²

¹Constraint Systems Laboratory

²Department of Computer Science and Eng., University of Nebraska-Lincoln, USA

[†]student, [‡]advisor, {dgeschwe|rwoodwar|choueiry|sscott}@cse.unl.edu

Abstract. Minimality, a highly desirable consistency property of Constraint Satisfaction Problems (CSPs), is in general too expensive to enforce. Previous work has shown the practical benefits of restricting minimality to the clusters of a tree decomposition, allowing us to solve many difficult problems in a backtrack-free manner. We explore two alternative algorithms for enforcing minimality whose performance widely vary from one instance to another. We advocate a fine-grain portfolio approach to dynamically choose, during lookahead, the most appropriate algorithm for a cluster. Our strategy operates by selecting among two algorithms for enforcing minimality and an algorithm that enforces the lowest-level of consistency, which, in our setting, is Generalized Arc Consistency. Empirical evaluation on benchmark problems shows a significant improvement both in terms of the number of instances solved and CPU time.

1 Introduction

Local consistency techniques are at the heart of Constraint Programming and constitute an invaluable tool for solving Constraint Satisfaction Problems (CSPs). On many problems, enforcing simple consistency properties, such as Generalized Arc Consistency (GAC) [27], during backtrack search can be sufficient to reduce the problem to a manageable state. However, some problems are more resilient and require stronger consistency properties to effectively filter them.

In this paper, we focus on constraint minimality as one such consistency property. A constraint is considered minimal if every tuple of the constraint can be extended to a complete solution to the CSP. Enforcing constraint minimality is prohibitively expensive because it involves enumerating many, if not all, of the solutions to the problem. However, it can be applied locally (that is, to a given subproblem) with some success [13, 23]. Following Karakashian et al. [23], we consider enforcing the property on clusters of the tree decomposition of the problem. This restriction (or localization) to the clusters of a tree decomposition can result in a strong filtering power at a manageable cost.

* Experiments were conducted at the Holland Computing Center facility of the University of Nebraska. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1041000 and NSF Grant No. RI-111795.

Previous work has proposed two algorithms for enforcing minimality: PERTUPLE [24] and ALLSOL [22]. Our contribution is the development of a portfolio approach for choosing between these two algorithms, as well as identifying when to forego both algorithms and instead use an algorithm for GAC [27, 7]. Our empirical evaluation shows that such a portfolio can solve significantly more problem instances than GAC, ALLSOL, or PERTUPLE alone, and in less runtime on average.

The paper is organized as follows. Section 2 reviews some necessary background material. Section 3 discusses enforcing minimality on the clusters of a tree decomposition. Section 4 discusses the construction of the portfolio. Section 5 describes experimental evaluations of the portfolio on benchmark problems. Finally, Section 6 presents our conclusions.

2 Background

The *Constraint Satisfaction Problem* (CSP) is denoted by $\mathcal{P} = (\mathcal{X}, \mathcal{D}, \mathcal{C})$. $\mathcal{X} = \{x_1, \dots, x_n\}$ is a set of n variables, each associated with a finite domain from $\mathcal{D} = \{D_1, \dots, D_n\}$. $\mathcal{C} = \{C_1, \dots, C_e\}$ is the set of constraints restricting how values may be assigned to variables. Each constraint covers some subset of the variables, known as the scope of the constraint. A solution to a CSP is an assignment to each variable a value from its domain such that all constraints are satisfied. Deciding the existence of a solution is an NP-complete problem.

A constraint C_i is defined by relation R_i over the $scope(C_i)$. In this paper, we consider relations expressed as a set of allowed tuples. Each relation R_i is a subset of the Cartesian product of the domains of the variables in the $scope(C_i)$. Each tuple in the relations represent an assignment of values to the respective variables that is consistent with the constraint.

Several graphical representations of a CSP exist. The constraint network of a binary CSP is a graph where the vertices represent the variables and the edges the binary constraints. The constraint network of a non-binary CSP is a *hypergraph*. In the hypergraph, the vertices represent the variables and the hyperedges the scopes of the constraints. In the *primal graph*, the vertices represent the variables, and the edges connect every two variables that appear in the scope of some constraint. In the *dual graph*, the vertices represent the constraints of the CSP, and the edges connect vertices corresponding to constraints whose scopes overlap. Finally, the *incidence graph* of a CSP is a bipartite graph where one set contains all the variables and the other all the constraints. An edge connects a variable and constraint if the variable appears in the scope of the constraint. The incidence graph is the same graph used in the hidden-variable encoding [33].

A *tree decomposition* of a CSP is a tree embedding of its constraint network. The tree nodes are *clusters* of variables and constraints from the CSP. A tree decomposition must satisfy two conditions: *a)* each constraint appears in at least one cluster and the variables in its scope must appear in this cluster; and *b)* for every variable, the clusters where the variable appears induce a connected subtree. Many techniques for generating a tree decomposition of a CSP exist [11,

21, 17]. We use an adaption for non-binary CSPs of the tree-clustering technique [11]. First, we triangulate the primal graph of the CSP using the min-fill heuristic [26]. Then, we identify the maximal cliques in the resulting chordal graph using the MAXCLIQUES algorithm [14], and use the identified maximal cliques to form the clusters of the tree decomposition. We build the tree by connecting the clusters using the JOINTREE algorithm [9]. In order to enhance constraint propagation between adjacent clusters, we use the projection schema described by Karakashian et al. [23].

A CSP is *arc consistent* (AC) if every value has a supporting value in all neighboring variables. A similar property for non-binary constraints is *generalized arc consistency* (GAC) [27]. Our experiments use the GAC-2001 algorithm [7]. *Minimality* requires that any tuple that satisfies a constraint appears in at least one solution to the CSP [29].

The two algorithms, ALLSOL and PERTUPLE, proposed by Karakashian [22], both compute the minimal relations. PERTUPLE performs a backtrack search on every tuple in every relation, trying to consistently extend it a tuple in each other relation in the CSP. If the search fails, the tuple is removed. Otherwise, the search stops after finding the first solution. Further, the solution is used as a support structure for all the tuples that appear in the solution, which are marked as ‘minimal.’ In contrast, ALLSOL conducts a single backtrack search over the tuples of the relations, finding all the solutions and marking as ‘minimal’ every tuple that appears in any solution. If PERTUPLE is interrupted at any point, any deleted tuple is guaranteed to be inconsistent. However, when interrupted, the effort invested by ALLSOL is lost. Whereas the space used for storing support structures in PERTUPLE constitutes a tradeoff between time and space, ALLSOL does not incur such an overhead.

Related work on computing minimality includes: [5, 6, 16]; portfolio approaches: [32, 15, 37, 30, 20, 2]; adaptive consistency: [10, 12, 28, 34, 31, 3, 36, 4, 35].

3 Enforcing Minimality in a Tree Decomposition

Karakashian never compared the performance of PERTUPLE and ALLSOL during search, but only on individual clusters [22] collected from tree decompositions of CSP instances. When exploiting a tree decomposition for lookahead, Karakashian et al. exclusively used PERTUPLE [23]. We consider three variations of their basic process:

1. During lookahead, we include the option of whether or not to enforce GAC over the entire CSP after processing any given cluster.
2. Every time a cluster is considered for consistency, we can consistently call a specific consistency algorithm, or use a classifier to determine whether to call ALLSOL or PERTUPLE on the cluster or to do ‘Neither.’
3. Finally, we include an optional ‘timeout’ setting for processing individual clusters. This timeout interrupts the consistency algorithm currently operating on the cluster when the set threshold is reached. In the case of PERTUPLE, the filtering done so far is preserved. For ALLSOL, it is lost.

FILTERCLUSTERS (Algorithm 1) implements the above strategies and controls how consistency is enforced and propagated. In addition to the clusters, FILTERCLUSTERS takes three parameters that implement the above described variations of the process. Table 1 lists the parameter settings that yield six algorithms.

Algorithm 1: FILTERCLUSTERS(*clusterOrder*, *classifier*, *interleaveGAC*, *timeout*)

Input: *clusterOrder*, *classifier*, *interleaveGAC*, *timeout*
Output: Entire problem is GAC with potentially minimal clusters

```

1  didFiltering  $\leftarrow$  true
2  passDidFiltering  $\leftarrow$  true
3  consistent  $\leftarrow$  true
4  (consistent, didFiltering)  $\leftarrow$  GAC()
5  if consistent = false then return false
6  while passDidFiltering do
7    passDidFiltering  $\leftarrow$  false
8    foreach cluster  $\in$  clusterOrder do
9      algo  $\leftarrow$  CLASSIFY(cluster, classifier)
10     if algo = 'AllSol' then
11       (consistent, didFiltering)  $\leftarrow$  ALLSOL(cluster, timeout)
12     else if algo = 'PerTuple' then
13       (consistent, didFiltering)  $\leftarrow$  PERTUPLE(cluster, timeout)
14     else didFiltering  $\leftarrow$  false
15     if consistent = false then return false
16     if didFiltering then passDidFiltering  $\leftarrow$  true
17     if interleaveGAC and didFiltering then
18       (consistent, didFiltering)  $\leftarrow$  GAC()
19       if consistent = false then return false
20   clusterOrder  $\leftarrow$  REVERSE(clusterOrder)
21 if interleaveGAC = false then
22   (consistent, didFiltering)  $\leftarrow$  GAC()
23   if consistent = false then return false
24 return true

```

Table 1. Parameter variations of FILTERCLUSTERS

Algorithm	<i>classifier</i>	<i>interleaveGAC</i>	<i>timeout</i>
ALLSOL	Always select 'AllSol'	<i>false</i>	∞
PERTUPLE	Always select 'PerTuple'	<i>false</i>	∞
ALLSOL ⁺	Always select 'AllSol'	<i>true</i>	1 (s)
PERTUPLE ⁺	Always select 'PerTuple'	<i>true</i>	1 (s)
RANDOM	Randomly select 'AllSol', 'PerTuple', or 'Neither'	<i>true</i>	1 (s)
DECTREE	Decision tree selects 'AllSol', 'PerTuple', or 'Neither'	<i>true</i>	1 (s)

FILTERCLUSTERS filters both the domains of the variables and the tuples of the relations. It may be applied as a preprocessing step as well as a look-ahead procedure during search. The foreach loop (line 8) processes clusters from the specified *clusterOrder*. In our setting, this ordering corresponds to the MAX-CLIQUEs ordering of the clusters, but we will investigate other priority orderings in the future. The outer while-loop (line 6) iterates until no further filtering can be achieved. At each pass, the direction of the cluster ordering is reversed to facilitate propagation (line 20). The classifier allows the selection of the most appropriate algorithm on a cluster by cluster basis (line 9). The option to run GAC (i.e., *interleaveGAC* = *true*) allows 'easy and quick' filtering, which may trigger rapid and effective propagation throughout the problem. The *timeout* option, which specifies a time limit in seconds, ensures that excessive time is not wasted on a single cluster, allowing us to recover from classification errors.

4 Building a Portfolio

Our algorithm portfolio must decide which of the two minimality algorithms to enforce on a cluster given that the performance of the two algorithms vary widely. Because both algorithms enforce the same consistency, the portfolio must select the fastest algorithm based on features extracted from the cluster being processed or, when both algorithms are too costly, it must choose to run neither.

Inspired by features that appeared in the literature [25, 1], we identified a selection of 73 features that attempt to capture the constraint-network structure and the relation properties of a problem instance. The majority of the features that we collect are aggregations of many data points. In general, we aggregate using the mean, coefficient of variation, minimum, maximum, and entropy (*). For some features, we report total sum ([†]) or \log_{10} of mean ([‡]).

- *CSP parameters*: number of variables; number of relations; number of tuples per relation^{†*}; domain size^{†*}; arity of relations^{*}; tightness of relations^{*}; relational linkage^{‡*}.
- *Dual-graph parameters*: density; vertex degree^{*}; vertex eccentricity^{*}; vertex-clustering coefficient^{*}.
- *Incidence-graph parameters*: density; vertex degree^{*}; vertex eccentricity^{*}.
- *Primal-graph parameters*: density; vertex degree^{*}; vertex eccentricity^{*}; vertex-clustering coefficient^{*}.

In order to train a classifier for the portfolio, we collected a large data set of runtimes for both algorithms. We took instances from 175 benchmarks and broke them down into clusters of a tree decomposition. We then sampled 9362 individual instances from these clusters, either randomly selecting 70 clusters from each of the 175 benchmarks or taking all the clusters of a benchmark when it has less than 70 clusters. We ran both ALLSOL and PERTUPLE on every cluster, while recording (for each cluster) a set of features as well as the runtimes of the algorithms. Figure 1 shows the runtime distribution of the training instances. Although there are substantially more instances favoring PERTUPLE, ALLSOL does have its niche of instances on which it completes in up to two orders of magnitude faster.

After collecting data from 9362 individual clusters, we trained a decision tree using the J48 algorithm of the Weka machine learning software suite [18]. Using the 9362 clusters as a training set, we labeled each cluster ‘AllSol’ when ALLSOL was the fastest, ‘PerTuple’ when PERTUPLE was fastest, and ‘Neither’ when neither algorithm completed within ten minutes. We weighted our instances to increase the importance of instances with a large difference in runtimes, computing the weight of an instance i using Equation (1) where $allSol(i)$, $perTuple(i)$ are the CPU time on i of ALLSOL and PERTUPLE, respectively.

$$weight(i) = \left[\left| \log_{10} \left(\frac{allSol(i)}{perTuple(i)} \right) \right| \cdot \left| \log_{10} (|allSol(i) - perTuple(i)| + 0.01) \right| \right] \quad (1)$$

We designed this weighting scheme to emphasize instances where the execution times differ greatly, both in the ratio and in the difference of their values. We assigned a weight of 20 to the ‘Neither’ instances after empirical testing.

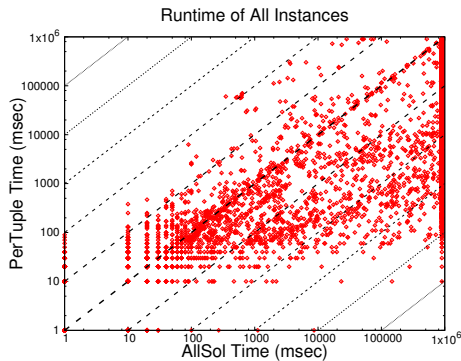


Fig. 1. Distribution of algorithm runtimes on single clusters

As a preliminary evaluation, we performed ten-fold cross validation using the collected instances. Our classifier achieved an unweighted accuracy of 80.1% and a weighted accuracy of 90.8%, which indicates the classifier is correctly handling the more heavily weighted instances.

5 Experimental Evaluation

In our evaluation, we use the six algorithms obtained by setting the parameters of FILTERCLUSTERS alongside GAC for real-full lookahead [19] in a backtrack search on a set of 1055 instances taken from 42 benchmarks from the XCSP library.¹ Our search procedure terminates after finding the first solution and uses the dom/deg *dynamic* ordering heuristic. Although dom/wdeg [8] can improve performance across the board, FILTERCLUSTERS is not yet equipped to take advantage of this heuristic. Our experiments run on a cluster computer with Intel Xeon E5-2650 v3 2.30GHz processors. Search is allowed to run for two hours (7200 sec) and given 12 GB of memory. To account for load variations on the cluster computer, we measure instruction count and convert it to runtime using a standardized measure of instructions per cycle and clock speed. We use a timeout of 1 second per-cluster because, based on the data from the 9362 clusters shown in Figure 1, this value strikes a good balance between completing clusters and not spending excessive time on any one cluster.

Table 2 summarizes the results (top) and provides detailed results per benchmark (bottom). We place the solvers into two categories. On one hand, the basic solvers, which include GAC, ALLSOL, and PERTUPLE. On the other hand, hybrid solvers, which include ALLSOL⁺, PERTUPLE⁺, RANDOM, and DECTREE. We compute the average CPU time only over instances completed by at least one of the solvers. Both timeouts and memouts (memory out) are considered 7200 seconds. Due to the randomness of RANDOM, we perform ten runs for each instance and report the median.

¹ <http://www.cril.univ-artois.fr/~lecoutre/benchmarks.html>

Overall, it is clear that DECTREE outperforms all solvers both in terms of the number of completed instances, and average and sum CPU time. It solves instances $2.2\times$ faster than GAC on average, and completes 135 more instances than GAC out of the 1055 tested.

RANDOM is surprisingly competitive with DECTREE. This fact is largely due to the stabilizing effect of the per-cluster timeout, which minimizes the time loss from poor classification decisions. We ran an experiment to assess the extent of this effect. We compared the performance of RANDOM and DECTREE with no per-cluster timeout. RANDOM completes only 484 instances whereas DECTREE completes 649, with an average CPU time across all instances completed by at least one solver of 2,955.9 seconds and 1,413.9 seconds, respectively. Thus, DECTREE makes substantially better decisions than a random choice.

The lower sections of Table 2 break down the performance of the seven solvers by benchmark. We identify three categories: benchmarks where the hybrid solvers outperform all others (top), those where hybrid and basic solvers perform equally well (middle), and finally those on which the basic solvers perform best (bottom). For each benchmark, we format in bold the smallest average runtime and all runtimes within 1 second or 5% of the best time.

In the top category, DECTREE and RANDOM are generally the best, but are outperformed by PERTUPLE⁺ on two benchmarks. The benchmarks in the middle category seem to be solvable relatively fast by most solvers and have few timeouts (except GAC). The basic solvers outperform the others on the benchmarks in the bottom-most category. Those benchmarks tend to be memory intensive: indeed PERTUPLE, PERTUPLE⁺, and DECTREE have many memouts.

6 Conclusions

We advocate a portfolio method for enforcing constraint minimality on the clusters of a tree decomposition, making minimality even more beneficial in practice by selectively applying it during problem solving. We provide three improvements in the application of constraint minimality: a classifier for choosing when to run ALLSOL, PERTUPLE, or neither, the use of GAC prior to every cluster being processed, and a timeout mechanism to prevent getting stuck on a single cluster. Our approach yields more problem completions and faster runtimes than lookahead with a simple GAC, PERTUPLE, or ALLSOL.

As a continuation of our approach, we plan to use a classifier that estimates the runtime of a consistency algorithm and the amount of filtering it could achieve in order to dynamically set-up the timeout threshold. Such a classifier would allow us to allocate more time for running the algorithm when significant filtering can be expected and less time when there is little prospect for filtering. We also want to extend our approach to estimating the memory overhead for running a given consistency algorithm and for selecting the most appropriate bolstering schema at the separators [23].

References

1. Roberto Amadini, Maurizio Gabbrielli, and Jacopo Mauro. An Enhanced Features Extractor for a Portfolio of Constraint Solvers. In *Proc. of ACM SAC 2014*, pages 1357–1359. ACM, 2014.
2. Roberto Amadini, Maurizio Gabbrielli, and Jacopo Mauro. A Multicore Tool for Constraint Solving. In *Proc. of IJCAI 2015*, pages 232–238. AAAI Press, 2015.
3. Amine Balafrej, Christian Bessiere, Remi Coletta, and El Houssine Bouyahf. Adaptive Parameterized Consistency. In *Proc. of CP 2013*, pages 143–158, 2013.
4. Amine Balafrej, Christian Bessiere, and Anastasia Paparrizou. Multi-Armed Bandits for Adaptive Constraint Propagation. In *Proc. of IJCAI 2015*, 2015.
5. Kenneth M. Bayer, Martin Michalowski, Berthe Y. Choueiry, and Craig A. Knoblock. Reformulating CSPs for Scalability with Application to Geospatial Reasoning. In *Proc. of CP 2007*, volume 4741 of *LNCS*, pages 164–179. Springer, 2007.
6. Christian Bessiere, H el ene Fargier, and Christophe Lecoutre. Global Inverse Consistency for Interactive Constraint Satisfaction. In *Proc. of CP 2013*, volume 8124 of *LNCS*, pages 159–174, 2013.
7. Christian Bessiere, Jean-Charles R egin, Roland H.C. Yap, and Yuanlin Zhang. An Optimal Coarse-Grained Arc Consistency Algorithm. *Artificial Intelligence*, 165(2):165–185, 2005.
8. Fr ed eric Boussemart, Fred Hemery, Christophe Lecoutre, and Lakhdar Sais. Boosting Systematic Search by Weighting Constraints. In *Proc. of ECAI 2004*, pages 146–150, 2004.
9. Rina Dechter. *Constraint Processing*. Morgan Kaufmann, 2003.
10. Rina Dechter and Judea Pearl. Network-Based Heuristics for Constraint-Satisfaction Problems. *Artificial Intelligence*, 34:1–38, 1988.
11. Rina Dechter and Judea Pearl. Tree Clustering for Constraint Networks. *Artificial Intelligence*, 38:353–366, 1989.
12. Susan Epstein, Richard Wallace, Eugene Freuder, and Xingjian Li. Learning Propagation Policies. In *Proc. of the Second International Workshop on Constraint Propagation and Implementation*, pages 1–15, 2005.
13. Eugene C. Freuder and Charles D. Elfe. Neighborhood Inverse Consistency Preprocessing. In *Proc. of AAAI 1996*, pages 202–208, 1996.
14. Martin C. Golumbic. *Algorithmic Graph Theory and Perfect Graphs*. Academic Press Inc., 1980.
15. Carla P. Gomes and Bart Selman. Algorithm Portfolios. *Artificial Intelligence*, 126(1):43–62, 2001.
16. Georg Gottlob. On Minimal Constraint Networks. In *Proc. of CP 2011*, volume 6876 of *LNCS*, pages 325–339. Springer, 2011.
17. Georg Gottlob, Nicola Leone, and Francesco Scarcello. A Comparison of Structural CSP Decomposition Methods. In *Proc. of IJCAI 1999*, pages 394–399, 1999.
18. Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H. Witten. The WEKA Data Mining Software: An Update. *ACM SIGKDD Explorations Newsletter*, 11(1):10–18, 2009.
19. Robert M. Haralick and Gordon L. Elliott. Increasing Tree Search Efficiency for Constraint Satisfaction Problems. *Artificial Intelligence*, 14:263–313, 1980.
20. Barry Hurley, Lars Kotthoff, Yuri Malitsky, and Barry O’Sullivan. Proteus: A Hierarchical Portfolio of Solvers and Transformations. In *Proc. of CPAIOR 2014*, pages 301–317, 2014.

21. Peter G. Jeavons, David A. Cohen, and Marc Gyssens. A Structural Decomposition for Hypergraphs. *Contemporary Mathematics*, 178:161–177, 1994.
22. Shant Karakashian. *Practical Tractability of CSPs by Higher Level Consistency and Tree Decomposition*. PhD thesis, CSE, UNL, Lincoln, NE, May 2013.
23. Shant Karakashian, Robert Woodward, and Berthe Y. Choueiry. Improving the Performance of Consistency Algorithms by Localizing and Bolstering Propagation in a Tree Decomposition. In *Proc. of AAAI 2013*, pages 466–473, 2013.
24. Shant Karakashian, Robert Woodward, Christopher Reeson, Berthe Y. Choueiry, and Christian Bessiere. A First Practical Algorithm for High Levels of Relational Consistency. In *Proc. of AAAI 2010*, pages 101–107, 2010.
25. Shant Karakashian, Robert J. Woodward, Berthe Y. Choueiry, and Stephen D. Scott. Algorithms for the Minimal Network of a CSP and a Classifier for Choosing Between Them. Technical Report TR-UNL-CSE-2012-0007, CSE, University of Nebraska-Lincoln, Lincoln, NE, 2012.
26. U. Kjærulff. Triangulation of Graphs - Algorithms Giving Small Total State Space. Research Report R-90-09, Aalborg University, Denmark, 1990.
27. Alan K. Mackworth. On Reading Sketch Maps. In *Proc. of IJCAI 1977*, pages 598–606, 1977.
28. Deepak Mehta and Marc R.C. van Dongen. Probabilistic Consistency Boosts MAC and SAC. In *Proc. of IJCAI 2007*, pages 143–148, 2007.
29. Ugo Montanari. Networks of Constraints: Fundamental Properties and Application to Picture Processing. *Information Sciences*, 7:95–132, 1974.
30. Eoin O’Mahony, Emmanuel Hebrard, Alan Holland, Conor Nugent, and Barry O’Sullivan. Using Case-based Reasoning in an Algorithm Portfolio for Constraint Solving. In *Proc. of the Irish Conference on AI and Cognitive Science*, pages 210–216, 2008.
31. Anastasia Paparrizou and Kostas Stergiou. Evaluating Simple Fully Automated Heuristics for Adaptive Constraint Propagation. In *Proc. of IEEE ICTAI 2012*, pages 880–885, 2012.
32. John R. Rice. The Algorithm Selection Problem. *Advances in Computers*, 15:65–118, 1976.
33. Francesca Rossi, Charles Petrie, and Vasant Dhar. On the Equivalence of Constraint Satisfaction Problems. In *Proc. of ECAI 1990*, pages 550–556, 1990.
34. Kostas Stergiou. Heuristics for Dynamically Adapting Propagation in Constraint Satisfaction Problems. *AI Communications*, 22(3):125–141, 2009.
35. Robert J. Woodward, Shant Karakashian, Berthe Y. Choueiry, and Christian Bessiere. Solving Difficult CSPs with Relational Neighborhood Inverse Consistency. In *Proc. of AAAI 2011*, pages 112–119, 2011.
36. Robert J. Woodward, Anthony Schneider, Berthe Y. Choueiry, and Christian Bessiere. Adaptive Parameterized Consistency for Non-Binary CSPs by Counting Supports. In *Proc. of CP 2014*, volume 8656 of *Lecture Notes in Computer Science*, pages 755–764, 2014.
37. Lin Xu, Frank Hutter, Holger H. Hoos, and Kevin Leyton-Brown. SATzilla: Portfolio-based Algorithm Selection for SAT. *JAIR*, 32(1):565–606, 2008.