Project no.:	**PIRSES-GA-2011-295261**
Project full title:	**Mobility between Europe and Argentina applying Logics to Systems**
Project Acronym:	**MEALS**
Deliverable no.:	**2.2 / 1**
Title of Deliverable:	**SyMT: finding symmetries in SMT formulas**

Contractual Date of Delivery to the CEC:	**30-Sep-2013**
Actual Date of Delivery to the CEC:	**30-Sep-2013**
Organisation name of lead contractor for this deliverable:	**UBA**
Author(s):	**Carlos Areces, David Déharbe, Pascal Fontaine, Ezequiel Orbe**
Participants(s):	**UNC, UBA, INR, SAU**
Work package contributing to the deliverable:	**WP2**
Nature:	**R**
Dissemination Level:	**Public**
Total number of pages:	**10**
Start date of project:	1 Oct. 2011 **Duration:** 48 month

Abstract:

The QF_UF category of the SMT-LIB test set contains many formulas with symmetries, and breaking these symmetries results in an important speedup. We here propose SyMT, a simple tool based on graph automorphism detection algorithms to find out symmetries in SMT formulas. SyMT helps SMT users by highlighting the symmetries in their formulas, giving thus hints on how they can improve them to enforce the SMT solver to examine one path out of many symmetric ones in the search tree. The classic propositional symmetry breaking technique can be lifted to SMT and yield a generic technique to break the symmetries found by SyMT.
Experiments on a large part of the SMT-LIB show that symmetries are pervasive in most categories.

Note:

This deliverable is based on material that has been published in the Proceedings of the 11th International Workshop on Satisfiability Modulo Theories, Helsinki, Finland, July 2013.

This project has received funding from the European Union Seventh Framework Programme (FP7 2007-2013) under Grant Agreement Nr. 295261.

Contents

1 Introduction

Consider a propositional formula $\varphi(p, q)$ with propositional variables p and q, and symmetric by permutation of p and q. Propositional symmetry breaking [12] eliminates symmetry, e.g. by adding clause $p \Rightarrow q$, since there is a Boolean model of $\varphi(p, q)$ if and only if there is a model such that $p \Rightarrow q$. It is not necessary to search for models such that $p \wedge \neg q$, and so the search space can be reduced. Now consider the first-order formula $\varphi(f(a) = f(b), a = b)$ with the standard interpretation of equality. It is clear that there exists no model such that $f(a) \neq f(b) \wedge a = b$ holds, although $f(a) = f(b) \wedge a \neq b$ is satisfiable; if $\varphi(p, q)$ has only models such that exactly one proposition in $\{p, q\}$ is true, $\varphi(f(a) = f(b), a = b) \wedge (f(a) = f(b) \Rightarrow a = b)$ is unsatisfiable. This simple example shows it is not sound to break symmetry of an SMT[1] formula based on the symmetry of its propositional structure alone. Essentially the problem is that the abstraction does not take the theory into account. However, we show in the paper that it is sound to break symmetries stemming from permutation of uninterpreted symbols, similarly to what is done for propositional logic.

As previous results suggest [8], exploiting symmetries in SMT formulas can lead to an impressive decrease in the size of the search space, and thus to a considerable increase in efficiency. Techniques described in [8] are, however, highly heuristic and vulnerable to formula rewriting. Graph automorphism detection algorithms [11, 9, 10] have been used to find symmetries in propositional formulas. We provide a tool, based on those techniques, to discover symmetries in SMT formulas. The tool does not break the symmetries automatically though. There would indeed exist many heuristic choices for symmetry breaking and the SMT user is in the best position to make the right choices, based on the application.

Outline. We first give a formal basis for symmetry breaking in SMT, then present the implementation of SyMT, our tool for detecting symmetries in SMT formulas. Some statistics on symmetry detection on a large part of the SMT-LIB [5] are given. They clearly show that (1) graph automorphism algorithms scale for SMT formulas, and (2) the SMT-LIB contain many highly symmetric formulas.

2 Symmetries in SMT

We assume knowledge of basic notions of permutation group theory, such as generator and cyclic form. We use the standard notions of multi-sorted logic, term, formula, and interpretation commonly used in the context of SMT. A theory is a set of interpretations. Consider a finite set $\mathcal{S}$ of uninterpreted symbols (constants, functions or predicates), and a bijective function σ on $\mathcal{S}$, that maps every symbol to a symbol of the same sort (i.e., arity and sorts of arguments and image should match). Function σ extends naturally to terms and formulas, and $t\sigma$ denotes σ applied to term or formula t, just like a higher-order substitution would, considering symbols in $\mathcal{S}$ as

[1] SMT stands for Satisfiability Modulo Theories, see [4] for a thorough introduction.

variables. σ can also be applied on an interpretation $\mathcal{I}$ to yield interpretation $\mathcal{I}\sigma$ similar to $\mathcal{I}$ except that $\mathcal{I}\sigma[s'] = \mathcal{I}[s]$ whenever $s\sigma = s'$. The identity function is denoted σ_I.

We say that σ is a symmetry for formula φ if $\varphi\sigma$ is syntactically equal to φ up to satisfiability preserving rewritings, e.g. using commutativity of some interpreted symbols. Notice that if σ is a symmetry for φ, so is any of its powers σ^i, and in particular σ^{-1} is also a symmetry of φ since there exists n such that $\sigma^n = \sigma_I$. The case where σ is its own inverse ($\sigma^2 = \sigma_I$) is a particular, though extremely frequent, case. It occurs when there is a group that contains all permutations of elements in a subset of $\mathcal{S}$. In our experiments on the SMT-LIB test bed, we have observed that most symmetry groups found have a set of generators that are their own inverse. Consider a symmetry σ such that $\sigma^2 = \sigma_I$ for a formula φ. For every interpretation $\mathcal{I}$ of φ we have $\mathcal{I}\sigma[\varphi] = \mathcal{I}[\varphi]$ (using straightforward structural induction). Consider now a set of atoms (not necessarily simple propositional variables) $p_1, \ldots p_n$ and their image $q_1 = p_1\sigma, \ldots q_n = p_n\sigma$. If φ is satisfiable in a model $\mathcal{M}$ then there exists a model of φ that furthermore satisfies the following formulas for $i \in \{1..n\}$:

$$\psi_i =_{\text{def}} \Big(\bigwedge_{1 \le j < i} p_j \equiv q_j\Big) \Rightarrow (p_i \Rightarrow q_i).$$

This model is indeed either $\mathcal{M}$ or $\mathcal{M}\sigma$. Assume k is the smallest value for which $\mathcal{M}[p_k] \neq \mathcal{M}[q_k]$, and consider ψ_k. If $\mathcal{M}[p_k] = \bot$ and $\mathcal{M}[q_k] = \top$ then $\mathcal{M}$ satisfies ψ_k, as well as all ψ_i with $i \neq k$. Now, if $\mathcal{M}[p_k] = \top$ and $\mathcal{M}[q_k] = \bot$ then $\mathcal{M}\sigma$ is a model of φ such that $\mathcal{M}\sigma[p_i] = \mathcal{M}\sigma[q_i]$ for $i < k$ and $\mathcal{M}\sigma[p_k] = \top$ and $\mathcal{M}\sigma[q_k] = \bot$. The model $\mathcal{M}\sigma$ of φ thus satisfies ψ_i for $i \in \{1..n\}$.

It is well known (see, e.g., [12]) that the formulas ψ_i can serve to break symmetry for propositional formulas. The above shows that this extends to SMT. This leaves out, however, many choices for the set of atoms p_i: the insight of the SMT user is usually necessary to make the best choice.

3 SyMT Implementation

SyMT is a command line tool implemented in C that detects symmetries in SMT formulas, taking into account the commutativity of conjunction, disjunction, addition, multiplication and equality. Given an input SMT formula, SyMT proceeds by creating a colored graph from it and then uses a graph automorphism component to detect the generators of the automorphism group of the colored graph. In particular, SyMT uses Saucy 3.0 [10] as the graph automorphism component. Integration with Saucy is done via Saucy's C API. SyMT also provides simplification capabilities on the input formulas, some of which involve using theory reasoning (and thus may unfortunately fail on large instances). Simplification of the input formula is important because it may uncover hidden symmetries and remove trivial symmetries, e.g., symmetries that do not involve uninterpreted symbols.

Example 1. Hereunder is the command line and output of SyMT on a formula of the QF_UF category of SMT-LIB:

```
./SyMT --enable-simp smt-lib2/QF_UF/NEQ/NEQ004_size4.smt2
(p7 p9)(c12 c13)
(c_3 c_1)
(c_2 c_1)
(c_0 c_1)
```

SyMT finds four generators for the symmetry group, and prints them in cyclic form. There is the full group of permutations of `c_0`, `c_1`, `c_2`, `c_3`, generated by the last three generators, as well as a further symmetry that permutes unary predicates `p7` and `p9`, while in the same time permutes `c12` and `c13`. This last symmetry was not detected with the heuristic techniques of [8].

Reduction to the colored graph automorphism problem is the most successful technique for detecting symmetries in propositional formulas in clausal form, primarily due to the availability of efficient tools to detect graph automorphisms (e.g., [11, 9, 10]) that are fast and easy to integrate. Several reductions from propositional formulas to colored graphs have been proposed [6, 7, 1], all based on the same idea: to use the formula to construct a colored graph whose automorphism group is isomorphic to the symmetry group of the formula. Also, extensions to other logics, e.g., QBF [3] and modal logics [2], have been proposed, further showing the applicability of this technique.

We now present the reduction algorithm to colored graphs for SMT formulas. The reduction is as a two-stage process. First, SyMT constructs the syntax direct acyclic graph of the formula with some additional nodes. Second, colors are introduced, to avoid spurious symmetries. Colors are represented as natural numbers. Let φ be an SMT formula. The colored graph $G(\varphi)$ is constructed recursively as follows (= and other predicates, and propositional symbols are considered as functions and constants ranging over Booleans):

- **Graph Construction:**

  1. For each symbol, add a unique *symbol* node.
  2. For each (constant or propositional) term without argument, the *root* node is the symbol node introduced above.
  3. For each term $f(t_1, \ldots, t_n)$ of arity $n > 0$,
     (a) Add a *root* node for $f(t_1, \ldots, t_n)$. Add an edge from the root node to the (unique) symbol node for f.
     (b) If the function is commutative (e.g. $\wedge$, $\vee$, $\equiv$, $=$, $+$, $*$), add an edge from the root node to the root node of t_i ($i \in \{1..n\}$). Quantifiers, as commutative operators, are handled similarly (coloring discriminates the matrix).
     (c) If the function is not commutative:
        i. For each argument t_i, add an *argument node* and an edge from this node to the root node of t_i.
        ii. Add an edge from the argument node of t_i to the argument node of t_{i+1} ($1 \leq i < n$). These edges represent the ordering of the arguments in $f(t_1, \ldots, t_n)$.
        iii. Add an edge from the root node to the argument node of t_1.

- **Graph Coloring:**

  1. Argument nodes are assigned a specific, unique color.
  2. Uninterpreted symbol nodes and root nodes are assigned a color based on their sort (Boolean being considered as any other sort).
  3. Each interpreted symbol node is assigned a unique color.

Example 2. Consider formula $\varphi = p(f(a,b)) \vee p(f(b,a)) \vee p(g(a,b)) \vee p(g(b,a))$, where p is a unary predicate and f, g, a and b are uninterpreted symbols. The associated colored graph, $G(\varphi)$, is shown in Figure 1 (colors are represented by numeric labels and node shapes in the figure).

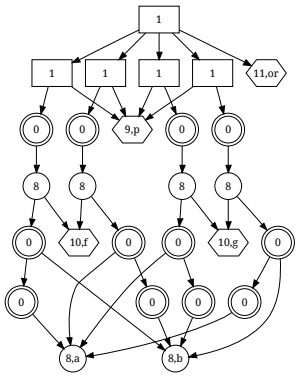

Generator 1: $(a\ b)$
Generator 2: $(f\ g)$

Figure 1: Graph representation of φ.

4 Symmetries in SMT-LIB

We test SyMT against 19 categories[2] from SMT-LIB [5] to investigate the existence of symmetries and evaluate the efficiency of our tool. All tests are run on an Intel Xeon X3440 with 16GB, using the four cores simultaneously and we report the cumulative core time (roughly 4 times the CPU time). Three different configurations of SyMT were tested. Configuration 1 has no simplification: the formula is parsed and converted to a graph for automorphism detection. Configuration 2 uses trivial syntactic simplifications. Configuration 3 enables stronger simplifications, using an SMT engine, e.g., simplification of atoms implied by unit clauses. Configuration 2 may fail (with no symmetry reported) because the simplification algorithm used is not linear with respect to the input formula. However it often reveals symmetries hidden by irrelevant garbage easily removed by the simplification procedure. Configuration 3 is likely to fail on very large formulas, but again, it may reveal hidden symmetries. Simplification sometimes reduces a formula to false, in which case no symmetry is reported.

[2]Bit vectors are not supported by our parser.

Category	#Inst	#Sym[1]	#Sym[2]	#Sym[3]	#Sym[P]	Avg[GS]	Time
AUFLIA	6480	6212	6231	5796	6251	134.00	347.14
AUFLIRA	19917	15779	16475	12046	16476	1.08	6.65
AUFNIRA	989	985	985	902	985	1.00	0.33
QF_AUFLIA	1140	534	603	91	613	1.19	0.58
QF_AX	551	280	280	22	280	1.15	0.35
QF_IDL	1749	346	658	747	840	12750.60	60.18
QF_LIA	5938	715	1165	475	1185	110.95	97.04
QF_LRA	634	99	176	212	247	40.46	2.52
QF_NIA	530	167	169	168	169	5.98	2.64
QF_NRA	166	9	43	43	43	1.00	0.19
QF_RDL	255	41	41	62	62	180.20	7.61
QF_UF	6647	250	544	543	544	83.47	26.87
QF_UFIDL	431	32	200	198	225	1.19	1.95
QF_UFLIA	564	0	198	198	198	0.00	0.36
UFNIA	1796	1062	1061	1048	1062	47.08	471.02

Table 1: Symmetries in SMT-LIB

Among the 19 analyzed categories, three (LRA, QF_UFLRA, QF_UFNRA) do not reveal symmetries with SyMT. Of the only five formulas in UFLRA, one has symmetries. The others 14 categories presented numerous symmetries in at least one of the tested configurations. Table 1 summarizes the results obtained for these 14 categories. For each category we report the number of instances (#Inst), the number of instances that have symmetries for the various simplification configurations (#Sym[1], #Sym[2] and #Sym[3]), the number of instances that have symmetries in at least one of the configurations (#Sym[P]), the average logarithm in base 2 of the size of the symmetry group (Avg[GS]) for Configuration 1, and the total time in seconds required to analyze all the instances (Time) also for Configuration 1. It is clear from Table 1 that the SMT-LIB has many highly symmetric formulas, in most categories. The cumulative time required to build the graph and detect the symmetries is negligible in all categories. We do not output the times for other configurations since there are timeouts and time is dominantly spent in the simplification modules, so these numbers give little insight about symmetry detection itself.

Results on QF_UF require a comment. It seems that Saucy (the graph isomorphism tools used in SyMT) is not complete and does not exhibit all symmetries that are guessed by the simple heuristic in [8]. For QF_UF, we actually discovered more symmetries using Bliss [9] as a back-end, but for licensing reasons SyMT cannot include this tool. We are investigating solutions.

5 Conclusions and future work

We presented SyMT, a tool to detect symmetries in SMT formulas. SyMT is based on the reduction of the symmetry detection problem to graph automorphism detection. We presented

the corresponding graph construction algorithm and showed that symmetry detection scales on SMT formulas by providing experimental results on executions of the tool on many SMT-LIB categories. We also showed that propositional symmetry breaking can be lifted to the SMT case, which provides a simple symmetry breaking mechanism for SMT.

In future work we will address the issue of symmetry breaking. We want to study the structures of symmetry groups found by SyMT. A deeper understanding of these structures may provide useful information to develop generic symmetry breaking mechanisms. We also believe that, to fully exploit the presence of symmetries in formulas, *ad hoc*, application-tailored, heuristics are also necessary. We will use SyMT to mine the SMT-LIB to find symmetries, and we will devise appropriate heuristics integrated into an SMT symmetry breaking pre-processor. We expect this will result in a significant speed up for solving the formulas in the repository, since our experiments show symmetries are pervasive in many SMT test sets. We plan to carry out a similar analysis on the TPTP library [13].

We are aware that symmetry breaking is essentially heuristic, and a compilation of *ad hoc* heuristics would not be a silver bullet: the expertise of the user is generally the best approach to break symmetries. The current version of SyMT already provides the SMT users with a simple, yet powerful, tool to detect symmetries.

The tool and its source are available for download under the BSD License at `http://www.veriT-solver.org/SyMT`. It uses the Saucy 3.0 source code, distributed under its own specific license.

Bibliography

[1] F. Aloul, A. Ramani, I. Markov, and K. Sakallah. Solving difficult instances of Boolean satisfiability in the presence of symmetry. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 22(9):1117–1137, 2003.

[2] C. Areces, G. Hoffmann, and E. Orbe. Symmetries in modal logics: A coinductive approach. In *Proc. of the 7th Workshop on Logical and Semantic Frameworks, with Applications (LSFA 2012)*, Rio de Janeiro, September 2012.

[3] G. Audemard, B. Mazure, and L. Sais. Dealing with symmetries in quantified Boolean formulas. In *Proc. of SAT'04*, pages 257–262, 2004.

[4] C. Barrett, R. Sebastiani, S. Seshia, and C. Tinelli. Satisfiability modulo theories. In A. Biere, M. Heule, H. van Maaren, and T. Walsh, editors, *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*, chapter 26, pages 825–885. IOS Press, Feb. 2009.

[5] C. Barrett, A. Stump, and C. Tinelli. The Satisfiability Modulo Theories Library (SMT-LIB). `www.SMT-LIB.org`, 2010.

[6] J. Crawford. A theoretical analysis of reasoning by symmetry in first-order logic. In *Proc. of AAAI Workshop on Tractable Reasoning*, pages 17–22, 1992.

[7] J. Crawford, M. Ginsberg, E. Luks, and A. Roy. Symmetry-breaking predicates for search problems. In L. Carlucci Aiello, J. Doyle, and S. Shapiro, editors, *KR*, pages 148–159. Morgan Kaufmann, 1996.

[8] D. Déharbe, P. Fontaine, S. Merz, and B. Woltzenlogel Paleo. Exploiting symmetry in SMT problems. In N. Bjørner and V. Sofronie-Stokkermans, editors, *CADE*, volume 6803 of *LNCS*, pages 222–236. Springer, 2011.

[9] T. Junttila and P. Kaski. Engineering an efficient canonical labeling tool for large and sparse graphs. In D. Applegate, G. Brodat, D. Panario, and R. Sedgewick, editors, *Proc. of the 9th Workshop on Algorithm Engineering and Experiments and the 4th Workshop on Analytic Algorithms and Combinatorics*. SIAM, 2007.

[10] H. Katebi, K. Sakallah, and I. Markov. Conflict anticipation in the search for graph automorphisms. In N. Bjørner and A. Voronkov, editors, *LPAR*, volume 7180 of *LNCS*, pages 243–257. Springer, 2012.

[11] B. McKay. Nauty user's guide. Technical report, Australian National University, Computer Science Department, 1990.

[12] K. Sakallah. Symmetry and satisfiability. In A. Biere, M. Heule, H. van Maaren, and T. Walsh, editors, *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*, pages 289–338. IOS Press, 2009.

[13] G. Sutcliffe. The TPTP Problem Library and Associated Infrastructure: The FOF and CNF Parts, v3.5.0. *Journal of Automated Reasoning*, 43(4):337–362, 2009.

MEALS Partner Abbreviations

SAU: Saarland University, D

RWT: RWTH Aachen University, D

TUD: Technische Universität Dresden, D

INR: Institut National de Recherche en Informatique et en Automatique, FR

IMP: Imperial College of Science, Technology and Medicine, UK

ULEIC: University of Leicester, UK

TUE: Technische Universiteit Eindhoven, NL

UNC: Universidad Nacional de Córdoba, AR

UBA: Universidad de Buenos Aires, AR

UNR: Universidad Nacional de Río Cuarto, AR

ITBA: Instituto Técnológico Buenos Aires, AR