Q3B: An Efficient BDD-based SMT Solver for Quantified Bit-Vectors*

Martin Jonáš and Jan Strejček

Masaryk University, Brno, Czech Republic
{xjonas,strejcek}@fi.muni.cz

Abstract. We present the first stable release of our tool Q3B for deciding satisfiability of quantified bit-vector formulas. Unlike other state-of-the-art solvers for this problem, Q3B is based on translation of a formula to a BDD that represents models of the formula. The tool also employs advanced formula simplifications and approximations by effective bit-width reduction and by abstraction of bit-vector operations. The paper focuses on the architecture and implementation aspects of the tool, and provides a brief experimental comparison with its competitors.

1 Introduction

Advances in solving formula *satisfiability modulo theories* (SMT) achieved during the last few decades enabled significant progress and practical applications in the area of automated analysis, testing, and verification of various systems. In the case of software and hardware systems, the most relevant theory is the *theory of fixed-sized bit-vectors*, as these systems work with inputs expressed as bit-vectors (i.e., sequences of bits) and perform bitwise and arithmetic operations on bit-vectors. The quantifier-free fragment of this theory is supported by many general-purpose SMT solvers, such as CVC4 [1], MathSAT [7], Yices [10], or Z3 [9] and also by several dedicated solvers, such as Boolector [21] or STP [12]. However, there are some use-cases where quantifier-free formulas are not natural or expressive enough. For example, formulas containing quantifiers arise naturally when expressing loop invariants, ranking functions, loop summaries, or when checking equivalence of two symbolically described sets of states [13, 24, 8, 17, 18]. In the following, we focus on SMT solvers for *quantified* bit-vector formulas. In particular, this paper describes the state-of-the-art SMT solver Q3B including its implementation and the inner workings.

Solving of quantified bit-vector formulas was first supported by Z3 in 2013 [25] and for a limited set of *exists/forall* formulas with only a single quantifier alternation by Yices in 2015 [11]. Both of these solvers decide quantified formulas by *quantifier instantiation*, in which universally quantified variables in the Skolemized formula are repeatedly instantiated by ground terms until the resulting quantifier-free formula is unsatisfiable or a model of the original formula is found. In 2016, we proposed a different approach for solving quantified bit-vector

* This work has been supported by Czech Science Foundation, grant GA18-02177S.

formulas: by using binary decision diagrams (BDDs) and approximations [14]. For evaluation of this approach, we implemented an experimental SMT solver called Q3B, which outperformed both Z3 and Yices. Next solver that was able to solve quantified bit-vector formulas was Boolector in 2017, using also an approach based on quantifier instantiation [22]. Unlike Z3, in which the universally quantified variables are instantiated only by constants or subterms of the original formula, Boolector uses a counterexample-guided synthesis approach, in which a suitable ground term for instantiation is synthesized based on the defined grammar. Thanks to this, Boolector was able to outperform Q3B and Z3 on certain classes of formulas. More recently, in 2018, support of quantified bit-vector formulas has also been implemented into CVC4 [20]. The approach of CVC4 is also based on quantifier instantiation, but instead of synthesizing terms given by the grammar as Boolector, CVC4 uses predetermined rules based on invertibility conditions, which directly give terms that can prune many spurious models without using potentially expensive counterexample-guided synthesis. The authors of CVC4 have shown that this approach outperforms Z3, CVC4, and the original Q3B. However, Q3B has been substantially improved since the original experimental version. In 2017, we extended it with simplifications of quantified bit-vector formulas using unconstrained variables [15]. Further, in 2018, we added the experimental implementation of abstractions of bit-vector operations [16]. With these techniques, Q3B is able to decide more formulas than Z3, Boolector, and CVC4. Besides the theoretical improvements, Q3B was also improved in terms of stability, ease of use, technical parts of the implementation, and compliance with the SMT-LIB standard. This tool paper presents the result of these improvements: Q3B 1.0, the first stable version of Q3B.

We briefly summarize the SMT solving approach of Q3B. As in most of modern SMT solvers, the input formula is first simplified using satisfiability-preserving transformations that may reduce the size and complexity of the formula. The simplified formula is then converted to a binary decision diagram (BDD) that represents all assignments satisfying the formula, i.e., the *models* of the formula. If the BDD represents at least one model, we say that the BDD is *satisfiable* and it implies satisfiability of the formula. If the BDD represents the empty set of models, we say that it is *unsatisfiable* and so is the formula. Unfortunately, there are formulas for which the corresponding BDD (or some of the intermediate BDDs that appear during its computation) is necessarily exponential in the number of bits in the formula. For example, this is the case for formulas that contain multiplication of two bit-vector variables [5]. To be able to deal with such formulas, Q3B computes in parallel also BDDs underapproximating and overapproximating the original set of models, i.e., BDDs representing subsets and supersets of the original set of models, respectively. The approximating BDDs may be much smaller in size than the precise BDD, especially if the approximation is very rough. Still, they can be used to decide satisfiability of the original formula. If an overapproximating BDD is unsatisfiable, the original formula is also unsatisfiable. If the overapproximating BDD is satisfiable, we take one of its models, i.e., an assignment to the top-level existential variables of the formula, and check whether it

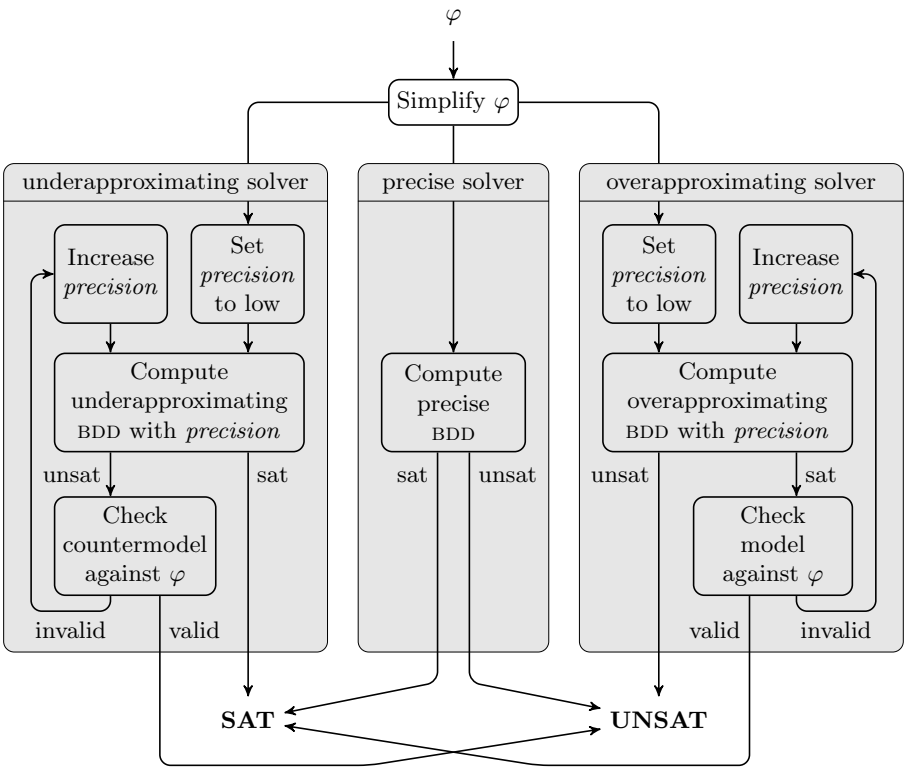

Fig. 1: High-level overview of the SMT solving approach used by Q3B. The three shaded areas are executed in parallel and the first result is returned.

is a model of the original formula. If the answer is positive, the original formula is satisfiable. In the other case, we build a more precise overapproximating BDD. Underapproximating BDDs are utilized analogously. The only difference is that for unsatisfiable underapproximating BDD, we check the validity of a countermodel, i.e., an assignment to the top-level universal variables that makes the formula unsatisfiable. The approach is depicted in Figure 1.

Q3B currently supports two ways of computing the approximating BDDs from the input formula. First of these are *variable bit-width approximations* in which the *effective bit-width* of some variables is reduced. In other words, some of the variables are represented by fewer bits and the rest of the bits is set to zero bits, one bits, or the sign bit of the reduced variable. This approach was originally used by the SMT solvers UCLID [6] and Boolector [21]. Q3B extends this approach to quantified formulas: if bit-widths of only existentially quantified variables are reduced, the resulting BDD is underapproximating; if bit-widths of only universally quantified variables are reduced, the resulting BDD is overapproximating. The second way to obtain an approximation is *bit-vector operation abstraction* [16], during which the individual bit-vector operations may not compute all bits of

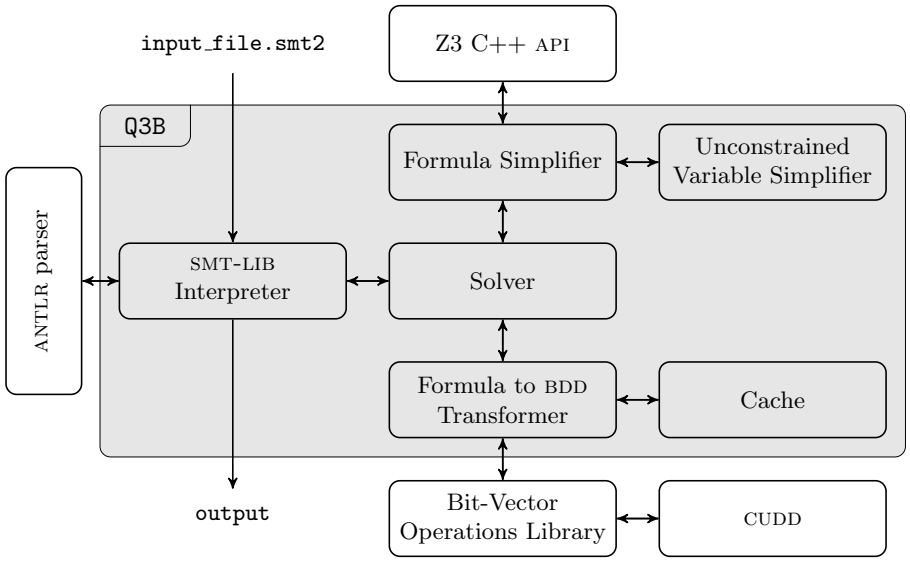

Fig. 2: Architecture of Q3B. Components in the shaded box are parts of Q3B, the other components are external.

the result, but produce some *do-not-know bits* if the resulting BDDs would exceed a given number of nodes. An underapproximating BDD then represents assignments that satisfy the formula for all possible values of these do-not-know bits. Analogously, an overapproximating BDD represents all assignments that satisfy the formula for some value of the do-not-know bits. Q3B also supports a combination of these two methods, in which both the effective bit-with of variables is reduced and the limit on the size of BDDs is imposed. During an approximation refinement, either the effective bit-width or the size limit is increased, based on the detected cause of the imprecision.

2 Architecture

This section describes the internal architecture of Q3B. The overall structure including internal and external components and the interactions between them is depicted in Figure 2. We explain the purpose of the internal components:

SMT-LIB Interpreter (implemented in `SMTLIBInterpreter.cpp`) reads the input file in the SMT-LIB format [3], which is the standard input format for SMT solvers. The interpreter executes all the commands from the file. In particular, it maintains the assertion stack and the options set by the user, calls solver when `check-sat` command is issued, and queries `Solver` if the user requires the model with the command `get-model`.

Formula Simplifier (implemented in `FormulaSimplifier.cpp`) provides interface for all applied formula simplifications, in particular miniscoping, conversion to negation normal form, pure literal elimination, equality propagation, constructive equality resolution (CER) [14], destructive equality resolution (DER) [25], simple theory-related rewriting, and simplifications using unconstrained variables. Most of these simplifications are implemented directly in this component; only CER, DER, and majority of the theory-related rewritings are performed by calling Z3 API and simplifications using unconstrained variables are implemented in a separate component of Q3B. The simplifier also converts top-level existential variables to uninterpreted constants, so their values are also included in a model. Some simplifications that could change models of the formula are disabled if the user enables model generation, i.e., sets `:produce-models` to `true`.

Unconstrained Variable Simplifier (implemented in `UnconstrainedVariableSimplifier.cpp`) provides simplifications of formulas that contain unconstrained variables, i.e., variables that occur only once in the formula. Besides previously published unconstrained variable simplifications [15], which were present in the previous versions of Q3B, this component now also provides new *goal-directed* simplifications of formulas with unconstrained variables. In these simplifications, we aim to determine whether a subterm containing an unconstrained variable should be minimized, maximized, sign minimized, or sign maximized in order to satisfy the formula. If the subterm should be minimized and contains an unconstrained variable, the term is replaced by a simpler term that gives the minimal result that can be achieved by any value of the unconstrained variable. Similarly for maximization, sign minimization, and sign maximization.

Solver (implemented in `Solver.cpp`) is the central component of our tool. It calls formula simplifier and then creates three threads for the precise solver, the underapproximating solver, and the overapproximating solver. It also controls the approximation refinement loops of the approximating solvers. Finally, it returns the result of the fastest thread and stores the respective model, if the result was `sat`.

Formula to BDD Transformer (implemented in the file `ExprToBDDTransformer.cpp`) performs the actual conversion of a formula to a BDD. Each subterm of the input formula is converted to a vector of BDDs (if the subterm's sort is a bit-vector of width n then the constructed vector contains n BDDs, each BDD represents one bit of the subterm). Further, each subformula of the input formula is converted to a BDD. These conversions proceed by a straightforward bottom-up recursion on the formula syntax tree. The transformer component calls an external library to compute the effect of logical and bit-vector operations on BDDs and vectors of BDDs, respectively. Besides the precise conversion, the transformer can also construct overapproximating and underapproximating BDDs. Precision of approximations depends on parameters set by the solver component.

Cache (implemented as a part of `ExprToBDDTransformer.cpp`) maintains for each converted subformula and subterm the corresponding BDD or a vector

of BDDs, respectively. Each of the three solvers has its own cache. When an approximating solver increases precision of the approximation, entries of its cache that can be affected by the precision change are invalidated. All the caches are internally implemented by hash-tables.

3 Implementation

Q3B is implemented in C++17, is open-source and available under MIT license on GitHub: `https://github.com/martinjonas/Q3B`. The project development process includes continuous integration and automatic regression tests.

Q3B relies on several external libraries and tools. For representation and manipulation with BDDs, Q3B uses the open-source library CUDD 3.0 [23]. Since CUDD does not support bit-vector operations, we use the library by Peter Navrátil [19] that implements bit-vector operations on top of CUDD. The algorithms in this library are inspired by the ones in the BDD library BuDDy[1] and they provide a decent performance. Nevertheless, we have further improved its performance by several modifications. In particular, we added a specific code for handling expensive operations like bit-vector multiplication and division when arguments contain constant BDDs. This for example considerably speeds up multiplication whenever one argument contains many constant zero bits, which is a frequent case when we use the variable bit-width approximation fixing some bits to zero. Further, we have fixed few incorrectly implemented bit-vector operations in the original library. Finally, we have extended the library with the support for do-not-know bits in inputs of the bit-vector operations and we have implemented abstract versions of arithmetic operations that can produce do-not-know bits when the result exceeds a given number of BDD nodes.

For parsing the input formulas in SMT-LIB format, Q3B uses ANTLR parser generated from the grammar[2] for SMT-LIB 2.6 [2]. We have modified the grammar to correctly handle bit-vector numerals and to support `push` and `pop` commands without numerical argument. The parser allows Q3B to support all bit-vector operations and almost all SMT-LIB commands except `get-assertions`, `get-assignment`, `get-proof`, `get-unsat-assumptions`, `get-unsat-core`, and all the commands that work with algebraic data-types. This is in sharp contrast with the previous experimental versions of Q3B, which only collected all the assertions from the input file and performed the satisfiability check regardless of the rest of the commands and of the presence of the `check-sat` command. The reason for this was that the older versions parsed the input file using the Z3 C++ API, which can provide only the list of assertions, not the rest of the SMT-LIB script. Thanks to the new parser, Q3B 1.0 can also provide the user with a model of a satisfiable formula after calling `get-model`; this important aspect of other SMT solvers was completely missing in the previous versions.

On the other hand, C++ API of the solver Z3 is still used for internal representation of parsed formulas. The Z3 C++ API is also used to perform manipu-

[1] `https://sourceforge.net/projects/buddy/`
[2] `https://github.com/julianthome/smtlibv2-grammar`

lations with formulas, such as substitution of values for variables, and some of the formula simplifications. Note that these are the only uses of Z3 API in Q3B during solving the formula; no actual SMT- or SAT-solving capabilities of Z3 are used during the solving process.

Some classes of Q3B, in particular `Solver`, `FormulaSimplifier`, and `UnconstrainedVariableSimplifier`, expose a public C++ API that can be used by external tools for SMT solving or just performing formula simplifications. For example, `Solver` exposes method `Solve(formula, approximationType)`, which can be used to decide satisfiability by the precise solver, the underapproximating solver, or the overapproximating solver. `Solver` also exposes the method `SolveParallel(formula)`, which simplifies the input formula and runs all three of these solvers in parallel and returns the first result as depicted in Figure 1.

4 Experimental Evaluation

We have evaluated the performance of QB3 1.0 and compared it to the latest versions of SMT solvers Boolector (v3.0), CVC4 (v1.6), and Z3 (v4.8.4). All tools were used with their default settings except for CVC4, where we used the same settings as in the paper that introduces quantified bit-vector solving in CVC4 [20], since they give better results than the default CVC4 settings. As the benchmark set, we have used all 5751 quantified bit-vector formulas from the SMT-LIB repository. The benchmarks are divided into 8 distinct families of formulas. We have executed each solver on each benchmark with CPU time limit 20 minutes and RAM limit of 8 GiB. All the experiments were performed in a Ubuntu 16.04 virtual machine within a computer equipped with Intel(R) Core(TM) i7-8700 CPU @ 3.20GHz CPU and 32 GiB of RAM. For reliable benchmarking we employed BenchExec [4], a tool that allocates specified resources for a program execution and precisely measures their usage. All scripts used for running benchmarks and processing their results, together with detailed descriptions and some additional results not presented in the paper, are available online[3].

Table 1 shows the numbers of benchmarks in each benchmark family solved by the individual solvers. Q3B is able to solve the most benchmarks in benchmark families *2017-Preiner-scholl-smt08*, *2017-Preiner-tptp*, *2017-Preiner-UltimateAutomizer*, *2018-Preiner-cav18*, and *wintersteiger*, and it is competitive in the remaining families. In total, Q3B also solves more formulas than each of the other solvers: 116 more than Boolector, 83 more than CVC4, and 139 more than Z3. Although the numbers of solved formulas for the solvers seem fairly similar, the cross-comparison in Table 2 shows that the differences among the individual solvers are actually larger. For each other solver, there are at least 143 benchmarks that can be solved by Q3B but not by the other solver. We think this shows the importance of developing an SMT solver based on BDDs and approximations besides the solvers based on quantifier instantiation.

[3] `https://github.com/martinjonas/q3b-artifact`

Family	Total	Boolector	CVC4	Q3B	Z3
2017-Preiner-keymaera	4035	4022	3998	4009	**4031**
2017-Preiner-psyco	194	193	190	182	**194**
2017-Preiner-scholl-smt08	374	312	248	**319**	272
2017-Preiner-tptp	73	69	**73**	**73**	**73**
2017-Preiner-UltimateAutomizer	153	152	151	**153**	**153**
20170501-Heizmann-UltimateAutomizer	131	30	**128**	124	32
2018-Preiner-cav18	600	553	**565**	**565**	553
wintersteiger	191	163	174	**185**	163
Total	5751	5494	5527	**5610**	5471
CPU time [s]		7793	5877	19853	**4055**

Table 1: For each solver and benchmark family, the table shows the number of benchmarks from the given family solved by the given solver. The column *Total* shows the total number of benchmarks in the given family. The last line provides the total CPU times for the benchmarks solved by all four solvers.

	Boolector	CVC4	Q3B	Z3	Uniquely solved
Boolector	0	123	69	78	8
CVC4	156	0	60	171	6
Q3B	185	143	0	208	25
Z3	55	115	69	0	6

Table 2: For all pairs of the solvers, the table shows the number of benchmarks that were solved by the solver in the corresponding row, but not by the solver in the corresponding column. The column *Uniquely solved* shows the number of benchmarks that were solved only by the given solver.

5 Conclusions and Future Work

We have described the architecture and inner workings of the first stable version of the state-of-the-art SMT solver Q3B. Experimental evaluation on all quantified bit-vector formulas from SMT-LIB repository shows that this solver slightly outperforms other state-of-the-art solvers for such formulas.

As future work, we would like to drop the dependency on the Z3 API: namely to implement our own representation of formulas and reimplement all the simplifications currently outsourced to Z3 API directly in Q3B. We also plan to extend some simplifications with an additional bookkeeping needed to construct a model of the original formula. With these extensions, all simplifications could be used even if the user wants to get a model of the formula. We would also like to implement production of unsatisfiable cores since they are also valuable for software verification.

References

1. Clark Barrett, Christopher L. Conway, Morgan Deters, Liana Hadarean, Dejan Jovanovic, Tim King, Andrew Reynolds, and Cesare Tinelli. CVC4. In *Computer Aided Verification - 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings*, pages 171–177, 2011.
2. Clark Barrett, Pascal Fontaine, and Cesare Tinelli. The SMT-LIB Standard: Version 2.6. Technical report, Department of Computer Science, The University of Iowa, 2017. Available at `www.SMT-LIB.org`.
3. Clark Barrett, Aaron Stump, and Cesare Tinelli. The SMT-LIB Standard: Version 2.0. In A. Gupta and D. Kroening, editors, *Proceedings of the 8th International Workshop on Satisfiability Modulo Theories (Edinburgh, UK)*, 2010.
4. Dirk Beyer, Stefan Löwe, and Philipp Wendler. Benchmarking and Resource Measurement. In *Model Checking Software - 22nd International Symposium, SPIN 2015, Proceedings*, volume 9232 of *Lecture Notes in Computer Science*, pages 160–178. Springer, 2015.
5. Randal E. Bryant. On the Complexity of VLSI Implementations and Graph Representations of Boolean Functions with Application to Integer Multiplication. *IEEE Trans. Comput.*, 40(2):205–213, 1991.
6. Randal E. Bryant, Daniel Kroening, Joël Ouaknine, Sanjit A. Seshia, Ofer Strichman, and Bryan A. Brady. An abstraction-based decision procedure for bit-vector arithmetic. *STTT*, 11(2):95–104, 2009.
7. Alessandro Cimatti, Alberto Griggio, Bastiaan Joost Schaafsma, and Roberto Sebastiani. The MathSAT5 SMT Solver. In *Tools and Algorithms for the Construction and Analysis of Systems - 19th International Conference, TACAS 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings*, pages 93–107, 2013.
8. Byron Cook, Daniel Kroening, Philipp Rümmer, and Christoph M. Wintersteiger. Ranking function synthesis for bit-vector relations. *Formal Methods in System Design*, 43(1):93–120, 2013.
9. Leonardo Mendonça de Moura and Nikolaj Bjørner. Z3: An efficient SMT solver. In *Tools and Algorithms for the Construction and Analysis of Systems, 14th International Conference, TACAS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29-April 6, 2008. Proceedings*, volume 4963 of *Lecture Notes in Computer Science*, pages 337–340. Springer, 2008.
10. Bruno Dutertre. Yices 2.2. In *Computer Aided Verification - 26th International Conference, CAV 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 18-22, 2014. Proceedings*, pages 737–744, 2014.
11. Bruno Dutertre. Solving Exists/Forall Problems with Yices. In *Workshop on satisfiability modulo theories*, 2015.
12. Vijay Ganesh and David L. Dill. A Decision Procedure for Bit-Vectors and Arrays. In *Computer Aided Verification, 19th International Conference, CAV 2007, Berlin, Germany, July 3-7, 2007, Proceedings*, pages 519–531, 2007.
13. Sumit Gulwani, Saurabh Srivastava, and Ramarathnam Venkatesan. Constraint-Based Invariant Inference over Predicate Abstraction. In *Verification, Model Checking, and Abstract Interpretation, 10th International Conference, VMCAI 2009, Savannah, GA, USA, January 18-20, 2009. Proceedings*, pages 120–135, 2009.

14. Martin Jonáš and Jan Strejček. Solving Quantified Bit-Vector Formulas Using Binary Decision Diagrams. In *Theory and Applications of Satisfiability Testing - SAT 2016 - 19th International Conference, Bordeaux, France, July 5-8, 2016, Proceedings*, volume 9710 of *Lecture Notes in Computer Science*, pages 267–283. Springer, 2016.
15. Martin Jonáš and Jan Strejček. On Simplification of Formulas with Unconstrained Variables and Quantifiers. In *Theory and Applications of Satisfiability Testing - SAT 2017 - 20th International Conference, Melbourne, VIC, Australia, August 28 - September 1, 2017, Proceedings*, pages 364–379, 2017.
16. Martin Jonáš and Jan Strejček. Abstraction of Bit-Vector Operations for BDD-Based SMT Solvers. In *Theoretical Aspects of Computing - ICTAC 2018 - 15th International Colloquium, Stellenbosch, South Africa, October 16-19, 2018, Proceedings*, pages 273–291, 2018.
17. Daniel Kroening, Matt Lewis, and Georg Weissenbacher. Under-Approximating Loops in C Programs for Fast Counterexample Detection. In *Computer Aided Verification - 25th International Conference, CAV 2013*, volume 8044 of *LNCS*, pages 381–396. Springer, 2013.
18. Jan Mrázek, Petr Bauch, Henrich Lauko, and Jiří Barnat. SymDIVINE: Tool for control-explicit data-symbolic state space exploration. In *Model Checking Software - 23rd International Symposium, SPIN 2016, Co-located with ETAPS 2016, Eindhoven, The Netherlands, April 7-8, 2016, Proceedings*, pages 208–213, 2016.
19. Peter Navrátil. Adding Support for Bit-Vectors to BDD Libraries CUDD and Sylvan. Bachelor's thesis, Masaryk University, Faculty of Informatics, Brno, 2018.
20. Aina Niemetz, Mathias Preiner, Andrew Reynolds, Clark Barrett, and Cesare Tinelli. Solving quantified bit-vectors using invertibility conditions. In *Computer Aided Verification - 30th International Conference, CAV 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 14-17, 2018, Proceedings, Part II*, pages 236–255, 2018.
21. Aina Niemetz, Mathias Preiner, Clifford Wolf, and Armin Biere. Btor2 , BtorMC and Boolector 3.0. In *Computer Aided Verification - 30th International Conference, CAV 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 14-17, 2018, Proceedings, Part I*, pages 587–595, 2018.
22. Mathias Preiner, Aina Niemetz, and Armin Biere. Counterexample-Guided Model Synthesis. In *Tools and Algorithms for the Construction and Analysis of Systems - 23rd International Conference, TACAS 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings, Part I*, volume 10205 of *Lecture Notes in Computer Science*, pages 264–280. Springer, 2017.
23. Fabio Somenzi. CUDD: CU Decision Diagram Package Release 3.0.0. *University of Colorado at Boulder*, 2015.
24. Saurabh Srivastava, Sumit Gulwani, and Jeffrey S. Foster. From program verification to program synthesis. In *Proceedings of the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2010, Madrid, Spain, January 17-23, 2010*, pages 313–326, 2010.
25. Christoph M. Wintersteiger, Youssef Hamadi, and Leonardo Mendonça de Moura. Efficiently solving quantified bit-vector formulas. *Formal Methods in System Design*, 42(1):3–23, 2013.