

Symbolic computation and satisfiability checking

Davenport, J. H., England, M., Griggio, A., Sturm, T. & Tinelli, C.

Author post-print (accepted) deposited by Coventry University's Repository

Original citation & hyperlink:

Davenport, JH, England, M, Griggio, A, Sturm, T & Tinelli, C 2020, 'Symbolic computation and satisfiability checking', *Journal of Symbolic Computation*, vol. 100, pp. 1-10.

<https://dx.doi.org/10.1016/j.jsc.2019.07.017>

DOI 10.1016/j.jsc.2019.07.017

ISSN 0747-7171

Publisher: Elsevier

**NOTICE: this is the author's version of a work that was accepted for publication in *Journal of Symbolic Computation*. Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. A definitive version was subsequently published in *Journal of Symbolic Computation*, 100, (2020)
DOI: 10.1016/j.jsc.2019.07.017**

© 2020, Elsevier. Licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International

<http://creativecommons.org/licenses/by-nc-nd/4.0/>

Copyright © and Moral Rights are retained by the author(s) and/ or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This item cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder(s). The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holders.

This document is the author's post-print version, incorporating any revisions agreed during the peer-review process. Some differences between the published version and this version may remain and you are advised to consult the published version if you wish to cite from it.

Symbolic Computation and Satisfiability Checking Editorial

James H. Davenport

Faculty of Science, University of Bath, UK

Matthew England

Faculty of Engineering, Environment and Computing, Coventry University, UK

Alberto Griggio

Embedded Systems Unit, Fondazione Bruno Kessler, Italy

Thomas Sturm

*CNRS, Inria, and the University of Lorraine, France;
MPI Informatics and Saarland University, Germany*

Cesare Tinelli

Department of Computer Science, University of Iowa, USA

Abstract

The two communities of Symbolic Computation and Satisfiability Checking have recently found themselves tackling similar problems and having a growing interest in each other's technology. This special issue presents articles whose contribution is of interest to, and is influenced by, both communities. Given the context of this journal we start this editorial with a more thorough overview of Satisfiability Checking, and then turn to Symbolic Computation and the potentials and challenges for collaboration. The collection of articles in this issue is evidence of the already existing fruitful work at the intersection of these communities.

Email addresses: J.H.Davenport@bath.ac.uk (James H. Davenport),
Matthew.England@coventry.ac.uk (Matthew England), Griggio@fbk.eu (Alberto Griggio), thomas@thomas-sturm.de (Thomas Sturm), cesare-tinelli@uiowa.edu (Cesare Tinelli)

Keywords: satisfiability checking, symbolic computation, community integration

1. Two Communities

1.1. The SAT Problem

The Satisfiability Checking community is focussed on the problem of determining whether a given logical formula is satisfiable, that is, whether there is an interpretation of certain uninterpreted symbols that evaluates the formula to `true`. In the case where the uninterpreted symbols are all Boolean¹ variables this becomes the classic SAT Problem, the first problem to be proven NP-Complete (Cook, 1971).

Despite its theoretical complexity, today instances of the SAT problem involving hundreds of thousands of variables are solved routinely. This is usually done by dedicated *SAT-solvers*, many of which are available free and open source. SAT-solvers are now integrated in a wide variety of industrial domains, perhaps most notably formal verification. A textbook charting the rise of the SAT-solver is the Handbook of Satisfiability by Biere et al. (2009).

Modern SAT-solvers mostly rely on a combination of enumeration, propagation, and resolution. *Enumeration* refers to the natural solution of enumerating all possible variable assignments. In the case of an unsatisfiable formula with n variables this means checking 2^n assignments. Propagation and resolution help tackle this exponential search space: we introduce them next alongside the algorithms which first employed them.

1.1.1. The DPLL Algorithm

Propagation was the key feature of the DPLL algorithm by Davis et al. (1962). We first transform² the formula into *conjunctive normal form*: the formula becomes a conjunction of clauses, where each clause is a disjunction of literals (variables or their negation). We then proceed with an enumerative search but after each variable is set we seek implications for other variables. Such implications happen if the latest assignment renders a clause *unit* (all

¹For our purposes the term Boolean algebra refers to propositional logic.

²Naively, this might be expensive, but the transformation of Tseitin (1983) (reprint of work from the 1960s) is in PTIME. The transformed formula is equi-satisfiable with the original (not necessarily equivalent).

unassigned literals except one become **false**). Furthermore, if all literals in a clause are rendered **false** by an assignment, then no further analysis is required in this part of the search space. We instead backtrack to the last free decision (one made as part of the search rather than propagation).

Example 1. Consider the formula

$$(\alpha \vee \beta \vee \neg\gamma) \wedge (\neg\alpha \vee \neg\beta) \wedge (\beta \vee \gamma) \wedge (\neg\alpha \vee \beta \vee \neg\gamma), \quad (1)$$

which is already in conjunctive normal form with the parentheses indicating the four clauses. An initial assignment of α to **true** would render the second clause unit and so propagation would set β to **false**. This in turn renders both the third and fourth clauses unit, however, propagating on one of them will give an assignment which renders the other **false**. So either way we backtrack to our choice and reassign α to **false**.

This time there is no unit clause and so we have to make a choice for β . If we choose to assign β to **true** then the formula is already satisfied regardless of the value for γ ; while if we assigned β to **false** then propagation would lead to a conflict and backtracking to pick β to be **true**.

1.1.2. The CDCL Algorithm

Propagation avoids exploring every branch of the search space but we may still end up replaying similar analyses in separate branches. The next milestone in the development of SAT-solvers was the idea of using resolution to avoid such repetition, as described by the conflict driven clause learning (CDCL) algorithm (Marques-Silva and Sakallah, 1999).

The main idea is to use *resolution* to learn the reason for a conflict and generate a new clause from that which we add to the input formula so that similar conflicts are avoided later in the search. This is done by means of a resolvent: given two clauses, one which has x as a literal and one with $\neg x$, the resolvent is the disjunction of all the other literals from both clauses. All assignments that satisfy both clauses must also satisfy the resolvent (since one of the other literals must have been satisfied).

Example 2. Consider a formula $(1) \wedge \psi$ where (1) was the formula from Example 1 and ψ another formula in conjunctive normal form involving more variables than just α, β, γ . In the previous example we learnt that there was no satisfying solution to (1) with α assigned to **true**. We want to avoid repeating this analysis in different branches of the search space defined by

the other variables in ψ . Resolution would ensure this: before the initial backtrack in Example 1 we saw that the third and fourth clauses conflicted – they wanted to propagate different assignments for γ . Thus we take the resolvent of these clauses to form $\neg\alpha \vee \beta$. This in turn causes a conflict with the second clause and their resolvent is simply $\neg\alpha$. Thus we add this to the input to ensure from then on that we only assign α to **false**.

The process of deriving resolvent clauses by analysing conflicts is the main idea of the CDCL algorithm, and together with propagation, it forms the backbone of modern SAT solvers. In fact, the CDCL algorithm can be seen as an interplay between two phases, *model search* and *proof search*, which drive each other in a feedback loop: model search tries to find a solution for the input formula, by applying decisions and propagation; when a conflict arises, the information about which propagations were performed is used by the proof search phase to learn new clauses (generated via resolution), which in turn will drive the model search phase to explore a different portion of the search space.

1.1.3. Modern SAT-solvers

We have presented above the two most important ideas to tackle the SAT problem, but there are many other innovations used in modern SAT-solvers. There is substantial engineering such as the *two watched literals* rule, heuristics for variable selection / branching (e.g. the VSIDS variable activity score), restart strategies, and techniques for clause database management (e.g. the LBD measure). Also of great importance are the preprocessing and inprocessing procedures, such as variable elimination, subsumption, and blocked clauses elimination. The best single reference for an overview of these topics is the Handbook of Satisfiability Biere et al. (2009).

1.2. Satisfiability Modulo Theories

The success of modern SAT solvers led naturally to the idea of applying such techniques beyond propositional calculus. The approach has been best formalised by the paradigm of Satisfiability Modulo Theories (SMT), which leads to SMT-solvers. These take as input a logical formula whose atoms are statements in a theory domain (or perhaps several different domains).

The SMT-solver will usually first convert the formula to conjunctive normal form and build its *Boolean skeleton* by replacing each theory atom by a new Boolean variable. This skeleton is passed to a SAT-solver: if it is found

to be unsatisfiable then the conclusion can be applied to the original formula also. To conclude, in contrast, a satisfying solution, one must check at some point the validity of the solution from the SAT-solver in the theory. This is usually done by calling a *theory solver*: a procedure for deciding if the set of theory constraints implied by the interpretation of Boolean variables is consistent. If not consistent then the theory solver is expected to offer an explanation: a statement that is a valid in the theory stating that a, hopefully small, subset of the constraints cannot hold together. This forms a new clause for the Boolean skeleton, generalising the inconsistency just encountered, to ensure future candidate solutions are not invalid in a similar manner. This way we are running a loop of calls to the SAT-solver and theory solver until, ideally, either a satisfying solution is found or unsatisfiability is concluded. In practice, one also admits incomplete theories and theory solvers so that SMT solvers possibly return *unknown* as a final result.

Example 3. Consider the formula

$$(x > 0 \vee x = -1) \wedge (x^2 < 0 \vee x^2 > 9).$$

The Boolean skeleton is $(a \vee b) \wedge (c \vee d)$ where new Boolean variables a, b, c, d represent the truth of the four sign conditions involving a theory variable x . A SAT-solver may propose the solution to this skeleton which sets a and c to **true** and b and d to **false**. A theory solver would determine that the set of constraints this produces is incompatible and produce the explanation that we cannot have both a and c **true** at once. Thus the clause $(\neg a \vee \neg c)$ is added to the input. Further analysis by the SAT-solver may then swap the Boolean assignment for c and d . This time the theory solver should conclude the resulting set of conditions is satisfiable in the theory (by any $x > 3$).

The above description is of the DPLL(T) architecture for SMT, and the example implemented the *lazy* approach, where the theory solver is only consulted after a full Boolean assignment has been obtained. There can be less lazy versions, with the theory solver consulted after each assignment. The chapter in the Handbook of Satisfiability by Barrett et al. (2009) is a key reference for this SMT paradigm.

There are other approaches to generalising SAT-solving to theory domains. The *eager* approach, introduced in Lahiri and Seshia (2004), translates the original formula to an equisatisfiable Boolean formula in a single step, which is then given to a standard SAT solver. *Programmatic SAT* introduced by Ganesh et al. (2012) calls user specified code instead of a particular

theory solver. The *model constructing satisfiability calculus* (MCSAT) introduced by de Moura and Jovanović (2013) extends propagation and resolution into the theory domain: it allows guesses, not only for truth values of the theory constraints, but also for guesses for the actual theory variables, followed by propagation techniques to drive the search for other theory variables away from unsatisfiable parts of the search space.

The theory solvers discussed above implement decision procedures to decide whether a set of constraints is compatible in the given domain. For linear real arithmetic the simplex method could be used for example, while in linear integer arithmetic cutting planes or interval constraint propagation could be used. The textbook by Kroening and Strichman (2013) offers a good outline of the various procedures commonly employed by SMT. For some domains the procedures in question were developed within the field of Symbolic Computation and have their natural home in computer algebra systems.

1.3. Symbolic Computation

The use of computers to do exact mathematics dates back almost as far as computing itself with the first theses appearing in 1953. Early successes include algorithms to perform automatic integration which could even certify when something could not be integrated in closed form (Risch, 1969). New methods for polynomial factorisation accompanied the first general purpose algebraic software in the 1960s. One of the most impactful contributions to the field was Gröbner Bases, developed by Buchberger (1965), which allowed for the efficient solution of many problems for polynomials over algebraically closed fields.

The scope of the field of Symbolic Computation was set out in the editorial of the first volume of this journal (Buchberger, 1985) and is summarised as “*the algorithmic solution of problems dealing with symbolic objects*”. Although the aforementioned editorial explicitly included both computational algebra and computational logic over time the Symbolic Computation field has seemed to focus more on the algebraic. Indeed, the implementation tools are almost exclusively referred to as *computer algebra systems*.

Today, these vary from large commercial general purpose systems such as **Maple**, and **Mathematica** to free and open source alternatives such as **Maxima**, **Reduce**, **Sage** and **Singular**. There are also a wide variety of specialised systems dedicated to subfields including **Macaulay2** for algebraic geometry, **GAP**

for group theory, `CoCoA` for commutative algebra, `Qepcad` for real quantifier elimination, and `PariGP` for number theory.

1.4. *Difficulties with Computer Algebra as Theory Solver*

The history of Symbolic Computation is longer than that of Satisfiability Checking. The first SAT solver following modern design principles did not appear until the 1990s and by the time that SMT was actively developed many of today's popular computer algebra systems were already well established. So it should seem natural that SMT-solvers consider using a computer algebra system as a theory solver. However, this did not happen, for a number of reasons beyond the lack of familiarity of the communities. Computer algebra systems tend to be large standalone systems aimed at interactive problem solving with a human, while SMT-solvers tend to be made up from small, flexible and highly efficient libraries. Hence it is not trivial for an SMT-solver to share data-structures with a computer algebra system as it does with its typical solvers. More fundamentally though, the actual algorithms employed by computer algebra systems lacked in functionality necessary for successful use in SMT. To be *SMT-compliant* a theory solver must satisfy a number of criteria (Ábrahám, 2015; Ábrahám et al., 2016):

- The solver should work incrementally, i.e. after determining the consistency of a constraint set there should be the option to add an additional constraint and recheck, reusing as much computation as possible.
- The solver should have the ability to support backtracking, i.e. after determining the consistency of a constraint set there should be the option to remove a constraint and recheck, again, reusing as much computation as possible.
- In the case of unsatisfiability, the solver is expected to offer an explanation as to why the constraints cannot be satisfied, e.g. a small subset which are on their own incompatible.

These requirements are not commonly met by the algorithms in computer algebra systems, where the intended use case is usually to perform one computation to solve a single problem. To make a Symbolic Computation algorithm SMT-compliant requires much work, certainly at the implementation level, e.g. the data-structures used, but more importantly also in the mathematical theory.

2. SC-Square

We use SC² (SC-Square) to refer to the collaboration between the two communities of Satisfiability Checking and Symbolic Computation.

2.1. EU Project

SC² originates from a European Union Horizon 2020 project³ that ran from 2016 to 2018 with the aim of bridging the gap between these communities. The project consortium consisted of: University of Bath (UK), RWTH Aachen (Germany), Fondazione Bruno Kessler (Italy), the University of Genoa (Italy), Maplesoft Europe Ltd, the University of Lorraine (France), Coventry University (UK), the University of Oxford (UK), the University of Kassel (Germany), the Max Planck Institute for Informatics (Germany), and Johannes Kepler University of Linz (Austria). In addition, the project had over 50 associates from around the world.

The main project goals were around community building, in particular through the funding of workshops and summer schools. The 2017 SC² Summer School in Saarbrücken, Germany included 12 extended lectures on the foundations of the two fields, key algorithms, and applications. The latter were presented by speakers from industry describing the use of such technology to verify software and systems in industries as diverse as automotive and finance. The slides and videos of some talks from the Summer School are available online⁴.

The existence and maintenance of a common specification language SMT-LIB (Barrett et al., 2016), with its accompanying collections of benchmarks and annual competitions, is seen as an important accelerator on research in Satisfiability Checking. Another key output from the project was work to extend SMT-LIB to better support the areas of relevance to SC² (a natural choice since there was no computer algebra alternative). The upcoming SMT-LIB 3.0 is set to include a variety of features relevant to Symbolic Computation such as new functions, quantifier elimination, real algebraic numbers, and optimisation. There have been 10 separate contributions of benchmarks linked to the SC² project, accounting for a third of the recent library growth. This has included a more diverse range of applications such as those from the life sciences (Sturm, 2017) and economics (Mulligan et al.,

³<http://www.sc-square.org/EU-CSA.html>

⁴<http://www.sc-square.org/CSA/school/lectures.html>

2018). Also, **Maple** now supports the SMT-LIB language and integration with SAT/SMT solvers (Forrest, 2017) as does **Reduce** which further, has a fully featured CDCL solver implemented and won the NRA category of the 2017 SMT competition⁵.

The project also funded a variety of proof of concept studies including on integrations: between computer algebra system Redlog and SMT-solver VeriT, between computer algebra systems CoCoA and Maple and SMT-solver SMT-RAT, and between computer algebra system CoCoA and SMT-solver MathSAT; as well as development on new incremental versions of symbolic algorithms.

2.2. Workshop Series

The SC² project initiated an annual International Workshop on Satisfiability Checking and Symbolic Computation:

- SC² 2016: 24th September 2016 in Timisoara, Romania. As part of the International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC 2016). Proceedings edited by Ábrahám and Fontaine (2016).
- SC² 2017: 29th July 2017 in Kaiserslautern, Germany. Adjacent to the International Symposium on Symbolic and Algebraic Computation (ISSAC 2017). Proceedings edited by England and Ganesh (2017).
- SC² 2018: 11th July 2018 in Oxford, UK. Part of the Federated Logic Conference (FLoC 2018). Proceedings edited by Bigatti and Brain (2018).
- SC² 2019: 10th July 2019 in Bern, Switzerland. Part of the SIAM Conference on Applied Algebraic Geometry. Proceedings to appear, edited by Abbott and Griggio.

Although the EU project finished in 2018 there is considerable effort being spent to continue the workshop series. A constitution has been approved which sets out requirements: such as the conference alternating between affiliation with conferences in each community; and the need for two chairs each year, one from each community.

⁵<http://smtcomp.sourceforge.net/2017/results-NRA.shtml>

Most articles in this special issue are based on preliminary work presented at an SC² workshop, or one of the other SC² special sessions that took place at: the 2016 Conference on Applications of Computer Algebra (ACA 2016); the 2016 International Workshop on Computer Algebra in Scientific Computing (CASC 2016); the 2017 International Conference on Mathematical Aspects of Computer and Information Sciences (MACIS 2017); and the 2018 International Congress on Mathematical Software (ICMS 2018).

3. Articles of this Special Issue

The special issue comprises seven articles. A call was made following the 2017 SC² workshop, and submission was open to all work on the the intersection of the two communities (no requirement for affiliation with the EU project). Of the seven articles, four have authors funded by the project (although all had taken part in the project workshops and special sessions). An average of 3.3 reviews per article was obtained with all reviewed by at least one expert in each of the two communities. All accepted articles underwent a multi-stage revision process to take into account the high standards of this journal.

3.1. *Non-linear real arithmetic*

SMT-LIB organises the theory solver domains into logics. From these, the one which has probably found the greatest intersection with Symbolic Computation has been non-linear real arithmetic. Here the atoms of the logical formulas are sign-conditions on non-linear multivariate polynomials, whose variables take real values. Questions in this theory are decidable, a result uncovered by Tarski (1948).

The decision procedures most commonly used for non-linear real arithmetic include *(Partial) Cylindrical Algebraic Decomposition* (Collins, 1975; Collins and Hong, 1991) and *Virtual Substitution* (Weispfenning, 1997; Kosta, 2016), familiar to readers of this journal as key algorithms of Symbolic Computation, developed for quantifier elimination in real closed fields. Given a non-linear arithmetic formula as described above which is preceded with quantifications on some of the variables involved, *Quantifier Elimination* (QE) means to find an equivalent unquantified version. For example, QE would transform $\exists x, ax^2 + bx + c = 0 \wedge a \neq 0$ to the equivalent unquantified statement $b^2 - 4ac \geq 0$. Satisfiability Checking is hence a sub-problem of QE for the case where all variables are existentially quantified (so the answer is

boolean rather than an expression in the unquantified variables). Although SMT-LIB is now expanding to encompass QE also.

Two articles of the present special issue concern Cylindrical Algebraic Decomposition (CAD). A CAD is a decomposition of ordered \mathbb{R}^n space into cells arranged cylindrically, meaning the projections of any pair of cells with respect to the variable ordering are either equal or disjoint. The projections form an induced CAD of the lower dimensional space. The cells are (semi)-algebraic meaning each can be described with a finite sequence of polynomial constraints. To study a logical formula for SMT or QE a CAD must be produced *truth-invariant* to the formula (so the formula is either `true` or `false` on each cell). We can then check satisfiability by testing a finite number of sample points, and in the case of QE form the equivalent formulae from the cell descriptions. CAD traditionally consists of two stages: first projection, which identifies polynomials which will form boundaries in the decomposition; and then lifting, where the decomposition is produced by repeated use of real root isolation on these polynomials rendered univariate by substitution at sample points of the CAD for the space below.

The article by Abraham and Kremer describes how they have adapted the original CAD algorithm to satisfy the requirements for SMT-compliance outlined in Section 1.4. Their algorithm builds a CAD incrementally, by performing one projection step at a time and then moving to lifting to perform the additional real root isolation and cell decomposition required by the new additions. They maintain a history of computation allowing for backtracking, developing data structures quite different to existing CAD implementations. Their work is implemented in their SMT-solver `SMT-RAT`.

The article by England, Bradford and Davenport considers how the process of CAD construction can be improved when the input contains equational constraints (equations which must be satisfied for the formula to hold). Intuitively each of these decrease the solution space by one dimension, and the authors show that a corresponding decrease in the theoretical complexity of CAD can be achieved. The work may act as a tutorial on CAD with equational constraints with all necessary background and the current state of the art.

The article by Brown and Vale-Enriquez presents work for a partial solver for non-linear real arithmetic, i.e. a solver that can sometimes decide the satisfiability, and if not can at least simplify the input formula to ease the use of something like CAD. The partial solver is based on existing simplification algorithms, which (like the CAD of Abraham and Kremer) have had

to be heavily modified for SMT-compliance. These simplification algorithms have also been adapted to produce the desired explanations for unsatisfiability required for SMT-compliance. This work is implemented in the `Tarksi` system, developed by Brown to implement the non-uniform CAD (NuCAD) algorithm developed by Brown (2015, 2016) which is an example of work in Symbolic Computation inspired by techniques from Satisfiability Checking (most notably those of Jovanovic and de Moura (2012)).

3.2. *Linear Arithmetic*

The article by Bromberger, Sturm and Weidenbach presents a new decision procedure for linear integer arithmetic. Although this theory is decidable, the authors explain that the leading solvers currently neglect termination on some classes of problems in favour of efficiency on all other problems. As an alternative, the authors present an extension to an existing calculus so it becomes complete and terminating but stays efficient. The procedure is based on a mixture of model-driven reasoning and quantifier elimination techniques.

The article by Abbott, Bigatti, Palezzato and Robbiano concerns the computation of minimal polynomials in P/I where P is a polynomial ring and I an ideal. Such non-linear objects can capture fundamental information about linear algebra. The work can also improve the computation of Gröbner Bases, which although a technique for algebraically closed fields such as \mathbb{C} is regularly used as an initial unsatisfiability check for real arithmetic, and is well known as a useful pre-processing for CAD (Huang et al., 2016). Their algorithms can also assist with polynomial factorization, useful for many procedures (notably CAD, where experience shows that it improves performance even when not logically necessary). The authors implement their work in the computer algebra system `CoCoA` but also `CoCoALib`: a free library of C++ code that can be particularly easily integrated with other C-like systems.

3.3. *Symbolic Computation for Better SAT-solvers*

The article by Horáček and Kreuzer is the only one in the issue to focus on the Boolean domain, as in Section 1.1. The idea is to use the equivalence between the Boolean SAT-problem and polynomial system solving over a finite field. The former may be attacked with SAT-solvers as described above while the latter by techniques such as Gröbner Bases, which originate from Symbolic Computation. The algebraic solver can be called to support the SAT-solver at selected points, or be in run in parallel, periodically exchanging

information. The article in this issue is focussed on the method to convert a formula from the conjunctive normal form (CNF) used by the SAT-solver to the algebraic normal form (ANF) required for the algebraic techniques. They develop a blockwise algorithm to producing fewer and lower degree polynomials.

3.4. SC² Combinations for Combinatorics

The article by Bright, Kotsireas and Ganesh investigates the Williamson Conjecture, an open problem in combinatorics on the existence of certain matrices, useful for error correcting codes. Through a combination of SAT-solver and computer algebra system the authors are able to enumerate all cases up to a much larger dimension than achieved before, uncover new cases, and gain new insight into the distribution of these matrices. Rather than the standard SMT approach outlined in Section 1.2 they use programmatic SAT: here rather than a full CNF encoding instances can be a set of CNF clauses and user code that encodes constraints too cumbersome to be written in CNF. The authors have applied these techniques to a variety of other combinatorial problems: an intersection of the two communities not predicted by the EU grant but documented by the workshop series.

Acknowledgements

Much of the editorial work for this special issue, as well as the work for many of the articles, was supported by EU H2020-FETOPEN-2016-2017-CSA project SC² (712689). The editors thank all the external reviewers for their detailed and thoughtful comment on the articles in this special issue.

References

- Ábrahám, E., 2015. Building bridges between symbolic computation and satisfiability checking. In: Proceedings of the 2015 International Symposium on Symbolic and Algebraic Computation. ISSAC '15. ACM, pp. 1–6.
URL <https://doi.org/10.1145/2755996.2756636>
- Ábrahám, E., Abbott, J., Becker, B., Bigatti, A., Brain, M., Buchberger, B., Cimatti, A., Davenport, J., England, M., Fontaine, P., Forrest, S., Griggio, A., Kroening, D., Seiler, W., Sturm, T., 2016. SC²: Satisfiability checking meets symbolic computation. In: Kohlhase, M., Johansson, M., Miller, B., de Moura, L., Tompa, F. (Eds.), Intelligent Computer Mathematics: Proceedings CICM 2016. Vol. 9791 of Lecture Notes in Computer Science. Springer International Publishing, pp. 28–43.
URL https://doi.org/10.1007/978-3-319-42547-4_3
- Ábrahám, E., Fontaine, P. (Eds.), 2016. Proceedings of the 1st Workshop on Satisfiability Checking and Symbolic Computation (SC² 2016). No. 1804 in CEUR Workshop Proceedings. Springer.
URL <http://ceur-ws.org/Vol-1804/>
- Barrett, C., Fontaine, P., Tinelli, C., 2016. The Satisfiability Modulo Theories Library (SMT-LIB). Online Resource. URL <http://www.SMT-LIB.org>.
- Barrett, C., Sebastiani, R., Seshia, S., Tinelli, C., 2009. Satisfiability modulo theories. In: Biere, A., Heule, M., van Maaren, H., Walsh, T. (Eds.), Handbook of Satisfiability (Volume 185 Frontiers in Artificial Intelligence and Applications), Chapter 26. IOS Press, pp. 825–885.
- Biere, A., Biere, A., Heule, M., van Maaren, H., Walsh, T., 2009. Handbook of Satisfiability: Volume 185 Frontiers in Artificial Intelligence and Applications. IOS Press, Amsterdam, The Netherlands.
- Bigatti, A., Brain, M. (Eds.), 2018. Proceedings of the 3rd Workshop on Satisfiability Checking and Symbolic Computation (SC² 2018). No. 2189 in CEUR Workshop Proceedings. Springer.
URL <http://ceur-ws.org/Vol-2189/>
- Brown, C., 2015. Open non-uniform cylindrical algebraic decompositions. In: Proceedings of the 2015 International Symposium on Symbolic and Algebraic Computation. ISSAC '15. ACM, pp. 85–92.
URL <https://doi.org/10.1145/2755996.2756654>
- Brown, C., 2016. Bridging two communities to solve real problems. In: 18th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC '16). IEEE, pp. 11–14.
URL <https://doi.org/10.1109/SYNASC.2016.015>

- Buchberger, B., 1965. Ein algorithmus zum auffinden der basiselemente des restklassenrings nach einem nulldimensionalen polynomideal. (an algorithm for finding the basis elements of the residue class ring of a zerodimensional polynomial ideal). PhD thesis, Mathematical Institute, University of Innsbruck, see (Buchberger, 2006) for an English translation.
- Buchberger, B., 1985. Symbolic computation (an editorial). *Journal of Symbolic Computation* 1 (1), 1 – 6.
URL [https://doi.org/10.1016/S0747-7171\(85\)80025-0](https://doi.org/10.1016/S0747-7171(85)80025-0)
- Buchberger, B., 2006. Bruno Buchberger’s PhD thesis (1965): An algorithm for finding the basis elements of the residue class ring of a zero dimensional polynomial ideal. *Journal of Symbolic Computation* 41 (3-4), 475–511.
URL <https://doi.org/10.1016/j.jsc.2005.09.007>
- Caviness, B., Johnson, J., 1998. Quantifier Elimination and Cylindrical Algebraic Decomposition. *Texts & Monographs in Symbolic Computation*. Springer-Verlag.
URL <https://doi.org/10.1007/978-3-7091-9459-1>
- Collins, G., 1975. Quantifier elimination for real closed fields by cylindrical algebraic decomposition. In: *Proceedings of the 2nd GI Conference on Automata Theory and Formal Languages*. Springer-Verlag (reprinted in the collection Caviness and Johnson (1998)), pp. 134–183.
URL https://doi.org/10.1007/3-540-07407-4_17
- Collins, G., Hong, H., 1991. Partial cylindrical algebraic decomposition for quantifier elimination. *Journal of Symbolic Computation* 12, 299–328.
URL [https://doi.org/10.1016/S0747-7171\(08\)80152-6](https://doi.org/10.1016/S0747-7171(08)80152-6)
- Cook, S. A., 1971. The complexity of theorem-proving procedures. In: *Proceedings of the Third Annual ACM Symposium on Theory of Computing, STOC '71*. ACM, New York, NY, USA, pp. 151–158.
URL <http://doi.acm.org/10.1145/800157.805047>
- Davis, M., Logemann, G., Loveland, D., 1962. A machine program for theorem-proving. *Commun. ACM* 5 (7), 394–397.
URL <https://doi.org/10.1145/368273.368557>
- de Moura, L., Jovanović, D., 2013. A model-constructing satisfiability calculus. In: *Giacobazzi, R., Bertone, J., Mastroeni, I. (Eds.), Verification, Model Checking, and Abstract Interpretation (Proc. VMCAI 2013)*. Springer Berlin Heidelberg, pp. 1–12.
- England, M., Ganesh, V. (Eds.), 2017. *Proceedings of the 2nd Workshop on Satisfiability Checking and Symbolic Computation (SC² 2017)*. No. 1974 in *CEUR Workshop Proceedings*. Springer.
URL <http://ceur-ws.org/Vol-1974/>

- Forrest, S., 2017. Integration of smt-lib support into maple. In: England, M., Ganesh, V. (Eds.), Proceedings of the 2nd International Workshop on Satisfiability Checking and Symbolic Computation (SC² 2017). No. 1974 in CEUR Workshop Proceedings. p. 6.
URL <http://ceur-ws.org/Vol-1974/>
- Ganesh, V., O'Donnell, C. W., Soos, M., Devadas, S., Rinard, M. C., Solar-Lezama, A., 2012. Lynx: A programmatic sat solver for the rna-folding problem. In: Cimatti, A., Sebastiani, R. (Eds.), Theory and Applications of Satisfiability Testing – SAT 2012. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 143–156.
URL https://doi.org/10.1007/978-3-642-31612-8_12
- Huang, Z., England, M., Davenport, J., Paulson, L., 2016. Using machine learning to decide when to precondition cylindrical algebraic decomposition with Groebner bases. In: 18th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC '16). IEEE, pp. 45–52.
URL <https://doi.org/10.1109/SYNASC.2016.020>
- Jovanovic, D., de Moura, L., 2012. Solving non-linear arithmetic. In: Gramlich, B., Miller, D., Sattler, U. (Eds.), Automated Reasoning: 6th International Joint Conference (IJ-CAR). Vol. 7364 of Lecture Notes in Computer Science. Springer, pp. 339–354.
URL https://doi.org/10.1007/978-3-642-31365-3_27
- Kosta, M., 2016. New concepts for real quantifier elimination by virtual substitution. Ph.D. thesis, Saarland University. Available at: <http://scidok.sulb.uni-saarland.de/volltexte/2016/6716>.
- Kroening, D., Strichman, O., 2013. Decision Procedures: An Algorithmic Point of View. Springer, New York.
- Lahiri, S. K., Seshia, S. A., 2004. The UCLID decision procedure. In: Alur, R., Peled, D. A. (Eds.), Computer Aided Verification, 16th International Conference, CAV 2004, Boston, MA, USA, July 13-17, 2004, Proceedings. Vol. 3114 of Lecture Notes in Computer Science. Springer, pp. 475–478.
URL https://doi.org/10.1007/978-3-540-27813-9_40
- Marques-Silva, J., Sakallah, K., 1999. GRASP: A search algorithm for propositional satisfiability. Computers, IEEE Transactions on 48 (5), 506–521.
URL <https://doi.org/10.1109/12.769433>
- Moskewicz, M., Madigan, C., Zhao, Y., Zhang, L., Malik, S., 2001. Chaff: Engineering an Efficient SAT Solver. In: Proceedings 38th Design Automation Conference. pp. 530–535.
- Mulligan, C., Bradford, R., Davenport, J., England, M., Tonks, Z., 2018. Non-linear real arithmetic benchmarks derived from automated reasoning in economics. In: Bigatti, A., Brain, M. (Eds.), Proceedings of the 3rd Workshop on Satisfiability Checking and Symbolic Computation (SC² 2018). No. 2189 in CEUR Workshop Proceedings. pp. 48–60.
URL <http://ceur-ws.org/Vol-2189/>

- Risch, R., 1969. The problem of integration in finite terms. *Trans. A.M.S.* 139, 167–189.
- Sturm, T., 2017. A survey of some methods for real quantifier elimination, decision, and satisfiability and their applications. *Mathematics in Computer Science* 11 (3), 483–502.
URL <https://doi.org/10.1007/s11786-017-0319-z>
- Tarski, A., 1948. *A Decision Method For Elementary Algebra And Geometry*. RAND Corporation, Santa Monica, CA, see (Tarski, 1998) for a readily available reprint.
- Tarski, A., 1998. A decision method for elementary algebra and geometry. In: Caviness, B., Johnson, J. (Eds.), *Quantifier Elimination and Cylindrical Algebraic Decomposition*. Texts and Monographs in Symbolic Computation. Springer-Verlag, pp. 24–84.
URL https://doi.org/10.1007/978-3-7091-9459-1_3
- Tseitin, G., 1983. On the complexity of derivation in propositional calculus. In: Siekmann, J., Wrightson, G. (Eds.), *Automation of Reasoning: 2: Classical Papers on Computational Logic 1967–1970*. Springer Berlin Heidelberg, pp. 466–483.
URL https://doi.org/10.1007/978-3-642-81955-1_28
- Weispfenning, V., 1997. Quantifier elimination for real algebra — the quadratic case and beyond. *Applicable Algebra in Engineering, Communication and Computing* 8 (2), 85–101.
URL <https://doi.org/10.1007/s002000050055>