

SC²: Satisfiability Checking meets Symbolic Computation (Project Paper)

Erika Ábrahám¹, Bernd Becker², Anna Bigatti³, Bruno Buchberger⁴,
Alessandro Cimatti⁵, James H. Davenport⁶, Matthew England⁷,
Pascal Fontaine⁸, Stephen Forrest⁹, Daniel Kroening¹⁰,
Werner Seiler¹¹, and Thomas Sturm¹²

¹ RWTH Aachen University, Aachen, Germany

² Albert Ludwigs University, Freiburg, Germany

³ University of Genoa, Genoa, Italy

⁴ Johannes Kepler University, Linz, Austria

⁵ Fondazione Bruno Kessler, Trento, Italy

⁶ University of Bath, Bath, U.K.

⁷ Coventry University, Coventry, U.K.

⁸ LORIA, Nancy, France

⁹ Maplesoft Europe Ltd

¹⁰ University of Oxford, Oxford, U.K.

¹¹ University of Kassel, Kassel, Germany

¹² Max-Planck Institute for Informatics

Abstract. *Symbolic Computation* and *Satisfiability Checking* are two research areas, both having their individual scientific focus but sharing also common interests in the development, implementation and application of decision procedures for arithmetic theories. Despite their commonalities, the two communities are rather weakly connected. The aim of our newly accepted SC² project (H2020-FETOPEN-CSA) is to strengthen the connection between these communities by creating common platforms, initiating interaction and exchange, identifying common challenges, and developing a common roadmap from theory along the way to tools and (industrial) applications. In this paper we report on the aims and on the first activities of this project, and formalise some relevant challenges for the unified SC² community.

Keywords: Logical Problems, Satisfiability Checking, Satisfiability Modulo Theories Solving, Symbolic Computation, Computer Algebra Systems

1 Introduction

The use of advanced methods to solve practical and industrially relevant problems by computers has a long history. While it is customary to think that “computers are getting faster” (and indeed, they were, and are still getting more powerful in terms of multicores etc.), the progress in algorithms and software

has been even greater. One of the leaders in the field of linear and mixed integer programming points out [Bix15, slide 37] that you would be over 400 times better off running today’s algorithms and software on a 1991 computer than you would running 1991 software on today’s computer. The practice is heavily inspired by the theory: [Bix15, slide 31] shows that the biggest version-on-version performance advance in software was caused by “mining the theory”. *But* this progress has been in what is, mathematically, quite a limited domain: that of linear programming, possibly where some of the variables are integer-valued.

There has been also much progress in the use of computers to solve hard non-linear algebraic¹³ problems. This is the area generally called *Symbolic Computation* (or *Computer Algebra*), and can solve non-linear problems over both the real and complex numbers, though generally with very different techniques. This has produced many new applications and surprising developments: in an area everyone believed was solved, non-linear solving over the reals (using cylindrical algebraic decomposition — CAD) has recently found a new algorithm for computing square roots [EH14]. CAD is another area where practice is (sometimes) well ahead of theory: the theory [DH88, BD07] states that the complexity is doubly exponential in the number of variables, but useful problems can still be solved in practice ([AMIA14] points out that CAD is the most significant engine in the “Todai robot” project).

Independently and contemporaneously, there has been much practical progress in solving the SAT problem, i.e., checking the satisfiability of logical problems over the Boolean domain. The SAT problem is known to be NP-complete [Coo71]. Nevertheless, the *Satisfiability Checking* [BBH⁺09] community developed SAT solvers which can successfully handle inputs with millions of Boolean variables. Among other industrial applications, these tools are now at the heart of many techniques for verification and security of computer systems.

Driven by this success, big efforts were made to enrich propositional SAT-solving with solver modules for different theories. Highly interesting techniques were implemented in *SAT-modulo-theories (SMT) solvers* [BSST09, KS08] for checking easier theories, but the development for quantifier-free non-linear real and integer arithmetic¹³ is still in its infancy.

Figure 1 shows a non-exhaustive history of tool developments in these two areas. It illustrates nicely the historically deeper roots of computer algebra systems, but also the high intensity of research in both areas. The resulting tools are successfully applied in several academic and industrial areas, however, the current state is still not satisfactorily, as described in [PQR09]:

“Despite substantial advances in verification technology, complexity issues with classical decision procedures are still a major obstacle for formal verification of real-world applications, e.g., in automotive and avionic industries.”

¹³ It is usual in the SMT community to refer to these constraints as *arithmetic*. But, as they involve quantities as yet unknown, manipulating them is *algebra*. Hence both words occur, with essentially the same meaning, throughout this document.

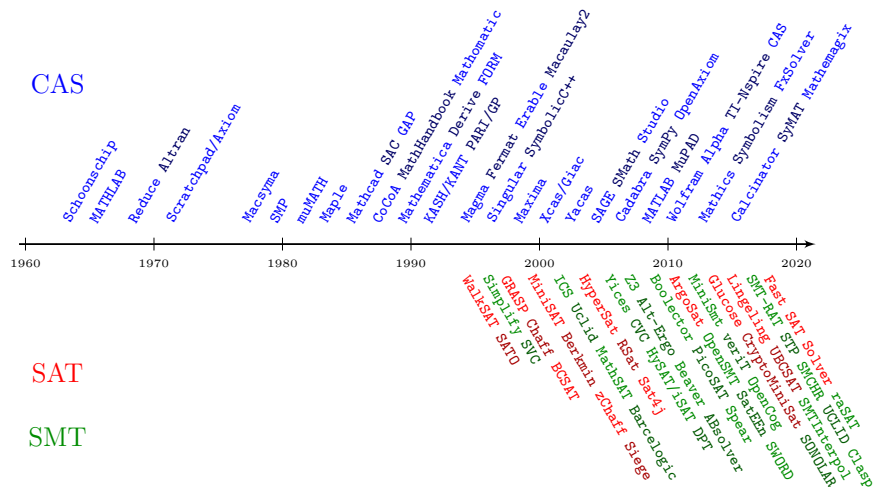


Fig. 1. History of some computer algebra systems and SAT/SMT solvers (not exhaustive; years approximate first release as far as known and as positioning allowed) [Á15]

Both communities address similar problems and share the challenge to improve their solutions to achieve applicability on complex large-scale applications. However, the Symbolic Computation community and the Satisfiability Checking community are largely in their own silos and traditionally do not interact much with each other.

To connect these communities, we successfully applied for a European Horizon 2020 *Coordination and Support Action*, with an envisaged project start in July 2016. The overall aim of this project is to create a new research community bridging the gap between Satisfiability Checking and Symbolic Computation, whose members will ultimately be well informed about both fields, and thus able to combine the knowledge and techniques of both fields to develop new research and to resolve problems (both academic and industrial) currently beyond the scope of either individual field. We call the new community SC^2 , as it will join the communities of **S**atisfiability **C**hecking and **S**ymbolic **C**omputation.

The contributions of this paper are twofold: Firstly, we discuss the potentials of closer connection and more intensive exchange between the two communities, and list a number of challenges that are currently out of reach but could be attacked by a unified SC^2 community (Section 3). Secondly, we discuss what is needed to trigger and support these developments, and describe the actions of our project to satisfy these needs (Section 4).

2 Background

Before describing our project, we give a short description of the state-of-the-art in Satisfiability Checking and Symbolic Computation. Parts of this section are taken from [Á15].

2.1 Symbolic Computation and Computer Algebra Systems

Computer Algebra, the use of computers to do algebra rather than simply arithmetic, is almost as old as computing itself, with the first theses dating back to 1953. This initial work consisted of programs to do one thing, but the focus soon moved on to ‘systems’, capable of doing a variety of tasks. One early one was Collins’ system **SAC** [Col71], written in Fortran. Many of the early systems were written in LISP, largely because of its support for garbage collection and large integers. The group at M.I.T., very influential in the 1960s and 70s, developed **Macsyma** [MF71]. This system now exists in several versions, even on mobile phones¹⁴. At about the same time, Hearn developed **Reduce** [Hea05], and shortly after a group at IBM Yorktown Heights produced **SCRATCHPAD**, then **AXIOM** [JS92], a system that attempted to match the generality of Mathematics with ‘generic programming’ to allow algorithms to be programmed in the generality in which they are conventionally stated, e.g. “polynomials over a ring”.

Symbolic Computation was initially seen as part of Artificial Intelligence, with major triumphs such as [Sla61] being “A Heuristic Program that Solves Symbolic Integration Problems in Freshman Calculus”, firmly in the AI camp. By the end of the 1960s, this approach to integration had been replaced by an algorithm [Mos67], which had the great advantage that, when backed up with a suitable completeness theorem [Ris69] it could *prove* unintegrability: “there is no formula made up of exponentials, logarithms and algebraic functions which differentiates to e^{-x^2} ”, in other words “ e^{-x^2} is unintegrable”.

The 1960s and 1970s also saw great advances in other areas. We had much more efficient algorithms to replace naïve use of Euclid’s algorithm for greatest common divisor computation (and hence the simplification of fractions), far better algorithms than the search-based methods for polynomial factorisation, and so on. All this meant that Symbolic Computation firmly moved into the camps of algorithmics and complexity theory, and the dominant question became “what is the worst-case complexity of this algorithm”.

Gröbner bases. One great success of this period was the method of *Gröbner bases* [Buc65]. This allows effective, and in many cases efficient, solution of many problems of polynomials over algebraically-closed fields (typically the complex numbers, though applications over finite fields and in cryptography abound). This notion paved the way for the discovery of numerous effective methods for polynomial ideals; many applications in other areas of Mathematics quickly followed. Buchberger’s algorithm for computing a Gröbner basis is a prime example of the huge gulf that can separate an abstract algorithm from a usable efficient implementation. Over the fifty years since its initial publication, research into the algorithm’s behaviour has produced several significant improvements: the modern refined version is typically thousands of times faster than the original. The search for further improvements continues today.

The remarkable computational utility of Gröbner bases prompted the development of a number of distinct, independent implementations of refined versions

¹⁴ <https://sites.google.com/site/maximaonandroid/>.

of Buchberger’s algorithm. The main commercial general-purpose computer algebra systems (including `MAGMA`, `Maple`, `Mathematica`) can all compute Gröbner bases; researchers needing the flexibility and ability to experiment with new algorithms also use specialised computer algebra systems such as `Singular` [DGPS15], `Macaulay/Macaulay2` [GS] and `CoCoA/CoCoALib` [ABL] which are freely downloadable from their respective websites.

Cylindrical algebraic decomposition. Another great success of the 1970s was the development of *cylindrical algebraic decomposition* (CAD) in [Col75]. This replaced the indescribably (no finite tower of exponentials bounds the complexity) complex methods of Tarski for real algebraic geometry. A CAD is a decomposition of \mathbb{R}^n into cells arranged cylindrically (meaning their projections are equal or disjoint) and described by semi-algebraic sets. For a detailed description of modern CAD, see [BDE⁺16].

Hong created a C version of both the `SAC` library and the comprehensive CAD code, which is now open-source and freely available as `SACLIB` and `QEPCAD-B` [Bro03]. Another example is the `RedLog` package [DS97] of the computer algebra system `Reduce`, which offers an optimised combination of the cylindrical algebraic decomposition with virtual substitution (see below) and Gröbner bases methods.

Virtual substitution To mention a last algorithm, the *virtual substitution* method [Wei97] focuses on non-linear real arithmetic formulas where the degree of the quantified variables is not too large. Although the method can be generalised to arbitrary degrees, current implementations are typically limited to input, where the total degree of the quantified variables does not exceed 2. In practice, this limitation is somewhat mildened by employing powerful heuristics like systematic *degree shifts* or polynomial factorisation. The key idea is to eliminate existential quantifiers in favour of *finite* disjunctions plugging in test terms that are derived from the considered formula.

These methods and their numerous refinements constitute the usual tool box of state-of-the-art computer algebra systems, and enable them to tackle hard arithmetic problems.

2.2 Satisfiability Checking

In the ’60s, another line of research on *Satisfiability Checking* [BBH⁺09] for *propositional logic* started its career. The first idea used *resolution* for quantifier elimination [DP60], and had serious problems with the explosion of the memory requirements with increasing problem size. Another research line [DLL62] suggested a combination of *enumeration* and *Boolean constraint propagation* (BCP). A major improvement was achieved in the ’90s by *combining* the two approaches, leading to *conflict-driven clause-learning* and *non-chronological backtracking* [MSS99]. Later on, this impressive progress was continued by novel efficient implementation techniques (e.g., sophisticated decision heuristics, two-watched-literal scheme, restarts, cache performance, etc.), resulting in numerous powerful *SAT solvers*.

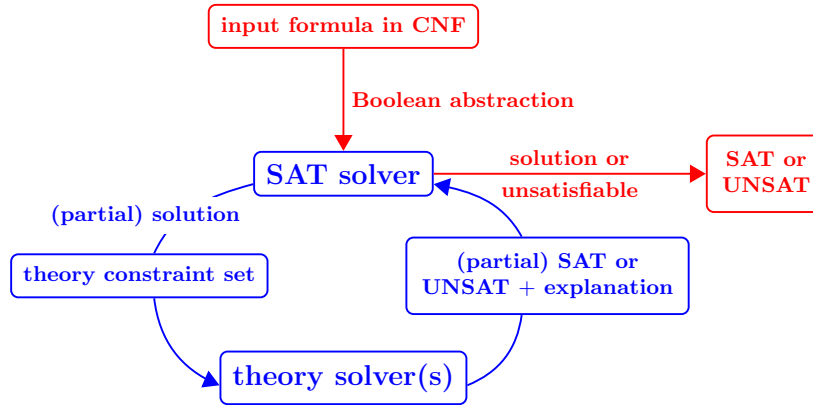


Fig. 2. The functioning of SMT solvers

Driven by this success, big efforts were made to enrich propositional SAT-solving with solver modules for different existentially quantified theories. Highly interesting techniques were implemented in *SAT-modulo-theories (SMT) solvers* for checking, e.g., equality logic with uninterpreted functions, array theory, bit-vector arithmetic and quantifier-free linear real and integer arithmetic, but the development for quantifier-free non-linear real and integer arithmetic is still in its infancy. For further reading, see, e.g., [BSST09, KS08].

Modern *SMT solvers* typically combine a *SAT solver* with one or more *theory solvers* as illustrated in Figure 2. First the input formula is transformed into conjunctive normal form (CNF), a conjunction of disjunctions (clauses); this transformation can be done in linear time and space using Tseitin’s transformation on the cost of additional variables. Next, the resulting CNF is abstracted to a pure Boolean propositional logic formula by replacing each theory constraint by a fresh Boolean proposition. Intuitively, the truth value of each fresh proposition defines whether the theory constraint, which it substitutes, holds. The SAT solver tries to find solutions for this propositional abstraction and during solving it consults the theory solver(s) to check the consistency of the theory constraints that should hold according to the current values of the abstraction variables.

On the one hand, theory solvers only need to check *conjunctions (sets)* of theory constraints, instead of arbitrary Boolean combinations. On the other hand, theory solvers should have the following properties for being *SMT-compliant*:

- They should work *incrementally*, i.e., after they determine the consistency of a constraint set, they should be able to take delivery of some additional constraints and re-check the extended set, thereby making use of results from the previous check.
- In case of unsatisfiability, they should be able to return an *explanation* for inconsistency, e.g., by a preferably small inconsistent subset of the constraints.
- They should support *backtracking*, i.e., the removal of previously added constraints.

Optimally, theory solvers should also be able to provide a *satisfying solution*, if the problem is satisfiable, and a *proof of unsatisfiability* for the explanation, if the problem is unsatisfiable.

A great advantage of the SMT technology is that it can employ decision procedures not only in isolation, but also *in combination*. For example, solving non-linear arithmetic formulas can often be speeded up by first checking linear abstractions or linear problem parts using more efficient decision procedures, before applying heavier procedures. Additionally, theories can also be combined already in the input language of SMT solvers. For example, deductive program verification techniques generate verification conditions, which might refer to arrays, bit-vectors as well as integers; in such cases, dedicated SMT solvers can apply several decision procedures for different theories in combination.

When combining decision procedures, *incomplete* but *efficient* procedures are also valuable, if they guarantee termination but not necessarily return a conclusive answer. Such incomplete methods are frequently applied in SMT solving, a typical example being interval constraint propagation, based on interval arithmetic. Some solvers combine such incomplete methods with complete decision procedures, in order to guarantee the solution of the problem, while increasing efficiency. Other solvers even sacrifice completeness and might return a “don’t know” answer, but still they are able to solve certain extremely large problems, which are out of reach for complete methods, very fast. Furthermore, incomplete procedures are the only way to support problems from undecidable theories, like formulas containing exponential or trigonometric functions.

SAT and SMT solvers are tuned for efficiency. Combining complete and incomplete decision procedures, making use of efficient heuristics, learning not only propositional facts but also (Boolean abstractions of) theory lemmas at the SAT level allow modern SMT solvers to solve relevant large-size problems with tens of thousands of variables, which could not be solved before by single decision procedures in isolation. For some example applications see, e.g., [BKM14].

Examples for solvers that are able to cope with linear arithmetic problems (either in a complete or in an incomplete manner) are **Alt-Ergo** [CIM13], **CVC4** [BCD⁺11], **iSAT3** [FHT⁺07,SKB13], **MathSAT** [CGSS13], **OpenSMT2** [BPST10], **SMT-RAT** [CKJ⁺15], **veriT** [BCd⁺09], **Yices2** [DdM06], and **Z3** [dMB08]. A further interesting SMT-approach for linear integer arithmetic is proposed in [BSW15].

Much less activity can be observed for SMT solvers for non-linear arithmetic. A few SMT tools embedded some (complete as well as incomplete) decision procedures. Such a solver is **iSAT3**, which uses interval constraint propagation. The SMT solver **MiniSmt** [ZM10] tries to reduce non-linear real arithmetic problems to linear real arithmetic and can solve only satisfiable instances this way. We are aware of only two SMT solvers that are complete for non-linear real arithmetic: Firstly, the prominent **Z3** solver developed at Microsoft Research, which uses an elegant SMT-adaptation of the cylindrical algebraic decomposition method [JdM12]. Secondly, **SMT-RAT** [CKJ⁺15], using solver modules for simplex, the cylindrical algebraic decomposition, the virtual substitution method,

Gröbner bases, interval constraint propagation, branch and bound, and their strategic combination [dMP13].

Even fewer SMT solvers are available for non-linear integer arithmetic, which is undecidable in general. A linearisation approach was proposed in [BLNM⁺09]. The SMT solving spin-off of **AProVE** [CFF⁺13] uses bit-blasting. To our knowledge, **Z3** implements a combination of linearisation and bit-blasting. **iSAT3** uses interval constraint propagation, whereas **Alt-Ergo** combines the idea of [BCC⁺12] with an axiom-based version of interval constraint propagation. **SMT-RAT** can tackle this theory using a generalised branch-and-bound technique.

The increasing variety of the theories considered by SMT solvers made an urgent need for a common input language. The *SMT-LIB* initiative [BST10] defined a *standard input format* for SMT solvers with a first release in 2004, and provides a large and still increasing number of *benchmarks*, systematically collected for all supported theories. *SMT-LIB* also enabled the start of *SMT competitions*; the first one took place in 2005 with 12 participating solvers in 7 divisions (theories, theory combinations, or fragments thereof) on 1360 benchmarks, which increased in 2014 to 20 solvers competing in 32 divisions on 67426 benchmarks. The *SMT-LIB* standard and the competitions not only intensified the SMT research activities, but also gave visibility and acceptance for SMT solving in computer science and beyond. Once a problem is formulated in the *SMT-LIB* language, the user can employ *any* SMT solver to solve the problem.

3 Some Scientific Challenges and Opportunities

On the one hand, SMT solving has its strength in efficient techniques for exploring Boolean structures, learning, combining solving techniques, and developing dedicated heuristics, but its current focus lies on easier theories and it makes use of Symbolic Computation results only in a rather naive way. There are fast SMT solvers available for the satisfiability check of linear real and integer arithmetic problems, but just a few can handle non-linear arithmetic. On the other hand, Symbolic Computation is strong in providing powerful procedures for sets (conjunctions) of arithmetic constraints, but it does not exploit the achievements in SMT solving for efficiently handling logical fragments, using heuristics and learning to speed-up the search for satisfying solutions.

The Satisfiability Checking community would definitely profit from further exploiting Symbolic Computation achievements and adapt and extend them to comply with the requirements on embedding in the SMT context. However, it is a highly challenging task, as it requires a deep understanding of complex mathematical problems, whose embedding in SMT solving is far from trivial and needs their adaptation and extension.

Symmetrically, Symbolic Computation could profit from exploiting successful SMT ideas, but it requires expertise in efficient solver technologies and their implementation, like dedicated data structures, sophisticated heuristics, effective learning techniques, and approaches for incrementality and explanation generation in theory solving modules.

In this section we describe some ideas how algorithms and tools from both communities could be made more powerful by exploiting scientific exchange and technology transfer.

3.1 Symbolic Computation Techniques for Satisfiability Checking

Many practical decision procedures, designed by the Symbolic Computation community, are implemented in computer algebra systems (*e.g.* linear real and integer arithmetic, non-linear real arithmetic, linear programming, quantified formulas, Gröbner and involutive bases). To use them in a Satisfiability Checking context, some scientific and engineering issues need solutions, notably to find new ways of *incremental* solving, *explaining* unsatisfiability and generating *lemmas*.

Whereas for linear real arithmetic useful procedures have been adopted to satisfy the requirements for SMT embedding, for non-linear arithmetic there are still many opportunities remain to be explored. For example, there are (to the best of our knowledge) just two SMT solvers, **Z3** and **SMT-RAT**, which make use of the CAD method, both in a different way: **Z3** uses a very elegant solution to explore the state space by a close integration of theory decisions and theory propagation in the Boolean SAT search, and construct CAD only partially to explain conflicts in the above search (more precisely, to compute a semi-algebraic description of CAD cells that do not satisfy a given sign condition). In contrast, **SMT-RAT** implements an incremental version of the CAD method, which works hand-in-hand with the search at the logical level. For this latter approach, the power of heuristics (variable and polynomial ordering for incremental projection, choice and order of sample points for lifting) and the generation of lemmas (most importantly the computation of explanations for unsatisfiability) is still far from being fully exploited.

There are still great potentials for improvements not only for the SMT-embedding of the CAD method, but also other non-linear arithmetic decision procedures like the virtual substitution method or Gröbner bases, and their strategic combination with each other and further light-weight methods such as interval constraint propagation.

Another important aspect is the Symbolic Computation community's expertise in simplification and preprocessing. The complexity of the problems this community handles is often extremely high, and no practical procedure would exist without significant techniques to prepare the input problems. Such techniques do exist too for Satisfiability Checking, but they rather focus on easier theories. A transfer of the *savoir-faire* in simplification and preprocessing for non-linear real and integer arithmetic would certainly be highly profitable.

3.2 Satisfiability Checking Techniques for Symbolic Computation

A key ingredient in the success of the Satisfiability Checking tools is the use of *learning* and *non-chronological backtracking* techniques to speed up the search through tree-shaped search structures. Traditionally (and in the majority of

cases) CAD proceeds through a two stage process: projecting the problem through lower real dimensions; then lifting: incrementally building a solution in increasing real dimensions. An alternative approach using triangular decomposition was introduced in [CMXY09] where first the complex space is cylindrically decomposed and then refined to a CAD of the real space, where all data is in a tree-shaped structure.

Other techniques are certainly amenable for learning with the non-chronological backtracking approach. For instance, first prototypes integrating CDCL-style learning techniques with virtual substitution for linear quantifier elimination have been successfully created and studied. Integration of learning techniques with the computation of comprehensive Gröbner bases [Wei92] should also be investigated.

Incrementality, which played an important role in the success of Satisfiability Checking, may also be used to make Symbolic Computation techniques more efficient. The alternative CAD construction method described above is also incremental in nature [BCD⁺14] and so may offer one option here. There exist algorithms for computing Gröbner bases which exploit known mathematical facts about the ideal generated by the basis like its Hilbert function or some syzygies. Traditionally, this has been seen as a way to speed up computations. However, these approaches can naturally be adapted into incremental algorithms.

A central aspect of Satisfiability Modulo Theories is the combination frameworks for theories and decision procedures. Combining theories in Symbolic Computation (combined real/floating point theories, interval constraint propagation with other arithmetic theories) might also bring a number of advantages and possibilities, for instance, more expressive languages, or efficiency due to hierarchical reasoning. While combination of theories in Symbolic Computation are typically very specific and ad-hoc, the SMT community systematically uses the generic Nelson–Oppen framework [NO79] for disjoint theories. Such a framework can of course not be used as it is, but it might be an inspiration for a modular approach in Symbolic Computation.

3.3 Standard Languages and Benchmarks

The initiation and maintenance of a common problem specification language standard *SMT-LIB* [BST10] and of competitions form an important part of the Satisfiability Checking community effort. Besides providing a stimulating event for tool developers to exhibit their systems, the competitions are also a vehicle for publishing practical progress. Competition results are advertised, and consulted by users to pick the best tools and techniques to solve problems.

The Symbolic Computation community does not have a similar tradition, and indeed, to quote one major system developer: “it is very hard to get any practical improvements published — the reviewers will often say this is not hard science”. Although it is not good to only focus on a small library of benchmarks and have the competition as sole goal, competitions do have a tremendously positive effect on tools and techniques, as witnessed in the Satisfiability Checking community,

especially if the competition challenges are concrete industrial challenges. Such driving forces could be also established in Symbolic Computation.

Though in Satisfiability Checking the standard input language allowed to provide large benchmark sets to the community, benchmarks for non-linear arithmetic theories are still rare, and harder to describe without ambiguity. Therefore, also the Satisfiability Checking community would profit from a common standard with an increased number of non-linear arithmetic benchmarks.

4 Project Actions

The solution of challenging problems, like mentioned in the previous section, could get within reach, when supported by a stronger collaboration between these research areas, creating an infrastructure for dialogue and knowledge transfer. However, the research areas of Satisfiability Checking and Symbolic Computation are still quite disconnected, as reflected in their communication platforms and support structures. Symbolic Computation has its own conferences (ACA, CASC, ISSAC, etc.), several dedicated journals, and the SIGSAM forum. Dually, Satisfiability Checking is supported by its own conferences and journals, the SatLive forum to keep up-to-date with research, SMT standards, and SAT- and SMT-solver competitions.

The main aims of our project are to create *communication platforms* and propose *standards* to enable the interaction between the two communities, and to use these platforms to *initiate discussions and cooperations* and to *identify potentials, challenges and obstacles* for future research and practical applications. In the following we shortly describe the actions of our SC^2 project to achieve these goals.

Communication platforms To create a unified SC^2 community, we will initiate platforms to support the interaction of currently disjoint groups. We organised a Dagstuhl Seminar *Symbolic Computation and Satisfiability Checking*¹⁵ November 15-20, 2015, which already led to numerous interesting discussions and interactions. At CASC 2016, we will organise a *topical session* devoted to topics from the cross-community SC^2 area. Furthermore, we will establish a workshop series in the area of SC^2 , covering the interests of both communities, and having its first edition affiliated with SYNASC 2016. These workshops will serve as platforms for scientific exchange, discussion and cooperation within and between the currently disjoint communities. To support and attract young new community members, we will organise a dedicated summer school aimed at interested young researchers from SC^2 areas, with courses specifically tailored to their needs.

Research roadmap The above platforms will initiate cross-community interactions, and help to clearly identify unused potentials. We aim at initiating discussions on what the communities can learn from each other, what are the common challenges which they can solve together, what Satisfiability Checking could learn

¹⁵ <http://www.dagstuhl.de/en/program/calendar/semhp/?semnr=15471>

from Symbolic Computation achievements, and which Satisfiability Checking results could be adapted to improve Symbolic Computation solutions.

Our long-term objective is to create a research roadmap of potentials and challenges, both to the two traditional subject silos, but also challenges that only the new joined SC^2 community can address. This roadmap should identify, within the problems currently faced in the industry, the particular points that can be expected to be solved by the SC^2 community in the short and middle term, and will provide recommendations for spin off projects.

Standards, benchmarks, competitions We ambition to create a standard problem specification language capable of representing common problems of the SC^2 community. We plan on extending the *SMT-LIB* language, which is already mature and fully accepted among the SMT (Satisfiability Checking) community, to handle features needed for the Symbolic Computation community. This will be done in a modular way, with a particular focus on extensibility for new features.

Agreeing on a common language, and being able to share challenging problems is an essential aspect for building a dynamic community. This will foster further discussions and uncover problems that can be solved by the SC^2 community altogether, set clear challenges on which various approaches can be evaluated, classify the approaches according to their strength and weaknesses on the various kinds of problems. Mixed approaches will naturally emerge, to tackle problems exhibiting several orthogonal difficulties. The standard could also serve as a communication protocol for platforms mixing tools, to build meta-tools to solve large and difficult problems out of reach of current techniques often specialised to just one kind of job.

How to become an associate? This project cannot reach its aims by involving just a small number of core project members. To be able to cover sufficiently wide research and application areas and to take into account their needs and interests, there are currently 37 SC^2 *associates* from both research communities as well as from industry. Our associates will be regularly informed about the project activities and they will be invited to take part in the corresponding events.

The SC^2 Coordination and Support Action will be an optimal platform for industrial and academic partners and associates to form smaller working groups and initiate specific projects. If you would like to participate in the project as an associate, please contact the Project Coordinator James Davenport¹⁶.

5 Conclusions and Future Work

In this paper we gave a short description of the aims and actions of our upcoming EU Coordination and Support Action SC^2 .

The SC^2 project will maintain a website (<http://www.sc-square.org>) making readily accessible all the public information of the project (e.g. contact information, details of past and forthcoming SC^2 workshops and other similar events).

¹⁶ Email contact: J.H.Davenport@bath.ac.uk

Acknowledgements

We are grateful for support by the H2020-FETOPEN-2016-2017-CSA project SC² (712689) and the ANR project SMaRT. Earlier work in this area was also supported by the EPSRC grant EP/J003247/1.

References

- [Á15] E. Abraham. Building bridges between symbolic computation and satisfiability checking. In *Proceedings ISSAC 2015*, pages 1–6. ACM, 2015.
- [ABL] J. Abbott, A. M. Bigatti, and G. Lagorio. CoCoA-5: A system for doing computations in commutative algebra. <http://cocoa.dima.unige.it>.
- [AMIA14] N. H. Arai, T. Matsuzaki, H. Iwane, and H. Anai. Mathematics by machine. In *Proceedings ISSAC 2014*, pages 1–8. ACM, 2014.
- [BBH⁺09] A. Biere, A. Biere, M. Heule, H. van Maaren, and T. Walsh. *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, 2009.
- [BCC⁺12] F. Bobot, S. Conchon, E. Contejean, A. Mahboubi, A. Mebsout, and G. Melquiond. A simplex-based extension of Fourier-Motzkin for solving linear integer arithmetic. In *Proceedings IJCAR 2012*, volume 7364 of *LNCS*, pages 67–81. Springer, 2012.
- [BCd⁺09] T. Bouton, D. Caminha, B. de Oliveira, D. Déharbe, and P. Fontaine. veriT: An open, trustable and efficient SMT-solver. In *Proceedings CADE-22*, volume 5663 of *LNCS*, pages 151–156. Springer, 2009.
- [BCD⁺11] C. Barrett, C. L. Conway, M. Deters, L. Hadarean, T. King D. Jovanović, A. Reynolds, and C. Tinelli. CVC4. In *Proceedings CAV 2011*, volume 6806 of *LNCS*, pages 171–177. Springer, 2011.
- [BCD⁺14] R. J. Bradford, C. Chen, J. H. Davenport, M. England, M. Moreno Maza, and D. J. Wilson. Truth table invariant cylindrical algebraic decomposition by regular chains. In *Proceedings CASC 2014*, volume 8660 of *LNCS*, pages 44–58. Springer, 2014.
- [BD07] C. W. Brown and J. H. Davenport. The complexity of quantifier elimination and cylindrical algebraic decomposition. In *Proceedings ISSAC 2007*, pages 54–60. ACM, 2007.
- [BDE⁺16] R.J. Bradford, J.H. Davenport, M. England, S. McCallum, and D.J. Wilson. Truth Table Invariant Cylindrical Algebraic Decomposition. *J. Symbolic Computation*, 76:1–35, 2016.
- [Bix15] R. E. Bixby. Computational progress in linear and mixed integer programming. *Presentation at ICIAM 2015*, 2015.
- [BKM14] C. Barrett, D. Kroening, and T. Melham. Problem solving for the 21st century: Efficient solvers for satisfiability modulo theories. Technical Report 3, London Mathematical Society and Smith Institute for Industrial Mathematics and System Engineering, 2014. Knowledge Transfer Report.
- [BLNM⁺09] C. Borralleras, S. Lucas, R. Navarro-Marset, E. Rodriguez-Carbonell, and A. Rubio. Solving non-linear polynomial arithmetic via SAT modulo linear arithmetic. In *Proceedings CADE-22*, volume 5663 of *LNCS*, pages 294–305. Springer, 2009.
- [BPST10] R. Bruttomesso, E. Pek, N. Sharygina, and A. Tsitovich. The OpenSMT2 solver. In *Proceedings TACAS 2010*, volume 6015 of *LNCS*, pages 150–153. Springer, 2010.

- [Bro03] C. W. Brown. QEPCAD B: A program for computing with semi-algebraic sets using CADs. *ACM SIGSAM Bulletin*, 37(4):97–108, 2003.
- [BSST09] C. Barrett, R. Sebastiani, S. A. Seshia, and C. Tinelli. Satisfiability modulo theories. In *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*, chapter 26, pages 825–885. IOS Press, 2009.
- [BST10] C. Barrett, A. Stump, and C. Tinelli. The satisfiability modulo theories library (*SMT-LIB*). www.SMT-LIB.org, 2010.
- [BSW15] M. Bromberger, T. Sturm, and C. Weidenbach. Linear integer arithmetic revisited. In *Proceedings CADE-25*, volume 9195 of *LNCS*, pages 623–637. Springer, 2015.
- [Buc65] B. Buchberger. *Ein Algorithmus zum Auffinden des basiselemente des Restklassenringes nach einem nulldimensionalen Polynomideal*. PhD thesis, University of Innsbruck, 1965. English translation: *J. Symbolic Computation*, 41:475–511, 2006.
- [CFF⁺13] M. Codish, Y. Fekete, C. Fuhs, J. Giesl, and J. Waldmann. Exotic semi-ring constraints. In *Proceedings SMT 2013*, volume 20 of *EPiC Series*, pages 88–97. EasyChair, 2013.
- [CGSS13] A. Cimatti, A. Griggio, B. Schaafsma, and R. Sebastiani. The **MathSAT5** smt solver. In *Proceedings TACAS 2013*, volume 7795 of *LNCS*, pages 93–107. Springer, 2013.
- [CIM13] S. Conchon, M. Iguernelala, and A. Mebsout. A collaborative framework for non-linear integer arithmetic reasoning in **Alt-Ergo**. In *Proceedings SYNASC 2013*, pages 161–168. IEEE, 2013.
- [CKJ⁺15] F. Corzilius, G. Kremer, S. Junges, S. Schupp, and E. Ábrahám. **SMT-RAT**: An open source C++ toolbox for strategic and parallel SMT solving. In *Proc. SAT 2015*, volume 9340 of *LNCS*, pages 360–368. Springer, 2015.
- [CMXY09] C. Chen, M. Moreno Maza, B. Xia, and L. Yang. Computing cylindrical algebraic decomposition via triangular decomposition. In *Proceedings ISSAC 2009*, pages 95–102. ACM, 2009.
- [Col71] G. E. Collins. The SAC-1 system: An introduction and survey. In *Proceedings SYMSAC 1971*, pages 144–152. ACM, 1971.
- [Col75] G. E. Collins. Quantifier elimination for real closed fields by cylindrical algebraic decomposition. In *Automata Theory and Formal Languages*, volume 33 of *LNCS*, pages 134–183. Springer, 1975.
- [Coo71] S. A. Cook. The complexity of theorem-proving procedures. In *Proceedings STOC 1971*, pages 151–158. ACM, 1971.
- [DdM06] B. Dutertre and L. de Moura. A fast linear-arithmetic solver for DPLL(T). In *Proc. CAV 2006*, volume 4144 of *LNCS*, pages 81–94. Springer, 2006.
- [DGPS15] W. Decker, G.-M. Greuel, G. Pfister, and H. Schönemann. **Singular** 4-0-2 — A computer algebra system for polynomial computations. <http://www.singular.uni-kl.de>, 2015.
- [DH88] J. H. Davenport and J. Heintz. Real quantifier elimination is doubly exponential. *J. Symbolic Computation*, 5:29–35, 1988.
- [DLL62] M. Davis, G. Logemann, and D. Loveland. A machine program for theorem-proving. *Communications of the ACM*, 5(7):394–397, 1962.
- [dMB08] L. M. de Moura and N. Bjørner. **Z3**: An efficient SMT solver. In *Proceedings TACAS 2008*, volume 4963 of *LNCS*, pages 337–340. Springer, 2008.
- [dMP13] L. de Moura and G. O. Passmore. The strategy challenge in SMT solving. In *Automated Reasoning and Mathematics*, pages 15–44. Springer, 2013.
- [DP60] M. Davis and H. Putnam. A computing procedure for quantification theory. *Journal of the ACM*, 7(3):201–215, 1960.

- [DS97] A. Dolzmann and T. Sturm. **Redlog**: Computer algebra meets computer logic. *ACM SIGSAM Bulletin*, 31(2):2–9, 1997.
- [EH14] M. Eraso and H. Hong. Synthesis of optimal numerical algorithms using real quantifier elimination (Case study: Square root computation). In *Proceedings ISSAC 2014*, pages 162–169. ACM, 2014.
- [FHT⁺07] M. Fränzle, C. Herde, T. Teige, S. Ratschan, and T. Schubert. Efficient solving of large non-linear arithmetic constraint systems with complex Boolean structure. *Journal on Satisfiability, Boolean Modeling and Computation*, 1(3-4):209–236, 2007.
- [GS] D. R. Grayson and M. E. Stillman. **Macaulay2**, a software system for research in algebraic geometry. <http://www.math.uiuc.edu/Macaulay2/>.
- [Hea05] A. C. Hearn. **REDUCE**: The first forty years. In *Proceedings A3L*, pages 19–24. Books on Demand GmbH, 2005.
- [JdM12] D. Jovanović and L. de Moura. Solving non-linear arithmetic. In *Proceedings IJCAR 2012*, volume 7364 of *LNAI*, pages 339–354. Springer, 2012.
- [JS92] R. D. Jenks and R. S. Sutor. *AXIOM: The Scientific Computation System*. Springer, 1992.
- [KS08] D. Kroening and O. Strichman. *Decision Procedures: An Algorithmic Point of View*. Springer, 2008.
- [MF71] W. A. Martin and R. J. Fateman. The **Macsyma** system. In *Proceedings SYMSAC 1971*, pages 59–75. ACM, 1971.
- [Mos67] J. Moses. *Symbolic Integration*. PhD thesis, M.I.T. & MAC TR-47, 1967.
- [MSS99] J. P. Marques-Silva and K. A. Sakallah. **GRASP**: A search algorithm for propositional satisfiability. *IEEE Trans. Computers*, 48:506–521, 1999.
- [NO79] G. Nelson and D. C. Oppen. Simplifications by cooperating decision procedures. *ACM Transactions on Programming Languages and Systems*, 1(2):245–257, 1979.
- [PQR09] A. Platzer, J.-D. Quesel, and P. Rümmer. Real world verification. In *Proceedings CADE-22*, pages 485–501. ACM, 2009.
- [Ris69] R. H. Risch. The problem of integration in finite terms. *Transactions of the American Mathematical Society*, 139:167–189, 1969.
- [SKB13] K. Scheibler, S. Kupferschmid, and B. Becker. Recent improvements in the SMT solver **iSAT**. In *Proceedings MBMV 2013*, pages 231–241. Institut für Angewandte Mikroelektronik und Datentechnik, Fakultät für Informatik und Elektrotechnik, Universität Rostock, 2013.
- [Sla61] J. Slagle. *A Heuristic Program that Solves Symbolic Integration Problems in Freshman Calculus*. PhD thesis, Harvard University, 1961.
- [Wei92] V. Weispfenning. Comprehensive Gröbner bases. *J. Symbolic Computation*, 14(1):1–29, 1992.
- [Wei97] V. Weispfenning. Quantifier elimination for real algebra – The quadratic case and beyond. *Applicable Algebra in Engineering, Communication and Computing*, 8(2):85–101, 1997.
- [ZM10] H. Zankl and A. Middeldorp. Satisfiability of non-linear (ir)rational arithmetic. In *Proceedings LPAR 2010*, volume 6355 of *LNAI*, pages 481–500. Springer, 2010.