



Cooperating Techniques for Solving Nonlinear Real Arithmetic in the *cvc5* SMT Solver (System Description)

Gereon Kremer¹, Andrew Reynolds², Clark Barrett¹,
and Cesare Tinelli²

¹ Stanford University, Stanford, USA

² The University of Iowa, Iowa City, USA
andrew.j.reynolds@gmail.com

Abstract. The *cvc5* SMT solver solves quantifier-free nonlinear real arithmetic problems by combining the cylindrical algebraic coverings method with incremental linearization in an abstraction-refinement loop. The result is a complete algebraic decision procedure that leverages efficient heuristics for refining candidate models. Furthermore, it can be used with quantifiers, integer variables, and in combination with other theories. We describe the overall framework, individual solving techniques, and a number of implementation details. We demonstrate its effectiveness with an evaluation on the SMT-LIB benchmarks.

Keywords: Satisfiability modulo theories · Nonlinear real arithmetic · Abstraction refinement · Cylindrical algebraic coverings

1 Introduction

SMT solvers are used as back-end engines for a wide variety of academic and industrial applications [2, 19, 20]. Efficient reasoning in the theory of real arithmetic is crucial for many such applications [5, 8]. While modern SMT solvers have been shown to be quite effective at reasoning about *linear* real arithmetic problems [21, 43], *nonlinear* problems are typically much more difficult. This is not surprising, given that the worst-case complexity for deciding the satisfiability of nonlinear real arithmetic formulas is doubly-exponential in the number of variables in the formula [15]. Nevertheless, a variety of techniques have been proposed and implemented, each attempting to target a class of formulas for which reasonable performance can be observed in practice.

Related Work. All complete decision procedures for nonlinear real arithmetic (or the *theory of the reals*) originate in computer algebra, the most prominent being cylindrical algebraic decomposition (CAD) [11]. While alternatives exist [6, 25, 41], they have not seen much use [27], and CAD-based methods are the only sound and complete methods in practical use today. CAD-based methods used

in modern SMT solvers include incremental CAD implementations [34, 36] and cylindrical algebraic coverings [3], both of which are integrated in the traditional CDCL(T) framework for SMT [40].

In contrast, the NLSAT [30] calculus and the generalized MCSAT [28, 39] framework provide for a much tighter integration of a conflict-driven CAD-based theory solver into a theory-aware core solver. This has been the dominant approach over the last decade due to its strong performance in practice. However, it has the significant disadvantage of being difficult to integrate with CDCL(T)-based frameworks for theory combination.

A number of *incomplete* techniques are also used by various SMT solvers: incremental linearization [9] gradually refines an abstraction of the nonlinear formula obtained via a naive linearization by refuting spurious models of the abstraction; interval constraint propagation [24, 36, 45] employs interval arithmetic to narrow down the search space; subtropical satisfiability [22] provides sufficient linear conditions for nonlinear solutions in the exponent space of the polynomials; and virtual substitution [12, 31, 46] makes use of parametric solution formulas for polynomials of bounded degree. Though all of these techniques have limitations, each of them is useful for certain subclasses of nonlinear real arithmetic or in combination with other techniques.

Contributions. We present an integration of cylindrical algebraic coverings and incremental linearization, implemented in the `cvc5` SMT solver. Crucial to the success of the integration is an abstraction-refinement loop used to combine the two techniques cooperatively. The solution is effective in practice, as witnessed by the fact that `cvc5` won the nonlinear real arithmetic category of SMT-COMP 2021 [44], the first time a non-MCSAT-based technique has won since 2013. Our integrated technique also has the advantage of being very flexible: in particular, it fits into the regular CDCL(T) schema for theory solvers and theory combination, it supports (mixed) integer problems, and it can be easily extended using further subsolvers that support additional arithmetic operators beyond the scope of traditional algebraic routines (e.g., transcendental functions).

2 Nonlinear Solving Techniques

The nonlinear arithmetic solver implemented in `cvc5` generally follows the abstraction-refinement framework introduced by Cimatti et al. [9] and depicted in Fig. 1. The input assertions are first checked by the linear arithmetic solver, where they are linearized implicitly by treating every application of multiplication as if it were an arithmetic variable. For example, given input assertions $x \cdot y > 0 \wedge x > 1 \wedge y < 0$, the linear solver treats the expression $x \cdot y$ as a variable. It may then find the (spurious) model: $x \mapsto 2$, $y \mapsto -1$, and $x \cdot y \mapsto 1$. We call the candidate model returned by the linear arithmetic solver, where applications of multiplication are treated as variables, a *linear model*. If a linear model does not exist, i.e., the input is unsatisfiable according to the linear solver, the linear solver generates a conflict that is immediately returned to the CDCL(T) engine.




Fig. 1. Structural overview of the nonlinear solver

When a linear model does exist, we check whether it already satisfies the input assertions or try to *repair* it to do so. We only apply a few very simple heuristics for repairs such as updating the value for z in the presence of a constraint like $z = x \cdot y$ based on the values of x and y .

If the model can not be repaired, we refine the abstraction for the linear solver [9]. This step constructs lemmas, or conflicts, based on the input assertions and the linear model, to advance the solving process by blocking either the current linear model or the current Boolean model, that is, the propositional assignment generated by the SMT solver’s SAT engine. The Boolean model is usually eliminated only by the coverings approach, while the incremental linearization technique generates lemmas with new literals that target the linear model, e.g., the lemma $x > 0 \wedge y < 0 \Rightarrow x \cdot y < 0$ in the example above. We next describe our implementation of cylindrical algebraic coverings and incremental linearization, and how they are combined in `cvc5`.

2.1 Cylindrical Algebraic Coverings

Cylindrical algebraic coverings is a technique recently proposed by Ábrahám et al. [3] and is heavily inspired by CAD. While the way the computation proceeds is very different from traditional CAD, and instead somewhat similar to NLSAT [30], their mathematical underpinnings are essentially identical. The cylindrical algebraic coverings subsolver in `cvc5` closely follows the presentation in [3]. Below, we discuss some differences and extensions. For this discussion, we must refer the reader to [3] for the relevant background material because of space constraints. We note that `cvc5` relies on the `libpoly` library [29] to provide most of the computational infrastructure for algebraic reasoning.

Square-Free Basis. As with most CAD projection schemas, the set of projection polynomials needs to be a square-free basis when computing the characterization for an interval in [3, Algorithm 4]. However, the resultants computed in this algorithm combine polynomials from different sets, which are not necessarily coprime. The remedy is to either make these sets of polynomials pairwise square-free or to fully factor all projection polynomials. We adopt the former approach.

Starting Model. Although the linear model may not satisfy the nonlinear constraints, we may expect it to be in the vicinity of a proper model. We thus

optionally use the linear model as an *initial assignment* for the cylindrical algebraic coverings algorithm in one of two ways: either using it initially in the search and discarding it as soon as it conflicts; or using it whenever possible, even if it leads to a conflict in another branch of the search. Unfortunately, neither technique has any discernible impact in our experiments.

Interval Pruning. As already noted in [3], a covering may contain two kinds of redundant intervals: intervals fully contained in another interval, or intervals contained in the union of other intervals. Removing the former kind of redundancies is not only clearly beneficial, but also required for how the characterizations are computed. It is not clear, however, if it is worthwhile to remove redundancies of the second kind because, while it can simplify the characterization locally, it may also make the resulting interval smaller, slowing down the overall solving process. Note that there may not be a unique redundant interval: e.g., if multiple intervals overlap, it may be possible to remove one of two intervals, but not both of them. We have implemented a simple heuristic to detect redundancies of the second kind, always removing the smallest interval with respect to the interval ordering given in [3]. Even if these redundancies occur in about 7.5% of all `QF_NRA` benchmarks, using this technique has only a very limited impact. It may be that for certain kinds of benchmarks, underrepresented in `SMT-LIB`, the technique is valuable. Or it may be that some variation of the technique is more broadly helpful. These are interesting directions for future work.

Lifting and Coefficient Selection with Lazard. The original cylindrical algebraic coverings technique is based on McCallum’s projection operator [37], which is particularly well-studied, but also (refutationally) unsound: polynomial nullification may occur when computing the real roots, possibly leading to the loss of real roots and thus solution candidates. One then needs to check for these cases and fall back to a more conservative, albeit more costly, projection schema such as those due to Collins [11] or Hong [26].

Lazard’s projection schema [35], which has been proven correct only recently [38], provides very small projection sets and is both sound and complete. This comes at the price of a different mathematical background and a modified lifting procedure, which corresponds to a modified procedure for real root isolation. Although the local projections employed in cylindrical algebraic coverings have not been formally verified for Lazard’s projection schema yet, we expect no significant issues there. Adopting it seems to be a logical improvement, as already mentioned in [3]. The modified real root isolation procedure is a significant hurdle in practice, as it requires additional nontrivial algorithms [32, Section 5.3.2]. We implemented it using `CoCoALib` [1] in `cvc5` [33], achieving soundness without any discernible negative performance impact.

Using Lazard’s projection schema, for all its benefits, may seem questionable for the following reasons: (i) the unsoundness of McCallum’s projection operator is virtually never witnessed in practice [32, 33, Section 6.5], and (ii) the projection sets computed by Lazard’s and McCallum’s projection operator are identical on more than 99.5% on all of `QF_NRA` [33]. We argue, though, that working in

the domain of formal verification warrants the effort of obtaining a (provably) correct result, especially if it does not incur a performance overhead.

Proof Generation. Recently, generating formal proofs to certify the result of SMT solvers has become an area of focus. In particular, there is a large and ongoing effort to produce proofs in `cvc5`. The incremental linearization approach can be seen as an oracle which produces lemmas that are easy to prove individually, so `cvc5` does generate proofs for them; the complex part is finding those lemmas and making sure they actually help the solver make progress.

The situation is very different for cylindrical algebraic coverings: the produced lemma is the infeasible subset, and we usually have no simpler proof than the computations relying on CAD theory. That said, cylindrical algebraic coverings appear to be more amenable to automatic proof generation than traditional CAD-based approaches [4, 14]. In fact, although making these proofs detailed enough for automated verification is still an open problem, they are already broken into smaller parts that closely follow the tree-shaped computation of the algorithm. This allows `cvc5` to produce at least a proof skeleton in that case.

2.2 Incremental Linearization

Our theory solver for nonlinear (real) arithmetic optionally uses lemma schemas following the incremental linearization approaches described by Cimatti et al. [9] and Reynolds et al. [42]. These schemas incrementally refine candidate models from the linear arithmetic solver by introducing selected quantifier-free lemmas that express properties of multiplication, such as signedness (e.g., $x > 0 \wedge y > 0 \Rightarrow x \cdot y > 0$) or monotonicity (e.g., $|x| > |y| \Rightarrow x \cdot x > y \cdot y$). They are generated as needed to refute spurious models that violate these properties.

Most lemma schemas built-in in `cvc5` are crafted so as to avoid introducing new monomial terms or coefficients, since that could lead to non-termination in the CDCL(T) search. As a notable exception, we rely on a lemma schema for *tangent planes* for multiplication [9], which can be used to refute the candidate model for any application of the multiplication operator \cdot whose value in the linear model is inconsistent with the standard interpretation of \cdot . Note that since these lemmas depend upon the current model value chosen for arithmetic variables, tangent plane lemmas may introduce an unbounded number of new literals into the search. The set of lemma schemas used by the solver is user-configurable, as described in the following section.

2.3 Strategy

The overall theory solver for nonlinear arithmetic is built from several subsolvers, implementing the techniques described above, using a rather naive strategy, as summarized in Algorithm 1. After a spurious linear model has been constructed that cannot be repaired, we first apply a subset of the lemma schemas that do not introduce an unbounded number of new terms (with procedure `IncLinearizationLight`); then, we continue with the remaining lemma schemas

```

1 Function NlSolve(assertions)
2   if not LinearSolve(assertions) then return linear conflict
3   M = linear model for assertions
4   if RepairModel(assertions, M) then return repaired model
5   if IncLinearizationLight(assertions, M) then return lemmas
6   if IncLinearizationFull(assertions, M) then return lemmas
7   return Coverings(assertions, M)

```

Algorithm 1: Strategy for nonlinear arithmetic solver

(with procedure `IncLinearizationFull`); finally, we resort to the coverings solver which is guaranteed to find either a conflict or a model. Internally, each procedure sequentially tries its assigned lemma schemas from [9, 42] until it constructs a lemma that can block the spurious model.

The approach is dynamically configured based on input options and the logic of the input formula. For example, by default, we disable `IncLinearizationFull` for `QF_NRA` as it tends to diverge in cases where the coverings solver quickly terminates.

2.4 Beyond QF_NRA

The presented solver primarily targets quantifier-free nonlinear real arithmetic, but is used also in the presence of quantifiers and with multiple theories.

Quantified Logics. Solving quantified logics for nonlinear arithmetic requires solving quantifier-free subproblems, and thus any improvement to quantifier-free solving also benefits solving with quantifiers. In practice, however, the instantiation heuristics are just as important for overall solver performance.

Multiple Theories. The theory combination framework as implemented in `cvc5` requires evaluating equalities over the combined model. To support this functionality, real algebraic numbers had to be properly integrated into the entire solver; in particular, the ability to compute with these numbers could not be local to the cylindrical algebraic coverings module or even the nonlinear solver.

3 Experimental Results

We evaluate our implementation within `cvc5` (commit id 449dd7e) in comparison with other SMT solvers on all 11552 benchmarks in the quantifier-free nonlinear real arithmetic (`QF_NRA`) logic of SMT-LIB. We consider three configurations of `cvc5`, each of which runs a subset of steps from Algorithm 1. All the configurations run lines 2–4. In addition, `cvc5.cov` runs line 7, `cvc5.inclin` runs lines 5 and 6, and `cvc5` runs lines 5 and 7. All experiments were conducted on Intel Xeon E5-2637v4 CPUs with a time limit of 20 min and 8 GB memory.

We compare `cvc5` with recent versions of all other SMT solvers that participated in the `QF_NRA` logic of SMT-COMP 2021 [44]: `MathSAT` 5.6.6 [10], `SMT-RAT` 19.10.560 [13], `veriT` [7] (`veriT+raSAT+Redlog`), `Yices2` 2.6.4 [18] (`Yices-QS` for

quantified logics), and `z3` 4.8.14 [16]. `MathSAT` employs an abstraction-refinement mechanism very similar to the one described in Sect. 2.2; `veriT` [23] forwards nonlinear arithmetic problems to the external tools `raSAT` [45], which uses interval constraint propagation, and `Redlog/Reduce` [17], which focuses on virtual substitution and cylindrical algebraic decomposition; `SMT-RAT`, `Yices2`, and `z3` all implement some variant of `MCSAT` [30]. Note that `SMT-RAT` also implements the cylindrical algebraic coverings approach, but it is less effective than `SMT-RAT`'s adaptation of `MCSAT` [3].

				Beyond QF_NRA		sat unsat solved		
QF_NRA				NRA	Yices2	z3	3817	4048
	sat	unsat	solved		z3	236	3812	4048
<code>cvc5</code>	5137	5596	10733		<code>cvc5.cov</code>	236	3809	4045
<code>Yices2</code>	4966	5450	10416		<code>cvc5</code>	221	3809	4030
<code>z3</code>	5136	5207	10343		<code>cvc5.inclin</code>	120	3786	3906
<code>cvc5.cov</code>	5001	5077	10078	QF_UFNRA	z3	24	11	35
<code>SMT-RAT</code>	4828	5038	9866		<code>Yices2</code>	23	11	34
<code>veriT</code>	4522	5034	9556		<code>cvc5</code>	20	11	31
<code>MathSAT</code>	3645	5357	9002		<code>cvc5.inclin</code>	12	11	23
<code>cvc5.inclin</code>	3421	5376	8797		<code>cvc5.cov</code>	2	11	13

(a)

(b)

Fig. 2. (a) Experiments for QF_NRA (b) Experiments for NRA and QF_UFNRA

Figure 2a shows that `cvc5` significantly outperforms all other QF_NRA solvers. Both the coverings approach (`cvc5.cov`) and the incremental linearization approach (`cvc5.inclin`) contribute substantially to the overall performance of the unified solver in `cvc5`, with coverings solving many satisfiable instances, and incremental linearization helping on unsatisfiable ones. Even though `cvc5.inclin` closely follows [9], it outperforms `MathSAT` on unsatisfiable benchmarks, those where `cvc5` relies on incremental linearization the most.

Comparing `cvc5` and `Yices2` is particularly interesting, as the coverings approach in `cvc5` and the NLSAT solver in `Yices2` both rely on `libpoly` [29], thus using the same implementation of algebraic numbers and operations over them. Our integration of incremental linearization and algebraic coverings is compatible with the traditional CDCL(T) framework and outperforms the alternative NLSAT approach, which is specially tailored to nonlinear real arithmetic.

Going beyond QF_NRA, we also evaluate the performance of our solver in the context of theory combination (with all 37 benchmarks from QF_UFNRA) and quantifiers (with all 4058 benchmarks from NRA). There, `cvc5` is a close runner-up to `Yices2` and `z3`, thanks to the coverings subsolver which significantly improves `cvc5`'s performance. We conjecture that the remaining gap is due to components other than the nonlinear arithmetic solver, such as the solver for equality and uninterpreted functions, details of theory combination, or quantifier instantiation

heuristics. Interestingly, the sets of unsolved instances in **NRA** are almost disjoint for **cvc5.cov**, **Yices2** and **z3**, indicating that each tool could solve the remaining benchmarks with reasonable extra effort.

4 Conclusion

We have presented an approach for solving quantifier-free nonlinear real arithmetic problems that combines previous approaches based on incremental linearization [9] and cylindrical algebraic coverings [3] into one coherent abstraction-refinement loop. The resulting implementation is very effective, outperforming other state-of-the-art solver implementations, and integrates seamlessly in the CDCL(T) framework.

The general approach also applies to integer problems, quantified formulas, and instances with multiple theories, and can additionally be used in combination with transcendental functions [9] and bitwise conjunction for integers [47]. Further evaluations of these combinations are left to future work.

References

1. Abbott, J., Bigatti, A.M., Palezzato, E.: New in CoCoA-5.2.4 and CoCoALib-0.99600 for SC-square. In: Satisfiability Checking and Symbolic Computation. CEUR Workshop Proceedings, vol. 2189, pp. 88–94. CEUR-WS.org (2018). <http://ceur-ws.org/Vol-2189/paper4.pdf>
2. Ábrahám, E., Corzilius, F., Johnsen, E.B., Kremer, G., Mauro, J.: Zephyrus2: on the fly deployment optimization using SMT and CP technologies. In: Fränzle, M., Kapur, D., Zhan, N. (eds.) SETTA 2016. LNCS, vol. 9984, pp. 229–245. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-47677-3_15
3. Ábrahám, E., Davenport, J.H., England, M., Kremer, G.: Deciding the consistency of non-linear real arithmetic constraints with a conflict driven search using cylindrical algebraic coverings. *J. Logic. Algebr. Methods Program.* **119**(100633) (2021). <https://doi.org/10.1016/j.jlamp.2020.100633>
4. Ábrahám, E., Davenport, J.H., England, M., Kremer, G.: Proving UNSAT in SMT: the case of quantifier free non-linear real arithmetic. arXiv preprint [arXiv:2108.05320](https://arxiv.org/abs/2108.05320) (2021)
5. Arnett, T.J., Cook, B., Clark, M., Rattan, K.: Fuzzy logic controller stability analysis using a satisfiability modulo theories approach. In: 19th AIAA Non-Deterministic Approaches Conference, p. 1773 (2017)
6. Basu, S., Pollack, R., Roy, M.F.: On the combinatorial and algebraic complexity of quantifier elimination. *J. ACM* **43**, 1002–1045 (1996). <https://doi.org/10.1145/235809.235813>
7. Bouton, T., Caminha B. de Oliveira, D., Déharbe, D., Fontaine, P.: veriT: an open, trustable and efficient SMT-solver. In: Schmidt, R.A. (ed.) CADE 2009. LNCS (LNAI), vol. 5663, pp. 151–156. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-02959-2_12
8. Cashmore, M., Magazzeni, D., Zehtabi, P.: Planning for hybrid systems via satisfiability modulo theories. *J. Artif. Intell. Res.* **67**, 235–283 (2020). <https://doi.org/10.1613/jair.1.11751>

9. Cimatti, A., Griggio, A., Irfan, A., Roveri, M., Sebastiani, R.: Incremental linearization for satisfiability and verification modulo nonlinear arithmetic and transcendental functions. *ACM Trans. Comput. Logic* **19**, 19:1–19:52 (2018). <https://doi.org/10.1145/3230639>
10. Cimatti, A., Griggio, A., Schaafsma, B.J., Sebastiani, R.: The MathSAT5 SMT solver. In: Piterman, N., Smolka, S.A. (eds.) *TACAS 2013*. LNCS, vol. 7795, pp. 93–107. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36742-7_7
11. Collins, G.E.: Quantifier elimination for real closed fields by cylindrical algebraic decomposition. In: Brakhage, H. (ed.) *GI-Fachtagung 1975*. LNCS, vol. 33, pp. 134–183. Springer, Heidelberg (1975). https://doi.org/10.1007/3-540-07407-4_17
12. Corzilius, F., Abraham, E.: Virtual substitution for SMT-solving. In: Owe, O., Steffen, M., Telle, J.A. (eds.) *FCT 2011*. LNCS, vol. 6914, pp. 360–371. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22953-4_31
13. Corzilius, F., Kremer, G., Junges, S., Schupp, S., Abraham, E.: SMT-RAT: an open source C++ toolbox for strategic and parallel SMT solving. In: Heule, M., Weaver, S. (eds.) *SAT 2015*. LNCS, vol. 9340, pp. 360–368. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24318-4_26
14. Davenport, J., England, M., Kremer, G., Tonks, Z., et al.: New opportunities for the formal proof of computational real geometry? arXiv preprint [arXiv:2004.04034](https://arxiv.org/abs/2004.04034) (2020)
15. Davenport, J.H., Heintz, J.: Real quantifier elimination is doubly exponential. *J. Symb. Comput.* **5**(1), 29–35 (1988). [https://doi.org/10.1016/S0747-7171\(88\)80004-X](https://doi.org/10.1016/S0747-7171(88)80004-X)
16. de Moura, L., Bjørner, N.: Z3: an efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) *TACAS 2008*. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78800-3_24
17. Dolzmann, A., Sturm, T.: REDLOG: computer algebra meets computer logic. *ACM SIGSAM Bull.* **31**(2), 2–9 (1997). <https://doi.org/10.1145/261320.261324>
18. Dutertre, B.: Yices 2.2. In: Biere, A., Bloem, R. (eds.) *CAV 2014*. LNCS, vol. 8559, pp. 737–744. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-08867-9_49
19. Ermon, S., Le Bras, R., Gomes, C.P., Selman, B., van Dover, R.B.: SMT-aided combinatorial materials discovery. In: Cimatti, A., Sebastiani, R. (eds.) *SAT 2012*. LNCS, vol. 7317, pp. 172–185. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-31612-8_14
20. Fagerberg, R., Flamm, C., Merkle, D., Peters, P.: Exploring chemistry using SMT. In: Milano, M. (ed.) *CP 2012*. LNCS, pp. 900–915. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33558-7_64
21. Faure, G., Nieuwenhuis, R., Oliveras, A., Rodríguez-Carbonell, E.: SAT modulo the theory of linear arithmetic: exact, inexact and commercial solvers. In: Kleine Büning, H., Zhao, X. (eds.) *SAT 2008*. LNCS, vol. 4996, pp. 77–90. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-79719-7_8
22. Fontaine, P., Ogawa, M., Sturm, T., Vu, X.T.: Subtropical satisfiability. In: Dixon, C., Finger, M. (eds.) *FroCoS 2017*. LNCS (LNAI), vol. 10483, pp. 189–206. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66167-4_11
23. Fontaine, P., Ogawa, M., Sturm, T., Vu, X.T., et al.: Wrapping computer algebra is surprisingly successful for non-linear SMT. In: *SC-square 2018-Third International Workshop on Satisfiability Checking and Symbolic Computation* (2018)
24. Gao, S., Kong, S., Clarke, E.M.: dReal: an SMT solver for nonlinear theories over the reals. In: Bonacina, M.P. (ed.) *CADE 2013*. LNCS (LNAI), vol. 7898, pp. 208–214. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38574-2_14

25. Grigor'ev, D.Y., Vorobjov, N.: Solving systems of polynomial inequalities in subexponential time. *J. Symb. Comput.* **5**, 37–64 (1988). [https://doi.org/10.1016/S0747-7171\(88\)80005-1](https://doi.org/10.1016/S0747-7171(88)80005-1)
26. Hong, H.: An improvement of the projection operator in cylindrical algebraic decomposition. In: *International Symposium on Symbolic and Algebraic Computation*, pp. 261–264 (1990). <https://doi.org/10.1145/96877.96943>
27. Hong, H.: Comparison of several decision algorithms for the existential theory of the reals. RES report, Johannes Kepler University (1991)
28. Jovanović, D., Barrett, C., de Moura, L.: The design and implementation of the model constructing satisfiability calculus. In: *Formal Methods in Computer-Aided Design*, pp. 173–180. IEEE (2013). <https://doi.org/10.1109/FMCAD.2013.7027033>
29. Jovanovic, D., Dutertre, B.: LibPoly: a library for reasoning about polynomials. In: *Satisfiability Modulo Theories. CEUR Workshop Proceedings*, vol. 1889. CEUR-WS.org (2017). <http://ceur-ws.org/Vol-1889/paper3.pdf>
30. Jovanović, D., de Moura, L.: Solving non-linear arithmetic. In: Gramlich, B., Miller, D., Sattler, U. (eds.) *IJCAR 2012. LNCS (LNAI)*, vol. 7364, pp. 339–354. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-31365-3_27
31. Košta, M., Sturm, T.: A generalized framework for virtual substitution. *CoRR* abs/1501.05826 (2015)
32. Kremer, G.: Cylindrical algebraic decomposition for nonlinear arithmetic problems. Ph.D. thesis, RWTH Aachen University (2020). <https://doi.org/10.18154/RWTH-2020-05913>
33. Kremer, G., Brandt, J.: Implementing arithmetic over algebraic numbers: a tutorial for Lazard's lifting scheme in CAD. In: *Symbolic and Numeric Algorithms for Scientific Computing*, pp. 4–10 (2021). <https://doi.org/10.1109/SYNASC54541.2021.00013>
34. Kremer, G., Ábrahám, E.: Fully incremental cylindrical algebraic decomposition. *J. Symb. Comput.* **100**, 11–37 (2020). <https://doi.org/10.1016/j.jsc.2019.07.018>
35. Lazard, D.: An improved projection for cylindrical algebraic decomposition. In: Bajaj, C.L. (ed.) *Algebraic Geometry and Its Applications*, pp. 467–476. Springer, New York (1994). https://doi.org/10.1007/978-1-4612-2628-4_29
36. Loup, U., Scheibler, K., Corzilius, F., Ábrahám, E., Becker, B.: A symbiosis of interval constraint propagation and cylindrical algebraic decomposition. In: Bonacina, M.P. (ed.) *CADE 2013. LNCS (LNAI)*, vol. 7898, pp. 193–207. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38574-2_13
37. McCallum, S.: An improved projection operation for cylindrical algebraic decomposition. In: Caviness, B.F. (ed.) *EUROCAL 1985. LNCS*, vol. 204, pp. 277–278. Springer, Heidelberg (1985). https://doi.org/10.1007/3-540-15984-3_277
38. McCallum, S., Parusiński, A., Paunescu, L.: Validity proof of Lazard's method for cad construction. *J. Symb. Comput.* **92**, 52–69 (2019). <https://doi.org/10.1016/j.jsc.2017.12.002>
39. de Moura, L., Jovanović, D.: A model-constructing satisfiability calculus. In: Giacobazzi, R., Berdine, J., Mastroeni, I. (eds.) *VMCAI 2013. LNCS*, vol. 7737, pp. 1–12. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-35873-9_1
40. Nieuwenhuis, R., Oliveras, A., Tinelli, C.: Solving SAT and SAT modulo theories: from an abstract Davis-Putnam-Logemann-Loveland procedure to DPLL(T). *J. ACM* **53**(6), 937–977 (2006). <https://doi.org/10.1145/1217856.1217859>
41. Renegar, J.: A faster PSPACE algorithm for deciding the existential theory of the reals. In: *Symposium on Foundations of Computer Science*, pp. 291–295 (1988). <https://doi.org/10.1109/SFCS.1988.21945>

42. Reynolds, A., Tinelli, C., Jovanović, D., Barrett, C.: Designing theory solvers with extensions. In: Dixon, C., Finger, M. (eds.) *FroCoS 2017*. LNCS (LNAI), vol. 10483, pp. 22–40. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66167-4_2
43. Roselli, S.F., Bengtsson, K., Åkesson, K.: SMT solvers for job-shop scheduling problems: models comparison and performance evaluation. In: *International Conference on Automation Science and Engineering (CASE)*, pp. 547–552 (2018). <https://doi.org/10.1109/COASE.2018.8560344>
44. SMT-COMP 2021 (2021). <https://smt-comp.github.io/2021/>
45. Tung, V.X., Van Khanh, T., Ogawa, M.: raSAT: an SMT solver for polynomial constraints. In: Olivetti, N., Tiwari, A. (eds.) *IJCAR 2016*. LNCS (LNAI), vol. 9706, pp. 228–237. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-40229-1_16
46. Weispfenning, V.: Quantifier elimination for real algebra—the quadratic case and beyond. *Appl. Algebra Eng. Commun. Comput.* **8**(2), 85–101 (1997). <https://doi.org/10.1007/s002000050055>
47. Zohar, Y., et al.: Bit-precise reasoning via Int-blasting. In: Finkbeiner, B., Wies, T. (eds.) *VMCAI 2022*. LNCS, vol. 13182, pp. 496–518. Springer, Cham (2022). https://doi.org/10.1007/978-3-030-94583-1_24

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

