SMT and Quantifier Elimination: the Nonlinear Real Arithmetic case

James Davenport
masjhd@bath.ac.uk

University of Bath
Partially supported by EPSRC grant EP/T015713

16 October 2023

Setting

Quite often Satisfiability Modulo Theories (SMT), but also more general Quantifier Elimination (QE) settings.

Table: My dictionary

SMT	Computer Algebra
Real Arithmetic	Polynomial Algebra
Note that neither really likes division	
SATisfiable	A witness to the Variety $\neq \emptyset$
UNSATisfiable	Variety $= \emptyset$
Quantifier-free	All variables under $\exists$

Cylindrical Algebraic Decomposition

Problem (Quantifier Elimination)

Given a quantified statement about polynomials $f_i \in \mathbf{Q}[x_1, \ldots, x_n]$

$$\Phi_j := Q_{j+1}x_{j+1} \cdots Q_n x_n \Phi(f_i) \qquad Q_i \in \{\forall, \exists\} \tag{1}$$

produce an equivalent $\Psi(g_i) : g_i \in \mathbf{Q}[x_1, \ldots, x_j]$*: "equivalent"* $\equiv$ *"same real solutions".*

Solution [Col75]: produce a Cylindrical Algebraic Decomposition of $\mathbf{R}^n$ such that each f_i is sign-invariant on each cell, and the cells are *cylindrical*: $\forall i, \alpha, \beta$ the projections $P_{x_1,\ldots,x_i}(C_\alpha)$ and $P_{x_1,\ldots,x_i}(C_\beta)$ are equal or disjoint. Each cell C_i has a sample point s_i (again cylindrical) and then the truth of Φ in a cell is the truth at a sample point, and $\forall x_r$ becomes $\bigwedge_{x_r \text{ samples}}$ etc.

Plus/Minus of CAD

+ Solves the problem given, e.g. $\forall x \exists y f > 0 \wedge (g = 0 \vee h < 0)$
− The same structure solves all other problems with the same polynomials and order of quantified variables, e.g. $\forall y f = 0 \vee (g < 0 \wedge h > 0)$
− Current algorithms can be misled by spurious solutions. Consider $\{x^2 + y^2 - 2, (x-6)^2 + y^2 - 2\}$. Because $x = 3, y = \pm\sqrt{-7}$ is a common zero, current algorithms wrongly regard $x = 3$ as a critical point (which it would be over $\mathbf{C}^2$).

The original complexity

When Collins [Col75] produced his Cylindrical Algebraic Decomposition algorithm, the complexity was $O\left(d^{2^{2n+8}} m^{2^{n+6}}\right) l^3 k$, where n is the number of variables, d the maximum degree of any input polynomial in any variable, m the number of polynomials occurring in the input, k the number of occurrences of polynomials (essentially the length) and l the maximum coefficient length.

From now on omit l, k, and assume classical arithmetic.

Given m polynomials of degree d in x_n, we consider P_C:

1. $O(md)$ coefficients (degree $\leq d$)
2. $O(md)$ discriminants and subdiscriminants (degree $\leq 2d^2$)
3. $O(m^2d)$ resultants and subresultants (degree $\leq 2d^2$)

Then make square-free etc., and repeat.

$$(m, d) \Rightarrow (m^2d, 2d^2) \Rightarrow (2m^4d^4, 8d^4) \Rightarrow (32m^8d^{12}, 128d^8) \Rightarrow \cdots$$

This feed from d to m causes the $d^{2^{2n+O(1)}}$.

McCallum's Notational Idea [McC84]

Problem (Square-free Decomposition)

Generally a good idea, and often necessary. But one polynomial of degree d might become $O(\sqrt{d})$ polynomials, but on the other hand the degree might not reduce. Hence (m, d) gets worse.

Say that a set of polynomials is (M, D) if it can be partitioned into $\le M$ sets, with the sum of the degrees in each set $\le D$. This *is* preserved under square-free, relatively prime, and even complete factorisation, and behaves well w.r.t. resultants etc.

Why the subresultants? McCallum's solution [McC84]

Essentially because the vanishing of $\mathrm{res}(f, g)$ at $(\alpha_1, \ldots, \alpha_n)$ means that f and g cross above there, but the multiplicity of the crossing is determined by the vanishing of subresultants.
Hence we may need the subresultants to determine the finer points of the geometry if the resultant vanishes on a set of positive dimension.
Given (M, D) polynomials in x_n, we consider P_M:

1. (MD, D) coefficients (equally, (M, D^2))
2. $(M, 2D^2)$ discriminants
3. $(O(M^2), 2D^2)$ resultants

$(O(M^2), 2D^2)$ in all

Ths works for *order-invariance*, rather than sign-invariance, as long as no polynomial, original or computed, is identically zero on a set of positive dimension ("well-oriented").
Note the curiosity that a stronger result has a faster algorithm.

Lower bounds

Suppose $\Phi_0(x,y)$ defines $y = f_0(x)$. Let $\Phi_i(x_i,y_i) :=$

$$\exists z_i \forall x_{i-1}, y_{i-1} \left[\begin{array}{c} (y_{i-1} = y_i \wedge x_{i-1} = z_i) \\ \vee \\ (y_{i-1} = z_i \wedge x_{i-1} = x_i) \end{array} \right] \Rightarrow \Phi_{i-1}(x_{i-1} y_{i-1}). \tag{2}$$

Then $\Phi_i(x,y)$ defines $y = f_i(x) = f_{i-1}(f_{i-1}(x))$.
Using this "trick", we build large formulae quickly:

[DH88] $d^{2^{n/5+O(1)}}$: (split) complexes, $f_0 := (y_{\Re} + i y_{\Im}) = (x_{\Re} + i x_{\Im})^4 - 1$

[BD07] $m^{2^{n/3+O(1)}}$: reals, $f_0 := y = \begin{cases} 2x & (x < \frac{1}{2}) \\ 2 - 2x & (x \geq \frac{1}{2}) \end{cases}$

[BD07] Hence doubly exponential even for factored sparse polynomials.

Note that we have $O(n)$ alternations of quantifiers: this is necessary [Bas99, for example]

But straight SMT is purely $\exists$

Hence these bounds don't apply.

However, as long as we are using repeated resultants, the degree will grow doubly exponentially. There are alternatives to cylindrical algebraic decomposition.

- Virtual Term Substitution [Wei88, Wei97, Koš16]: eliminates $\exists y \Phi(x_1, \ldots, x_n, y)$ to $\Psi(x_1, \ldots, x_n)$ *provided* y occurs at most linear/quadratic/cubic in Φ.

  The degrees in Ψ may be the square of the degrees in Φ, so it's not as applicable as it looks.

  * Implemented as a pre-processor to Lazard-CAD in Maple [Ton21].

- QE by Comprehensive Gröbner Systems (CGS) [Wei98] (with a recent exploration in [FIS15]). Implemented in SYNRAC (only?), but fast: [Ton21].

  Almost nothing is known about the complexity of CGS.

  $\surd$ Operates block-at-a-time, and is fast in practice [Ton21].

But isn't Bézout's degree bound singly exponential in n?

Indeed so, but it applies to $\exists x_2 \ldots \exists x_n f_1 = 0 \wedge \cdots f_n = 0$.
[McC99] showed that Quantifier Elimination on

$$Q_{j+1}x_{j+1} \cdots Q_n x_n \left(f = 0 \wedge \Phi(g_i)\right) \qquad Q_i \in \{\forall, \exists\} \tag{3}$$

allowed reducing the double exponent of m by 1.
Extended by [BDE+16] to cases where $f = 0$ only governed parts of the formula
Also [McC01] extended to

$$Q_{j+1}x_{j+1} \cdots Q_n x_n f_1 = 0 \wedge \cdots \wedge f_r = 0 \wedge \Phi(g_i) \tag{4}$$

and, under assumptions of primitivity, [EBD15] used this to reduce the double exponent of m by r.
But the double exponent of d is still there, and this conflicts with Bézout.

Consider $\operatorname{res}_y(\operatorname{res}_x(f_1, f_2), \operatorname{res}_x(f_1, f_3))$. This has degree $O(d^4)$, again apparently contradicting Bézout. Consider the roots

$O(d^3)$ z: $\exists y, x : f_1(x, y, z) = f_2(x, y, z) = f_3(x, y, z)$

$O(d^4)$ z: $\exists y, x_1, x_2 : \begin{array}{c} f_1(x_1, y, z) = f_2(x_1, y, z) \\ \wedge f_1(x_2, y, z) = f_3(x_2, y, z) \end{array}$

These last are (generally) not roots of $\operatorname{res}_y(\operatorname{res}_x(f_1, f_2), \operatorname{res}_x(f_2, f_3))$

Hence a potentially complicated scheme of gcds of resultants

BB Instead, compute a Gröbner base of the f_i

But Aren't Gröbner bases doubly exponential?

Yes but only in the codimension [MR13], so we require that the f_i really reduce the dimension (and we can't extend this to the partial equation constraint setting of [BDE+16])

And we require that all the polynomials thus appearing are primitive.

Complexity of Gröbner Bases etc.

But — Aren't Gröbner bases doubly exponential?

Yes — but only in the codimension [MR13]

Resultant — If k polynomials determine a variety of co-dimension k, then the multiresultant has singly exponential degree.

Issue — The problem seems to be embedded components, as in [MM82, MR13], so maybe we should rule these out.

Weak asymptotic complexity? — As in [AL17], but I have no proof.

However — [Chi09] claims "Double-exponential lower bound for the degree of any system of generators of a polynomial prime ideal". I have found nobody who understands this paper.

Indeed, it's certainly a tedious constraint.
The key construct from lower bounds in (2) was

$$L_i := (y_{i-1} = y_i \wedge x_{i-1} = z_i) \vee (y_{i-1} = z_i \wedge x_{i-1} = x_i) \qquad (5)$$

This can be rewritten as $L'_i :=$

$$\left[\begin{array}{c} (y_{i-1} - y_i)(y_{i-1} - z_i) = 0 \wedge \underbrace{(y_{i-1} - y_i)(x_{i-1} - x_i)}_{\text{imprimitive}} = 0 \\ \wedge (x_{i-1} - z_i)(y_{i-1} - z_i) = 0 \wedge (x_{i-1} - z_i)(x_{i-1} - x_i) = 0 \end{array}\right] \qquad (6)$$

Let $Q_i := \exists z_i \forall x_{i-1}, y_{i-1}$ and consider $Q_i L_i \Rightarrow (Q_{i-1} L_{i-1} \Rightarrow \Phi_{i-2})$.
We can rewrite this as

$$Q_i Q_{i-1} \neg L'_i \vee \neg L'_{-1} \vee \Phi_{i-2}, \qquad (7)$$

and its negation is

$$\neg \Phi_i := \overline{Q}_i \overline{Q}_{i-1} L'_i \wedge L'_{-1} \wedge \neg \Phi_{i-2}, \qquad (8)$$

so the [DH88, BD07] examples are purely conjunctions of imprimitive equational constraints [DE16].

The Lazard projection [Laz94, MPP19]

P_L is very similar to P_M (only needs leading and trailing coefficients).

What is guaranteed is Lazard-invariance, not order-invariance.

Like order-invariance, Lazard-invariance is stronger than sign-invariance.

The lifting process is different: if a polynomial is nullified, we divide *its evaluation on the nullifying variety* through by the nullifying multiple (and therefore locally lift w.r.t. a different polynomial).

Does any of this equational constraint work generalise to the Lazard projection? Apparently so [Nai21].

There's a further improvement to the Lazard projection in [BM20], which if anything makes the equational constraint work more efficient [DNSU23].

For purely existential problems $\exists x_k, \ldots, x_n \Phi$.
$\sigma_{i,j} \in \{=, <, \leq, >, \geq\}$, but for exposition, assume all $\sigma_{i,j} \in \{<, >\}$.

1. $\Phi = (p_{1,1}\sigma_{1,1}0 \wedge \cdots) \vee (p_{2,1}\sigma_{2,1}0 \wedge \cdots) \vee \cdots$
2. Commute $\exists$ and $\vee$ and treat each disjunct Φ_i separately

So we don't care where $p_{1,1}$ and $p_{2,1}$ meet. Doesn't change asymptotics, but may well be useful in practice.

3. Choose a sample point $(s_1, \ldots, s_n^{(1)})$.
4. If this satisfies Φ_i return SAT (and witness)
5. Otherwise $\exists j : p_{i,j}(s_1, \ldots, s_n^{(1)}) \not\sigma_{i,j}0$. Remember j with $(s_1, \ldots, s_n^{(1)})$.
6. Compute largest interval $I_{n,1} = (l, u)$ such that $\forall x_n \in (l, u) p_{i,j}(s_1, \ldots, x_n) \not\sigma_{i,j}0$.
7. If $I_{n,1} \neq \mathbf{R}$ choose $s_n^{(2)} \notin I_{n,1}$. If $(s_1, \ldots, s_n^{(2)})$ satisfies Φ_i return SAT (and witness).
8. Repeat steps 5–7 until $(s_1, \ldots, s_{n-1}, \mathbf{R})$ is covered.
9. Some intervals might be redundant, so prune

Cylindrical Algebraic Coverings II [ADEK21]

10. Each of $I_{n,i}$ defines an oval in $(s_1, \ldots, s_{n-2}, x, y)$ space which cover $(s_1, \ldots, s_{n-1}, \mathbf{R})$.
11. Compute largest interval $I_{n-1,1} = (l, u)$ such that $\forall x_{n-1} \in (l, u)$ the $I_{n,i}$ cover $(s_1, \ldots, s_{n-2}, x_{n-1}, \mathbf{R})$.
12. If $I_{n-1,1} \neq \mathbf{R}$ choose a different value of s_{n-1}, $\notin I_{n-1,1}$ and repeat steps 5–9 for this value of s_{n-1}.
13. Repeat steps 4–12 until $(s_1, \ldots, s_{n-2}, \mathbf{R})$ is covered.
14. Repeat, decreasing the dimension, until we're covered the whole of the x_1-axis (or we get SAT).

Termination isn't entirely obvious, but each cell we compute contains at least one cell (the cell its sample point is in) from a CAD for the same polynomials, and the CAD itself is finite.
But the intervals $I_{k,i}$ have endpoints which are roots of iterated resultants, so degree dependence is still doubly exponential.
Open: can we improve with multi-resultant theory?

A note on division [Kov23]

Obvious answer to division: If A contains a denominator D, replace by $D \neq 0 \wedge A'$, where A' is denominator-cleared version of A.
Let $F := \forall x (0 \leq 1/x^2)$.
This converts to $F_1 := \forall x (x^2 \neq 0 \wedge 0 \leq 1)$, which is **false**.
But $\neg F = \exists x (0 > 1/x^2)$
This converts to $\exists x (x^2 \neq 0 \wedge 0 > 1)$, which is also **false**.
In JHD's view, we cannot simply "guard away" the problem that $1/x^2$ is genuinely undefined at $x = 0$, and the guarding process is inserting $x^2 \neq 0 \wedge \cdots$ in both F and $\neg F$.
In this case we should probably have

$$F := \forall x \left(0 \leq \begin{cases} 1/x^2 & x \neq 0 \\ \infty & x = 0 \end{cases} \right)$$

An alternative is to say that we didn't mean to consider the exceptional case at all, hence replacing A by $D = 0 \vee A'$ **under** $\forall$.
This "solution" doesn't scale well to mixed quantifiers, though.

Conclusions

1. The true complexity of quantifier elimination largely comes from the logical structure, especially alternation of quantifiers.
2. Imprimitive polynomials implicitly encode an $\vee$, hence logical structure.
3. The definition of cylindricity means that the results must be applicable to all quantifier structures (with the variables in the same order).
4. However, while the worst case is very bad, there is a lot that can be done.
5. Standard "Satisfiability Modulo Theories" will always produce conjunctions of elementary formulae, so this special case is worth optimising. Should be particularly suited to QE by CGS [Wei98] or CAC [ADEK21].

E. Ábrahám, J.H. Davenport, M. England, and G. Kremer.
Deciding the Consistency of Non-Linear Real Arithmetic Constraints with a Conflict Driven Search Using Cylindrical Algebraic Coverings.
Journal of Logical and Algebraic Methods in Programming, 119, 2021.
doi:10.1016/j.jlamp.2020.100633.

D. Amelunxen and M. Lotz.
Average-case complexity without the black swans.
J. Complexity, 41:82–101, 2017.

S. Basu.
New results on quantifier elimination over real closed fields and applications to constraint databases.
J. ACM, 46:537–555, 1999.

C.W. Brown and J.H. Davenport.
The Complexity of Quantifier Elimination and Cylindrical Algebraic Decomposition.
In C.W. Brown, editor, *Proceedings ISSAC 2007*, pages 54–60, 2007.
doi:10.1145/1277548.1277557.

R.J. Bradford, J.H. Davenport, M. England, S. McCallum, and D.J. Wilson.
Truth table invariant cylindrical algebraic decomposition.
J. Symbolic Comp., 76:1–35, 2016.

L. Busé and B. Mourrain.
Explicit Factors of some Iterated Resultants and Discriminants.
Math. Comp., 78:345–386, 2009.
doi:10.1090/S0025-5718-08-02111-X.

Bibliography III

C.W. Brown and S. McCallum.
Enhancements to Lazard's Method for Cylindrical Algebraic Decomposition.
In F. Boulier, M. England, T.M. Sadykov, and E.V. Vorozhtsov, editors, *Computer Algebra in Scientific Computing CASC 2020*, volume 12291 of *Springer Lecture Notes in Computer Science*, pages 129–149, 2020.
doi:https://doi.org/10.1007/978-3-030-60026-6_8.

A.L. Chistov.
Double-exponential lower bound for the degree of any system of generators of a polynomial prime ideal.
St. Petersburg Math. J., 20:983–1001, 2009.

Bibliography IV

G.E. Collins.
Quantifier Elimination for Real Closed Fields by Cylindrical Algebraic Decomposition.
In H. Brakhage, editor, *Proceedings 2nd. GI Conference Automata Theory & Formal Languages*, volume 33 of *Springer Lecture Notes in Computer Science*, pages 134–183, 1975.
doi:10.1007/3-540-07407-4_17.

J.H. Davenport and M. England.
Need Polynomial Systems be Doubly-exponential?
In G-M. Greuel, T. Koch, P. Paule, and A. Sommese, editors, *International Congress on Mathematical Software ICMS 2016*, volume 9725 of *Springer Lecture Notes in Computer Science*, pages 157–164, 2016.
doi:10.1007/978-3-319-42432-3_20.

Bibliography V

J.H. Davenport and J. Heintz.
Real Quantifier Elimination is Doubly Exponential.
J. Symbolic Comp., 5:29–35, 1988.

J.H. Davenport, A.S. Nair, G.K. Sankaran, and A.K. Uncu.
Lazard-style CAD and Equational Constraints.
In G. Jeronimo, editor, *Proceedings ISSAC 2023*, pages 218–226, 2023.

M. England, R. Bradford, and J.H. Davenport.
Improving the Use of Equational Constraints in Cylindrical Algebraic Decomposition.
In D. Robertz, editor, *Proceedings ISSAC 2015*, pages 165–172, 2015.
doi:10.1145/2755996.2756678.

M. England and J.H. Davenport.
The Complexity of Cylindrical Algebraic Decomposition with Respect to Polynomial Degree.
In V.P. Gerdt, W. Koepf, W.M. Seiler, and E.V. Vorozhtsov, editors, *Proceedings CASC 2016*, volume 9890 of *Springer Lecture Notes in Computer Science*, pages 172–192. Springer, 2016.
doi:10.1007/978-3-319-45641-6_12.

R. Fukasaku, H. Iwane, and Y. Sato.
Real Quantifier Elimination by Computation of Comprehensive Gröbner Systems.
In D. Robertz, editor, *Proceedings ISSAC 2015*, pages 173–180, 2015.

M. Košta.
New concepts for real quantifier elimination by virtual substitution.
PhD thesis, Universität des Saarlandes, 2016.

Z. Kovacs.
Why adding non-vanishing conditions on all denominators is problematic.
Commnication to Chris Brown, 2023.

D. Lazard.
An Improved Projection Operator for Cylindrical Algebraic Decomposition.
In C.L. Bajaj, editor, *Proceedings Algebraic Geometry and its Applications: Collections of Papers from Shreeram S. Abhyankar's 60th Birthday Conference*, pages 467–476, 1994.

S. McCallum.
An Improved Projection Operation for Cylindrical Algebraic Decomposition.
PhD thesis, University of Wisconsin-Madison Computer Science, 1984.

S. McCallum.
On Projection in CAD-Based Quantifier Elimination with Equational Constraints.
In S. Dooley, editor, *Proceedings ISSAC '99*, pages 145–149, 1999.

S. McCallum.
On Propagation of Equational Constraints in CAD-Based Quantifier Elimination.
In B. Mourrain, editor, *Proceedings ISSAC 2001*, pages 223–230, 2001.
doi:10.1145/384101.384132.

E. Mayr and A. Meyer.
The Complexity of the Word Problem for Commutative Semi-groups and Polynomial Ideals.
Adv. in Math., 46:305–329, 1982.

S. McCallum, A. Parusiński, and L. Paunescu.
Validity proof of Lazard's method for CAD construction.
J. Symbolic Comp., 92:52–69, 2019.

Bibliography X

E.W. Mayr and S. Ritscher.
Dimension-dependent bounds for Gröbner bases of polynomial ideals.
J. Symbolic Comp., 49:78–94, 2013.

A.S. Nair.
Curtains in Cylindrical Algebraic Decomposition.
PhD thesis, University of Bath, 2021.
URL: https://researchportal.bath.ac.uk/en/studentTheses/curtains-in-cylindrical-algebraic-decomposition.

Bibliography XI

Z. Tonks.
Poly-algorithmic Techniques in Real Quantifier Elimination.
PhD thesis, University of Bath, 2021.
URL: `https://researchportal.bath.ac.uk/en/studentTheses/poly-algorithmic-techniques-in-real-quantifier-eliminat`

V. Weispfenning.
The Complexity of Linear Problems in Fields.
J. Symbolic Comp., 5:3–27, 1988.

V. Weispfenning.
Quantifier elimination for real algebra — the quadratic case and beyond.
AAECC, 8:85–101, 1997.

V. Weispfenning.
A New Approach to Quantifier Elimination for Real Algebra.
In B.F. Caviness and J.R. Johnson, editors, *Quantifier Elimination and Cylindrical Algebraic Decomposition*, pages 376–392. Springer-Verlag, 1998.