Varieties of Doubly-Exponential behaviour in Quantifier Elimination and Cylindrical Algebraic Decomposition

James Davenport[1]
Job at https://www.bath.ac.uk/jobs/Vacancy.aspx?ref=CC9078

University of Bath

February 2022

[1] Partially Supported by EPSRC Grant EP/T015713/1

Plan of Talk

1. Introduction
2. Doubly Exponential?
3. Resultant-based projection
4. Chordality
5. Equational Constraints
6. Virtual Term Substitution
7. Comprehensive Gröbner Bases
8. Regular Chains
9. Summary

Notation

- d The maximum degree (in each variable separately) of the input polynomials.
- l The maximum bit-length of the integer coefficients
- m The number of (distinct) polynomials.
- n The number of variables.
- a The number of alternations of quantifiers. $a \leq n - 1$.
- q The number of equational constraints.
- This is the standard theory setting. Real problems tend to involve rational functions, and rational, or even algebraic, numbers.
- (M, D) At most M sets, each of combined degree $\leq D$ [McC84].

Doubly Exponential?

The complexity of QE (and hence CAD) is doubly exponential in n, more precisely $d^{2^{e_d}} m^{2^{e_m}}$ where e_d and e_m depend non-trivially on n (or on a).

[Col75] $e_m \le n + O(1)$; $e_d \le \log_2 3n + O(1)$.

[McC84] Both bounded by $n + O(1)$, conditional on no (awkward) nullification.

[Laz94] (justified by [MPP19]) $n + O(1)$ unconditionally.

[DH88] both e_d and e_m were at least $n/5 + O(1)$, with a being $\Theta(n)$ (in fact $2n/5 + O(1)$)

[BD07] (again with a being $\Theta(n)$, this time $2n/3 + O(1)$) that e_m was at least $n/3 + O(1)$, even if $d = 1$.

[BD07] e_m was at least $n/3 + O(1)$, even if $d = 1$ (again with a being $\Theta(n)$, this time $2n/3 + O(1)$).

Numerous heuristics [HEW$^+$19, e.g.], generally based on degrees of polynomials, for choosing order of elimination etc..

Graph Theory to the rescue?

Instead of considering degrees of the polynomials in F, consider the graph $\mathcal{G}(F)$ on $\{x_1, \ldots, x_n\}$ with an edge betwen (x_i, x_j) iff there is a polynomial in F containing both x_i and x_j.
Connectedness?

Gröbner If $\mathcal{G}(F)$ is not connected, the problems are independent, and [Buc79, Criterion 1] will treat them as such.

CAD Essentially independent, but this is hard to describe: we have "the outer product" of the two (or more) CADs. We definitely need to project one component at a time.

Graph Theory to the rescue continued

A graph $\mathcal{G}$ is *chordal* if every every > 3-cycle has a chord. Equivalently, every induced cycle has length 3. Every graph $\mathcal{G}$ has a chordal completion $\overline{\mathcal{G}}$.

Minimum chordal completion is NP-complete [Yan81], but that doesn't really worry me.

If this is the complete graph, then graph theory doesn't seem to help us: the exciting case is when $\overline{\mathcal{G}}$ is smaller.

An ordering $\succ$ on the vertices $x_1, \ldots, x_n$ is a *perfect elimination ordering* if $\forall i$ x_i and its neighbours $x_j : x_j \prec x_i$ form a clique. This, and chordality, can be found efficiently [RTL76].

Graph Theory to the rescue continued

Non-trivial chordality has been exploited.

Regular Chains [Che20] shows how it can be exploited efficiently.

Gröbner Bases [CP16] consider "chordal elimination". The challenge here is that an S-polynomial can introduce new edges in $\mathcal{G}$.

CAD [LXZZ21] consider chordality here, ordering x_i in a perfect elimination ordering.

Here e_d (and I think e_m) becomes the "tree depth" $\leq n$, assuming that these paths are compatible with any quantifier structure.

What we currently lack is any view of how common in practice these non-trivial chordal structures are.

Equational Constraints and e_m

Any CAD algorithm based on iterated resultants is bound to have $e_m = n + O(1)$ in the worst case, because this is how the number of polynomials grows as we take iterated resultants and discriminants: from m to $\frac{m(m+1)}{2}$ as we eliminate one variable.

Starting with [McC99], we explore how, in constructing a CAD to do QE for $f(\mathbf{x}) = 0 \wedge \Phi(g_i(\mathbf{x}))$, i.e. a CAD of $f(\mathbf{x}) = 0$ rather than the whole of $\mathbb{R}^n$, it may not be necessary to consider $\mathrm{res}_x(g_i, g_j)$, but merely the $\mathrm{res}_x(f, g_i)$. Geometrically, we do not care how g_i and g_j interact off the variety, and algebraically we have rules for commuting resultants/discriminants.

If applicable (these ideas were developed for the McCallum projection, i.e. no nullification, and adapting to Lazard is challenging [Nai21]), these reduce e_m from $n + O(1)$ to $n - q + O(1)$.

There's a snag if $\mathrm{res}(f_i, f_j)$ (the derived equational constraint in fewer variables) has non-trivial content, which corresponds to $\vee$ — back to QE?

Iterated Resultants and e_d

Any CAD algorithm based on iterated resultants is bound to have $e_d = n + O(1)$ in the worst case, because this is how resultant degrees grow.

If f, g, h have degree d, then $\mathrm{res}_x(f, g)$ has degree $2d^2$ and $P_z := \mathrm{res}_y(\mathrm{res}_x(f, g), \mathrm{res}_x(f, h))$ has degree $8d^4$. This despite the fact that Bézout says there are $O(d^3)$ common zeroes.

$P(z)$ has as roots, not just the z-coordinates of common zeros $\{z : \exists y \exists x f(x, y, z) = g(x, y, z) = h(x, y, z)\}$, but also $\{z : \exists y\, (\exists x_1 f(x_1, y, z) = g(x_1, y, z) \wedge \exists x_2 f(x_2, y, z) = h(x_2, y, z))\}$ — spurious zeroes.

[BM09] show that there is a suitable multivariate resultant which has the "right" degree.

Equational Constraints

We have seen that equational constraints can reduce e_m. But they can also reduce e_d as well.
[ED16, DE16, EBD20] consider the use of either multivariate resultants [BM09] or Gröbner bases, and show that, under generic assumptions, this will also reduce e_d to $n - q + o(1)$.
We need the "generic assumptions", as there are issues when the resultants (or Gröbner basis elements) are not primitive [DE16].
Nevertheless, all these techniques bring substantial improvements in practice.

VTS=Virtual Term Substitution, [Wei88] for linear

Here $\cdots Q_n y_n \Phi(y_1,\ldots,y_n)$ in which y_n occurs linearly can be replaced by $\cdots\hat{\Phi}(y_1,\ldots,y_{n-1})$. This was extended in [Wei94, Wei97] to the quadratic case and beyond, with details of the cubic case being in [Koš16]. An extension to unbounded degree is given in [LPJ14].

A crude description would be "substituting in the critical values and their neighbours", but the details are more subtle, hence Weispfenning's concept of *virtual* term substitution.

In particular, if y_n occurs quadratically, with corresponding critical values $y_n=\frac{1}{2a}\left(-b\pm\sqrt{b^2-4ac}\right)$, there might be 0, 1 or 2 critical values, and we also need to worry about the case $a=0$: hence VTS has substitutions with guards, and the result of eliminating an $\exists$ quantifier, and hence a block of $\exists$, is a disjunction, often large.

However, VTS treats $\forall$ as $\neg_1\exists\neg_2$, so $\neg_2$ turns the disjunction into a conjunction, processing the $\exists$ builds a further disjunction on top of this, which $\neg_1$ turns back into a conjunction.

Each could have exponential blowup, so a $\geq 2^{2^a}$ behaviour for e_m.

The key idea is this. We consider an "innermost block" in this form:

$$\exists\overline{x}\left(\begin{array}{c} f_1(\overline{y},\overline{x})=0\wedge\cdots f_r(\overline{y},\overline{x})=0\wedge \\ p_1(\overline{y},\overline{x})>0\wedge\cdots p_s(\overline{y},\overline{x})>0\wedge \\ q_1(\overline{y},\overline{x})\neq 0\wedge\cdots q_t(\overline{y},\overline{x})\neq 0 \end{array}\right)$$

where $\overline{y}$ represents the remaining variables, and $f_i, p_j, q_k \in \mathbb{Q}[\overline{y},\overline{x}]\setminus\mathbb{Q}[\overline{y}]$. We introduce new variables $\overline{z}$ and $\overline{w}$, with $\overline{z},\overline{w}\succ\overline{x}$, and consider the polynomials

$$\{f_1,\ldots,f_r,\underbrace{z_1^2p_1-1,\ldots,z_s^2p_s-1}_{\text{forcing positive}},\underbrace{w_1q_1-1,\ldots,w_tq_t-1}_{\text{forcing nonzero}}\}.$$

Let $\mathcal{G}=(S_i,G_i)$ be a Comprehensive Gröbner System (with parameters $\overline{y}$) for this so that $\overline{y}$ space is partitioned by the S_i. We claim each G_i will be
$\{f_1',\ldots,f_{r'}',u_1z_1^2-p_1',\ldots,u_sz_s^2-p_s',v_1w_1-q_1',\ldots,v_tw_t-q_t'\}$.
Our answer will be $\bigvee_i \Psi_i(S_i,G_i)$: next two slides explain Ψ_i.

G_i zero-dimensional ($\overline{z}, \overline{w}$ irrelevant for dimension)

If $G_i = (1)$ then we return `false`. Otherwise recall
$G_i = \{f'_1, \ldots, f'_{r'}, u_1 z_1^2 - p'_1, \ldots, u_s z_s^2 - p'_s, v_1 w_1 - q'_1, \ldots, v_t w_t - q'_t\}$.
Let $I = \langle f'_1, \ldots, f'_{r'} \rangle$,

$$\chi(x) = \prod_{(e_1,\ldots,e_s)\in\{0,1\}^s} \chi^I_{(p'_1/u_1)^{e_1},\cdots,(p'_s/u_s)^{e_s}}(x) = x^{2^s d} + \sum_0^{2^s d-1} a_i x^i.$$

The answer is $\Psi_i := \mathcal{F}(S_i) \wedge I_{2^s d}(a_i)$.
JHD: at least that's my reconstruction. I can't see where the w_i (the $\neq 0$) terms come in. Also, the subscript of $\chi^I_{\ldots}$, the characteristic polynomial of $M^I_{\ldots}$, is not a polynomial.

$\exists\phi$: $G_i > 0$-dimensional ($\overline{z}, \overline{w}$ irrelevant for dimension)

$\overline{u} :=$ maximal independent variables $(\overline{x}, G_i, \succ)$. (B)
If $\overline{u} = \overline{x}$ return SYNRAC($\mathcal{F}(S) \wedge \exists\overline{x}\phi$) [Wei98]
$\overline{x}' := \overline{x} \setminus \overline{u}$; $\phi_1 := \texttt{Free}(\phi, \overline{x}')$; $\phi_2 := \texttt{NonFree}(\phi, \overline{x}')$;
$\varphi := \phi_1 \wedge$Recurse(S_i,$\exists\overline{x}'\phi_2$) (1)(A)
JHD: I think this means φ now only contains $\overline{u}$-variables
Let $\varphi_1 \vee \cdots \vee \varphi_l$ be a disjunctive normal form of φ. (C)
for $1 \le j \le l$ **do**
  $\varphi_j^{(1)} := \texttt{Free}(\varphi, \overline{u})$; $\varphi_j^{(2)} := \texttt{NonFree}(\varphi_j, \overline{u})$;
  $\psi_j := \varphi_j^{(1)} \wedge$Recurse($S_i$,$\exists\overline{u}\phi_j^{(2)}$) (2)(E)
Return $\Psi := \mathcal{F}(S_i) \wedge (\psi_1 \vee \cdots \vee \psi_l)$
JHD: "Recurse" goes right back to the MainQE, note that call (1) has pushed the $\overline{u}$-variables into being parameters (I think) (D). But somehow S_i gets lost in these recursions: I hope I've added it in the right place. Their Theorem 16 states that this does terminate — far from obvious (F).

- (A) Recursing with S is, I think, my interpolation to make sense of the recursions we'll see later. S initially is $\mathbb{R}^{\#\overline{y}}$.
- (B) There's a lot of freedom here: ML?
- (C) Note that our main recursion is on ϕ in conjunctive normmal form (CNF), whereas here we convert to disjunctive normal form (DNF) and implicitly back at the end of the block. Since CNF$\leftrightarrow$DNF naïvely is exponential, this would provide an exponential blowup at each $\exists/\forall$ boundary, similar to [DH88].
- (D) Therefore this recursion is on strictly fewer variables, since dim > 0.
- (E) Therefore this recursion is on strictly fewer variables, since $\overline{u} \neq \overline{x}$. $\varphi_j^{(1)}$ is free of $\overline{u}$ by construction, and free of $\overline{x}'$ since it comes from ϕ_1, so actually belongs in an outer block. We might ask why such things exist, but they could be generated by the recursion.
- (F) But the two previous notes are probably key.

Complexity of CGB

I know no results on the complexity of Comprehensive Gröbner Bases.
Since we are doing Gröbner Bases, we might *hope for* singly exponential behaviour at each block, and hence $e_d = O(a)$ rather than $O(n)$, but worst-case Gröbner bases can be doubly exponential [MR13]. *If* we get $O(a)$ behaviour, though, this does not depend on having a lot of equational constraints.
We are doing CNF/DNF conversions at each quantifier alternation, as with VTS, so this could be expected to give us $e_m = O(a)$ rather than $O(n)$.

Regular Chains [CM14b, CM14a]

Regular Chains/Triangular Decompositions are an alternative to Gröbner bases, and write the solution as a union of triangular sets.

Very little is known about the complexity of Triangular Decompositions. I *believe* that the upper bounds for Gröbner bases [Dub90, etc.] still apply, but I haven't seen a formal proof.

In the presence of equational constraints [BCD$^+$14], we should get the same improvement as Gröbner bases deliver.

There is probably a relationship between the different triangular sets in a Triangular Decomposition and the sets S_i in a Comprehensive Gröbner Basis, but again I don't know what this is.

~~Black swans~~ [AL17]

"Average-case complexity without the black swans": i.e. without an exponentially-rare family that is worse than exponentially bad.

Definition

For $k \in \mathbf{N}$ let (M_k, μ_k) be a probability space and let $T_k : M_k \to \mathbb{R}$ be a μ_k-measurable function. We say that the family $\{T_k\}$ has a weak expectation of $O(f(k))$ if there exists a family of sets of exceptional inputs,$E_k \subseteq M_k$, such that $\mu_k(E_k) = e^{-\Omega(k)}$ and the conditional expectation $E[T_k(x)|x \notin E_k]$ is bounded by $O(f(k))$.

- Condition numbers inversely proportional to a distance to a homogeneous algebraic set of ill-posed inputs;
- Renegar's condition number for conic optimization;
- The running time of power iteration for computing a leading eigenvector of a Hermitian matrix.

? Any such result in our area?

Summary

Idea	e_m	e_d
Collins	$n + O(1)$	$(\log_2 3)n + O(1)$
McCallum (but nullification)	$n + O(1)$	$n + O(1)$
Lazard [MPP19]	$n + O(1)$	$n + O(1)$
Equational Constraints	(?) $n - q + O(1)$	(?) $n - q + O(1)$
Virtual Term Substitution	(?) $O(a)$	challenges
Comprehensive Gröbner Bases	(??) $O(a)$	(??) $O(a)$–$n + O(1)$
Regular Chains	(??) $n - q + O(1)$	(??) $n - q + O(1)$

But Virtual Term Substitution (where applicable), Comprehensive Gröbner Bases and Regular Chains all seem to be very fast in practice.

Conclusions/Questions

1. We need to understand the complexity of Virtual Term Substitution.
2. What about unbounded degree: [LPJ14]? It is restricted to univariates — is this inherent?
3. We need to understand the complexity of Comprehensive Gröbner Bases.
4. We need to understand the complexity of Regular Chains.
5. We need to understand the inter-relationships between these methods.
6. Are there any "weak average case complexity" results? The examples of [BD07, DH88] seem very special.

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

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*, pages 44–58, 2014.

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.

Bibliography II

L. Busé and B. Mourrain.
Explicit factors of some iterated resultants and discriminants.
Math. Comp., 78:345–386, 2009.

B. Buchberger.
A Criterion for Detecting Unnecessary Reductions in the Construction of Groebner Bases.
In *Proceedings EUROSAM 79*, pages 3–21, 1979.

Changbo Chen.
Chordality Preserving Incremental Triangular Decomposition and Its Implementation.
In A.M. Bigatti, J. Carette, J.H. Davenport, M. Joswig, and T. de Wolff, editors, *Mathematical Software — ICMS 2020*, volume 12097 of *Springer Lecture Notes in Computer Science*, pages 27–38. Springer, 2020.

URL: https://www.researchgate.net/publication/342758264_Chordality_Preserving_Incremental_Triangular_Decomposition_and_Its_Implementation, doi:10.1007/978-3-030-52200-1_3.

C. Chen and M. Moreno Maza.
Cylindrical Algebraic Decomposition in the RegularChains Library.
In *Proceedings Mathematical Software — ICMS 2014*, pages 425–433, 2014.

C. Chen and M. Moreno Maza.
Quantifier Elimination by Cylindrical Algebraic Decomposition Based on Regular Chains.
In K. Nabeshima, editor, *Proceedings ISSAC 2014*, pages 91–98, 2014.

Bibliography IV

G.E. Collins.
Quantifier Elimination for Real Closed Fields by Cylindrical Algebraic Decomposition.
In *Proceedings 2nd. GI Conference Automata Theory & Formal Languages*, pages 134–183, 1975.

D. Cifuentes and P. Parrilo.
Exploiting chordal structure in polynomial ideals: A Grobner bases approach.
SIAM Journal on Discrete Mathematics, 30:1534–1570, 2016.

J.H. Davenport and M. England.
Need Polynomial Systems be Doubly-exponential?
In *Proceedings ICMS 2016*, pages 157–164, 2016.

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

Bibliography V

T.W. Dubé.
The structure of polynomial ideals and Gröbner Bases.
SIAM J. Comp., 19:750–753, 1990.

M. England, R.J. Bradford, and J.H. Davenport.
Cylindrical Algebraic Decomposition with Equational Constraints.
In J.H. Davenport, M. England, A. Griggio, T. Sturm, and C. Tinelli, editors, *Symbolic Computation and Satisfiability Checking: special issue of Journal of Symbolic Computation*, volume 100, pages 38–71. 2020.

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*, Springer Lecture Notes in Computer Science 9890, pages 172–192. Springer, 2016.
URL: `http://arxiv.org/abs/1605.02494`, `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.

Z. Huang, M. England, D. Wilson, J.H. Davenport, and L.C. Paulson.
Using Machine Learning to Improve Cylindrical Algebraic Decomposition.
Mathematics in Computer Science, 11:461–488, 2019.

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

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.

K. Liiva, G.O. Passmore, and P.B. Jackson.
A note on real quantifier elimination by virtual term substitution of unbounded degree.
https://homepages.inf.ed.ac.uk/pbj/papers/pas14.pdf, 2014.

H. Li, B. Xia, H. Zhang, and T. Zheng.
Choosing the Variable Ordering for Cylindrical Algebraic Decomposition via Exploiting Chordal Structure.
ISSAC '21: Proceedings of the 2021 International Symposium on Symbolic and Algebraic Computation, pages 281–288, 2021.

Bibliography IX

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

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.

Donald J Rose, R Endre Tarjan, and George S Lueker.
Algorithmic aspects of vertex elimination on graphs.
SIAM Journal on computing, 5(2):266–283, 1976.

Bibliography XI

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

V. Weispfenning.
Quantifier elimination for real algebra — the cubic case.
In *Proceedings ISSAC 1994*, pages 258–263, 1994.

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.
Quantifier Elimination and Cylindrical Algebraic Decomposition, pages 376–392, 1998.

Mihalis Yannakakis.
Computing the minimum fill-in is NP-complete.
SIAM Journal on Algebraic Discrete Methods, 2(1):77–79, 1981.