NEW ZEALAND JOURNAL OF MATHEMATICS
Volume 52 (2021), 233–257
https://doi.org/10.53733/92

SMALE'S 17TH PROBLEM: ADVANCES AND OPEN DIRECTIONS

Felipe Cucker

(Received 4 March, 2021)

Dedicated to Vaughan Jones (1952-2020). In memoriam.

Abstract. We give an overview of Smale's 17th problem describing the context in which Smale proposed it, the ideas that led to its solution, and the extensions and subsequent progress after this solution.

Contents

1. Foreword

In June 1999 I received an invitation letter to give a short course on a summer school. The letter was signed by Vaughan Jones, then the chairman of the NZMRI (New Zealand Mathematical Research Institute). New Zealand being in the south, the summer school was going to be in January. I immediately accepted. I had never been in New Zealand, but I knew that it shared some similarities with my home country (Uruguay). My stay at Kaikoura was delightful. The gentle beauty

2000 *Mathematics Subject Classification* 65H20, 65Y20.
Key words and phrases: polynomial equation solving, homotopy methods, approximate zero, complexity, polynomial time.
Partially supported by GRF grant CityU 11302321.

of the town, the vicinity of Marlborough's vineyards, the whales in the neighboring waters, the kindness of the local people in whose house I stayed, ... And the school was no less fascinating. Not only because of the interest of its students and the quality of the short courses I took advantage to attend but also because of its setting: lectures and lunches were held at the local Marae. More than twenty years later I still look at it as a unique experience.

January 2000 was also the beginning of the 21st century. So, while we were enjoying the school at Kaikoura, the presses of the American Mathematical Society were at work to print a singular volume [**4**]. Commissioned by the International Mathematical Union, and to celebrate the year 2000 as the Year of Mathematics, the book acknowledges its inspiration on the list of problems proposed by Hilbert in 1900 for the mathematicians of the 20th century. Thus, along with articles describing the state of mathematics at the end of the century a few articles proposed, in the spirit of Hilbert, lists of problems for the mathematicians of the 21st century. Among the latter one was written by Vaughan [**14**]; it describes ten open problems related to his research interests. Another was writen by Steve Smale [**27**][1]. Similar in character to Vaughan's it describes eighteen problems related to Smale's research interests[2]. The nature of the 17th had a large overlap with the subject of my course at Kaikoura. Our aim in this article is to describe —with a focus on ideas rather than on technical details— what this problem is, what the work of Smale leading to it was, how it was solved a few years ago, and what are the questions left open after this solution. It is a fitting update to the course I gave at Vaughan's invitation.

2. Smale 17th Problem

2.1. The background. Work on the solution of polynomial equations may be traced back to the Babylonians, 4000 years ago, who solved some quadratic equations arising from problems relating areas and sides of rectangles [**13**]. The general solution to cubic and quartic equations is generally attributed to Ferrari and Cardano in the 16th century. Solutions could be written in terms of square and cubic roots and involved complex numbers. It took almost three centuries to prove that a similar result cannot be achieved for the general quintic equation [**1**].

Nonetheless, the solutions exist. The Fundamental Theorem of Algebra states that a polynomial equation of degree d has exactly d zeros in $\mathbb{C}$ when counted with their multiplicity. An equally fundamental result in algebraic geometry, known as the Bézout Theorem, extends the FTA to (square) systems of polynomials. It needs a caveat though. The system $\{X_1 + X_2 = 0, X_1 + X_2 = 1\}$ has, obviously, no zeros: the two lines are parallel. But if one "adds points at infinity" to the complex plane $\mathbb{C}^2$ these two lines do meet at one such point. Recall, the *complex projective space* $\mathbb{P}^n$ is the quotient $\mathbb{C}^{n+1}/\sim$ where $x \sim y$ if $x = \lambda y$ for some $\lambda \in \mathbb{C}$, $\lambda \neq 0$. Also, a polynomial $f \in \mathbb{C}[X_0, X_1, \ldots, X_n]$ is *homogeneous* of degree d when all its monomials have degree d. In this case, if $x \in \mathbb{C}^{n+1}$ is such that $f(x) = 0$ then, for all $\lambda \in \mathbb{C}$, $f(\lambda x) = \lambda^d f(x) = 0$. Hence we can talk about the zeros of f in $\mathbb{P}^n$ and, by extension, to the zeros in $\mathbb{P}^n$ of a system $F = (f_1, \ldots, f_n)$ of homogeneous polynomials (possibly of different degrees). Given a point $(x_0, \ldots, x_n) \in \mathbb{C}^{n+1} \setminus \{0\}$

[1]Smale's paper had actually been first published in the *Mathematical Intelligencer* [**29**].

[2]Of the eighteen, three (The Riemann Hypothesis, the Poincaré Conjecture and the question "Is $\mathsf{P} = \mathsf{NP}$?") were well-known at the time; the other fifteen were new.

we denote by $[x_0 : \cdots : x_n]$ its class in $\mathbb{P}^n$. In all what follows we will choose representatives $(x_0, \ldots, x_n)$ of a point $[x] \in \mathbb{P}^n$ satisfying that $\|x\| = 1$. This has no loss of generality and simplifies a number of expressions.

Now let $\mathbf{d} := (d_1, \ldots, d_n) \in \mathbb{N}^n$ and $\mathcal{H}_{\mathbf{d}}$ be the linear space of the systems $F = (f_1, \ldots, f_n)$ with $f_i \in \mathbb{C}[X_0, \ldots, X_n]$ homogeneous of degree d_i. Generically (and by this we mean outside a subset of smaller dimension in $\mathcal{H}_{\mathbf{d}}$) the system F has a finite number of zeros in $\mathbb{P}^n$. Bézout Theorem states that in this case, and when counted with multiplicity, this number is $\mathcal{D} := d_1 d_2 \cdots d_n$. The question underlying Smale's 17th problem is the following:

Can one efficiently compute (any) one zero of a given $F \in \mathcal{H}_{\mathbf{d}}$? (S17)

2.2. The Bézout series. During the early 1990s, in a series of papers known as "the Bézout series" [**22, 23, 24, 26, 25**], Mike Shub and Steve Smale laid down the foundations of an approach to answer this question. We won't attempt to summarize these papers but limit ourselves to describe some of the main ingredients in this approach.

2.2.1. *Approximate zeros.* We already mentioned that the solutions of a general quintic equation (in one variable) cannot be expressed by radicals. Obviously, neither can the zeros of a system $F \in \mathcal{H}_{\mathbf{d}}$. The most common way out of this obstacle is to compute a (sufficiently good) approximation of such a zero. Arguably the most common definition of approximation is through a bound on the distance to a true zero. If $d_{\mathbb{P}}$ denotes the Riemannian distance in $\mathbb{P}^n$ (i.e., the angle), we say that $z \in \mathbb{P}^n$ is an *ε-approximate zero* of F when there is a zero $\zeta \in \mathbb{P}^n$ of F such that $d_{\mathbb{P}}(z, \zeta) \leq \varepsilon$. A shortcoming of this notion is the dependence on ε: being an approximate zero in this sense is not an absolute notion. The definition adopted in the Bézout series, which goes back to [**28**], is different and relies on Newton's method. The latter can be extended to the projective space. Given $z \in \mathbb{P}^n$ and $F \in \mathcal{H}_{\mathbf{d}}$, the *Newton iterate* of F at z is

$$\mathrm{N}_F(z) := z - \mathrm{D}_z F|_{\mathrm{T}_z}^{-1} F(z). \tag{1}$$

Here T_z is the tangent space to $\mathbb{P}^n$ at z (i.e., the orthogonal complement $z^\perp$ to z in $\mathbb{C}^{n+1}$) and $\mathrm{D}_z F|_{\mathrm{T}_z}$ is the restriction to this space of the derivative of F at z. A remarkable feature of Newton's method is its quadratic convergence. Let ζ be a non-singular (i.e., not multiple) zero of F and z be a point sufficiently close to ζ. Then,

$$d_{\mathbb{P}}(\mathrm{N}_F^k(z), \zeta) \leq d_{\mathbb{P}}(\mathrm{N}_F(z), \zeta)\, 2^{1-2^k} \tag{2}$$

where N_F^k is the kth iterate of (1). Shub and Smale call a point z satisfying (2) an *approximate zero* of F and the point ζ its *associated zero*. Note, this is a qualitative notion; there is no dependence on a parameter measuring how well the point approximates the zero. It is also a strong property. Once a point ζ is an approximate zero of F, inequality (2) allows one to get an ε-approximation with $\log_2(1 + \log_2 \varepsilon)$ Newton steps.

Two questions are naturally posed. Firstly, what does 'sufficiently close' mean? Can one provide estimates for this notion? The answer goes back (again) to [**28**].

For a point $\zeta \in \mathbb{P}^n$ define

$$\gamma(F,\zeta) := \sup_{k\geq 2} \left\| \mathrm{D}_\zeta F|_{\mathrm{T}_\zeta}^{-1} \frac{\mathrm{D}^k_\zeta F}{k!} \right\|^{\frac{1}{k-1}} \tag{3}$$

where $\mathrm{D}^k_\zeta F$ is the kth derivative of F at ζ and the norm is the operator norm. The *γ-Theorem* of Smale states that if $F(\zeta) = 0$ and $d_{\mathbb{P}}(z,\zeta)\gamma(F,\zeta) \leq \frac{1}{6}$ then z is an approximate zero of F with associated zero ζ (see, e.g., [**17**, Thm. 12]).

Secondly, in the absence of a true zero ζ at hand, can one certify that a point z is an approximate zero of F? To answer this question let

$$\beta(F,z) := \left\| \mathrm{D}_z F|_{\mathrm{T}_z}^{-1} F(z) \right\|$$

be the length of the Newton step in (1) and

$$\alpha(F,z) := \beta(F,z)\gamma(F,z).$$

Smale's *α-Theorem* states that there exists a constant α_0 such that if $\alpha(F,z) \leq \alpha_0$ then z is an approximate zero of F. A proof with $\alpha_0 = 0.02$ is in [**10**, Thm. 19.9].

The goal in (S17) is therefore not to compute a true zero but an approximate zero as above.

2.2.2. *Linear homotopies.* The algorithmic scheme considered in the Bézout series was a linear homotopy. Its general idea (which goes back at least to Lahaye [**15**]) can be simply described. Assume you want to compute the zeros of the polynomial $F = X^3 - X^2 - 5X + 4$ and further assume that you aren't aware of the formula to do so. You may try to use the fact that you certainly know the zeros of $G = X^3 - X$. To do so, consider the segment (in the space of polynomials of degree at most 3)

$$Q_t := (1-t)G + tF, \qquad \text{with } t \in [0,1]. \tag{4}$$

We observe that its extremities are $Q_0 = G$ and $Q_1 = F$. Hence, when t varies from 0 to 1, we may expect the zeros of Q_t to vary from the zeros of G ($1, 0$ and -1) to those of F. Starting with one such zero ζ_0 of G (say $\zeta_0 = 1$) this will induce a curve $\mathscr{C} = (Q_t, \zeta_t)_{t\in[0,1]}$ with $Q_t(\zeta_t) = 0$ for all $t \in [0,1]$ and if we can "follow" this curve via a finite sequence $\{Q_{t_i}, z_i\}_{i=0,\dots,k}$ such that $0 = t_0 < t_1 < \cdots < t_k = 1$ and z_i is a good approximation of ζ_{t_i} then we will end up with z_k, a good approximation to the zero ζ_1 of $Q_1 = F$. Figure 1 attempts to convey this idea.

A few comments are in order. The first is that the same idea applies for systems in $\mathcal{H}_{\mathbf{d}}$; we now have n homogeneous polynomials in $n+1$ variables and $\mathcal{D}$ zeros in $\mathbb{P}^n$ instead of 3 in $\mathbb{C}$ but the basic idea is the same. The second is that the "lifting" of the segment $[Q_0, Q_1]$ to the curve $\mathscr{C} \subset \mathcal{H}_{\mathbf{d}} \times \mathbb{P}^n$ is well-defined provided $\mathrm{D}_{\zeta_t} Q_t|_{\mathrm{T}_{\zeta_t}}$ is invertible for all $t \in [0,1]$. Recall that the *discriminant variety* in $\mathcal{H}_{\mathbf{d}}$ is the set

$$\Sigma := \{F \in \mathcal{H}_{\mathbf{d}} \mid \exists \zeta \in \mathbb{P}^n \text{ s.t. } F(\zeta) = 0 \text{ and } \mathsf{rank}(\mathrm{D}_\zeta F|_{\mathrm{T}_\zeta}) < n\}. \tag{5}$$

Systems outside Σ are those having $\mathcal{D}$ different smooth zeros. Then, as long as the segment $[Q_0, Q_1]$ does not intersect Σ, the $\mathcal{D}$ zeros of G induce $\mathcal{D}$ curves in $\mathcal{H}_{\mathbf{d}} \times \mathbb{P}^n$ which do not intersect one another. Furthermore, it is known that Σ is a complex hypersurface in $\mathcal{H}_{\mathbf{d}}$. Hence, it has real codimension 2 in $\mathcal{H}_{\mathbf{d}}$. This implies that, generically, the segment $[Q_0, Q_1]$ does not intersect Σ.

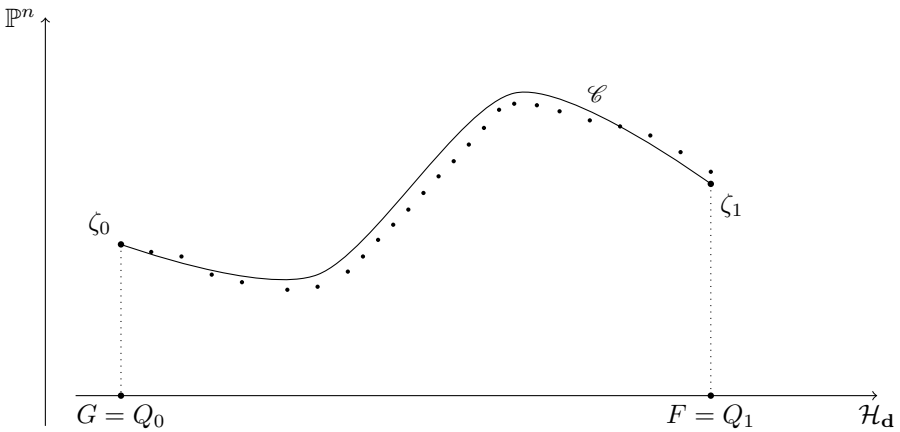

Figure 1. The curve $\mathscr{C}$, the initial pair (G, ζ_0), the target pair (F, ζ_1) and the points (Q_{t_i}, z_i) that "follow" $\mathscr{C}$.

The bottomline is, linear homotopy, in general, works. But for some unlucky choices of F and G it may go awfully wrong. Consequently, to give a formal meaning to the notion of "efficiently compute" we need a framework other than the worst-case scenario.

2.3. Probabilities and complexity. To talk about computational efficiency supposes a *cost measure* at hand. As we are considering numerical algorithms, that is, algorithms whose basic operations are arithmetic operations and comparisons between (idealized) real numbers the simplest, and most natural, cost measure for a computation with input F is the number $\mathsf{cost}(F)$ of such operations performed along the computation. The *complexity of an algorithm* is a function relating the cost of the algorithm's computations to the size of the input data.

In our case, the size N of a system $F \in \mathcal{H}_{\mathbf{d}}$ is the number of complex numbers we need to describe the system. That is, the number of coefficients in its polynomials. This gives us

$$N = \sum_{i=1}^{n} N_i \qquad \text{where} \qquad N_i = \binom{n+d_i}{n}. \tag{6}$$

We have mentioned that we need a framework for complexity other than the worst-case scenario. The one proposed by Smale, which is widely used, is that of *average complexity*. To describe it we need some structure on $\mathcal{H}_{\mathbf{d}}$.

We have mentioned that $\mathcal{H}_{\mathbf{d}}$ is a linear space. We can turn it into an inner product space by endowing it with the *Weyl inner product*. This is a dot product in a scaled monomial basis which has the property of being *unitarily invariant*. That is, for all $F, G \in \mathcal{H}_{\mathbf{d}}$ and all unitary transformation $u \in U(n+1)$, we have $\langle F \circ u, G \circ u \rangle = \langle F, G \rangle$. See [**10**, §16.1] for details. In all what follows, for a system $F \in \mathcal{H}_{\mathbf{d}}$, $\|F\|$ will denote its Weyl norm $\langle F, F \rangle^{1/2}$.

The Weyl norm induces a unit sphere in $\mathcal{H}_{\mathbf{d}}$, in what follows denoted by $\mathbb{S}(\mathcal{H}_{\mathbf{d}})$, which we can endow with the uniform distribution. The latter, in turn, allows one

to define the *average cost* of an algorithm (taking inputs in $\mathbb{S}(\mathcal{H}_{\mathbf{d}})$) to be

$$\mathop{\mathbb{E}}_{F\sim\mathbb{S}(\mathcal{H}_{\mathbf{d}})} \mathsf{cost}(F). \tag{7}$$

2.4. The problem's statement. We can now formally state Smale's 17th problem. This statement relies on the usual understanding of "efficient" as solvable with polynomially bounded average cost. The question in (S17) thus becomes

Is there an algorithm that, given an $F \in \mathbb{S}(\mathcal{H}_{\mathbf{d}})$ randomly drawn from the uniform distribution, halts with probability one returning an approximate zero of F, and whose average cost is bounded by a polynomial in N?

3. The Solution

To turn the linear homotopy scheme into a bona fide algorithm we need to specify how do we choose the initial pair (G, ζ_0) and how do we compute the sequences $t_0, t_1, \ldots, t_k$ and $z_0, \ldots, z_k$. Let us start with the latter. Before doing so, however, we observe that it makes sense to work with the angle $d_{\mathbb{S}}(Q, H)$ between systems $Q, H \in \mathcal{H}_{\mathbf{d}}$ instead of the distance $\|Q - H\|$. This is so because there are no changes in the zeros of a system when the system is multiplied by a constant. Consequently, we will consider the segment $[G, F]$ parameterized by a fraction τ of the angle $d_{\mathbb{S}}(G, F)$ instead of a fraction t of the distance $\|F - G\|$ as in (4). That is, we will consider $\{Q_\tau\}_{\tau\in[0,1]}$ where Q_τ is the only system in $[G, F]$ satisfying $d_{\mathbb{S}}(G, Q_\tau) = \tau\, d_{\mathbb{S}}(G, F)$.

3.1. Condition numbers. Assume, for the time being, that we are given F, G and ζ_0 with $G(\zeta_0) = 0$. A first remark is that we won't attempt to find a universal partition of $[0, 1]$ that we can use for all triples (F, G, ζ_0) but rather that we will proceed dynamically: at the ith step, we have computed a pair (Q_i, z_i) for which we know that z_i is an approximate zero of Q_i with associated zero ζ_i (we should write Q_{τ_i} and ζ_{τ_i} but we simplify the notation for ease of reading). Next assume that

$$d_{\mathbb{P}}(z_i, \zeta_i) \leq \frac{1}{12\gamma(Q_i, \zeta_i)}. \tag{8}$$

Note that this is stronger, by a factor of 2, than the condition required in Smale's γ-Theorem (cf. §2.2.1). Given a $\Delta\tau \leq \frac{d_{\mathbb{S}}(Q_i,F)}{d_{\mathbb{S}}(G,F)}$ we let Q_{i+1} be the only system in $[Q_i, F]$ satisfying $d_{\mathbb{S}}(Q_i, Q_{i+1}) = \Delta\tau\, d_{\mathbb{S}}(G, F)$.

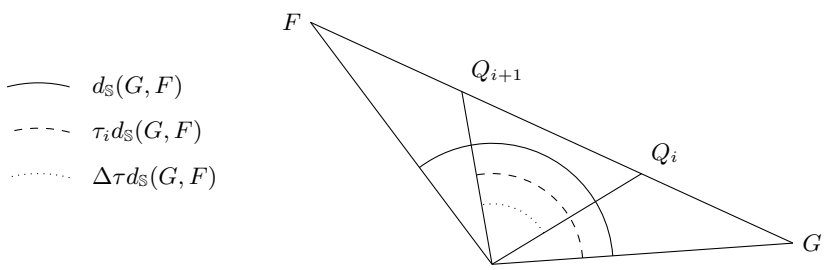

Figure 2. The system Q_{i+1} given by Q_i and the step-length $\Delta\tau$.

The goal is to compute a bound $\mathcal{B}(Q_i, z_i)$ such that

$$\Delta\tau \leq \mathcal{B}(Q_i, z_i) \implies d_{\mathbb{P}}(z_i, \zeta_{i+1}) \leq \frac{1}{6\gamma(Q_{i+1}, \zeta_{i+1})}. \tag{9}$$

If we manage to find $\mathcal{B}$ such that (9) holds, then z_i is an approximate zero of Q_{i+1} with associated zero ζ_{i+1} (by Smale's γ-Theorem) and then, taking

$$z_{i+1} := \mathrm{N}_{Q_{i+1}}(z_i)$$

we have (by (2))

$$d_{\mathbb{P}}(z_{i+1}, \zeta_{i+1}) \leq \frac{1}{2} d_{\mathbb{P}}(z_i, \zeta_{i+1}) \leq \frac{1}{12\gamma(Q_{i+1}, \zeta_{i+1})},$$

i.e., that (8) holds, and we can iterate the process.

Note that inequality (8) is trivial for $i = 0$ as $z_0 = \zeta_0$ and hence $d_{\mathbb{S}}(z_0, \zeta_0) = 0$. This allows us to begin the process. Figure 3 shows one step of the homotopy.

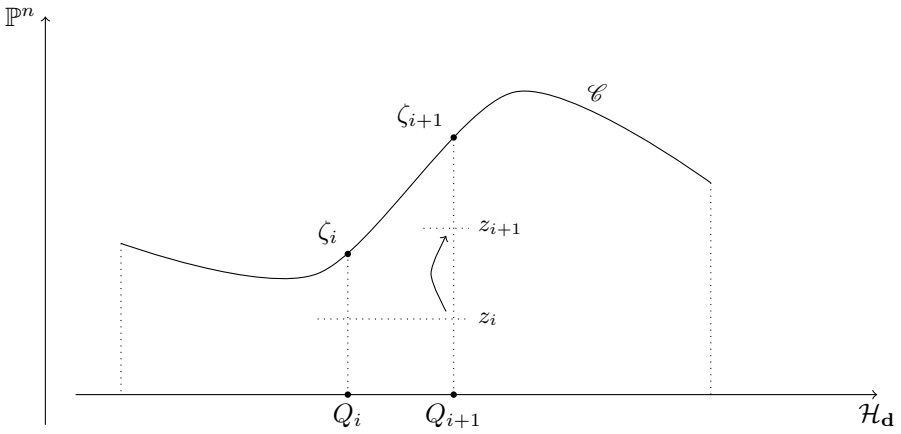

Figure 3. Computation of next pair $(Q_{t_{i+1}, z_{i+1}})$.

How do we obtain $\mathcal{B}$ satisfying (9)?

Recall, we want $d_{\mathbb{P}}(z_i, \zeta_{i+1}) \leq \frac{1}{6\gamma(Q_{i+1}, \zeta_{i+1})}$. We begin by noting that

$$d_{\mathbb{P}}(z_i, \zeta_{i+1}) \leq d_{\mathbb{P}}(z_i, \zeta_i) + d_{\mathbb{P}}(\zeta_i, \zeta_{i+1}). \tag{10}$$

Inequality (8) already bounds the first term in the right-hand side. In addition, Lipschitz bounds for γ^{-1} allow one to show, for Q_{i+1} sufficiently close to Q_i and ζ_{i+1} sufficiently close to ζ_i, that

$$d_{\mathbb{P}}(z_i, \zeta_i) \leq \frac{1}{12\gamma(Q_i, \zeta_i)} \leq \frac{1}{c_1\gamma(Q_{i+1}, \zeta_{i+1})} \qquad \text{for some } c_1 \in (6, 12). \tag{11}$$

Let us assume for a while that this holds true and focus on the second term in the right-hand side of (10). To bound this term define, for $F \in \mathcal{H}_{\mathbf{d}}$ and $z \in \mathbb{P}^n$, the *condition number* of F at z to be

$$\mu_{\mathrm{norm}}(F, z) := \|F\| \left\| \mathrm{D}_z F|_{\mathrm{T}_z}^{-1} \mathsf{diag}(\sqrt{d_i}) \right\|.$$

This is a condition number in the classical sense: if $F(\zeta) = 0$, then $\mu_{\text{norm}}(F, \zeta)$ bounds the first-order variation of ζ in terms of the variation of F[3]. Hence, we may expect that

$$d_{\mathbb{P}}(\zeta_i, \zeta_{i+1}) \leq (1 + c_2)\Delta\tau\mu_{\text{norm}}(Q_i, \zeta_i) \tag{12}$$

for some small constant c_2. For a given c_3, if we want the left-hand side above to be at most $\frac{1}{c_3\gamma(Q_{i+1},\zeta_{i+1})}$ we should impose

$$\Delta\tau \leq \frac{1}{c_3(1 + c_2)\gamma(Q_{i+1}, \zeta_{i+1})\mu_{\text{norm}}(Q_i, \zeta_i)}. \tag{13}$$

At this stage we use a fundamental result, known as the Higher Derivative Estimate, which states that, for all $F \in \mathcal{H}_{\mathbf{d}}$ and $z \in \mathbb{P}^n$,

$$\gamma(F, z) \leq \frac{1}{2}D^{3/2}\mu_{\text{norm}}(F, z), \tag{14}$$

where $D := \max\{d_1, \ldots, d_n\}$, to replace (13) by

$$\Delta\tau \leq \frac{2}{c_3(1 + c_2)D^{3/2}\mu_{\text{norm}}(Q_{i+1}, \zeta_{i+1})\mu_{\text{norm}}(Q_i, \zeta_i)}.$$

The use of Lipschitz bounds, now for μ_{norm}^{-1}, allows one to obtain

$$\max\{\mu_{\text{norm}}(Q_{i+1}, \zeta_{i+1}), \mu_{\text{norm}}(Q_i, \zeta_i)\} \leq (1 + c_4)\mu_{\text{norm}}(Q_i, z_i) \tag{15}$$

for some small c_4. We can conclude that taking

$$\Delta\tau \leq \mathcal{B}(Q_i, z_i) := \frac{2}{c_3(1 + c_2)(1 + c_4)^2D^{3/2}\mu_{\text{norm}}(Q_i, z_i)^2} \tag{16}$$

we ensure that

$$d_{\mathbb{P}}(\zeta_i, \zeta_{i+1}) \leq \frac{1}{c_3\gamma(Q_{i+1}, \zeta_{i+1})}. \tag{17}$$

Finally, by chosing $c_3 := \frac{6c_1}{c_1-6}$ we obtain from (10), (11), and (17) that

$$d_{\mathbb{P}}(z_i, \zeta_{i+1}) \leq \frac{1}{c_1\gamma(Q_{i+1}, \zeta_{i+1})} + \frac{1}{c_3\gamma(Q_{i+1}, \zeta_{i+1})} = \frac{1}{6\gamma(Q_{i+1}, \zeta_{i+1})}.$$

This general argument comes from [**21**] where the value of the constants c_i is not worked out[4]. Detailed proofs are in [**8, 9**] and [**10**, §17.1]. In the latter it is shown that one can actually chose them so that both (11) and (15) hold and that with that choice (16) becomes

$$\mathcal{B}(Q_i, z_i) := \frac{0.0085}{d_{\mathbb{S}}(G, F)D^{3/2}\mu_{\text{norm}}(Q_i, z_i)} \tag{18}$$

This results in an easy-to-describe algorithm LinHom that takes as input a triple (F, G, ζ_0) with $F, G \in \mathcal{H}_{\mathbf{d}}$ and $G(\zeta_0) = 0$.

---

[3]Actually, the true condition number would be $\mu(F, z)$ (see [**10**, Corollary 16.14]) which is defined in the same manner but without the diagonal matrix $\mathsf{diag}(\sqrt{d_i})$. The fact that μ_{norm} is unitarily invariant (whereas μ isn't), however, make $\mu_{\text{norm}}(F, z)$ more convenient to use.

[4]The author of [**21**] writes "In previous papers we have paid careful attention to the constants. In this paper we are more cavalier."

Algorithm LinHom
input $F, G \in \mathcal{H}_{\mathbf{d}}$ and $\zeta_0 \in \mathbb{P}^n$ such that $G(\zeta_0) = 0$

$\tau := 0$, $Q := G$, $z := \zeta_0$
`repeat`
  $\Delta\tau := \frac{0.0085}{d_{\mathbb{S}}(F,G)D^{3/2}\mu^2_{\mathrm{norm}}(Q,z)}$
  $\tau := \min\{1, \tau + \Delta\tau\}$
  $Q := Q_\tau$
  $z := N_Q(z)$
`until` $\tau = 1$
`RETURN` z

With some (but not much) additional work the arguments above show the following result (Proposition 17.3 in [**10**]).

Proposition 3.1. *Suppose that* $[G, F]$ *does not intersect* Σ. *Then, the execution of* LinHom(F, G, ζ_0) *stops after at most* K *steps with*

$$K = K(F, G, \zeta_0) \leq 188 D^{3/2} d_{\mathbb{S}}(G, F) \int_0^1 \mu^2_{\mathrm{norm}}(Q_\tau, \zeta_\tau) \mathrm{d}\tau.$$

The returned point z *is an approximate zero of* F. □

3.2. A probabilistic solution. The first breakthrough towards a solution of Smale's 17th problem was produced by Carlos Beltrán and Luis Miguel Pardo [**7**]. It consisted of a randomized algorithm, in the sequel BP, to generate the initial pair (G, ζ_0). In the course of its execution (with input $(n, \mathbf{d})$) this algorithm draws real numbers from the standard normal distribution and, because of this, its output (G, ζ_0) is random as well. If we define the *solution variety* to be

$$\mathcal{V}_{\mathbb{S}} := \{(G, \zeta) \in \mathbb{S}(\mathcal{H}_{\mathbf{d}}) \times \mathbb{P}^n \mid G(\zeta) = 0\} \tag{19}$$

then the distribution ρ_{std} induced on $\mathcal{V}_{\mathbb{S}}$ by the ouputs of BP can be easily described as follows:

(**i**) draw G from the uniform distribution on $\mathbb{S}(\mathcal{H}_{\mathbf{d}})$
(**ii**) draw ζ from the (discrete) uniform distribution among the $\mathcal{D}$ zeros of G (20)

A surprising feature of BP is that it constructs a polynomial system along with one of its zeros without ever solving a (nonlinear) polynomial system. Roughly speaking, it first draws the "linear part" of the system, then computes a zero of this linear part, and finally adds a suitable complement of higher degree terms. For a description of BP, its properties and its cost, see [**8, 9**] or [**10**, §17.6].

A zero-finding algorithm becomes clear, let's call it LV (from *Las Vegas*, as randomized algorithms as ours are called this way). On input $F \in \mathbb{S}(\mathcal{H}_{\mathbf{d}})$, we first draw a pair $(G, \zeta_0) \in \mathcal{V}_{\mathbb{S}}$ from ρ_{std} via a call to BP. We then run the linear homotopy on the triple (F, G, ζ_0). This returns an approximate zero of F. To understand its cost we next make two fundamental observations.

Firstly, if F, G are independently drawn from $\mathbb{S}(\mathcal{H}_{\mathbf{d}})$ and $\tau \in [0,1]$ is fixed, then the system $Q_\tau \in [G,F]$ given by

$$d_{\mathbb{S}}(G,Q_\tau) = \tau d_{\mathbb{S}}(G,F)$$

is also uniformly distributed in $\mathbb{S}(\mathcal{H}_{\mathbf{d}})$.

Secondly, let

$$\mu^2_{\text{avg}}(F) := \frac{1}{\mathcal{D}} \sum_{\zeta | F(\zeta)=0} \mu^2_{\text{norm}}(F,\zeta).$$

Then [**8**, Thm. 23],

$$\mathop{\mathbb{E}}_{Q\sim\mathbb{S}(\mathcal{H}_{\mathbf{d}})} \mu^2_{\text{avg}}(Q) \leq nN. \tag{21}$$

These facts allow us to bound the average number of steps of LV. Indeed,

$$\begin{aligned}
& \mathop{\mathbb{E}}_{F\sim\mathbb{S}(\mathcal{H}_{\mathbf{d}})} \mathop{\mathbb{E}}_{(G,\zeta_0)\sim\rho_{\text{std}}} K(F,G,\zeta_0) \\
\underset{\text{Prop. 3.1}}{\leq} \quad & 188D^{3/2} \mathop{\mathbb{E}}_{F\sim\mathbb{S}(\mathcal{H}_{\mathbf{d}})} \mathop{\mathbb{E}}_{(G,\zeta_0)\sim\rho_{\text{std}}} d_{\mathbb{S}}(G,F) \int_0^1 \mu^2_{\text{norm}}(Q_\tau,\zeta_\tau)\mathrm{d}\tau \\
= \quad & 188D^{3/2} \mathop{\mathbb{E}}_{F\sim\mathbb{S}(\mathcal{H}_{\mathbf{d}})} \mathop{\mathbb{E}}_{G\sim\mathbb{S}(\mathcal{H}_{\mathbf{d}})} d_{\mathbb{S}}(G,F) \int_0^1 \frac{1}{\mathcal{D}} \sum_{\zeta | Q_\tau(\zeta)=0} \mu^2_{\text{norm}}(Q_\tau,\zeta)\mathrm{d}\tau \\
\leq \quad & 188\,\pi D^{3/2} \int_0^1 \mathop{\mathbb{E}}_{F\sim\mathbb{S}(\mathcal{H}_{\mathbf{d}})} \mathop{\mathbb{E}}_{G\sim\mathbb{S}(\mathcal{H}_{\mathbf{d}})} \mu^2_{\text{avg}}(Q_\tau)\mathrm{d}\tau \\
= \quad & 188\,\pi D^{3/2} \int_0^1 \mathop{\mathbb{E}}_{Q\sim\mathbb{S}(\mathcal{H}_{\mathbf{d}})} \mu^2_{\text{avg}}(Q)\mathrm{d}\tau \\
\underset{(21)}{\leq} \quad & 188\,\pi D^{3/2} \int_0^1 nN\mathrm{d}\tau \;=\; 188\pi D^{3/2} nN.
\end{aligned}$$

In addition to this, the cost of each step, which is dominated by the computation of the Newton iteration, is $\mathcal{O}(N+n^3)$, which is $\mathcal{O}(N)$ if we assume that $d_i \geq 2$ for $i = 1,\dots,d$. This yields a total average cost of $\mathcal{O}(nD^{3/2}N^2)$ for LV (as the cost for the execution of BP is dominated by that of the linear homotopy). This $\mathcal{O}(nD^{3/2}N^2)$ average cost was independently obtained in [**8**] and [**9**].

A few last remarks are necessary. The average complexity bound above considers two different sorts of average. On the one hand, the initial pair (G,ζ_0) is produced by a randomized algorithm and hence, the algorithm LV itself is randomized. On the other hand, the input system F is considered random to derive average complexity bounds.

Algorithm LV is of Las Vegas type: for a given input $F \in \mathbb{S}(\mathcal{H}_{\mathbf{d}})$, if the algorithm halts, it returns an approximate zero of F (and it halts with probability 1 for all F outside a set of measure zero). Yet, its running time (its cost) is a random variable as it depends on (G,ζ_0). The argument above does not provide a bound for the *randomized cost* $\mathsf{randcost}(F)$ of LV on input F, which is defined as $\mathbb{E}_{(G,\zeta_0)\sim\rho_{\text{std}}} K(F,G,\zeta_0)$. But it does so for the average of these randomized costs over $F \in \mathbb{S}(\mathcal{H}_{\mathbf{d}})$.

Because of the probabilistic nature of LV, the bound above was not a solution to Smale's 17th problem. The latter asked for a deterministic algorithm. More work would be needed to reach this solution.

3.3. A near solution and other advances. The arguments above, which follow the exposition in [**8**], relie on (among other things) the fact that when F and G are uniformly drawn from $\mathbb{S}(\mathcal{H}_{\mathbf{d}})$, so is, for any $\tau \in [0,1]$, the system G_τ. But as soon as these distributions are not identical, technical difficulties may arise. This is at the root of the different choice of setting taken in [**9**], where the same $\mathcal{O}(D^{3/2}nN^2)$ bound for LV is proved as well. The setting here considers Gaussian distributions on $\mathcal{H}_{\mathbf{d}}$.

The Weyl norm induces a standard Gaussian distribution $N(0,\mathrm{I})$ on $\mathcal{H}_{\mathbf{d}}$ via the density

$$\rho(F) := \frac{1}{(2\pi)^N} e^{-\frac{\|F\|^2}{2}}$$

which in turn induces an *average cost* of an algorithm (taking inputs in $\mathcal{H}_{\mathbf{d}}$)

$$\mathop{\mathbb{E}}_{F\sim N(0,\mathrm{I})} \mathsf{cost}(F). \tag{22}$$

This quantity is not different to that in (7). Indeed, the standard Gaussian distribution on $\mathbb{R}^{2N}$ induces, via the bijection

$$\begin{array}{rcl} \mathbb{R}^{2N}\setminus\{0\} & \to & \mathbb{S}(\mathbb{R}^{2N})\times(0,\infty) \\ F & \mapsto & \left(\dfrac{F}{\|F\|}, \|F\|^2\right), \end{array}$$

the uniform distribution on $\mathbb{S}(\mathbb{R}^{2N})$ and a χ^2-distribution with $2N$ degrees of freedom on $(0,\infty)$. Moreover these two components are independent [**10**, Prop. 2.19]. Using the fact that the right-hand side in Proposition 3.1 is scale-invariant on F it immediately follows that

$$\mathop{\mathbb{E}}_{F\sim\mathbb{S}(\mathcal{H}_{\mathbf{d}})} \mathsf{cost}(F) = \mathop{\mathbb{E}}_{F\sim N(0,\mathrm{I})} \mathsf{cost}(F).$$

The fact that it is scale-invariant on G allows one to use a version of BP that returns a pair in

$$\mathcal{V}_{\mathcal{H}} := \{(G,\zeta) \in \mathcal{H}_{\mathbf{d}} \times \mathbb{P}^n \mid G(\zeta) = 0\} \tag{23}$$

instead of a pair in $\mathcal{V}_{\mathbb{S}}$. Furthermore, if F and G are Gaussians, then so is

$$Q_t = tF + (1-t)G$$

for every $t \in [0,1]$. We cannot however directly use these facts in conjunction with Proposition 3.1 because there τ parameterizes a fraction of the angle $d_{\mathbb{S}}(G,F)$ and now t parameterizes a fraction of $\|F-G\|$. But it is easy to go from one to the other via the change of variables

$$t = \frac{\|G\|}{\|F\|\sin\alpha\cot(\tau\alpha) - \|F\|\cos\alpha + \|G\|}$$

where $\alpha = d_{\mathbb{S}}(G,F)$. Then Proposition 3.1 takes the following form.

Corollary 3.2. *Suppose that $[G,F]$ does not intersect Σ. Then the execution of $\mathsf{LinHom}(F,G,\zeta_0)$ stops after at most K steps with*

$$K = K(F,G,\zeta_0) \le 188 D^{3/2}\|F\|\|G\| \int_0^1 \frac{\mu^2_{\mathrm{norm}}(Q_t,\zeta_t)}{\|Q_t\|^2}\mathrm{d}t.$$

The returned point z is an approximate zero of F. □

We won't describe the derivation of the bound for

$$\mathop{\mathbb{E}}_{F\sim N(0,\mathrm{I})}\ \mathop{\mathbb{E}}_{(G,\zeta_0)\sim\rho_{\mathsf{std}}} K(F,G,\zeta_0).$$

Suffice it to mention that, although different in the details, it is similar to the one described in §3.2 and it yields the same $\mathcal{O}(D^{3/2}nN)$ bound.

One of the differences is the need to estimate $\mathbb{E}_{Q\sim N(0,\sigma^2\mathrm{I})}\frac{\mu^2_{\mathrm{avg}}(Q)}{\|\bar{Q}\|^2}$ instead of $\mathbb{E}_{Q\sim\mathbb{S}(\mathcal{H}_{\mathbf{d}})}\mu^2_{\mathrm{avg}}(Q)$ (cf. (21)). It is here when one notes that if any of F or G is not centered (i.e., have a mean different from 0), then Q_t is still Gaussian (even though not centered). In particular, this is true when either F or G are fixed. If F, for instance, is fixed, then Q_t is distributed as $N(F,(1-t)^2\mathrm{I})$. This fact was the trigger to extend (21) not only to the Gaussian setting but also to the non-centered case. The main technical result in [**9**] (see Theorem 3.6 there or [**10**, Thm. 18.4]) shows that for all $\bar{Q}\in\mathcal{H}_{\mathbf{d}}$ and $\sigma>0$ we have

$$\mathop{\mathbb{E}}_{Q\sim N(\bar{Q},\sigma^2\mathrm{I})}\frac{\mu^2_{\mathrm{avg}}(Q)}{\|Q\|^2}\leq\frac{e(n+1)}{2\sigma^2}. \tag{24}$$

With this *smoothed analysis* at hand, and some additional work, two different complexity results followed by fixing, respectively, F and (G,ζ_0).

3.3.1. *Instance complexity.* Fixing $F\in\mathcal{H}_{\mathbf{d}}$ and looking at the cost of LV with input F lead us to the quantity $\mathsf{randcost}(F)$. The first of the two results mentioned above ([**9**, Thm. 3.7] or [**10**, Thm. 18.2]) shows that, for algorithm LV and $F\in\mathcal{H}_{\mathbf{d}}$,

$$\mathsf{randcost}(F)\leq\mathcal{O}(D^3nN^2\mu^2_{\max}(F))$$

where $\mu_{\max}(F)=\max_{\zeta|F(\zeta)=0}\mu_{\mathrm{norm}}(F)$. Reasonably enough, the randomized cost for F depends on how well-conditioned the zeros of F are.

3.3.2. *Deterministic algorithms.* If we fix the initial pair (G,ζ_0) instead, and we take the average of the cost of the linear homotopy for an input F drawn from $N(0,\mathrm{I})$ we obtain our second result, namely

$$\mathop{\mathbb{E}}_{F\sim N(0,\mathrm{I})}K(F,G,\zeta_0)=\mathcal{O}(D^3nN\mu^2_{\max}(G))$$

which gives an average total cost of $\mathcal{O}(D^3nN^2\mu^2_{\max}(G))$. This bound deceptively suggests that the solution to Smale's 17th problem is near; it suffices to construct, for a given pair $(n,\mathbf{d})$, a system $G\in\mathcal{H}_{\mathbf{d}}$ with $\mu_{\max}(G)=N^{\mathcal{O}(1)}$. We say 'deceptively' because this innocent-looking task proved to be remarkably difficult. A first attempt to use this bound took as G the system given by

$$G_i:=\frac{1}{\sqrt{2n}}(X_0^{d_i}-X_i^{d_i}),\qquad i=1,\dots,n$$

and as ζ_0 the point $[(1,\dots,1)]$. One can show that $\mu^2_{\max}(G)\leq 2(n+1)^D$ and, together with a different algorithmic approach for high-degree systems (those with $D>n$), the following result ([**9**, Thm. 3.9] or [**10**, Thm. 18.3]).

Theorem 3.3. *There is a deterministic algorithm that on input $F\in\mathcal{H}_{\mathbf{d}}$ computes an approximate zero of F with average cost $N^{\mathcal{O}(\log\log N)}$. Moreover, if we restrict data to systems satisfying*

$$D\leq n^{\frac{1}{1+\varepsilon}}\qquad or\qquad D\geq n^{1+\varepsilon}$$

for some fixed $\varepsilon > 0$, then the average time of the algorithm is polynomial in the input size N. □

The quest for a construction of a good initial pair (G, ζ_0), raised already in the Bézout series, would not be closed until very recently [**6**]. But by then Smale's 17th problem had already been solved.

3.4. The solution. The final solution of Smale's 17th problem by Pierre Lairez [**16**] relies on a beautifully ingenious idea. Beltrán and Pardo had avoided constructing a pair (G, ζ_0) that would work for all inputs $F \in \mathbb{S}(\mathcal{H}_{\mathbf{d}})$ via a randomized construction: the pair (G, ζ_0) depended on random reals drawn from a random number generator. Lairez's solution uses this construction with a twist: the random numbers are extracted from F itself.

To better understand this idea, let's consider a situation simplified to the extreme. Take a random number $F \in [0, 1]$ from the uniform distribution (by glueing the extremes of this interval we can look at it as the circle $\mathbb{S}^1$, a simplified version of $\mathbb{S}(\mathcal{H}_{\mathbf{d}})$). Fix a number $\ell \in \mathbb{N}$. We can associate to F the numbers

$$F^\circ := 2^{-\ell}\lfloor 2^\ell F \rfloor \qquad \text{and} \qquad R^\circ := 2^\ell(F - F^\circ).$$

The first, the *truncation* of F, approximates F by taking the first ℓ bits of its base-2 expansion. It satisfies $F - F^\circ \leq 2^{-\ell}$. The second, the *fractional part* of F, is uniformly distributed in $[0, 1]$ and independent of F°. Lairez devised a similar procedure now for systems $F \in \mathbb{S}(\mathcal{H}_{\mathbf{d}})$ drawn from the uniform distribution. The resulting systems F° and R° are both in $\mathbb{S}(\mathcal{H}_{\mathbf{d}})$ and satisfy that F° approximates F, and that R° is *nearly* uniform and *nearly* independent from F°.

The fractional part R° is then used as random source to produce a pair $(G, \zeta_0) \sim \rho_{\mathsf{std}}$ by a version of the BP algorithm which replaces the set R of numbers obtained from a random number generator by the coefficients of R°. Once this done, a linear homotopy is performed with initial pair (G, ζ_0) and target system F°. Had we have started with $F \sim \mathbb{S}(\mathcal{H}_{\mathbf{d}})$ and the set R (the random numbers in the call to BP) independent of F, the systems F and G would be uniform in $\mathbb{S}(\mathcal{H}_{\mathbf{d}})$ and independent, and the execution $\mathsf{LinHom}(F, G, \zeta_0)$ would end on an approximate zero of F. As it happens, we are starting with F° and R° nearly independent —but not exactly so— and F° close to F —but not exactly there—.

The fact we want to rely on is that the larger is ℓ the better are these approximations. And that we can quantify this dependence. Lairez used this fact to devise a simple test that ensures that the curve $\mathscr{C}^\circ$ associated to the triple (F°, G, ζ_0) is close enough to the curve $\mathscr{C}$ associated to (F, G, ζ_0), that the approximate zero of F° obtained by following $\mathscr{C}^\circ$ is also an approximate zero of F, and that the average complexity analysis can still be carried on. It suffices to ensure that, at every step of the linear homotopy,

$$D^{3/2}\mu^2_{\mathrm{norm}}(Q_i, z_i)\varrho \leq \frac{1}{151} \tag{25}$$

where ϱ is an easy-to-compute quantity satisfying $d_{\mathbb{S}}(F, F^\circ) \leq \varrho$. Hence, ϱ (and a fortiori $d_{\mathbb{S}}(F, F^\circ)$) must be smaller than $(151\, D^{3/2}\mu^2_{\mathrm{norm}}(Q, z))^{-1}$ all along (a neighborhood of) $\mathscr{C}^\circ$. As we don't know a priori a bound for the maximum of these μ^2_{norm}s, the final algorithm starts with an initial value for ℓ, computes the associated F° and R°, computes (Q, ζ_0) from R°, and calls for $\mathsf{LinHom}(F^\circ, G, \zeta_0)$.

If condition (25) is satisfied at all steps, then the returned point in $\mathbb{P}^n$ is an approximate zero of F. If at some step (25) is violated then the algorithm replaces ℓ by 2ℓ and starts again. The complexity analysis in [**16**] shows that the average total cost of this deterministic procedure is $\mathcal{O}(nD^{3/2}N^2)$. This settles Smale's 17th problem.

4. Subsequent Progress

The positive answer in Lairez's paper did not bring to a stop the interest on Smale's 17th problem. On the contrary, a number of questions naturally arose.

4.1. Eigenpair computations. One such question dealt with applying the ideas for the design and analysis of linear homotopies to specific systems of equations which do not entirely fit the framework above. A clear example is the computation of eigenpairs. Recall, an *eigenpair* of a matrix $A \in \mathbb{C}^{n\times n}$ is a pair $(\lambda, v) \in \mathbb{C}\times\mathbb{P}^{n-1}$ satisfying $Av = \lambda v$. This equality amounts to the n equations

$$\begin{array}{rcl} A_1 v & = & \lambda v_1 \\ & \vdots & \\ A_n v & = & \lambda v_n \end{array} \tag{26}$$

where A_j denotes the jth row of A. These equations are linear in the variables $v_1, \ldots, v_n$ and quadratic in the set of all the variables. There are n equations with (generically) a finite number of of solutions in $\mathbb{C}\times\mathbb{P}^{n-1}$ but, in glaring contrast with the general framework, we expect system (26) to have only n solutions, not 2^n. A tailor-made approach is necessary.

Algorithms computing eigenpairs have existed for long. But their analysis, notwithstanding their practical performance, has been wanting. This moved James Demmel [**12**, pg. 139] to write

> *So the problem of devising an algorithm [for the eigenvalue problem] that is numerically stable and globally (and quickly!) convergent remains open.*

An answer to this problem was given in [**3**], where the general lines in the previous sections —devising an appropriate condition number and a step-length computation in terms of it, analysing the number of steps of the linear homotopy as in Proposition 3.1, performing a smoothed analysis of the condition number as in (24), and exhibiting a good initial triple (A_0, λ_0, v_0) (which in this context turns out to be easy!)— allowed one to show an $\mathcal{O}(n^7)$ bound for the average cost of computing eigenpairs. We note that this algorithm, even though its observed average complexity in simulations is $\mathcal{O}(n^{3.66})$, is not efficient when compared with the eigensolvers available in packages such as `Matlab` or `Julia`. But it answers Demmel's question in the sense that it was the first eigensolver for which correctness, numerical stability, and a polynomially bounded average cost could be established.

4.2. Rigid homotopies. Probably the most obvious question raised by the solution to Smale's 17th problem dealt with possible complexity improvements. Because just reading the system F takes time N, a lower bound of $\Omega(N)$ is clear. The upper bound for the average cost of LV is $\mathcal{O}(nD^{3/2}N^2)$. We observe that in this bound the crucial parameter is N. To see why, think on the case $n = D$. In this case, $N \geq \binom{n+D}{n} = \binom{2n}{n} \sim \frac{4^n}{\sqrt{\pi n}}$. That is, N is exponentially larger than n (and D). We therefore want to understand how small can we make the exponent of N. We

already noted that the cost of each iteration of the homotopy is $\mathcal{O}(N)$. This cannot be improved. The question then becomes, can one decrease the expected number of iterations (without paying for it in the cost of each iteration)?

A first answer towards this goal was given in [**2**] where a more elaborated computation of $\Delta\tau$, which draw from ideas in [**3**], allowed for a $\mathcal{O}(D^{3/2}nN^{1/2})$ bound for the expected number of iterations (and hence an $\mathcal{O}(D^{3/2}nN^{3/2})$ bound for the average total cost).

An optimal answer was provided by Lairez [**17**]. The involved ideas were new and touched both the algorithm and its analysis.

Let us begin with the algorithmic ideas. The innovation here is to perform a homotopy that keeps invariant the shape of the hypersurfaces $f_1 = 0, \dots, f_n = 0$ (hence the name *rigid*) continuously deforming only their position within $\mathbb{P}^n$. To describe how this is done, denote by $U(n+1)$ the group of unitary matrices of dimension $n+1$. This group has an obvious action on $\mathbb{P}^n$; for any $u \in U(n+1)$ and $z \in \mathbb{C}^{n+1} \setminus \{0\}$ we define $u([z]) := [u(z)]$. It also acts on the space H_d of homogeneous polynomials of degree d in $X_0, \dots, X_n$. If $f \in H_d$ and $u \in U(n+1)$ then $uf := f \circ u^{-1}$ is easily seen to be in H_d. Its zero set $Z(uf)$ is just a rotation of $Z(f)$; indeed, $Z(uf) = u(Z(f))$. This action extends componentwise to systems. Let $\mathcal{U} := U(n+1)^n$. If $F \in \mathcal{H}_{\mathbf{d}}$ and $\mathbf{u} = (u_1, \dots, u_n) \in \mathcal{U}$ then $\mathbf{u}f := (u_1 f_1, \dots, u_n f_n) \in \mathcal{H}_{\mathbf{d}}$.

Assume next that we are given a pair $(\mathbf{v}, \zeta_0) \in \mathcal{U} \times \mathbb{P}^n$ such that $Q_0(\zeta_0) = 0$ where $Q_0 := \mathbf{v}F$. We can then start a homotopy with the pair (Q_0, ζ_0). The breakthrough lies on the fact that the segment we will lift is not the segment $[Q_0, F] \subset \mathcal{H}_{\mathbf{d}}$ but the geodesic $[\mathbf{v}, \mathbf{I}] \subset \mathcal{U}$. Here $\mathbf{v} = (v_1, \dots, v_n)$ and $\mathbf{I} = (\mathrm{I}, \dots, \mathrm{I})$. This lifting lives in the solution variety (compare with (19) and (23))

$$\mathcal{V}_{\mathcal{U}} := \{(\mathbf{u}, \zeta) \in \mathcal{U} \times \mathbb{P}^n \mid (f_i \circ u_i)(\zeta) = 0,\ i = 1, \dots, n\}.$$

In practice, we don't need to use the geodesic $[\mathbf{u}, \mathbf{I}]$; we can use instead any path $\{\mathbf{u}_t\}_{t \in [0,T]}$ with endpoints $\mathbf{v}$ and $\mathbf{I}$ and satisfying a few conditions (Lairez imposes being 1-Lipschitz and having length at most $4n$). The corresponding systems in the homotopy are then

$$Q_t := \mathbf{u}_t F = ((\mathbf{u}_t)_1 f_1, \dots, (\mathbf{u}_t)_n f_n).$$

Figure 4 aims at depicting the process.

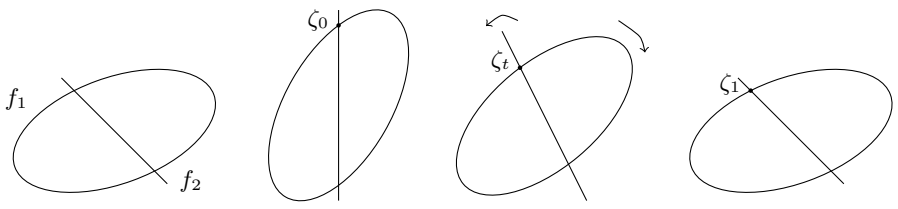

Figure 4. A rigid homotopy. **Left.** The given hypersurfaces. **Second.** The rotated hypersurfaces with a common zero. **Third.** An intermediate stage in the homotopy. **Right.** The final stage of the homotopy: the original hypersurfaces with a common zero.

The general arrangement in Lairez's algorithm is otherwise the same as in LV. We first generate the random pair $(\mathbf{u}_0, \zeta_0) \in \mathcal{V}_{\mathcal{U}}$. To do so Lairez develops in this context a version of the BP algorithm which draws $\mathcal{O}(n^3)$ real numbers from the standard Gaussian distribution and performs $\mathcal{O}(n^4 + nD^4)$ arithmetic operations. The distribution $\rho_{\mathcal{U}}$ in $\mathcal{V}_{\mathcal{U}}$ induced by this version of BP is similar to ρ_{std}; it just replaces $G \in \mathbb{S}(\mathcal{H}_{\mathbf{d}})$ by $\mathbf{u} \in \mathcal{U}$ in (20(i)), and G by $\mathbf{u}F$ in (20(ii)).

After which, the homotopy proceeds as in the linear case, with steps now in the path $\{\mathbf{u}_t\}_{t\in[0,T]} \subset \mathcal{U}$ instead of in the segment $[G, F]$.

The computation of the length of these steps is what brings us to the innovations in the analysis of the homotopy. Lairez observed that in the derivation of $\mathcal{B}$ leading to (18) two inequalities are coarse. Firstly, (12), which bounds the variation of the zero ζ as a worst-case variation. Secondly, the Higher Derivative Estimate (14) that replaces γ by μ_{norm}. In both cases the use of $\mu_{\mathrm{norm}}(F, z)$ is computationally practical but leads to coarse bounds. One would like to use measures serving the same purposes, yielding finer bounds, and being also inexpensive to compute. To achieve this, Lairez splits condition at a point into two components. Given $F \in \mathcal{H}_{\mathbf{d}}$ and $z \in \mathbb{P}^n$, we let

$$F_z := \left(\frac{f_1}{\|\mathrm{D}_z f_1\|}, \dots, \frac{f_n}{\|\mathrm{D}_z f_n\|} \right).$$

The *incidence condition number* of F at z is then given by

$$\kappa(F, z) := \left\| (\mathrm{D}_z F_z)^\dagger \right\|.$$

Here $\dagger$ denotes Moore-Penrose inverse. When $F(z) = 0$ this quantity depends only on the angles made by the tangent spaces at z of the n hypersurfaces $Z(f_i) := \{z \in \mathbb{P}^n \mid f_i(z) = 0\}$. Moreover, $\kappa(\mathbf{u}F, z)$ satisfies a version of (12) w.r.t. variations of $\mathbf{u}$ in $\mathcal{U}$ (see [**17**, Lem. 16]). That is, we have

$$d_{\mathbb{P}}(\zeta_i, \zeta_{i+1}) \lesssim \Delta t\, \kappa(\mathbf{u}_i F, \zeta_i). \tag{27}$$

For the other ingredient, we define, for $f \in H_d$ and $z \in \mathbb{P}^n$,

$$\gamma_{\mathrm{Frob}}(f, z) := \begin{cases} \sup_{k\geq 2} \left(\frac{1}{k!}\|\mathrm{D}_z f\|^{-1}\|\mathrm{D}_z^k f\|_{\mathrm{Frob}}\right)^{\frac{1}{k-1}} & \text{if } \mathrm{D}_z f \text{ is nonzero} \\ \infty & \text{otherwise} \end{cases}$$

and we let

$$\bar{\gamma}^2_{\mathrm{Frob}}(F, z) := \gamma^2_{\mathrm{Frob}}(f_1, z) + \cdots + \gamma^2_{\mathrm{Frob}}(f_n, z).$$

In contrast with $\kappa(F, z)$, $\bar{\gamma}_{\mathrm{Frob}}(F, z)$ does not depend on how the hypersurfaces $Z(f_i)$ intersect at z but only on the geometry of each individual hypersurface at z.

The *split Frobenius γ number*

$$\hat{\gamma}_{\mathrm{Frob}}(F, z) := \kappa(F, z)\, \bar{\gamma}_{\mathrm{Frob}}(F, z) \tag{28}$$

captures both aspects. Moreover, it bounds $\gamma(F, z)$ as we have

$$\gamma(F, z) \leq \hat{\gamma}_{\mathrm{Frob}}(F, z) \tag{29}$$

thus providing us with a replacement for (14) which, we will see, turns out to be better for our purposes.

With (27) and (29) respectively replacing (12) and (14) we can reason as in §3.1 and use now Lipschitz bounds for κ and γ_{Frob} to derive an expression for the step

length of the form

$$\Delta t = \frac{c}{\kappa(\mathbf{u}_i F, z_i)\,\hat{\gamma}_{\mathrm{Frob}}(\mathbf{u}_i F, z_i)}$$

for some constant c. This, in turn, leads to a bound for the number of steps in the homotopy of the form (compare with Proposition 3.1)

$$K = K_F(\mathbf{u}, \zeta_0) \leq C \int_0^T \kappa(\mathbf{u}_t F, \zeta_t)\,\hat{\gamma}_{\mathrm{Frob}}(\mathbf{u}_t F, \zeta_t)\mathrm{d}t \tag{30}$$

for some constant C. Furthermore, Lairez [**17**, Prop. 17] shows that, for a given $F \in \mathcal{H}_{\mathbf{d}}$ with square-free components,

$$\mathop{\mathbb{E}}_{(\mathbf{u},\zeta)\sim\rho_{\mathcal{U}}} \kappa^2(\mathbf{u}F, \zeta) \leq 6n^2 \tag{31}$$

and that [**17**, Lem 38], for a Gaussian $f \in H_d$,

$$\mathop{\mathbb{E}}_{f\sim N(0,\mathrm{I})}\ \mathop{\mathbb{E}}_{\zeta\sim Z(f)} \gamma^2_{\mathrm{Frob}}(f, \zeta) \leq \frac{1}{4} d^3(d+n).$$

We can use this inequality to bound, for a Gaussian $F \in \mathcal{H}_{\mathbf{d}}$, the value of $\mathbb{E}_{F\sim N(0,\mathrm{I})}\,\mathbb{E}_{(\mathbf{u},\zeta)\sim\rho_{\mathcal{U}}}\,\bar{\gamma}^2_{\mathrm{Frob}}(\mathbf{u}F, \zeta)$. Indeed, if $(\mathbf{u}, \zeta) \in \mathcal{V}_{\mathcal{U}}$ is $\rho_{\mathcal{U}}$-distributed then $u_1\zeta, \ldots, u_n\zeta$ are zeros of $f_1, \ldots, f_n$, uniformly distributed in $Z(f_1), \ldots, Z(f_n)$, respectively, and independent [**17**, Thm. 8]. Also, if F is Gaussian, so are its components f_i. It follows that

$$\begin{aligned}
\mathop{\mathbb{E}}_{F\sim N(0,\mathrm{I})}\ \mathop{\mathbb{E}}_{(\mathbf{u},\zeta)\sim\rho_{\mathcal{U}}} \bar{\gamma}^2_{\mathrm{Frob}}(\mathbf{u}F, \zeta) &= \mathop{\mathbb{E}}_{(\mathbf{u},\zeta)\sim\rho_{\mathcal{U}}} \sum_{i=1}^n \mathop{\mathbb{E}}_{f_i\sim N(0,\mathrm{I})} \gamma^2_{\mathrm{Frob}}(u_i f_i, \zeta) \\
&= \mathop{\mathbb{E}}_{(\mathbf{u},\zeta)\sim\rho_{\mathcal{U}}} \sum_{i=1}^n \mathop{\mathbb{E}}_{f_i\sim N(0,\mathrm{I})} \gamma^2_{\mathrm{Frob}}(f_i, u_i\zeta) \\
&= \sum_{i=1}^n \mathop{\mathbb{E}}_{f_i\sim N(0,\mathrm{I})}\ \mathop{\mathbb{E}}_{\zeta_i\sim Z(f_i)} \gamma^2_{\mathrm{Frob}}(f_i, \zeta_i) \\
&\leq \frac{1}{4} n D^3(D+n).
\end{aligned} \tag{32}$$

At this stage, the reasoning continues as in §3.2. One firstly shows that when F is Gaussian (or, equivalently, when $F \sim \mathbb{S}(\mathcal{H}_{\mathbf{d}})$) and $(\mathbf{u}, \zeta) \sim \rho_{\mathcal{U}}$, the pair $(\mathbf{u}_t, \zeta_t)$ in (30) is also $\rho_{\mathcal{U}}$-distributed in $\mathcal{V}_{\mathcal{U}}$. Hence

$$\begin{aligned}
\mathop{\mathbb{E}}_{F\sim N(0,\mathrm{I})}\ \mathop{\mathbb{E}}_{(\mathbf{u},\zeta)\sim\rho_{\mathcal{U}}} K &\leq C \mathop{\mathbb{E}}_{F\sim N(0,\mathrm{I})} \int_0^T \mathop{\mathbb{E}}_{(\mathbf{u}_t,\zeta_t)\sim\rho_{\mathcal{U}}} \kappa(\mathbf{u}_t F, \zeta_t)\,\hat{\gamma}_{\mathrm{Frob}}(\mathbf{u}_t F, \zeta_t)\mathrm{d}t \\
&\leq 4Cn \mathop{\mathbb{E}}_{F\sim N(0,\mathrm{I})}\ \mathop{\mathbb{E}}_{(\mathbf{u},\zeta)\sim\rho_{\mathcal{U}}} \kappa(\mathbf{u}F, \zeta)\,\hat{\gamma}_{\mathrm{Frob}}(\mathbf{u}F, \zeta) \\
&\underset{(28)}{=} 4Cn \mathop{\mathbb{E}}_{F\sim N(0,\mathrm{I})}\ \mathop{\mathbb{E}}_{(\mathbf{u},\zeta)\sim\rho_{\mathcal{U}}} \kappa^2(\mathbf{u}F, \zeta)\,\bar{\gamma}_{\mathrm{Frob}}(\mathbf{u}F, \zeta) \\
&= 4Cn \mathop{\mathbb{E}}_{F\sim N(0,\mathrm{I})}\ \mathop{\mathbb{E}}_{(\mathbf{u},\zeta)\sim\rho_{\mathcal{U}}} \kappa^2(\mathbf{u}F, \zeta) \\
&\qquad \mathop{\mathbb{E}}_{F\sim N(0,\mathrm{I})}\ \mathop{\mathbb{E}}_{(\mathbf{u},\zeta)\sim\rho_{\mathcal{U}}} \bar{\gamma}_{\mathrm{Frob}}(\mathbf{u}F, \zeta) \\
&\underset{(31-32)}{\leq} 4Cn \cdot 6n^2 \cdot \sqrt{\frac{1}{4} n D^3(D+n)} \\
&= \mathcal{O}(n^4 D^2).
\end{aligned} \tag{33}$$

For the second equality we used that κ and $\bar{\gamma}_{\mathrm{Frob}}$ are independent. For the last inequality we used that F has square-free components almost surely (along with (31)) and Jensen's inequality (along with (32)).

This is the bound we wished for. It shows that the average number of steps in the rigid homotopy is polynomial in n and D. It only remains to be seen that the cost of each such step is still linear in N. Lairez [**17**, Cor. 34] shows that this is the case: the cost of each step is $\mathcal{O}(n^2D^2N)$. It may be worth to describe, nonetheless, how $\gamma_{\mathrm{Frob}}(f_i, z)$ is computed for a given $f_i \in H_{d_i}$ and a point $z \in \mathbb{C}^{n+1}$ as this computation is not straightforward.

The crucial observation is the following. Let $g(X) \in H_{d_i}$ be the shifted polynomial $f_i(z + X)$. This is no longer a homogeneous polynomial. For $k = 0, \dots, d_i$ let $g^{[k]}$ denote the homogeneous component of degree k of g. Then $\mathrm{D}^k_z f = \mathrm{D}^k_0 g$ and [**17**, Lem. 30]

$$\left\| \frac{1}{k!} \mathrm{D}^k_z f \right\|_{\mathrm{Frob}} = \left\| \frac{1}{k!} \mathrm{D}^k_0 g \right\|_{\mathrm{Frob}} = \left\| g^{[k]} \right\| \tag{34}$$

where, we emphasize, the norm of the last term is the Weyl norm. This equality provides a way to compute the left-hand side: compute g, then its component $g^{[k]}$, and finally the Weyl norm of the latter. The cost of this preocedure is dominated by the cost of the computation of g. Let $N_i = \dim_{\mathbb{C}} H_{d_i} = \binom{n+d_i}{n}$. A naive approach to compute g leads to an $\Omega(N_i^2)$ which is too large. But Lairez devised a procedure to compute g whose cost is $\mathcal{O}(nd_i^2N_i)$ [**17**, Prop. 32]. Since the computation of $\|\mathrm{D}_z f_i\|^{-1}$ has cost $\mathcal{O}(N_i)$ it follows that we can compute $\gamma_{\mathrm{Frob}}(f_i, z)$ with cost $\mathcal{O}(nd_i^2N_i)$. Furthermore, if $u_i \in U(n+1)$ we have $\gamma_{\mathrm{Frob}}(u_i f_i, z) = \gamma_{\mathrm{Frob}}(f_i, u_i z)$ so that we can compute $\gamma_{\mathrm{Frob}}(u_i f_i, z)$ with the same cost. We conclude that, given $F \in \mathcal{H}_{\mathbf{d}}$, $\mathbf{u} \in \mathcal{U}$, and $z \in \mathbb{C}^{n+1}$, we can compute $\bar{\gamma}_{\mathrm{Frob}}(\mathbf{u}F, z)$ with cost $\mathcal{O}(nD^2N)$.

The other ingredients in the computation of the step-length being simpler, one concludes the analysis with a total cost for this computation of $\mathcal{O}(n^2D^2N)$. In conjunction with the bound for the average number of steps in (33), and the $\mathcal{O}(n^4 + nD^4)$ cost of computing the initial pair, we end up with a $\mathcal{O}(n^6D^4N)$ for the average total cost of computing one zero of F. The exponent 1 on N is optimal.

4.3. Structured systems. Consider the polynomial

$$a_0X_0^3X_1^2X_2^3 + a_1X_1^5X_2^3 + a_2X_0^8 + a_3X_0X_1^2X_2^5. \tag{35}$$

We can describe it via the list

$$\{(a_0, (3, 2, 3)), (a_1, (0, 5, 3)), (a_2, (8, 0, 0)), (a_3, (1, 2, 5))\} \tag{36}$$

which, even with the integers $3, 2, 3, \dots$ written in binary, is substantially shorter than the list of all the $\binom{8+3}{3} = 165$ coefficients of a generic polynomial in H_8 (in X_0, X_1, X_2). The list in (36) is a *sparse encoding* of the polynomial above; the list of all coefficients (in some preestablished order) is known as *dense encoding*. It is the one we have been considering in all the previous sections and has a size of $\Theta(\dim_{\mathbb{C}} H_d)$. It is easy to see that the sparse encoding of a polynomial cannot be much bigger than its dense encoding but can be way smaller. A price for this succinctness is fragility: a number of operations with polynomials destroy sparseness. Think for instance in quotient and remainder applied to the polynomials $X_0^n - X_1^n$ and $X_0 - X_1$. They are both sparse (in the sense of having only two monomials)

but their quotient is $X_0^{n-1} + X_0^{n-2}X_1 + \cdots + X_1^{n-1}$ which has n monomials. Only writing this quotient has a high cost in the sparse size of the data, in contrast with its cost as a function on the dense size, which is small.

Another way to describe a polynomial, which can be even more succint than a sparse encoding, is as a *straight-line program* (SLP for short). The sequence of operations (starting with the variables X_0 and X_1)

$$\begin{aligned} X_0 + X_1 &\to (X_0 + X_1)^2 \to (X_0 + X_1)^4 \to (X_0 + X_1)^8 \to (X_0 + X_1)^{16} \\ &\vdots \\ &\to (X_0 + X_1)^{2^n} = \sum_{j=0}^{2^n} \binom{2^n}{j} X_0^{2^n - j} X_1^j \end{aligned}$$

performs $n + 1$ operations and ends up in a polynomial f of degree 2^n for which both the dense and sparse encodings have size $\Omega(2^n)$. We can encode f via the sequence of operations above and this encoding has size $\mathcal{O}(n)$. Not unexpectedly, the catalog of operations which are expensive in terms of the SLP encoding is bigger. The only operation that is clearly efficiently performed in terms of the SLP size, besides arithmetic operations with polynomials, is evaluation: given an SLP in n variables $X_1, \ldots, X_n$ encoding a polynomial f and a point $z \in \mathbb{C}^n$, the computation of $f(z) \in \mathbb{C}$ takes time linear in the SLP size. A fundamental result of Baur and Strassen [**5**] shows that we can compute all of $f(z), \partial_{X_1} f(z), \ldots, \partial_{X_n} f(z)$ with essentially the same cost.

The question which is naturally posed is the following:

What is the cost of computing a zero of a system in terms of its sparse (or SLP) size? (S17-Sparse)

Before attempting any answer to this question it is worth to look at it with more detail. Consider the sparse size, to fix ideas. A first remark is that the set of polynomials in H_d having some coefficient equal to 0 has measure zero in H_d. Consequently, the average cost over $\mathcal{H}_{\mathbf{d}}$ in terms of the sparse size of systems in $\mathcal{H}_{\mathbf{d}}$ will be the same as that in terms of its dense size. One may refine the question above and consider only polynomials with a given monomial structure. For instance, one of the polynomials may be imposed to have the form in (35). The average (regarding this polynomial) would then be taken, not over the whole of H_8, but over the tuples (a_0, a_1, a_2, a_3) of coefficients, say from a Gaussian distribution on $\mathbb{C}^4$. Other components of the input system $F \in \mathcal{H}_{\mathbf{d}}$ may also have a (possibly different) fixed monomial structure.

This determines a space of inputs which is a linear subspace $\mathcal{L}$ of $\mathcal{H}_{\mathbf{d}}$ over which a Gaussian measure is naturally defined (or a uniform measure on its associated unit sphere). And suggests that an approach as described in Sections 2 and 3 could be possible. Unfortunately though, such an approach is plagued with difficulties.

Firstly, one observes that it may happen that all systems in $\mathcal{L}$ are singular. Indeed, this is the case, for instance, if a component of F has the form

$$\begin{aligned} & a_0 X_0^3 X_1^2 X_2^3 + a_1 X_0^2 X_1^3 X_2^3 + a_2 X_0^8 + a_3 X_0^2 X_1 X_2^5 \\ = \; & X_0^2 (a_0 X_0 X_1^2 X_2^3 + a_1 X_1^4 X_2^2 + a_2 X_0^6 + a_3 X_1 X_2^5) \end{aligned}$$

because all the zeros of F on the hypersurface $\{X_0 = 0\}$ are double. For systems in such an $\mathcal{L}$ any linear homotopy will loop forever and the average complexity of the algorithm will be infinity.

Secondly, $\mathcal{L}$ is not unitary invariant. In general, given $F \in \mathcal{L}$ and $\mathbf{u} \in \mathcal{U}$, we don't have $\mathbf{u}F \in \mathcal{L}$. This deprives us from a very powerful technical tool.

Nonetheless research in homotopy methods for sparse systems has been carried out for decades. And even though there are no conclusive results, advances have been made. The state of the art in this quest can be found in the papers [**18, 19**] and in the references therein.

4.4. Low-complexity systems. The framework of rigid homotopies introduced by Lairez naturally considers (possibly small) subsets of $\mathcal{H}_{\mathbf{d}}$. Indeed, for a given input $F \in \mathcal{H}_{\mathbf{d}}$ the space where the homotopy path leading to F lives in is not the whole of $\mathcal{H}_{\mathbf{d}}$ (as in Sections 2 and 3) but its subset $\mathcal{U}F := \{\mathbf{u}F \mid \mathbf{u} \in \mathcal{U}\}$. The initial system G is not arbitrary in $\mathcal{H}_{\mathbf{d}}$ but of the form $\mathbf{v}F$ for some arbitrary $\mathbf{v} \in \mathcal{U}$. In addition, if $L(F)$ denotes the evluation complexity of F (that is, the length of the shortest SLP computing F) then $L(\mathbf{u}F) \leq L(F) + (n+1)^4$. In particular, if F has a low evaluation complexity then all systems $\mathbf{u}_t F$ in the homotopy path have low evaluation complexity as well.

These considerations are the motivation behind [**11**], where the goal is to devise an efficient rigid homotopy for low-complexity systems. On a first approach, a system $F \in \mathcal{H}_{\mathbf{d}}$ is fixed and random inputs are considered with the form $\mathbf{u}F$ where $\mathbf{u}$ is drawn from the uniform distribution in $\mathcal{U}$. This mimics the framework in §3.2 replacing random $F, G \in \mathbb{S}(\mathcal{H}_{\mathbf{d}})$ by random $\mathbf{u}, \mathbf{v} \in \mathcal{U}$.

At a first glance, it would seem that the algorithm in [**17**] can be applied without modifications. Unfortunately, a careful look at the details shows that there is a difficulty. For $f \in H_d$, the computation of $\gamma_{\mathrm{Frob}}(f, z)$ described at the end of §4.2 relies on computing Weyl norms $\|g_k\|$ for the homogeneous components g_k of the polynomial $g := f(z + X)$. If f is given by a short SLP then g is given by an equally short SLP. This is clear. Lemma 3.4 in [**11**] shows that we can actually do more: for any $w \in \mathbb{C}^{n+1}$ we can compute the values $g_0(w), \ldots, g_d(w)$ with cost $\mathcal{O}(dL(f) + n + \log d)$. So, the evaluation complexity of the g_ks is also small. We are left with the problem of computing their Weyl norms (obviously, without expanding the g_ks to obtain their coefficients; this would be too expensive). The way out lies on (yet another) useful property of the Weyl norm, namely, that for $g \in H_d$,

$$\|g\|^2 = \binom{n+d}{d} \mathop{\mathbb{E}}_{w \sim \mathbb{S}(\mathbb{C}^{n+1})} |g(w)|^2. \tag{37}$$

This allows for a randomized algorithm to approximate $\|g\|$: sample s points $w_1, \ldots, w_s$ from $\mathbb{S}(\mathbb{C}^{n+1})$, compute the empirical average $M := \frac{1}{s}\sum |g(w_i)|^2$ and return $\sqrt{\binom{n+d}{d} M}$. The cost of this algorithm is low (it relies on evaluating g) but one pays for it both in terms of precision —we only obtain an approximation of $\|g\|$— and of certainty —we only have a probabilistic guarantee of correctness. More precisely, the following is Theorem 3.3 in [**11**].

Theorem 4.1. *There is a randomized algorithm which, given $f \in \mathbb{C}[X_0, \ldots, X_n]$ as a black-box function, an upper bound d on its degree, a point $z \in \mathbb{C}^{n+1}$, and*

some $\varepsilon > 0$, returns a number $\mathsf{G} \geq 0$ satisfying

$$\gamma_{\mathrm{Frob}}(f,z) \leq \mathsf{G} \leq 192n^2 d\gamma_{\mathrm{Frob}}(f,z)$$

with probability at least $1-\varepsilon$. The computation takes $\mathcal{O}\big(d\log(\frac{d}{\varepsilon})(L(f)+n+\log d)\big)$ operations. □

The algorithms in §3.2, §3.3, and §4.2 are all of *Las Vegas* type. This means that they are randomized algorithms (they draw random numbers during their execution) whose output is guaranteed to be correct: only their running time is random (for each input data). The algorithm in Theorem 4.1 is also a randomized one but of *Monte Carlo* type. Its running time is deterministic but its output is correct only with a given probability.

One can replace in the rigid homotopy the computation of $\gamma_{\mathrm{Frob}}(f,z)$ described at the end of §4.2 by the randomized algorithm in Theorem 4.1. In doing so, however, the rigid homotopy may be affected in two ways. If the γ_{Frob}s computed along the homotopy are overestimated (the upper bound in Theorem 4.1 does not hold) then the corresponding step-lengths are too short and the complexity bounds no longer hold. If, instead, some γ_{Frob} is underestimated (the lower bound in Theorem 4.1 does not hold) then the correctness of the rigid continuation is compromised. The algorithm devised in [**11**], called there BoostBlackBoxSolve, deals with both problems. The next result (Theorem 1.1 in [**11**]) states its properties regarding termination and correctness.

Theorem 4.2. *Let $F = (f_1, \ldots, f_n)$ be a homogeneous polynomial system with degrees at most D in $n+1$ variables having only regular zeros. On input F, given as a black-box evaluation program, and $\varepsilon > 0$, Algorithm* BoostBlackBoxSolve *terminates almost surely and computes a point $z \in \mathbb{P}^n$, which is an approximate zero of F with probability at least $1-\varepsilon$.* □

Because of its reliance on the Monte-Carlo computation of the γ_{Frob}s, Algorithm BoostBlackBoxSolve is itself of Monte Carlo type: the output is only correct with given probability. Nonetheless, this probability is bounded below independently of F. The algorithm succeeds with probability at least $1-\varepsilon$ for any input F with regular zeros (of course, it may not terminate if this regularity hypothesis is not satisfied as the homotopy path may lead to a singular zero).

We are next interested on the average cost of BoostBlackBoxSolve over random data of the form $\mathbf{u}F$ for a fixed F and random $\mathbf{u} \in \mathcal{U}$. As F itself is fixed, we should expect this cost to depend with the geometry of F. We can express this dependence in terms of γ_{Frob}. We define, for $f \in H_d$,

$$\Gamma(f)^2 := \mathop{\mathbb{E}}_{z\in Z(f)} \gamma_{\mathrm{Frob}}(f,z)^2$$

where the distribution on $Z(f) \subset \mathbb{P}(\mathbb{C}^{n+1})$ is the uniform, and, for $F \in \mathcal{H}_{\mathbf{d}}$,

$$\Gamma(F) := \big(\Gamma(f_1)^2 + \cdots + \Gamma(f_n)^2\big)^{1/2}.$$

This quantity is a measure of regularity (or conditioning) of the set of hypersurfaces $\{Z(f_1), \ldots, Z(f_n)\}$. If all these hypersurfaces have only algebraically regular points then $\Gamma(F) < \infty$. But the converse is not true. A necessary and sufficient condition for the finiteness of $\Gamma(F)$ is yet to be found.

The complexity of BoostBlackBoxSolve, is bounded in the next result [**11**, Thm. 1.2].

Theorem 4.3. *Let $F = (f_1, \dots, f_n)$ be a homogeneous polynomial system with degrees at most D in $n+1$ variables. Assume that, for all $i \leq n$, f_i is square-free. Let $\mathbf{u} \in \mathcal{U}$ be uniformly distributed. Then, on input $\mathbf{u}F$, given as a black-box evaluation program, and $\varepsilon \in [0, \frac{1}{4}]$, Algorithm* BoostBlackBoxSolve *terminates after*

$$(n, D)^{\mathcal{O}(1)} \cdot L(F) \cdot \left(\Gamma(F) \log \Gamma(F) + \log \log \varepsilon^{-1}\right)$$

operations on average. "On average" refers to expectation over both the random draws made by the algorithm and the random variable $\mathbf{u}$, but F is fixed. □

The hypothesis on F (the fact that all the $Z(f_i)$ are regular) implies that each f_i is square-free. Hence, by Theorem 4.2, ensures that the algorithm terminates almost surely. It further implies that the quantity $\Gamma(F)$ is finite and, therefore, so is the bound in Theorem 4.3. We note that this bound is polynomial in n and D, linear in $L(F)$, almost linear in $\Gamma(F)$, and linear in $\log \log \varepsilon^{-1}$. The latter means that we can take $\varepsilon = 2^{-2^{100}}$ without afterthoughts as to the effect on the running's cost. For all practical purposes the algorithm behaves as one with certified outputs.

4.5. Algebraic branching programs. Theorem 4.3 can be extended to subsets $\mathscr{M} \subset \mathcal{H}_{\mathbf{d}}$ other than $\mathcal{U}F$ (for a given $F \in \mathcal{H}_{\mathbf{d}}$). The crucial requirement on $\mathscr{M}$ is that it is unitarily invariant in the following sense: $\mathscr{M}$ is endowed with a probability distribution such that for all $F \in \mathscr{M}$ and all $\mathbf{u} \in \mathcal{U}$, $\mathbf{u}F \in \mathscr{M}$, and F and $\mathbf{u}F$ are identically distributed. Trivially, the set $\mathcal{U}F$ is unitarily invariant.

Final average complexity bounds will now depend not on $\Gamma(F)$ for a particular F but on the quantity $\mathbb{E}_{f \in \mathscr{M}} \Gamma(F)^2$. The second main result in [**11**] estimates this quantity for a family of well-known subsets $\mathscr{M}$ given by SLPs known as Algebraic Branching Programs (ABP) introduced by Nissan [**20**].

To understand ABPs let us consider a simple example. Consider the matrices

$$A_1 = \begin{bmatrix} 6X_0 + X_1 & -3X_0 + 2X_1 & 5X_0 \\ 4X_0 - 3X_1 & 2X_0 + 7X_1 & 6X_0 - 5X_1 \end{bmatrix}$$

and

$$A_2 = \begin{bmatrix} 3X_0 + 2X_1 & -X_0 + 4X_1 \\ X_0 - X_1 & 2X_0 - 5X_1 \\ 2X_0 - X_1 & -3X_0 - X_1 \end{bmatrix}$$

whose entries are linear forms in (X_0, X_1). We can take the product

$$A_1 \cdot A_2 = \begin{bmatrix} 25X_0^2 + 15X_0X_1 & -27X_0^2 + 37X_0X_1 - 6X_1^2 \\ 26X_0^2 - 12X_0X_1 - 8X_1^2 & -18X_0^2 + 32X_0X_1 - 42X_1^2 \end{bmatrix}$$

as well as its trace

$$f_{\mathbf{A}}(X) = 7X_0^2 + 47X_0X_1 - 42X_1^2.$$

The coefficients of this polynomial in H_2 are simple functions of the entries in A_1 and A_2. In this example, the former is an element in $\mathbb{C}^3$ and the latter an element in $\mathbb{C}^{24}$. Clearly, the map $(A_1, A_2) \mapsto f_{\mathbf{A}}$ is surjective. But this needs not to be so.

Let $\mathbf{r} = (r_1, \dots, r_d) \in \mathbb{N}^d$, $r_0 = r_d$, and $M_{\mathbf{r}}(n+1)$ be the space of d-tuples of matrices $\mathbf{A} = (A_1(X), \dots, A_d(X))$ where $A_j(X)$ is a $r_{j-1} \times r_j$ matrix whose entries

are linear forms on $(X_0, \dots, X_n)$. This is a complex linear space of dimension $(n+1)\sum_{j=1}^{d} r_{j-1} r_j$.

For $\mathbf{A} \in M_{\mathbf{r}}(n+1)$ we consider the homogeneous polynomial of degree d

$$f_{\mathbf{A}}(X) := \operatorname{tr}(A_1(X) \cdot \ldots \cdot A_d(X)).$$

For large enough $r_1, \dots, r_d$ the map $\mathbf{A} \mapsto f_{\mathbf{A}}$ is a surjection $M_{\mathbf{r}}(n+1) \to H_d$, but for smaller values $\dim_{\mathbb{C}} M_{\mathbf{r}}(n+1) < \dim_{\mathbb{C}} H_d$. Tuples of polynomials described by ABPs give us a natural example of an $\mathscr{M} \subset \mathcal{H}_{\mathbf{d}}$ for which we can actually efficiently run the rigid homotopy.

To show this, we endow $M_{\mathbf{r}}(n+1)$ with the Hermitian norm

$$\|\mathbf{A}\|^2 := \sum_{j=1}^{d} \sum_{i=0}^{n} \|A_j(e_i)\|^2_{\mathrm{Frob}}$$

where e_i is the ith unit vector in $\mathbb{C}^{n+1}$. With this norm at hand we can define the standard Gaussian distribution on $M_{\mathbf{r}}(n+1)$, via the density $\pi^{\dim_{\mathbb{C}} M_{\mathbf{r}}(n+1)} e^{-\|\mathbf{A}\|^2}$. We say that an ABP is *irreducible* when $r_1, \dots, r_{d-1} \geq 2$. The second main result in [**11**], Theorem 1.5, is then the bound for irreducible random ABPs

$$\mathop{\mathbb{E}}_{\mathbf{A} \sim M_{\mathbf{r}}(n+1)} \Gamma(f_{\mathbf{A}})^2 \leq \frac{3}{4} d^3 (d+n) \log d. \tag{38}$$

Because the distribution of a Gaussian ABP is unitarily invariant one can derive from Theorems 4.2 and 4.3 the following [**11**, Corollary 1.6].

Corollary 4.4. *Let $f_1, \dots, f_n$ be given by independent irreducible Gaussian ABPs of degree at most D and evaluation complexity at most L, and $\varepsilon \in [0, \frac{1}{4}]$. Then* BoostBlackBoxSolve *returns an approximate zero of $F = (f_1, \dots, f_n)$ with probability at least $1 - \varepsilon$ in*

$$(n, D)^{\mathcal{O}(1)} L \log\log \varepsilon^{-1}$$

operation on average. □

It is remarkable that the values r_j do not occur on the right-hand side of (38). They do occur, however in the bound in Corollary 4.4 as, it is easy to see using iterated matrix multiplication, $L = \mathcal{O}(nDr^3)$ where $r = \max r_j$.

References

[1] N.H. Abel. Mémoire sur les équations algébriques, ou l'on démontre l'impossibilité de la résolution de l'équation générale du cinquième degré. In L. Sylow and S. Lie, editors, *Œuvres Complètes de Niels Henrik Abel*, volume I, pages 28–33. Grøndahl & Søn, 1881.

[2] D. Armentano, C. Beltrán, P. Bürgisser, F. Cucker, and M. Shub. Condition length and complexity for the solution of polynomial systems. *Found. Comput. Math.*, 16:1401–1422, 2016.

[3] D. Armentano, C. Beltrán, P. Bürgisser, F. Cucker, and M. Shub. A stable, polynomial-time algorithm for the eigenpair problem. *J. Europ. Math. Soc.*, 20:1375–1437, 2018.

[4] V. Arnold, M. Atiyah, P. Lax, and B. Mazur, editors. *Mathematics: Frontiers and Perspectives*. AMS, 2000.

[5] W. Baur and V. Strassen. The complexity of partial derivatives. *Theoret. Comput. Sci.*, 22(3):317–330, 1983.
[6] C. Beltrán, U. Etayo, J. Marzo, and J. Ortega-Cerdá. A sequence of polynomials with optimal condition number. *J. Amer. Math. Soc.*, 34:219–244, 2021.
[7] C. Beltrán and L.M. Pardo. On Smale's 17 problem: a probabilistic positive solution. *Found. Comput. Math.*, 8:1–43, 2008.
[8] C. Beltrán and L.M. Pardo. Fast linear homotopy to find approximate zeros of polynomial systems. *Found. Comput. Math.*, 11(1):95–129, 2011.
[9] P. Bürgisser and F. Cucker. On a problem posed by Steve Smale. *Annals of Mathematics*, 174:1785–1836, 2011.
[10] P. Bürgisser and F. Cucker. *Condition*, volume 349 of *Grundlehren der mathematischen Wissenschaften*. Springer-Verlag, Berlin, 2013.
[11] P. Bürgisser, F. Cucker, and P. Lairez. Rigid continuation paths II. Structured polynomial systems. Preprint. Available at `arXiv:2010.10997`, 2021.
[12] J.W. Demmel. *Applied Numerical Linear Algebra*. SIAM, 1997.
[13] J. Friberg. A geometric algorithm with solutions to quadratic equations in a sumerian juridical document from Ur III Umma. *Cuneiform Digital Library Journal*, 3, 2009.
[14] V. Jones. Ten problems. In [**4**], pg. 79–91.
[15] E. Lahaye. Une méthode de resolution d'une categorie d'equations transcendantes. *C. R. Acad. Sci. Paris*, 198:1840–1842, 1934.
[16] P. Lairez. A deterministic algorithm to compute approximate roots of polynomial systems in polynomial average time. *Found. Comput. Math.*, 17(5):1265–1292, 2017.
[17] P. Lairez. Rigid continuation paths I. Quasilinear average complexity for solving polynomial systems. *J. Amer. Math. Soc.*, 33(2):487–526, 2020.
[18] G. Malajovich. Complexity of sparse polynomial solving: homotopy on toric varieties and the condition metric. *Found. Comput. Math.*, 19(1):1–53, 2019.
[19] G. Malajovich. Complexity of sparse polynomial solving 2: renormalization. Preprint. Available at `arXiv:2005.01223`, 2020.
[20] N. Nisan. Lower bounds for non-commutative computation. In *Proceedings of the Twenty-Third Annual ACM Symposium on Theory of Computing*, pages 410–418. ACM, 1991.
[21] M. Shub. Complexity of Bézout's Theorem VI geodesics in the condition (number) metric. *Found. Comput. Math.*, 9:171–178, 2009.
[22] M. Shub and S. Smale. Complexity of Bézout's Theorem I: geometric aspects. *J. Amer. Math. Soc.*, 6:459–501, 1993.
[23] M. Shub and S. Smale. Complexity of Bézout's Theorem II: volumes and probabilities. In F. Eyssette and A. Galligo, editors, *Computational Algebraic Geometry*, volume 109 of *Progress in Mathematics*, pages 267–285. Birkhäuser, 1993.
[24] M. Shub and S. Smale. Complexity of Bézout's Theorem III: condition number and packing. *Journal of Complexity*, 9:4–14, 1993.
[25] M. Shub and S. Smale. Complexity of Bézout's Theorem V: polynomial time. *Theor. Comput. Sci.*, 133:141–164, 1994.
[26] M. Shub and S. Smale. Complexity of Bézout's Theorem IV: probability of success; extensions. *SIAM J. of Numer. Anal.*, 33:128–148, 1996.

[27] S. Smale. Mathematical problems for the next century. In [**4**], pg. 271–294.
[28] S. Smale. Newton's method estimates from data at one point. In R. Ewing, K. Gross, and C. Martin, editors, *The Merging of Disciplines: New Directions in Pure, Applied, and Computational Mathematics*. Springer-Verlag, 1986.
[29] S. Smale. Mathematical problems for the next century. *Mathematical Intelligencer*, 20:7–15, 1998.

Dept. of Mathematics, City University of Hong Kong, Kowloon Tong, Hong Kong
macucker@cityu.edu.hk