

Amortized Bound for Root Isolation via Sturm Sequences

Zilin Du, Vikram Sharma and Chee K. Yap

Abstract. This paper presents two results on the complexity of root isolation via Sturm sequences. Both results exploit amortization arguments.

For a square-free polynomial $A(X)$ of degree d with L -bit integer coefficients, we use an amortization argument to show that all the roots, real or complex, can be isolated using at most $O(dL + d \lg d)$ Sturm probes. This extends Davenport's result for the case of isolating all real roots.

We also show that a relatively straightforward algorithm, based on the classical subresultant PQS, allows us to evaluate the Sturm sequence of $A(X)$ at rational $\tilde{O}(dL)$ -bit values in time $\tilde{O}(d^3 L)$; here the \tilde{O} -notation means we ignore logarithmic factors. Again, an amortization argument is used. We provide a family of examples to show that such amortization is necessary.

Mathematics Subject Classification (2000). Primary 68Q25; Secondary 34A34; Tertiary 26C10.

Keywords. Sturm sequence, Davenport-Mahler bound, subresultant, complexity, root isolation, separation bound.

1. Introduction

Sturm sequences are a classical tool for real root isolation [CL83]. We recall the main steps of the standard real root isolation algorithm based upon Sturm sequences [CL83]: Let $A(X)$ be a square-free integer polynomial of degree d with L -bit coefficients.

- (1) Compute the Sturm sequence of $A(X)$.
- (2) Compute an interval $(-B, B)$ containing all real zeros of $A(X)$. Initialize a queue Q with the interval $(-B, B)$.

The work is supported by NSF Grant #CCF-043836. A preliminary version of this work appeared at the International Workshop on Symbolic-Numeric Computation (SNC 2005), Xi'an, China, July 19–21, 2005.

- (3) While Q is non-empty, do: extract an interval I from Q and compute the number of zeros in I using the Sturm sequence. If I has one zero, output I . If I has no zeros, discard I . Otherwise, split I into two open intervals I_L, I_R , at the midpoint $m(I)$ of I . Check if $m(I)$ is a zero. If so, output $m(I)$. Push I_L and I_R into Q .

The original bound for this algorithm is $\tilde{O}(d^7L^3)$ from Collins-Loos [CL83]. Davenport [Dav85, Prop. 3'] stated a complexity bound of $\tilde{O}(d^4L^2)$. In this paper, the \tilde{O} -notation means that we are ignoring logarithmic factors. These complexity bounds are estimated from the following three bounds:

- (I) The complexity of computing the Sturm sequence.
- (II) The complexity of evaluating a Sturm sequence at a given point.
- (III) Bound on the number of bisections needed to isolate all the roots .

The complexity of (I) is $\tilde{O}(d^4L^2)$ in [CL83], but this has been improved to $\tilde{O}(d^2L)$ in [LR01, Rei97]. The best complexity bound for (II) is $\tilde{O}(d^3L)$ assuming (as we may in root isolation) that the evaluation point is a rational number with bit size $\tilde{O}(dL)$. Although Davenport [Dav85] stated this bound, the first published proofs was given by Reischert [Rei97], and independently, by Lickteig-Roy [LR01]. The bound for (III) is $\tilde{O}(d^2L)$ in Collins-Loos [CL83]. Davenport improved the bound in (III) by a factor of d to $O(dL + d \lg d)$. The overall complexity of real root isolation using Sturm sequences is the product of the bounds in (II) and (III). Thus the best current bound for real root isolation via Sturm sequences is $\tilde{O}(d^4L^2)$.

Our Contribution. In this paper we give amortization arguments which achieve the above bounds for (II) and (III). Both results use amortization analysis, a technique that is common in discrete algorithms [CLRS01, Chap. 17].

For (II), we give an approach that is much simpler than Reischert or Lickteig-Roy. We rely only on the standard theory of Subresultant polynomial remainder sequences (PRS). But instead of the PRS, we use another idea that goes back to Strassen [Str83], which represents the PRS by its polynomial quotient sequence (PQS). We then give an amortized argument for the straightforward evaluation of this PQS. We also give a family of examples to show that a non-amortized worst case bound will not do.

For (III), we give a charging scheme argument that leads to a slightly sharper bound than that of Davenport. But the main benefit of our argument is its extendibility to the case of isolating complex roots of a polynomial; it is not obvious how to extend Davenport's argument to this case. In particular, we show that the number of Sturm sequence evaluations are $O(dL + d \lg d)$ even for the case of isolating complex roots of $A(X)$. We think our argument gives some insights into how the distances between various roots actually affect the complexity of Step (III). There is a key difference between Sturm's method applied to isolating complex roots, for instance in [Pin76, Wil78], as opposed to the case of real roots: in the complex case, one has to re-compute a Sturm sequence at each probe. This

drawback can be overcome by applying the two-dimensional Sturm sequences of Hermite (cf. Pedersen [Ped90], or an alternative by Milne [Mil92]).

The Complexity of Root Isolation. The complexity results for root isolation via Sturm sequences is inferior to the $\tilde{O}(d^3L)$ bound obtained by Schönhage [Sch82]. Nevertheless, there are some advantages in the Sturm approach: Schönhage’s algorithm simultaneously approximates all the roots of a polynomial, but these approximations may not represent isolations until the root separation bound is achieved. In contrast, the Sturm approach can isolate any subset of roots in a suitable region, or the i ’th largest real root for any chosen i or range of i ’s. The Sturm method is ideally suited for root isolation, a problem that is distinct from root approximation.

There are many other results [Ren87, KS94, NR96, Pan96] on the complexity of root approximation that do not directly depend on root isolation. These methods, like Schönhage’s, simultaneously approximate all the complex roots of a polynomial. For instance, the bit complexity for the Neff-Pan algorithm is $\tilde{O}(d^3L + d\mu)$ where μ is the desired relative precision in each complex zero. If we choose $\mu = \tilde{O}(dL)$ which is the root separation bound, we are assured of isolating all the roots. In any case, these bounds do not improve Schönhage’s bound. In comparing complexity bounds, we must take in account normalization assumptions. E.g., Pan [Pan96] normalized the polynomials so that all its zeros lie in the unit circle. This transforms a polynomial $A(X)$ with L -bit coefficients into a normalized polynomial with dL -bit coefficients.

2. Efficient Evaluation of Sturm Sequences: Simplified Approach

In this section, we address the complexity of Step (II) in the introduction. In particular, we must evaluate Sturm sequences at rational values of X with bit sizes proportional to the logarithm of the root separation bound; the latter we know is bounded by $O(d(L + \lg d))$. Also, a rational number has L -bits if its numerator and denominators are at most L -bit integers.

Recall that Reischert [Rei97] and Lickteig-Roy [LR01] have showed the complexity of Step (II) as $\tilde{O}(d^3L)$. However, their approaches are fairly complicated and require specialized algorithms (Reischert uses a generalized form of the half-GCD algorithm and Lickteig-Roy use computation over the rational field \mathbb{Q}). We now show how a fairly straightforward algorithm that achieves the same bounds.

Let $A(X), B(X) \in \mathbb{Z}[X]$ where $d = \deg(A) > \deg(B)$ and the bit lengths of the coefficients of A, B are at most L . Recall the notion [Yap00, p. 83] of a **polynomial remainder sequence** (PRS) of (A, B) **based on a sequence** $(\beta_1, \dots, \beta_{k-1})$ where $\beta_i \in \mathbb{Z}$: this is a sequence

$$(A_0, A_1, \dots, A_k) \tag{2.1}$$

of polynomials such that $A_0 = A$ and $A_1 = B$, and for $i = 1, \dots, k$, there exists $Q_i \in \mathbb{Z}[X]$ such that

$$\beta_i A_{i+1} = a_i^{\delta_i+1} A_{i-1} - Q_i A_i \tag{2.2}$$

where a_i is the leading coefficient of A_i and $\delta_i = \deg(A_{i-1}) - \deg(A_i) = \deg(Q_i) > 0$, with the termination condition that $A_{k+1} = \beta_k = 0$. The key problem in PRS is to devise effective methods for computing the β_i 's so that the bit size of coefficients of the A_i remain polynomial in d and L . In particular, the **subresultant PRS** from Collins [Yap00, p. 89] achieves this with bit sizes of coefficients bounded by $\tilde{O}(dL)$.

Let the ‘‘bit size’’ of $A(X)$ be $\lg H(A(X))$ where $H(A(X))$ is the **height** of $A(X)$, i.e., the maximum of the absolute value of the coefficients of $A(X)$ [Yap00, p. 23]. An alternative representation for the PRS uses the following concept. Let us define the **polynomial quotient sequence** (PQS) of (A, B) **based on a sequence** $(\beta_1, \dots, \beta_{k-1})$ to be a sequence

$$(A_0, A_1, Q_1, Q_2, \dots, Q_{k-1}) \tag{2.3}$$

where the Q_i 's are defined as in (2.2). Note that the number of coefficients in the PRS (2.1) may be $\Omega(d^2)$. Thus if it is used as a Sturm sequence for A_0, A_1 , evaluating this sequence at any value of X may require $\Omega(d^2)$ arithmetic operations. In contrast, if we only store the PQS in (2.3) and also $(\beta_1, \dots, \beta_{k-1})$, we can easily evaluate the Sturm sequence at any X using only $O(d)$ arithmetic operations; the reason is that $\sum_{i=1}^{k-1} \deg(Q_i)$ is only d . This advantage in the number of arithmetic operations has been noted by many authors including [LR01, Rei97]. However, when we consider bit complexity, it is no longer clear whether we still have an advantage by a factor of d .

It is not hard to see from the definition of a PQS that the bit sizes of the Q_i 's are $\tilde{O}(d^2L)$. More precisely:

Lemma 1. *The bit sizes of the coefficients $Q_i(X)$ is bounded by $O(\delta_i dL)$.*

Proof. From (2.2) we know that Q_i is the pseudo-quotient of A_{i-1} divided by A_i . By [Yap00, Lemma 3.8], the coefficients of Q_i is obtained as principal minors of a matrix M of size $\delta_i + 1 \times \deg(A_{i-1}) + 1$. The first row of M contains the coefficients of A_{i-1} and the remaining rows contains shifted coefficients of A_i . Since each entry has $\tilde{O}(dL)$ bits, the minors have the stated bound. \square

The following example shows that the bit sizes of coefficients of the Q_i 's can be $\Omega(d^2L)$. Consider the subresultant PRS for the polynomials

$$A_0(X) = aX^{3d}, \quad A_1(X) = bX^{2d} + c.$$

Then we have $A_2(X) = (-1)^d b^d a c X^d$, $\beta_1 = (-1)^{d+1}$, $Q_1(X) = b^d a X^d$, $A_3(X) = (-1)^d b^{d-1} (ac)^{d+1}$, $\beta_2 = (-1)^d b^{d^2+1}$. $Q_2 = b^{d^2+1} (ac)^d X^d$.

It follows that a naive worst case bound of $\tilde{O}(dL)$ on the coefficients of the PQS is wrong.

Now we show the desired bound on the complexity of evaluating Sturm sequences using PQS.

Theorem 2. *Let $A, B \in \mathbb{Q}[X]$ have degree d , and let its coefficients be L -bit rationals. Let (A_0, A_1, \dots, A_h) be the subresultant PRS of A, B , based on $(\beta_1, \dots, \beta_{h-1})$. Also, let a_i be the leading coefficient of A_i .*

- (i) *We can compute (A_0, \dots, A_h) , $(\beta_1, \dots, \beta_{h-1})$ and also (a_0, \dots, a_h) in time $\tilde{O}(d^3L)$.*
- (ii) *We can evaluate the Subresultant PRS of A, B at any $\tilde{O}(dL)$ -bit rational value in time $\tilde{O}(d^3L)$.*

Proof.

- (i) This is done by a straightforward evaluation of the subresultant PRS [Yap00, Sect. 3.5, p. 89].
- (ii) Let x be an $\tilde{O}(dL)$ -bit rational. To compute $A_i(x)$ for $i = 0, \dots, h$, we use the following steps:
 - Step (a): Evaluate $A_0(x), A_1(x)$.
 - Step (b): Evaluate $Q_1(x), \dots, Q_{h-1}(x)$.
 - Step (c): For $i=1$ to $h-1$, compute $A_{i+1}(x)$ as $(a_i^{\delta_i+1}A_{i-1}(x) - Q_i(x)A_i(x))/\beta_i$.

All the polynomial evaluations must be done using Horner's rule in order for our bounds to be valid. Step (a) is $\tilde{O}(d^2L)$. For Step (b), the complexity of evaluating Q_i at x is $\tilde{O}(\delta_i^2 dL)$ (by Lemma 1). Summing over all i 's, we obtain an overall complexity of $\sum_{i=1}^{h-1} \tilde{O}(\delta_i^2 dL) = \tilde{O}(d^3L)$, since $\sum_{i=1}^{h-1} \delta_i \leq d$. Finally, for step (c), we note that each of the quantities $a_i^{\delta_i+1}$, $Q_i(x)$, $A_i(x)$, and hence β_i , is $\tilde{O}(d^2L)$ -bit rationals. Thus, the cost of computing each $A_i(x)$ is $\tilde{O}(d^2L)$ and so the overall cost of step (c) is $\tilde{O}(d^3L)$. \square

The above proof amounts to a simple algorithm for achieving $\tilde{O}(d^3L)$ bit complexity for Step (II). The amortization argument amounts to exploiting the inequality $\sum_i \delta_i \leq d$.

3. The Davenport-Mahler Bound

The basic inequality that our amortized analysis will exploit is the Davenport-Mahler theorem (Theorem 3). This theorem gives a lower bound on the product of differences of certain pairs of roots of a polynomial $A(X) = a_d \prod_{i=1}^d (X - \alpha_i)$ in terms of its **discriminant** $\text{disc}(A) = a_d^{2d-2} \prod_{1 \leq i < j \leq d} (\alpha_i - \alpha_j)^2$ and **Mahler measure** $M(A) = |a_d| \prod_{i=1}^d \max\{1, |\alpha_i|\}$, see [Yap00, 6.6, 4.5] [Mc99, 1.5, 2.1]. The literature has several variants of this theorem that use the same proof but formulate different conditions on how roots may be paired so that the proof works. We give the most general condition supported by the proof. It is equivalent to Johnson's formulation [Joh98] and generalizes Davenport's original formulation [Dav85, Prop. I.5.8].

Theorem 3. *Let $A(X) = a_d \prod_{i=1}^d (X - \alpha_i)$ be a square-free complex polynomial of degree d . Let $G = (V, E)$ be a directed graph whose nodes $\{v_1, \dots, v_k\}$ are a subset of the roots of $A(X)$ such that:*

1. If $(\alpha, \beta) \in E$ then $|\alpha| \leq |\beta|$;
2. G is acyclic;
3. The in-degree of any node is at most 1.

If exactly m of the nodes have in-degree 1, then

$$\prod_{(v_i, v_j) \in E} |v_i - v_j| \geq \sqrt{|\text{disc}(A)|} \cdot M(A)^{-(d-1)} \cdot (d/\sqrt{3})^{-m} \cdot d^{-d/2}. \quad (3.1)$$

Proof. See [ESY06, Theorem 3.1]. □

Remark 3.1. Suppose the edge set of a graph $G = (V, E)$, $V \subseteq \{\alpha_1, \dots, \alpha_d\}$, can be partitioned into k disjoint edge sets $E = E_1 \cup \dots \cup E_k$ such that each of the graphs $G_1 = (V, E_1), \dots, G_k = (V, E_k)$ satisfies the properties in the theorem above, then $\prod_{(u, v) \in E} |u - v|$ is bounded from below by the product of the bounds corresponding to each G_i .

The following lemma gives us an upper and lower bound on the product of k intervals defined by the real roots of a polynomial $A(X)$. For any polynomial $A(X)$, let $\text{lead}(A)$ be its leading coefficient.

Lemma 4. *Let $A(X) \in \mathbb{R}[X]$ be a square-free polynomial of degree d . If $\alpha_1 < \beta_1 \leq \alpha_2 < \beta_2 \leq \alpha_3 < \dots < \beta_k$ are real zeros of $A(X)$ then*

$$\prod_{i=1}^k |\alpha_i - \beta_i| \begin{cases} \leq & M(A)/\text{lead}(A), \\ \geq & M(A)^{-d+1} d^{-d/2} (\sqrt{3}/d)^k. \end{cases} \quad (3.2)$$

Proof. Let $\gamma_1, \dots, \gamma_d$ be all the (not necessarily distinct) zeros of $A(X)$. First we prove that $M(A)$ is an upper bound on the product $\prod_{i=1}^k |\alpha_i - \beta_i|$. We consider two possibilities.

Case A: Suppose there exists an $h = 1, \dots, k$ such that $\alpha_h < 0 < \beta_h$. Let $\gamma_1, \dots, \gamma_d$ denote all the distinct roots in the set $\{\alpha_i, \beta_i : i = 1, \dots, k\}$. Thus, $k+1 \leq d \leq 2k$. We have

$$\begin{aligned} \prod_{i=1}^k |\beta_i - \alpha_i| &= \left(\prod_{i=1}^{h-1} (\beta_i - \alpha_i) \right) \cdot |\alpha_h - \beta_h| \cdot \left(\prod_{i=h+1}^k (\beta_i - \alpha_i) \right) \\ &\leq \left(\prod_{i=1}^{h-1} |\alpha_i| \right) \cdot (|\alpha_h| + |\beta_h|) \cdot \left(\prod_{i=h+1}^k |\beta_i| \right) \\ &\leq \left(\prod_{i=1}^d \max\{1, |\gamma_i|\} \right) \end{aligned}$$

which, in turn, is bounded by $M(A)/|\text{lead}(A)|$.

Case B: Suppose $\beta_i \leq 0 \leq \alpha_{i+1}$ for some $i = 0, 1, \dots, k+1$ (with $\beta_0 = -\infty, \alpha_{k+1} = \infty$). The above argument can easily be adapted to this case as well.

The lower bound follows from [Dav85, Prop. 8]. □

4. Amortized Bound on Number of Probes to Isolate Real Roots

Suppose (a, b) is an open interval containing k roots,

$$a < \alpha_1 < \alpha_2 < \dots < \alpha_k < b. \tag{4.1}$$

The values of these roots are unknown, and our goal is to “locate” them by isolating intervals. We will bound the size of the binary search tree of the algorithm described in introduction in terms of the amortized bound given in (3.2). The bounds expressed in (3.1) and (3.2) are amortized because they are better than the worst case bound obtained by taking the product of the worst case for each gap, i.e., the root separation bound. We next formalize our problem and give a general framework of an algorithm that encompasses the one based upon Sturm sequences.

All our intervals are either open intervals $I = (c, d)$ or exact intervals $I = [c, c]$ for some $a \leq c < d \leq b$. The **width** of I is $w(I) = d - c$ for open intervals and $w(I) = 0$ for exact intervals. Also, let $\#(I)$ denote the number of roots in I . If $I = [c, c]$, we write $\#(c)$ instead of $\#([c, c])$. Clearly $\#(c) = 0$ or 1 . We call I an **isolating interval** if $\#(I) = 1$.

The **Real Root Isolation Problem** for an interval $I = (a, b)$ is that of finding a set of $\#(I)$ pairwise disjoint isolating intervals containing the real roots in I . To solve this problem, we consider algorithm that make “probes”. Each **probe** is defined by an input open interval I and the **result** of the probe is the pair $(\#(I_L), \#(I_R))$ where $m = m(I) = (c + d)/2$, $I_L = (c, m)$ and $I_R = (m, d)$. Note that $\#(m) = \#(I) - \#(I_L) - \#(I_R)$.

In the following, let $\tau > 0$ be an arbitrary **threshold parameter**. For instance, we may choose $\tau = 1$. Relative to τ , we define an interval I as **small** or **big** depending on whether $w(I) < \tau$ or $w(I) \geq \tau$.

Let I have roots as in (4.1). A **segment** of I is an interval of the form $\sigma_i = (\alpha_i, \alpha_{i+1})$ for $i = 0, 1, \dots, k$, with $a = \alpha_0$ and $b = \alpha_{k+1}$. Let $\bar{\sigma}(I) = \{\sigma_0, \sigma_1, \dots, \sigma_k\}$ denote the set of segments of I . Call σ_0 and σ_k the **outer segments**; all others are **inner segments**. Define $\Delta(I) = \Delta^+ + \Delta^-$ where

$$\begin{aligned} \Delta^+ = \Delta^+(I) &:= \sum_{\sigma_i \text{ is big}} 2 \lg(w(\sigma_i)/\tau), \\ \Delta^- = \Delta^-(I) &:= \sum_{\substack{\sigma_i \text{ is a small} \\ \text{inner segment}}} (1 - \lg(w(\sigma_i)/\tau)). \end{aligned} \tag{4.2}$$

Thus, Δ^+, Δ^- are the summations over big and small segments (respectively). By definition, Δ^+ (resp., Δ^-) is equal to 0 if there are no big (resp., small) segments.

Theorem 5. *Given an interval $I = (a, b)$ and also $k = \#(I)$, we can solve the real root isolation problem for I using at most $k - 1 + \Delta(I)$ probes.*

Proof. We first present the algorithm. This is a standard binary search: if $\#(I) \leq 1$, the problem is trivial. Otherwise, we initialize a queue to contain just the interval

$I = (a, b)$. In the general step, we extract an interval $J = (c, d)$ from the queue and probe J . The probe returns the pair $(\#(J_L), \#(J_R))$ as above. Inductively, we may assume that $\#(J)$ is known. If $\#(J_L) + \#(J_R) < \#(J)$, we can output $m(J)$ as a root. For each $i = L, R$, we have three possibilities: if $\#(J_i) = 0$, we discard J_i ; if $\#(J_i) = 1$, we output J_i ; if $\#(J_i) > 1$, we put J_i into the queue. The algorithm halts when the queue is empty. It is clear that the output is a complete list of isolating intervals.

We now prove that this algorithm halts after at most $k - 1 + \Delta(I)$ probes using an amortization argument. Let J be an interval that was probed by the algorithm. Probe J produces the subintervals J_L, J_R . We call J a **splitting** probe if $m(J)$ is a root, or if both J_L and J_R are placed in the queue; otherwise J is **non-splitting**. Clearly, there are at most $k - 1$ splitting probes. It remains to prove that at most $\Delta(I)$ probes are non-splitting. This is done by introducing a charging scheme for such probes.

Let $J = (c, d)$ be a non-splitting probe. Then $m(J) = (c + d)/2$ belongs to an outer segment $\sigma \in \bar{\sigma}(J)$. We shall charge probe J to a segment $\sigma' \in \bar{\sigma}(I)$, defined as follows:

$$\sigma' = \begin{cases} \text{the segment in } \bar{\sigma}(I) \text{ that contains } \sigma & \text{if } w(\sigma) \geq \tau, \\ \text{the inner segment in } \bar{\sigma}(J) \text{ adjacent to } \sigma & \text{if } w(\sigma) < \tau. \end{cases}$$

σ' has the following properties:

- It is uniquely defined.
- It is a segment in $\bar{\sigma}(I)$, even when $w(\sigma) < \tau$.
- It is big if σ is big, which is clear.
- It is small if σ is small: to see this, note that $w(\sigma') \subset w(J)$ and also $w(J) = 2w(\sigma) < 2\tau$. Note that σ' is an inner segment.

We now consider two cases:

Case (A): σ' is small. We show that σ' is charged at most $1 - \lg(w(\sigma')/\tau)$ times. Let J_1, \dots, J_ℓ be the probe (intervals) that are charged to σ' , and we may assume $\sigma' \subset J_1 \subset J_2 \subset \dots \subset J_\ell$. Thus

$$w(\sigma') < w(J_1) \leq 2^{-1}w(J_2) \leq \dots \leq 2^{-\ell+1}w(J_\ell) < 2^{-\ell+1}\tau.$$

Hence $\ell < 1 - \lg(w(\sigma')/\tau)$.

Case (B): σ' is big. We show that σ' is charged at most $2\lg(w(\sigma')/\tau)$ times. At the first probe J whereby $m(J) \in \sigma'$, the segment σ' is split into two halves σ'_L and σ'_R . By symmetry, consider σ'_L . Subsequently, suppose J_1, \dots, J_ℓ are all the probe (intervals) that are charged to σ' via σ'_L . More precisely, assume $\sigma'_L \subset J_1$ and $J_\ell \subset J_{\ell-1} \subset \dots \subset J_1$. Then we have

$$w(\sigma'_L) = w(J_1 \cap \sigma') \geq 2w(J_2 \cap \sigma') \geq \dots \geq 2^{\ell-1}w(J_\ell \cap \sigma') \geq 2^{\ell-1}\tau.$$

This proves that $\ell \leq \lg(2w(\sigma'_L)/\tau)$. By also taking into account the charges via σ'_R , the total charges on σ' is at most $\lg(4w(\sigma'_L)w(\sigma'_R)/\tau^2) \leq 2\lg(w(\sigma')/\tau)$, using the fact that $w(\sigma'_L) + w(\sigma'_R) = w(\sigma')$. \square

Implementing the probe model. We discuss how our probe model can be implemented. If $J = (c, d)$ and $\#(J)$ is known, then the probe amounts to performing a “Sturm query” for the interval $[c, m(J)]$ and returns $(\#(J_L), \#(J_R))$. This means that we compute the Sturm sequence for the polynomial $A(X)$ and evaluate the number of sign variations at c and at $m(J)$, and take their difference. Let $V(c)$ denote the number of sign variations of the Sturm sequence evaluated at c . Indeed, if we assume that inductively, we already know $V(c)$ and $V(d)$, then this probe can be done by just computing $V(m(J))$.

Normally, the sign variation $V(x)$ is assumed to be well-defined only when x is not a root of $A(X)$. However, in case $A(X)$ is square-free, $V(x)$ is well-defined even when x is a root. To see this, let $(A_0(X), A_1(X), \dots, A_h(X))$ be the Sturm sequence with $A_0(X) = A(X)$, $A_1(X) = dA/dX$. Then $V(x)$ is the number of sign variations in the sequence of numbers $(A_0(x), A_1(x), \dots, A_h(x))$. But even when $A_0(x) = 0$, the sequence cannot vanish identically because $A(X)$ is square-free. Moreover, since $\text{sign}(A_0(x^+)) = \text{sign}(A_1(x))$, we have

$$\text{sign}(A_0(x^+), A_1(x^+), \dots, A_h(x^+)) = \text{sign}(A_1(x), A_1(x), \dots, A_h(x))$$

and hence $V(x^+) = V(x)$. Similarly, $V(x^-) = 1 + V(x^+)$. Hence $\#(J) = V(c^+) - V(d^-)$ where $J = (c, d)$. Thus our probe model can be implemented with a single sign-variation computation.

5. Complexity of Real Root Isolation

Using the bounds from the previous section we derive an a priori bound on the number of Sturm probes as given in Theorem 5.

Theorem 6. *Let $A(X) \in \mathbb{R}[X]$ be a square-free polynomial of degree d . Then we can isolate all the real roots of $A(X)$ using at most*

$$1.5d \lg d + (d + 1) \lg M(A) + 2d + \frac{1}{2} \lg \frac{1}{|\text{disc}(A)|} + \lg \frac{M(A)}{|\text{lead}(A)|}$$

probes.

Proof. Without loss of generality, we can assume that $w(I)$, where $I = (-B, B)$ is the initial interval, is bounded by $2M(A)/|\text{lead}(A)|$. In this proof, we further assume $\tau = 1$. Let $\alpha_1 < \dots < \alpha_k$ be the k real zeros of $A(X)$. The segments $\bar{\sigma}(I)$ are defined by these α 's and also the two outer segments $(-B, \alpha_1)$ and (α_k, B) . If any of these outer segments has width $< \tau$, then they are never charged. It is also not hard to see that the total combined charges to these 2 outer segments is at most $\lg(M(A)/|\text{lead}(A)|) + 1$, and they are counted as part of Δ^+ .

We first invoke the upper bound in Lemma 4: choose the intervals (α_i, β_i) (for $i = 1, \dots, m$) in this Lemma to correspond to the big segments in $\bar{\sigma}(I)$, not counting any big outer segments; the upper bound in Lemma 4 then implies

$$\prod_{i=1}^m |\alpha_i - \beta_i| \leq M(A)/\text{lead}(A).$$

But the definition of Δ^+ includes any big outer segments and from the preceding remarks we know that the outer segments are charged at most $1 + \lg \frac{M(A)}{|\text{lead}(A)|}$. Hence we conclude that $\Delta^+ \leq (\lg \frac{M(A)}{|\text{lead}(A)|} + 1) + 2 \lg M(A)$.

We next invoke the lower bound in Lemma 4 by choosing the intervals (α_i, β_i) (for $i = 1, \dots, m$), $m \leq k$, in this Lemma to correspond to the small segments in $\bar{\sigma}(I)$, not counting any small outer segments. Then the lower bound in Lemma 4 implies $\Delta^- \leq (k - 1) + (d - 1) \lg M(A) + (3d/2) \lg d + \frac{1}{2} \lg \frac{1}{|\text{disc}(A)|}$. Therefore,

$$\begin{aligned} (k - 1) + \Delta &\leq (d - 1) + \Delta^+ + \Delta^- \\ &\leq (d - 1) + \lg \frac{M(A)}{|\text{lead}(A)|} + 2 \lg M(A) + 1 + \\ &\quad (d - 1) + (d - 1) \lg M(A) + 1.5d \lg d + \frac{1}{2} \lg \frac{1}{|\text{disc}(A)|} \\ &< (d + 1) \lg M(A) + 1.5d \lg d + 2d + \lg \frac{M(A)}{|\text{lead}(A)|} + \frac{1}{2} \lg \frac{1}{|\text{disc}(A)|}. \end{aligned}$$

□

Corollary 7. *We can isolate all the real roots of a square-free integer polynomial of degree d and coefficients of bit length L using at most $dL + 2d \lg d + O(d + L)$ Sturm probes.*

Proof. We have the following observations for $A(X)$:

1. From Landau’s inequality (cf. [Yap00, Lem. 4.14(i)]) $M(A) \leq \|A\|_2$; furthermore $\|A\|_2 \leq (d + 1)^{\frac{1}{2}} 2^L$.
2. Since $A(X)$ is an integer polynomial $|\text{lead}(A)| \geq 1$ and as it is square-free $|\text{disc}(A)| \geq 1$.

Applying these observations to the result of the above theorem gives the following upper bound on the number of Sturm probes:

$$1.5d \lg d + (d + 1) \left[\frac{1}{2} \lg(d + 1) + L \right] + 2d + \frac{1}{2}(d + 1) + L.$$

But this is clearly bounded by $dL + 2d \lg d + O(d + L)$. □

Remark 5.1. This result saves a term of dL as compared to [Dav85, Prop. 2].

From Theorem 2 we have the complexity of evaluating the Sturm sequence of $A(X)$ at an input of bit size dL as $\tilde{O}(d^3 L)$.

Hence we have following bit complexity for real root isolation:

Corollary 8. *If $A(X)$ is a degree d square-free integer polynomial with L -bit coefficients, then we can isolate all the real roots of $A(X)$ in time $\tilde{O}(d^4 L^2)$.*

Remark 5.2. Is our choice of $\tau = 1$ optimal? Consider the two extreme choices. If we choose τ to be 2^{L+1} , then all segments are small. Thus,

$$\Delta = \Delta^- \leq (d - 1)(2 + 2L + 0.5 \lg d) + (3d/2) \lg d.$$

This is asymptotically the same as the bound in our theorem, but slightly worse with an additive term of dL . If we choose τ to be the root separation bound, i.e., $\tau = 2^{-Ld}d^{-d}$, then all segments are big. Hence we have

$$\Delta = \Delta^+ \leq 2L + \lg d + 2Ld^2 + 2d^2 \lg d.$$

This has an additive $d^2(L + \lg d)$ term, which is asymptotically worse than our theorem. At any rate, this suggests that τ should be chosen to have a balance between the number of big and small segments.

6. Amortized Bound on Number of Probes to Isolate Complex Roots

We can extend the result of Sect. 4 to the case of isolating any set of roots in a rectangular region S in the complex plane. More precisely, S is a half-open set including its northern and western edges but omitting its southern and eastern edges, i.e., the set $\{(x, y) : a \leq x < b, c < y \leq d\}$.

Let $w(S)$ represent the edge length of S and $\#(S)$ the number of roots in S . We call S an **isolating square** if $\#(S) = 1$. We allow S to be a single point, in which case $w(S) = 0$ and $\#(S) = 1$ or 0.

The **Root Isolation Problem** for a square S is to find a set of $\#(S)$ pairwise disjoint squares containing all the roots in S .

A **probe** in this setting is defined by a half-open square S and the result of the probe is the number of roots in the four smaller squares, S_i ($i = 1, 2, 3, 4$), obtained by the segments joining the mid-points of the edges of S . Note that each of these smaller squares are half-open and they partition S . In this paper we do not discuss the implementation of this probe model, though similar implementations can be found in [Pin76, Wil78, Yap00, Pan97].

A probe S is called a **splitting probe** if either the centre of the square C is a root or $\#(S) \neq \#(S_i)$ for any $i = 1, \dots, 4$; otherwise, we call it a **non-splitting probe**.

Suppose S contains k roots $\alpha_1, \dots, \alpha_k$ in \mathbb{C} . For every root α_i inside S we define the following four pairs:

- $\sigma_{i,NE} = (\alpha_i, \alpha_{NE(i)})$, where $\alpha_{NE(i)}$ is the nearest root or corner of S lying north-east of α_i .
- $\sigma_{i,NW} = (\alpha_i, \alpha_{NW(i)})$, where $\alpha_{NW(i)}$ is the nearest root or corner of S lying north-west of α_i .
- $\sigma_{i,SE} = (\alpha_i, \alpha_{SE(i)})$, where $\alpha_{SE(i)}$ is the nearest root or corner of S lying south-east of α_i .
- $\sigma_{i,SW} = (\alpha_i, \alpha_{SW(i)})$, where $\alpha_{SW(i)}$ is the nearest root or corner of S lying south-west of α_i .

These pairs certainly exist; however, they may not be unique. We view these pairs as directed edges coming out from α_i . We also insist that these edges are never horizontal or vertical. This ensures that the four edges issuing out of each

root are distinct (so our graph is simple). The set of edges in the directed graph so obtained for S will be represented as $\overline{\sigma}(S)$. By an **outer edge** of $\overline{\sigma}(S)$ we mean an edge connecting a root with a vertex of S ; the remaining edges will be called **inner edges** and their set will be denoted by $I(S)$. For each $\sigma' \in \overline{\sigma}(S)$, let $w(\sigma')$ denote the Euclidean length of the edge σ' . We will write “ $\sigma_{i,*}$ ” for any of the four edges coming out of α_i . Figure 1(a) illustrates the graph $\overline{\sigma}(S)$, where the dots represent the roots in S .




FIGURE 1. (a) The directed graph $\overline{\sigma}(S)$. (b) Charging scheme probe S'

Similar to Sect. 4 we define a threshold parameter $\tau > 0$ and define an edge $\sigma_{i,*}$ to be big or small depending on whether $w(\sigma_{i,*}) \geq \tau$ or $w(\sigma_{i,*}) < \tau$. We extend the definitions of (4.2) to this context,

$$\begin{aligned} \Delta^+ = \Delta^+(S) &:= \sum_{\sigma' \in \overline{\sigma}(S): \sigma' \text{ is big}} (1 + \lg(w(\sigma')/\tau)), \\ \Delta^- = \Delta^-(S) &:= \sum_{\sigma' \in \overline{\sigma}(S) \cap I(S): \sigma' \text{ is small}} (1 - \lg(w(\sigma')/\tau)). \end{aligned} \tag{6.1}$$

Let $\Delta(S) := \Delta^+ + \Delta^-$.

Analogous to Theorem 5 we show the following:

Theorem 9. *Given a square S and $k = \#(S)$, we can solve the root isolation problem in at most $1.5k - 1 + \Delta(S)$ probes.*

Proof. The algorithm based upon our probe model performs a four-way search. If $\#(S) \leq 1$ then we are done. Otherwise, initialize a queue containing S . In the general step, we extract a square S' from the queue and probe it. The consequence of the probe is the quadruple $(\#(S_1), \#(S_2), \#(S_3), \#(S_4))$ where S_i 's are as shown

in Figure 1(b). We enqueue S_i if $\#(S_i) > 1$, output it if $\#(S_i) = 1$, and discard it otherwise. The algorithm terminates when the queue is empty.

It is clear that there can be at most $k - 1$ splitting probes. We now devise a charging scheme, analogous to what was done earlier, to account for the non-splitting probes. The analogy between the two charging schemes is the following: In Sect. 4 the charging of big segments was counting the number of intervals whose midpoint is contained within that segment; now the charging counts the number of squares that are sufficiently large and whose edges intersect the segment. The charging of small segments, in the same section, was counting the number of intervals that are sufficiently small and which contain the segment; now we count the number of squares in place of intervals.

Consider a non-splitting probe S' with center C . Since this probe is non-splitting all the roots inside S' must lie in one of S_i , $i = 1, 2, 3, 4$; suppose they lie in the square S_4 south-east of C . Let α_i be a root nearest to C . Thus the edge $\sigma_{i,NW}$ must intersect either the top or the left edge of S' and hence

$$w(\sigma_{i,NW} \cap S') \geq w(S_4) = \frac{1}{2}w(S'). \tag{6.2}$$

Furthermore, since S_4 is not an isolating square there is another root β inside S_4 such that one of the edges from α_i , apart from the north-west one, points to β ; denote this edge by e . Then

$$w(e) \leq \sqrt{2}w(S_4) = w(S')/\sqrt{2}. \tag{6.3}$$

These notations are illustrated in Figure 1(b).

Now we describe our charging scheme:

1. If $w(S') \geq 2\tau$ then we charge the probe S' to the corresponding $\sigma_{i,NW}$. This means $\sigma' := \sigma_{i,NW}$ is big.
2. If $w(S') < \sqrt{2}\tau$ then we charge S' to the edge e . This means $\sigma' := e$ is small and cannot be an outer edge.
3. If $\sqrt{2}\tau \leq w(S') < 2\tau$ then we count an additional charge. However, there are at most $k/2$ such charges, since S' is not an isolating square and the next probe S_4 is such that $w(S_4) < \sqrt{2}\tau$.

Thus we have the following cases:

Case (A): σ' is small. Suppose S_1, \dots, S_ℓ are the probes charged to σ' . Assume $S_\ell \subset S_{\ell-1} \subset \dots \subset S_1$ and that σ' is in S_ℓ . Thus from (6.3) we have

$$\sqrt{2}w(\sigma') \leq w(S_\ell) \leq 2^{-1}w(S_{\ell-1}) \leq \dots \leq 2^{-\ell+1}w(S_1) < 2^{-\ell+1}\sqrt{2}\tau,$$

and hence $\ell < 1 - \lg(w(\sigma')/\tau)$.

Case (B): σ' is big. Let S_1, S_2, \dots, S_ℓ be all the probes charge to σ' , such that $S_\ell \subset S_{\ell-1} \subset \dots \subset S_2 \subset S_1$. Then from (6.2) we have

$$w(\sigma') \geq w(\sigma' \cap S_1) \geq \frac{1}{2}w(S_1) \geq w(S_2) \geq \dots \geq 2^{\ell-2}w(S_\ell) \geq 2^{\ell-1}\tau,$$

and hence $\ell \leq 1 + \lg(w(\sigma')/\tau)$. □

7. Complexity of Complex Root Isolation

To bound the complexity of isolating all the roots of a degree d square-free polynomial $A(X)$ we need to bound the number of probes to achieve this. Let S be a half-open square containing all the roots of A ; from Cauchy’s bound we know that $w(S) \leq 2(1 + \|A\|_\infty)$. According to Theorem 9 we need to bound $\Delta^+(S)$ and $\Delta^-(S)$ for $\tau = 1$.

Lemma 10. *Let $A(X)$ be a degree d square-free integer polynomial with L -bit coefficients. Then we have $\Delta^-(S) = O(dL + d \lg d)$ and $\Delta^+(S) = O(dL)$.*

Proof. Suppose $\alpha_1, \dots, \alpha_d$ are the distinct roots of $A(X)$ in S . We consider the upper bounds for $\Delta^-(S)$ and $\Delta^+(S)$ in turn:

Bound on $\Delta^-(S)$. For a root α_i let $\alpha_{N(i)}$ denote the nearest root to it. Consider the directed graph $G = (V, E)$ obtained from the pairs $(\alpha_i, \alpha_{N(i)})$. It is not hard to see that the cycles in G have the property that all the edges have the same length. Thus we can break any cycle of length greater than two to a set of cycles of length less than or equal to two. Now construct two graphs $G_1 = (V, E_1)$ and $G_2 = (V, E_2)$ from G where G_2 contains one of the edges from each cycle in G and G_1 contains all the remaining edges of G ; thus $E = E_1 \cup E_2$. Both G_1 and G_2 are DAG’s. Moreover, the vertices in G_2 have in-degree at-most one. We re-orient the edge $(\alpha, \beta) \in E_2$ such that $|\alpha| \leq |\beta|$; it is clear that this re-orientation does not affect the in-degree of the vertices. Now we apply Theorem 3 to G_2 to obtain:

$$\prod_{(\alpha, \beta) \in E_2} |\alpha - \beta| \geq \sqrt{|\text{disc}(A)|} \cdot M(A)^{-(d-1)} \cdot (d/\sqrt{3})^{-d} \cdot d^{-d/2}. \tag{7.1}$$

The in-degree of the vertices of G_1 can be at most six since any root of $A(X)$ can have at most six neighboring roots that are all closest to it and to each other. Now we re-orient the edge $(\alpha, \beta) \in E_1$ such that $|\alpha| \leq |\beta|$. As a result of this re-orientation the in-degree of any vertex in G_1 is bounded by three. The reason is that any circle that passes through the centre of the hexagon can contain at most three vertices since the vertices are diametrically opposite each other. Thus from Remark 3.1 we have:

$$\prod_{(\alpha, \beta) \in E_1} |\alpha - \beta| \geq |\text{disc}(A)|^{1.5} \cdot M(A)^{-(3d-3)} \cdot (d/\sqrt{3})^{-3d} \cdot d^{-1.5d}. \tag{7.2}$$

Combining this with (7.1) we have:

$$\prod_{(\alpha, \beta) \in E} |\alpha - \beta| \geq |\text{disc}(A)|^2 \cdot M(A)^{-(4d-4)} \cdot (d/\sqrt{3})^{-4d} \cdot d^{-2d}. \tag{7.3}$$

We can partition the set of roots in V into four different types: $R_i, i = 1, 2, 3, 4$, be the set of roots which have exactly i inner edges where the edge $(\alpha, \beta) \in E$ is oriented such that $|\alpha| \leq |\beta|$. Then an upper bound on $\Delta^-(S)$ translates into a lower bound on $\prod w(\sigma')$, where σ' are small inner edges, but this

product is equal to the product of the small inner edges corresponding to each root in R_i , $i = 1, 2, 3, 4$. This latter product is clearly greater than

$$\prod_{\alpha_i \in R_1} |\alpha_i - \alpha_{N(i)}| \cdot \prod_{\alpha_i \in R_2} |\alpha_i - \alpha_{N(i)}|^2 \cdot \prod_{\alpha_i \in R_3} |\alpha_i - \alpha_{N(i)}|^3 \cdot \prod_{\alpha_i \in R_4} |\alpha_i - \alpha_{N(i)}|^4,$$

since each inner edge is at least the distance to the nearest root. From (7.3) and (6.1), we thus have

$$\Delta^-(S) = O(dL + d \lg d),$$

since $\lg M(A) \leq L + 0.5 \lg(d + 1)$.

Bound on $\Delta^+(S)$. Since S contains all the roots, we know $w(\sigma_{i,*})$, $i = 1, \dots, d$, is less than the diagonal of S , which instead is less than 2^{L+3} . Thus for any root α_i we have

$$w(\sigma_{i,NW})w(\sigma_{i,NE})w(\sigma_{i,SW})w(\sigma_{i,SE}) \leq 2^{4L+12},$$

and from (6.1) we have

$$\begin{aligned} \Delta^+(S) &\leq d + \sum_{\sigma' \in \overline{\sigma}(S): \sigma' \text{ is big}} \lg w(\sigma') \\ &\leq 4dL + 13d, \end{aligned}$$

which clearly is $O(dL)$. □

The bounds from the above lemma combined with Theorem 9 give the following:

Theorem 11. *Let $A(X)$ be a square-free integer polynomial of degree d and L -bit coefficients. Then the number of probes to isolate all the roots of $A(X)$ are at most $O(dL + d \lg d)$.*

Recall the distinction between the Sturm based approach for the case of real roots and complex roots. Since in the latter one has to compute a new Sturm sequence at each probe, this can increase the bit size of the coefficients of these polynomials by one, so that the final probe can possibly involve polynomials having bit size of the coefficients $\tilde{O}(dL)$. The evaluation of the Sturm sequence corresponding to these polynomials can take $\tilde{O}(d^4L^2)$, and hence the total complexity to isolate all roots of a degree d square-free integer polynomial with L -bit coefficients can potentially be $\tilde{O}(d^5L^3)$. Hence Theorem 11 seems to be of theoretical interest only.

8. Conclusion

The two contributions of this paper are (1) a simplified approach for achieving the best known complexity bound in evaluating Sturm sequences, and (2) a new probe complexity bound for isolating complex roots. The common theme in these two results is the use of amortization arguments.

References

- [CL83] G.E. Collins and R. Loos. Real zeros of polynomials. In B. Buchberger, G. E. Collins and R. Loos, editors, *Computer Algebra*, pages 83–94. Springer-Verlag, 2nd edition, 1983.
- [CLRS01] T.H. Corman, C.E. Leiserson, R.L. Rivest and C. Stein. *Introduction to Algorithms*. The MIT Press and McGraw-Hill Book Company, Cambridge, Massachusetts and New York, second edition, 2001.
- [Dav85] J.H. Davenport. Computer algebra for cylindrical algebraic decomposition. Technical report, The Royal Institute of Technology, Dept. of Numerical Analysis and Computing Science, S-100 44, Stockholm, Sweden, 1985. Reprinted as: Tech. Report 88-10, School of Mathematical Sciences, Univ. of Bath, Claverton Down, bath BA2 7AY, England.
- [ESY06] A. Eigenwillig, V. Sharma and Chee Yap. Almost tight complexity bounds for the Descartes method. In *Proc. Int'l Symp. Symbolic and Algebraic Computation (ISSAC'06)*, 2006. To appear. Genova, Italy. Jul 9-12, 2006.
- [Joh98] J.R. Johnson. Algorithms for polynomial real root isolation. In B.F. Caviness and J.R. Johnson, editors, *Quantifier Elimination and Cylindrical Algebraic Decomposition*, Texts and monographs in Symbolic Computation, pages 269–299. Springer, 1998.
- [KS94] M.-H. Kim and S. Sutherland. Polynomial root-finding algorithms and branched covers. *SIAM J. Computing*, 23:415–436, 1994.
- [LR01] T. Lickteig and M.-F. Roy. Sylvester-Habicht sequences and fast Cauchy index computation. *J. of Symbolic Computation*, 31:315–341, 2001.
- [Mc99] M. Mignotte and D. Ştefănescu. *Polynomials: An Algorithmic Approach*. Springer, 1999.
- [Mil92] P.S. Milne. On the solutions of a set of polynomial equations. In B. R. Donald, D. Kapur, and J. L. Mundy, editors, *Symbolic and Numerical Computation for Artificial Intelligence*, pages 89–102. Academic Press, London, 1992.
- [NR96] C.A. Neff and J.H. Reif. An efficient algorithm for the complex roots problem. *J. Complexity*, 12(3):81–115, 1996.
- [Pan96] V.Y. Pan. Optimal and nearly optimal algorithms for approximating polynomial zeros. *Computers Mathematics and Applications*, 31(12):97–138, 1996.
- [Pan97] V.Y. Pan. Solving a polynomial equation: some history and recent progress. *SIAM Review*, 39(2):187–220, 1997.
- [Ped90] P. Pedersen. Counting real zeroes. Technical Report 243, Courant Institute of Mathematical Sci., Robotics Lab., New York Univ., 1990. PhD Thesis, Courant Institute, New York University.
- [Pin76] J.R. Pinkert. An exact method for finding the roots of a complex polynomial. *ACM Trans. on Math. Software*, 2:351–363, 1976.
- [Rei97] D. Reischert. Asymptotically fast computation of subresultants. In *ISSAC 97*, pages 233–240, 1997. Maui, Hawaii.
- [Ren87] J. Renegar. On the worst-case arithmetic complexity of approximating zeros of polynomials. *Journal of Complexity*, 3:90–113, 1987.

- [Sch82] A. Schönhage. The fundamental theorem of algebra in terms of computational complexity, 1982. Manuscript, Department of Mathematics, University of Tübingen.
- [Str83] V. Strassen. The computational complexity of continued fractions. *SIAM J. Computing*, 12:1–27, 1983.
- [Wil78] H.S. Wilf. A global bisection algorithm for computing the zeros of polynomials in the complex plane. *Journal of the ACM*, 25:415–420, 1978.
- [Yap00] C.K. Yap. *Fundamental Problems of Algorithmic Algebra*. Oxford University Press, 2000.

Zilin Du, Vikram Sharma and Chee K. Yap
Courant Institute of Mathematical Sciences
New York University
New York, NY 10012, USA
e-mail: zilin@cs.nyu.edu
sharma@cs.nyu.edu
yap@cs.nyu.edu