Sturm Theory

Counting the number of real roots of a polynomial over a given interval was one of the most fundametal problems of the 19th Century. Some of the greatest names in Mathematics had worked on it, starting from Descartes' with his rule of signs, Newton on counting imaginar roots, to methods of Budan and Fourier. However, all these approaches only gave an upper estimate on the number of roots and not the exact count. It was Charles-François Sturm in 1829 who gave a method to derive the exact count.

Let $A(x) = \sum_{i=0}^{n} a_i x^i$ be a degree n polynomial with real coefficients. Further suppose $A(x)$ is square-free, i.e., $\texttt{GCD}(A, A')$ is a constant. Consider the prs

$$\overline{A} := (A_0, A_1, \ldots, A_k)$$

where $A_0 := A$ and $A_1 := A'$. We know that the following recurrence holds

$$A_{i-1} = Q_i A_i + A_{i+1}. \tag{1}$$

Let (a, b) be an interval in which we want to count the roots of $A(x)$. For any point $\alpha \in \mathbb{R}$, define $\mathrm{Var}_{\overline{A}}(\alpha)$ as the number of sign changes, i.e. change from positive to negative or vice versa, in the sequence obtained from $\overline{A}$ after dropping all the zeros. Sturm's theorem roughly states that

$$\mathrm{Var}_{\overline{A}}[a, b] := \mathrm{Var}_{\overline{A}}(a) - \mathrm{Var}_{\overline{A}}(b)$$

is the number of real roots of $A(x)$ in (a, b). Clearly, the number $\mathrm{Var}_{\overline{A}}(a)$ changes iff a passes through a root of one of members of the sequence $\overline{A}$. So it makes sense to study the behaviour of the sequence $\overline{A}$ in the neighborhoods of the roots of $A_i(x)$, $i = 0, \ldots, k$. Let us break However, the sequence $\overline{A}$ is not quite the correct one and we will see why.

We have the following observations about the sequence $\overline{A}$ w.r.t. α:

LEMMA 1. *Let $\overline{A}$ be the prs of A and A'.*

1. *No two consecutive elements of $\overline{A}$ vanish at α.*
2. *A_k does not vanish at α.*
3. *If $A_i(\alpha) = 0$ then $A_{i-1}(\alpha) A_{i+1}(\alpha) > 0$.*

The first two statements are a consequence of the fact that $A(x)$ is square-free. The last follows immediately from (1).

We now consider the quantitty $\mathrm{Var}_{\overline{A}}(\alpha)$ as we move across the roots of A_i's. There are two essentially two cases to consider:

Case 1. α **is a root of** $A(x)$**:** For x in the vicinity of α we have by the mean value theorem

$$A(x) = A'(\eta)(x - \alpha)$$

where η is between x and α. Since A is square-free the sign of the derivative A' in a small neighborhood of the root remains constant (either positive or negative). However, from the above equation it follows that to the left of α we have $A(x)A'(\eta) < 0$ and to the right of α we have $A(x)A'(\eta) > 0$. *Thus* $\mathrm{Var}_{\overline{A}}(\alpha)$ *drops by one as we move across* α*.*

Case 2. α **is a root of** $A_i(x)$**:** Note that α can be a multiple root of $A_i(x)$, $0 < i < k$ (since $A(x)$ is square-free $A_k(x)$ is a constant). For simplicity's sake let's assume α is a simple root of $A_i(x)$. Then we know that in a small neighborhood around α, A_{i-1} and A_{i+1} have the same sign. Now when we go across α, A_i changes sign, and in doing so we either add increase $\mathrm{Var}_{\overline{A}}$ by two or decrease it by two.

Thus we have shown that $\mathrm{Var}_{\overline{A}}[a, b]$ has the same parity as the number of real roots in (a, b). How do we avoid this difference of two? The crucial observation of Sturm was to change the sequence $\overline{A}$ by flipping signs of the remainders: a sequence $\overline{A}$ is called a **Sturm sequence** for A if the following recurrence is satisfied

$$A_{i-1} = A_i Q_i - A_{i+1}, \tag{2}$$

where $A_0 := A(x)$ and $A_1 := A'(x)$. Thus a Sturm sequence is obtained from a prs of $A(x)$ and $A'(x)$ by flipping the signs of the remainders. Now we have the following analogue to Lemma 1.

LEMMA 2. *Let $\overline{A}$ be the Sturm sequence for A. Then*

1. *No two consecutive elements of $\overline{A}$ vanish at α.*
2. *A_k does not vanish at α.*
3. *If $A_i(\alpha) = 0$ then $A_{i-1}(\alpha)A_{i+1}(\alpha) < 0$.*

The argument in Case 1. remains unchanged, i.e., we still loose one sign variation in going across a root of $A(x)$. In Case 2., however, we do not loose any sign variation, because in the neighborhood of α, the sign of A_{i-1} and A_{i+1} is opposite, and hence whether or not A_i changes sign across α, we do not loose any sign variation. Thus we have the following classic result.

THEOREM 3. *Given a square-free polynomial $A(x) \in \mathbb{R}[x]$ let $\overline{A}$ be its Sturm sequence. Then $\mathrm{Var}_{\overline{A}}[a, b]$ is the number of real roots of $A(x)$ in the interval (a, b), where a, b are not roots of $A(x)$.*

Note: It's not clear in the first place how the drop in the sign changes are stored in the sequence $\overline{A}$. What is interesting is that between two roots α, β of A, even though A and A' have the same sign to the right of α, by the time we reach β their sign differs, which it must by Rolle's theorem, so that we again have a drop of sign across β. In between the two roots there is no change in sign variation, but the distribution of signs across A_i's, $0 < i < k$, changes.

Is the restriction of square-freeness really necessary? The surprising thing is that it is not. Suppose α is a root of multiplicity m of $A(x)$. Then we know that $(x - \alpha)^{m-1}$ divides $A'(x)$ and hence $(x - \alpha)^{m-1}$ divides $\mathtt{GCD}(A, A')$. This implies that $(x - \alpha)^{m-1}$ divides all the polynomials in the Sturm sequence. Define ϕ_i, $i = 0, \ldots, k$, s.t. $A_i = (x - \alpha)^{m-1}\phi_i$; further define $\overline{\phi}$ as the sequence $(\phi_0, \ldots, \phi_k)$. Then Lemma 2 applies to the sequence $\overline{\phi}$. We claim that the sequence $\overline{\phi}$ drops a sign variation as we move across α. The two cases again apply. Since α occurs with multiplicity one in ϕ_0 zero in ϕ_1, we always loose a sign between ϕ_0 and ϕ_1 when going across α. If α is a root of some ϕ_i, $i > 1$, then the sign of ϕ_{i-1} and ϕ_{i+1} is different, so we do not loose any sign variations when ϕ_i changes sign across α. But what is the relation between $\mathrm{Var}_{\overline{\phi}}$ and $\mathrm{Var}_{\overline{A}}$ at some point x? The two quantities are the same because $A_i = (x - \alpha)^{m-1}\phi_i$, for $i = 0, \ldots, k$, i.e., as we have have scaled all the polynomials by the same quantity so the sign variations remain the same. Thus we have completed the proof of the following famous result.

THEOREM 4 (Sturm's Theorem). *Given a polynomial $A(x) \in \mathbb{R}[x]$ let $\overline{A}$ be its Sturm sequence. Then $\mathrm{Var}_{\overline{A}}[a, b]$ is the number of distinct real roots of $A(x)$ in the interval (a, b), where a, b are not roots of $A(x)$.*

Remark: The theorem applies to integer polynomials, where the Sturm sequence is obtained by changing signs of the remainders in the Subresultant PRS algorithm.

We can generalize the notion of a sturm sequence to any two polynomials $A, B \in \mathbb{R}[x]$, as the sequence $(A_0, A_1, \ldots, A_k)$ satisfying (2), where $A_0 := A, A_1 := B$. Note that if $\deg(A) < \deg(B)$ then $A_2 = -A$, and subsequently the sequence proceeds as the standard sturm sequence. Given an interval $[a, b]$, the concept of $\mathrm{Var}_{A,B}[a, b]$ is well defined.

1 Algorithm for Isolating Real Roots

In this section we describe and analyse an algorithm for isolating real roots of an integer polynomial $A(x)$. The algorithm is a straightforward bisection algorithm.

INPUT: An integer polynomial $A(x)$.
OUTPUT: A sequence of pairwise disjoint intervals $(I_1, \ldots, I_\ell)$ such that each interval contains exactly one real root of $A(x)$ and together the account for all the real roots of $A(x)$.

```
1.  Compute the sturm sequence Ā of A(x).
2.  Compute an interval (−B, B) containing all the real roots.
3.  Initialize a queue Q ← (−B, B).
4.  While Q is not empty do:
5.      Remove an interval I from Q.
        If Var_Ā(I) == 1 then output I.
        Else If Var_Ā(I) == 0 then discard I.
        Else
            If midpoint of I is a root then output [m, m].
            Split I into two equal halves and push them onto Q.
```

There are three quantities that we have to bound:

Step 1. Complexity of computing the Sturm sequence in step 1.

Step 2. Complexity of evaluating a sturm sequence at a given point.

Step 3. The size of the subdivision tree.

Complexity for Step 1: A half-gcd approach can be used to evaluate the subresultants. Very crudely speaking this takes $O(M(n)\log n)$ operations on numbers of bit-size $O(nL)$. Thus the overall complexity is $O(M(n)\log n \cdot M(nL)) = \widetilde{O}(n^2L)$, where $\widetilde{O}$ means we ignore logarithmic factors.

Complexity for Step 2: The right approach to evaluate a Sturm sequence is to store A, A' and the pseudoquotients sequence. It can be shown that the bit-complexity of the pseudo-quotients is $O(\delta_i nL)$. The overall complexity for evaluating the sturm sequence at a rational of bit-length $O(nL)$ is $O(\sum_{i=1}^{k}\delta_i^2 nL) = \widetilde{O}(n^3L)$.

In the remaining section we bound the size of the subdivision tree. In particular, we will show that the size is $\widetilde{O}(nL)$. This is surprising because a straightforward estimate would be $\widetilde{O}(n^2L)$: the worst case depth of the tree is $\widetilde{O}(nL)$, and there are at most n leaves in the tree. The improvement is obtained by amortizing over all the leaves simultaneously. Multiplying the tree size, with the worst case complexity of each node from Step 3 is $\widetilde{O}(n^3L)$, we obtain that the overall complexity of the algorithm is $\widetilde{O}(n^4L^2)$.

1.1 Size of the subdivision tree

Let T be the subdivision tree of the algorithm. We first prune all the leaves from T to obtained T'. Since T is a binary tree $|T| \le 2|T'|$. We will thus bound $|T'|$.

Consider an interval J associated with a leaf of T'. Since J was not terminal in T, we know that there are at most two distinct roots $\alpha_J, \beta_J \in J$. Thus $w(J) \ge |\alpha_J - \beta_J|$. The number of nodes along the path from the root of the tree T', with the associated interval I_0, to J is $\lg w(I_0)/w(J)$. From the lower bound on $w(J)$ it follows that the depth of this path is bounded by $\lg w(I_0)/|\alpha_J - \beta_J|$. Summing over all the leaves J in T' we have the following bound:

$$|T'| \le \sum_J \lg \frac{w(I_0)}{|\alpha_J - \beta_J|}. \tag{3}$$

Since there can be at most n real roots we further have

$$|T'| \le n \lg w(I_0) - \sum_J \lg |\alpha_J - \beta_J| = n \lg w(I_0) - \lg \prod_J |\alpha_J - \beta_J|.$$

To derive an upper bound on $|T'|$, we thus need to derive a lower bound on $\prod_J |\alpha_J - \beta_J|$. Note that for two leaves J, J' of T' the associated root pairs are distinct. Applying the Davenport-Mahler-Mignotte bound to the square-free part of $A(x)$, i.e., to $A(x)/\texttt{GCD}(A, A')$, it follows that

$$\prod_J |\alpha_J - \beta_J| \geq 2^{-\binom{n}{2}} M(A)^{2(1-n)}.$$

Thus

$$|T'| = O(n(L+n)).$$

By a more careful analysis, the bound can be improved to $O(n(L + \log n))$. Furthermore, it can be shown that for the Mignotte polynomials this is the best possible.

2 Isolating Complex Roots

A sturm sequence helped us to count the exact number of real roots in any inteval $I \subseteq \mathbb{R}$. To isolate complex roots, we need a similar analogue that can help us count the exact number of complex roots in any rectangular region in $\mathbb{C}$. We first develop such a function, and then show how to compute it. This section is based on Herb Wilf's paper [1].

Consider a point $\alpha \in \mathbb{C}$ and box $B \in \mathbb{C}$. Suppose a point z moves on the boundary of B, ∂B, in a counter clockwise manner, then how does $\arg(z - \alpha)$ change? For any function $f : \mathbb{C} \to \mathbb{C}$, let Let $\Delta_B \arg f(z)$ represent the change of argument of $f(z)$, as z goes around ∂B. There are two cases to consider: if α is inside B then the argument changes by 2π, as z makes one loop around B; if α is outside B then the argument does not change. Thus

$$\frac{1}{2\pi}\Delta_B \arg(z - \alpha) = \begin{cases} 1 & \text{if } \alpha \in B; \\ 0 & \text{if } \alpha \notin B. \end{cases} \tag{4}$$

Now suppose that we have two points α, β, then it is clear that for any point $z \in \mathbb{C}$,

$$\arg[(z - \alpha)(z - \beta)] = \arg(z - \alpha) + \arg(z - \beta). \tag{5}$$

For any polynomial $A(z) \in \mathbb{C}[z]$, applying (4) and (5) to its linear factors $(x - \alpha_i)$ it follows that as z varies over ∂B in a counter-clockwise direction

$$\frac{1}{2\pi}\Delta_B \arg A(z) = \text{ Number of roots of } A(z) \text{ inside } B. \tag{6}$$

This observation is called **the principle of argument**. For correctness, we must assume that no root of $A(z)$ lies on the ∂B. Thus if we can compute $\Delta_B \arg A(z)$ then we can count the exact number of roots of $A(z)$ inside B.

Let's focus on $\Delta_B \arg A(z)$. The curve traced by $A(z)$, as z varies around B in a counter-clockwise direction, is a closed curve (not necessarily simple, i.e., it may self-intersect). The change in the argument of $A(z)$ is closely related to the winding number of $A(z)$ around the origin; in fact, it is 2π times the winding number. However, there is another combinatorial way to express the change in argument which is based upon the following observation: as the curve $A(z)$ goes from $+ve$ y-axis to $-ve$ y-axis crossing over the $-ve$ x-axis, the argument increases by π; if it continues further from $-ve$ y-axis to $+ve$ y-axis crossing over the $+ve$ x-axis, then the argument increases by π; if, however, we go in the other direction, i.e., from $-ve$ y-axis to $+ve$ y-axis crossing over the $-ve$ x-axis, then the argument decreases by π; we also decrease by π when going from $+ve$ y-axis to $-ve$ y-axis crossing over the $+ve$ x-axis. A more succinct way of expressing this observation is to define

$$N_+^- := \#\left\{\text{ continuos arcs of } A(z) \text{ that go from a } -\text{ve quadrant to } +\text{ve quadrant and cross x-axis}\right\},$$

and similarly

$$N_+^- := \#\left\{\text{ continuos arcs of } A(z) \text{ that go from a } +\text{ve quadrant to } -\text{ve quadrant and cross x-axis}\right\}.$$

Then

$$\Delta_B \arg A(z) = \pi(N_-^+ - N_+^-). \tag{7}$$

See Figure **??** for an illustration of the concepts above. We will say type-1 arcs for those arcs that contribute to the count N_-^+, and type-2 arcs for those that contribute to N_+^-. To compute N_-^+ and N_+^-, we first study the curve $A(z)$ as z traverses around one edge of B.

Consider an edge $ab \subset \mathbb{C}$ of B. Any point z on this edge can be expressed as $z = a + (b-a)t$, where $t \in [0,1]$. Thus

$$A(z) = A(a + (b-a)t) = A_R(t) + iA_I(t)$$

where $A_{R/I}(t)$ are polynomials with real coefficients. We further assume that A_R and A_I are relatively prime, which implies that they don't have any common real roots. Define the rational function

$$\rho(t) := \frac{A_I(t)}{A_R(t)}. \tag{8}$$

Note that, $\arg A(t) = \arctan \rho(t)$, thus we have to study the rational function $\rho(t)$. Let $N[ab]_-^+$ denote the number of type-1 arcs the curve $A(z)$ has as z goes from a to b; similarly define $N[ab]_+^-$; thus N_-^+ is roughly the sum of these counts for the four edges of B. Note that the starting and ending point of these arcs correspond to two consecutive roots of $A_R(t)$ in $[0,1]$. So let

$$t_0 := 0 \le t_1 < t_2 < \cdots < t_k \le 1$$

be the k roots of A_R in $[0,1]$. Let t_i^+ denote a point just greater than t_i and t_i^- a point just smaller than t_i; corresponding to these points in the nieghboirhood of t_i, we have the value of the function $\rho(t_i^\pm)$ appropriately defined by a one-sided limit. We will express $N[ab]_-^+$ in terms of the signs of $\rho(t_i^\pm)$. Note that an arc of type-1 occurs when $\rho(t_i^+) < 0$ and $\rho(t_{i+1}^-) > 0$. Similarly, an arc of type-2 occurs when $\rho(t_i^+) > 0$ and $\rho(t_{i+1}^-) > 0$. Since we are interested in the difference $N[ab]_-^+ - N[ab]_+^-$, a neat way to count arcs of either type is the following term

$$\left(\frac{\mathtt{sign}(\rho(t_{i+1}^-)) - \mathtt{sign}(\rho(t_i^+))}{2}\right)$$

that takes 1 for arcs of type-1 and -1 for type-2 arcs; it vanishes if both the signs are the same. Thus

$$\sum_{i=1}^{k-1}\left(\frac{\mathtt{sign}(\rho(t_{i+1}^-)) - \mathtt{sign}(\rho(t_i^+))}{2}\right)$$

gives us the excess of type-1 arcs over type-2 arcs. The problem is at the starting and ending points. Since we are counting arcs along individual edges, it may be that the curve $A(z)$ as z goes from the edge ab to edge cd combine to give us an arc of one of the two types. To do this accounting correctly, we further add $\mathtt{sign}(\rho(t_1^-))/2$ and $\mathtt{sign}(\rho(t_k^+))/2$, and thus we obtain that the excess of type-1 arcs over type-2 arcs along an edge ab of B is

$$\begin{aligned} N[ab]_-^+ - N[ab]_+^- &= \frac{\mathtt{sign}(\rho(t_1^-))}{2} + \sum_{i=1}^{k-1}\left(\frac{\mathtt{sign}(\rho(t_{i+1}^-)) - \mathtt{sign}(\rho(t_i^+))}{2}\right) + \frac{\mathtt{sign}(\rho(t_k^+))}{2} \\ &= \sum_{i=1}^{k}\left(\frac{\mathtt{sign}(\rho(t_{i+1}^-)) - \mathtt{sign}(\rho(t_i^+))}{2}\right). \end{aligned} \tag{9}$$

To denote the dependence on the edges of B, let ρ_j be the rational function obtained by expanding $A(z)$ along the jth edge, and let k_j the number of poles of γ_{ab} in $[0,1]$ (observe that t_i's are the poles of $\gamma(t)$). Then from (6) and (7) it follows that

$$N_-^+ - N_+^- = \sum_{j=1}^{4}\sum_{i=1}^{k_j}\left(\frac{\mathtt{sign}(\rho_j(t_{i+1,j}^-)) - \mathtt{sign}(\rho_j(t_{i,j}^+))}{2}\right) \tag{10}$$

where $t_{i,j}$ denotes the poles of γ_j. Our aim, therefore, is to compute

$$\sum_{i=1}^{k}\left(\frac{\mathtt{sign}(\rho(t_{i+1}^-)) - \mathtt{sign}(\rho(t_i^+))}{2}\right).$$

This sum is called the **Cauchy index** of ρ over the interval $[0, 1]$, denoted as $I_0^1\rho(t)$. It should be interpreted as follows: the fraction is positive for those poles t_i for which the sign across t_i changes from positive to negative, as t increases from 0 to 1; similarly, it is negative for those poles t_i for which the sign across t_i changes from negative to positive, as t increases from 0 to 1; thus the sum is how much in excess the first type of poles are over the second type of poles. The surprising thing is that the Cauchy index of ρ can be computed from a generalization of sturm sequences for the relatively prime polynomials $A_R(t)$ and $A_I(t)$.[1] More precisely, we will show in the next section that

$$I_0^1\rho(t) = \mathrm{Var}_{A_I, A_R}[0, 1]. \tag{11}$$

This equation along with (10) implies the following result.

THEOREM 5. *Suppose $A(z)$ has no boundaries on the boundary of a box B. Then the number of zeros of $A(z)$ in the interior of B is exactly*

$$\frac{1}{2}\sum_{j=1}^{4} I_0^1\left(\frac{A_I^{(j)}}{A_R^{(j)}}\right)$$

where $A_I^{(i)}, A_R^{(i)}$ are the polynomials corresponding to the jth edge of B.

We defer the proof (11) to the next section, and instead describe an algorithm to isolate the complex roots of $A(z)$. Given the theorem above the algorithm is pretty straightforward.

INPUT: A polynomial $A(z) \in \mathbb{C}[z]$.
OUTPUT: Isolating boxes for all the distinct roots of $A(z)$.
1. Compute a box B_0 containing all the roots of A; use Cauchy's bound.
2. Initialize a queue $Q \leftarrow B_0$.
3. While Q is not empty do
   - Remove a box B from Q.
   - Compute the real $A_R(t)$ and the imaginary part $A_I(t)$ corresponding to each edge of B:
   - Compute the Cauchy index $I_0^1\rho(t)$ for each of the four edges
   - Add the Cauchy indices to get an estimate on the number of roots in B.
   - If the estimate is one then output B; if zero then discard.
   - If estimate > 1 then subdivided B into four boxes of equal size and push these boxes onto Q.

¶1. The Size of the Subdivision Tree The analysis for the size of the subdivision tree T is similar to what we had earlier. Let T' be the tree obtained by pruning all the leaves of T. Then we know that for a box B associated with a leaf of T', there must be two roots $\alpha_B, \beta_B \in B$. Thus $2w(B) \geq |\alpha_B - \beta_B|$; note the two, because the roots may be diagonally placed in B. Let $N(\alpha)$ denote a root of A closest to α. Then we have the lower bound

$$2w(B) \geq |\alpha_B - N(\alpha_B)|.$$

Proceeding as was done earlier, we have that

$$|T'| \leq n \log w(B_0) - n \log_B \prod |\alpha_B - N(\alpha_B)|.$$

For deriving a lower bound on the product, we can again use the discriminant based approach, to obtain that

$$\prod |\alpha_B - N(\alpha_B)| \geq M(A)^{2(1-n)} 2^{\binom{n}{2}} \sqrt{\mathrm{disc}(A)}.$$

Thus we obtain taht

$$|T'| \leq \widetilde{O}(n(L+n)).$$

A crucial difference between the algorithms for complex root isolation and the real root isolation is that the former needs to compute four new sturm sequences for every box, unlike the latter where we used just one sturm sequence. In practice, this is very costly to do.

[1] One of the main contributions of Sturm was to show that the Cauchy index can be computed from Sturm sequences for the case of rational functions.

2.1 Computing Cauchy Index

We want to show that the Cauchy index of a rational function $\rho := A/B$, where $A, B \in \mathbb{R}[x]$ are two relatively prime polynomials, w.r.t. an interval $[a, b]$ satisfies

$$I_a^b \rho(t) = -\mathrm{Var}_{A,B}[a, b]. \tag{12}$$

The proof of this relation is very similar to the proof of the Sturm theorem Theorem 3. Let $a \leq \gamma_1 \leq \gamma_2 \leq \cdots \leq \gamma_k \leq b$, be the roots of $B(x)$. By definition of Cauchy index we know that

$$I_a^b \rho = \sum_{i=1}^{k} \left(\frac{\mathtt{sign}(\rho(\gamma_i^-)) - \mathtt{sign}(\rho(\gamma_i^+))}{2} \right).$$

The roots γ of B can be characterized into three types: type-1 are those for which the fraction on the RHS is positive, type-2 are those for which the fraction is negative, and type-3 are those for which it is zero. We have to show that $-\mathrm{Var}_{A,B}[a, b]$ counts these types correctly, namely for type-1 roots $\mathrm{Var}_{A,B}$ drops by one, for type-2 roots the sign $\mathrm{Var}_{A,B}$ increases by one, and for type-3 root $\mathrm{Var}_{A,B}$ does not change.

Consider the sturm sequence $(A, B, A_2, \ldots, A_k)$ for (A, B). We have the following three cases correspoding to the three types of roots of B:

1. For a type-1 pole the sign of A, B is the same to the left of γ and becomes the same to the right of γ. Thus $\mathrm{Var}_{A,B}$ increases by one across γ.

2. For a type-2 pole the sign of A, B is different to the left of γ and becomes the same to the right of γ. Thus $\mathrm{Var}_{A,B}$ decreases by one across γ.

3. For a type-3 pole the sign of A, B is the same to the left of γ and remains the same to the right of γ. Thus $\mathrm{Var}_{A,B}$ does not change across γ; note the signs of A and B might flip, but they flip in the same way.

References

[1] H. S. Wilf. A global bisection algorithm for computing the zeros of polynomials in the complex plane. *J. ACM*, 25(3):415–420, 1978.