

Quadratic Interval Refinement for Real Roots

John Abbott
Dipartimento di Matematica
Università di Genova
Genova 16146, Italy
abbott@dimma.unige.it

Abstract

We present a new algorithm for refining a real interval containing a single real root: the new method combines characteristics of the classical Bisection algorithm and Newton's Iteration. Our method exhibits quadratic convergence when refining isolating intervals of simple roots of polynomials (and other well-behaved functions). We assume the use of arbitrary precision rational arithmetic. Unlike Newton's Iteration our method does not need to evaluate the derivative.

1 Introduction

Typically the process of approximating the real roots of a univariate polynomial comprises two phases: *root isolation* where the distinct roots are separated into disjoint intervals, followed by *root refinement* where the approximations of the roots are improved until they are within specified limits. This paper is concerned with refinement, and assumes that isolation has already taken place (see for instance [1]).

There are many methods for root refinement. Newton's Iteration is a well-known example, and dates back over three hundred years. Bisection is another refinement method, slower than Newton's Iteration but more robust. This article presents a variant of Regula Falsi which combines the robustness of Bisection with the rapid convergence of Newton's Iteration. The algorithm is implemented as an integral part of CoCoA (from version 4.3 onwards). We call the new algorithm **Quadratic Interval Refinement**, or simply QIR.

Prerequisites for using QIR are knowledge of an initial interval I for which the function has opposite signs at the two end points (both must be rational), and a procedure which evaluates f at a given rational point with arbitrary precision. Naturally, the function f must be continuous on I . It is best if the interval I contains a single simple root of f ; furthermore, if f has a valid Taylor expansion about that single simple root then convergence of QIR will ultimately be quadratic. If I contains several roots then QIR will eventually discard all but one of them as refinement proceeds.

1.1 Notes

This algorithm was originally presented as a poster at the ISSAC 2006 conference; the software implementation was released publicly in late 2003 (when CoCoA 4.3 was first distributed).

An interesting and practical generalization of our algorithm to the case of approximate arithmetic is given in [6], together with a detailed bit-complexity analysis. They report a significant improvement in speed over the exact arithmetic version described here.

1.2 Notation

For a real number x we write $\text{round}(x)$ to mean the integer closest to x ; either candidate may be chosen in the case of a tie. We write $I(a, b)$ to denote the open interval $\{x \in \mathbb{R} : a < x < b\}$.

We abuse terminology harmlessly by taking **isolating interval** to mean an interval $I(a, b)$ for which $f(a)$ and $f(b)$ have opposite sign — we do not require that the interval contain just a single root of f .

By **Newton's Iteration** we refer to the root approximation method starting from some given x_0 and generating successive iterates using the formula $x_n = x_{n-1} - f(x_{n-1})/f'(x_{n-1})$. This method is sometimes called *Newton-Raphson Iteration*.

By **Bisection** we refer to the interval refinement method where the interval $I(a, b)$ is refined to either $I(a, \frac{1}{2}(a+b))$ or $I(\frac{1}{2}(a+b), b)$, the choice depending on the sign of $f(\frac{1}{2}(a+b))$. This method is sometimes also called *Binary Chop*.

By **Regula Falsi** we refer to the interval refinement method where linear interpolation is used to predict the position of the root, and then *one* end of the interval is moved to this predicted value so that the new interval still contains a sign change.

2 Motivation

Why invent a new method when Newton's Iteration works so well? QIR offers a simple unified approach which can refine any isolating interval with a rate of convergence and a computational cost comparable to that of Newton's Iteration. Thus QIR has several advantages over Newton's Iteration:

- (a) the criteria guaranteeing convergence are straightforward and easy to verify;
- (b) there is no need to evaluate the derivative of the function whose zeroes we are approximating;
- (c) the successive approximants have simple denominators;
- (d) roots are approximated by intervals containing them, so the accuracy of an approximation is quite explicit.

A significant difficulty with Newton's Iteration is obtaining a suitable starting value given an isolating interval: this is the essence of point (a). It may be necessary to use some other interval refinement method prior to using Newton's Iteration. Point (b) poses no problem when approximating roots of explicit polynomials since the derivative can readily be evaluated, but in other cases it may not be easy to obtain values of the derivative. Point (c) is relevant only when using arbitrary precision rational arithmetic. Point (d) is important when a guarantee of the accuracy obtained is desired. With some cunning one can alleviate weaknesses (c) and (d) of Newton's Iteration: *e.g.* replacing approximants by similar values having a simpler denominator of about the right size, and using the derivative to estimate the width of an interval containing the root.

Example QIR can never fail, whereas Newton’s Iteration can. The polynomial $f(x) = x^3 - x + 0.7$ has a single real root, but Newton’s Iteration fails to converge when given a starting value x_0 in the range $-0.1 < x_0 < 0.1$; so bad starting points are not so rare.

Note that one should avoid naive use of Newton’s Iteration with exact rational arithmetic because the sizes of the numerators and denominators of successive approximants can increase rapidly: typically the k -th approximant will have numerator (and also denominator) containing $O(d^k)$ digits where d is the degree of the polynomial whose root is being sought.

Example The polynomial $x^5 - 2$ contains a real root in the interval $I(1, 2)$. To estimate the root with an error less than 2^{-32} starting from the initial approximation $x_0 = 1$ requires five Newton iterations, and produces an approximant having a denominator with over 500 digits. Alternatively starting from $x_0 = 2$ requires seven iterations, and the final approximant has a denominator with over 17000 digits. In contrast, QIR obtains an interval of width $\leq 2^{-32}$ after six iterations, and no value in the computation has more than 50 digits.

3 Description of the Method

QIR works by repeatedly narrowing the isolating interval until the width is smaller than the prescribed limit. Each individual narrowing step receives an **initial interval** and a **refinement factor** $N \in \mathbb{N}$ by which it must reduce the interval width; it uses *discretized* linear interpolation to guess a narrower interval containing the root. If the guess was good, the interval is updated; if not, the interval is unchanged, and the narrowing step returns an indication of failure. The refinement factor is increased after every successful narrowing step.

A narrowing step conceptually divides its initial interval into N consecutive equal width sub-intervals. It then uses linear interpolation to locate the sub-interval in which it guesses the root to lie. The guess is tested by evaluation at the end points of the chosen sub-interval. If the guess was lucky then the initial interval is replaced by the chosen sub-interval; otherwise an indication of failure is returned. The case $N = 4$ is handled specially: the initial interval is *always* replaced by the correct sub-interval, *but* failure is indicated if linear interpolation led to a bad guess.

The refinement factor is varied according to the following rule:

- after a successful narrowing step, $N \leftarrow N^2$
- after a failed narrowing step, if $N > 4$ then $N \leftarrow \sqrt{N}$

This strategy is inspired by the knowledge that linear interpolation produces approximations whose errors decrease roughly quadratically when sufficiently close to a root of a well-behaved function.

Obviously, if one of the evaluation points happens to be the exact root ξ then this value is returned as an exact root. This is possible only if ξ is rational and with “suitable” denominator.

3.1 An Illustrative Example

We present here an example where QIR appears to achieve quadratic convergence immediately (but only due to “good fortune”), then later it is forced to reduce the refinement factor back to 4 for a few iterations before true quadratic convergence can begin. Note how few evaluations are needed to attain quite high accuracy.

Example Given the polynomial $f = 10^{200} \cdot x^2 - 1$ the root isolator in CoCoA produces the interval $I(0, 2)$ for the positive root. Starting from this interval QIR initially enjoys seven **successes**, increasing the refinement factor to 4^{128} , before encountering a number of **failures** which force the refinement factor back down to 4. Then once the interval width has become small enough (in this case, less than 10^{-100} , after 24 iterations in `RefineInterval`) true quadratic convergence begins. For instance, an interval of width less than 10^{-1000} is obtained after a total of 34 iterations, *i.e.* at most 68 evaluations of f ; a further four iterations give an interval of width less than 10^{-10000} .

4 The Algorithm

Here we present explicitly the algorithm using a pseudo-language. The actual source code is contained in the CoCoA package `RealRoots.cpkg` and is publicly available as part of the standard distribution of CoCoA (from version 4.3 onwards) — see the web site [2].

4.1 Main routine: `RefineInterval`

INPUT: f the function whose zero is being sought,

$I(x_{lo}, x_{hi})$ an open interval with rational end points $x_{lo} < x_{hi}$, and for which $f(x_{lo}) \cdot f(x_{hi}) < 0$
 $\varepsilon > 0$ an upper bound for width of the refined interval.

OUTPUT: $I(\xi_{lo}, \xi_{hi})$ an open sub-interval of width at most ε having rational end points and such that $f(\xi_{lo}) \cdot f(\xi_{hi}) < 0$. Exceptionally the output may be the exact value of the root ξ .

(1) Initialize refinement factor $N \leftarrow 4$, and interval $I \leftarrow I(x_{lo}, x_{hi})$.

(2) While $width(I) > \varepsilon$ do steps (2.1)–(2.4).

(2.1) Apply `RefineIntervalByFactor` to f , I and N — this usually modifies the interval I .

(2.2) If an **exact root** ξ has been found, simply return ξ .

(2.3) If **success** was reported then replace $N \leftarrow N^2$.

(2.4) Otherwise **failure** was reported, so if $N > 4$ replace $N \leftarrow \sqrt{N}$.

(3) Return the interval I .

4.2 Auxiliary procedure: `RefineIntervalByFactor` when $N > 4$

INPUTS: f the function whose zero we are seeking,

$I = I(x_{lo}, x_{hi})$ an open interval with rational end points $x_{lo} < x_{hi}$ such that $f(x_{lo}) \cdot f(x_{hi}) < 0$
 $N \in \mathbb{N}$ the refinement factor — here we assume $N > 4$, see §4.3 below for the case $N = 4$.

OUTPUT: **success** or **failure**; if **success** then as a side effect the interval I is replaced by a open sub-interval $I(\xi_{lo}, \xi_{hi})$ having rational end points, whose width is $1/N$ times that of I , and such that $f(\xi_{lo}) \cdot f(\xi_{hi}) < 0$. Exceptionally, the output may be **exact root** ξ .

(1) Conceptually divide the interval I into N consecutive sub-intervals of equal width. Compute the width of each sub-interval: $w = (x_{hi} - x_{lo})/N$. Define $x_0 = x_{lo}$, and $x_k = x_{k-1} + w$ for $k = 1, 2, \dots, N$; so these x_i are the end points of the sub-intervals.

- (2) Using linear interpolation predict approximate position of the zero: *i.e.* determine the closest end point $\hat{x} = x_\kappa$ where $\kappa = \text{round}(N \cdot f(x_{lo}) / (f(x_{lo}) - f(x_{hi})))$
- (3) If $f(\hat{x}) = 0$ then return **exact root** \hat{x} .
- (4) Check whether our prediction was good; *i.e.* the function f changes sign in one of the sub-intervals having \hat{x} as an end point.
- (5) If $f(\hat{x})$ has the same sign as $f(x_{lo})$, check the interval to the right:
 - (5.1) If $f(\hat{x} + w) = 0$ then return **exact root** $\hat{x} + w$.
 - (5.2) If $f(\hat{x} + w)$ has the same sign as $f(\hat{x})$ then return **failure**.
 - (5.3) Otherwise replace $I \leftarrow I(\hat{x}, \hat{x} + w)$ and return **success**.
- (6) Otherwise $f(\hat{x})$ has the same sign as $f(x_{hi})$, so check the interval to the left:
 - If $f(\hat{x} - w) = 0$ then return **exact root** $\hat{x} - w$.
 - If $f(\hat{x} - w)$ has the same sign as $f(\hat{x})$ then return **failure**,
 - Otherwise replace $I \leftarrow I(\hat{x} - w, \hat{x})$ and return **success**.

4.3 Auxiliary procedure: RefineIntervalByFactor when $N = 4$

We give a verbal description rather than pseudo-code for `RefineIntervalByFactor` with $N = 4$. Use linear interpolation to predict which of the 5 sub-interval end points is closest to the root. Refine the input interval using Bisection twice. If one of the end points of the refined interval agrees with our prediction, return **success**, otherwise return **failure**. Note that the initial interval is *always* narrowed by a factor of 4.

4.4 Termination of the Algorithm

We show that the algorithm always terminates. In `RefineInterval` if $N = 4$ then the interval is always narrowed; if $N > 4$ then either the interval is narrowed (by more than a factor of 4), or N is reduced. Since `RefineInterval` terminates when the interval has width less than ε , and since each narrowing reduces the width by at least a factor of 4, there can be only a finite number of narrowings. If no narrowing occurs (*i.e.* $N > 4$ and **failure** was reported) then N is reduced to its own square-root, but as N is never smaller than 4, there can be only finitely many iterations in which the interval is not narrowed.

4.5 Quadratic Convergence

The mathematical justification underlying the convergence rate of QIR is very similar to that underlying Newton's Iteration. We assume that the isolating interval I contains a single simple root ξ and that the function f admits a Taylor expansion centred on ξ valid in an open neighbourhood: *i.e.* for sufficiently small δ we have

$$f(\xi + \delta) = C_1 \cdot \delta + C_2 \cdot \delta^k + O(\delta^{k+1})$$

for some exponent $k \geq 2$ and suitable constants C_1 and C_2 ; moreover $C_1 \neq 0$ since ξ is a simple root.

Let x_{lo} and x_{hi} be the two end points of the interval I and ε be its width. When ε is small enough we have:

$$\begin{aligned} f(x_{lo}) &= C_1 \cdot (x_{lo} - \xi) + C_2 \cdot (x_{lo} - \xi)^k + O(x_{lo} - \xi)^{k+1} \\ f(x_{hi}) &= C_1 \cdot (x_{hi} - \xi) + C_2 \cdot (x_{hi} - \xi)^k + O(x_{hi} - \xi)^{k+1} \end{aligned}$$

Estimating ξ by linear interpolation gives $\hat{\xi} = x_{lo} + \lambda \cdot (x_{hi} - x_{lo})$ where $\lambda = f(x_{lo}) / (f(x_{lo}) - f(x_{hi}))$. Observe that $f(x_{lo}) - f(x_{hi}) = C_1\varepsilon + C_2\eta + O(\varepsilon^{k+1})$ where $|\eta| = |(x_{lo} - \xi)^k - (x_{hi} - \xi)^k| \leq \varepsilon^k$. Substituting and simplifying gives $\hat{\xi} = x_{lo} + (\xi - x_{lo}) - \frac{C_2}{C_1}(x_{lo} - \xi)(\frac{\eta}{\varepsilon} - (x_{lo} - \xi)^{k-1}) + O(\varepsilon^{k+1})$. Regarding cubic and higher order terms in ε as negligible, this tells us that $|\hat{\xi} - \xi| < \varepsilon^2 \cdot 2C_2/C_1$.

Now we show that convergence of QIR is eventually quadratic. We say that an isolating interval I of width ε is **narrow enough** for f whenever $\varepsilon \cdot C_2/C_1 < 1/8$. Observe that, starting from any isolating interval, this condition will be satisfied after only finitely many iterations. Henceforth we shall assume I is narrow enough.

When applying the algorithm `RefineIntervalByFactor` to the parameters (f, I, N) we shall say that N is **small enough for I** if $N \cdot \varepsilon \cdot C_2/C_1 < 1/2$. Note that $N = 4$ is always small enough because we have assumed that I is narrow enough.

Suppose that algorithm `RefineIntervalByFactor` is called on (f, I, N) where N is small enough for I . The condition $\varepsilon \cdot C_2/C_1 < 1/8$ implies that linear interpolation produces an estimate $\hat{\xi}$ whose distance from ξ does not exceed $\varepsilon^2 \cdot 2C_2/C_1 < \varepsilon/N$. Hence step (2) of `RefineIntervalByFactor` makes a good prediction, and so **success** will be returned. Denote the narrowed interval by I' . Now, the next iteration of `RefineInterval` will call `RefineIntervalByFactor` on the parameters (f, I', N^2) . Since I' is narrower than I it is surely narrow enough for f . Moreover, one may easily verify that N^2 is small enough for I' . Hence, by induction, all subsequent calls to `RefineIntervalByFactor` will return **success**, and convergence is therefore quadratic.

In contrast, while still assuming that I is narrow enough for f , if `RefineIntervalByFactor` is called when N is not small enough for I then **failure** may occur. Since N is not small enough, we must have $N > 4$. So when **failure** occurs the refinement factor N is reduced. Hence, there can be only finitely many **failures** before N becomes small enough for I , and then guaranteed quadratic convergence ensues.

4.6 Weaknesses

A weakness of QIR as described here is the cost of evaluating f exactly at various rational points. For instance, if f is a high degree polynomial and the evaluation point x is a rational with large denominator then $f(x)$ is likely to have an enormous denominator: more or less the d -th power of the denominator of x where d is the degree of f .

An interesting solution would be to use high precision floating point arithmetic. If variable precision floating point numbers are available then it would be sufficient to evaluate f approximately to obtain values with at least $m + 2$ bits of precision where $m = \log_2(N)$ the logarithm of the refinement factor. If the relative errors in $f(x_{lo})$ and $f(x_{hi})$ are at most 2^{-m-2} then step (2) in `RefineIntervalByFactor` will compute the index κ differing by at most 1 from the value one would have obtained using exact rational arithmetic.

As presented here QIR may calculate an interval much narrower than the specified limit. This is a (minor) defect because it means the last two evaluations of f will be at rational numbers with

needlessly large denominators (with up to twice as many digits as truly required), and thus costing more than is strictly necessary. In practice one can simply reduce the value of the refinement factor N in the final iteration of `RefineInterval` to avoid the overshoot.

5 Experimental Results

We present here a small selection of experimental results. The computer used was a Macintosh G5 (running MacOS X 10.4) with 2GHz processors, and 3.5G bytes of RAM. The software used was CoCoA 4.6 which relies upon the excellent library GMP (version 4.2, see [3]) for arbitrary precision arithmetic. All timings are measured in seconds.

The real root isolation code in CoCoA is a *simplicistic* implementation of the method described in [1], which is based on Descartes’s Rule of Signs. Their implementation appears to be faster than ours in CoCoA.

We describe briefly the polynomials used for the timing tests. The first two polynomials are adapted from the FRISCO test suite (see [4] and [5]) since CoCoA cannot represent complex numbers. Both polynomials are designed to be challenging for a root isolator, and indeed the real root isolator implemented in CoCoA fails with “*Recursion too deep*”. However $I(0, 1)$ is a valid isolating interval for both polynomials, and it was used as input to QIR. The first polynomial is $f_1 = ((c^2x^2 - 3)^4 + c^4x^{18})(c^2x^2 - 3)$ for $c = 10^{100}$; there are four complex roots quite close to the real positive root, and this delays the onset of quadratic convergence. The second polynomial is $f_2 = x^{50} + (10^{50}x - 1)^3$ and has two complex roots very close to the real positive root again delaying the onset of quadratic convergence: note how quickly 10000 digits are obtained compared to the time needed for 1000 digits.

The third and fourth polynomials have a similar structure; all their roots are real by construction. f_3 has degree 32 and is the minimal polynomial of $\sqrt{176} + \sqrt{190} + \sqrt{195} + \sqrt{398} + \sqrt{1482}$, while f_4 has degree 128 and is the minimal polynomial of $\sqrt{99} + \sqrt{627} + \sqrt{661} + \sqrt{778} + \sqrt{929} + \sqrt{1366} + \sqrt{1992}$. These polynomials are interesting because they both have a root which is “almost integer”: f_3 has a root close to 45, and f_4 has a root close to 10.

In the table below, the column headed “#roots” indicates how many intervals were refined, the column headed “isolate” is the time taken by the root isolator, and the other columns give the times taken to refine all the intervals found by the isolator to a width not greater than ε .

Poly	#roots	isolate	$\varepsilon = 10^{-100}$	$\varepsilon = 10^{-1000}$	$\varepsilon = 10^{-10000}$
f_1	1	fail	0.0	2.4	2.9
f_2	1	fail	0.4	43	47
f_3	32	2.8	0.6	2.4	100
f_4	128	74	10	125	6600

The GMP library offers functions for computing the integer part of roots of big integers (namely `mpz_sqrt` and `mpz_root`); these functions use Newton’s Iteration specialized to the computation of k -th roots. We compared QIR with these functions in the following way. To compute d digits of $\sqrt{5}$ using `mpz_sqrt` we passed as argument the value 5×10^{2d} . Similarly to compute d digits of $\sqrt[3]{3}$ and $\sqrt[5]{2}$ we called `mpz_root` with arguments 3×10^{3d} and 2×10^{5d} respectively. The timings for QIR are for refining the root interval to a width not exceeding 10^{-d} — the initial isolating interval was

$[0, 4]$ in both cases. The square root function `mpz_sqrt` is especially efficient: *e.g.* it computed 10^7 digits of $\sqrt{5}$ in 8.2 seconds while QIR needed 35 seconds to achieve the same accuracy. In contrast, we see that QIR is quite competitive with `mpz_root`; here are the results:

Value approximated	QIR $d = 10^5$	QIR $d = 10^6$	GMP $d = 10^5$	GMP $d = 10^6$
$\sqrt[3]{3}$	0.4	5.1	0.8	17
$\sqrt[5]{2}$	1.1	12	1.3	26

6 Conclusions

QIR is a simple algorithm which combines the speed of Newton's Iteration with the robustness of Bisection. It is easy to implement, and the implementation supplied with CoCoA demonstrates its genuine competitiveness with Newton's Iteration (except perhaps for computing square roots).

References

- [1] F Roullier, P Zimmermann "Efficient Isolation of Polynomial Real Roots" *INRIA Rapport de Recherche 4113*, Feb 2001.
- [2] The CoCoA web site <http://cocoa.dima.unige.it/>
- [3] The GMP web site <http://www.swox.com/gmp/>
- [4] <http://www-sop.inria.fr/saga/POL/>
- [5] <http://www.dm.unipi.it/cluster-pages/mpsolve/bench.htm>
- [6] M Kerber, M Sagraloff *Efficient Real Root Approximation*, Proc. ISSAC 2011