Journal of Symbolic Computation 73 (2016) 46–86

Contents lists available at ScienceDirect

Journal of Symbolic Computation

www.elsevier.com/locate/jsc

Computing real roots of real polynomials

Michael Sagraloff [1], Kurt Mehlhorn [1]

MPI for Informatics, Campus E1 4, 66123 Saarbrücken, Germany

ARTICLE INFO

Article history:
Received 4 July 2014
Accepted 9 March 2015
Available online 16 March 2015

Keywords:
Root finding
Root isolation
Root refinement
Approximate arithmetic
Certified computation
Complexity analysis

ABSTRACT

Computing the roots of a univariate polynomial is a fundamental and long-studied problem of computational algebra with applications in mathematics, engineering, computer science, and the natural sciences. For isolating as well as for approximating all complex roots, the best algorithm known is based on an almost optimal method for approximate polynomial factorization, introduced by Pan in 2002. Pan's factorization algorithm goes back to the splitting circle method from Schönhage in 1982. The main drawbacks of Pan's method are that it is quite involved[2] and that all roots have to be computed at the same time. For the important special case, where only the real roots have to be computed, much simpler methods are used in practice; however, they considerably lag behind Pan's method with respect to complexity.

In this paper, we resolve this discrepancy by introducing a hybrid of the Descartes method and Newton iteration, denoted ANewDsc, which is simpler than Pan's method, but achieves a run-time comparable to it. Our algorithm computes isolating intervals for the real roots of any real square-free polynomial, given by an oracle that provides arbitrary good approximations of the polynomial's coefficients. ANewDsc can also be used to only isolate the roots in a given interval and to refine the isolating intervals to an arbitrary small size; it achieves near optimal complexity for the latter task.

E-mail addresses: msagralo@mpi-inf.mpg.de (M. Sagraloff), mehlhorn@mpi-inf.mpg.de (K. Mehlhorn).

[1] The author ordering deviates from the default alphabetic ordering used in Theoretical Computer Science, because the first author contributed significantly more to the paper than the second author.

[2] In Victor Pan's own words: "Our algorithms are quite involved, and their implementation would require a non-trivial work, incorporating numerous known implementation techniques and tricks". In fact, we are not aware of any implementation of Pan's method.

http://dx.doi.org/10.1016/j.jsc.2015.03.004
0747-7171/

1. Introduction

Computing the roots of a univariate polynomial is a fundamental problem in computational algebra. Many problems from mathematics, engineering, computer science, and the natural sciences can be reduced to solving a system of polynomial equations, which in turn reduces to solving a polynomial equation in one variable by means of elimination techniques such as resultants or Gröbner Bases. Hence, it is not surprising that this problem has been extensively studied and that numerous approaches have been developed; see McNamee (2002, 2007), McNamee and Pan (2013), Pan (1997) for an extensive (historical) treatment. Finding all roots of a polynomial and the approximate factorization of a polynomial into linear factors are closely related. The most efficient algorithm for approximate factorization is due to Pan (2002); it is based on the splitting circle method of Schönhage (1982) and considerably refines it. From an approximate factorization, one can derive arbitrary good approximations of all complex roots as well as corresponding isolating disks; e.g. see Emiris et al. (2014), Mehlhorn et al. (2015). The main drawbacks of Pan's algorithm are that it is quite involved (see Footnote 2) and that it necessarily computes all roots, i.e., cannot be used to only isolate the roots in a given region. It has not yet been implemented. In contrast, simpler approaches, namely Aberth's, Weierstrass–Durand–Kerner's and QR algorithms, found their way into popular packages such as MPSolve (Bini and Fiorentino, 2000) or `eigensolve` (Fortune, 2002), although their excellent empirical behavior has never been entirely verified in theory.

In parallel, there is steady ongoing research on the development of dedicated real roots solvers that also allow to search for the roots only in a given interval. Several methods (e.g. Sturm method, Descartes method, continued fraction method, Bolzano method) have been proposed, and there exist many corresponding implementations in computer algebra systems. With respect to computational complexity, all of these methods lag considerably behind the splitting circle approach. *In this paper, we resolve this discrepancy by introducing a hybrid of the Descartes method and Newton iteration, denoted* ANewDsc *(read approximate-arithmetic-Newton-Descartes). Our algorithm is simpler than Pan's algorithm, is already implemented with very promising results for polynomials with integer coefficients (Kobel et al.), and has a complexity comparable to that of Pan's method.*

1.1. Algorithm and results

Before discussing the related work in more detail, we first outline our algorithm and provide the main results. Given a square-free univariate polynomial P with real coefficients, the goal is to compute disjoint intervals on the real line such that all real roots are contained in the union of the intervals and each interval contains exactly one real root. The Descartes or Vincent–Collins–Akritas[3] method is a simple and popular algorithm for real root isolation. It starts with an open interval guaranteed to contain all real roots and repeatedly subdivides the interval into two open intervals and a split point. The split point is a root if and only if the polynomial evaluates to zero at the split point. For any interval I, Descartes' rule of signs (see Section 2.3) allows one to compute an integer v_I, which bounds the number m_I of real roots in I and is equal to m_I, if $v_I \leq 1$. The method discards intervals I with $v_I = 0$, outputs intervals I with $v_I = 1$ as isolating intervals for the unique real root contained in them, and splits intervals I with $v_I \geq 2$ further. The procedure is guaranteed to terminate for square-free polynomials, as $v_I = 0$, if the circumcircle of I (= the one-circle region of I) contains no root of p, and $v_I = 1$, if the union of the circumcircles of the two equilateral triangles with side I (the two-circle region of I) contains exactly one root of I, see Fig. 1 on page 57.

The advantages of the Descartes method are its simplicity and the fact that it applies to polynomials with real coefficients. The latter has to be taken with a grain of salt. The method uses the four basic arithmetic operations (requiring only divisions by two) and the sign-test for numbers in the field of coefficients. In particular, if the input polynomial has integer or rational coefficients, the

[3] Descartes did not formulate an algorithm for isolating the real roots of a polynomial but (only) a method for bounding the number of positive real roots of a univariate polynomial (Descartes' rule of signs). Collins and Akritas (1976) based on ideas going back to Vincent formulated a bisection algorithm based on Descartes' rule of signs.