Journal of Symbolic Computation 44 (2009) 1280–1284

Contents lists available at ScienceDirect

Journal of Symbolic Computation

journal homepage: www.elsevier.com/locate/jsc

The number of roots of a lacunary bivariate polynomial on a line

Martín Avendaño

Departamento de Matemática, Facultad de Ciencias Exactas y Naturales, Universidad de Buenos Aires, 1428 Buenos Aires, Argentina

ARTICLE INFO

Article history:
Received 11 November 2007
Accepted 22 February 2008
Available online 13 February 2009

Keywords:
Descartes' rule of signs
Fewnomials
Factorization of polynomials

ABSTRACT

We prove that a polynomial $f \in \mathbb{R}[x, y]$ with t non-zero terms, restricted to a real line $y = ax + b$, either has at most $6t - 4$ zeros or vanishes over the whole line. As a consequence, we derive an alternative algorithm for deciding whether a linear polynomial $y - ax - b \in K[x, y]$ divides a lacunary polynomial $f \in K[x, y]$, where K is a real number field. The number of bit operations performed by the algorithm is polynomial in the number of non-zero terms of f, in the logarithm of the degree of f, in the degree of the extension $K/\mathbb{Q}$ and in the logarithmic height of a, b and f.

1. Introduction

The famous Descartes' rule of signs, 1641, establishes that the number of positive real roots of a polynomial $f \in \mathbb{R}[x]$, counted with multiplicities, is bounded by the number of changes of signs in its ordered vector of coefficients, skipping the zeros. As a direct consequence, the number of real roots of f is bounded by $2t - 1$, where t is its number of non-zero terms (here all roots are counted with multiplicities, except 0 which is counted at most once).

There are not yet natural generalizations of Descartes' rule of signs for the multivariate setting, but a lot of work has been and is being done for estimating the number of real isolated or non-degenerate roots (that is, where the Jacobian does not vanish, thus implying that the root is isolated) of multivariate square systems of real polynomials, in terms of the number of variables and non-zero terms.

The main result in this direction is due to Khovanskii (1991). A simple version of it implies that a square system of n real polynomial equations in n indeterminates, involving a total t non-zero terms, has at most $(n + 1)^t 2^{t(t-1)/2}$ non-degenerate roots in the positive orthant. Further improvements of Khovanskii's result have been obtained by F. Bihan and F. Sottile; however the exponential dependence on the number of non-zero terms t cannot be avoided yet (Bihan, 2007; Bates et al., 2007).

E-mail address: avendano@math.tamu.edu.

0747-7171/$ – see front matter © 2009 Elsevier Ltd. All rights reserved.
doi:10.1016/j.jsc.2008.02.016

In Li et al. (2003), Li, Rojas and Wang (see also Perruci (2005)) studied particular cases of bivariate square systems and showed that the number of common isolated or non-degenerate roots of a trinomial and a polynomial with at most t non-zero terms, $t \geq 3$, is bounded above by $2^t - 2$.

Furthermore, Koushnirenko's Conjecture, formulated in the mid-1970s (which says that a square system of n real polynomial equations in n indeterminates such that the k-th polynomial has t_k non-zero terms should have at most $(t_1 - 1) \cdots (t_n - 1)$ non-degenerate roots in the positive orthant), turned out to be false, by the counter-example provided by B. Haas in 2002 for a system of two trinomials in two variables (Haas, 2002).

The main result of this article is a refinement of the result of Li et al. (2003) for the particular case where the trinomial is a linear polynomial. Without loss of generality we can assume that the linear polynomial is of the form $y - ax - b$ and we thus study the possible number of real roots of a bivariate polynomial on a line $y = ax + b$:

Theorem 1. *Let $f = \sum_{i=1}^{t} a_i x^{\alpha_i} y^{\beta_i} \in \mathbb{R}[x, y]$ be a polynomial with at most t non-zero terms, and let $a, b \in \mathbb{R}$. Set $g(x) = f(x, ax + b)$. Then either $g \equiv 0$ or g has at most $6t - 4$ real roots, counted with multiplicities except for the possible roots 0 and $-b/a$ that are counted at most once. Moreover, if $g \not\equiv 0$ and $b/a > 0$ (resp. $b/a < 0$), then g has at most $2t - 2$ real roots counted with multiplicities in each of the three intervals $(-\infty, -b/a)$, $(-b/a, 0)$ and $(0, +\infty)$ (resp. $(-\infty, 0)$, $(0, -b/a)$ and $(-b/a, +\infty)$).*

The assumption that f has integer exponents is strong but necessary in order to obtain the stated bounds. The following system (which is obtained from Gomez et al. (2007, Thm.1) by the monomial change of variables $x \leftarrow y^2$ and $y \leftarrow x^6 y$):

$$\begin{cases} y + \frac{44}{31}x - 1 = 0 \\ x^7 + \frac{44}{31}x^{15/4}y^{1/2} - x^{47/12}y^{1/6} + 1936254x^{-133/12}y^{133/6} = 0 \end{cases}$$

has exactly 7 roots in $\mathbb{R}_+^2$. This example shows a 4-nomial with *rational* exponents having $7 > 2 \cdot 4 - 2$ roots on the line $y = -\frac{44}{31}x + 1$ with $0 < x < \frac{31}{44}$.

As a consequence of our result we derive an alternative algorithm for checking whether a given linear form $y - ax - b$ divides a polynomial f in $K[x, y]$, where K is a real number field. The number of bit operations performed by the algorithm is polynomial in the degree $[K : \mathbb{Q}]$ of the field extension, in the number t of non-zero terms of f, in the logarithm of the degree of f and in the logarithmic height of a, b and f. We recall that the (non-logarithmic) absolute height is an extension of the notion of absolute value of rational integers to the algebraic numbers (Hindry and Silverman, 2000, Part B).

The first algorithm for this purpose can be deduced from a more general result by Kaltofen and Koiran (2005). They presented a polynomial-time algorithm for computing all linear factors of a lacunary bivariate polynomial. This result has been further generalized in Kaltofen and Koiran (2006) and Avendaño et al. (2007) to an algorithm that computes all the small degree factors of bivariate and multivariate lacunary polynomials. All these algorithms use a version of the "gap theorem" introduced by Cucker et al. (1999). Here, we reduce the problem to the univariate case by considering specializations $f(x, x^n)$ for small values of n.

2. Changes of signs

Definition 2. Let $f \in \mathbb{R}[x]$ be a non-zero polynomial. We denote by $V(f)$ the number of changes of signs in the sequence of coefficients of f skipping the zeros. We also set $V(0) = -2$.

Remark 3. Let $f \in \mathbb{R}[x]$. Then $V(kf(rx)) = V(f)$ for all $k \neq 0$ and $r > 0$.

Theorem 4 (*Descartes' Rule of Signs*). *Let $f \in \mathbb{R}[x]$ be a non-zero polynomial. Then f has at most $V(f)$ positive roots counted with multiplicities.*

Let us point out here that this famous theorem is a consequence of the following fact: if $f \in \mathbb{R}[x]$ is a non-zero polynomial and $r > 0$, then $V((x - r)f) \geq V(f) + 1$, i.e. if we add a positive root to a polynomial, then its number of changes of signs increases by at least 1. The next lemma gives an analogue of this property for negative roots.

Lemma 5. *Let $f \in \mathbb{R}[x]$ and let $r > 0$. Then $V((x+r)f) \leq V(f)$.*

Proof. By Remark 3 we have $V((x+r)f) = V((rx+r)f(rx)) = V((x+1)f(rx))$ and $V(f(rx)) = V(f)$. Therefore we only need to consider the case $r = 1$. We proceed by induction on the number t of non-zero terms of f. The theorem is trivial for $t = 0$ and $t = 1$. Now let us suppose that it holds for all $t \leq n$. Let $f \in \mathbb{R}[x]$ with $n+1$ non-zero monomials.

$$f = \sum_{i=1}^{n+1} a_i x^{\alpha_i} \quad \text{where } a_i \neq 0,\ \forall i, \quad \text{and} \quad 0 \leq \alpha_1 < \alpha_2 < \cdots < \alpha_{n+1} = d = \deg(f).$$

Let $g = \sum_{i=1}^{n} a_i x^{\alpha_i}$. By the inductive hypothesis we have $V((x+1)g) \leq V(g)$. First, we consider the case $\alpha_n < d-1$, i.e. when the terms of $(x+1)g$ do not overlap with those of $a_{n+1}x^d(x+1)$. There are two possibilities: if $a_n a_{n+1} > 0$, then $V((x+1)f) = V((x+1)g) \leq V(g) = V(f)$, and if $a_n a_{n+1} < 0$, then $V((x+1)f) = V((x+1)g) + 1 \leq V(g) + 1 = V(f)$. In both cases we have $V((x+1)f) \leq V(f)$. Now the only remaining case is $\alpha_n = d-1$. Here $(x+1)f$ and $(x+1)g$ only differ in their terms of degree d and $d+1$, as shown in the following table.

	x^d	x^{d+1}
$(x+1)g$	a_n	0
$(x+1)f$	$a_n + a_{n+1}$	a_{n+1}

If $a_n a_{n+1} > 0$, then $V(f) = V(g)$, and according to the table, we have $V((x+1)f) = V((x+1)g)$. Therefore $V((x+1)f) \leq V(f)$. On the other hand, if $a_n a_{n+1} < 0$, then $V(f) = V(g) + 1$, but we have three different possibilities for the table, depending on whether $|a_n|$ is greater than, equal to or less than $|a_{n+1}|$. Set s for the sign of a_n.

	x^d	x^{d+1}
$(x+1)g$	s	0
$(x+1)f$	s	$-s$

Case $|a_n| > |a_{n+1}|$

	x^d	x^{d+1}
$(x+1)g$	s	0
$(x+1)f$	0	$-s$

Case $|a_n| = |a_{n+1}|$

	x^d	x^{d+1}
$(x+1)g$	s	0
$(x+1)f$	$-s$	$-s$

Case $|a_n| < |a_{n+1}|$

The tables show that $V((x+1)f) \leq V((x+1)g) + 1$ for each of the three cases. Using the inductive hypothesis and $V(f) = V(g) + 1$, we conclude that $V((x+1)f) \leq V(f)$. □

Remark 6. Let $f, g \in \mathbb{R}[x]$ and suppose that g has t terms. Then $V(f+g) \leq V(f) + 2t$.

Note that the value of $V(0)$ is not relevant for Theorem 4 and Lemma 5. The only reason for defining $V(0) = -2$ is the previous remark (in the case $f = 0$ and $t = 1$).

Proposition 7. *Let $f \in \mathbb{R}[x, y]$ with t non-zero terms. Let $p = (x+r_1)\cdots(x+r_n) \in \mathbb{R}[x]$ where $r_i > 0$ for all $i = 1, \ldots, n$. Then*

$$V(f(x, p(x))) \leq 2t - 2.$$

Proof. We write $f = \sum_{i=1}^{n} a_i(x) y^{\alpha_i}$, where $0 \leq \alpha_1 < \cdots < \alpha_n$ and $a_i(x) \in \mathbb{R}[x]$, and we set $t_i > 0$, the number of non-zero terms of a_i. It is clear that $t = t_1 + \cdots + t_n$.

We define $f_k = \sum_{i=k}^{n} a_i(x) y^{\alpha_i - \alpha_k}$ for $k = 1, \ldots, n$ and $f_{n+1} = 0$. Lemma 5 and Remark 6 imply that the polynomials f_k satisfy:

- $f_{n+1} = 0 \Rightarrow V(f_{n+1}(x, p(x))) = -2$.
- $f_k = y^{\alpha_{k+1} - \alpha_k} f_{k+1} + a_k(x) \Rightarrow f_k(x, p(x)) = p(x)^{\alpha_{k+1} - \alpha_k} f_{k+1}(x, p(x)) + a_k(x) \Rightarrow V(f_k(x, p(x))) \leq V(f_{k+1}(x, p(x))) + 2t_k$.
- $f = y^{\alpha_1} f_1 \Rightarrow f(x, p(x)) = p(x)^{\alpha_1} f_1(x, p(x)) \Rightarrow V(f(x, p(x))) \leq V(f_1(x, p(x)))$.

Thus, we conclude that $V(f(x, p(x))) \leq -2 + 2(t_1 + \cdots + t_n) = 2t - 2$. □

Now we have all the tools needed to prove Theorem 1:

Proof. If $a = 0$ or $b = 0$, then $g \in \mathbb{R}[x]$ is a polynomial with at most t non-zero terms. Descartes' rule of signs implies that either $g \equiv 0$ or g has at most $2t-1 \leq 6t-4$ real roots (counted with multiplicities except for the possible root 0). In the case $a \neq 0$ and $b \neq 0$, the real roots of $f(x, ax+b)$ correspond one to one to the roots of $f(bx/a, b(x+1)) = \widehat{f}(x, x+1)$, where $\widehat{f} = \sum_{i=1}^{t} a_i a^{-\alpha_i} b^{\alpha_i+\beta_i} x^{\alpha_i} y^{\beta_i}$. Since this bijection preserves the multiplicity of the roots and maps the possible roots 0 and $-b/a$ of g to the roots 0 and -1 of $\widehat{f}(x, x+1)$, we only need to consider the case $a = b = 1$, i.e. $g = f(x, x+1)$. Suppose that $g \not\equiv 0$. Descartes' rule of signs and Proposition 7 imply that the number of positive roots of g counted with multiplicities is at most $2t-2$. On the other hand, the roots of g in $(-\infty, -1)$ correspond to the positive roots of $0 \not\equiv g(-1-x) = f(-1-x, -x) = f_1(x, x+1)$, where $f_1 = \sum_{i=1}^{t} a_i (-1)^{\alpha_i+\beta_i} x^{\beta_i} y^{\alpha_i}$. Therefore the number of roots (with multiplicities) of g in $(-\infty, -1)$ is also bounded above by $2t-2$. Finally, the roots of g in $(-1, 0)$ correspond to the positive roots of

$$0 \not\equiv (x+1)^{\deg(g)} g\left(\frac{-x}{x+1}\right) = (x+1)^{\deg(g)} f\left(\frac{-x}{x+1}, \frac{1}{x+1}\right) = f_2(x, x+1)$$

where $f_2 = \sum_{i=1}^{t} a_i (-1)^{\alpha_i} x^{\alpha_i} y^{\deg(g)-\alpha_i-\beta_i}$. Therefore there are at most $2t-2$ of such roots. Taking into account the possible roots 0 and -1, each one counted at most once, we conclude that g has at most $6t-4$ real roots. □

3. Checking linear factors of a bivariate polynomial

Proposition 8. *Let $f = \sum_{i=1}^{t} a_i x^{\alpha_i} y^{\beta_i} \in \mathbb{R}[x, y]$. Let $a, b \in \mathbb{R}$ be such that $b \neq |1-a|$. Then $y - ax - b \,|\, f \Leftrightarrow x^n - ax - b \,|\, f(x, x^n)$ for at least $6t-3$ odd integers $n \geq 3$.*

Proof. ($\Leftarrow$) : Suppose that $3 \leq n_1 < n_2 < \cdots < n_{6t-3}$ are $6t-3$ odd numbers such that $x^n - ax - b \,|\, f(x, x^n)$. Let $w_i \in \mathbb{R}$ be a root of $x^{n_i} - ax - b$ for each $1 \leq i \leq 6t-3$. Then $f(w_i, aw_i + b) = f(w_i, w_i^{n_i}) = 0$ for all $1 \leq i \leq 6t-3$. This means that $f(x, ax+b)$ has at least $6t-3$ real roots. Applying Theorem 1 we conclude that $f(x, ax+b) \equiv 0$, or simply $y - ax - b \,|\, f$. It only remains to prove that $w_i \neq w_j$ for all $i \neq j$. Actually, if $x^{n_i} - ax - b$ and $x^{n_j} - ax - b$ have a common root $w = w_i = w_j \in \mathbb{R}$, then $w^{n_i - n_j} = 1$ and therefore $w = \pm 1$. This would imply that $0 = w^{n_i} - aw - b = -b \pm (1-a)$, in contradiction with the hypothesis $b \neq |1-a|$. □

Note that if $(a, b) \neq (0, 1)$, then either $b \neq |1-a|$ or $b \neq |1+a|$.

Algorithm TEST

Input: A polynomial $f = \sum_{i=1}^{t} a_i x^{\alpha_i} y^{\beta_i} \in K[x, y]$ with t terms, encoded as a list of vectors $(a_i, \alpha_i, \beta_i) \in K \times \mathbb{N}_0 \times \mathbb{N}_0$ representing the monomials of f, and two numbers $a, b \in K$.

Output: True or False depending on whether $y - ax - b \,|\, f(x, y)$ or not.

Step 1: If $(a, b) = (0, 1)$, compute $f(x, 1)$. If this polynomial is zero, return True. Otherwise return False.

Step 2: If $b = |1-a|$ then replace f by $f(-x, y)$ and a by $-a$.

Step 3: For $n = 3, 5, 7, \ldots, 12t-5$ do

Step 3.1: If $f(x, x^n) \neq 0$ then

Step 3.1.1: Compute all the irreducible factors (with multiplicities) of $x^n - ax - b$ in $K[x]$ using a univariate dense factorization algorithm.

Step 3.1.2: Compute all the irreducible factors (with multiplicities) of $f(x, x^n)$ in $K[x]$ with degree $\leq n$ using a univariate lacunary factorization algorithm.

Step 3.1.3: If there is an irreducible factor in the first list that either does not belong to the second list or belongs but with less multiplicity, then return False.

Step 4: Return True.

The correctness of the algorithm is a consequence of Proposition 8. In order to estimate its complexity, we first state the following two famous results on the factorization of polynomials of univariate polynomials. See Hindry and Silverman (2000, Part B) for a definition of the absolute height.

Dense Factorization. Let K be a number field. Given $f \in K[x]$ of degree d and absolute height H, it is possible to compute all its irreducible factors in $K[x]$ with multiplicities in $[d\,[K:\mathbb{Q}]\,\log H]^{O(1)}$ bit operations (see Lenstra et al. (1982) for the rational case and (Landau, 1985) for the general case).

Lacunary Factorization. Let K be a number field. Given $f \in K[x]$ a polynomial of degree d, with at most t monomials and absolute height H, it is possible to find all its irreducible factors (with their corresponding multiplicities) in $K[x]$ of degree bounded by s in $[t\,s\,[K:\mathbb{Q}]\,\log d\,\log H]^{O(1)}$ bit operations (see Lenstra (1999)).

The complexity of the algorithm TEST is clearly dominated by its main loop (step 3), where it performs $6t-3$ calls to the dense and lacunary factorization algorithms in order to factorize $x^n - ax - b$ completely and find all the factors of degree bounded by n of $f(x, x^n)$. We have that $\deg(x^n - ax - b) = n \leq 12t - 5$ and $H(x^n - ax - b) \leq H(a)H(b)$; therefore the step 3.1.1 requires at most $[(6t-3)\,(12t-5)\,[K:\mathbb{Q}]\,\log(H(a)\,H(b))]^{O(1)}$ bit operations. On the other hand, we have that $f(x, x^n)$ is a lacunary polynomial with at most t non-zero terms, of degree bounded by $nd \leq (12t-5)d$ and absolute height bounded by $(2H(f))^t$ because the coefficients of $f(x, x^n)$ are sums of at most t coefficients of f. Thus, step 3.1.2 requires no more than $[(6t-3)\,t\,(12t-5)[K:\mathbb{Q}]\,\log(d(12t-5))\,\log(2H(f))^t]^{O(1)}$ bit operations. This proves that the total number of bit operations performed by the algorithm is polynomial in t, $\log(d)$, $[K:\mathbb{Q}]$ and $\log(H(a)H(b)H(f))$.

A natural question is whether it is possible not only to check linear factors but also to find them by means of such a simple algorithm. We are currently working on this and hope that this will be the subject of a forthcoming paper.

Acknowledgements

The author thanks Teresa Krick, Daniel Perruci, J. Maurice Rojas and Frank Sottile for reading an earlier version of this paper and for several useful discussions on fewnomial systems.

The author's research supported by grants PICT 33671/05 and UBACYT X-112, Argentina.

References

Avendaño, M., Krick, T., Sombra, M., 2007. Factoring bivariate lacunary polynomials. Journal of Complexity 23, 193–216.

Bates, D.J., Bihan, F., Sottile, F., 2007. Bounds on the number of real solutions to polynomial equations. International Mathematics Research Notices (114), 114–117.

Bihan, F., 2007. Sottile: New fewnomial upper bounds from Gale dual polynomial systems. Moscow Mathematics Journal 7 (3), 387–407.

Cucker, F., Koiran, P., Smale, S., 1999. A polynomial time algorithm for diophantine equations in one variable. JSC 27 (1), 21–29.

Gomez, J., Niles, A., Rojas, J.M., 2007. New complexity bounds for certain real fewnomial zero sets (extended abstract). In: Effective Methods in Algebraic Geometry MEGA'07.

Haas, B., 2002. A simple counter-example to Koushnirenko's conjecture. Beiträge zur Algebra und Geometrie 43 (1), 1–8.

Hindry, M., Silverman, J.H., 2000. Diophantine Geometry: An Introduction. Springer-Verlag, New York.

Kaltofen, E., Koiran, P., On the complexity of factoring bivariate supersparse (lacunary) polynomials. In: ISSAC'05 Proc. 2005 Internat. Symp. Symbolic Algebraic Comput. pp. 208–215.

Kaltofen, E., Koiran, P., Finding small degree factors of multivariate supersparse (lacunary) polynomials over algebraic number fields. In: ISSAC'06 Proc. 2006 Internat. Symp. Symbolic Algebraic Comput. pp. 162–168.

Khovanskii, A., 1991. Fewnomials. AMS press, Providence, Rhode Island.

Landau, S., 1985. Factoring polynomials over algebraic number fields. SIAM Journal on Computing 14, 184–195.

Lenstra, A.K., Lenstra, H.W., Lovasz, L., 1982. Factoring polynomials with rational coefficients. Mathematishe Annalen 261, 515–534.

Lenstra, H.W., 1999. Finding small degree factors of lacunary polynomials. Number Theory in Progress 1, 267–276.

Li, T.Y., Rojas, J.M., Wang, X., 2003. Counting real connected components of trinomial curves intersections and m-nomial hypersurfaces. Discrete and Computational Geometry 30 (3), 379–414.

Perruci, D., 2005. Some bounds for the number of connected components of real zero sets of sparse polynomials. Discrete and Computational Geometry 34 (3), 475–495.