Information Processing Letters 110 (2010) 898–901

Contents lists available at ScienceDirect

Information Processing Letters

www.elsevier.com/locate/ipl

A local decision test for sparse polynomials

Elena Grigorescu [a,1,*], Kyomin Jung [b,2], Ronitt Rubinfeld [a,c,3]

[a] Massachusetts Institute of Technology, USA
[b] KAIST, Republic of Korea
[c] University of Tel Aviv, Israel

ARTICLE INFO

Article history:
Received 17 September 2009
Received in revised form 6 July 2010
Accepted 9 July 2010
Available online 16 July 2010
Communicated by J. Torán

Keywords:
Randomized algorithms
Sparsity tests
Multivariate polynomials

ABSTRACT

An ℓ-*sparse* (multivariate) polynomial is a polynomial containing at most ℓ-monomials in its explicit description. We assume that a polynomial is implicitly represented as a black-box: on an input query from the domain, the black-box replies with the evaluation of the polynomial at that input. We provide an efficient, randomized algorithm, that can decide whether a polynomial $f : \mathbb{F}_q^n \to \mathbb{F}_q$ given as a black-box is ℓ-sparse or not, provided that q is large compared to the polynomial's total degree. The algorithm makes only $O(\ell)$ queries, which is independent of the domain size. The running time of our algorithm (in the bit-complexity model) is $poly(n, \log d, \ell)$, where d is an upper bound on the degree of each variable. Existing interpolation algorithms for polynomials in the same model run in time $poly(n, d, \ell)$. We provide a similar test for polynomials with integer coefficients.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

When dealing with massive data sets that encode information in a redundant manner, a typical goal is to efficiently retrieve qualitative properties of the encoded data. This task immediately excludes the possibility of reading the entire encoding, and brings up questions regarding the type of properties that can be decided from a very restricted view of the data. It is common to model such setting as a black-box interaction. Namely, the algorithm sends input-queries to a black-box that computes a function $f : D \to R$ (for some finite or infinite domain D and range R) and receives back the evaluations of f at those inputs. An algorithm is called *local* if the number of queries it makes is independent of the size of the domain.

In our model, the black-box encodes n-variate polynomials over $\mathbb{F}_q$ (where q is a prime power) or over $\mathbb{Z}$, and we are interested in deciding whether f is ℓ-sparse, or it is at least $(\ell+1)$-sparse in its explicit description as a sum of monomials. We propose a randomized, polynomial time algorithm that has the additional feature of locality, and we call such a test a *local decision test*.

The model considered in this work is related to the Property Testing model in which a local test distinguishes between objects satisfying a property from objects that are "far" from that property, under appropriate notions of distance. (For surveys on the area of Property Testing, see for example [11,20,21,7].) Note that unlike the Property Testing model, we are required to output "fail" even if a non-sparse polynomial is very "close" to some sparse polynomial (i.e. agrees on almost all inputs). From this point of view, our model is stronger, and thus certain tasks may require more computation.

1.1. Our results and techniques

We exhibit a local decision test which, on input ℓ, given query-access to a polynomial $f : \mathbb{F}_q^n \to \mathbb{F}_q$, of known degree d per individual variable, decides whether f is ℓ-sparse. We assume that d and n are small compared to

* Corresponding author.
E-mail addresses: elena_g@mit.edu (E. Grigorescu), kyomin@kaist.edu (K. Jung), ronitt@csail.mit.edu (R. Rubinfeld).

[1] Supported by NSF award CCF-0829672.
[2] Supported by the Engineering Research Center of Excellence Program of Korea MEST/National Research Foundation of Korea (NRF) (Grant 2009-0063242).
[3] Supported by NSF grants 0732334 and 0728645, Marie Curie Reintegration grant PIRG03-GA-2008-231077, and the Israel Science Foundation grant nos. 1147/09 and 1675/09.

0020-0190/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.ipl.2010.07.012

the field's size q. Alternately, this can be interpreted as allowing d and n to be large, but also allowing queries over some extension field of $\mathbb{F}_q$. The locality of the test is $2\ell+1$, which is independent of n or d. We remark that our query complexity is the same as in the interpolation algorithm of Ben-Or and Tiwari [2]. However, since testing for sparsity is easier a task than interpolation, we achieve a better dependence on the degree d: the running time of our algorithm is $poly(n, \log d, \ell)$, while the current interpolation algorithms (in the same black-box model as ours) run in time $poly(n, d, \ell)$.

Our proof relies on a common criteria for sparsity, based on so-called Wronskian matrices, often employed in the sparse interpolation literature [2,12,15]. We closely follow Ben-Or and Tiwari's [2] approach, associating a Wronskian matrix H to the polynomial f. The matrix H exhibits two nice properties: (1) its entries are evaluation points of f, and (2) for some carefully chosen evaluation points, its determinant is 0 if f is sparse, and non-zero otherwise. In [2] the matrix H is further manipulated in order to obtain the explicit monomial degrees of f.

Our main observation is that one can treat $\det H$ as a polynomial function which can be evaluated over the entire field. Using the sparsity criteria based on the value of $\det H$, our test simply chooses uniformly random evaluation points from the field and outputs a positive answer only when the determinant evaluates to 0.

In addition, we show a similar test for $f : \mathbb{Z}^n \to \mathbb{Z}$, when a bound L on the coefficients is known.

Its running time is $poly(n, \log d, \ell, \log L)$, and its query complexity is again $2\ell + 1$.

1.2. Related work

In the algebraic context, a property class is usually a collection of polynomials sharing a common feature, as for example, linearity, symmetry, homogeneity, or low degree. In these cases a natural test simply verifies the defining condition of the property, under a uniformly random choice of the evaluation points. Consider for instance the class of multivariate polynomials $\{f : \mathbb{F}_q^n \to \mathbb{F}_q \mid f(x_1, x_2, x_3, \ldots, x_n) \equiv f(x_2, x_1, x_3, \ldots, x_n)\}$, i.e. polynomials that are symmetric in the first two variables. A quick local decision test simply picks a n-tuple $(r_1, r_2, r_3, \ldots, r_n) \in \mathbb{F}_q^n$ and checks if $f(r_1, r_2, r_3, \ldots, r_n) = f(r_2, r_1, r_3, \ldots, r_n)$.

While the algebraic properties mentioned above have natural explicit tests, in the case of sparsity it is not a priori clear what a good test might look like. In fact the question has been previously explored in the Property Testing context by [5] and [6] who employed a heavy Fourier analytic machinery and ideas drawn from the area of "testing by implicit learning" both to describe the tests and to analyze them. In contrast, the decision test we propose has the merit of being rather clean and elegant: it simply computes the determinant of a certain matrix consisting of entries that are evaluations of f. In addition, while our model seems harder to test than their Property Testing model, the two running times as well as query complexities are comparable. Note however that in [6] the authors assume that the degree of f is unknown. Formally, for polynomials $f : \mathbb{F}_q^n \to \mathbb{F}_q$ [6] exhibited an algorithm which runs in time $n\,poly(\ell, \frac{1}{\epsilon})$, makes $poly(\ell, \frac{1}{\epsilon})$ queries and decides if f is an ℓ-sparse polynomial or differs in an ϵ fraction of inputs from any ℓ-sparse polynomial. We remark that in their case q is a constant.

The related interpolation problem has been intensely studied under a wide variety of models: deterministic [2, 17] and randomized [19], algebraic (with unit time per operation) [2,17] and bit-probe [19], over the integers [19, 2] as well as over fields [2,13], in the standard polynomial basis or over some non-standard bases [4,13,18,12]. Depending on the focus of the result, different tradeoffs between the number of black-box queries and the running time have been achieved. The initial Zippel's randomized algorithm over the integers runs in time (bit-probe model) $poly(n, d, \ell, \log L)$ (where L is an upper bound on the absolute value of the integer coefficients) and makes $nd\ell$ queries. Ben-Or and Tiwari's [2] deterministic interpolation algorithm runs also in $poly(n, d, \ell, \log L)$ time, and makes $2\ell + 1$ evaluations. Mansour [19], and later Alon and Mansour [1] exhibit a deterministic interpolation algorithm, which runs in time $poly(n, \log d, \ell, \log L)$. However, its query complexity depends on d and their black-box model is different from ours. Namely, even though the polynomial is defined over the integers, the black-box can be queried over the complexes, with some good accuracy parameter. Grigoriev, Karpinski and Singer [14] study the interpolation problem for rational functions and they achieve a number of queries linear in the number of variables, and the usual running time $poly(n, d, \ell, \log L)$. Recently, Garg and Schost [9], propose an interpolation result in the easier, "straight-line program" model, with time complexity $poly(n, \log d, \ell, T)$, where T is the length of the straight-line program.

2. Preliminaries

First, recall a standard definition. A polynomial $f(x_1, \ldots, x_n) = \sum_{i=1}^{k} a_i x_1^{d_{i1}} \ldots x_n^{d_{in}}$, with $a_i \neq 0$ for $i = 1, 2, \ldots, k$ is called *ℓ-sparse* if $k \leqslant \ell$.

We will also make use of the following well-known result.

Lemma 1 *(Schwartz–Zippel). Let $f \in \mathbb{F}(x_1, \ldots, x_n)$ have degree d in each variable, and let $S \subseteq \mathbb{F}$ be a finite set. Then*

$$\Pr_{r_i \in S}\big[f(r_1, r_2, \ldots, r_n) = 0\big] \leqslant \frac{dn}{|S|},$$

where the r_is are uniformly distributed over S.

Finally, we formally define a *local decision test*.

Definition 2 *(Local decision test).* An algorithm $\mathcal{T}$ is a *(s, ϵ, T)-local decision test* for a class $\mathcal{P}$ if, given black-box access to $f \in \mathbb{F}[x_1, \ldots, x_n]$ (denoted $\mathcal{T}^f$),

1. $\mathcal{T}^f$ makes s queries,
2. $\mathcal{T}^f$ runs in time T (in the bit-complexity model),
3. if $f \in \mathcal{P}$ $\Pr[\mathcal{T}^f \text{ accepts}] = 1$,
4. if $f \notin \mathcal{P}$ $\Pr[\mathcal{T}^f \text{ rejects}] \geqslant 1 - \epsilon$.

(1) Choose $u = (u_1, u_2, \ldots, u_n) \in \mathbb{F}_q^n$ uniformly at random.
(2) Let

$$H_{\ell+1}(f,u) = \begin{pmatrix} f(u^0) & f(u^1) & \ldots & f(u^{\ell}) \\ f(u^1) & f(u^2) & \ldots & f(u^{\ell+1}) \\ \ldots & \ldots & \ldots & \ldots \\ f(u^{\ell}) & f(u^{\ell+1}) & \ldots & f(u^{2\ell}) \end{pmatrix},$$

where $u^i = (u_1^i, u_2^i, \ldots, u_n^i)$, and compute $\det H_{\ell+1}(f,u)$.
(3) If $\det H_{\ell+1}(f,u) = 0$ output "f is ℓ-sparse".
(4) Else output "f is not ℓ-sparse".

Fig. 1. Local decision test for sparsity (over a finite field).

3. Local decision tests for sparsity

In this section we state and prove our main theorems. We start with the finite field case and then explain how to modify that test to work over the integers.

Theorem 3. *Let $0 < \epsilon < 1$, ℓ, n, d be positive integers, and $q > \frac{\ell(\ell+1)dn}{\epsilon}$ be a prime power. Then there exists a $(2\ell+1, \epsilon, \tilde{O}((\ell n + \ell^2)\log(\frac{\ell^2 dn}{\epsilon})))$-local decision test[4] for the class of ℓ-sparse polynomials $f \in \mathbb{F}_q[x_1, \ldots, x_n]$ with degrees at most d in each variable.*

Proof. We follow the same initial steps as in [2] and start by introducing some useful notations. For a given n-tuple $u = (u_1, u_2, \ldots, u_n) \in \mathbb{F}_q^n$, and a positive integer i, define the tuple $u^i = (u_1^i, u_2^i, \ldots, u_n^i)$.

Moreover, for a polynomial $f \in F_q[x_1, \ldots, x_n]$, n-tuple $u \in \mathbb{F}_q^n$, and integer $t > 0$, define the Hankel matrix associated to f at evaluation point u by

$$H_t(f,u) = \begin{pmatrix} f(u^0) & f(u^1) & \ldots & f(u^{t-1}) \\ f(u^1) & f(u^2) & \ldots & f(u^t) \\ \ldots & \ldots & \ldots & \ldots \\ f(u^{t-1}) & f(u^t) & \ldots & f(u^{2t-2}) \end{pmatrix}.$$

These objects constitute the main tool in Ben-Or and Tiwari's algorithm. In the interpolation context, the problem of determining the individual monomial degrees reduces to inverting such a Hankel matrix for some very special choice of the initial tuple u. In our case, we observe that most choices of $u \in \mathbb{F}_q^n$ can be used to describe the sparsity of f, by inspecting the determinant of $H_t(f,u)$.

Fig. 1 describes our test in the finite field case.

Analysis. It is clear that the test makes $2\ell+1$ queries. The proof of correctness, as well as the test's error probability, are based on our main observation that one can treat $\det H_{\ell+1}(f,u)$ as a multivariate polynomial in the variables $u_1, u_2, \ldots, u_n$. Since q is assumed to be large, the Schwartz–Zippel Lemma will then imply that the test fails with small probability.

[4] $\tilde{O}(F)$ refers to $F \cdot poly(\log F)$.

The sharp criteria for sparsity from [2] is only stated there for a fixed value of u. However, the result holds in general, for any value of $u \in \mathbb{F}_q^n$.

Lemma 4. *(See [2].) Let $f(x_1, \ldots, x_n) = \sum_{i=1}^{k} a_i M_i(x_1, \ldots, x_n)$, where $a_i \in \mathbb{F}_q$, and $M_i(x_1, \ldots, x_n)$ are the monomials of f. Let $u = (u_1, \ldots, u_n)$.*

1. *If $k < \ell+1$, then $\det H_{\ell+1}(f,u) \equiv 0$ (as a polynomial in $u_1, \ldots, u_n$).*
2. *If $k \geqslant \ell+1$, the following expression holds:*

$$\det H_{\ell+1}(f,u) = \sum_{S \subset [k],\, |S| = \ell+1} \prod_{i \in S} a_i \prod_{i<j,\, i,j \in S} \big(M_i(u) - M_j(u)\big)^2.$$

For the sake of intuition, we mention that the proof of Lemma 4 is based on the fact that $H_t(f,u)$ can be expressed as follows:

$$\begin{pmatrix} 1 & 1 & \ldots & 1 \\ M_1(u) & M_2(u) & \ldots & M_k(u) \\ M_1(u^2) & M_2(u^2) & \ldots & M_k(u^2) \\ \ldots & \ldots & \ldots & \ldots \\ M_1(u^{t-1}) & M_2(u^{t-1}) & \ldots & M_k(u^{t-1}) \end{pmatrix} \times \begin{pmatrix} a_1 & 0 & \ldots & 0 \\ 0 & a_2 & \ldots & 0 \\ 0 & 0 & \ldots & 0 \\ \ldots & \ldots & \ldots & \ldots \\ 0 & 0 & \ldots & a_k \end{pmatrix} \times \begin{pmatrix} 1 & M_1(u) & \ldots & M_1(u^{t-1}) \\ 1 & M_2(u) & \ldots & M_2(u^{t-1}) \\ 1 & M_3(u) & \ldots & M_3(u^{t-1}) \\ \ldots & \ldots & \ldots & \ldots \\ 1 & M_k(u) & \ldots & M_k(u^{t-1}) \end{pmatrix}.$$

From Lemma 4, if f is ℓ-sparse, then our test always outputs the right answer. Now suppose that f is at least $(\ell+1)$-sparse. In that case, $\det H_{\ell+1}(f,u)$ is a polynomial in $u_1, \ldots, u_n$, and each variable has degree at most $2\binom{\ell+1}{2}d$. Our test makes an error only when a uniformly random $u \in \mathbb{F}_q^n$ is a root of the n-variate polynomial $\det H_{\ell+1}(f,u)$. Applying the Schwartz–Zippel Lemma, it follows that this happens with probability at most ϵ, since $q > \frac{2\binom{\ell+1}{2}dn}{\epsilon}$.

We now argue the computation time in the bit-complexity model. Since the evaluations of f are provided by the black-box, the only operations performed by the testing algorithm are during the computation of $\det H_{\ell+1}(f,u)$. In general, computing the determinant of an $\ell \times \ell$ matrix takes time ℓ^{ω} (where $\omega \approx 2.37$ is the exponent of matrix multiplication). However, since $H_{\ell+1}(f,u)$ is a Hankel matrix, its determinant can be computed in time $O(\ell^2)$ (by [16]). The complexity of our test is then $\tilde{O}((\ell n + \ell^2)\log(\frac{\ell^2 dn}{\epsilon}))$. This follows by taking into account the input size of each black-box call and the fact

(1) Choose a prime number $p \geqslant \max\{\frac{\ell(\ell+1)dn}{\epsilon}, L\}$.
(2) For $i = 1, 2, \ldots n$, choose $u_i \in \mathbb{Z}_p$ uniformly at random.
Let $u = (u_1, u_2 \ldots, u_n)$.
(3) Let

$$H_{\ell+1}(f,u) = \begin{pmatrix} f(u^0) & f(u^1) & \ldots & f(u^\ell) \\ f(u^1) & f(u^2) & \ldots & f(u^{\ell+1}) \\ \ldots & \ldots & \ldots & \ldots \\ f(u^\ell) & f(u^{\ell+1}) & \ldots & f(u^{2\ell}) \end{pmatrix} \bmod p,$$

where $u^i = (u_1^i, u_2^i, \ldots, u_n^i)$, and compute $\det H_{\ell+1}(f,u) \bmod p$.
(4) If $\det H_{\ell+1}(f,u) = 0 \bmod p$, output "$f$ is ℓ-sparse".
(5) Else output "f is not ℓ-sparse".

Fig. 2. Local decision test for sparsity (over the integers).

that multiplication and addition take $O(\log q \log\log q)$ bit-operations [10] for rings of size q that support Fast Fourier Transform. □

A similar result holds for polynomials over $\mathbb{Z}$.

Theorem 5. *Let $0 < \epsilon < 1$, and ℓ, n, d, L be positive integers. Then there exists a $(2\ell+1, \epsilon, \tilde{O}((\ell n + \ell^2)\log(L + \frac{\ell^2 dn}{\epsilon})))$-local decision test for the class of ℓ-sparse polynomials $f \in \mathbb{Z}[x_1, \ldots, x_n]$ with degrees at most d in each variable, and whose coefficients are bounded in absolute value by L.*

Proof. The proof of Theorem 5 is very similar to the proof of Theorem 3. Here we pick a prime number p larger than $\max\{\frac{\ell(\ell+1)dn}{\epsilon}, L\}$ and compute modulo p. Fig. 2 describes our local decision test for sparsity of polynomials f over $\mathbb{Z}$.

First notice that Lemma 4 still holds over $\mathbb{Z}$. Since $p > L$ and p is prime, the coefficient of the highest degree monomial (in a lexicographic order of the variables) in $\det H_{\ell+1}(f, (x_1, x_2, \ldots, x_n))$ remains non-zero even after the modulo operation. Hence, if $\det H_{\ell+1}(f, (x_1, \ldots, x_n))$ is a non-zero polynomial over $\mathbb{Z}$, then $\det H_{\ell+1}(f, (x_1, \ldots, x_n)) \bmod p$ is a non-zero polynomial over $\mathbb{Z}_p$. The same argument as before finishes the proof. □

4. Conclusions and open problems

A variety of open questions could be interesting in the context of local decision testing for sparsity. First, our test works in a restrictive setting when the total degree is small, and one could try and remove this relaxation if possible. Furthermore, showing lower bounds on the number of evaluation points necessary to test sparsity is a direction we have not explored.

It would also be interesting to exhibit local decision tests for other problems in which the tests are not obvious from the definition of the property class. In particular, a good such example is the problem of testing "juntas" (polynomials that only depend on a few variables). This question has been recently completely resolved in the Property Testing model [8,3].

Acknowledgements

We thank Jonathan Kelner for pointing us to Ref. [16]. We are grateful to Madhu Sudan and Kevin Matulef for helpful conversations. We thank the anonymous reviewers for useful comments.

References

[1] Noga Alon, Yishay Mansour, Epsilon-discrepancy sets and their application for interpolation of sparse polynomials, Inform. Process. Lett. 54 (6) (1995) 337–342.
[2] Michael Ben-Or, Prasoon Tiwari, A deterministic algorithm for sparse multivariate polynomial interpolation, in: STOC, 1988, pp. 301–309.
[3] Eric Blais, Testing juntas nearly optimally, in: STOC, 2009, pp. 151–158.
[4] Michael Clausen, Andreas W.M. Dress, Johannes Grabmeier, Marek Karpinski, On zero-testing and interpolation of k-sparse multivariate polynomials over finite fields, Theoret. Comput. Sci. 84 (2) (1991) 151–164.
[5] Ilias Diakonikolas, Homin K. Lee, Kevin Matulef, Krzysztof Onak, Ronitt Rubinfeld, Rocco A. Servedio, Andrew Wan, Testing for concise representations, in: FOCS, 2007, pp. 549–558.
[6] Ilias Diakonikolas, Homin K. Lee, Kevin Matulef, Rocco A. Servedio, Andrew Wan, Efficiently testing sparse GF(2) polynomials, in: ICALP, 2008, pp. 502–514.
[7] Eldar Fischer, The art of uninformed decisions: A primer to property testing, Science 75 (2001) 97–126.
[8] Eldar Fischer, Guy Kindler, Dana Ron, Shmuel Safra, Alex Samorodnitsky, Testing juntas, J. Comput. System Sci. 68 (4) (2004) 753–787.
[9] Sanchit Garg, Éric Schost, Interpolation of polynomials given by straight-line programs, Theoret. Comput. Sci. 410 (27-29) (2009) 2659–2662.
[10] Joachim Von Zur Gathen, Jurgen Gerhard, Modern Computer Algebra, Cambridge University Press, New York, NY, USA, 2003.
[11] Oded Goldreich, Combinatorial property testing (a survey), in: Randomization Methods in Algorithm Design, 1998, pp. 45–60.
[12] Dima Grigoriev, Marek Karpinski, A zero-test and an interpolation algorithm for the shifted sparse polynomials, in: AAECC, 1993, pp. 162–169.
[13] Dima Grigoriev, Marek Karpinski, Michael F. Singer, Fast parallel algorithms for sparse multivariate polynomial interpolation over finite fields, SIAM J. Comput. 19 (6) (1990) 1059–1063.
[14] Dima Grigoriev, Marek Karpinski, Michael F. Singer, Interpolation of sparse rational functions without knowing bounds on exponents, in: FOCS, 1990, pp. 840–846.
[15] Dima Grigoriev, Yagati N. Lakshman, Algorithms for computing sparse shifts for multivariate polynomials, in: ISSAC, 1995, pp. 96–103.
[16] Li Gyun-y, A new algorithm for the inversion of Hankel and Toeplitz matrices, Ukrainian Math. J. 36 (6) (1984) 536–540.
[17] Erich Kaltofen, Yagati N. Lakshman, Improved sparse multivariate polynomial interpolation algorithms, in: ISSAC, 1988, pp. 467–474.
[18] Yagati N. Lakshman, B. David Saunders, Sparse polynomial interpolation in nonstandard bases, SIAM J. Comput. 24 (2) (1995) 387–397.
[19] Yishay Mansour, Randomized interpolation and approximation of sparse polynomials, SIAM J. Comput. 24 (2) (1995) 357–368.
[20] Dana Ron, Property Testing (A Tutorial), Handbook of Randomization, vol. 2, 2000.
[21] Dana Ron, Property testing: A learning theory perspective, Foundations and Trends in Machine Learning 1 (3) (2008) 307–402.