Information Processing Letters 142 (2019) 77–79

Contents lists available at ScienceDirect

Information Processing Letters

www.elsevier.com/locate/ipl

#P-completeness of counting roots of a sparse polynomial

Alexey Milovanov

National Research University Higher School of Economics, Russian Federation

ARTICLE INFO

Article history:
Received 16 January 2017
Received in revised form 13 September 2018
Accepted 13 September 2018
Available online 17 October 2018
Communicated by Łukasz Kowalik

Keywords:
Computational complexity
Finite fields
#P-completeness

ABSTRACT

It is known (from *Counting curves and their projections* by Joachim von zur Gathen, Marek Karpinski, Igor Shparlinski [1, part 4]) that counting the number of points on a curve $R(x, y) = 0$ where $R(x, y)$ is a sparse polynomial over $\mathbb{F}_q$ is #P-complete under randomized reductions.

We give a simple proof of a stronger result: counting roots of a sparse *univariate* polynomial over $\mathbb{F}_q$ is #P-complete under *deterministic* reductions.

1. Introduction and main result

We consider the field $\mathbb{F}_q$ where $q = p^n$ is a power of a prime number p. Elements of $\mathbb{F}_q$ can be represented as polynomials from $\mathbb{F}_p[x]$ modulo some irreducible polynomial of degree n. This polynomial can be found in polynomial (in p, n) time,[1] as well as the matrix that relates two representations corresponding to different irreducible polynomials, see [4]. Therefore, for a fixed p, we do not need to specify the choice of an irreducible polynomial when speaking about algorithmic problems dealing with elements of $\mathbb{F}_{p^n}$.

Fix a prime number p. Consider the following counting problem: given an integer n in the unary representation and a polynomial $R \in \mathbb{F}_{p^n}[x]$, find the number of R's roots in $\mathbb{F}_{p^n}$. The polynomial R is given in a sparse representation,[2] as a list of monomials; each monomial $a_k x^k$ is presented as a pair (k in binary, a_k as an element of $\mathbb{F}_{p^n}$). This problem is called SparcePolynomialRoots-p. Our main result is the following statement:

Theorem 1. *For every prime p the problem* SparsePolynomialRoots-*p is #P-complete under deterministic reductions.*

Note that this implies the result from [1] mentioned in the Abstract.

2. Proof of the main result

We use #3SAT (counting the number of satisfying assignments for a 3-CNF) as a standard #P-complete problem. Consider some 3-CNF S. Each clause in S has the form $L_1 \vee L_2 \vee L_3$ where L_i are literals (i.e., variables or negations of variables). Then we construct a system S' of polynomial equations whose solutions correspond to the satisfying assignments for S. For each propositional variable x_i we have an equation $x_i^2 = x_i$ that guarantees that $x_i = 0$ (False) or $x_i = 1$ (True); the literal $\neg x_i$ is now $1 - x_i$, and each disjunction $L_1 \vee L_2 \vee L_3$ is converted to a polynomial equation $(1 - L_1) \cdot (1 - L_2) \cdot (1 - L_3) = 0$. Note that the correspondence between the satisfying assignments for S and solutions of S' works for every field; we use it for the field $\mathbb{F}_p$.

E-mail address: amilovanov@hse.ru.

[1] Note that the time here is polynomial in p, not in $\log p$; for us this is enough, since p is fixed in our statements.

[2] The same problem for a polynomial presented as a list of *all* coefficients, including zeros (each coefficient takes $\Theta(n)$ bits), is solvable in polynomial time for multivariate polynomials and a fixed p and the number of variables, see, e.g., [3].

https://doi.org/10.1016/j.ipl.2018.09.008
0020-0190/

We want the number of equations in S' to be smaller than the number of the variables (we need this for some technical reasons). To achieve this, we add $2t$ dummy variables $x_{m+1}, \dots x_{m+2t}$ (to the existing variables $x_1, \dots, x_m$) and t new equations that guarantee (using Fermat's little theorem) that all dummy variable are zeros: $(1 - x_{m+1}^{p-1}) \cdot (1 - x_{m+2}^{p-1}) = 1$, $(1 - x_{m+3}^{p-1}) \cdot (1 - x_{m+4}^{p-1}) = 1$, etc. For large enough t we have more variables than equations. Note that this trick does not change the number of solutions. We keep the notation S' for the resulting system.

Let $x_1, \dots, x_n$ be the variables that appear in S'. These variables are considered as elements of $\mathbb{F}_p$. Now we reduce S' to one polynomial equation over $\mathbb{F}_{p^n}$. For that, we consider a basis $\omega_1, \dots, \omega_n$ of $\mathbb{F}_{p^n}$ over $\mathbb{F}_p$. Then every $x \in \mathbb{F}_{p^n}$ can be represented uniquely as

$$x = x_1\omega_1 + \dots + x_n\omega_n,$$

where $x_1, \dots, x_n \in \mathbb{F}_p$. First we transform the equations of S' into sparse polynomial equations with one variable $x \in \mathbb{F}_{p^n}$, and then show how the resulting system of polynomial equations in x can be replaced by one equation.

Now we implement this plan. We need to find sparse polynomials $f_i \in \mathbb{F}_{p^n}[x]$ such that $f_i(x) = x_i$. This is enough for our first step, since a product of a constant number of polynomials (three for the disjuctions and $O(p)$ for the additional equations; recall that p is a constant) in the sparse representation is again a polynomial in the sparse representation whose size is only polynomially bigger. The following lemma [5, Lemma 3.51] helps.

Lemma. *Assume that $\alpha_1, \dots, \alpha_k$ for some $k \leqslant n$ are elements of $\mathbb{F}_{p^n}$ that are linearly independent over $\mathbb{F}_p$. Then the determinant*

$$\begin{vmatrix} \alpha_1 & \alpha_1^p & \alpha_1^{p^2} & \dots & \alpha_1^{p^{k-1}} \\ \alpha_2 & \alpha_2^p & \alpha_2^{p^2} & \dots & \alpha_2^{p^{k-1}} \\ \dots & \dots & \dots & \dots & \dots \\ \alpha_k & \alpha_k^p & \alpha_k^{p^2} & \dots & \alpha_k^{p^{k-1}} \end{vmatrix}$$

is a non-zero element of $\mathbb{F}_{p^n}$.

For reader's convenience we reproduce the proof here.

Proof of the lemma. Consider this determinant as a function of α_1 when other α_i are fixed. In other words, consider the polynomial $P(x)$ that is obtained if we replace α_1 by x everywhere in the first row. We get a polynomial of degree (at most) p^{k-1}. The powers of x appearing in P are $1, p, p^2, \dots, p^{k-1}$, so this polynomial is linear as a function $\mathbb{F}_{p^n} \to \mathbb{F}_{p^n}$ if we consider $\mathbb{F}_{p^n}$ as a vector space over $\mathbb{F}_p$ (recall that $(a+b)^p = a^p + b^p$ in a field of characteristic p, and $\alpha^p = \alpha$ in $\mathbb{F}_p$). The polynomial $P(x)$ has roots $\alpha_2, \dots, \alpha_k$ (two equal rows guarantee the zero determinant); all p^{k-1} linear combinations of $\alpha_2, \dots, \alpha_k$ with $\mathbb{F}_p$-coefficients are also roots due to linearity. Reasoning by induction, we may assume that the leading coefficient of P, being the determinant of the same type for smaller k, is not zero. Then we know that P has no other roots, and $P(\alpha_1) \neq 0$, since α is not a linear combination of $\alpha_2, \dots, \alpha_k$. □

Now we define the polynomial

$$f_1(x) := c \begin{vmatrix} x & x^p & x^{p^2} & \dots & x^{p^{n-1}} \\ \omega_2 & \omega_2^p & \omega_2^{p^2} & \dots & \omega_2^{p^{n-1}} \\ \omega_3 & \omega_3^p & \omega_3^{p^2} & \dots & \omega_3^{p^{n-1}} \\ \dots & \dots & \dots & \dots & \dots \\ \omega_n & \omega_n^p & \omega_n^{p^2} & \dots & \omega_n^{p^{n-1}} \end{vmatrix}$$

for suitable $c \neq 0$. We know (see the proof of the lemma) that f_1 equals 0 on the linear combinations of $\omega_2, \dots, \omega_n$, i.e., on all elements with $x_1 = 0$. The lemma says that $f_1(\omega_1) \neq 0$, and linearity guarantees that f_1 has the same values on all elements x with $x_1 = 1$. Choose c to make $f_1(\omega_1)$ equal to 1. Linearity over $\mathbb{F}_p$ then guarantees that $f_1(x) = x_1$ for all $x \in \mathbb{F}_{p^n}$.

We have constructed the polynomial f_1; in the same way we construct $f_i(x) \in \mathbb{F}_{p^n}[x]$ such that $f_i(x) = x_i$ for $x = \sum x_i\omega_i$. In this way we reduce a polynomial equation over $\mathbb{F}_p$ with n variables to a univariate polynomial equation over $\mathbb{F}_{p^n}$.

What have we achieved? We know that the number of satisfying assignments for 3-CNF S (with Boolean variables) is equal to the number of solutions of the system of polynomial equations $P_1(x) = 0, P_2(x) = 0, \dots$ where P_k are some polynomials in $\mathbb{F}_{p_n}[x]$ and $x \in \mathbb{F}_{p^n}$. Each P_k is obtained from some equation in S' by replacing all x_i by $f_i(x)$. We can now replace the system by one equation

$$P_1(x)\omega_1 + P_2(x)\omega_2 + \dots = 0$$

in $\mathbb{F}_{p^n}$ using the fact that polynomials P_i have values 0 and 1 (being a product of two or three polynomials with this property). Here we use the specific properties of the system S', in particular, we use that the number of equations in S' is at most n (otherwise we cannot find enough linearly independent ω_i).

As we have discussed, each P_k has only polynomially many monomials in the sparse representation. Note also that the coefficients of all f_i (and therefore the coefficients of all P_k) can be computed in poly(size of S) time. Indeed, we need only to calculate the determinants that define the coefficients of f_i, and this is a polynomial task; note that the powers of $\omega_1, \dots, \omega_n$ can be computed by repeated squaring.

So, for each fixed p, we have constructed a deterministic polynomial reduction of #3-SAT to the problem of counting the number of roots for a univariate polynomial in $\mathbb{F}_{p^n}$ in a sparse representation. Theorem 1 is proven.

3. Related questions

Note that the reduction in the proof is *parsimonious*, i.e., every satisfying assignment of a 3-CNF S corresponds to a root of the sparse polynomial constructed starting from S. This is useful if we consider the following problems:

- SparsePolynomialRoot-p for fixed p: given n and a polynomial from $\mathbb{F}_{p^n}$ in sparse representation, find out if the given polynomial has a root in $\mathbb{F}_{p^n}$ or not.

- SparsePolynomialRootsParity-p: given n and a polynomial from $\mathbb{F}_{p^n}$ in sparse representation, return the *parity* of the number of roots of the given polynomial in $\mathbb{F}_{p^n}$.

Since 3SAT is NP-complete and ⊕3SAT is ⊕-complete, and since our reduction is parsimonious, we get the following corollaries.

Corollary 1 *([2]).* SparsePolynomialRoot-p *is* NP-*complete for every prime* p.

Corollary 2. SparsePolynomialParityRoots-p *is* ⊕P-*complete for every prime* p.

Remark. We may also consider the version of the problem where the input p is presented in binary. Of course, this problem is also #P-hard (because it is #P-hard for a fixed p). However, the membership in #P for this problem is an open question.

Acknowledgements

I would like to thank Alexander Shen for help in writing this paper.

This work is supported in parts by the RFBR grant 16-01-00362, by the Young Russian Mathematics award, MK-5379.2018.1 and the RaCAF ANR-15-CE40-0016-01 grants.

The study has also been funded by the Russian Academic Excellence Project '5-100'.

References

[1] J. von zur Gathen, M. Karpinski, I. Shparlinski, Counting curves and their projections, Comput. Complex. 6 (1996) 64–99.

[2] A. Kipnis, A. Shamir, Cryptanalysis of the HFE public key cryptosystem by relinearization, in: CRYPTO-99, in: Lecture Notes for Computer Science, vol. 1666, 1999, pp. 19–30.

[3] Alan G.B. Lauder, Daqing Wan, Counting points on varieties over finite fields of small characteristic, in: Algorithmic Number Theory: Lattices, Number Fields, Curves and Cryptography, in: Math. Sci. Res. Inst. Publ., vol. 44, Cambridge Univ. Press, Cambridge, 2008, pp. 579–612.

[4] H.W. Lenstra Jr., Finding isomorphisms between finite fields, Math. Comput. 56 (1991) 329–347.

[5] R. Lidl, H. Niederreiter, Introduction to Finite Fields and Their Applications, Cambridge University Press, 1986.