## Deterministic root finding over finite fields using Graeffe transforms

- Source link

Bruno Grenet, Joris van der Hoeven, Grégoire Lecerf
Institutions: University of Montpellier, École Polytechnique
Published on: 01 Jun 2016 - Applicable Algebra in Engineering, Communication and Computing (Springer Berlin Heidelberg)

Topics: Finite field, Deterministic algorithm, Multiplicative group, Properties of polynomial roots and Root-finding algorith m

Related papers:

- Fast Polynomial Factorization and Modular Composition
- Modern Computer Algebra
- Randomized Root Finding over Finite FFT-fields using Tangent Graeffe Transforms
- On the geometry of Graeffe Iteration
- On fast multiplication of polynomials over arbitrary algebras


# Deterministic root finding over finite fields using Graeffe transforms 

Bruno Grenet, Joris van der Hoeven, Grégoire Lecerf

## To cite this version:

Bruno Grenet, Joris van der Hoeven, Grégoire Lecerf. Deterministic root finding over finite fields using Graeffe transforms. 2015. hal-01104251

HAL Id: hal-01104251
https://hal.archives-ouvertes.fr/hal-01104251
Preprint submitted on 16 Jan 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Deterministic root finding over finite fields using Graeffe transforms 

Bruno Grenet ${ }^{A}$<br>Laboratoire d'informatique, de robotique et de microélectronique de Montpellier LIRMM, UMR 5506 CNRS, CC477<br>Université Montpellier 2 161, rue Ada 34095 Montpellier Cedex 5, France<br>Email: bruno.grenet@lirmm.fr<br>Joris van der Hoeven ${ }^{b}$, Grégoire Lecerf ${ }^{c}$<br>Laboratoire d'informatique de l'École polytechnique<br>LIX, UMR 7161 CNRS<br>Campus de l'École polytechnique<br>1, rue Honoré d'Estienne d'Orves<br>Bâtiment Alan Turing, CS35003<br>91120 Palaiseau, France<br>b. Email: vdhoeven@lix.polytechnique.fr<br>c. Email: lecerf@lix.polytechnique.fr

Version of January 13, 2015


#### Abstract

We design new deterministic algorithms, based on Graeffe transforms, to compute all the roots of a polynomial which splits over a finite field $\mathbb{F}_{q}$. Our algorithms were designed to be particularly efficient in the case when the cardinality $q-1$ of the multiplicative group of $\mathbb{F}_{q}$ is smooth. Such fields are often used in practice because they support fast discrete Fourier transforms. We also present a new nearly optimal algorithm for computing characteristic polynomials of multiplication endomorphisms in finite field extensions. This algorithm allows for the efficient computation of Graeffe transforms of arbitrary orders.


## 1. INTRODUCTION

Let $\mathbb{F}_{q}$ represent the finite field with $q=p^{k}$ elements, where $p$ is a prime number, and $k \geqslant 1$. Throughout this article, such a field is supposed to be described as a quotient of $\mathbb{F}_{p}[x]$ by a monic irreducible polynomial. Let $f \in \mathbb{F}_{q}[x]$ represent a separable monic polynomial of degree $d \geqslant 1$ which splits over $\mathbb{F}_{q}$, which means that all its irreducible factors have degree one and multiplicity one. In this article we are interested in computing all the roots of $f$.

### 1.1. Motivation

In a previous paper [20], we presented efficient randomized root finding algorithms. Such algorithms were needed for the efficient interpolation, into the standard monomial basis, of polynomials that are given through evaluation functions. This latter problem is also known

[^0]under the name sparse interpolation, and root finding often turns out to be a bottleneck, as indicated in [23]. In fact, in this case, the ground field can be chosen to be $\mathbb{F}_{p}$ with $p=M 2^{m}+1$, and where $2^{m}$ is taken to be much larger than the number of terms to be discovered. In order to minimize the size of $p$, making it fit into a machine register, we usually take $M=O(\log p)$ as small as possible. A typical example is $p=7 \cdot 2^{26}+1$. We informally refer to such primes as FFT primes.

While working on [20], it turned out that some of the new methods could also be used for the design of fast deterministic methods for computing all the roots of $f$. Even though randomized algorithms are more useful in practice, the existence of efficient deterministic algorithms remains interesting from a theoretical point of view. The goal of this paper is to report on our progress on this issue.

### 1.2. Notations and prerequisites

We will use the same notations as in [20], which we briefly recall now. The multiplicative group $\mathbb{F}_{q} \backslash\{0\}$ of $\mathbb{F}_{q}$ is written $\mathbb{F}_{q}^{*}$. Complexity bounds will frequently be expressed using the soft-Oh notation: $f(n) \in \tilde{O}(g(n))$ means that $f(n)=g(n) \log O(1) g(n)$. Given $x \in \mathbb{R}$, we write $\lfloor x\rfloor=\max \{k \in \mathbb{Z}: k \leqslant x\}$ and $\lceil x\rceil=\min \{k \in \mathbb{Z}: k \geqslant x\}$. The remainder of a division of $g$ by $f$ is denoted by $g$ rem $f$.

We write $\mathrm{M}(d)$ for the complexity of multiplication of polynomials of degree $<d$ over an arbitrary ring. We also write $\mathbf{I}(n)$ and $\mathbf{M}_{q}(d)$ for the bit complexities (in the Turing machine model) of $n$-bit integer multiplication and multiplication of polynomials of degrees $<d$ over $\mathbb{F}_{q}$. We make the customary assumption that $\mathrm{M}(d) / d, \mathrm{I}(n) / n$ and $\mathrm{M}_{q}(d) / d$ are increasing functions. Currently best known bounds [10, 21, 22] are $\mathrm{M}(d)=$ $O(d \log d \log \log d), \mathrm{I}(n)=O\left(n \log n 8^{\log ^{*} n}\right)$ and $\mathrm{M}_{q}(d)=O\left(d \log q \log (d \log q) 8^{\log }{ }^{*}(d \log q)\right)$, where $\log ^{*}$ stands for the iterated logarithm function $\log ^{*} n=\min \left\{k \in \mathbb{N}: \log { }^{k \times .} \log x \leqslant 1\right\}$.

We freely use the following classical facts: ring operations in $\mathbb{F}_{p} \operatorname{cost} O(l(\log p))$ and one division or inversion in $\mathbb{F}_{p}$ costs $O(I(\log p) \log \log p)$ [17, Chapter 11]. The ring operations in $\mathbb{F}_{q}$ cost at most $O\left(\mathrm{M}_{p}(k)\right)$, and inversions take at most $O\left(\mathrm{M}_{p}(k) \log k+\mathrm{I}(\log p) \log \log p\right)$. For convenience, $\mathrm{m}_{q}$ and $\mathrm{d}_{q}$ will respectively denote cost functions for the product and the inverse in $\mathbb{F}_{q}$. The gcd of two polynomials of degrees at most $d$ over $\mathbb{F}_{q}$ can be computed in time $O\left(\mathrm{M}_{q}(d) \log d\right)$ [17, Algorithm 11.4]. Given polynomials $f$ and $g_{1}, \ldots, g_{l}$ monic with $\operatorname{deg} f=d$ and $\operatorname{deg} g_{1}+\cdots+\operatorname{deg} g_{l}=O(d)$, all the remainders $f$ rem $g_{i}$ can be computed in time $O\left(\mathrm{M}_{q}(d) \log l\right)$ [17, Chapter 10]. The inverse problem, called Chinese remaindering, can be solved within a similar cost $O\left(\mathrm{M}_{q}(d) \log l+d \mathrm{~d}_{q}\right)$.

For an integer $n$, the largest prime dividing $n$ is written $S_{1}(n)$, and the second one $S_{2}(n)$. Pollard and Strassen have given a deterministic algorithm for factoring $p-1$, of bit complexity $O\left(\mathrm{M}\left(S_{2}(p-1)^{1 / 2}\right) \mathrm{I}(\log p) \log p\right)$ (see for instance [17, Corollary 19.4]). Ignoring smoothness considerations, the latter result has been improved into $O\left(I\left(p^{1 / 4} \log p\right)\right)$ in [6], and further refined to $O\left(\mathrm{I}\left(p^{1 / 4}(\log p) / \sqrt{\log \log p}\right)\right)$ in [13].

In this article, when needed, we consider that the factorization of $q-1$ and a primitive element of $\mathbb{F}_{q}^{*}$ have been precomputed once, and we discard the necessary underlying costs. In practice, if the factorization of $q-1$ is given, then it is straightforward to verify whether a given element is primitive. For known complexity bounds and historical details on these tasks, we refer to $[2,33,44]$.

### 1.3. Related work

It is now classical that polynomial factorization over $\mathbb{F}_{q}$ reduces to finding roots over $\mathbb{F}_{p}$ in deterministic polynomial time, uniformly in $d$ and $\log q$. But it is still unknown whether root finding can be solved in deterministic polynomial time or not, even under classical
conjectures in number theory. This problem can nevertheless be solved efficiently by randomized algorithms in average polynomial time. Concerning related results and historical surveys on this topic, the reader might consult [17, 18, 25, 33].

Seminal algorithms for polynomial factorization over finite fields are classically attributed to Berlekamp [3, 4], and Cantor and Zassenhaus [11], but central earlier ideas can be found in works of Gauss, Galois, Arwins, Faddeev and Skopin. Cantor-Zassenhaus' algorithm is randomized and well suited to compute roots of polynomials of degree $d$ that split over $\mathbb{F}_{q}$ in average time $O\left(\mathrm{M}_{q}(d)(\log q+\log d) \log d+d \mathrm{~d}_{q}\right)$. Of course, if $q=O(d)$ then an exhaustive search can be naively performed in time $O\left(\mathrm{M}_{q}(d) \log d\right)$ (the factor $\log d$ can be discarded if a primitive element of $\mathbb{F}_{q}^{*}$ is given, by means of [8, Proposition 3$]$ ), so that the cost of root finding simplifies to $O\left(\mathrm{M}_{q}(d) \log q \log d+d \mathrm{~d}_{q}\right)$. This classical approach is for instance implemented in the NTL library written by Shoup [45]. However neither Berlekamp's nor Cantor-Zassenhaus' algorithm seems to benefit of particular prime numbers such as FFT primes. Instead, alternative approaches have been proposed by Moenck [32], von zur Gathen [16], Mignotte and Schnorr [31], and then by Rónyai [37]. Recent slight improvements of Moenck's and Mignotte-Schnorr's algorithms may also be found in [20].

Deterministic polynomial time root finding has been tackled by several authors. Camion showed that, for a fixed finite field, one can pre-compute a suitable subset which allows to factor any polynomial in deterministic polynomial time [9]. Nevertheless the construction of such a subset is purely theoretical, and it is not clear that there exists an efficient algorithm to compute it, even for FFT prime fields.

Schoof, Rónyai, Huang, and Źrałek designed different methods for particular types of input polynomials according to their syntax or to properties of the Galois group of the lifted input polynomial over $\mathbb{Q}[24,37,38,40,46]$. Some of their algorithms require the general Riemann hypothesis (GRH) or the extended Riemann hypothesis (ERH). Sub-exponential algorithms are also known from the work of Evdokimov [14], which has recently led to new conjectures in terms of graphs [1]. Other conjectural algorithms can be found in [15, 39].

Another series of articles concern complexity bounds which are uniformly polynomial in the degree and in $S_{1}(q-1)$. Such algorithms are practical when $q-1=p_{1}^{m_{1}} \cdots p_{r}^{m_{r}}$ is sufficiently smooth. Assuming a primitive element of $\mathbb{F}_{q}^{*}$ is given, Mignotte and Schnorr [31] proposed a method based on a cascade of gcds, that needs $O\left(\mathrm{M}(d) \sum_{i=1}^{r} m_{i}(\log q+\right.$ $\left.p_{i} \log d\right)$ ) operations in $\mathbb{F}_{q}$. The computation of a primitive element might of course be expensive but it can be seen as a pre-computation. In fact, von zur Gathen proved the deterministic polynomial time equivalence (in terms of $S_{1}(p-1)$ and input size) between the computation of primitive elements and polynomial factorization [16].

Rónyai [37] obtained a polynomial complexity bound in $S_{1}(p-1), d$ and $\log p$ by means of linear algebra techniques, using primitive roots of orders lower than in [16], but he did not explicit the exponents in the bound. For $\mathbb{F}_{p}[x]$, Shoup [41] reached the bound $p^{1 / 2} \log ^{2} p \tilde{O}\left(d^{2}\right)$ in terms of the number of operations in $\mathbb{F}_{p}$, and then refined it to $p^{1 / 2} \log p \tilde{O}(d k \min (d, k))+\tilde{O}\left(d k^{2}\right) \log p$ for $\mathbb{F}_{q}[x]$, still in terms of operations in $\mathbb{F}_{p}[42]$. Finally Shoup proved the bound $S_{1}(p-1)^{1 / 2}(d \log p)^{O(1)}$ in $\mathbb{F}_{p}[x]$ under ERH [43].

### 1.4. Our contributions

The main contribution of this article is an efficient deterministic root finding method. The new algorithm is based on generalized Graeffe transforms and it is particularly efficient for finite fields $\mathbb{F}_{p}$ such that $p$ is an FFT prime. Roughly speaking, the generalized Graeffe transform replaces modular exponention as used in Cantor-Zassenhaus' algorithm, and gcd computations are changed to multi-point evaluations.

Thanks to a slight improvement of Shoup's adaptation of Pollard-Strassen's algorithm to find roots deterministically [43], we obtain a new deterministic complexity bound in terms of $S_{1}(p-1)^{1 / 2}$. In addition, in the smooth case $S_{1}(q-1)=O(\log q)$, the cost of our deterministic algorithm is softly equivalent to the one of Cantor-Zassenhaus.

The complexity of generalized Graeffe transforms is studied in Section 2. A first ingredient is the transposed modular composition algorithm of Kedlaya and Umans [27], also called modular power projections algorithm. The second ingredient is a special variant for finite fields of Newton-Girard's formulas for recovering a polynomial from the power sums of its roots. Such a variant had been previously designed by Bostan, GonzalezVega, Perdry and Schost [7] in the case of prime finite fields. We extend their results to finite fields from scratch, using very simple arguments, and independently of the framework designed in [12].

## 2. Generalized Graeffe transforms over finite fields

Classically, the Graeffe transform of a polynomial $g \in \mathbb{F}_{q}[x]$ of degree $d$ is the unique polynomial $h \in \mathbb{F}_{q}[x]$ satisfying $h\left(x^{2}\right)=g(x) g(-x)$. If $g(x)=\prod_{i=1}^{d}\left(\alpha_{i}-x\right)$, then $h(x)=\prod_{i=1}^{d}\left(\alpha_{i}^{2}-x\right)$. This construction can be extended to higher orders as follows: the generalized Graeffe transform of $g$ of order $\pi$, written $G_{\pi}(g)$, is defined as the resultant

$$
G_{\pi}(g)(x)=(-1)^{\pi d} \operatorname{Res}_{z}\left(g(z), z^{\pi}-x\right)
$$

If $g=\prod_{i=1}^{d}\left(\alpha_{i}-x\right)$, then $G_{\pi}(g)(x)=\prod_{i=1}^{d}\left(\alpha_{i}^{\pi}-x\right)$. Equivalently, $G_{\pi}(g)$ is the characteristic polynomial of multiplication by $x^{\pi}$ in $\mathbb{F}_{q}[x] /(g)$ (up to the sign). In this section we show how to compute Graeffe transforms efficiently. We start with a particularly simple and efficient algorithm for the smooth case, repeated from [20].

### 2.1. Smooth case

For our root finding algorithm, the most important case is when $q-1$ is smooth and the order $\pi$ of the generalized Graeffe transform divides $q-1$.

Proposition 1. Let $\pi_{1}, \ldots, \pi_{m}$ be integers $\geqslant 2$, such that $\chi=\pi_{1} \cdots \pi_{m}$ divides $q-1$, and let $\xi_{i}$ be given primitive roots of unity of order $\pi_{i}$, for all $i \in\{1, \ldots, m\}$. If $g$ is a monic polynomial in $\mathbb{F}_{q}[x]$ of degree $d$, then the generalized Graeffe transforms of orders $\pi_{1}$, $\pi_{1} \pi_{2}, \pi_{1} \pi_{2} \pi_{3}, \ldots, \chi$ of $g$ can be computed in time $O\left(m \mathrm{M}_{q}(\mu d)\right)$ or $\sigma \tilde{O}(d \log q \log \mu)$, where $\mu=\max \left(\pi_{1}, \ldots, \pi_{m}\right)$ and $\sigma=\pi_{1}+\cdots+\pi_{m}$.

Proof. Writing $g(x)=c \prod_{j=1}^{d}\left(\alpha_{j}-x\right)$ in an algebraic closure of $\mathbb{F}_{q}$, the Graeffe transform of $g$ of order $\pi_{i}$ is $h_{i}(x)=c^{\pi_{i}} \prod_{j=1}^{d}\left(\alpha_{j}^{\pi_{i}}-x\right)$. Consequently this leads to $h_{i}\left(x^{\pi_{i}}\right)=g(x) g\left(\xi_{i} x\right) g\left(\xi_{i}^{2} x\right) \cdots g\left(\xi_{i}^{\pi_{i}-1} x\right)$. Using the latter formula, by Lemma 2 below, the transform can be obtained in time $O\left(\mathrm{M}_{q}\left(\pi_{i} d\right)\right)$. Taking the sum over $i$ concludes the proof.

Lemma 2. Let $g$ be a polynomial of degree $d \geqslant 1$ in $\mathbb{F}_{q}[x]$, let $\alpha \in \mathbb{F}_{q}$, and let $l$ be an integer. Then the product $P_{l}(x)=g(x) g(\alpha x) g\left(\alpha^{2} x\right) \cdots g\left(\alpha^{l-1} x\right)$ can be computed in time $O\left(\mathrm{M}_{q}(l d)\right)$.

Proof. Let $h:=\lfloor l / 2\rfloor$. If $l$ is even, then $P_{l}(x)=P_{h}(x) P_{h}\left(\alpha^{h} x\right)$, otherwise we have $P_{l}(x)=P_{h}(x) P_{h}\left(\alpha^{h} x\right) g\left(\alpha^{l-1} x\right)$. These formulas lead to an algorithm with the claimed cost.

Notice that this lemma generalizes to arbitrary fields.

### 2.2. General case

In the general case, let $g \in \mathbb{F}_{q}[x]$ be monic of degree $d$. Given a polynomial $e$ of degree $<d$, we are interested in computing the characteristic polynomial $h$ of the multiplication by $e$ in $R:=\mathbb{F}_{q}[x] /(g)$. Our strategy essentially follows Le Verrier's method: we first compute the traces $\operatorname{Tr}\left(e^{i}\right)$ of the multiplication by $e^{i}$ in $R$ for all $i \in\{1, \ldots, d\}$, which are also called the power sums of the roots of the characteristic polynomial $h$. The generating series $\tau(x)=\sum_{i \geqslant 0} \operatorname{Tr}\left(e^{i}\right) x^{i} \in \mathbb{F}_{q}[[x]]$ of these power sums relates to the logarithmic derivative of the reverse polynomial $\eta=x^{d} h(1 / x)$ of $h$ via a first order differential equation

$$
\begin{equation*}
\frac{\eta^{\prime}(x)}{\eta(x)}=\tau(x)+O\left(x^{d}\right) \tag{1}
\end{equation*}
$$

This equation can be seen as a compact form of the classical Newton-Girard identities which relate power sums and symmetric polynomials. After computing $\tau$ it thus suffices to solve this equation in order to deduce $h$. If $p \geqslant d$, then this method is classical. If $p<d$, then $h$ cannot directly be retrieved in this way [35, Fact 2.1]. In the rest of this section, we will overcome this problem by computing with $p$-adic integers at a suitable finite precision $\kappa$.

### 2.3. Power sums

Let $\mathbb{Z} / p^{\kappa} \mathbb{Z}$ be the ring of truncated $p$-adic integers at precision $\kappa$. It is convenient to write $\operatorname{GR}\left(p^{\kappa}, k\right)$ for the Galois ring of characteristic $p^{\kappa}$ and degree $k$, which is classically defined as $\operatorname{GR}\left(p^{\kappa}, k\right)=\left(\left(\mathbb{Z} / p^{\kappa} \mathbb{Z}\right)[z]\right) /(\mu)$, where $\mu \in\left(\mathbb{Z} / p^{\kappa} \mathbb{Z}\right)[z]$ is a monic polynomial of degree $k$ and is irreducible modulo $p$.

Let $g$ represent a given monic polynomial of degree $d$, let $A=\operatorname{GR}\left(p^{\kappa}, k\right)[x] /(g)$. We write $\operatorname{Tr}: A \rightarrow \operatorname{GR}\left(p^{\kappa}, k\right)$ the classical trace function on $A$. In order to compute the power sums $\operatorname{Tr}(e), \operatorname{Tr}\left(e^{2}\right), \ldots, \operatorname{Tr}\left(e^{d}\right)$, we will use the following proposition, which is essentially a consequence of a result due to Kedlaya and Umans in [26, 27]:

Proposition 3. Let $\delta>0$ be a fixed constant. If $d^{1+\delta} \leqslant q-1$, then the traces $\operatorname{Tr}(e)$, $\operatorname{Tr}\left(e^{2}\right), \ldots, \operatorname{Tr}\left(e^{d}\right)$ for a given $e \in A$ can be computed in time $\left(d \log \left(q^{\kappa}\right)\right)^{1+\delta}$.

Proof. Let $\gamma(x)=x^{d} g(1 / x)$ denote the reverse polynomial of $g$. Since $g$ is monic we have $\gamma(0)=1$. In the canonical basis made of the classes of $1, x, \ldots, x^{d-1}$ in $A$, the trace function can be computed as the first $d$ terms of the generating series

$$
\frac{\gamma^{\prime}(x)}{\gamma(x)}=\operatorname{Tr}(x)+\operatorname{Tr}\left(x^{2}\right) x+\operatorname{Tr}\left(x^{3}\right) x^{2}+\cdots+\operatorname{Tr}\left(x^{d}\right) x^{d-1}+O\left(x^{d}\right)
$$

in softly linear time, namely $\tilde{O}(\kappa d \log q)$.
The cardinality of $A$ is $q^{\kappa}$, and $A$ contains at least $q-1$ invertible elements. Then, for every constant $\delta>0$ such that $d^{1+\delta} \leqslant q-1$, Theorem 7.7 of [27, p. 1792] provides us with an algorithm to compute $\operatorname{Tr}(e), \operatorname{Tr}\left(e^{2}\right), \ldots, \operatorname{Tr}\left(e^{d}\right)$ in time $d^{1+\delta} \log ^{1+o(1)}\left(q^{\kappa}\right)$.

Remark 4. For $d$ close to $q$ the proposition does not apply. A simple solution is to lift the computations in $\mathbb{F}_{q^{2}}$. Nevertheless, for finding the roots of $g$, this case is in fact favorable: it suffices to evaluate $g$ at all the elements of $\mathbb{F}_{q}$. Therefore we will not need to work in $\mathbb{F}_{q^{2}}$.

### 2.4. Newton-Girard's identities

We carry on with the notations of the previous subsection. Let $e$ be a given element in $A$, and let $\tau(x)=\sum_{i \geqslant 0} \operatorname{Tr}\left(e^{i}\right) x^{i} \in \operatorname{GR}\left(p^{\kappa}, k\right)[[x]]$ represent the generating series of the power sums of the roots of the characteristic polynomial $h$ of $e$. If $p \geqslant d$, then $h$ can be deduced from $\operatorname{Tr}(e), \operatorname{Tr}\left(e^{2}\right), \ldots, \operatorname{Tr}\left(e^{d}\right)$ in softly linear time, using Newton-Girard's identities (1), as for instance in [5, Corollary 1]. If $k=1$, then Theorem 1 of [7] solves the equation modulo $p$ whenever $\kappa \geqslant\lfloor\log d / \log p\rfloor+1$. This subsection is dedicated to the remaining case when $p<d$ and $k$ is general. Our new alternative proofs are also shorter than in [7].

Lemma 5. Let $\theta \in \operatorname{GR}\left(p^{\kappa}, k\right)[x]$ be of degree $\leqslant d$ satisfying $\theta(0)=1$, and $\theta^{\prime}(x) / \theta(x)=$ $\tau(x)+O\left(x^{d}\right)$. If $\kappa>\lfloor\log d / \log p\rfloor$, then $\theta$ coincides with $\eta$ modulo $p$ at precision $O\left(x^{d+1}\right)$.

Proof. Let $\psi=\theta / \eta$. Writing $\psi$ as $\sum_{i \geqslant 0} \psi_{i} x^{i}$, we deduce from $\psi^{\prime}=O\left(x^{d}\right)$ that $i \psi_{i}=0$ for all $i \in\{1, \ldots, d\}$, whence $p^{\kappa-\operatorname{val}_{p} i}$ divides $\psi_{i}$, where $\operatorname{val}_{p} i$ denotes the $p$-adic valuation of $i$. Since $\kappa>\lfloor\log d / \log p\rfloor$ is equivalent to $d<p^{\kappa}$, we have $\kappa-\operatorname{val}_{p} i>0$, which concludes the proof.

Denote cost of multiplication in $\operatorname{GR}\left(p^{\kappa}, k\right)[x]$ of polynomials of degrees $<d$ by $\mathrm{M}_{p^{\kappa}, k}(d)$. We have $\mathrm{M}_{p^{\kappa}, k}(d)=O\left(k \kappa d \log p \log (k \kappa d \log p) 8^{\log ^{*}(k \kappa d \log p)}\right)=\tilde{O}(\kappa d \log q)$ according to [22].

Proposition 6. Let $e \in A$, and let $\kappa=\lfloor\log d / \log p\rfloor+1$. From $\operatorname{Tr}\left(e^{1}\right), \ldots, \operatorname{Tr}\left(e^{d}\right)$, we can compute the characteristic polynomial $h$ of e modulo $p$ in time $O\left(\mathrm{M}_{p^{\kappa}, k}(d)\right)=\tilde{O}(d \log q)$.

Proof. We will show how to compute a series $\theta \in \operatorname{GR}\left(p^{\kappa}, k\right)[[x]]$ satisfying $\theta(0)=1$ and $\theta^{\prime}(x) / \theta(x)=\tau(x)+O\left(x^{d}\right)$, using Newton's method. In view of the above lemma, this will yield $\eta$ followed by $h$ modulo $p$.

Suppose that we are given $\theta$ such that $\theta(0)=1$ and $\theta^{\prime}(x) / \theta(x)=\tau(x)+O\left(x^{n-1}\right)$ for some integer $n \geqslant 1$. Then we claim that there exists $\tilde{\theta} \in \operatorname{GR}\left(p^{\kappa}, k\right)[[x]]$ of valuation in $x$ at least $n$ and such that

$$
\begin{equation*}
\tilde{\theta}^{\prime}(x)=\tau(x)-\frac{\theta^{\prime}(x)}{\theta(x)}+O\left(x^{2 n-1}\right) . \tag{2}
\end{equation*}
$$

Indeed, using that $\tau-\theta^{\prime} / \theta=O\left(x^{n-1}\right)$ and $\tilde{\theta}=O\left(x^{n}\right)$, this equation is equivalent to

$$
\frac{\tilde{\theta}^{\prime}}{1+\tilde{\theta}}=\left(\tau-\frac{\theta^{\prime}}{\theta}\right)\left(1-\tilde{\theta}+\tilde{\theta}^{2}+\cdots\right)+O\left(x^{2 n-1}\right)=\tau-\frac{\theta^{\prime}}{\theta}+O\left(x^{2 n-1}\right),
$$

and therefore equivalent to

$$
\frac{((1+\tilde{\theta}) \theta)^{\prime}}{(1+\tilde{\theta}) \theta}=\frac{\tilde{\theta}^{\prime}}{1+\tilde{\theta}}+\frac{\theta^{\prime}}{\theta}=\tau(x)+O\left(x^{2 n-1}\right) .
$$

This latter equation admits $\eta / \theta-1$ as a solution. Having proved our claim, we can compute $\tilde{\theta}$ as being any integral of $\tau-\theta^{\prime} / \theta$ modulo $O\left(x^{2 n}\right)$, and observe that $\hat{\theta}=(1+\tilde{\theta}) \theta$ satisfies $\hat{\theta}(0)=1$ and $\hat{\theta}^{\prime} / \hat{\theta}=\tau+O\left(x^{2 n-1}\right)$.

The cost of one iteration (2) is bounded by $O\left(\mathrm{M}_{p^{\kappa}, k}(n)\right)$. Starting with $n=1$ and $\theta=1$, we recursively apply the iteration for $n=1,2,4, \ldots, 2^{[\log \kappa / \log 2]-1}$, the end result being a $\theta \in \operatorname{GR}\left(p^{k}, k\right)[[x]]$ with $\theta^{\prime} / \theta=\tau+O\left(x^{d}\right)$. It is classical that the cost of the last iteration dominates the total cost, which is therefore bounded by $O\left(\mathrm{M}_{p^{\kappa}, k}(d)\right)$.

### 2.5. Characteristic polynomials

Now we come back to the original computational problem over $\mathbb{F}_{q}$.
Theorem 7. Let $g \in \mathbb{F}_{q}[x]$ be a monic polynomial of degree $d$, and let $e \in \mathbb{F}_{q}[x]$ be of degree $<d$. For every constant $\delta>0$ such that $d^{1+\delta} \leqslant q-1$, the characteristic polynomial $h$ of the multiplication by e in $\mathbb{F}_{q}[x] /(g)$ can be computed in time $(d \log q)^{1+\delta}$.

Proof. We embed $\mathbb{F}_{q}[x] /(g)$ into $A=\operatorname{GR}\left(p^{\kappa}, k\right)[x] /(g)$, where $\operatorname{GR}\left(p^{\kappa}, k\right)$ is constructed from $\left(\mathbb{Z} / p^{\kappa} \mathbb{Z}\right)[z] /(\mu)$, and $\mu \in\left(\mathbb{Z} / p^{\kappa} \mathbb{Z}\right)[z]$ is the monic defining polynomial of degree $k$ of $\mathbb{F}_{q}$ over $\mathbb{F}_{p}$.

Let $\delta>0$ be such that $d^{1+\delta} \leqslant q-1$. Proposition 3 allows us to compute the traces of the successive powers of $e$ seen in $A$ in time $\left(d \log \left(q^{\kappa}\right)\right)^{1+\delta}$. Then Proposition 6 allows to deduce $h$ modulo $p$ in softly linear time.

Remark 8. If the characteristic $p$ is larger than the multiplicities of the roots of $h$, then Pan's algorithm [35] does not involve $p$-adic arithmetic, but may require higher order traces.

Corollary 9. Let $g$ be a monic polynomial in $\mathbb{F}_{q}[x]$ of degree d. For every constant $\delta>0$ such that $d^{1+\delta} \leqslant q-1$, the generalized Graeffe transform $G_{\pi}(g)$ can be computed in time $(d \log q)^{1+\delta}+\tilde{O}(d \log q \log \pi)$.

Proof. We apply the latter theorem to $e(x)=x^{\pi}$ rem $g(x)$, that can be computed in time $O\left(\mathrm{M}_{q}(d) \log \pi\right)$.

### 2.6. Tangent Graeffe transforms

Whenever a Graeffe transform preserves the number of distinct simple roots, the so called tangent Graeffe transform can be used to directly recover the original simple roots from the simple roots of the transformed polynomial. Our randomized root finding algorithms from [20] heavily rely on this fact. Unfortunately, we have not found a way to exploit tangent Graeffe transforms in the deterministic setting. For completeness, we will nevertheless show that the results of this section extend to the tangent case.

Introducing a formal parameter $\varepsilon$ with $\varepsilon^{2}=0$, we define the generalized tangent Graeffe transform of $g$ of order $\pi$ as being $G_{\pi}(g(x+\varepsilon)) \in\left(\mathbb{F}_{q}[\varepsilon] /\left(\varepsilon^{2}\right)\right)[x]$. For any ring $R$, computations with "tangent numbers" in $R[\varepsilon] /\left(\varepsilon^{2}\right)$ can be done with constant overhead with respect to computations in $R$ (in the FFT model, the overhead is asymptotically limited to a factor of two).

In the context of the Graeffe method, the tangent transform is classical (for instance, see [30, 34] for history, references, and use in numerical algorithms). The generalized tangent Graeffe transform can also be seen as the tangent characteristic polynomial of $x^{\pi}$ modulo $g(x+\varepsilon)$, and this construction is attributed to Kronecker in algebraic geometry [19, 28].

Proposition 1 from Section 2.1 admits a straightforward generalization to tangent Graeffe transforms with a constant overhead. In order to generalize Proposition 3, we introduce $A_{\varepsilon}=\operatorname{GR}\left(p^{\kappa}, k\right)[[\varepsilon]][x] /(g(x+\varepsilon))$, where $\operatorname{GR}\left(p^{\kappa}, k\right)=\left(\left(\mathbb{Z} / p^{\kappa} \mathbb{Z}\right)[z]\right) /(\mu)$, where $\mu \in\left(\mathbb{Z} / p^{k} \mathbb{Z}\right)[z]$ is monic of degree $k$ and irreducible modulo $p$. Let $\operatorname{Tr}_{\varepsilon}$ represent the trace function in $A_{\varepsilon}$.

Proposition 10. Let $\delta>0$ be a fixed constant. If $d^{1+\delta} \leqslant q-1$, then the traces $\operatorname{Tr}_{\varepsilon}(e+\varepsilon \bar{e})$, $\operatorname{Tr}_{\varepsilon}\left((e+\varepsilon \bar{e})^{2}\right), \ldots, \operatorname{Tr}_{\varepsilon}\left((e+\varepsilon \bar{e})^{d}\right)$ modulo $\varepsilon^{2}$, for a given $e+\varepsilon \bar{e} \in A_{\varepsilon}$, can be computed in time $\left(d \log \left(q^{\kappa}\right)\right)^{1+\delta}$.

Proof. The value $\operatorname{Tr}_{\varepsilon}\left((e+\varepsilon \bar{e})^{i}\right)$ equals the trace of $(e(x-\varepsilon)+\varepsilon \bar{e}(x-\varepsilon))^{i}$ computed in $\operatorname{GR}\left(p^{\kappa}, k\right)[[\varepsilon]][x] /(g)$. Therefore we obtain:

$$
\operatorname{Tr}_{\varepsilon}\left((e(x)+\varepsilon \bar{e}(x))^{i}\right)=\operatorname{Tr}\left((e(x-\varepsilon)+\varepsilon \bar{e}(x-\varepsilon))^{i}\right)=\operatorname{Tr}\left(e^{i}(x)+i \varepsilon\left(\bar{e}(x)-e^{\prime}(x)\right) e^{i-1}(x)\right) .
$$

We therefore use Theorem 7.7 of [27, p. 1792] twice over $A=\operatorname{GR}\left(p^{k}, k\right)[x] /(g)$ : once with $\operatorname{Tr}$, and once for $\operatorname{Tr} \circ \varphi$, where $\varphi$ is the product by $\bar{e}-e^{\prime}$.

Theorem 11. Let $g \in \mathbb{F}_{q}[x]$ be a monic polynomial of degree d, and let e, $\bar{e} \in \mathbb{F}_{q}[x]$ be of degrees $<d$. For every constant $\delta>0$ such that $d^{1+\delta} \leqslant q-1$, the characteristic polynomial $h(x)+\varepsilon \bar{h}(x)$ of the multiplication by $e(x)+\varepsilon \bar{e}(x)$ in $\mathbb{F}_{q}[[\varepsilon]][x] /(g(x+\varepsilon))$ can be computed in time $(d \log q)^{1+\delta}$.

Proof. We embed $\left.\mathbb{F}_{q}[\varepsilon \varepsilon]\right][x] /(g)$ into $A_{\varepsilon}=\operatorname{GR}\left(p^{\kappa}, k\right)[[\varepsilon]][x] /(g(x+\varepsilon))$, where $\operatorname{GR}\left(p^{\kappa}, k\right)$ is constructed from $\left(\mathbb{Z} / p^{\kappa} \mathbb{Z}\right)[z] /(\mu)$, and $\mu \in\left(\mathbb{Z} / p^{\kappa} \mathbb{Z}\right)[z]$ is the monic defining polynomial of degree $k$ of $\mathbb{F}_{q}$ over $\mathbb{F}_{p}$.

Let $\delta>0$ be such that $d^{1+\delta} \leqslant q-1$. Proposition 10 allows us to compute the traces of the successive powers of $e+\varepsilon \bar{e}$ seen in $A_{\varepsilon}$ in time $\left(d \log \left(q^{\kappa}\right)\right)^{1+\delta}$. Then Proposition 6 allows us to deduce $h+\varepsilon \bar{h}$ modulo $p$ in softly linear time, mutatis mutandis.

Corollary 12. Let $g$ be a monic polynomial in $\mathbb{F}_{q}[x]$ of degree d. For every constant $\delta>0$ such that $d^{1+\delta} \leqslant q-1$, the generalized tangent Graeffe transform $G_{\pi}(g(x+\varepsilon))$ can be computed in time $(d \log q)^{1+\delta}+\tilde{O}(d \log q \log \pi)$.

Proof. We apply the latter theorem to $e(x)+\varepsilon \bar{e}(x)=x^{\pi}$ rem $g(x+\varepsilon)$, that can be computed in time $O\left(\mathrm{M}_{q}(d) \log \pi\right)$.

## 3. Root finding based on Graeffe transforms

Our new root finding algorithm is technically reminiscent of numerical solvers: we make use of the Graeffe transform instead of the modular exponentiation, and replace gcds by multi-point evaluations. We begin this section with a first simple version of the algorithm, and then introduce new ingredients in order to improve it. Throughout this section, the quantity $\delta>0$ represents a constant that can be fixed arbitrarily small.

### 3.1. Simple version of the deterministic algorithm

We assume given integers $\pi_{i} \geqslant 2$ such that $q-1=\rho \pi_{1} \cdots \pi_{m}$. We let $\chi=\pi_{1} \cdots \pi_{m}$, and $\sigma=\pi_{1}+\cdots+\pi_{m}$ as previously. The core algorithm of this section computes a sequence of Graeffe transforms, and then deduces inductively their roots starting from the last transform and finishing to the input polynomial.

## Algorithm 13.

Input. $f \in \mathbb{F}_{q}[x]$ of degree $d$, monic, separable, which splits over $\mathbb{F}_{q}$, and such that $f(0) \neq 0$; a primitive root $\xi$ of unity of order $\chi$.
Output. $e_{1}, \ldots, e_{s}$ in $\{0, \ldots, \chi-1\}$, such that $f$ divides $\prod_{i=1}^{s}\left(x^{\rho}-\xi^{e_{i}}\right)$.

1. Compute the Graeffe transform $h_{0}$ of order $\rho$ of $f$, and recursively compute $h_{i}$ as the Graeffe transform of order $\pi_{i}$ of $h_{i-1}$, for all $1 \leqslant i \leqslant m-1$.
2. Initialize $E$ with the list $[0]$.
3. For $i$ from $m$ down to 1 do
a. Replace $E$ by the concatenation of $\left[(e+j \chi) / \pi_{i} \mid e \in E\right]$ for $j$ from 0 to $\pi_{i}-1$.
b. If $E$ has cardinality more than $d$ then remove the elements $e$ from $E$ such that $h_{i-1}\left(\xi^{e}\right) \neq 0$, using a multi-point evaluation to compute the $h_{i-1}\left(\xi^{e}\right)$.

## 4. Return $E$.

Proposition 14. Algorithm 13 is correct. If $d^{1+\delta} \leqslant q-1$ then it costs $\sigma O\left(\mathrm{M}_{q}(d) \log d\right)+$ $\left(d \log ^{2} q\right)^{1+\delta}+O\left(m d \mathbf{I}(\log \chi)+m \mathrm{~m}_{q} d \log \chi\right)=\sigma \tilde{O}(d \log q)+\left(d \log ^{2} q\right)^{1+\delta}+\tilde{O}\left(d \log ^{3} q\right)$. Whenever a primitive root of unity of order $\rho$ is given, and $\mu:=\max \left(\rho, \pi_{1}, \ldots, \pi_{m}\right)=$ $O(\log q)$, this bound further reduces to $\tilde{O}\left(d \log ^{3} q\right)$.

Proof. We prove the correctness by descending induction on $i$ from $m$ down to 1 . At the end of iteration $i$ of the loop of step 3 , we claim that $f$ divides $\prod_{e \in E}\left(x^{\rho \pi_{1} \cdots \pi_{i-1}}-\xi^{e}\right)$, and that the integers $e$ in $E$ are in $\{0, \ldots, \chi-1\}$ and are divisible by $\rho \pi_{1} \cdots \pi_{i-1}$. These properties all hold for $i=m+1$ when entering step 3 . By induction, we may assume that these properties hold with $i+1$ when entering the loop at level $i$. Let $e \in E$, and let $\alpha$ be a root of $f$ such that $\alpha^{\rho \pi_{1} \cdots \pi_{i}}=\xi^{e}$. Since $e$ is divisible by $\pi_{i}$, there exists $j \in\left\{0, \ldots, \pi_{i}-1\right\}$ such that $\alpha^{\rho \pi_{1} \cdots \pi_{i-1}}=\xi^{(e+j \chi) / \pi_{i}}$. This proves the correctness.

From now on we assume that $d^{1+\delta} \leqslant q-1$. As to the complexity, we can compute the $\chi / \pi_{i}$ for $i \in\{1, \ldots, m\}$ in time $O(m \mathbf{I}(\log \chi))$ as follows: we first compute $\pi_{1}, \pi_{1} \pi_{2}$, $\pi_{1} \pi_{2} \pi_{3}, \ldots, \pi_{1} \pi_{2} \cdots \pi_{m-1}$, and $\pi_{m}, \pi_{m-1} \pi_{m}, \pi_{m-2} \pi_{m-1} \pi_{m}, \ldots, \pi_{2} \cdots \pi_{m-1} \pi_{m}$, and then deduce each $\chi / \pi_{i}$ by multiplying $\pi_{1} \cdots \pi_{i-1}$ and $\pi_{i+1} \cdots \pi_{m}$.

Step 1 requires time $\left(d \log ^{2} q\right)^{1+\delta}$ by Corollary 9. In step 3.a, we compute all the $e / \pi_{i}$ in time $O(d \mathrm{I}(\log \chi))$, and then all the $\xi^{e / \pi_{i}}$ and $\xi^{\chi / \pi_{i}}$ by means of $O(d \log \chi)$ operations in $\mathbb{F}_{q}$. Deducing all the $\xi^{(e+j \chi) / \pi_{i}}$ takes $O\left(\pi_{i} d\right)$ additional products in $\mathbb{F}_{q}$. The multi-point evaluations can be done in time $O\left(\pi_{i} \mathrm{M}_{q}(d) \log d\right)$.

Finally, when $\mu=O(\log q)$, the term $\left(d \log ^{2} q\right)^{1+\delta}$ can be replaced by $(\rho+$ $\sigma) \tilde{O}(d \log q)=\tilde{O}\left(\sigma d \log q+d \log ^{2} q\right)=\tilde{O}\left(d \log ^{3} q\right)$ in view of Proposition 1.

The cost in the latter proposition is asymptotically higher than the one of MignotteSchnorr's algorithm [20, Section 2], because of the term $\left(d \log ^{2} q\right)^{1+\delta}$. In the smooth case when $\mu=O(\log q)$, the asymptotic complexity bounds become similar.

The rest of this section is devoted to improve Algorithm 13. In order to diminish the dependence on $\log q$, we stop representing roots by their logarithms as in MignotteSchnorr's algorithm. Instead we shall compute discrete logarithms only when necessary. Furthermore, we will show how to decrease the dependence on the $\pi_{i}$ to $\sqrt{\pi_{i}}$, by using the "baby-step giant-step" technique.

### 3.2. Reverse splitting of Graeffe transforms

Let $g$ be a monic polynomial which splits over $\mathbb{F}_{q}$. Once the roots $a_{1}, \ldots, a_{l}$ of the Graeffe transform of order $\pi$ of $g$ have been recovered, we may decompose $g$ into $l$ factors $g_{1}, \ldots, g_{l}$ such that the $\pi$-th powers of the roots of $g_{i}$ all coincide to $a_{i}$, for all $i \in\{1, \ldots, l\}$. In short, we have $\operatorname{Gr}_{\pi}\left(g_{i}\right)(x) \backsim\left(x-a_{i}\right)^{\operatorname{deg} g_{i}}$. This decomposition may be achieved efficiently via the following divide and conquer algorithm.

## Algorithm 15.

Input. $g \in \mathbb{F}_{q}[x]$ of degree $d$, monic, separable, which splits over $\mathbb{F}_{q}$, and such that $g(0) \neq 0$; the pairwise distinct roots $a_{1}, \ldots, a_{l}$ of the Graeffe transform of a given order $\pi \geqslant 2$ of $g$.
Output. The sequence of the monic $\operatorname{gcd}\left(x^{\pi}-a_{i}, g(x)\right)$ for all $i \in\{1, \ldots, l\}$.

1. If $l=1$ then return $g$.
2. Let $h:=\lfloor l / 2\rfloor$.
3. Compute $e(x):=\prod_{i=1}^{h}\left(x-a_{i}\right)$.
4. Compute $g_{1}(x):=\operatorname{gcd}\left(e\left(x^{\pi}\right), g(x)\right)$ and $g_{2}(x):=g(x) / g_{1}(x)$.
5. Call recursively the algorithm with input $g_{1}, a_{1}, \ldots, a_{h}$ and $g_{2}, a_{h+1}, \ldots, a_{l}$, and return the concatenation of the polynomial sequences returned so.

Proposition 16. Algorithm 15 is correct and costs $\pi \tilde{O}(d \log q \log \pi)$ or $(d \log q)^{1+\delta}+$ $\tilde{O}(d \log \pi \log q)$, for any constant $\delta>0$.

Proof. If $l=1$, then the $\pi$-th powers of the roots of $g$ are all equal to $a_{1}$, whence the correctness. If $l>1$, then the roots of $g_{1}$ are exactly those of $g$ whose $\pi$-th powers are in the set $\left\{a_{1}, \ldots, a_{h}\right\}$. Consequently, the $\pi$-th powers of the roots of $g_{2}$ are all in the set $\left\{a_{h+1}, \ldots, a_{l}\right\}$. This completes the proof of the correctness by induction.

Step 3 can be performed in softly linear time by the subproduct tree technique (for the sake of efficiency the subproduct tree should be shared between all the recursive calls).

On the one hand, computing $g_{1}$ can be done naively in time $O(\mathrm{M}(\pi d)+\mathrm{M}(d) \log d)=$ $\pi \tilde{O}(d \log q \log \pi)$. On the other hand, one may proceed as follows: compute $r(x):=$ $x^{\pi}$ rem $g(x)$ in time $\tilde{O}(d \log \pi \log q)$, and then use [27, Corollary 7.2, p. 1789] to obtain $e(r(x))$ rem $g(x)$ in time $(d \log q)^{1+\delta}$.

The sum of the degrees of the input polynomials $g$ at the same depth of the recursive calls is at most $d$, and the maximal depth is $O(\log l)=O(\log d)$. This yields the complexity bounds of the algorithm.

### 3.3. Finding roots of several polynomials in a given set

Let $g_{1}, \ldots, g_{l}$ be polynomials of respective degrees $d_{1}, \ldots, d_{l}$, and let $d:=d_{1}+\cdots+d_{l}$. For given subset $\mathcal{A}$ of $\mathbb{F}_{q}$ of cardinality $\leqslant d$, we are interested in finding the roots of $g_{i}$ in $\mathcal{A}$ for each $i \in\{1, \ldots, l\}$ simultaneously. We do not make additional assumptions, and in particular the $g_{i}$ are allowed to share common roots. A natural divide and conquer approach applies, as explained in the following algorithm.

## Algorithm 17.

Input. $g_{1}, \ldots, g_{l} \in \mathbb{F}_{q}[x]$ of degrees $\geqslant 1$; a subset $\mathcal{A}$ of $\mathbb{F}_{q}$ of cardinality at most $d:=$ $\Sigma_{i=1}^{l} \operatorname{deg} g_{i}$.
Output. The sequence of the sets of the roots of $g_{1}, \ldots, g_{l}$ in $\mathcal{A}$.

1. If $l=1$, then evaluate $g_{1}$ at all the points of $\mathcal{A}$, and return those which are roots.
2. Let $h:=\lfloor l / 2\rfloor$.
3. Compute $e_{1}:=\prod_{i=1}^{h} g_{i}$ and $e_{2}:=\prod_{i=h+1}^{l} g_{i}$.
4. Compute the subset $\mathcal{A}_{1}$ of points of $\mathcal{A}$ which are roots of $e_{1}$.
5. Compute the subset $\mathcal{A}_{2}$ of points of $\mathcal{A}$ which are roots of $e_{2}$.
6. Call recursively the algorithm with input $g_{1}, \ldots, g_{h}, \mathcal{A}_{1}$ and $g_{h+1}, \ldots, g_{l}, \mathcal{A}_{2}$, and return the concatenation of the returned sequences of sets of roots.

Proposition 18. Algorithm 17 is correct and costs $O\left(\mathrm{M}_{q}(d) \log d \log l\right)=\tilde{O}(d \log q)$.
Proof. For the correctness, it is clear that the roots of $g_{1}, \ldots, g_{h}$ in $\mathcal{A}$ are to be found in $\mathcal{A}_{1}$ - idem for $\mathcal{A}_{2}$. The recursive calls are legitimate since the cardinality of $\mathcal{A}_{1}$ (resp. of $\mathcal{A}_{2}$ ) cannot exceed $\operatorname{deg} g_{1}$ (resp. $\operatorname{deg} g_{2}$ ).

Steps 1 to 5 take $O\left(\mathrm{M}_{q}(d) \log d\right)$. The sum of the degrees of the input polynomials at the same depth of the recursive calls is at most $d$, and the maximal depth is $O(\log l)$. This yields the complexity bound, thanks to the super-additivity of $\mathrm{M}_{q}$.

Notice that this lemma generalizes to arbitrary fields.

### 3.4. Simultaneous baby-step giant-step root finding

Instead of using multi-point evaluations in step 3.b of Algorithm 13, we may split the separable part of $h_{i-1}$ thanks to Algorithm 15. We will detail later how this situation further reduces to computing all the roots of several polynomials of degree sum at most $d$, and having roots of order dividing $\pi$. To perform efficient root finding in this case, we adapt Pollard-Strassen's aforementioned strategy for integer factoring as follows.

## Algorithm 19.

Input. $g_{1}, \ldots, g_{l} \in \mathbb{F}_{q}[x]$ of degrees $\geqslant 1$, monic, separable, which split over $\mathbb{F}_{q}$, and such that $g_{j}(0) \neq 0$ holds for all $j \in\{1, \ldots, l\}$; a multiple $\pi$ of the orders of the roots of all the $g_{j}$; a primitive root $\xi$ of unity of order $\pi$.
Output. The sets $L_{j}$ of the $\xi$-logarithms of the roots of $g_{j}$ for all $j \in\{1, \ldots, l\}$.

1. Let $d:=\sum_{j=1}^{l} \operatorname{deg} g_{j}$, let $s:=\lfloor\sqrt{\pi d}\rfloor$, and $t:=\lceil\pi / s\rceil$.
2. Compute $h(x):=\prod_{j=0}^{s-1}\left(\xi^{j} x-1\right)$.
3. For all $j \in\{1, \ldots, l\}$ and all $i \in\{0, \ldots, t-1\}$, compute $r_{j, i}(x):=h(x) \operatorname{rem} g_{j}\left(\xi^{s i} x\right)$.
4. For all $j \in\{1, \ldots, l\}$, compute the set $I_{j}$ of indices $i \in\{0, \ldots, t-1\}$ such that $r_{j, i}(x)$ has a proper $\operatorname{gcd} e_{j, i}(x)$ with $g_{j}\left(\xi^{s i} x\right)$.
5. For all $j \in\{1, \ldots, l\}$ and all $i \in I_{j}$, compute all the $\xi$-logarithms $E_{j, i}$ of the roots of $e_{j, i}$ in $\mathcal{A}=\left\{1, \xi, \xi^{2}, \ldots, \xi^{s-1}\right\}$ by calling Algorithm $17\lceil s / d\rceil$ times on subsets of $\mathcal{A}$ of cardinalities at most $d$.
6. Return the sets $L_{j}:=\bigcup_{i \in I_{j}}\left(s i+E_{j, i}\right)$, for all $j \in\{1, \ldots, l\}$.

Proposition 20. Algorithm 19 is correct and costs $\tilde{O}(\sqrt{\pi d} \log q)+\tilde{O}(d \log q)$.
Proof. Of course, the returned values in $L_{j}$ are $\xi$-logarithms of roots of $g_{j}$. Conversely, let $\xi^{u}$ be a root of $g_{j}$ with $0 \leqslant u \leqslant \pi-1$. There exists a unique $i$ in $\{0, \ldots, t-1\}$ such that $i s \leqslant u<(i+1) s$, since $s t \geqslant \pi$. Then $\xi^{u-i s}$ is a common root of $g_{j}\left(\xi^{s i} x\right)$ and $h(x)$. This proves the correctness.

By Lemma 2, step 2 executes in time $O\left(\mathrm{M}_{q}(s)\right)=\tilde{O}(\sqrt{\pi d} \log q)$. Since $d t=s+O(d)$, step 3 costs

$$
O\left(\mathrm{M}_{q}(s) \log (s)+\mathrm{M}_{q}(d) \log d\right)=\tilde{O}(\sqrt{\pi d} \log q)+\tilde{O}(d \log q) .
$$

Using the super-additivity of $\mathrm{M}_{q}$, step 4 needs $O\left(t \mathrm{M}_{q}(d) \log d\right)=\tilde{O}(\sqrt{\pi d} \log q)+$ $\tilde{O}(d \log q)$. Since the sum of the degrees of all the polynomials $e_{j, i}$ is at most $d$, Proposition 18 implies a cost $\tilde{O}(s \log q+d \log q)$ for step 5 .

If $l=1$, then Algorithm 19 slightly improves Shoup's variant of Pollard-Strassen's algorithm described in [43], which takes $s=\sqrt{\pi}$ independently of $d$; taking $s=\sqrt{\pi d}$ turns out to be a better trade-off between the baby steps and the giant steps. We also notice that this trade-off might be further improved by introducing logarithm factors in $s$.

### 3.5. Fast discrete logarithm in smooth finite fields

Let us recall that the term $\tilde{O}\left(d \log ^{3} q\right)$ in the complexity of Algoritm 13 is due to computing roots from their logarithms. In the next subsection we will proceed differently and will compute logarithms only when needed. This is why we now address the discrete logarithm problem. Let us recall that the seminal discrete logarithm method in the smooth case of $\mathbb{F}_{p}^{*}$ goes back to [36] and performs $\sqrt{S_{1}(p-1)} O(\log p)$ operations in $\mathbb{F}_{p}$, making use internally of Shanks' baby-step giant-step paradigm. We recall this algorithm for completeness, adapting it to our context of Turing machines.

## Algorithm 21.

Input. A primitive root $\xi \in \mathbb{F}_{q}^{*}$ of order $\pi \geqslant 2$ of unity; $a$ in the multiplicative subgroup generated by $\xi$.
Output. The $\xi$-logarithm of $a$ in $\{0, \ldots, \pi-1\}$.

1. Let $s:=\lfloor\sqrt{\pi}\rfloor$ and $t:=\lceil\pi / s\rceil$.
2. Compute $1, \xi, \xi^{2}, \ldots, \xi^{s-1}$.
3. Compute $a, a \xi^{-s}, a \xi^{-2 s}, \ldots, a \xi^{-(t-1) s}$.
4. Sort the array made of the pairs $\left(i, \xi^{i}\right)$ for $i \in\{0, \ldots, s-1\}$ and $\left(j, a \xi^{-j s}\right)$ for $j \in\{0, \ldots, t-1\}$, accordingly to the lexicographic order of the second entries of the pairs seen as $k$-tuples in $\{0, \ldots, p-1\}^{k}$.
5. Find the first two consecutive pairs $\left(e_{1}, \xi^{e_{1}}\right)$ and ( $\left.e_{2}, a \xi^{-e_{2} s}\right)$ whose second entries coincide.
6. Return $e_{1}+s e_{2}$ rem $\pi$.

Proposition 22. Algorithm 21 is correct and requires a running time $\sqrt{\pi} \tilde{O}(\log q)+$ $\sqrt{\pi} \log \pi O(\log q)$.

Proof. Since $s t \geqslant \pi$, the discrete $\operatorname{logarithm} e=\log _{\xi} a$ can be uniquely written as $e=e_{2} s+e_{1}$ with $0 \leqslant e_{1}<s$ and $0 \leqslant e_{2}<t$. This proves the correctness.

Steps 2 and 3 perform $O(\sqrt{\pi})$ products in $\mathbb{F}_{q}$ and one inversion. Using the merge-sort algorithm, step 4 costs $O(\sqrt{\pi} \log q \log \pi)$.

The following divide and conquer approach allows us to compute logarithms for composite orders $\chi=\pi_{1} \cdots \pi_{m}$. Independently from the previous algorithm, it leads to a good complexity bound in terms of $\sigma=\pi_{1}+\cdots+\pi_{m}$. This approach might already be known, but we could not find a reference to it in the literature.

## Algorithm 23.

Input. $\xi \in \mathbb{F}_{q}^{*}$ of order $\chi=\pi_{1} \cdots \pi_{m}$, with given $\pi_{i} \geqslant 2 ; a$ in the multiplicative subgroup generated by $\xi$.
Output. The $\xi$-logarithm of $a$ in $\{0, \ldots, \chi-1\}$.

1. If $m=1$, then find and return $e \in\left\{0, \ldots, \pi_{1}-1\right\}$ such that $a=\xi^{e}$.
2. Let $m_{1}:=\lfloor m / 2\rfloor, m_{2}:=m-m_{1}, \chi_{1}:=\pi_{1} \cdots \pi_{m_{1}}$ and $\chi_{2}:=\chi / \chi_{1}$.
3. Recursively compute $e_{1}:=\log _{\xi^{\chi_{1}}} a^{\chi_{1}}$.
4. Recursively compute $e_{2}:=\log _{\xi^{\chi_{2}}}\left(a / \xi^{e_{1}}\right)$.
5. Return $e_{1}+\chi_{2} e_{2}$.

Proposition 24. Algorithm 23 is correct and executes in time $O\left(\sigma \mathrm{~m}_{q}+\mathrm{d}_{q} \log \chi+\right.$ $\left.\mathrm{m}_{q} \log \chi \log \log \chi+\mathrm{I}(\log \chi) \log \log \chi\right)=\sigma \tilde{O}(\log q)+\tilde{O}(\log q \log \chi)$. By using Algorithm 21 in step 1 , the cost becomes $\sqrt{\mu} \tilde{O}\left(\log ^{2} q\right)$, where $\mu:=\max \left(\pi_{1}, \ldots, \pi_{m}\right)$.

Proof. We prove the correctness by induction. The case $m=1$ is clear. Assume $m \geqslant 2$. In step 3 , we have $e_{1} \in\left\{0, \ldots, \chi_{2}-1\right\}$. By definition of $e_{1}$, we have $\left(a / \xi^{e_{1}}\right)^{\chi_{1}}=a^{\chi_{1}} / \xi^{\chi_{1} e_{1}}=1$ in step 4 , whence $a / \xi^{e_{1}}$ indeed lies in the multiplicative subgroup of order $\chi_{1}$ generated by $\xi^{\chi_{2}}$. By definition of $e_{2}$, we have $a / \xi^{e_{1}}=\left(\xi^{\chi_{2}}\right)^{e_{2}}$, whence $a=\xi^{e_{1}+\chi_{2} e_{2}}$, which completes the induction.

As to the complexity bound, let $T(\chi)$ denote the cost of the algorithm. We share the subproduct tree built from $\pi_{1}, \ldots, \pi_{m}$, in time $O(\mathrm{I}(\log \chi) \log m)$, between the recursive calls. Using binary powering in steps 3 and 4, we obtain the recursion relation

$$
T(\chi) \leqslant O\left(\mathrm{~m}_{q} \log \chi+\mathrm{d}_{q}\right)+T\left(\chi_{1}\right)+T\left(\chi_{2}\right)
$$

which leads to the bound $T(\chi) \leqslant \sum_{i=1}^{m} T\left(\pi_{i}\right)+O\left(m \mathrm{~d}_{q}+\mathrm{m}_{q} \log \chi \log m\right)$.
Using a naive exhaustive search in step 1 yields $T\left(\pi_{i}\right)=\pi_{i} \mathrm{~m}_{q}$, whence $T(\chi)=O\left(\sigma \mathrm{~m}_{q}+\right.$ $\left.m \mathrm{~d}_{q}+\mathrm{m}_{q} \log \chi \log m+\mathrm{I}(\log \chi) \log m\right)$. The first bound thus follows using $m=O(\log \chi)$. The second bound follows from Proposition 22 that gives $T\left(\pi_{i}\right)=\sqrt{\pi_{i}} \log \pi_{i} \tilde{O}(\log q)$.

Remark 25. As a byproduct of the complexity analysis, we remark that, given a primitive root of unity $\xi$ of order $\chi=\pi_{1} \cdots \pi_{m}$, where the $\pi_{i}$ are given integers $\geqslant 2$, one may compute primitive roots of unity $\xi_{i}$ for all orders $\pi_{i}$ in time $\tilde{O}(\log q \log \chi)$.

### 3.6. Squarefree part

Over any field $\mathbb{K}$ the separable factorization of a univariate polynomial $g$ of degree $d$ can computed by means of $O(\mathrm{M}(d) \log d)$ operations in $\mathbb{K}[29]$. If $\mathbb{K}$ is the finite field $\mathbb{F}_{q}$ then it is possible to deduce the squarefree factorization of $g$ with $O(d)$ additional extractions of $p$-th roots [29, Section 4]. This result is also left as an exercise in [17, Chapter 14, Exercise 14.30]. In the special case when $k=1$ and $p>d$, no root extraction is necessary and the squarefree part of $g$ is obtained as $g / \operatorname{gcd}\left(g, g^{\prime}\right)$. In the next subsection we will need the following bound, taking the number of simple and multiple roots into account.

Proposition 26. Let $g \in \mathbb{F}_{q}[x]$ be of degree d. In an algebraic closure of $\mathbb{F}_{q}$, let $\delta_{1}$ be the number of simple roots of $g$, let $\delta \geqslant 2$ be the number of its multiple roots, and let $\Delta_{\geqslant 2}=d-\delta_{1}$ be the number of multiple roots counted with multiplicities. Then the squarefree factorization and the squarefree part of $f$ may be computed in time $\tilde{O}(d \log q)+\Delta \geqslant 2 \tilde{O}\left(\log ^{2} q\right)$.

Proof. Let $\left(g_{1}, q_{1}, \nu_{1}\right), \ldots,\left(g_{l}, q_{l}, \nu_{l}\right)$ in $\left(\mathbb{F}_{q}[x] \backslash \mathbb{F}_{q}\right) \times\left\{1, p, p^{2}, \ldots,\right\} \times \mathbb{N}$ represent the separable factorization of $g\left[29\right.$, Introduction], so that we have $g(x)=g_{1}\left(x^{q_{1}}\right)^{\nu_{1}} \cdots g_{l}\left(x^{q_{l}}\right)^{\nu_{l}}$. Since $\mathbb{F}_{q}$ is perfect, the factors $g_{i}\left(x^{q_{i}}\right)$ with $q_{i}=\nu_{i}=1$ contain all the simple roots of $g$. Letting $h_{1}=\prod_{q_{i}=\nu_{i}=1} g_{i}$ and $h \geqslant 2=g / h_{1}$, we have $\operatorname{deg} h_{1}=\delta_{1}$ and $\operatorname{deg} h \geqslant 2=\Delta_{\geqslant 2}$. Following [29, Section 4], the actual number of $p$-th root extractions is $O\left(\Delta_{\geqslant 2}\right)$. Each such extraction amounts to $O(\log q)$ products in $\mathbb{F}_{q}$. In this way we obtain the squarefree factorization of $g$.

### 3.7. Improved root finding algorithm

We are now ready to combine all the previous algorithms of this section in order to improve Algorithm 13. Recall that we are given integers $\pi_{i} \geqslant 2$ such that $q-1=\rho \pi_{1} \cdots \pi_{m}$, and that we let $\chi=\pi_{1} \cdots \pi_{m}$. The following algorithm is parametrized by subroutines used internally. We distinguish three cases of use, leading to complexity bounds in terms of $\sigma$, of $\mu$, and in the smooth case of $S_{1}(q-1)=O(1)$.

## Algorithm 27.

Input. $f \in \mathbb{F}_{q}[x]$, monic, separable, which splits over $\mathbb{F}_{q}$, of degree $d<q$, such that $f(0) \neq 0$; a primitive root $\zeta$ of $\mathbb{F}_{q}^{*}$.
Output. The roots of $\operatorname{Gr}_{\rho}(f)$.

1. Compute $h_{0}$ as the Graeffe transform of $f$ of order $\rho$. Compute $\tilde{h}_{0}$ as the squarefree part of $h_{0}$. Recursively compute $h_{i}$ as the Graeffe transform of order $\pi_{i}$ of $\tilde{h}_{i-1}$, and $\tilde{h}_{i}$ as the squarefree part of $h_{i}$, for all $i$ from 1 to $m-1$.
2. Initialize $Z$ with $\{1\}$, and compute $\xi=\zeta^{\rho}$.
3. For $i$ from $m$ down to 1 do
a. Use Algorithm 15 with input $\tilde{h}_{i-1}$ and $Z=\left\{a_{1}, \ldots, a_{l}\right\}$, and let $g_{1}, \ldots, g_{l}$ be the polynomials obtained in return.
b. Initialize $W_{0}$ with the roots of all the $g_{j}$ of degree 1 , for $j \in\{1, \ldots, l\}$.
c. For all $j \in\{1, \ldots, l\}$ such that $\operatorname{deg} g_{j} \geqslant 2$, compute the $\xi$-logarithm $e_{j}$ of $a_{j}$ with Algorithm 23 (resp. with Algorithms 23 and 21).
d. Use Algorithm 17 (resp. Algorithm 19) to compute the sets of roots $W_{j}$ of $g_{j}\left(\xi^{e_{j} / \pi_{i}} x\right)$ for all $j \in\{1, \ldots, l\}$ such that $\operatorname{deg} g_{j} \geqslant 2$.
e. $Z:=W_{0} \cup \bigcup_{j}\left(\xi^{e_{j} / \pi_{i}} W_{j}\right)$.

## 4. Return $Z$.

Proposition 28. Algorithm 27 is correct and costs $(\sigma+\rho) \tilde{O}(d \log q \log (\mu \rho))+$ $\tilde{O}\left(d \log ^{2} q\right)$. In addition, for any fixed $\delta>0$, the cost is also bounded by $\sigma \tilde{O}(d \log q)+$ $\left(d \log ^{2} q\right)^{1+\delta}$ or $\tilde{O}\left(\sqrt{\mu} d \log ^{2} q\right)+\left(d \log ^{2} q\right)^{1+\delta}$, whenever $d^{1+\delta} \leqslant q-1$ and where $\mu=\max \left(\pi_{1}, \ldots, \pi_{m}\right)$.

Proof. Let introduce $h_{m}:=\operatorname{Gr}_{\pi_{m}}\left(\tilde{h}_{m-1}\right)$, which does not need to be computed. When entering the loop in step 3 , the set $Z$ contains the single root of $h_{m}$. We prove by descending induction on $i$ from $m$ down to 1 that $Z$ contains the roots of $h_{i}$ when entering step $i$ of the loop. Assuming this hypothesis holds for $i \leqslant m$, conditions of Algorithm 15 are met, and finding the roots of $\tilde{h}_{i-1}$ reduces to finding the roots of $g_{1}, \ldots, g_{l}$. The $g_{j}$ of degree 1 contribute to roots of $\tilde{h}_{i-1}$ in a straightforward manner. For the other $g_{j}$, of degree at least 2 , it turns out that the roots of $g_{j}\left(\xi^{e_{j} / \pi_{i}} x\right)$ have order dividing $\pi_{i}$, and thus Algorithms 17 or 19 produce the remaining roots of $\tilde{h}_{i-1}$. This shows the correctness.

First all, the primitive roots of unity of orders $\pi_{i}$ and $\rho$ needed during the execution of the algorithm can be obtained in time $\tilde{O}\left(\log ^{2} q\right)$ by Remark 25. The Graeffe transforms in step 1 may execute in time $(\sigma+\rho) \tilde{O}(d \log q \log (\mu \rho))$ by Proposition 1.

Let $d_{i}:=\operatorname{deg} h_{i}$, let $\delta_{1, i}$ be the number of simple roots of $h_{i}$, let $\delta_{\geqslant 2, i}$ be the number of multiple roots of $h_{i}$, and let $\Delta_{\geqslant 2, i}=d_{i}-\delta_{1, i}$ be the number of multiple roots of $h_{i}$ counted with multiplicities. By Proposition 26, computing the squarefree part $\tilde{h}_{i}$ of $h_{i}$ takes time $\tilde{O}(d \log q)+\Delta_{\geqslant 2, i} \tilde{O}\left(\log ^{2} q\right)$. From $\operatorname{deg} \tilde{h}_{i}=\delta_{1, i}+\delta_{\geqslant 2, i}$ and $\operatorname{deg} \tilde{h}_{i-1}=d_{i}$, we obtain $\operatorname{deg} \tilde{h}_{i-1}-\operatorname{deg} \tilde{h}_{i}=d_{i}-\delta_{1, i}-\delta_{\geqslant 2, i}=\Delta_{\geqslant 2, i}-\delta_{\geqslant 2, i}$. Summing these equalities over $i$ leads to $\operatorname{deg} \tilde{h}_{0}-\operatorname{deg} \tilde{h}_{m}=\sum_{i=1}^{m} \Delta_{\geqslant 2, i}-\sum_{i=1}^{m} \delta_{\geqslant 2, i} . \operatorname{Using} \Delta_{\geqslant 2, i} \geqslant 2 \delta_{\geqslant 2, i}, \operatorname{deg} \tilde{h}_{0}=d$, and $\operatorname{deg} \tilde{h}_{m}=1$, we deduce that $\sum_{i=1}^{m} \delta_{\geqslant 2, i} \leqslant d-1$, and then that $\sum_{i=1}^{m} \Delta_{\geqslant 2, i} \leqslant 2(d-1)$. Consequently, the total cost for computing all the squarefree parts drops to $\tilde{O}\left(d \log ^{2} q\right)$.

Step 3.a may take $\pi_{i} \tilde{O}\left(d \log q \log \pi_{i}\right)$ by Proposition 16 , which yields a total cost $\sigma \tilde{O}(d \log q \log \mu)$. In step 3.c, Algorithm 23 is called $O\left(\delta_{\geqslant 2, i}\right)$ times. Consequently, the total cost of this step is $\sigma d \tilde{O}(\log q)+d \tilde{O}\left(\log ^{2} q\right)$ by Proposition 24. If we use Algorithm 17 in step 3.d, then the cost is $\left\lceil\pi_{i} / d\right\rceil \tilde{O}(d \log q)=\pi_{i} \tilde{O}\left(\log ^{O(1)} d \log q\right)+\tilde{O}(d \log q)$ by Proposition 18. The sum of these costs is bounded by $\sigma \tilde{O}(d \log q)$. So far, this establishes the first complexity bound.

From now on assume that $d^{1+\delta} \leqslant q-1$. The cost for the Graeffe transforms in step 1 may drop into $\left(d \log ^{2} q\right)^{1+\delta}$ by Corollary 9. Step 3 .a may run in time $(d \log q)^{1+\delta}+$ $\tilde{O}\left(d \log \pi_{i} \log q\right)$ by Proposition 16. The total cost of this step thus amounts to $\left(d \log ^{2} q\right)^{1+\delta}$. Putting these bounds together yields the second complexity.

For the third complexity bound, we use Algorithm 23 combined with Algorithm 21 in step 3.c, so that Proposition 24 ensures a time $\sqrt{\mu} d \tilde{O}\left(\log ^{2} q\right)$. In addition we use Algorithm 19 in step 3.d, and Proposition 20 gives a total cost $\tilde{O}\left(\sqrt{\mu d} \log ^{2} q\right)+$ $\tilde{O}\left(d \log ^{2} q\right)$.

Recall that $S_{1}(q-1)$ represents the largest prime number in the irreducible factorization of $q-1$. If $S_{1}(q-1)=O(\log q)$ and if $\sigma+\rho=O(\log q)$, then the first complexity bound of Proposition 28 rewrites into $\tilde{O}\left(d \log ^{2} q\right)$, which is thus softly equivalent to the average cost of Cantor-Zassenhaus' algorithm. We are now able to state the main result of this section.

Theorem 29. Given a constant $\delta>0$, the irreducible factorization of $q-1$, a primitive element of $\mathbb{F}_{q}^{*}$, and $f \in \mathbb{F}_{q}[x]$ of degree d, monic, separable, which splits over $\mathbb{F}_{q}$, the roots of $f$ can be computed in time $\tilde{O}\left(\sqrt{S_{1}(q-1)} d \log ^{2} q\right)+\left(d \log ^{2} q\right)^{1+\delta}$.

Proof. Without loss of generality we can assume that $f(0) \neq 0$ and that $d \leqslant q-1$. If $d^{1+\delta}>$ $q-1$, then it suffices to evaluate $f$ at all the elements of $\mathbb{F}_{q}^{*}$ in time $\tilde{O}(d+q)=O\left(d^{1+\delta}\right)$. Consequently, for when $d^{1+\delta} \leqslant q-1$, we appeal to the preceding proposition.

Let us finally mention that, in general, the present method improves [16, Theorem 4.1] (see complexity details at the end of the proof). Nevertheless, if $k=1$, and if $S_{1}(p-1)$ is close to $p$, then Theorem 29 does not improve [42, Theorem 4.1(2)].

## Bibliography

[1] M. Arora, G. Ivanyos, M. Karpinski, and N. Saxena. Deterministic polynomial factoring and association schemes. LMS J. Comput. Math., 17(1):123-140, 2014.
[2] E. Bach. Comments on search procedures for primitive roots. Math. Comp., 66:1719-1727, 1997.
[3] E. R. Berlekamp. Factoring polynomials over finite fields. Bell System Tech. J., 46:1853-1859, 1967.
[4] E. R. Berlekamp. Factoring polynomials over large finite fields. Math. Comp., 24:713-735, 1970.
[5] A. Bostan, Ph. Flajolet, B. Salvy, and É. Schost. Fast computation of special resultants. J. Symbolic Comput., 41(1):1-29, 2006.
[6] A. Bostan, P. Gaudry, and É. Schost. Linear recurrences with polynomial coefficients and application to integer factorization and Cartier-Manin operator. SIAM J. Comput., 36(6):1777-1806, 2007.
[7] A. Bostan, L. Gonzales-Vega, H. Perdry, and É. Schost. Complexity issues on Newton sums of polynomials. Distributed in the digital proceedings of MEGA'05, 2005.
[8] A. Bostan and É. Schost. Polynomial evaluation and interpolation on special sets of points. J. Complexity, 21(4):420-446, 2005.
[9] P. Camion. A deterministic algorithm for factorizing polynomials of $\boldsymbol{F}_{q}[X]$. In Combinatorial mathematics (Marseille-Luminy, 1981), volume 75 of North-Holland Math. Stud., pages 149-157. NorthHolland, Amsterdam, 1983.
[10] D. G. Cantor and E. Kaltofen. On fast multiplication of polynomials over arbitrary algebras. Acta Infor., 28(7):693-701, 1991.
[11] D. G. Cantor and H. Zassenhaus. A new algorithm for factoring polynomials over finite fields. Math. Comp., 36(154):587-592, 1981.
[12] X. Caruso, D. Roe, and T. Vaccon. Tracking p-adic precision. LMS J. Comput. Math., 17:274-294, 2014. Special Issue A, Algorithmic Number Theory Symposium XI.
[13] E. Costa and D. Harvey. Faster deterministic integer factorization. Math. Comp., 83(285):339-345, 2014.
[14] S. Evdokimov. Factorization of polynomials over finite fields in subexponential time under GRH. In Algorithmic number theory (Ithaca, NY, 1994), volume 877 of Lecture Notes in Comput. Sci., pages 209-219. Springer, Berlin, 1994.
[15] Shuhong Gao. On the deterministic complexity of factoring polynomials. J. Symbolic Comput., 31(12): $19-36,2001$.
[16] J. von zur Gathen. Factoring polynomials and primitive elements for special primes. Theoret. Comput. Sci., 52(1-2):77-89, 1987.
[17] J. von zur Gathen and J. Gerhard. Modern computer algebra. Cambridge University Press, 2nd edition, 2003.
[18] J. von zur Gathen and D. Panario. Factoring polynomials over finite fields: A survey. J. Symbolic Comput., 31(1-2):3-17, 2001.
[19] M. Giusti, G. Lecerf, and B. Salvy. A Gröbner free alternative for polynomial system solving. J. Complexity, $17(1): 154-211,2001$.
[20] B. Grenet, J. van der Hoeven, and G. Lecerf. Randomized root finding over finite fields using tangent Graeffe transforms. Technical report, École polytechnique, 2015.
[21] D. Harvey, J. van der Hoeven, and G. Lecerf. Even faster integer multiplication. http://arxiv.org/ abs/1407.3360, 2014.
[22] D. Harvey, J. van der Hoeven, and G. Lecerf. Faster polynomial multiplication over finite fields. http://arxiv.org/abs/1407.3361, 2014.
[23] J. van der Hoeven and G. Lecerf. Sparse polynomial interpolation in practice. ACM Commun. Comput. Algebra, 48(4), 2014. In section "ISSAC 2014 Software Presentations".
[24] Ming-Deh A. Huang. Generalized Riemann hypothesis and factoring polynomials over finite fields. J. Algorithms, 12(3):464-481, 1991.
[25] E. Kaltofen. Polynomial factorization: a success story. In Hoon Hong, editor, ISSAC '03: Proceedings of the 2003 International Symposium on Symbolic and Algebraic Computation, pages 3-4. ACM Press, 2003.
[26] K. S. Kedlaya and C. Umans. Fast modular composition in any characteristic. In A. Z. Broder et al., editors, 49th Annual IEEE Symposium on Foundations of Computer Science 2008 (FOCS '08), pages 146-155. IEEE, 2008.
[27] K. S. Kedlaya and C. Umans. Fast polynomial factorization and modular composition. SIAM J. Comput., 40(6):1767-1802, 2011.
[28] L. Kronecker. Grundzüge einer arithmetischen Theorie der algebraischen Grössen. J. reine angew. Math., 92:1-122, 1882.
[29] G. Lecerf. Fast separable factorization and applications. Appl. Alg. Eng. Comm. Comp., 19(2):135160, 2008.
[30] G. Malajovich and J. P. Zubelli. Tangent Graeffe iteration. Numer. Math., 89(4):749-782, 2001.
[31] M. Mignotte and C. Schnorr. Calcul déterministe des racines d'un polynôme dans un corps fini. C. R. Acad. Sci. Paris Sér. I Math., 306(12):467-472, 1988.
[32] R. T. Moenck. On the efficiency of algorithms for polynomial factoring. Math. Comp., 31:235-250, 1977.
[33] G. L. Mullen and D. Panario. Handbook of Finite Fields. Discrete Mathematics and Its Applications. Chapman and Hall/CRC, 2013.
[34] V. Pan. Solving a polynomial equation: Some history and recent progress. SIAM Rev., 39(2):187220, 1997.
[35] V. Pan. New techniques for the computation of linear recurrence coefficients. Finite Fields Appl., 6:93-118, 2000.
[36] S. C. Pohlig and M. E. Hellman. An improved algorithm for computing logarithms over $\operatorname{GF}(p)$ and its cryptographic significance (corresp.). IEEE Trans. Inform. Theory, 24(1):106-110, 1978.
[37] L. Rónyai. Factoring polynomials modulo special primes. Combinatorica, 9(2):199-206, 1989.
[38] L. Rónyai. Galois groups and factoring polynomials over finite fields. SIAM J. Disc. Math., 5(3):345365, 1992.
[39] Chandan Saha. Factoring polynomials over finite fields using balance test. In S. Albers and P. Weil, editors, 25th International Symposium on Theoretical Aspects of Computer Science, volume 1 of Leibniz International Proceedings in Informatics (LIPIcs), pages 609-620, Dagstuhl, Germany, 2008. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik. http://drops.dagstuhl.de/opus/volltexte/ 2008/1323.
[40] R. Schoof. Elliptic curves over finite fields and the computation of square roots mod p. Math. Comp., 44(170):483-494, 1985.
[41] V. Shoup. On the deterministic complexity of factoring polynomials over finite fields. Inform. Process. Lett., 33(5):261-267, 1990.
[42] V. Shoup. A fast deterministic algorithm for factoring polynomials over finite fields of small characteristic. In S. M. Watt, editor, ISSAC '91: Proceedings of the 1991 International Symposium on Symbolic and Algebraic Computation, pages 14-21. ACM Press, 1991.
[43] V. Shoup. Smoothness and factoring polynomials over finite fields. Inform. Process. Lett., 38(1):3942, 1991.
[44] V. Shoup. Searching for primitive roots in finite fields. Math. Comp., 58:369-380, 1992.
[45] V. Shoup. NTL: A Library for doing Number Theory, 2014. Software, version 8.0.0. http:// www. shoup.net/ntl.
[46] B. Źrałek. Using partial smoothness of $p-1$ for factoring polynomials modulo p. Math. Comp., 79:2353-2359, 2010.


[^0]:    A. Partially supported by a LIX-Qualcomm ${ }^{\circledR}$-Carnot postdoctoral fellowship.

