Taking Roots over High Extensions
of Finite Fields

Javad Doliskani
jdoliska@uwo.ca

Éric Schost
eschost@uwo.ca

Abstract

We present a new algorithm for computing m-th roots over the finite field $\mathbb{F}_q$, where $q = p^n$, with p a prime, and m any positive integer. In the particular case $m = 2$, the cost of the new algorithm is an expected $O(\mathsf{M}(n)\log(p) + \mathsf{C}(n)\log(n))$ operations in $\mathbb{F}_p$, where $\mathsf{M}(n)$ and $\mathsf{C}(n)$ are bounds for the cost of polynomial multiplication and modular polynomial composition. Known results give $\mathsf{M}(n) = O(n\log(n)\log\log(n))$ and $\mathsf{C}(n) = O(n^{1.67})$, so our algorithm is subquadratic in n.

Keywords. Root extraction; square roots; finite field arithmetic.

Mathematics Subject Classification 2010. Primary 11Y16, 12Y05, Secondary 68W30.

1 Introduction

Beside its intrinsic interest, computing m-th roots over finite fields (for m an integer at least equal to 2) has found many applications in computer science. Our own interest comes from elliptic and hyperelliptic curve cryptography; there, square root computations show up in pairing-based cryptography [3] or point-counting problems [8].

Our result in this paper is a new algorithm for computing m-th roots in a degree n extension $\mathbb{F}_q$ of the prime field $\mathbb{F}_p$, with p a prime. Our emphasis is on the case where p is thought to be small, and the degree n grows. Roughly speaking, we reduce the problem to m-th root extraction in a lower degree extension of $\mathbb{F}_p$ (when $m = 2$, we actually reduce the problem to square root extraction over $\mathbb{F}_p$ itself).

Our complexity model. It is possible to describe the algorithm in an abstract manner, independently of the choice of a basis of $\mathbb{F}_q$ over $\mathbb{F}_p$. However, to give concrete complexity estimates, we have to decide which representation we use, the most usual choices being monomial and normal bases. We choose to use a monomial basis, since in particular our implementation is based on the library NTL [20], which uses this representation. Thus, the finite field $\mathbb{F}_q = \mathbb{F}_{p^n}$ is represented as $\mathbb{F}_p[X]/\langle f \rangle$, for some monic irreducible polynomial $f \in \mathbb{F}_p[X]$ of degree n; elements of $\mathbb{F}_q$ are represented as polynomials in $\mathbb{F}_p[X]$ of degree less than n. We will briefly mention the normal basis representation later on.

1

The costs of all algorithms are measured in number of operations $+, \times, \div$ in the base field $\mathbb{F}_p$ (that is, we are using an algebraic complexity model).

We shall denote upper bounds for the cost of *polynomial multiplication* and *modular composition* by respectively $\mathsf{M}(n)$ and $\mathsf{C}(n)$. This means that over any field $\mathbb{K}$, we can multiply polynomials of degree n in $\mathbb{K}[X]$ in $\mathsf{M}(n)$ base field operations, and that we can compute $f(g) \mod h$ in $\mathsf{C}(n)$ operations in $\mathbb{K}$, when f, g, h are degree n polynomials. We additionally require that both M and C are super-linear functions, as in [23, Chapter 8], and that $\mathsf{M}(n) = O(\mathsf{C}(n))$. In particular, since we work in the monomial basis, multiplications and inversions in $\mathbb{F}_q$ can be done in respectively $O(\mathsf{M}(n))$ and $O(\mathsf{M}(n)\log(n))$ operations in $\mathbb{F}_p$, see again [23].

The best known bound for $\mathsf{M}(n)$ is $O(n\log(n)\log\log(n))$, achieved by using Fast Fourier Transform [17, 5]. The most well-known bound for $\mathsf{C}(n)$ is $O(n^{(\omega+1)/2})$, due to Brent and Kung [4], where ω is such that matrices of size n over any field $\mathbb{K}$ can be multiplied in $O(n^\omega)$ operations in $\mathbb{K}$; this estimate assumes that $\omega > 2$, otherwise some logarithmic terms may appear. Using the algorithm of Coppersmith and Winograd [6], we can take $\omega \leq 2.37$ and thus $\mathsf{C}(n) = O(n^{1.69})$; an algorithm by Huang and Pan [10] actually achieves a slightly better exponent of 1.67, by means of rectangular matrix multiplication.

Main result. We will focus in this paper on the case of t-th root extraction, where t is a prime divisor of $q - 1$; the general case of m-th root extraction, with m arbitrary, can easily be reduced to this case (see the discussion after Theorem 1).

The core of our algorithm is a reduction of t-th root extraction in $\mathbb{F}_q$ to t-th root extraction in an extension of $\mathbb{F}_p$ of smaller degree. Our algorithm is probabilistic of Las Vegas type, so its running time is given as an expected number of operations. With this convention, our main result is the following.

Theorem 1. *Let t be a prime factor of $q - 1$, with $q = p^n$, and let s be the order of p in $\mathbb{Z}/t\mathbb{Z}$. Given $a \in \mathbb{F}_q^*$, one can decide if a is a t-th power in $\mathbb{F}_q^*$, and if so compute one of its t-th roots, by means of the following operations:*

- *an expected $O(s\mathsf{M}(n)\log(p) + \mathsf{C}(n)\log(n))$ operations in $\mathbb{F}_p$;*

- *a t-th root extraction in $\mathbb{F}_{p^s}$.*

Thus, we replace t-th root extraction in a degree n extension by a t-th root extraction in an extension of degree $s \leq \min(n, t)$. The extension degree s is the largest one for which t still divides $p^s - 1$, so iterating the process does not bring any improvement: the t-th root extraction in $\mathbb{F}_{p^s}$ must be dealt with by another algorithm. The smaller s is, the better.

A useful special case is $t = 2$, that is, we are taking square roots; the assumption that t divides $q - 1$ is then satisfied for all odd primes p and all n. In this case, we have $s = 1$, so the second step amounts to square root extraction in $\mathbb{F}_p$. Since this can be done in $O(\log(p))$ expected operations in $\mathbb{F}_p$, the total running time of the algorithm is an expected $O(\mathsf{M}(n)\log(p) + \mathsf{C}(n)\log(n))$ operations in $\mathbb{F}_p$.

A previous algorithm by Kaltofen and Shoup [12] allows one to compute t-th roots in $\mathbb{F}_{p^n}$ in expected time $O((\mathsf{M}(t)\mathsf{M}(n)\log(p) + t\mathsf{C}(n) + \mathsf{C}(t)\mathsf{M}(n))\log(n))$; we discuss it further in the next section. This algorithm requires no assumption on t, so it can be used in our

algorithm in the case $s > 1$, for t-th root extraction in $\mathbb{F}_{p^s}$. Then, its expected running time is $O((\mathsf{M}(t)\mathsf{M}(s)\log(p) + t\mathsf{C}(s) + \mathsf{C}(t)\mathsf{M}(s))\log(s))$.

The strategy of using Theorem 1 to reduce from $\mathbb{F}_q$ to $\mathbb{F}_{p^s}$ then using the Kaltofen-Shoup algorithm over $\mathbb{F}_{p^s}$ is never more expensive than using the Kaltofen-Shoup algorithm directly over $\mathbb{F}_q$. For $t = O(1)$, both strategies are within a constant factor; but even for the smallest case $t = 2$, our algorithm has advantages (as explained in the last section). For larger t, the gap in our favor will increase for cases when s is small (such as when t divides $p - 1$, corresponding to $s = 1$).

Finally, let us go back to the remark above, that for any m, one can reduce m-th root extraction of $a \in \mathbb{F}_q^*$ to computing t-th roots, with t dividing $q - 1$; this is well known, see for instance [1, Chapter 7.3]. We write $m = uv$ with $(v, q - 1) = 1$ and $t \mid q - 1$ for every prime divisor t of u, and we assume that a is indeed an m-th power.

- We first compute the v-th root a_0 of a as $a_0 = a^{v^{-1} \bmod q-1}$ by computing the inverse ℓ of $v \bmod q - 1$, and computing an ℓ-th power in $\mathbb{F}_q$. This takes $O(n\mathsf{M}(n)\log(p))$ operations in $\mathbb{F}_p$.

- Let $u = \prod_{i=1}^{d} m_i^{\alpha_i}$ be the prime factorization of u, which we assume is given to us. Then, for $k = 1, \ldots, \alpha_1$, we compute an m_1-th root a_k of a_{k-1} using Theorem 1, so that a_{α_1} is an $m_1^{\alpha_1}$-th root of a_0.

  One should be careful in the choice of the m_1-th roots (which are not unique), so as to ensure that each a_k is indeed an u/m_1^i-th power: if the given a_k is not such a power, we can multiply it by a m_1-th root of unity until we find a suitable one. The root of unity can be found by the algorithm of Theorem 1.

  Once we know a_{α_1}, the same process can be applied to compute an $m_2^{\alpha_2}$-th root of a_{α_1}, and so on.

The first step, taking a root of order v, may actually be the bottleneck of this scheme. When v is small compared to n, it may be better to use here as well the algorithm by Kaltofen and Shoup mentioned above.

Organization of the paper. The next section reviews and discusses known algorithms; Section 3 gives the details of the root extraction algorithm and some experimental results. In all the paper, $(\mathbb{F}_q^*)^t$ denotes the set of t-th powers in $\mathbb{F}_q^*$.

2 Previous work

Let t be a prime factor of $q - 1$. In the rest of this section, we discuss previous algorithms for t-th root extraction, with a special focus on the case $t = 2$ (square roots), which has attracted most attention in the literature.

We shall see in Section 3 that given such a prime t, the cost of testing for t-th power is always dominated by the t-th root extraction; thus, for an input $a \in \mathbb{F}_q^*$, we always assume that $a \in (\mathbb{F}_q^*)^t$.

All algorithms discussed below rely on some form of exponentiation in $\mathbb{F}_q$, or in an extension of $\mathbb{F}_q$, with exponents that grow linearly with q. As a result, a direct implementation using binary powering uses $O(\log(q))$ multiplications in $\mathbb{F}_q$, that is, $O(n\mathsf{M}(n)\log(p))$ operations in $\mathbb{F}_p$. Even using fast multiplication, this is quadratic in n; alternative techniques should be used to perform the exponentiation, when possible.

Some special cases of square root computation. If G is a group with an odd order s, then the mapping $f : G \to G$, $f(a) = a^2$ is an automorphism of G; hence, every element $a \in G$ has a unique square root, which is $a^{(s+1)/2}$. Thus, if $q \equiv 3 \pmod 4$, the square root of any $a \in (\mathbb{F}_q^*)^2$ is $a^{(q+1)/4}$; this is because $(\mathbb{F}_q^*)^2$ is a group of odd order $(q-1)/2$.

More complex schemes allow one to compute square roots for some increasingly restricted classes of prime powers q. The following table summarizes results known to us; in each case, the algorithm uses $O(1)$ exponentiations and $O(1)$ additions / multiplications in $\mathbb{F}_q$. The table indicates what exponents are used in the exponentiations.

Table 1: Some special cases for square roots

Algorithm	q	exponent
folklore	3 (mod 4)	$(q+1)/4$
Atkin	5 (mod 8)	$(q-5)/8$
Müller [14]	9 (mod 16)	$(q-1)/4$ and $(q-9)/16$
Kong *et al.* [13]	9 (mod 16)	$(q-9)/8$ and $(q-9)/16$

As was said above, using a direct binary powering approach to exponentiation, all these algorithms use $O(n\mathsf{M}(n)\log(p))$ operations in $\mathbb{F}_p$.

Cipolla's square root algorithm. To compute the square root of $a \in (\mathbb{F}_q^*)^2$, Cipolla's algorithm uses an element b in $\mathbb{F}_q$ such that $b^2 - 4a$ is not a square in $\mathbb{F}_q$. Then, the polynomial $f(Y) = Y^2 - bY + a$ is irreducible over $\mathbb{F}_q$, hence $\mathbb{K} = \mathbb{F}_q[Y]/\langle f \rangle$ is a field. Let y be the residue class of Y modulo $\langle f \rangle$. Since f is the minimal polynomial of y over $\mathbb{F}_q$, $N_{\mathbb{K}/\mathbb{F}_q}(y) = a$, ensuring that $\sqrt{a} = Y^{(q+1)/2} \bmod (Y^2 - bY + a)$.

Finding a quadratic nonresidue of the form $b^2 - 4a$ by choosing a random $b \in \mathbb{F}_q$ requires an expected $O(1)$ attempts [1, page 158]. The quadratic residue test, and the norm computation take $O(\mathsf{M}(n)\log(n) + \log(p))$ and $O(n\mathsf{M}(n)\log(p))$ multiplications in $\mathbb{F}_p$ respectively. Therefore, the cost of the algorithm is an expected $O(n\mathsf{M}(n)\log(p))$ operations in $\mathbb{F}_p$.

Algorithms extending Cipolla's to the computation of t-th roots in $\mathbb{F}_p$, where t is a prime factor of $p-1$, are in [24, 25, 15].

The Tonelli-Shanks algorithm. We will describe the algorithm in the case of square roots, although the ideas extend to higher orders. Tonelli's algorithm [21] and Shanks'

improvement [18] use discrete logarithms to reduce the problem to a subgroup of $\mathbb{F}_q^*$ of odd order. Let $q - 1 = 2^r \ell$ with $(\ell, 2) = 1$ and let H be the unique subgroup of $\mathbb{F}_q^*$ of order ℓ. Assume we find a quadratic nonresidue $g \in \mathbb{F}_q^*$; then, the square root of $a \in \mathbb{F}_q^*$ can be computed as follows: we can express a as $g^s h \in g^s H$ by solving a discrete logarithm in $\mathbb{F}_q^*/H$; s is necessarily even, so that $\sqrt{a} = g^{s/2} h^{(\ell+1)/2}$.

According to [16], the discrete logarithm requires $O(r^2 \mathsf{M}(n))$ multiplications in $\mathbb{F}_p$; all other steps take $O(n \mathsf{M}(n) \log(p))$ operations in $\mathbb{F}_p$. Hence, the expected running time of the algorithm is $O((r^2 + n \log(p)) \mathsf{M}(n))$ operations in $\mathbb{F}_p$. Thus, the efficiency of this algorithm depends on the structure of $\mathbb{F}_q^*$: there exists an infinite sequence of primes for which the cost is $O(n^2 \mathsf{M}(n) \log(p)^2)$, see [22].

Improving the exponentiation. All algorithms seen before use at best $O(n\mathsf{M}(n)\log(p))$ operations in $\mathbb{F}_p$, because of the exponentiation. Using ideas going back to [11], Barreto *et al.* [3] observed that for some of the cases seen above, the exponentiation can be improved, giving a cost subquadratic in n.

For instance, when taking square roots with $q = 3 \pmod 4$, the exponentiation $a^{(q+1)/4}$ can be reduced to computing $a^{1+u+\cdots+u^{(n-3)/2}}$, with $u = p^2$, plus two (cheap) exponentiations with exponents $p(p-1)$ and $(p+1)/4$. The special form of the exponent $1+u+\cdots+u^{(n-3)/2}$ makes it possible to apply a binary powering approach, involving $O(\log(n))$ multiplications and exponentiations, with exponents that are powers of p.

Further examples for square roots are discussed in [13, 9], covering the entries of Table 1. These references assume a normal basis representation; using (as we do) the monomial basis and modular composition techniques (which will be explained in the next section), the costs become $O(\mathsf{M}(n)\log(p) + \mathsf{C}(n)\log(n))$. Some cases of higher index roots are in [2]: if t is a factor of $p - 1$, such that t^2 does not divide $p - 1$, and if $\gcd(n, t) = 1$, then t-th root extraction can be done using $O(t\mathsf{M}(n)\log(p) + \mathsf{C}(n)\log(n))$ operations in $\mathbb{F}_p$.

Kaltofen and Shoup's algorithm. Finally, we mention what is, as far as we know, the only algorithm achieving an expected subquadratic running time in n (using the monomial basis representation), without any assumption on p.

Consider a factor t of $q - 1$. To compute a t-th root of $a \in (\mathbb{F}_q^*)^t$, the idea is simply to factor $Y^t - a \in \mathbb{F}_q[Y]$ using polynomial factorization techniques. Since we know that a is a t-th power, this polynomial splits into linear factors, so we can use an Equal Degree Factorization (EDF) algorithm.

A specialized EDF algorithm, dedicated to the case of high-degree extension of a given base field, was proposed by Kaltofen and Shoup [12]. It mainly reduces to the computation of a trace-like quantity $b + b^p + \cdots + b^{p^{n-1}}$, where b is a random element in $\mathbb{F}_q[Y]/\langle Y^t - a \rangle$. Using a binary powering technique similar to the one of the previous paragraph, this results in an expected running time of $O((\mathsf{M}(t)\mathsf{M}(n)\log(p) + t\mathsf{C}(n) + \mathsf{C}(t)\mathsf{M}(n))\log(n))$ operations in $\mathbb{F}_p$; remark that this estimate is faster than what is stated in [12] by a factor $\log(t)$, since here we only need one root, instead of the whole factorization. In the particular case $t = 2$, this becomes $O((\mathsf{M}(n)\log(p) + \mathsf{C}(n))\log(n))$. This achieves a running time subquadratic in n.

This idea actually allows one to compute a t-th root, for arbitrary t: starting from the

polynomial $Y^t - a$, we apply the above algorithm to $\gcd(Y^t - a, Y^q - Y)$; computing Y^q modulo $Y^t - a$ can be done by the same binary powering techniques.

3   A new root extraction algorithm

In this section, we focus on t-th root extraction in $\mathbb{F}_q$, for t a prime dividing $q - 1$ (as we saw in Section 1, m-th root extraction, for an integer $m \geq 2$, reduces to taking t-th roots, where t is a prime factor of m dividing $q - 1$).

The algorithm we present uses the trace $\mathbb{F}_q \to \mathbb{F}_{q'}$, for some subfield $\mathbb{F}_{q'} \subset \mathbb{F}_q$ to reduce t-th root extraction in $\mathbb{F}_q$ to t-th root extraction in $\mathbb{F}_{q'}$. We assume as before that the field $\mathbb{F}_q$ is represented by a quotient $\mathbb{F}_p[X]/\langle f \rangle$, with $f(X) \in \mathbb{F}_p[X]$ a monic irreducible polynomial of degree n. We let x be the residue class of X modulo $\langle f \rangle$.

Since we will handle both $\mathbb{F}_q$ and $\mathbb{F}_{q'}$, conversions may be needed. We recall that the minimal polynomial $g \in \mathbb{F}_p[Z]$ of an element $b \in \mathbb{F}_q$ can be computed in $O(\mathsf{C}(n)+\mathsf{M}(n)\log(n))$ operations in $\mathbb{F}_p$ [19]. Then, $\mathbb{F}_{q'} = \mathbb{F}_p[Z]/\langle g \rangle$ is a subfield of $\mathbb{F}_q = \mathbb{F}_p[X]/\langle f \rangle$; given $r \in \mathbb{F}_{q'}$, written as a polynomial in Z, we obtain its representation on the monomial basis of $\mathbb{F}_q$ by means of a modular composition, in time $\mathsf{C}(n)$. We will write this operation $\mathrm{Embed}(r, \mathbb{F}_q)$. Note that when b is in $\mathbb{F}_p$, all these operations are actually free.

3.1   An auxiliary algorithm

We first discuss a binary powering algorithm to solve the following problem. Starting from $\lambda \in \mathbb{F}_q$, we are going to compute

$$\alpha_i(\lambda) = \lambda^{1+p^s} + \lambda^{1+p^s+p^{2s}} + \cdots + \lambda^{1+p^s+p^{2s}+\cdots+p^{is}}$$

for given integers $i, s > 0$. This question is similar to (but distinct from) some exponentiations and trace-like computations we discussed before; our solution will be a similar binary powering approach, which will perform $O(\log(i))$ multiplications and exponentiations by powers of p. Let

$$\xi_i = x^{p^{is}}, \quad \zeta_i(\lambda) = \lambda^{p^s+p^{2s}+\cdots+p^{is}} \quad \text{and} \quad \delta_i(\lambda) = \lambda^{p^s} + \lambda^{p^s+p^{2s}} + \cdots + \lambda^{p^s+p^{2s}+\cdots+p^{is}},$$

where all quantities are computed in $\mathbb{F}_q$, that is, modulo f; for simplicity, in this paragraph, we will write α_i, ζ_i and δ_i. Note that $\alpha_i = \lambda\delta_i$, and that we have the following relations:

$$\xi_1 = x^{p^s}, \quad \zeta_1 = \lambda^{p^s}, \quad \delta_1 = \lambda^{p^s}$$

and

$$\xi_i = \begin{cases} \xi_{i/2}^{p^{is/2}} & \text{if } i \text{ is even} \\ \xi_{i-1}^{p^s} & \text{if } i \text{ is odd,} \end{cases} \quad \zeta_i = \begin{cases} \zeta_{i/2}\zeta_{i/2}^{p^{is/2}} & \text{if } i \text{ is even} \\ \zeta_1\zeta_{i-1}^{p^s} & \text{if } i \text{ is odd,} \end{cases} \quad \delta_i = \begin{cases} \delta_{i/2} + \zeta_{i/2}\delta_{i/2}^{p^{is/2}} & \text{if } i \text{ is even} \\ \delta_{i-1} + \zeta_i & \text{if } i \text{ is odd.} \end{cases}$$

Because we are working in a monomial basis, computing exponentiations to powers of p is not a trivial task; we will perform them using the following modular composition technique from [7].

6

Take $j \geq 0$ and $r \in \mathbb{F}_q$, and let R and Ξ_j be the canonical preimages of respectively r and ξ_j in $\mathbb{F}_p[X]$; then, we have

$$r^{p^{js}} = R(\Xi_j) \bmod f,$$

see for instance [23, Chapter 14.7]. We will simply write this as $r(\xi_j)$, and note that it can be computed using one modular composition, in time $\mathsf{C}(n)$. These remarks give us the following recursive algorithm, where we assume that $\xi_1 = x^{p^s}$ and $\zeta_1 = \lambda^{p^s}$ are already known.

Algorithm 1 XiZetaDelta$(\lambda, i, \xi_1, \zeta_1)$

Input: λ, a positive integer i, $\xi_1 = x^{p^s}$, $\zeta_1 = \lambda^{p^s}$
Output: ξ_i, ζ_i, δ_i
  1. **if** $i = 1$ **then**
  2.   **return** ξ_1, ζ_1, ζ_1
  3. **end if**
  4. $j \leftarrow \lfloor i/2 \rfloor$
  5. $\xi_j, \zeta_j, \delta_j \leftarrow$ XiZetaDelta$(\lambda, j, \xi_1, \zeta_1)$
  6. $\xi_{2j} \leftarrow \xi_j(\xi_j)$
  7. $\zeta_{2j} \leftarrow \zeta_j \cdot \zeta_j(\xi_j)$
  8. $\delta_{2j} \leftarrow \delta_j + \zeta_j \delta_j(\xi_j)$
  9. **if** i is even **then**
 10.   **return** ξ_{2j}, ζ_{2j}, δ_{2j}
 11. **end if**
 12. $\xi_i \leftarrow \xi_{2j}(\xi_1)$
 13. $\zeta_i \leftarrow \zeta_1 \cdot \zeta_{2j}(\xi_1)$
 14. $\delta_i \leftarrow \delta_{2j} + \zeta_i$
 15. **return** ξ_i, ζ_i, δ_i

We deduce the following algorithm for computing $\alpha_i(\lambda)$.

Algorithm 2 Alpha(λ, i)

Input: λ, a positive integer i
Output: α_i
  1. $\xi_1 \leftarrow x^{p^s}$
  2. $\zeta_1 \leftarrow \lambda^{p^s}$
  3. $\xi_i, \zeta_i, \delta_i \leftarrow$ XiZetaDelta$(\lambda, i, \xi_1, \zeta_1)$
  4. **return** $\lambda \delta_i$

Proposition 2. *Algorithm 2 computes $\alpha_i(\lambda)$ using $O(\mathsf{C}(n)\log(is) + \mathsf{M}(n)\log(p))$ operations in $\mathbb{F}_p$.*

Proof. To compute x^{p^s} and λ^{p^s} we first compute x^p and λ^p using $O(\log(p))$ multiplications in $\mathbb{F}_q$, and then do $O(\log(s))$ modular compositions modulo f. The depth of the recursion in Algorithm 1 is $O(\log(i))$; each recursive call involves $O(1)$ additions, multiplications and modular compositions modulo f, for a total time of $O(\mathsf{C}(n))$ per recursive call. $\qquad\square$

As said before, the algorithm can also be implemented using a normal basis representation. Then, exponentiations to powers of p become trivial, but multiplication becomes more difficult. We leave these considerations to the reader.

3.2 Taking t-th roots

We will now give our root extraction algorithm. As said before, we now let t be a prime factor of $q - 1$, and we let s be the order of p in $\mathbb{Z}/t\mathbb{Z}$. Then s divides n, say $n = s\ell$.

We first explain how to test for t-th powers. Testing whether $a \in \mathbb{F}_q^*$ is a t-th power is equivalent to testing whether $a^{(q-1)/t} = 1$. Let $\zeta = a^{(p^s-1)/t}$; then $a^{(q-1)/t} = \zeta^{1+p^s+\cdots+p^{(\ell-1)s}}$. Computing ζ requires $O(s\mathsf{M}(n)\log(p))$, and computing $\zeta^{1+p^s+\cdots+p^{(\ell-1)s}}$ using Algorithm 1 requires $O(\mathsf{C}(n)\log(n) + \mathsf{M}(n)\log(p))$ operations in $\mathbb{F}_p$. Therefore, testing for a t-th power takes $O(\mathsf{C}(n)\log(n) + s\mathsf{M}(n)\log(p))$ operations in $\mathbb{F}_p$.

In the particular case when t divides $p - 1$, we can actually do better: we have $a^{(q-1)/t} = \mathrm{res}(f, a)^{(p-1)/t}$, where $\mathrm{res}(\cdot, \cdot)$ is the resultant function. The resultant can be computed using $O(\mathsf{M}(n)\log(n))$ operations in $\mathbb{F}_p$, so the whole test can be done using $O(\mathsf{M}(n)\log(n)+\log(p))$ operations in $\mathbb{F}_p$.

In any case, we can now assume that we are given $a \in (\mathbb{F}_q^*)^t$, with t-th root $\gamma \in \mathbb{F}_q$. Defining $\beta = T_{\mathbb{F}_q/\mathbb{F}_{q'}}(\gamma)$, where $T_{\mathbb{F}_q/\mathbb{F}_{q'}} : \mathbb{F}_q \to \mathbb{F}_{q'}$ is the trace linear form and $q' = p^s$, we have

$$
\begin{aligned}
\beta &= \sum_{i=0}^{\ell-1} \gamma^{p^{is}} \\
&= \gamma(1 + \gamma^{p^s-1} + \gamma^{p^{2s}-1} + \cdots + \gamma^{p^{(\ell-1)s}-1}) \\
&= \gamma(1 + a^{(p^s-1)/t} + a^{(p^{2s}-1)/t} + \cdots + a^{(p^{(\ell-1)s}-1)/t}).
\end{aligned} \tag{1}
$$

Let $b = 1 + a^{(p^s-1)/t} + a^{(p^{2s}-1)/t} + \cdots + a^{(p^{(\ell-1)s}-1)/t}$, so that Equation (1) gives $\beta = \gamma b$. Taking the t-th power in both sides results in the equation $\beta^t = ab^t$ over $\mathbb{F}_{q'}$. Since we know a, and we can compute b, we can thus determine β by t-th root extraction in $\mathbb{F}_{q'}$. Then, if we assume that $b \neq 0$ (or equivalently that $\beta \neq 0$), we deduce $\gamma = \beta b^{-1}$; to resolve the issue that β may be zero, we will replace a by $a' = ac^t$, for a random element $c \in \mathbb{F}_q^*$.

Computing the t-th root of $a'b^t$ in $\mathbb{F}_{q'}$ is done as follows. We first compute the minimal polynomial $g \in \mathbb{F}_p[Z]$ of $a'b^t$, and let z be the residue class of Z in $\mathbb{F}_p[Z]/\langle g \rangle$. Then, we compute a t-th root r of z in $\mathbb{F}_p[Z]/\langle g \rangle$, and embed r in $\mathbb{F}_q$. The computation of r is done by a black-box t-th root extraction algorithm, denoted by $r \mapsto r^{1/t}$.

It remains to explain how to compute b efficiently. Let $\lambda = a^{(p^s-1)/t}$; then, one verifies that $b = 1 + \lambda + \alpha_{\ell-2}(\lambda)$, so we can use the algorithm of the previous subsection. Putting all together, we obtain the following algorithm

---

Algorithm 3 t-th root in $\mathbb{F}_q^*$

Input: $a \in (\mathbb{F}_q^*)^t$
Output: a t-th root of a
 1. $s \leftarrow$ the order of p in $\mathbb{Z}/t\mathbb{Z}$
 2. $\ell \leftarrow n/s$
 3. **repeat**
 4.   choose a random $c \in \mathbb{F}_q$
 5. $a' \leftarrow ac^t$
 6. $\lambda \leftarrow a'^{(p^s-1)/t}$

7. $b \leftarrow 1 + \lambda + \mathrm{Alpha}(\lambda, \ell - 2)$
8. **until** $b \neq 0$
9. $g \leftarrow \mathrm{MinimalPolynomial}(a'b^t)$
10. $\beta \leftarrow z^{1/t}$ in $\mathbb{F}_q[Z]/\langle g \rangle$
11. **return** $\mathrm{Embed}(\beta, \mathbb{F}_q)b^{-1}c^{-1}$

The following proposition proves Theorem 1.

Proposition 3. *Algorithm 3 computes a t-th root of a using an expected $O(s\mathsf{M}(n)\log(p) + \mathsf{C}(n)\log(n))$ operations in $\mathbb{F}_p$, plus a t-th root extraction in $\mathbb{F}_{q'}$.*

Proof. Note first that $\beta = 0$ means that $T_{\mathbb{F}_q/\mathbb{F}_{q'}}(\gamma c) = 0$. There are q/q' values of c for which this is the case, so we expect to have to choose $O(1)$ elements in $\mathbb{F}_q$ before exiting the repeat … until loop. Each pass in the loop uses $O(s\mathsf{M}(n)\log(p))$ operations in $\mathbb{F}_p$ to compute a' and λ, and $O(\mathsf{C}(n)\log(n) + \mathsf{M}(n)\log(p))$ operations in $\mathbb{F}_p$ to compute b.

Given a' and b, one obtains b^t and $a'b^t$ using another $O(s\mathsf{M}(n)\log(p))$ operations in $\mathbb{F}_p$; then, computing g takes time $O(\mathsf{C}(n) + \mathsf{M}(n)\log(n))$.

After the black-box call to t-th root extraction modulo g, embedding β in $\mathbb{F}_q$ takes time $\mathsf{C}(n)$. We can then deduce γ by two divisions in $\mathbb{F}_q$, using $O(\mathsf{M}(n)\log(n))$ operations in $\mathbb{F}_p$; this is negligible compared to the cost of all modular compositions. $\square$

3.3   Experimental results

We have implemented our root extraction algorithm, in the case $m = 2$ (that is, we are taking square roots); our implementation is based on Shoup's NTL [20]. Figure 1 compares our algorithm to Cipolla's and Tonelli-Shanks' algorithms over $\mathbb{F}_q$, with $q = p^n$, for the randomly selected prime $p = 348975609381470925634534573457497$, and different values of the extension degree n.
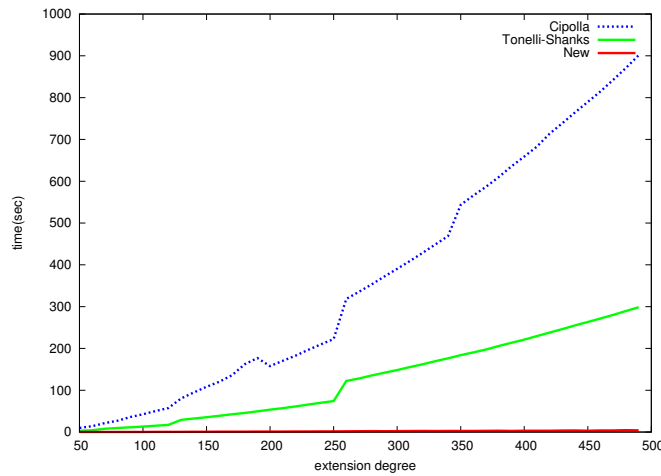

Figure 1: Our square root algorithm vs. Cipolla's and Tonelli-Shanks' algorithms.

9

Remember that the bottleneck in Cipolla's and Tonelli-Shanks' algorithms is the exponentiation, which takes $O(n\mathsf{M}(n)\log(p))$ operation in $\mathbb{F}_p$. As it turns out, NTL's implementation of modular composition has $\omega = 2$; this means that with this implementation we have $\mathsf{C}(n) = O(n^2)$, and our algorithm takes expected time $O(\mathsf{M}(n)\log(p) + n^2\log(n))$. Although this implementation is not subquadratic in n, it remains faster than Cipolla's and Tonelli-Shanks' algorithms, in theory and in practice.

Next, Figure 2 compares our NTL implementation of the EDF algorithm proposed by Kaltofen and Shoup, and our square root algorithm (note that the range of reachable degrees is much larger that in the first figure). We have ran the algorithms for several random elements for each extension degree. The vertical dashed lines, and the green line show the running time range, and the average running time of Kaltofen and Shoup algorithm respectively. In the case of our algorithm (the red graph), the vertical ranges are invisible because the deviation from the average is $\approx 10^{-2}$ seconds.
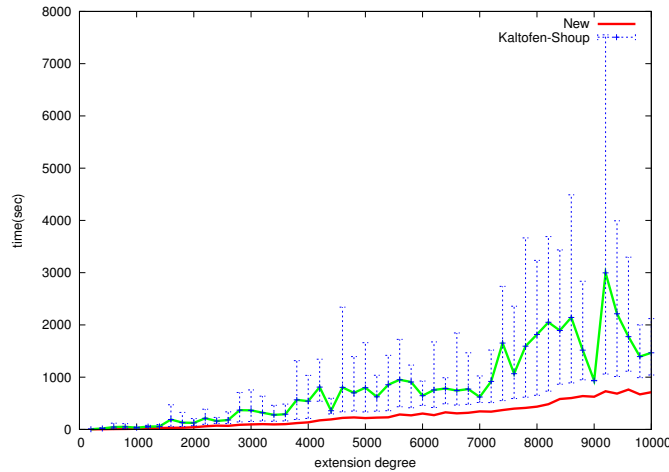

Figure 2: Our algorithm vs. Kaltofen and Shoup's algorithm.

This time, the results are closer. Nevertheless, it appears that the running time of our algorithm behaves more "smoothly", in the sense that random choices seem to have less influence. This is indeed the case. The random choice in Kaltofen and Shoup's algorithm succeeds with probability about $1/2$; in case of failure, we have to run to whole algorithm again. In our case, our choice of an element c in $\mathbb{F}_q^*$ fails with probability $1/p \ll 1/2$; then, there is still randomness involved in the t-th root extraction in $\mathbb{F}_p$, but this step was negligible in the range of parameters where our experiments were conducted.

Another way to express this is to compare the standard deviations in the running times of both algorithms. In the case of Kaltofen-Shoup's algorithm, the standard deviation is about $1/\sqrt{2}$ of the average running time of the whole algorithm. For our algorithm, the standard deviation is no more than $1/\sqrt{p}$ of the average running time of the trace-like computation (which is the dominant part), plus $1/\sqrt{2}$ of the average running time of the root extraction in $\mathbb{F}_p$ (which is cheap).

References

[1] E. Bach and J. Shallit. *Algorithmic Number Theory; Volume I: Efficient Algorithms*. The MIT Press, 1996.

[2] P. Barreto and J. Voloch. Efficient computation of roots in finite fields. *Designs, Codes and Cryptography*, 39:275–280, 2006.

[3] P. S. L. M. Barreto, H. Y. Kim, B. Lynn, and M. Scott. Efficient algorithms for pairing-based cryptosystems. In *Advances in cryptology—CRYPTO 2002*, volume 2442 of *Lecture Notes in Comput. Sci.*, pages 354–368. Springer, 2002.

[4] R. P. Brent and H. T. Kung. Fast algorithms for manipulating formal power series. *Journal of the Association for Computing Machinery*, 25(4):581–595, 1978.

[5] D. G. Cantor and E. Kaltofen. On fast multiplication of polynomials over arbitrary algebras. *Acta Informatica*, 28(7):693–701, 1991.

[6] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *J. Symb. Comp*, 9(3):251–280, 1990.

[7] J. von zur Gathen and V. Shoup. Computing Frobenius maps and factoring polynomials. *Comput. Complexity*, 2(3):187–224, 1992.

[8] P. Gaudry and É. Schost. Point-counting in genus 2 over prime fields, 2011.

[9] D.-G. Han, D. Choi, and H. Kim. Improved computation of square roots in specific finite fields. *IEEE Trans. Comput.*, 58:188–196, 2009.

[10] X. Huang and V. Y. Pan. Fast rectangular matrix multiplication and applications. *J. Complexity*, 14(2):257–299, 1998.

[11] T. Itoh and S. Tsujii. A fast algorithm for computing multiplicative inverses in $GF(2^m)$ using normal bases. *Inform. and Comput.*, 78(3):171–177, 1988.

[12] E. Kaltofen and V. Shoup. Fast polynomial factorization over high algebraic extensions of finite fields. In *ISSAC'97*, pages 184–188. ACM, 1997.

[13] F. Kong, Z. Cai, J. Yu, and D. Li. Improved generalized Atkin algorithm for computing square roots in finite fields. *Inform. Process. Lett.*, 98(1):1–5, 2006.

[14] S. Müller. On the computation of square roots in finite fields. *Des. Codes Cryptogr.*, 31(3):301–312, 2004.

[15] N. Nishihara, R. Harasawa, Y. Sueyoshi, and A. Kudo. A remark on the computation of cube roots in finite fields. Cryptology ePrint Archive, Report 2009/457, 2009. `http://eprint.iacr.org/`.

[16] S. C. Pohlig and M. E. Hellman. An improved algorithm for computing logarithms over gf(p) and its cryptographic significance. *IEEE Transactions on Information Theory*, IT-24:106–110, 1978.

[17] A. Schönhage and V. Strassen. Schnelle Multiplikation grosser Zahlen. *Computing*, 7:281–292, 1971.

[18] D. Shanks. Five number-theoretic algorithms. In *Proceedings of the Second Manitoba Conference on Numerical Mathematics*, pages 51–70, 1972.

[19] V. Shoup. Fast Construction of Irreducible Polynomials over Finite Fields". *Journal of Symbolic Computation*, 17(5):371–391, May 1994.

[20] V. Shoup. A library for doing number theory (NTL). `http://www.shoup.net/ntl/`, 2009.

[21] A. Tonelli. Bemerkung über die Auflösung quadratischer Congruenzen. *Göttinger Nachrichten*, pages 344–346, 1891.

[22] G. Tornaría. Square roots modulo p. In *LATIN 2002: Theoretical Informatics*, volume 2286 of *Lecture Notes in Comput. Sci.*, pages 430–434. Springer, 2002.

[23] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, 2003.

[24] H. C. Williams. Some algorithms for solving $x^q \equiv N \pmod{p}$. In *Proceedings of the Third Southeastern Conference on Combinatorics, Graph Theory and Computing (Florida Atlantic Univ., Boca Raton, Fla., 1972)*, pages 451–462. Florida Atlantic Univ., 1972.

[25] K. S. Williams and K. Hardy. A refinement of H. C. Williams' qth root algorithm. *Math. Comp.*, 61(203):475–483, 1993.