## 4 Arithmetic in finite fields II

We saw in Lecture 3 how to efficiently multiply integers, and, using Kronecker substitution, how to efficiently multiply polynomials with integer coefficients. This gives us what we need to multiply elements in finite fields, provided that we have a way to reduce the result to our standard representations of $\mathbb{F}_{p} \simeq \mathbb{Z} / p \mathbb{Z}$ and $\mathbb{F}_{q} \simeq \mathbb{F}_{p}[x] /(f)$, using integers in $[0, p-1]$ and polynomials of degree less than $\operatorname{deg} f$, respectively. In both cases we use Euclidean division.

### 4.1 Euclidean division

Given integers $a, b>0$, we wish to compute the unique integers $q, r \geq 0$ for which

$$
a=b q+r \quad(0 \leq r<b) .
$$

We have $q=\lfloor a / b\rfloor$ and $r=a \bmod b$. It is enough to compute $q$, since we can then compute $r=a-b q$. To compute $q$, we determine a sufficiently precise approximation $c \approx 1 / b$ and obtain $q$ by computing $c a$ and rounding down to the nearest integer.

We recall Newton's method for finding the root of a real-valued function $f(x)$. We start with an initial approximation $x_{0}$, and at each step, we refine the approximation $x_{i}$ by computing the $x$-coordinate $x_{i+1}$ of the point where the tangent line through $\left(x_{i}, f\left(x_{i}\right)\right)$ intersects the $x$-axis, via

$$
x_{i+1}:=x_{i}-\frac{f\left(x_{i}\right)}{f^{\prime}\left(x_{i}\right)} .
$$

To compute $c \approx 1 / b$, we apply this to $f(x)=1 / x-b$, using the Netwon iteration

$$
x_{i+1}=x_{i}-\frac{f\left(x_{i}\right)}{f^{\prime}\left(x_{i}\right)}=x_{i}-\frac{\frac{1}{x_{i}}-b}{-\frac{1}{x_{i}^{2}}}=2 x_{i}-b x_{i}^{2} .
$$

As an example, let us approximate $1 / b=1 / 123456789$. For the sake of illustration we work in base 10 , but in an actual implementation would use base 2 , or base $2^{w}$, where $w$ is the word size.

$$
\begin{aligned}
x_{0} & =1 \times 10^{-8} \\
x_{1} & =2\left(1 \times 10^{-8}\right)-\left(1.2 \times 10^{8}\right)\left(1 \times 10^{-8}\right)^{2} \\
& =0.80 \times 10^{-8} \\
x_{2} & =2\left(0.80 \times 10^{-8}\right)-\left(1.234 \times 10^{8}\right)\left(0.80 \times 10^{-8}\right)^{2} \\
& =0.8102 \times 10^{-8} \\
x_{3} & =2\left(0.8102 \times 10^{-8}\right)-\left(1.2345678 \times 10^{8}\right)\left(0.8102 \times 10^{-8}\right)^{2} \\
& =0.81000002 \times 10^{-8} .
\end{aligned}
$$

Note that we double the precision we are using at each step, and each $x_{i}$ is correct up to an error in its last decimal place. The value $x_{3}$ suffices to correctly compute $\lfloor a / b\rfloor$ for $a \leq 10^{15}$.

To analyze the complexity of this approach, let us assume that $b$ has $n$ bits and $a$ has at most $2 n$ bits; this is precisely the situation we will encounter when we wish to reduce
the product of two integers in $[0, p-1]$ modulo $p$. During the Newton iteration to compute $c \approx 1 / b$, the size of the integers involved doubles with each step, and the cost of the arithmetic operations grows at least linearly. The total cost is thus at most twice the cost of the last step, which is $\mathrm{M}(n)+O(n)$; note that all operations can be performed using integers by shifting the operands appropriately. Thus we can compute $c \approx 1 / b$ in time $2 \mathrm{M}(n)+O(n)$. We can then compute $c a \approx a / b$, round to the nearest integer, and compute $r=a-b q$ using at most $4 \mathrm{M}(n)+O(n)$ bit operations.

With a slightly more sophisticated version of this approach it is possible to compute $r$ in time $3 \mathrm{M}(n)+O(n)$. and if we expect to repeatedly perform Euclidean division with the same denominator, as when working in $\mathbb{F}_{p}$, the cost of all subsequent reductions can be reduced to $\mathrm{M}(n)+O(n)$. This uses what is known as Barret Reduction; see [2, Algorithm 10.17]. $\frac{1}{2}$ In any case, we obtain the following bound for multiplication in $\mathbb{F}_{p}$ using our standard representation as integers in $[0, p-1]$.

Theorem 4.1. The time to multiply two elements of $\mathbb{F}_{p}$ is $O(\mathrm{M}(n))$, where $n=\lg p$.
There is an analogous version of this algorithm above for polynomials that uses the exact same Newton iteration $x_{i+1}=2 x_{i}-b x_{i}^{2}$, where $b$ and the $x_{i}$ are now polynomials. Rather than working with Laurent polynomials (the polynomial version of approximating a rational number with a truncated decimal expansion), it is simpler to reverse the polynomials and work modulo a sufficiently large power of $x$, doubling the power of $x$ with each Newton iteration. More precisely, we have the following algorithm, which combines Algorithms 9.3 and 9.5 from [3]. For any polynomial $f(x)$ we write $\operatorname{rev} f$ for the polynomial $x^{\operatorname{deg} f} f\left(\frac{1}{x}\right)$; this simply reverses the coefficients of $f$.

Algorithm 4.2 (Fast Euclidean division of polynomials). Given $a, b \in \mathbb{F}_{p}[x]$ with $b$ monic, compute $q, r \in \mathbb{F}_{p}[x]$ such that $a=q b+r$ with $\operatorname{deg} r<\operatorname{deg} b$ as follows:

1. If $\operatorname{deg} a<\operatorname{deg} b$ then return $q=0$ and $r=a$.
2. Let $m=\operatorname{deg} a-\operatorname{deg} b$ and $k=\lceil\lg m+1\rceil$.
3. Let $f=\operatorname{rev}(b)$ (reverse the coefficients of $b$ ).
4. Compute $g_{0}=1, g_{i}=\left(2 g_{i-1}-f g_{i-1}^{2}\right) \bmod x^{2^{i}}$ for $i$ from 1 to $k$. (this yields $f g_{k} \equiv 1 \bmod x^{m+1}$ ).
5. Set $s=\operatorname{rev}(a) g_{k} \bmod x^{m+1}\left(\right.$ now $\left.\operatorname{rev}(b) s \equiv \operatorname{rev}(a) \bmod x^{m+1}\right)$.
6. Return $q=x^{m-\operatorname{deg} s} \operatorname{rev}(s)$ and $r=a-b q$.

As in the integer case, the work is dominated by the last iteration in step 4, which involves multiplying polynomials in $\mathbb{F}_{p}[x]$. To multiply elements of $\mathbb{F}_{q} \simeq \mathbb{F}_{p}[x] /(f)$ represented as polynomials of degree less than $d=\operatorname{deg} f$, we compute the product $a$ in $\mathbb{F}[x]$ and then reduce modulo $b=f$, and the degree of the polynomials involved are all $O(d)$. With Kronecker substitution, we can reduce these polynomial multiplications to integer multiplications, and obtain the following result.

Theorem 4.3. Let $q=p^{d}$ be a prime power, and assume that either $\log d=O(\log p)$ or $p=O(1)$. The time to multiply two elements of $\mathbb{F}_{q}$ is $O(\mathrm{M}(n))$, where $n=\lg q$.

[^0]Remark 4.4. The constraints on the relative growth rate on $p$ and $d$ in the theorem above are present only so that we can easily express our bounds in terms of the bound $\mathrm{M}(n)$ for multiplying integers. In fact, for all the bounds currently known for $\mathrm{M}(n)$, Theorem 4.3 holds uniformly, without any assumptions about the relative growth rate of $p$ and $d$. More precisely, it follows from the recent result of [5] that for any prime power $q$ we can multiply two elements in $\mathbb{F}_{q}$ in time $O\left(n \log n 8^{\log ^{*} n}\right)$, no matter how $q=p^{d}$ tends to infinity; but the proof of this requires more than just Kronecker substitution.

### 4.2 Extended Euclidean algorithm

We recall the Euclidean algorithm for computing the greatest common divisor of positive integers $a$ and $b$. For $a>b$ we repeatedly apply

$$
\operatorname{gcd}(a, b)=\operatorname{gcd}(b, a \bmod b),
$$

where we take $a \bmod b$ to be the unique integer $r \in[0, b-1]$ congruent to $a$ modulo $b$.
To compute the multiplicative inverse of an integer modulo a prime, we use the extended Euclidean algorithm, which expresses $\operatorname{gcd}(a, b)$ as a linear combination

$$
\operatorname{gcd}(a, b)=a s+b t
$$

with $|s| \leq b / \operatorname{gcd}(a, b)$ and $|t| \leq a / \operatorname{gcd}(a, b)$. If $a$ is prime, we obtain $a s+b t=1$, and $t$ is the inverse of $b$ modulo $a$. To compute the integers $s$ and $t$ we use the following algorithm. First, let

$$
R_{1}=\left[\begin{array}{l}
a \\
b
\end{array}\right], \quad S_{1}=\left[\begin{array}{l}
1 \\
0
\end{array}\right], \quad T_{1}=\left[\begin{array}{l}
0 \\
1
\end{array}\right]
$$

and note that $R_{1}=a S_{1}+b T_{1}$. We then compute

$$
Q_{i}=\left[\begin{array}{cc}
0 & 1 \\
1 & -q_{i}
\end{array}\right], \quad R_{i+1}=Q_{i} R_{i}, \quad S_{i+1}=Q_{i} S_{i}, \quad T_{i+i}=Q_{i} T_{i},
$$

where $q_{i}$ is the quotient $\left\lfloor R_{i, 1} / R_{i, 2}\right\rfloor$ obtained via Euclidean division. Note that applying the linear transformation $Q_{i}$ to both sides of $R_{i}=a S_{i}+b T_{i}$ ensures $R_{i+1}=a S_{i+1}+b T_{i+1}$. The algorithm terminates when the $k$ th step where $R_{k, 2}$ becomes zero, at which point we have

$$
R_{k}=\left[\begin{array}{l}
d \\
0
\end{array}\right], \quad S_{k}=\left[\begin{array}{c}
s \\
\pm b
\end{array}\right], \quad T_{k}=\left[\begin{array}{c}
t \\
\mp a
\end{array}\right]
$$

with $\operatorname{gcd}(a, b)=d=s a+t b$. As an example, with $a=1009$ and $b=789$ we have

| $r$ | $q$ | $s$ | $t$ |
| ---: | ---: | ---: | ---: |
| 1009 |  | 1 | 0 |
| 789 | 1 | 0 | 1 |
| 220 | 3 | 1 | -1 |
| 129 | 1 | -3 | 4 |
| 91 | 1 | 4 | -5 |
| 38 | 2 | -7 | 9 |
| 15 | 2 | 18 | -23 |
| 8 | 1 | -43 | 55 |
| 7 | 1 | 61 | -78 |
| 1 | 7 | -104 | 133 |
| 0 |  | 789 | -1009 |

From the second-to-last line with $s=-104$ and $t=133$ we see that

$$
1=-104 \cdot 1009+133 \cdot 789,
$$

and therefore 133 is the inverse of 789 modulo 1009 (and $-104 \equiv 685$ is the inverse of 1009 modulo 789).

It is clear that the $r$ is reduced by a factor of at least 2 every two steps, thus the total number of iterations is $O(n)$, and each step involves Euclidean division, whose cost is bounded by $O(\mathrm{M}(n))$. This yields a complexity of $O(n \mathrm{M}(n))$, but a more careful analysis shows that it is actually $O\left(n^{2}\right)$, even if schoolbook multiplication is used (the key point is that the total size of all the $q_{i}$ is $O(n)$ bits).

This can be further improved using the fast Euclidean algorithm, which uses a divide-and-conquer approach to compute the product $Q=Q_{k-1} \cdots Q_{1}$ by splitting the product in half and recursively computing each half using what is known as a half-gcd algorithm. One can then compute $R_{k}=Q R_{1}, S_{k}=Q S_{1}$, and $T_{k}=Q T_{1}$. The details are somewhat involved (care must be taken when determining how to split the product in a way that balances the work evenly and $k$ is not known ahead of time), but this yields a recursive running time of

$$
T(n)=2 T(n / 2)+O(\mathrm{M}(n))=O(\mathrm{M}(n) \log n) ;
$$

see [3, Chapter 11] for further information.
Theorem 4.5. Let $p$ be a prime. The time to invert an element of $\mathbb{F}_{p}^{\times}$is $O(\mathrm{M}(n) \log n)$, where $n=\lg p$.

The extended Euclidean algorithm works in any Euclidean ring, that is, a ring with a norm function that allows us to use Euclidean division to write $a=q b+r$ with $r$ of norm strictly less than $b$ (for any nonzero $b$ ). This includes polynomial rings, in which the norm of a polynomial is simply its degree. Thus we can compute the inverse of a polynomial modulo another polynomial, provided the two polynomials are relatively prime. One issue that arises when working in Euclidean rings other than $\mathbb{Z}$ is that there may be units (invertible elements) other than $\pm 1$, and the gcd is only defined up to a unit. In the case of the polynomial ring $\mathbb{F}_{p}[x]$, every element of $\mathbb{F}_{p}^{\times}$is a unit, and with the fast Euclidean algorithm in $\mathbb{F}_{p}[x]$ one typically normalizes the intermediate results by making the polynomials monic at each step; this involves computing the inverse of the leading coefficient in $\mathbb{F}_{p}$. If $\mathbb{F}_{q}=\mathbb{F}_{p}[x] /(f)$ with $\operatorname{deg} f=d$, one can then bound the time to compute an inverse in $\mathbb{F}_{q}$ by $O(\mathrm{M}(d) \log d)$, operations in $\mathbb{F}_{p}$, of which $O(d)$ are inversions; see [3, Thm. 11.10(i)]. This gives a bit complexity of

$$
O(\mathrm{M}(d) \mathrm{M}(\log p) \log d+d \mathrm{M}(\log p) \log \log p)
$$

but with Kronecker substitution we can sharpen this to

$$
O(\mathrm{M}(d(\log p+\log d)) \log d+d \mathrm{M}(\log p) \log \log p)
$$

In essentially all the cases that arise in practice, either $\log d=O(\log p)$ (this certainly applies in large characteristic) and $\log p=O(1)$ (small characteristic), this bound simplifies to $O(\mathrm{M}(n) \log n)$, where $n=\lg q=d \lg p$ is the number of bits in $q$, the same result we obtained for the case where $q=p$ is prime.

Theorem 4.6. Let $q=p^{d}$ be a prime power and assume that either $\log d=O(\log p)$ or $p=O(1)$. The time to invert an element of $\mathbb{F}_{q}^{\times}$is $O(\mathrm{M}(n) \log n)$, where $n=\lg q$.

### 4.3 Exponentiation (scalar multiplication)

Let $a$ be a positive integer. In a multiplicative group, the computation

$$
g^{a}=\underbrace{g g \cdots g}_{a}
$$

is known as exponentiation. In an additive group, this is equivalent to

$$
a g=\underbrace{g+g+\cdots+g}_{a},
$$

and is called scalar multiplication. The same algorithms are used in both cases, and most of these algorithms were first developed in a multiplicative setting (the multiplicative group of a finite field) and are called exponentiation algorithms. It is actually more convenient to describe the algorithms using additive notation (fewer superscripts), so we will do so.

The oldest and most commonly used exponentiation algorithm is the "double-and-add" method, also known as left-to-right binary exponentiation. Given an element $P$ of an additive group and a positive integer $a$ with binary representation $a=\sum 2^{i} a_{i}$, we compute the scalar multiple $Q=a P$ as follows:

```
def double_and_add (P, a):
    a=a.digits(2); n=len(a) # represent a in binary using n bits
    Q=P; i=n-2 # start l bit below the high bit
    while i >= 0:
        Q=Q+Q # double
        if a[i]==1: Q = Q + P # add
        i -= 1
    return Q
```

Alternatively, we may use the "add-and-double" method, also known as right-to-left binary exponentiation.

```
def add_and_double (P, a):
    a=a.digits(2); n=len(a) # represent a in binary using n bits
    Q=0; R=P; i=0 # start with the low bit
    while i < n:
        if a[i]==1: Q = Q + R # add
        R=R+R # double
        i += 1
    return Q
```

The number of group operations required is effectively the same for both algorithms. If we ignore the first addition in the add_and_double algorithm (which could be replaced by an assignment, since initially $Q=0$ ), both algorithms use precisely

$$
n+\mathrm{wt}(a)-2 \leq 2 n=O(n)
$$

group operations, where $\operatorname{wt}(a)=\#\left\{a_{i}: a_{i}=1\right\}$ is the Hamming weight of $a$, the number of 1 's in its binary representation. Up to the constant factor 2 , this is asymptotically optimal, and it implies that exponentiation in a finite field $\mathbb{F}_{q}$ has complexity $O(n \mathrm{M}(n))$ with $n=\lg q$; this assumes the exponent is less than $q$, but note that we can always reduce the exponent modulo $q-1$, the order of the cyclic group $\mathbb{F}_{q}^{\times}$. Provided the bit-size of the exponent
is $O\left(n^{2}\right)$, the $O\left(\mathrm{M}\left(n^{2}\right)\right)$ time to reduce the exponent modulo $q-1$ will be majorized by the $O(n \mathrm{M}(n))$ time to perform the exponentiation.

Notwithstanding the fact that the simple double-and-add algorithm is within a factor of 2 of the best possible, researchers have gone to great lengths to eliminate this factor of 2 , and to take advantage of situations where either the base or the exponent is fixed, and there are a wide variety of optimizations that are used in practice; see [2, Ch. 9] and [4] Here we give just one example, windowed exponentiation, which is able to reduce the constant factor from 2 to an essentially optimal $1+o(1)$.

### 4.3.1 Fixed-window exponentiation

Let the positive integer $s$ be a window size and write $a$ as

$$
a=\sum a_{i} 2^{s i}, \quad\left(0 \leq a_{i}<2^{s}\right) .
$$

This is equivalent to writing $a$ in base $2^{s}$. With fixed-window exponentiation, one first precomputes multiples $d P$ for each of the "digits" $d \in\left[0,2^{s}-1\right]$ that may appear in the base- $2^{s}$ expansion of $a$. One then uses a left-to-right approach as in the double-and-add algorithm, except now we double $s$ times and add the appropriate multiple $a_{i} P$.

```
def fixed_window (P, a, s):
    a=a.digits(2^s) # write a in base 2^s
    R = [0*P, P]
    for i in range(2,2^s): R.append(R[-1]+P) # precompute digits
    Q = R[a[-1]] # copy the top digit
    i = len(a)-2
    while i >= 0:
        for j in range(0,s): Q = Q + Q # double s times
        Q = Q + R[a[i]] # add the next digit
        i -= 1
    return Q
```

In the algorithm above we precompute multiples of $P$ for every possible digit that might occur. As an optimization one could examine the base- $2^{s}$ representation of $a$ and only precompute the multiples of $P$ that are actually needed.

Let $n$ be the number of bits in $a$ and let $m=\lceil n / s\rceil$ be the number of base- $2^{s}$ digits $a_{i}$. The precomputation step uses $2^{s}-2$ additions (we get $0 P$ and $1 P$ for free), there are $m-1$ additions of multiples of $P$ corresponding to digits $a_{i}$ (when $a_{i}=0$ these cost nothing), and there are a total of $(m-1) s$ doublings. This yields an upper bound of

$$
2^{s}-2+m-1+(m-1) s \approx 2^{s}+n / s+n
$$

group operations. If we choose $s=\lg n-\lg \lg n$, we obtain the bound

$$
n / \lg n+n /(\lg n-\lg \lg n)+n=n+O(n / \log n)
$$

which is $(1+o(1)) n$ group operations.

### 4.3.2 Sliding-window exponentiation

The sliding-window algorithm modifies the fixed-window algorithm by "sliding" over blocks of 0 s in the binary representation of $a$. There is still a window size $s$, but $a$ is no longer
treated as an integer written in a fixed base $2^{s}$. Instead, the algorithm scans the bits of the exponent from left to right, assembling "digits" of at most $s$ bits with both high and low bits set: with a sliding window of size 3 the bit-string 110100010010100 could be broken up as $11|0| 1|000| 1|00| 101 \mid 00$ with 4 nonzero digits, where as a fixed window approach would use $110|100| 010|010| 100$ with 5 nonzero digits.

This improves the fixed-window approach in two ways: first, it is only necessarily to precompute odd digits, and second, depending on the pattern of bits in $a$, sliding over the zeros may reduce the number of digits used, as in the example above. In any case, the sliding-window approach is never worse than the fixed-window approach, and for $s>2$ it is always better.

### 4.4 Root-finding in finite fields

Let $f(x)$ be a polynomial in $\mathbb{F}_{q}[x]$ of degree $d$. We wish to find a solution to $f(x)=0$ that lies in $\mathbb{F}_{q}$. As an important special case, this will allow us to compute square roots using $f(x)=x^{2}-a$, and, more generally, $r$ th roots. ${ }^{2}$

The algorithm we give here was originally proposed by Berlekamp for prime fields [1], and then refined and extended by Rabin [6], whose presentation we follow here. The algorithm is probabilistic, and is one of the best examples of how randomness can be exploited in a number-theoretic setting. As we will see, it is quite efficient, with an expected running time that is quasi-quadratic in the size of the input. By contrast, no deterministic polynomialtime algorithm for root-finding is known, not even for computing square roots. $\underline{3}^{-}$

### 4.4.1 Randomized algorithms

Probabilistic algorithms are typically classified as one of two types: Monte Carlo or Las Vegas. Monte Carlo algorithms are randomized algorithms whose output may be incorrect, depending on random choices that are made, but whose running time is bounded by a function of its input size, independent of any random choices. The probability of error is required to be less than $1 / 2-\epsilon$, for some $\epsilon>0$, and can be made arbitrarily small be running the algorithm repeatedly and using the output that occurs most often. In contrast, a Las Vegas algorithm always produces a correct output, but its running time may depend on random choices; we do require that its expected running time is finite. As a trivial example, consider an algorithm to compute $a+b$ that first flips a coin repeatedly until it gets a head and then computes $a+b$ and outputs the result. The running time of this algorithm may be arbitrarily long, even when computing $1+1=2$, but its expected running time is $O(n)$, where $n$ is the size of the inputs.

Las Vegas algorithms are generally preferred, particularly in mathematical applications. Note that any Monte Carlo algorithm whose output can be verified can always be converted to a Las Vegas algorithm (just run the algorithm repeatedly until you get an answer that is verifiably correct). The root-finding algorithm we present here is a Las Vegas algorithm.

[^1]
### 4.4.2 GCD splitting

Recall from the previous lecture that we defined the finite field $\mathbb{F}_{q}$ to be the splitting field of $x^{q}-x$ over its prime field $\mathbb{F}_{p}$; this definition also applies when $q=p$ is prime (since $x^{p}-x$ splits completely in $\mathbb{F}_{p}$ ), and in every case, the elements of $\mathbb{F}_{q}$ are precisely the roots of $x^{q}-x$. This means that if $k$ is any extension of $\mathbb{F}_{q}$ (for example, an algebraic closure), then each $\alpha \in k$ lies in $\mathbb{F}_{q}$ if and only if it is a root of $x^{q}-x$. Thus if $f(x)$ is an arbitrary polynomial in $\mathbb{F}_{q}[x]$, then the polynomial

$$
g(x):=\operatorname{gcd}\left(f(x), x^{q}-x\right)
$$

has the form $g(x)=\prod_{i}\left(x-\alpha_{i}\right)$, where the $\alpha_{i}$ range over all the distinct roots of $f$ that lie in $\mathbb{F}_{q}$. If $f$ has no roots in $\mathbb{F}_{q}$ then $g$ will have degree 0 (so $g=1$ if we assume our gcds are monic). Otherwise we can reduce the problem of finding a root of $f$ to the problem of finding a root of $g$, where $g$ has distinct roots that are all known to lie in $\mathbb{F}_{q}$.

In order to compute $g=\operatorname{gcd}\left(f, x^{q}-x\right)$ efficiently, we generally do not $\frac{4}{}$ compute $x^{q}-x$ and then take the gcd with $f(x)$; this would take time exponential in $n=\log q$. Instead, we compute $x^{q} \bmod f$ by exponentiating the polynomial $x$ to the $q$ th power in the ring $\mathbb{F}_{q}[x] /(f)$, whose elements are uniquely represented by polynomials of degree less than $d=\operatorname{deg} f$. Each multiplication in this ring involves the computation of a product in $\mathbb{F}_{q}[x]$ followed by a reduction modulo $f$; note that we do not assume $\mathbb{F}_{q}[x] /(f)$ is a field. This reduction is achieved using Euclidean division, and can be accomplished using two polynomial multiplications once an approximation to $1 / f$ has been precomputed, see $\S \underline{4.1}$, and is within a constant factor of the time to multiply two polynomials of degree $d$ in any case. The total cost of each multiplication in $\mathbb{F}_{q}[x] /(f)$ is thus $O(\mathrm{M}(d(n+\log d)))$, assuming that we use Kronecker substitution to multiply polynomials. The time to compute $x^{q} \bmod f$ using any of the exponentiation algorithms described in $\S 4.3$ is then $O(n \mathrm{M}(d(n+\log d)))$.

Once we have computed $x^{q} \bmod f$, we subtract $x$ and compute $g=\operatorname{gcd}\left(f, x^{q}-x\right)$. Using the fast Euclidean algorithm, this takes $O(\mathrm{M}(d(n+\log d)) \log d)$ time. Thus the total time to compute $g$ is $O(\mathrm{M}(d(n+\log d))(n+\log d))$, which we will see dominates the cost of root-finding. Typically $\log d=O(n)$ and this simplifies to $O(n \mathrm{M}(d n))$.

Having computed $g$, we seek to factor it into two polynomials of lower degree by again applying a gcd, with the goal of eventually obtaining a linear factor, which will yield a root.

Assuming that $q$ is odd (which we do), we may factor the polynomial $x^{q}-x$ as

$$
x^{q}-x=x\left(x^{s}-1\right)\left(x^{s}+1\right),
$$

where $s=(q-1) / 2$. Ignoring the root 0 (which we can easily check separately), this factorization splits $\mathbb{F}_{q}^{\times}$precisely in half: the roots of $x^{s}-1$ are the elements of $\mathbb{F}_{q}^{\times}$that are squares in $\mathbb{F}_{q}^{\times}$, and the roots of $x^{s}+1$ are the elements of $\mathbb{F}_{q}^{\times}$that are not. Recall that $\mathbb{F}_{q}^{\times}$is a cyclic group of order $q-1$, which is even (since $q$ is odd), thus the squares in $\mathbb{F}_{q}^{\times}$are the elements that are even powers of a generator, equivalently, elements whose order divides $(q-1) / 2$. If we compute

$$
h(x)=\operatorname{gcd}\left(g(x), x^{s}-1\right),
$$

we obtain a divisor of $g$ whose roots are precisely the roots of $g$ that are squares in $\mathbb{F}_{q}^{\times}$. If we suppose that the roots of $g$ are as likely to be squares as not, we should expect the degree

[^2]of $h$ to be approximately half the degree of $g$. And so long as the degree of $h$ is strictly between 0 and $\operatorname{deg} g$, one of $h$ or $g / h$ is a polynomial of degree at most half the degree of $g$, whose roots are all roots of our original polynomial $f$.

To make further progress, and to obtain an algorithm that is guaranteed to work no matter how the roots of $g$ are distributed in $\mathbb{F}_{q}$, we take a probabilistic approach. Rather than using the fixed polynomial $x^{s}-1$, we consider random polynomials of the form

$$
(x+\delta)^{s}-1
$$

where $\delta$ is uniformly distributed over $\mathbb{F}_{q}$. We claim that if $\alpha$ and $\beta$ are any two nonzero roots of $g$, then with probability $1 / 2$, exactly one of these is a root $(x+\delta)^{s}-1$. It follows from this claim that so long as $g$ has at least 2 distinct nonzero roots, the probability that the polynomial $h(x)=\operatorname{gcd}\left(g(x),(x+\delta)^{s}+1\right)$ is a proper divisor of $g$ is at least $1 / 2$.

Let us say that two elements $\alpha, \beta \in \mathbb{F}_{q}$ are of different type if they are both nonzero and $\alpha^{s} \neq \beta^{s}$. Our claim is an immediate consequence of the following theorem from [6].

Theorem 4.7 (Rabin). For every pair of distinct $\alpha, \beta \in \mathbb{F}_{q}$ we have

$$
\#\left\{\delta \in \mathbb{F}_{q}: \alpha+\delta \text { and } \beta+\delta \text { are of different type }\right\}=\frac{q-1}{2}
$$

Proof. Consider the map $\phi(\delta)=\frac{\alpha+\delta}{\beta+\delta}$, defined for $\delta \neq-\beta$. We claim that $\phi$ is a bijection form the set $\mathbb{F}_{q} \backslash\{-\beta\}$ to the set $\mathbb{F}_{q} \backslash\{1\}$. The sets are the same size, so we just need to show surjectivity. Let $\gamma \in \mathbb{F}_{q}-\{1\}$, then we wish to find a solution $x \neq-\beta$ to $\gamma=\frac{\alpha+x}{\beta+x}$. We have $\gamma(\beta+x)=\alpha+x$ which means $x-\gamma x=\gamma \beta-\alpha$. This yields $x=\frac{\gamma \beta-\alpha}{1-\gamma}$, which is not equal to $-\beta$, since $\alpha \neq \beta$. Thus $\phi$ is surjective.

We now note that

$$
\phi(\delta)^{s}=\frac{(\alpha+\delta)^{s}}{(\beta+\delta)^{s}}
$$

is -1 if and only if $\alpha+\delta$ and $\beta+\delta$ are of different type. The elements $\gamma=\phi(\delta)$ for which $\gamma^{s}=-1$ are precisely the non-residues in $\mathbb{F}_{q} \backslash\{1\}$, of which there are exactly $(q-1) / 2$.

We now give the algorithm.

Algorithm 4.8 (Root-finding). Given a polynomial $f \in \mathbb{F}_{q}[x]$, output an element $r \in \mathbb{F}_{q}$ such that $f(r)=0$, or null if no such $r$ exists.

1. If $f(0)=0$ then return 0 .
2. Compute $g=\operatorname{gcd}\left(f, x^{q}-x\right)$.
3. If $\operatorname{deg} g=0$ then return null.
4. While $\operatorname{deg} g>1$ :
a. Pick a random $\delta \in \mathbb{F}_{q}$.
b. Compute $h=\operatorname{gcd}\left(g,(x+\delta)^{s}-1\right)$.
c. If $0<\operatorname{deg} h<\operatorname{deg} g$ then replace $g$ by $h$ or $g / h$, whichever has lower degree.
5. Return $r=-b / a$, where $g(x)=a x+b$.

It is clear that the output of the algorithm is always correct, since every root of the polynomial $g$ computed in step 2 is a root of $f$, and when $g$ is updated in step 4 c it is always replaced by a proper divisor. We now consider its complexity.

### 4.4.3 Complexity

It follows from Theorem 4.7 that the polynomial $h$ computed in step 4 b is a proper divisor of $g$ with probability at least $1 / 2$, since $g$ has at least two distinct nonzero roots $\alpha, \beta \in \mathbb{F}_{q}$. Thus the expected number of iterations needed to obtain a proper factor $h$ of $g$ is bounded by 2 , and the expected cost of obtaining such an $h$ is $O(\mathrm{M}(e(n+\log e))(n+\log e))$, where $n=\log q$ and $e=\operatorname{deg} g$, and this dominates the cost of the division in step 4c.

Each time $g$ is updated in step 4 c its degree is reduced by at least a factor of 2 . It follows that the expected total cost of step 4 is within a constant factor of the expected time to compute initial polynomial $g$, which is $O(\mathrm{M}(d(n+\log d))(n+\log d))$, where $d=\operatorname{deg} f$, which simplifies to $O(n \mathrm{M}(d n))$ if we assume $\log d=O(n)$; this holds in all the applications we are interested in.

### 4.4.4 Finding all roots

The algorithm can easily be modified to find all the distinct roots of $f$, by modifying step 4c to recursively find the roots of both $h$ and $g / h$. In this case the amount of work done at each level of the recursion tree is bounded by $O(\mathrm{M}(d(n+\log d))(n+\log d))$. Bounding the depth of the recursion is somewhat more involved, but one can show that with very high probability the degrees of $h$ and $g / h$ are approximately equal and that the expected depth of the recursion is $O(\log d)$. Thus we can find all the distinct roots of $f$ in

$$
O(\mathrm{M}(d(n+\log d))(n+\log d) \log d)
$$

expected time. When $\log d=O(n)$ this becomes $O(n \mathrm{M}(d n) \log d)$.
We can also determine the multiplicity of each root of $f$ : assuming $\operatorname{deg} f$ is less than the characteristic of $\mathbb{F}_{q}$, a root $\alpha$ of $f$ occurs with multiplicity $k$ if and only if $\alpha$ is a root of $f^{(k)}$ but not a root of $f^{(k+1)}$, where $f^{(k)}$ denotes the $k$ th derivative of $f$. The time to perform this computation is negligible compared to the time to find the distinct roots.

### 4.4.5 Summary

The table below summarizes the costs of the various arithmetic operations we have considered, both in the integer ring $\mathbb{Z}$, and in a finite field $\mathbb{F}_{q}$, where we assume $q=p^{e}$ with either $p=O(1)$ or $e=O(1)$.

|  | integers $\mathbb{Z}$ | finite field $\mathbb{F}_{q}$ |
| :--- | :--- | :--- |
| addition/subtraction | $O(n)$ | $O(n)$ |
| multiplication | $\mathrm{M}(n)$ | $O(\mathrm{M}(n))$ |
| Euclidean division (reduction) | $O(\mathrm{M}(n))$ | $O(\mathrm{M}(n))$ |
| extended gcd (inversion) | $O(\mathrm{M}(n) \log n)$ | $O(\mathrm{M}(n) \log n)$ |
| exponentiation |  | $O(n \mathrm{M}(n))$ |
| square-roots (probabilistic) |  | $O(n \mathrm{M}(n))$ |
| root-finding (probabilistic) |  | $O(\mathrm{M}(d(n+\log d))(n+\log d))$ |

For integers, $n$ is the bit-length, and for the finite field $\mathbb{F}_{q}$ we let $n=\log q$. In the case of root-finding, $d$ is the degree of the polynomial, and for both square-roots and root-finding these are bounds on the expected running time of a Las Vegas algorithm. The bound for exponentiation assumes the bit-length of the exponent is $O\left(n^{2}\right)$.

## References

[1] Elwyn R. Berlekamp, Factoring polynomials over large finite fields, Mathematics of Computation 24 (1970), 713-735.
[2] Henri Cohen et al., Handbook of elliptic and hyperelliptic curve cryptography, CRC Press, 2006.
[3] Joachim von zur Gathen and Jürgen Gerhard, Modern computer algebra, third edition, Cambridge University Press, 2013.
[4] Daniel M. Gordon, A survey of fast exponentiation methods, Journal of Algorithms 27 (1998), 129-146.
[5] David Harvey, Joris van der Hoeven, and Grégoire Lecerf, Faster polynomial multiplication over finite fields, arXiv:1407.3361.
[6] Michael O. Rabin, Probabilistic algorithms in finite fields, SIAM Journal of Computing 9 (1980), 273-280.

MIT OpenCourseWare
http://ocw.mit.edu

### 18.783 Elliptic Curves

Spring 2015

For information about citing these materials or our Terms of Use, visit: http://ocw.mit.edu/terms.


[^0]:    ${ }^{1}$ The algorithm given in [2] for the precomputation step of Barret reduction uses Newton iteration with a fixed precision, which is asymptotically suboptimal; it is better to use the varying precision approach described above. But in practice the precomputation cost is usually not a major concern.

[^1]:    ${ }^{2}$ An entirely different approach to computing $r$ th roots using discrete logarithms is explored in Problem Set 2. It has better constant factors when the $r$-power torsion subgroup of $\mathbb{F}_{q}^{*}$ is small (which is usually the case), but is asymptotically slower then the algorithm presented here in the worst case.
    ${ }^{3}$ Deterministic polynomial-time bounds for root-finding can be proved in various special cases, including the computation of square-roots, if one assumes a generalization of the Riemann hypothesis.

[^2]:    ${ }^{4}$ The exception is when $d>q$, but in this case computing $\operatorname{gcd}\left(f(x), x^{q}-x\right)$ takes $O(\mathrm{M}(d(n+\log d) \log d)$ time, which turns out to be the same bound that we get for computing $x^{q} \bmod f(x)$ in any case.

