

CDM

Applications of Finite Fields

KLAUS SUTNER

CARNEGIE MELLON UNIVERSITY

FALL 2023



This lecture contains a number of “dirty tricks,” rather than the supremely elegant and deeply rooted mathematical ideas that you have become accustomed to in this class.

Think of it as the engineering approach: some clever trick may just be an opportunistic hack, but, boy, does it ever work well.

- 1 Implementing Finite Fields**
- 2 Discrete Logarithms**
- 3 Advanced Encryption Standard**

In order to implement a finite field we need a data structure to represent field elements and operations on this data structure to perform

- addition
- multiplication
- reciprocals, division
- exponentiation

Subtraction is essentially the same as addition and requires no special attention.

But, anything to do with multiplication is by no means trivial. In fact, there is an annual international conference that is dedicated to the implementation of finite fields.

- Theory of finite field arithmetic:
 - Bases (canonical, normal, dual, weakly dual, triangular ...)
 - Polynomial factorization, irreducible polynomials
 - Primitive elements
 - Prime fields, binary fields, extension fields, composite fields, tower fields ...
 - Elliptic and Hyperelliptic curves
- Hardware/Software implementation of finite field arithmetic:
 - Optimal arithmetic modules
 - Design and implementation of finite field arithmetic processors
 - Design and implementation of arithmetic algorithms
 - Pseudorandom number generators
 - Hardware/Software Co-design
 - IP (Intellectual Property) components
 - Field programmable and reconfigurable systems
- Applications:
 - Cryptography
 - Communication systems
 - Error correcting codes
 - Quantum computing

Our characterization of finite fields as quotients of polynomial rings provides a general method of implementation, and even a reasonably efficient one.

Still, it is a good idea to organize things into several categories, with different algorithmic answers.

- Prime fields \mathbb{Z}_p
- Characteristic 2, \mathbb{F}_{2^k}
- General case \mathbb{F}_{p^k}

For simplicity let's assume that the characteristic p is reasonably small so that arithmetic in \mathbb{Z}_p is $O(1)$.

Note that we are not necessarily looking for asymptotic speed-ups here, constant factors can be quite interesting.

We know how to handle $\mathbb{F}_p = \mathbb{Z}_p$: we need standard addition and multiplication in combination with remainder computations. The Extended Euclidean Algorithm can be used to compute inverses. No problem.

Yet even in \mathbb{Z}_p there is room for some clever computational tricks.

Here is one way to lower the cost of multiplication in the field.

This requires a bit of pre-computation and is only of interest when lots of field multiplications are needed. Suppose we have characteristic $p > 2$.

The key is to use a strange representation of \mathbb{F}^\times that seems to just make a mess.

- Pick $R > p$ coprime, typically $R = 2^k$.
- Represent modular number x by $\widehat{R}(x)$ where

$$\begin{array}{ccc} \widehat{R} : \mathbb{Z}_p & \longrightarrow & \mathbb{Z}_p \\ x & \longmapsto & R \cdot x \bmod p \end{array}$$

Since p and R are coprime, the map \widehat{R} is a bijection: we are just permuting the field elements.

Careful, though, \widehat{R} is not a field isomorphism. Actually, for addition, there is no problem:

$$\widehat{R}(a + b) = \widehat{R}(a) + \widehat{R}(b) \pmod{p}$$

but, unfortunately:

$$\widehat{R}(a \cdot b) = \widehat{R}(a) \cdot \widehat{R}(b) \cdot R^{-1} \pmod{p}$$

To get mileage out of \widehat{R} we need a cheap way to perform reductions:

$$x \rightsquigarrow xR^{-1} \pmod{p}$$

Cheap here just means that we should do better than doing vanilla \pmod{p} operations.

How can we do the reduction cheaply? First, precompute $\alpha = -p^{-1} \bmod R$. Suppose we have an integer x where $0 \leq x < Rp$.

- Compute $x' = x\alpha \bmod R$.
- Then $(x + x'p)/R$ is an integer and

$$(x + x'p)/R = xR^{-1} \bmod p.$$

Recall that $R = 2^k$, so these operations only require integer multiplication and some shifting.

This is more compelling by looking at the actual code.

Here is a typical implementation for $R = 2^{16}$ and some prime p , $2 < p < R$.

```
#define MASK  65535UL      // 2^16 - 1
#define SHFT  16;

// precompute alpha (Euclidean algorithm)

x0  = x & MASK;
x0  = (x0 * alpha) & MASK;
x   += x0 * p;
x   >>= SHFT;
return( x > p ? x-p : x );
```

The only expensive operations are two multiplications.

Let $p = 17$ and pick $R = 64$ so that $\alpha = 15$.

Here are the field elements $\neq 0$ and their representations:

a	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
$\widehat{R}(a)$	13	9	5	1	14	10	6	2	15	11	7	3	16	12	8	4

For example, for $5 \times 7 = 1 \pmod{p}$ we get $x = 14 \times 6 = 84$ and thus

$$x_0 = 84 \times 15 \pmod{64} = 44$$

$$(x + x_0 p) / R = 832 / 64 = 13$$

Indeed 13 corresponds to 1 in our representation, so everything is fine.

Given a generator g for \mathbb{F}^\times in a reasonably small field we can pre-compute and store a logarithm table

$$(h, i) \in \mathbb{F}^\times \times \mathbb{N} \quad \text{where } h = g^i$$

We also need the inverse table with entries (i, h) . For example, for \mathbb{F}_{2^8} the tables require only 512 bytes of memory.

Multiplication is then reduced to table lookups and some (machine-) integer addition and thus very fast.

This technique is used for example in the cryptographic system AES, see below.

These log tables can also be used to speed up computation in larger fields of characteristic 2: instead of dealing directly with \mathbb{F}_{2^k} we think of it as an extension of \mathbb{F}_{2^ℓ} where ℓ is small. We have a tower of fields

$$\mathbb{F}_2 \subseteq \mathbb{F}_{2^\ell} \subseteq \mathbb{F}_{2^k}$$

For example, suppose we are calculating in \mathbb{F}_{2^8} , so a typical element is $a = x^7 + x^6 + x^3 + 1$ or, as a coefficient vector over \mathbb{F}_2 , $(1, 1, 0, 0, 1, 0, 0, 1)$.

By grouping the coefficients into blocks of 2 we get $(3, 0, 2, 1)$ and can think of a as $3z^3 + 2z + 1 \in \mathbb{F}_{2^2}[z]$.

One can verify that ℓ needs to be a divisor of k for \mathbb{F}_{2^ℓ} to be a subfield of \mathbb{F}_{2^k} .

A popular choice is $\ell = 8$ so that the coefficients in the intermediate fields are bytes.

One uses a log table for the intermediate field \mathbb{F}_{2^ℓ} , so arithmetic there is very fast (comparable to the prime field \mathbb{F}_2).

The main field can now be implemented by using polynomials over $\mathbb{F}_{2^\ell}[x]$ of degree less than k/ℓ rather than the original degree less than k .

So there is a trade-off: the ground field becomes more complicated (though we can keep arithmetic there very fast using tables), but the degree of the polynomials decreases, so we are dealing with smaller bases.

Finding “optimal” implementations of finite fields is not easy.

1 **Implementing Finite Fields**

2 **Discrete Logarithms**

3 **Advanced Encryption Standard**

Suppose G is some cyclic finite group with generator g and cardinality n . We can easily exponentiate in G , the operation

$$e \rightsquigarrow g^e$$

takes $O(M \log e)$ steps where M is the cost of a single multiplication in G . But going backwards is apparently hard in many groups:

Given $a \in G$, find e such that $g^e = a$.

This is known as the **discrete logarithm problem**.

Of course, $e = \log_g a$ is trivially computable by a brute force search, but we are here interested in efficient computation when n is large.

In 1976, Whit Diffie and Martin Hellman seized on this apparent difficulty to propose a cryptographic scheme that promises

Secure communication using only insecure channels.

This almost seems logically impossible: if the eavesdropper has complete knowledge about the encryption method used and has full access to the communication channel it would seem that we cannot keep any secret.

Note that this is the idea that gave rise to RSA.

The group used in Diffie/Hellman is the multiplicative group of a finite field \mathbb{F}^* : we can choose the size and implement the group operation quite efficiently.

- Alice and Bob agree on generator g in some finite field $\mathbb{F} = \mathbb{F}_{p^k}$.
- Alice generates random number x , computes $a = g^x$ in \mathbb{F} .
Alice sends a to Bob.
- Bob generates random number y , computes $b = g^y$ in \mathbb{F} .
Bob sends b to Alice.
- Both Alice and Bob can now compute

$$c = g^{xy} = a^y = b^x$$

and use it as a secret key (for some other encryption algorithm).

Evilstdoer Charlie knows the algorithm, \mathbb{F} , g , a and b , but **not** x and y .

Apparently, Charlie cannot determine c without a huge search, so we only need to make \mathbb{F} large enough to foil his efforts.

Diffie/Hellman seems to hold the promise of secure communication over utterly insecure channels.

Great, but now there is an immediate challenge: can one break Diffie/Hellman, at least in some special cases.

“Break” would mean that we are happy to invest quite a bit of computation, just not the full brute-force exponential search that seems to be necessary to destroy the schema.

In other words: is there anything we can do to speed up computation of logarithms in a finite field?

Suppose g is a generator of the multiplicative subgroup and $a \neq 0$ is some element in a field of size $q = p^k$.

Let $m = \lceil \sqrt{q} \rceil$ and compute two lists

- ag^{-i} where $0 \leq i < m$, and
- g^{mj} where $0 \leq j < m$.

Then check for a common entry in the two lists: this produces $ag^{-i} = g^{mj}$, whence $a = g^{mj+i}$.

So we have essentially written the logarithm as a two-digit number in base m .


To check for a common element we can use hashing. Of course, one need not wait till the lists are complete to check for a match.

The baby-step/giant-step method requires $O(\sqrt{q} \log q)$ time and space $O(\sqrt{q})$.

This may not seem overly impressive, but it is a huge improvement over the standard $O(q)$ time algorithm (though that runs constant space).

Note that a cryptographic attack may well be worth this much computation. The NSA sure won't mind.

This should be called Pollard's Lasso method (in particular since the second algorithm in the paper is about "catching kangaroos"), but it's too late now.



A Rohrschach test:

- If you have a classical education, you will see a ρ .
- If you're a cowboy, you will see a lasso.

The motivation for this method is a bit strange. Consider a random function $f : A \rightarrow A$ where A has size n .

Then the expected value of some key parameters of the functional digraph of f are as follows:

# components	$\frac{1}{2} \log n$
# leaf nodes	$e^{-1}n$
# recurrent nodes	$\sqrt{\pi n/2}$
transient length	$\sqrt{\pi n/8}$
period length	$\sqrt{\pi n/8}$

The expected lengths of the longest transient/cycle are also $c_{1/2}\sqrt{n}$ where $c_1 \approx 1.74$ and $c_2 \approx 0.78$.

For simplicity we can think of the expected value of transient length t and period length p of a random point a in A as \sqrt{n} .

We know an elegant algorithm to compute these parameters: Floyd's trick. More precisely, we can compute t and p in expected time $O(\sqrt{n})$ using $O(1)$ space (we only need to store a small constant number of elements in A).

Wild Idea:

Can we compute a sequence (x_i) of elements as in Shanks' algorithm that behaves like a (pseudo-) random sequence so that $x_i = x_{2i}$ allows us to compute a discrete logarithm?

We need a “random” map.

To this end we first split the group G into three sets G_1 , G_2 and G_3 of approximately equal size (sets, not subgroups, so this will be easy in practical situations). Any ham-fisted approach will do.

Now, given a generator g and some element a , define $f : G \rightarrow G$ as follows:

$$f(x) = \begin{cases} gx & \text{if } x \in G_1, \\ x^2 & \text{if } x \in G_2, \\ ax & \text{otherwise.} \end{cases}$$

Of course, f is perfectly deterministic (at least given the partition of G).

Consider the orbit (x_i) of 1 under f , $x_0 = 1$.

Clearly, all the elements have the form $a^{\alpha_i} g^{\beta_i}$ and the exponents are updated according to

$$(\alpha_{i+1}, \beta_{i+1}) = \begin{cases} (\alpha_i, \beta_i + 1) & \text{if } x \in G_1, \\ (2\alpha_i, 2\beta_i) & \text{if } x \in G_2, \\ (\alpha_i + 1, \beta_i) & \text{otherwise.} \end{cases}$$

Since the partition of G is random, the three steps are chosen randomly (more or less).

Moreover, using Floyd's method we can find the minimal index e such that $x_e = x_{2e}$, i.e.,

$$a^{\alpha_e} g^{\beta_e} = a^{\alpha_{2e}} g^{\beta_{2e}}$$

But then

$$a^{\alpha_e - \alpha_{2e}} = g^{\beta_{2e} - \beta_e}$$

This equality does not directly solve the discrete logarithm problem but it can help a lot to compute the discrete logarithm.

Note that for cryptographic application any such weakness is potentially fatal: a good method must be secure under any and all circumstances.

Consider the multiplicative group of \mathbb{Z}_p where $p = 999959$, and let $g = 7$ and $a = 3$.

Running the algorithm produces $e = 1174$ and $x_e = 11400$, plus the identity

$$3^{310686} = 7^{764000} \pmod{p}$$

A classical case of “close, but no cigar”: we want to compute the logarithm of 3 with respect to generator 7, some η such that: $3 = 7^\eta \pmod{p}$.

But perhaps we squeeze a bit more information out of this identity? The identity lives over \mathbb{Z}_p^* .


Use the Extended Euclidean algorithm to get

$$\begin{aligned}\gcd(310686, p-1) &= 2 \\ &= 148845 \cdot 310686 - 46246 \cdot 999958\end{aligned}$$

Then raise 3^{310686} to the 148845 power mod p to obtain

$$\begin{aligned}3^2 &= 7^{356324} \pmod{p} \\ 3 &= \pm 7^{178162} \pmod{p}\end{aligned}$$

We can simply check the two cases and find that in \mathbb{Z}_p : $\log_7 3 = 178162$.



A plot of the orbit of 1 given our a and g .

1 **Implementing Finite Fields**

2 **Discrete Logarithms**

3 **Advanced Encryption Standard**

The now classical **DES** (data encryption standard) was officially adopted in 1977. It is based on—rather too short—keys of length 56 bits and has since fallen prey to Moore's law: DES can now be broken in a distributed attack in a matter of hours.

In September 1997, NIST issued a Federal Register notice soliciting an unclassified, publicly disclosed encryption algorithm.

15 candidate algorithms were submitted and closely scrutinized. In 2000 the NIST selected the **Rijndael** algorithm by Joan Daemen and Vincent Rijmen as the new standard.

It is now enshrined in the Federal Information Processing Standard (FIPS) for the Advanced Encryption Standard, FIPS-197, see


<http://csrc.nist.gov/CryptoToolkit/aes>

DES and AES are both based on ideas that first appeared in **Lucifer**, an encryption system developed mainly by IBM: the input message is chopped into fixed-size block of bits, which are then clobbered using a **key**.

- Use combinatorics and algebra to mangle a block of bits.
- Cleverly incorporate the key into this mangling process.
- Use multiple rounds to make sure the final result is sufficiently complicated.

Each round uses a **subkey** (aka **round key**) that is generated from the main key.

And, of course, everything has to be easily reversible when the key is known.



Each coding box is relatively simple, and may not provide a safe encoding. But a sufficiently long chain is hard to crack without knowledge of the master key.

There are several patented variants, all based on substitution-permutation networks (so-called [Feistel networks](#)) that mangle bits and mix in the key in some clever way. Up to 16 rounds are used to foil attacks.

The block size and key size vary from 48 to 128 bits.

A pleasant feature is that decryption is very similar to encryption, so hardware can be reused.

Alas, the devil is in the details[†], and Lucifer suffered from security issues.

[†]No pun intended.

DES encrypts blocks of 64 bits, using a key of 56 bits.

A 64-bit input block is permuted and then split into two 32-bit blocks (L_0, R_0) .

These blocks are then mangled in several rounds according to

$$(L_{i+1}, R_{i+1}) = (R_i, L_i \oplus f(R_i, K_i))$$

Here K_i is a key derived from the original key $K \in \mathbf{2}^{56}$ and $f : \mathbf{2}^{32+48} \rightarrow \mathbf{2}^{32}$ is a carefully constructed Boolean map. Tempting, but we won't go there.

The final output is then obtained from (L_{16}, R_{16}) .

$$2^{56} \quad 7.2 \times 10^{16}$$

$$2^{128} \quad 3.4 \times 10^{38}$$

$$2^{256} \quad 1.2 \times 10^{77}$$

Even with a million processors, the longer keys cannot be brute-forced.



AES encrypts blocks of 128 bits, using a cipher key of 128 (or 192, 256 bits).
Bit-sequences in AES are always divided into bytes, 8-bit blocks.

Finite fields and/or polynomials are used in two places:

- We can think of these bytes as coefficient vectors of elements in \mathbb{F}_{2^8} where the irreducible polynomial for the multiplicative structure is chosen to be

$$f(x) = x^8 + x^4 + x^3 + x^2 + 1$$

- 4-byte vectors are construed as polynomials in $\mathbb{F}_{2^8}[z]/(z^4 + 1)$.

The algorithm first xors with a subkey, and then proceeds in 10 rounds (actually, the number depends on the key size, but let's just focus on 128-bit keys). As in DES, each round mangles the bits some more (the final round is slightly different, but we will ignore this).

Abstractly, a single round looks like so:

- byte substitution
- shifting rows
- mixing columns
- add key

The row/column terminology comes from thinking of the initial input as being given by a 4×4 matrix of bytes (for a total of 128 bits; in reality the input is broken into corresponding pieces).

$$\begin{pmatrix} a_{0,0} & a_{0,1} & a_{0,2} & a_{0,3} \\ a_{1,0} & a_{1,1} & a_{1,2} & a_{1,3} \\ a_{2,0} & a_{2,1} & a_{2,2} & a_{2,3} \\ a_{3,0} & a_{3,1} & a_{3,2} & a_{3,3} \end{pmatrix}$$

So the current state is described by such a matrix which we may think of as a 4×4 matrix over our favorite field \mathbb{F}_{2^8} .



Define the **patched inverse** of an element $a \in \mathbb{B}$ to be

$$\bar{a} = \begin{cases} 0 & \text{if } a = 0, \\ a^{-1} & \text{otherwise.} \end{cases}$$

Define an 8×8 bit-matrix and 8-bit vector as follows

$$A = \begin{pmatrix} 1 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 & 0 & 1 & 1 & 1 \\ 1 & 1 & 1 & 0 & 0 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 & 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 \end{pmatrix} \quad v = \begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 1 \\ 1 \\ 0 \end{pmatrix}$$

Then byte substitution is given by the (almost) affine, reversible map $a \mapsto A\bar{a} + v$, applied to each byte separately.



On the right, we simply use a rather than the patched inverse as on the left.

Replace the state matrix by the row-shifted version

$$\begin{pmatrix} a_{0,0} & a_{0,1} & a_{0,2} & a_{0,3} \\ a_{1,1} & a_{1,2} & a_{1,3} & a_{1,0} \\ a_{2,2} & a_{2,3} & a_{2,0} & a_{2,1} \\ a_{3,3} & a_{3,0} & a_{3,1} & a_{3,2} \end{pmatrix}$$

More precisely, we shift the i th row by i places to the left (assuming 0-indexing).

The next operation is a bit more complicated.

For phase (3) we will use the same notation as in the Rijndael specification.

To denote an element of \mathbb{B} , we think of a byte as two hexadecimal digits.

So, for example, D4 corresponds to the field element

$$z^7 + z^6 + z^4 + z^2 \text{ mod } \tau$$

in polynomial notation, or 11010100 as a coefficient vector.

Since the byte field \mathbb{B} is quite small, we can easily use lookup tables to make all the field operations very fast.

Consider the polynomial (coefficients are written as two hex digits)

$$g(z) = 03z^3 + 01z^2 + 01z + 02 \in \mathbb{B}[z]$$

We can think of each column in the state matrix as another polynomial in $\mathbb{B}[z]$, so in this phase we multiply the column polynomial by g , and then reduce modulo $z^4 + 1$.

Since these operations are all linear, this all comes down to a single matrix multiplication over \mathbb{B} :

$$\mathbf{c} \rightsquigarrow \begin{pmatrix} 02 & 03 & 01 & 01 \\ 01 & 02 & 03 & 01 \\ 01 & 01 & 02 & 03 \\ 03 & 01 & 01 & 02 \end{pmatrix} \mathbf{c}$$

Well ...

Take a look at the NIST specifications for cryptography and Rijndael:

[ToolKit](#)

[Rijndael](#)

There are lots of implementation details as well as a careful discussion how to decrypt Rijndael encrypted message.

One pleasant aspect: the decryption operations are very similar to encryption, so essentially the same hardware can be used.

Claim: All four phases are reversible.

At each round, a subkey is xor-ed with the state matrix; all the subkeys are 128 bits. We will not discuss how the subkeys are generated from the original key. Incidentally, proper key management is a huge problem in cryptography.

The documents at the links also contain lots of implementation details as well as a careful discussion how to decrypt Rijndael encrypted message. Note that the inverse operations are very similar to the encryption operations, so essentially the same hardware can be used.

The xor operation with the keys is self-inverse (we're in characteristic 2), so step 4 in each round is easily undone by anyone with access to the key.

States 1–3 are also invertible, none of the operations loses information.

As with DES, the decryption process is quite similar to encryption.

With appropriate CPU support the throughput is quite high (hundreds of MB/s).