MODELLING FINITE FIELDS

Hendrik Lenstra

Mathematisch Instituut
Universiteit Leiden

Finite fields

A *finite field* is a finite set E equipped with elements $0,\, 1 \in E$ and maps $+,\, \cdot : E \times E \to E$ such that for all $a,\, b,\, c \in E$ one has

$$\begin{aligned} (a \cdot b) \cdot c = a \cdot (b \cdot c), &\quad (a+b)+c = a+(b+c), \\ \exists d : d + a = 0, &\quad (\exists e : e \cdot a = 1) \Leftrightarrow a \neq 0, \\ 1 \cdot a = a, &\quad (a+b) \cdot c = (a \cdot c) + (b \cdot c), \\ 0 + a = a, &\quad a \cdot (b+c) = (a \cdot b) + (a \cdot c). \end{aligned}$$

Two magic squares of Lee Sallows

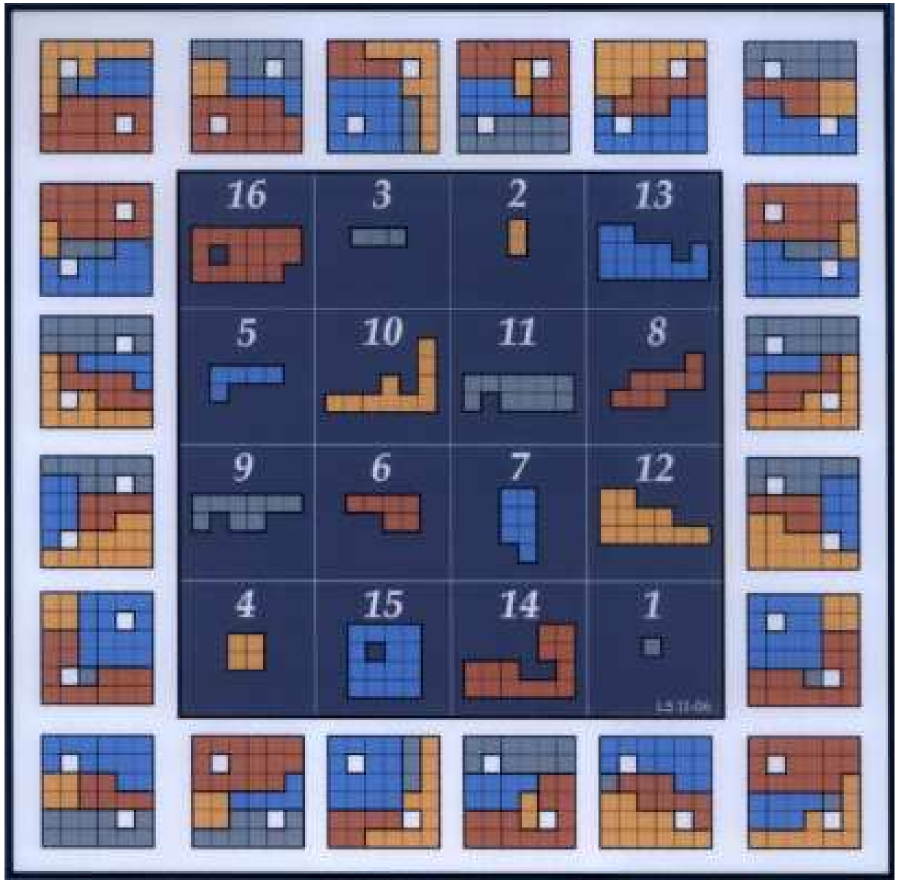

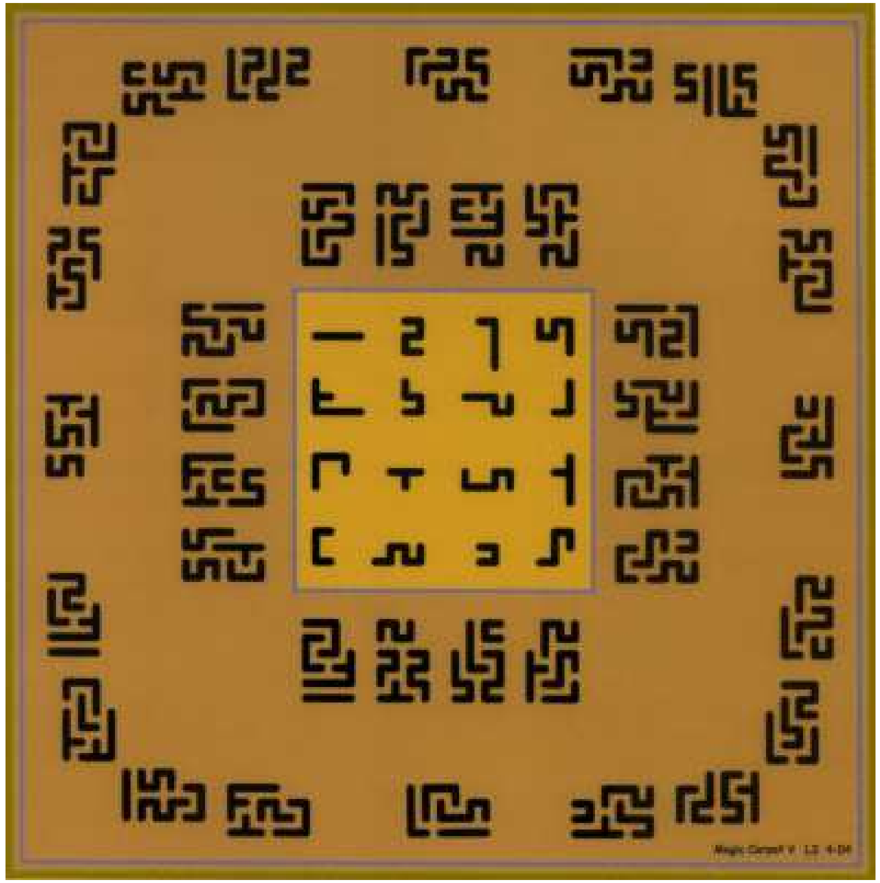

Prime fields

Example: for p prime, $\mathbf{F}_p = \mathbf{Z}/p\mathbf{Z} = \{0, 1, \ldots, p-1\}$ is a field of size p.

Prime fields

Example: for p prime, $\mathbf{F}_p = \mathbf{Z}/p\mathbf{Z} = \{0, 1, \ldots, p-1\}$ is a field of size p.

Let E be a finite field. The subset $\{1 + 1 + \ldots + 1\}$ is the *prime field* of E. It may be identified with $\mathbf{F}_p$ for a unique prime p, the *characteristic* $\operatorname{char} E$ of E.

Finite fields everywhere

Finite fields occur in

- finite group theory,
- algebraic number theory,
- statistics,
- combinatorics,
- algebraic geometry,
- coding theory,
- cryptography,
- ...

Degree and cardinality

Let E be a finite field, and $p = \operatorname{char} E$.

The *degree* $\deg E$ of E is the least number of generators of the additive group of E, which is the same as $\dim_{\mathbf{F}_p} E$.

If $\deg E = n$ then $\# E = p^n$.

A field of size 4

Any set $\{0, 1, \alpha, \beta\}$ of size 4 has exactly one field structure with zero element 0 and unit element 1. Notation: $\mathbf{F}_4$.

Addition: $x + x = 0$ for all x, and any two of $\{1, \alpha, \beta\}$ add up to the third. *Multiplication*: $\alpha^2 = \alpha^{-1} = \beta$.

One has $\operatorname{char} \mathbf{F}_4 = \deg \mathbf{F}_4 = 2$.

Other quadratic finite fields

Let p be an odd prime, and let $c \in \mathbf{F}_p = \mathbf{Z}/p\mathbf{Z}$ be such that

$$c^{(p-1)/2} = -1 \, (= p - 1).$$

Then the set $\mathbf{F}_p \oplus \mathbf{F}_p\sqrt{c}$ consisting of the p^2 expressions $\{a + b\sqrt{c}\}$ with a, $b \in \mathbf{F}_p$ is a field, the multiplication being determined by $\sqrt{c}^2 = c$.

It has characteristic p and degree 2.

Classifying finite fields

Theorem (E. Galois, 1830; E. H. Moore, 1893)**.**
There is a bijective map

$$\{\text{finite fields}\}/\cong \;\longrightarrow\; \{\text{primes}\} \times \mathbf{Z}_{>0}$$

sending $[E]$ *to* $(\operatorname{char} E, \deg E)$.

A field of size p^n is denoted by $\mathbf{F}_{p^n}$ or $\mathrm{GF}(p^n)$.

Founding fathers

Évariste Galois

(1811–1832)

Eliakim Hastings Moore

(1862–1932)

Classifying finite fields

Theorem (E. Galois, 1830; E. H. Moore, 1893)**.** *There is a bijective map*

$$\{\text{finite fields}\}/\cong \;\longrightarrow\; \{\text{primes}\} \times \mathbf{Z}_{>0}$$

sending $[E]$ *to* $(\operatorname{char} E, \deg E)$*.*

A field of size p^n is denoted by $\mathbf{F}_{p^n}$ or $\mathrm{GF}(p^n)$.

The number of isomorphisms between two fields of size p^n equals n, so for $n \geq 2$ a field of size p^n is not *uniquely unique.*

Modelling $\mathbf{F}_{p^n}$

- $\mathbf{F}_{p^n}$ = any set of size p^n,
  addition and multiplication by table look-up;
- $\mathbf{F}_{p^n} = \{\infty\} \amalg \left(\mathbf{Z}/(p^n - 1)\mathbf{Z}\right)$,
  multiplication = addition modulo $p^n - 1$,
  $x \mapsto x + 1$ by table look-up (*Zech logarithm*),
  $a + b = (ab^{-1} + 1) \cdot b$ for $b \neq 0$.

Vector space models

- $n = 1$: $\mathbf{F}_p = \mathbf{Z}/p\mathbf{Z} = \{0, 1, \ldots, p-1\}$,
  addition and multiplication modulo p;
- general n: $\mathbf{F}_{p^n} = (\mathbf{Z}/p\mathbf{Z})^n = \bigoplus_{i=0}^{n-1} \mathbf{F}_p \cdot e_i$,
  addition is vector addition,
  multiplication is determined by

  $$e_i \cdot e_j = \sum_{k=0}^{n-1} a_{ijk} e_k$$

  for certain $a_{ijk} \in \mathbf{F}_p$.

Special cases

- $\mathbf{F}_{p^n} = \mathbf{F}_p[X]/(f)$, where $f \in \mathbf{F}_p[X]$ is monic of degree n and irreducible, with basis $\{X^i \bmod f : 0 \leq i < n\}$;
- towers or tensor products of such fields;
- subfields of fields given by vector space models.

Explicit models

An *explicit model* for a field of size p^n is a field with additive group $\mathbf{F}_p^n = \bigoplus_{i=0}^{n-1} \mathbf{F}_p \cdot e_i$, where $\mathbf{F}_p = \mathbf{Z}/p\mathbf{Z}$.

Such a model is numerically specified by the system $(a_{ijk})_{i,j,k=0}^{n-1}$ of elements $a_{ijk} \in \mathbf{F}_p$ satisfying

$$e_i \cdot e_j = \sum_{k=0}^{n-1} a_{ijk} e_k \quad \text{for all } i,\ j.$$

Space: $O(n^3 \log p)$.

Example

For odd p, the field

$$\mathbf{F}_{p^2} = \mathbf{F}_p \oplus \mathbf{F}_p\sqrt{c}$$

(where $c \in \mathbf{F}_p$ satisfies $c^{(p-1)/2} = -1$)
is specified by

$$a_{000} = a_{011} = a_{101} = 1,$$

$$a_{110} = c,$$

$$a_{ijk} = 0 \text{ whenever } i + j + k \text{ is odd.}$$

A converse

Exercise. If $(a_{ijk})_{i,j,k=0}^{1}$ defines a field of size p^2, with p odd, and

$$b_{ij} = \textstyle\sum_{0\le k,l\le 1} a_{ijk}a_{kll},$$

$$c = b_{00}b_{11} - b_{01}b_{10} \in \mathbf{F}_p,$$

then one has $c^{(p-1)/2} = -1$.

A converse

Exercise. If $(a_{ijk})_{i,j,k=0}^{1}$ defines a field of size p^2, with p odd, and

$$b_{ij} = \sum_{0 \le k,l \le 1} a_{ijk} a_{kll},$$

$$c = b_{00}b_{11} - b_{01}b_{10} \in \mathbf{F}_p,$$

then one has $c^{(p-1)/2} = -1$.

Conclusion. *Constructing* $\mathbf{F}_{p^2}$ *is "equivalent" to finding* $c \in \mathbf{F}_p$ *with* $c^{(p-1)/2} = -1$.

Finding a quadratic non-residue

For an odd prime p, the number of $c \in \mathbf{F}_p$ with $c^{(p-1)/2} = -1$ equals $(p-1)/2$.

Hence there is a probabilistic algorithm with polynomial expected run time that, given p, finds such an element c.

No deterministic polynomial-time algorithm for this problem is known.

Constructing finite fields

Conjecture. *For some $t \in \mathbf{R}_{>0}$, there is an algorithm that for given p, n constructs in time at most $(n + \log p)^t$ an explicit model for a field of size p^n.*

Constructing finite fields

Conjecture. *For some* $t \in \mathbf{Z}_{>0}$*, there is an algorithm that for given* p*,* n *constructs in time at most* $(n + \log p)^t$ *an explicit model for a field of size* p^n*.*

This is correct

- if a probabilistic algorithm is allowed,
- if GRH is true,
- if p is fixed.

Classifying finite fields

Theorem (E. Galois, 1830; E. H. Moore, 1893)**.**
There is a bijective map

$$\{\text{finite fields}\}/\cong \;\longrightarrow\; \{\text{primes}\} \times \mathbf{Z}_{>0}$$

sending $[E]$ *to* $(\operatorname{char} E, \deg E)$.

A field of size p^n is denoted by $\mathbf{F}_{p^n}$ or $\mathrm{GF}(p^n)$.

The number of isomorphisms between two fields of size p^n equals n, so for $n \geq 2$ a field of size p^n is not *uniquely unique*.

Isomorphisms of quadratic fields

Let p be an odd prime.

If c, $d \in \mathbf{F}_p$ satisfy $c^{(p-1)/2} = d^{(p-1)/2} = -1$, then the number of $e \in \mathbf{F}_p$ with $c = e^2 \cdot d$ equals 2, and for each such e the map

$$\mathbf{F}_p \oplus \mathbf{F}_p\sqrt{c} \to \mathbf{F}_p \oplus \mathbf{F}_p\sqrt{d}$$

$$a + b\sqrt{c} \mapsto a + be\sqrt{d}$$

is a field isomorphism.

What does the notation $\mathbf{F}_{p^n}$ *mean?*

- "the" finite field of size p^n, well-defined only up to isomorphism,
- a finite field of size p^n,
- $\{\alpha \in \bar{\mathbf{F}}_p : \alpha^{p^n} = \alpha\}$, where $\bar{\mathbf{F}}_p$ is an algebraic closure of $\mathbf{Z}/p\mathbf{Z}$.

What does the notation $\mathbf{F}_{p^n}$ mean?

- "the" finite field of size p^n, well-defined only up to isomorphism,
- a finite field of size p^n,
- $\{\alpha \in \bar{\mathbf{F}}_p : \alpha^{p^n} = \alpha\}$, where $\bar{\mathbf{F}}_p$ is an algebraic closure of $\mathbf{Z}/p\mathbf{Z}$.

Bourbaki: "par abus de langage".

M. Artin: "this notation is not too ambiguous".

Should we care?

What does the notation $\mathbf{C}$ *mean?*

Unsatisfactory definitions:

- "the" quadratic field extension of $\mathbf{R}$,
- "the" algebraic closure of $\mathbf{R}$.

Satisfactory definition:

- $\mathbf{C} = \mathbf{R}[X]/(X^2 + 1)$.

Three models for the field of complex numbers

- $\mathbf{R} \times \mathbf{R}$, with $(a, b) \cdot (c, d) = (ac - bd, ad + bc)$,

- $\{\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in M(2, \mathbf{R}) : a = d,\ b + c = 0\}$,

- $(\mathbf{R}1 \oplus \mathbf{R}\gamma \oplus \mathbf{R}\delta)/\mathbf{R}{\cdot}(1 + \gamma + \delta)$, with $\gamma^2 = \gamma^{-1} = \delta$.

Any two of these admit *two* $\mathbf{R}$-isomorphisms.

Finding consistent identifications

In each model, single out a special square root of -1.

Choose the isomorphism under which these special square roots correspond.

Finding consistent identifications

In each model, single out a special square root of -1.

Choose the isomorphism under which these special square roots correspond.

Equivalently: for each model, pick an isomorphism with the *standard model* $\mathbf{R}[X]/(X^2 + 1)$, and let the isomorphisms pass through the standard model.

Why define $\mathbf{F}_{p^n}$*?*

Three computer-related reasons:

- it helps finding consistent isomorphisms between finite fields of the same size;
- it is convenient in computer algebra systems;
- formal correctness enhances computer-checkability.

Desirable properties of $\mathbf{F}_{p^n}$

(i) there are compatible embeddings $\mathbf{F}_{p^n} \subset \mathbf{F}_{p^m}$ for $n|m$;

(ii) $\mathbf{F}_{p^n}$ is easy to construct;

(iii) it is easy to identify any given field of size p^n with $\mathbf{F}_{p^n}$.

Definition with Conway polynomials

$\mathrm{GF}(p^n) = \mathbf{Z}[X]/(p, f_{p,n})$, where $f_{p,n} \in \mathbf{Z}[X]$ is the *Conway polynomial*, see

http://www.math.rwth-aachen.de/~Frank.Luebeck/data/ConwayPol/

Definition with Conway polynomials

$\mathrm{GF}(p^n) = \mathbf{Z}[X]/(p, f_{p,n})$, where $f_{p,n} \in \mathbf{Z}[X]$
is the *Conway polynomial*, see

http://www.math.rwth-aachen.de/
~Frank.Luebeck/data/ConwayPol/

$f_{p,n} = X^n - a_1 X^{n-1} + a_2 X^{n-2} - \ldots + (-1)^n a_n$,
with $(a_1, a_2, \ldots, a_n) \in \{0, 1, \ldots, p-1\}^n$
lexicographically minimal such that

- $\big(\mathbf{Z}[X]/(p, f_{p,n})\big)^* = \langle \bar{X} \rangle \cong \mathbf{Z}/(p^n - 1)\mathbf{Z}$,
- $f_{p,d}(X^{(p^n-1)/(p^d-1)}) \in (p, f_{p,n})$ for each $d|n$.

Desirable properties of $\mathbf{F}_{p^n}$

(i) there are compatible embeddings $\mathbf{F}_{p^n} \subset \mathbf{F}_{p^m}$ for $n|m$;

(ii) $\mathbf{F}_{p^n}$ is easy to construct;

(iii) it is easy to identify any given field of size p^n with $\mathbf{F}_{p^n}$.

How do Conway polynomials score?

The fields $\mathrm{GF}(p^n)$ as just defined satisfy (i), they do *not* satisfy (ii), but once $\mathrm{GF}(p^n)$ has been constructed, it satisfies (iii).

Due to their algorithmic inaccessibility, Conway polynomials need to be replaced.

Existence

Theorem (Bart de Smit & HWL)**.**
One can define explicit models $\mathbf{F}_{p^n}$*,*
one for each pair (p, n)*, such that*
(i)*,* (ii)*, and* (iii) *are satisfied.*

Desirable properties of $\mathbf{F}_{p^n}$

(i) there are compatible embeddings $\mathbf{F}_{p^n} \subset \mathbf{F}_{p^m}$ for $n|m$;

(ii) $\mathbf{F}_{p^n}$ is easy to construct;

(iii) it is easy to identify any given field of size p^n with $\mathbf{F}_{p^n}$.

Existence and uniqueness

Theorem (Bart de Smit & HWL). *One can define explicit models* $\mathbf{F}_{p^n}$*, one for each pair* (p, n)*, such that* (i), (ii), *and* (iii) *are satisfied.*

There is a sense in which the sequence $(\mathbf{F}_{p^n})_{p,n}$ of explicit models is uniquely determined.

Property (ii) *in the quadratic case*

Theorem. *There is a probabilistic algorithm with polynomial expected run time that, on input an odd prime p, finds $c \in \mathbf{F}_p$ with $c^{(p-1)/2} = -1$, and that finds the same c when called twice for the same p.*

Property (ii) *in the quadratic case*

Theorem. *There is a probabilistic algorithm with polynomial expected run time that, on input an odd prime p, finds $c \in \mathbf{F}_p$ with $c^{(p-1)/2} = -1$, and that finds the same c when called twice for the same p.*

The output of the algorithm on input p is called the *standard quadratic non-residue* modulo p, notation: $s(p)$.

Property (iii) *in the quadratic case*

Theorem. *There is a deterministic polynomial-time algorithm that, on input an odd prime p and an element $d \in \mathbf{F}_p$ with $d^{(p-1)/2} = -1$, computes $s(p)$ as well as $e \in \mathbf{F}_p$ with $s(p) = e^2 \cdot d$.*

Existence of s

Define

$$s(p) = \sqrt{\sqrt{\cdots\sqrt{1}}},$$

each squareroot being chosen in

$$\{(p+1)/2, \ldots, p-2, p-1\},$$

and the number of $\sqrt{\ }$-signs being the number of factors 2 in $p-1$.

One can show that s has all asserted properties.

A table of standard quadratic non-residues

p	$s(p)$	p	$s(p)$	p	$s(p)$	p	$s(p)$	p	$s(p)$
3	2	29	17	61	50	101	91	139	138
5	3	31	30	67	66	103	102	149	105
7	6	37	31	71	70	107	106	151	150
11	10	41	27	73	51	109	76	157	129
13	8	43	42	79	78	113	78	163	162
17	14	47	46	83	82	127	126	167	166
19	18	53	30	89	77	131	130	173	93
23	22	59	58	97	78	137	127	179	178

Uniqueness of s

Let $s'(p) \in \mathbf{F}_p$, $s'(p)^{(p-1)/2} = -1$,
for each odd prime p.

Theorem. *The function* s' *also has property* (iii) *if and only if there is a function* f *that can be computed in polynomial time such that for all* p *one has* $f\big(p, s(p)\big) \in \mathbf{F}_p$ *and*

$$s'(p) = f\big(p, s(p)\big)^2 \cdot s(p).$$

Property (iii) *in the quadratic case*

Theorem. *There is a deterministic polynomial-time algorithm that, on input an odd prime* p *and an element* $d \in \mathbf{F}_p$ *with* $d^{(p-1)/2} = -1$*, computes* $s(p)$ *as well as* $e \in \mathbf{F}_p$ *with* $s(p) = e^2 \cdot d$*.*

Standard models for finite fields

For p odd, write $\mathbf{F}_{p^2} = \mathbf{F}_p \cdot 1 \oplus \mathbf{F}_p \cdot \sqrt{s(p)}$. It is an explicit model for a field of size p^2, called the *standard model*.

For general p and n, one can define the *standard model* for a field of size p^n, notation: $\mathbf{F}_{p^n}$.

It is an explicit model, and the sequence $(\mathbf{F}_{p^n})_{p,n}$ has the desired properties.

Desired properties

(i) there are compatible embeddings $\mathbf{F}_{p^n} \subset \mathbf{F}_{p^m}$ for $n|m$;

(ii) $\mathbf{F}_{p^n}$ is easy to construct;

(iii) it is easy to identify any given field of size p^n with $\mathbf{F}_{p^n}$.

Existence of the standard models

See

http://www.math.leidenuniv.nl/
~desmit/papers/standard_models.pdf

(Bart de Smit & HWL).

Property (iii) *in general*

Main theorem (Bart de Smit & HWL)**.**
There is a polynomial-time algorithm that on input p, n, and an explicit model A for a field of size p^n, computes the standard model $\mathbf{F}_{p^n}$ as well as a field isomorphism $\mathbf{F}_{p^n} \to A$.

Two more lectures

Thursday:

fundamental algorithms

for finite fields.

Friday:

the structure of $\bar{\mathbf{F}}_p$,

construction of the

standard model.