Lattices and Cryptography

in a post-quantum world

Richard Pinch

Strategic Advisor, Mathematics & Security Research, GCHQ
Visiting fellow, Heilbronn Institute for Mathematical Research

20 March 2017

Disclaimer

This talk does not represent the position of Her Majesty's Government, GCHQ or NCSC. It does not constitute an invitation to tender.

Images from GCHQ archives are Crown Copyright and reproduced courtesy of Director GCHQ.

A brief history of cryptography

- Traditional cryptography — Secret key Julius Caesar 55BC ($+3 \bmod 24$) via one-time pad to the present day (AES, TLS/SSL, 3G, 4G)
- Modern cryptography — Public key Cliff Cocks 1972 (RSA) to the present day (RSA, DH, TLS/SSL)
- Post-modern cryptography — Quantum and post-quantum ~~2014~~ 2006 on

Traditional cryptography

Shared key *(symmetric)* cryptography achieves a number of objectives:

- Privacy/confidentiality — only a holder of the common secret can *read*
- Authentication — only a holder of the common secret can *write*
- Integrity — only a holder of the common secret can *change*
- Non-repudiation — only a holder of the common secret can *write*

The paradox of traditional cryptography

Shared key cryptography depends on sharing the secret key securely.

So how do we send the secret *without* the secure channel in the first place?

Computation

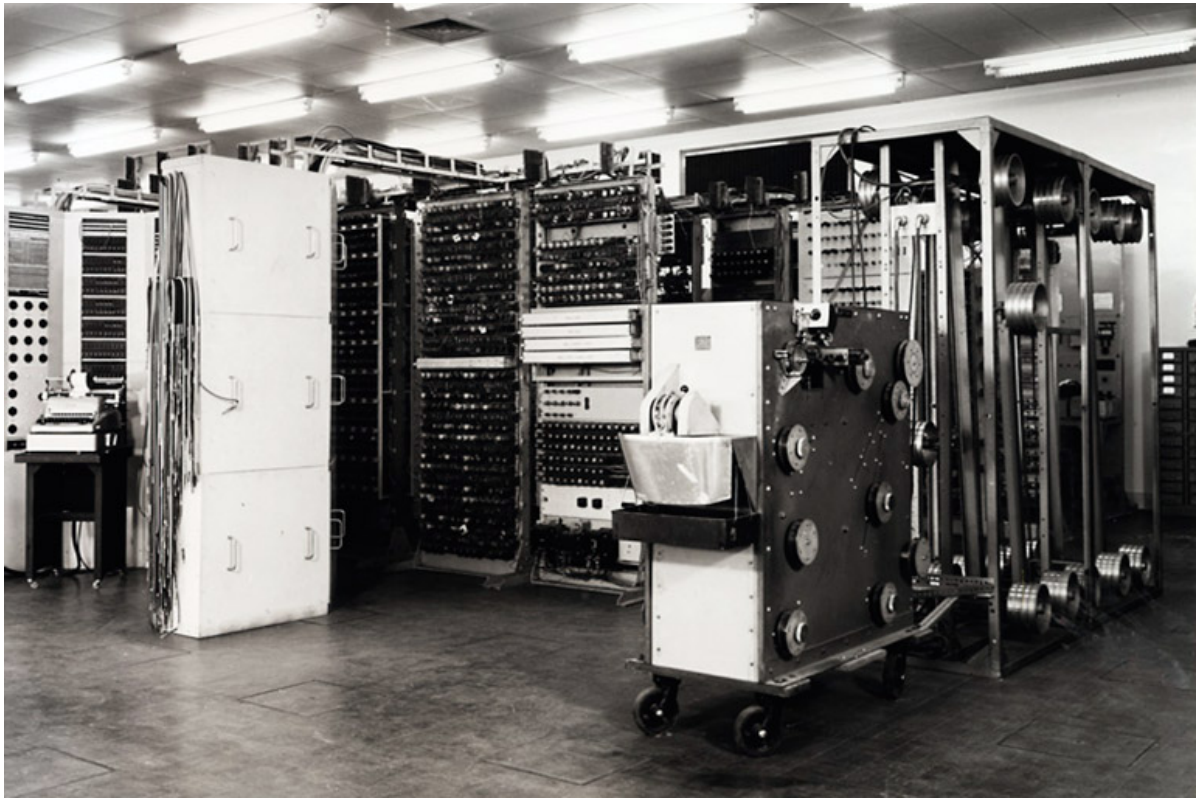

Non-secret encryption

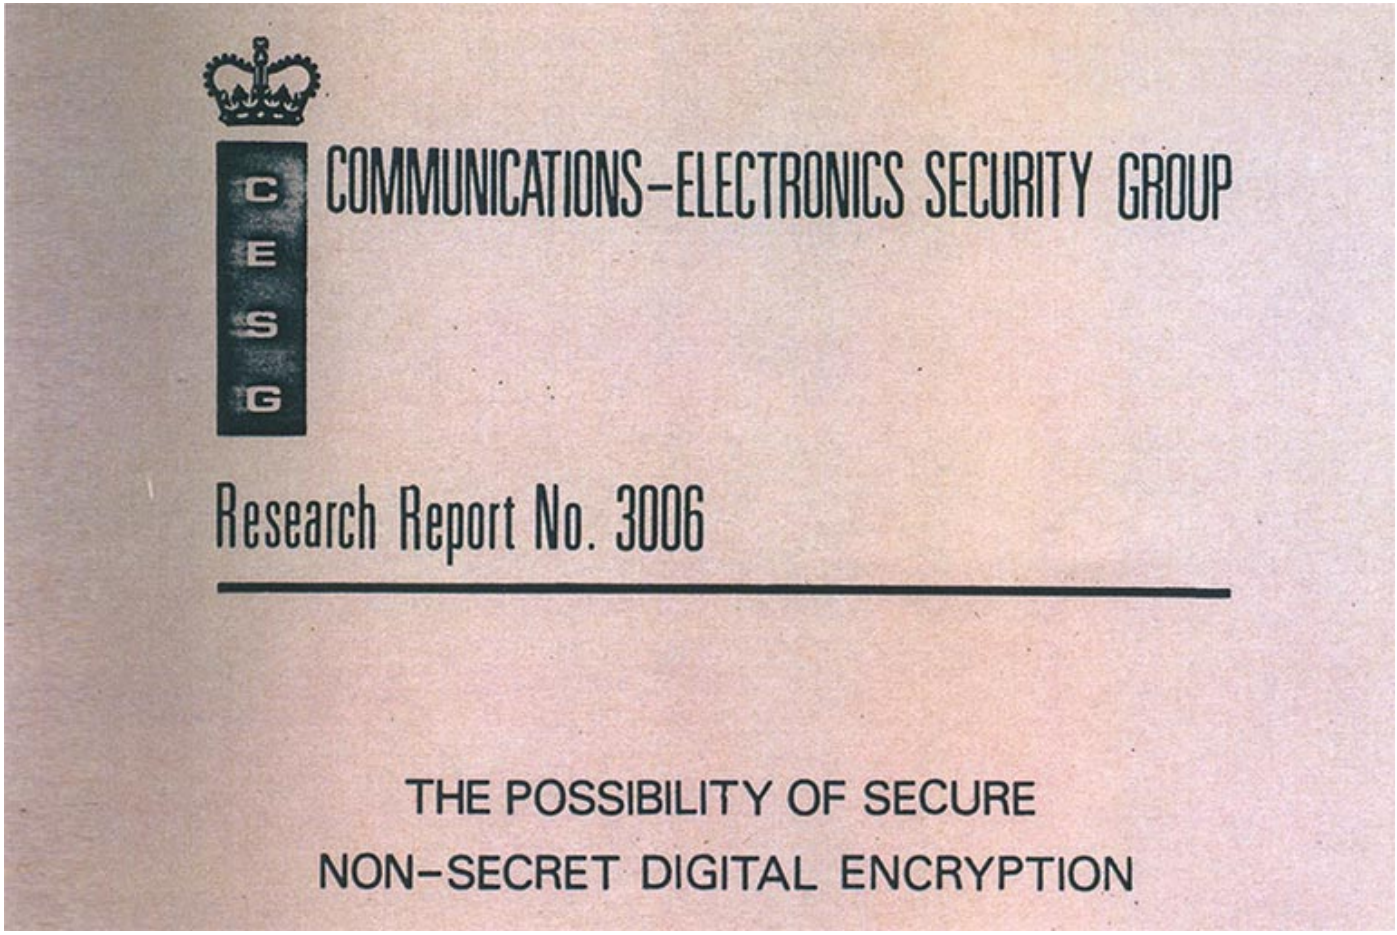

C E S G

COMMUNICATIONS-ELECTRONICS SECURITY GROUP

Research Report No. 3006

THE POSSIBILITY OF SECURE
NON-SECRET DIGITAL ENCRYPTION

Modern cryptography

Modern cryptography separates out these objectives

- Privacy/confidentiality — secret trapdoor function
- Authentication — proof of knowledge
- Integrity — one-way function (message digest)
- Non-repudiation — digital signature
- Sharing — **public key**

Hidden structure

Public-key *(asymmetric)* cryptography makes use of public data with hidden structure, the private or secret key.
It is believed that the recovery of the private key from the public data is "hard" and that this can be quantified.

- Cocks/RSA encryption — public: N, private: $N = pq$;
  - Believed that decrypting is equivalent (RP) to factoring
  - Believed to take time $\exp\left(1.923(\log N)^{\frac{1}{3}}(\log\log N)^{\frac{2}{3}}\right)$ to factor
- McEliece — public: code generator, private: cyclic or Goppa structure
  - Believed to be hard to decode without knowledge of structure
  - Believed to be hard to find stucture
- Lattice-based — public: lattice description, private: "good" basis
  - Believed to be hard to decrypt without knowledge of structure
  - Believed to be hard to find stucture

Other applications

- Key exchange and transport
- Authentication, proof of knowledge
- Privacy-enhanced computation
- Multi-party computation
- . . .

Quantum computing

Richard Feynman was one of the first to suggest the application of quantum states to computation.

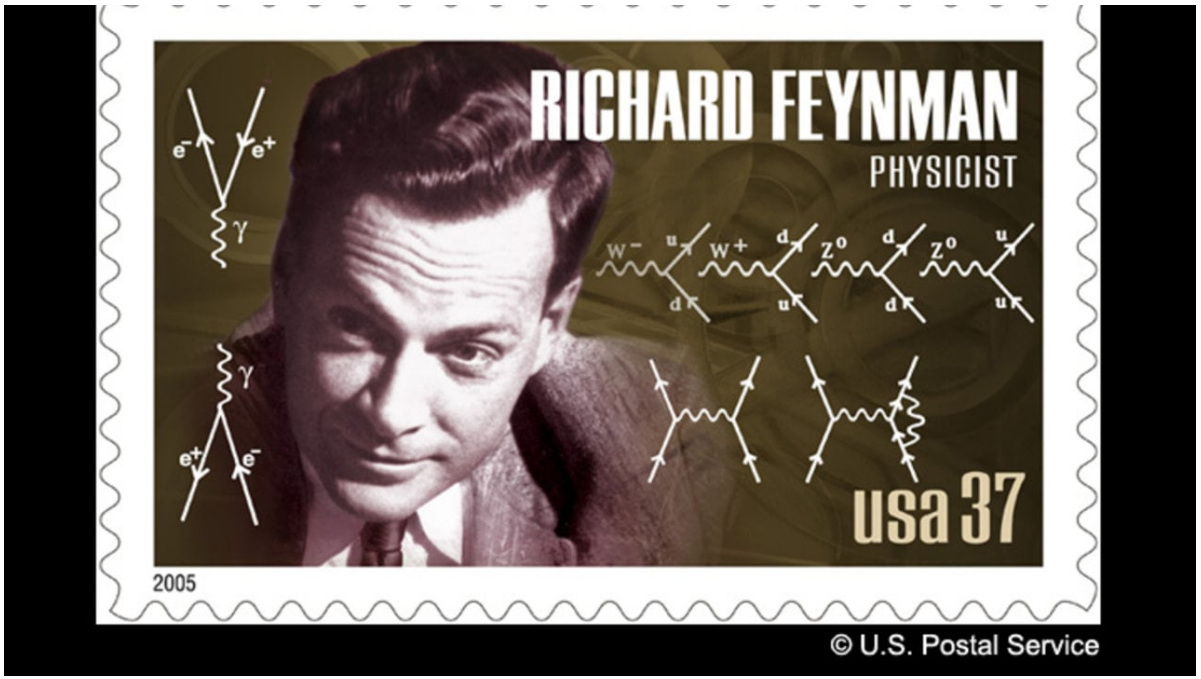

Quantum computation

A quantum computer is composed of quantum registers (entangled bits) and quantum circuits that process these bits.

A computation consists of preparing the states (feeding in the data) and observing the result (reading out the answer). A quantum algorithm is a way of arranging the computation so that the desired answer emerges with high probability.

Computability

The Analytical Engine has no pretensions whatever to originate anything. It can do whatever we know how to order it to perform.

Ada Lovelace

Comparison

Quantum computation is not a new paradigm in *computability* — it does not make it possible to compute anything that could not in principle also be computed on a Turing machine.

It is however a new paradigm in *complexity* — it makes it possible to solve certain kinds of problem in times significantly faster than classical computing.

Comparison

Quantum computation is not a new paradigm in *computability* — it does not make it possible to compute anything that could not in principle also be computed on a Turing machine.

It is however a new paradigm in *complexity* — it makes it possible to solve certain kinds of problem in times significantly faster than classical computing.

Quantum computation

Most current public key cryptosystems rely on the difficulty of two specific problems:

- Integer factorisation (Cocks/RSA,Rabin)
- Discrete logarithm in a prime field (Williamson/DH,DSA), other finite fields (XTR, CEILIDH) or an elliptic curve (ECC, Suite B, 25519)

Unfortunately, Shor's Algorithm for a quantum computer is effective against precisely these two classes of problems.
We also know that Grover's Algorithm is effective for speeding up general search problems.

Post-quantum

We therefore need

- To understand the effectiveness of probable models of quantum computation against existing cryptosystems, especially those implemented in standards or very widely used products
- To generate new proposals for cryptosystems which are not vulnerable to known attacks;
- To develop a robust understanding of the effective classical and quantum attacks on proposed systems and their parameters.

Quadratic forms and bilinear forms

We start with a real vector space V of finite dimension n.
A *bilinear form* is a map $b : V \times V \to \mathbb{R}$, linear in each argument separately.
It is *symmetric* if $b(x, y) = b(y, x)$ and *skew-symmetric (alternating)* if $b(x, y) = -b(y, x)$.
A *quadratic form* is a map $q : V \to \mathbb{R}$, scaling as $q(\lambda x) = \lambda^2 q(x)$ and such that the *polarisation*

$$b_q(x, y) = \tfrac{1}{2}(q(x + y) - q(x) - q(y))$$

is bilinear. Conversely if b is a symmetric bilinear form then $q(x) = b(x, x)$ is a quadratic form with $b_q = b$.
A quadratic form is *positive semidefinite* if $q(x) \geq 0$ and *definite* if $q(x) \neq 0$ for $x \neq 0$.

Bases and matrices

Let $(e) = (e_1, \ldots, e_n)$ be a basis for V with bilinear form b.
The *Gram matrix* for e is of b is the $n \times n$ matrix $G = b(e_i, e_j)$. If b is symmetric, so is G.
If P is a change of basis matrix from (e) to (f) then G transforms by $G \mapsto P^\top GP$: this is *congruence*.

Sylvester's Law of Inertia

Sylvester's Law of Inertia. Every real vector space with a quadratic form has a basis for which the Gram matrix is diagonal with only ± 1 or 0 on the diagonal: further, the number of $+1$, -1 and 0 is determined by the form. Equivalently: any real symmetric matrix is congruent to exactly one matrix with $+1$ then -1 then 0 on the diagonal.

The numbers of $(+, -, 0)$ on the diagonal is the *signature*.
A positive semidefinite form has signature $(r, 0, n - r)$: a positive definite form has signature $(n, 0, 0)$.

A *Euclidean space* is a space with a positive definite quadratic form: it has a basis which makes it a conventional Euclidean inner product space $(\mathbb{R}^n, \langle , \rangle)$.

Gram–Schmidt process

This is a procedure for generating an orthogonal or orthonormal basis for a real inner-product (Euclidean) space from a given (ordered) basis. At each step we project onto the span of the previously obtained vectors.
Let $v_1, \ldots, v_n$ be a basis for the inner-product space $(V, \langle , \rangle)$. Define a basis $v_1^\sharp, \ldots$ iteratively by

$$v_i^\sharp = v_i - \sum_{j=1}^{i-1} \frac{\left\langle v_i, v_j^\sharp \right\rangle}{\left\langle v_j^\sharp, v_j^\sharp \right\rangle} v_i$$

The vectors $(v_1^\sharp, \ldots, v_n^\sharp)$ form an orthogonal basis for $(V, \langle , \rangle)$ and can be normalised by dividing by the scale factors $\left\langle v_j^\sharp, v_j^\sharp \right\rangle$.

Gram–Schmidt and Cholesky

The change of basis matrix implied by the Gram–Schmidt process is triangular, and yields an orthonormal basis, one for which the Gram matrix is the identity.
This is the *Cholesky decomposition*: if G is symmetric definite then

$$G = U^\top \cdot U$$

where U is an upper triangular matrix containing the Gram–Schmidt coefficients.
Sylvester gave an intrinsic characterisation of positive definite symmetric matrix: all principal minors are positive.

Gram–Schmidt and QR

The Gram–Schmidt process also yields the *QR decomposition* for a square matrix A. Assume A is non-singular, and view A as a change of basis matrix on the Euclidean space. The new quadratic form has Gram matrix $G = A^\top A$ and this has a Cholesky decomposition

$$A^\top A = R^\top R$$

with R upper triangular. Write $Q = AR^{-1}$. Then $Q^\top Q = R^{-\top} A^\top A R^{-1} = I$: that is, Q is an orthogonal matrix.
We have

$$A = Q \cdot R$$

where Q is a matrix with orthonormal columns and R is upper triangular: indeed, the columns of Q are the orthonormal basis vectors resulting and the nonzero entries in R are the Gram–Schmidt coefficients.

Lattices

Geometric definition. A *lattice* is a discrete subgroup L of Euclidean space $(\mathbb{R}^n, |\cdot|)$ of maximal dimension / which spans the space over $\mathbb{R}$.

Algebraic definition. A *lattice* is a finite rank free $\mathbb{Z}$-module / $\mathbb{Z}^n$ equipped with a positive definite quadratic form q.

These definitions are equivalent.

Equivalent definitions

Geometric $\Rightarrow$ Algebraic.
A discrete subgroup of $\mathbb{R}^n$ cannot have rank greater than n. Hence L has rank n and Euclidean distance $|\cdot|^2$ induces a positive definite quadratic form on L.

Algebraic $\Rightarrow$ Geometric.
If L is free, it has a basis over $\mathbb{Z}$ which is then a basis for $L_{\mathbb{R}} = L \otimes \mathbb{R}$ over $\mathbb{R}$. The quadratic form extends to $L_{\mathbb{R}}$ and induces a symmetric bilinear form on $L_{\mathbb{R}}$ by polarisation, with a matrix which is diagonalisable over $\mathbb{R}$, giving an embedding of L into $(\mathbb{R}^n, |\cdot|)$.

The Gram matrix

In either view of a lattice it has a basis $e_1, \ldots, e_n$ and a symmetric bilinear form B with $B(x, x) = q(x) = |x|^2$.

The *Gram matrix* of B is the $n \times n$ matrix $B(e_i, e_j)$.

Equivalent bases

Two bases of L, say (e_i) and (f_j) are related by a change of basis matrix P which is integer and has an integer inverse, that is, $P \in \mathrm{SL}_n(\mathbb{Z})$.

The Gram matrices are related by $G_e = P^\top G_f P$.

The fundamental parallelepiped

Euclidean space is covered by translates of the *fundamental parallelepiped* for a basis (e_i)

$$D_e = \left\{ \sum_i x_i e_i : 0 \leq x_i < 1 \right\}$$

The *shape* of D depends on the choice of basis for L. The *volume* of D does not:

$$\operatorname{vol} D = |\det(e)|$$

which is $\sqrt{\det G}$ where G is the Gram matrix.

The ellipsoid

The ellipsoid attached to a basis (e_i) is defined as the subset of $\mathbb{R}^n$ for which

$$E_e = \left\{ x \in \mathbb{R}^n : x^\top G x \leq 1 \right\}$$

The *shape* of E depends on the choice of basis for L.

Choice of basis

Reduction means choosing a “good” basis, or a “good” shape of D preferably in an effective and efficient way.

A *canonical* or *standard* basis would be a specific choice of basis, one for each lattice, which would be unique and possible to compute effectively and efficiently.

Introduction and warning

Many of the ideas in lattices can be seen in the two-dimensional case. There is a very beautiful theory, rich in applications.

Beware! The high dimensional situation is often quite different.

Introduction and warning

Many of the ideas in lattices can be seen in the two-dimensional case. There is a very beautiful theory, rich in applications.

Beware! The high dimensional situation is often quite different.

Complex embedding

The two-dimensional Euclidean plane can be interpreted as the complex number Argand diagram.
A lattice described by two complex numbers z_1, z_2 can be described up to rotation and scaling (multiplication by z_1) by the single complex number $\tau = z_2/z_1$.
By choice of orientation we may assume that $\Im\tau > 0$. So lattice bases in two dimensions are encoded by points in the upper half-plane $\mathcal{H}$.
The basis transformations in $\mathrm{SL}_2(\mathbb{Z})$ act on $\mathcal{H}$ by Möbius transformations

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} : z \mapsto \frac{az+b}{cz+d}$$

The modular triangle

The action of $\mathrm{SL}_2(\mathbb{Z})$ on $\mathcal{H}$ has a *fundamental domain* (the "modular triangle") defined as

$$\Delta = \{z \in \mathcal{H} : |\Re z| \leq \tfrac{1}{2} \ , \ |z| \geq 1\}$$

The operations

$$T_{\pm} : z \mapsto z \pm 1$$

and

$$S : z \mapsto \frac{-1}{z}$$

correspond to elements of $\mathrm{SL}_2(\mathbb{Z})$ and successive applications of S, $T_{\pm}$ transform any element of $\mathcal{H}$ to an element of Δ.
Note that Δ is closed, and S, $T_{\pm}$ identity its edges, so that apart from the boundary cases, every element of $\mathcal{H}$ is transformed into just one element of $\mathcal{H}$.

Reduction

We define a lattice basis (z_1, z_2) to be *reduced* if the corresponding $\tau = z_2/z_1$ is in Δ.

- Every lattice has just one reduced basis (with certain exceptions)
- The reduced basis $(1, \tau)$ has the largest angle between basis vectors
- The reduced basis $(1, \tau)$ has the most nearly equal length basis vectors
- There is a fundamental domain for $\mathrm{SL}_2(\mathbb{Z})$
- There is an efficient algorithm for computing the reduced basis

Quadratic forms

The positive definite quadratic forms q on L can be expressed as $ax^2 + bxy + cy^2$ with a, b, c real, $a > 0$ and $b^2 - 4ac < 0$.
Reduction gives a form with $a \leq c$ and $|b| \leq a$. This can be taken as a definition of reduction.
The root of the associated quadratic lies in the fundamental domain.

Higher dimensions

Minkowski reduction

Let L be a lattice with quadratic form q. Minkowski reduction produces a basis (e) for L by taking e_1 to be the non-zero element with smallest q-value, and then extending by taking e_{j+1} to be the vector with smallest q-value that extends $e_1, \ldots, e_j$ as a lattice basis.

A Minkowski-reduced basis may be characterised by the property that if the coordinates $x_1, \ldots, x_n$ of an element $x = x_1 e_1 + \cdots + x_n e_n$ satisfy $\mathrm{hcf}\{x_k, \ldots, x_n\} = 1$ then $f(x) \geq f(e_k)$.

In general a lattice has a unique Minkowski-reduced basis up to sign changes; Minkowski reduction defines a fundamental domain for $\mathrm{SL}_n(\mathbb{Z})$, and this is defined by finitely many algebraic constraints.

Gram–Schmidt and reduction

NB: Gram–Schmidt is not a lattice procedure.

At each step we project onto the span of the previously obtained vectors. Let $v_1, \ldots, v_n$ be a basis for the inner-product space $(V, \langle , \rangle)$. Define a basis $v_1^\sharp, \ldots$ iteratively by

$$v_i^\sharp = v_i - \sum_{j=1}^{i-1} \mu_{ij} v_i$$

where

$$\mu_{ij} = \frac{\left\langle v_i, v_j^\sharp \right\rangle}{\left\langle v_j^\sharp, v_j^\sharp \right\rangle}$$

We want to find a basis for L for which —

- The diagonal terms r_{ii} are balanced, decrease as slowly as possible
- The off-diagonal terms are *size-reduced*, $|\mu_{ij}| \leq \frac{1}{2}$

Korkin–Zolotarev reduction

Let L be a lattice in $\mathbb{R}^n$. Let S_j be $\langle e_1, \ldots, e_j \rangle$ and P_j the orthogonal projection of L onto $S_j^\perp$. Choose e_{j+1} to have the shortest projection to P_j. The basis is *Korkin–Zolotarev reduced* if in addition it is size-reduced, $|\mu_{ij}| \leq \frac{1}{2}$.
In general a lattice has a unique Minkowski-reduced basis up to sign changes; similarly for KZ-reduced bases.
These definitions of reduction are related to the successive minima described above. If e is KZ-reduced then

$$\frac{4}{i+3}\lambda_i \leq |e_i| \leq \frac{i+3}{4}\lambda_i$$

Application to algebraic number theory

An algebraic number field K is a finite degree extension of the rational field $\mathbb{Q}$.
The *trace form* $\langle x, y \rangle \mapsto \mathrm{tr}_{K/\mathbb{Q}}(xy)$ is a symmetric bilinear form on K.
The *additive embedding* $a : K \hookrightarrow \mathbb{R}^r \oplus \mathbb{C}^s \approx \mathbb{R}^n$ where $n = r + 2s$.
The *multiplicative* or *logarithmic embedding* $\ell : K^* \hookrightarrow \mathbb{R}^r \oplus \mathbb{C}^s$.
The ring of integers $\mathcal{O}$ is a $\mathbb{Z}$-lattice with the trace form as a symmetric bilinear form: the additive embedding a is a realisation of this, and the covolume of $a(\mathcal{O})$ is the *discriminant* of $\mathcal{O}$.
The ideals of $\mathcal{O}$ inherit a lattice structure. An *ideal lattice* is any lattice isometric to an ideal in a number ring.
A principal ideal $(\alpha) = \alpha\mathcal{O}$. In general not all ideals are principal.
The norm of an ideal $\mathfrak{a}$ is the index $[\mathcal{O} : \mathfrak{a}]$. The covolume $\mathrm{vol}\,\mathfrak{a} = [\mathcal{O} : \mathfrak{a}]^2 \cdot \mathrm{vol}\,\mathcal{O}$.

Sketch of the finiteness of the class number

Two ideals $\mathfrak{a}, \mathfrak{b}$ of $\mathcal{O}$ are equivalent if there are elements $\alpha, \beta \in \mathcal{O}$ such that $(\alpha)\mathfrak{a} = (\beta)\mathfrak{b}$.
We claim that there are finitely many equivalence classes of ideals: indeed each class contains an ideal of norm
The principal step is that any ideal $\mathfrak{a}$ contains an element of norm at most

$$\frac{n!}{n^n}\left(\frac{4}{\pi}\right)^s N(\mathfrak{a})|\Delta|^{1/2}$$

Incidentally, this implies that $|\Delta| > 1$ for all number fields other than $\mathbf{Q}$.

Convex bodies

A region B in $\mathbb{R}^n$ is *convex* if $x, y \in B$ implies $tx + (1-t)y \in B$ for all $t \in [0,1]$, that is, B contains the entre line segment between any two of its points.
The closure of a convex body is convex, and it is often covenient to assume that convex bodies are closed.

Theorems of Blichfeld and Minkowski

Theorem. Let B be a convex body of volume V in $\mathbb{R}^n$ and L a lattice of covolume $< V$. Then B contains distinct points x, y with $x - y \in L$.
Idea of proof. Fix a fundamental parallelepiped Π for L and tile $\mathbb{R}^n$ with its translates. This tiling divides B into regions $B_n = B \cap (n + \Pi)$ for $n \in L$. The translates $B_n - n$ have total volume V greater than that of Π so cannot fit without overlap.

Theorem. Let B be a convex body in $\mathbb{R}^n$ of volume V which is symmetric about 0. Let L be a lattice of covolume $< 2^{-n}V$. Then L contains a point of B other than 0.

Successive minima

Let L be embedded in $\mathbb{R}^n$ and let $B_\lambda = \lambda B_1$ denote the ball of radius λ. The *successive minima* of L are the numbers

$$\lambda_i = \inf\{\lambda \in \mathbb{R} : L \cap B_\lambda \text{ has } i \text{ independent points}\}$$

so that λ_1 is the length of the shortest vector in L. Clearly $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$.

Minkowski's second theorem

Let L be a lattice of covolume D in $\mathbb{R}^n$. Then

$$\lambda_1 \leq \sqrt{n} D^{1/n}$$

This follows from the trivial estimate that vol $B_1 > (2/\sqrt{n})^n$ since B_1 contains a cube of side $2/sqrtn$.
Indeed

$$\lambda_1 \cdots \lambda_n \leq n^{n/2} D$$

Any good estimate for the volume of the unit ball, vol B_1 will do here.

Successive minima

In the opposite direction, we have

$$\lambda_1 \geq \left\langle x_1^\sharp, x_1^\sharp \right\rangle$$

where the $x^\sharp$ are the (unnormalised) Gram–Schmidt vectors.

Lattice problems

- **SVP**. Shortest vector problem.
  - Find a vector of length λ_1 in L.
- **CVP**. Closest vector problem.
  - Given a point of $\mathbb{R}^n$, find the closest element of L.

Problem and solution variants

- *Decision*. Decide whether or not the required answer exists.
- *Promise*. Find the required answer given that it exists.
- *Optimise*. Find the "merit" of the desired answer.
- *Approximate*. Find an approximation to the desired answer with "merit" within a factor of $1 - \epsilon$
- *Randomised*. Find the desired answer with a probability at least $1 - \epsilon$.

Random lattices

We need to understand what we mean by a “typical” or “random” lattice.
A basis for a lattice in $\mathbb{R}^n$ is an invertible matrix. Scaling we may take an element of $\mathrm{SL}_n(\mathbb{R})$. But different bases generate the same lattice, so we consider $\mathrm{SL}_n(\mathbb{Z})\backslash \mathrm{SL}_n(\mathbb{R})$.
There is a natural Haar measure on $\mathrm{SL}_n(\mathbb{R})$ and hence on this quotient giving it finite volume.
For example, there is a probabilistic analogue of Minkowski’s theorem: of C is convex, the probability that a random lattice of covolume 1 intersects C only in 0 is $\ll \frac{1}{\mathrm{vol}(C)}$.

Ajtai random lattices

Note. There is an other definition of random lattice.

Consider lattices which are sublattices of $\mathbb{Z}^n \hookrightarrow \mathbb{R}^n$.

An *Ajtai random lattice* is a q-ary lattice (contains $q\mathbb{Z}^n$) defined by random linear equations (parity checks) modulo q.

Ajtai showed that finding very short vectors in a random instance of such lattices is as hard as the problem for lattices in generality.

LLL reduction

The Lenstra–Lenstra–Lovasz basis reduction method uses a definition of reduction and gives an algorithm to achieve it which is provably efficient. Firstly we define an *LLL-reduced* lattice basis e_i with parameter m. Let the $e^\sharp$ be the Gram-Schmidt vectors corresponding to the e and write

$$\mu_{i,j} = \frac{\left\langle e_i, e_j^\sharp \right\rangle}{\left\langle e_j^\sharp, e_j^\sharp \right\rangle}$$

for $i > j$. We require

- The $\mu_{i,j} < \frac{1}{2}$;
- $\|e_{i+1}^\sharp\|^2 \geq (m - \frac{1}{4})\|e_i^\sharp\|^2$.

The LLL algorithm transforms any lattice basis into an LLL-reduced basis. The paramter m lies between $\frac{1}{4}$ and 1. A common choice is $m = \frac{3}{4}$.

The LLL algorithm

The algorithm applies alternate reduction and swap steps to a lattice basis (e) until it satisfies the LLL condition.
G. Perform Gram–Schmidt to compute the GS basis $(e^\sharp)$ and the $\mu_{i,j}$.
R. Reduction step. For $i = 2, \ldots n$ For $j = 1, \ldots, i-1$, replace e_i by

$$e_i - \left\lceil \frac{\left\langle e_i, e_j^\sharp \right\rangle}{\left\langle e_j^\sharp, e_j^\sharp \right\rangle} \right\rfloor e_j$$

where $\lceil \cdot \rfloor$ denotes nearest integer.
S. Swap step. Find the first i such that

$$m\|e_i^\sharp\|^2 > \|\mu_{i+1,i} e_i^\sharp + e_{i+1}^\sharp\|^2$$

If such an i is found, exchange e_i and e_{i+1} and return to **G**.
If no such i is found, then the basis is LLL-reduced, and stop.

Properties of LLL reduction

Let (e) be an LLL-reduced basis. Then

$$\|e_1\| \leq \left(\frac{2}{\sqrt{4m-1}}\right)^{n-1} \lambda_1$$

That is, the first basis vector is not far off being the shortest vector.

Properties of the LLL algorithm

Theorem. The LLL algorithm always terminates.
In particular, an LLL-reduced basis always exists.

Theorem. The number of iterations of the algorithm is bounded by a polynomial in the number of bits required to define the lattice.

Lattice-based cryptography

Some significant proposals

- NTRU
- SOLILOQUY
- LWE: learning with errors

NTRU

A cryptosystem working in the ring $R = \mathbb{F}_q[x]/\langle x^n - 1\rangle$ where q is a small prime. (Other polynomial families have been proposed.) There is an auxiliary small prime p.
The secret key consists of a pair of polynomials f, g with small coefficients $(0, \pm 1)$; the public key h is the quotient g/f in R. There is an auxiliary decryption key $f_p = f^{-1}$ in $\mathbb{F}_p[X]$.
Encrypt m: $z = r \cdot h + ((m + \pi) \bmod p) \in R$, where r is random and π is padding/check.
Decrypt z: $a = f \cdot z \in R$; $t = f_p \cdot a \bmod p$. Stripping out padding, t yields m.

Lattice attacks on NTRU

Lattice attack on f. Form a lattice of dimension $2n$ with generators rows of

$$L = \begin{pmatrix} \alpha I & H \\ 0 & qI \end{pmatrix}$$

where H has rows $h, Xh, X^2h, \ldots, X^{n-1}h$ all taken mod q.
This is the lattice of polynomial pairs (u, v) such that $gu \equiv \alpha fv \bmod q$ So the vector $(\alpha f, g)$ is in L. Note that there are "spurious" keys f' in this lattice will also decrypt (some) messages.

Lattice attack on m. A similar lattice containing $(\alpha m, r)$.

SOLILOQUY background

Let n be a prime and ζ a primitive n-th root of unity.

Let $K = \mathbb{Q}(\zeta)$ be the n-th cyclotomic field and $\mathcal{O} = \mathbb{Z}[\zeta]$ its ring of integers. Elements of $\mathcal{O}$ are monic polynomials of the form $\alpha = \sum_{i=1}^{n} a_i \zeta^i$.

For prime $p \equiv 1 \bmod n$ the principal ideal $p\mathcal{O}$ decomposes into a product of prime ideals $p\mathcal{O} = \prod_{i=1}^{n} \mathcal{P}_i$.

The prime ideals $\mathcal{P}_i$ are conjugates with norm p. They have a simple two-element representation $\mathcal{P} = p\mathcal{O} + (\zeta - c_i)\mathcal{O}$ where the c_i are n-th roots of unity modulo p.

We will be interested in the value $c \equiv 2^{(p-1)/n} \bmod p$ and its prime ideal $\mathcal{P} \equiv p\mathcal{O} + (\zeta - c)\mathcal{O}$.

SOLILOQUY

SOLIOQUY is a lattice-based cryptosystem developed by CESG around 2006 with a compact public key.
The private key is a small ring element

$$\alpha = \sum_{i=1}^{n} a_i \zeta^i.$$

such that the norm $p = N\alpha$ is prime and $c \not\equiv 1 \bmod p$.
The corresponding public key is p.
Encrypt m:

$$m = \sum_{i=0}^{n-1} e_i \zeta^i \mapsto z = \sum_{i=0}^{n-1} e_i c^i \bmod p$$

Decrypt z:

$$m = z - \lceil z\alpha^{-1} \rfloor \cdot \alpha$$

provided that m satisfied $\lceil m\alpha^{-1} \rfloor = 0$

SOLILOQUY surprise

There is an essentially quantum attack on SOLILOQUY.
There is a quantum algorithm to recover the private key.
CESG abandoned the development of SOLILOQUY in 2013 and do not recommend it for any practical applications.

SOLILOQUY surprise

There is an essentially quantum attack on SOLILOQUY.
There is a quantum algorithm to recover the private key.
CESG abandoned the development of SOLILOQUY in 2013 and do not recommend it for any practical applications.

Learning with errors

LWE is a public key-cryptosystem which recovers a secret vector m from a ciphertext which consists of a sequence of approximate modular linear equations $f \cdot m \sim r \bmod q$.
Secret key: matrix S; public key: matrix A and $P = AS + E$ where E is a random small "error" matrix.
Encrypt: to encrypt m in base t, choose random small garble a and send $z = (A^\top a, P^\top a + \lceil (q/t)m \rfloor)$.
Note that $z = (A^\top a, E^\top x + S^\top A^\top x)$ where $x = \lceil mq/t \rfloor)$.
Decrypt: given $z = (z_1, z_2)$, let $d = z_2 - S^\top z_1$: we have $d = E^\top + \lceil mq/t \rfloor$.
Since E is chosen from a small error distribution, m is recovered by rounding $(t/q)d$ to integer values.

Ring-LWE

As LWE but replace the matrices A and $P = AS + E$ by elements of a suitable ring $R = \mathbb{Z}[x]/\langle f \rangle$.
The lattice problems are now replaced by their ideal lattice analogues.

Applications of lattice reduction

- Modular knapsack
- Modular equations
- Factorisation

Modular knapsack

A proposed cryptosystem which is vulnerable to lattice reduction.
A *knapsack problem* is of the form: Given a set of positive integers w_i, find a subset summing to a given target T.
The *greedy algorithm* is: choose the largest unused weight $w_i \leq T$ and iterate on $T - w_i$. It works when the w_i increase rapidly (consider $w_i = 2^i$).
The Merkle–Hellman *modular knapsack* cryptosystem takes integers w_i as a public key. A message is encrypted as $z = \sum_{i \in m} w_i$. Decryption is solving the knapsack.
The hidden structure is an x and N so that the $x \cdot w_i$ form a rapidly increasing sequence taken modulo N.
This is vulnerable to a lattice reduction attack.

Modular knapsack lattice attack

Consider the rows

$$\begin{pmatrix} w_2 & w_3 & \dots & w_n & \left\lceil a_1^{1/n} \right\rfloor \\ w_1 & 0 & \dots & 0 & 0 \\ 0 & w_1 & \dots & 0 & 0 \\ & & & \ddots & 0 \\ 0 & 0 & \dots & -w_1 & 0C \end{pmatrix}$$

A short enough vector in this lattice, say $(k_2, \dots, k_n, k_1)$ yields a rational approximation say U/M to k_1/w_1 for which M and the $s_i = w_i U - k_1 M$ form a superincreasing modular weight system.

Modular equations

Consider a polynomial equation $f(x) \equiv 0 \bmod N$. Solving such problems in general implies factoring N: consider the equation $x^2 \equiv 1 \bmod N$.
The Coppersmith – Howgrave-Graham method solves such problems when it is known that there is a solution $x \bmod N$ for which x is small: $x \ll N^{1/d}$ where d is the degree of f.
Consider the lattice of all polynomial equations of bounded degree modulo N which have x as a root: such equations form a lattice with generators $f, xf, x^2 f, \ldots$.

Modular equations lattice

$$\begin{pmatrix} f_0 & f_1 & \dots & f_d & 0 & \dots & 0 \\ 0 & f_0 & \dots & f_{d-1} & f_d & \dots & 0 \\ & & & \vdots & & & \\ 0 & 0 & \dots & & & \dots & f_d \\ & & & & & & \\ N & 0 & \dots & & & \dots & 0 \\ 0 & N & \dots & & & \dots & 0 \\ & & & \ddots & & & \\ 0 & 0 & \dots & & & \dots & N \end{pmatrix}$$

A short vector in this lattice is a polynomial p which has x as a root modulo N but with $p(x)$ small enough that $p(x)$ is actually zero.
The solution x of $p(x) = 0$ can then be obtained by standard numerical techniques.

Partial factorisation

Given $N = pq$ and knowledge of sufficiently many high-order bits of p and of q, it is possible to recover the full values.
We use a two-variable polynomial of the form $(P + x)(Q + y) = N$ where P, Q are the known parts of the unknown factors p, q of N.

Any questions?

Any questions?

Any answers?

Any questions?

Any questions?

Any answers?