Diffie-Hellman Type Key Exchange before and with quantum computing

Gerhard Frey
University of Duisburg-Essen

RECSI 2022
19. 10.2022

We want

- secure messages, authenticate participants, sign documents

with simple protocols based on (hopefully) hard (mathematical) tasks.
For this you always need a key space from which the sender chooses a (secret) key which the receiver knows and uses for decryption. The method is thus *symmetric.*
We shall focus on the key exchange.

How to exchange keys?

formerly ...
it was simply assumed that the sender and receiver already shared a secret key. The group of participants was " overlookable and trustworthy" (one believed). If this was not yet the case, then the key transfer was done either

- in personal contact or
- again encrypted (e.g. via encrypted radio or coded machines), alternatively
- by a "trusted" central institution that distributed keys (how?).

nowadays ...

one has to deal with many and partly unknown participants and work in open networks like the Internet.

Just at the right time (1976) the work of

W. Diffie and M. E. Hellman: New directions in cryptography

was published. Each participant A in the system has **two** keys, a public key p_A and a secret key s_A.

Two participants A, B can compute a common key $s = s_{A,B} = s_{B,A}$ from (p_A, s_B) and (p_B, s_A) respectively, but it is very hard for all $C \neq A$ to get information about s_A.

We now specify the somewhat blurred expression "it's hard":

What is hard, what is fast? Let

$$f_N : A_N \to B_N$$

be mappings from sets A_N to sets B_N. The algorithmic evaluation $f_N(r)$ (at a random argument r) has (time/memory) complexity $F(N)$.

$$f_N = \mathcal{O}(g)$$

if there is $d \in \mathbb{R}_{>0}$ with $F(N) \leq d \cdot g(N)$. Define

$$L_N(\alpha, c) := exp(c \cdot \log(N)^\alpha \cdot \log\log(N)^{1-\alpha}).$$

Then the (probabilistic) asymptotic time (memory) complexity of the family f_N is called

- *polynomial* if $F_N = \mathcal{O}(L_N(0, c))$ ("*fast algorithm*")
- *exponential* if $F_N = \mathcal{O}(L_N(1, c))$ ("*hard algorithm*") and
- *subexponential* if there exists $0 < \alpha < 1$ with $F_N = \mathcal{O}(L_N(\alpha, c))$

in $\log(N)$. Subexponential complexity is a very interesting case between the two extremes.

Binary Complexity

We assume that the machines used are binary computers.
Algorithms are sequences of bit operations applied to bit strings. **Examples of algorithms with polynomial bit complexity:**

- the (extended) Euclidean algorithm and
- the exponentiation and inversion in finite groups of known order (expressed in costs for group operations).
- Hence: Exponentiation in $\mathbb{F}_q$ is polynomial in $\log(q)$, as is multiplication and the evaluation of polynomials of fixed degree.
- For given n and $x_0 \leq n$ prime to n and random $x \leq n$, the computation of a number k with $x \equiv k \cdot x_0 \mod n$ is polynomially in n.

Subexponential Bit Complexity

Factorization in $\mathbb{N}$
because of algorithms using the number field sieve with $\alpha = 1/3$.

Example for exponential bit complexity:

Definition

Let $(G,+)$ be a finite cyclic group with generator g_0 with its usual $\mathbb{Z}$-module-structure.

The discrete logarithm (DL) of $g \in G$ to g_0 is defined by

$$\log_{g_0}(g) = \min\{n \in \mathbb{N}; n \cdot g_0 = g\}.$$

Definition

A black-box group (G, m, i) is a set with two "oracles" $m(.,.)$ and $i(.)$ that provide to pairs $(g_1, g_2) \in G \times G$ the product $m(g_1, g_2)$ and to elements $g \in G$ the inverse $i(g)$ such that the group laws hold (including the existence of 0_G with $0_G = m(g, i(g))$ as neutral element).

Theorem (U. Maurer & S.Wolf (1998)

Let G be a black-box group of prime order q with generator g_0.

The discrete logarithm $\log_{g_0}$ in G has the (probabilistic) bit-complexity

$$P(q) = \Theta(q^{1/2}),$$

i.e. $q^{1/2} = \mathcal{O}(P(q))$ and $P(q) = \mathcal{O}(q^{1/2})$.

$P(q) = \mathcal{O}(q^{1/2})$ holds in all cyclic groups of order q (Pollard, Shanks).

Diffie-Hellman Key Exchange

Diffie-Hellman original

Observation: Let $q = p^d$ be a prime power, ℓ a prime dividing $q-1$, and $\zeta_\ell \in \mathbb{F}_q$ a ℓ-th unit root. Then

$$(\zeta_\ell^{k_1})^{k_2} = \zeta_\ell^{k_1 \cdot k_2} = (\zeta_\ell^{k_2})^{k_1}.$$

If one agrees (publicly) on an isomorphism

$$f_q : \mathbb{F}_q \to (\mathbb{Z}/p)[X]/(m(X)),$$

then P^1, P^2 can be **publicly** create a *"secret"*. P^i choosesl s_i (e.g. randomly in $0, ..., \ell - 1$) and sends

$$p_i := f_q(\zeta_\ell^{s_i}).$$

P^1 computes $s := p_2^{s_1}$ which is equal to $p_1^{s_2}$.
For security, it is necessary that the DL is hard in $(\mathbb{Z}/p)[X]/(m(X)^*[\ell]$.

Already **Gauß** was interested in the discrete logarithm in $\mathbb{F}_q^*$ and called it "index" (Disquisitiones Arithmeticae (1801)).
C.G. Jacobi calculated tables up to $p = 1000$ (1839).
Kraichik (1922) developed the *index-calculus algorithm*, which has been continuously reinvented and refined in cryptography from 1980 to the present (A. Joux, D. Robert and many others).
It has *subexponential* complexity (with relatively small constants), which become dramatically smaller when q is not a prime number. To be on the safe side, ℓ should be of size 4000 bits and q has to be "almost" prime.

Obvious generalization: DL systems to cyclic groups
$(C =< a_0 >, +)$ with $\ell \cdot a_0 = 0$.

$$f_C : C \hookrightarrow \mathbb{N}$$

such that $f_C(a_1 + a_2)$ is (quickly) computable from $f_C(a_1)$ and $f_C(a_2)$.
Thus $f_C(C)$ becomes a $\mathbb{Z}$-set with (fast) scalar multiplication.
Key Exchange:
P^1, P^2 randomly choose k_1, k_2 and compute and publish the public keys

$$p_i := f_C(k_i \cdot a_0))(i = 1, 2).$$

P^1 calculates the number $s := k_1 \cdot f_C(k_2 \cdot a_0)$ which equals $k_2 \cdot f_C(k_1 \cdot a_0)$.

The construction of cryptographically strong groups C (with numbering)
(in the world of binary computers)
is a very successful area of current public key cryptography.

Mathematical Task:
Construct (a family of) groups G with:

1. Elements in G can be presented in a compact way ($\mathcal{O}(\log(\mid G \mid))$ bits are needed)

2. Addition and inversion are given by algorithms, which can be implemented easily and efficiently and are very fast. (Complexity $\mathcal{O}(\log(\mid G \mid))$.)

3. The computation of the DL in G (for random elements) is (to the best of our knowledge) very difficult and therefore not feasible in practice (complexity ideally exponential in $\mid G \mid$), in particular $\mid G \mid$ must be divisible by a large prime number.

The horizon widens: Arithmetic Geometry

Idea: $\mathbb{F}_q^* = G_m(\mathbb{F}_q)$ is the Picard group of divisors of a projective curve (of arithmetic genus 1) over $\mathbb{F}_q$.

So: Find G as a subgroup of Pic_C^0 where C is a projective curve over a finite field.

One quickly sees that if one wants to achieve greater security than with classical DL, one must choose C as smooth projective curve with (geometric) genus $g_C \geq 1$.

Can we find families of curves such that for large subgroups of Pic_C^0 the conditions from above are satisfied?

The key for the arithmetic of curves C over fields K is the **Theorem of Riemann-Roch**. We formulate consequences of it.

Theorem

- C is over K birationally equivalent to a plane projective curve of degree $\mathcal{O}(g_C)$.
- In every divisor class of degree g_C there is a positive divisor.

It follows that, for $K = \mathbb{F}_q$, Pic_C^0 is a finite abelian group, and that the elements can be represented with a number of bits that is polynomial in g_C and $\log q$. The effectiveness of the addition is described by the very remarkable result of **F. Hess and C. Diem**.

Theorem

Let C be a curve of genus g_C over $\mathbb{F}_q$.

Then the addition in Pic_C^0 can be performed (probalistically) with a number of bit operations bounded (explicitly) polynomially in g_C (fixed for q) and $\log(q)$ (fixed for g_C).

Both the methods and the result are analogous to the results for computing ideal classes of number fields; the role of Minkowski's lattice point theorem is taken over by Riemann-Roch theorem.

Thus, subgroups of Pic_C^0 satisfy conditions 1) and 2) for moderate g_C if we succeed in finding curves whose Picard groups contain subgroups of large prime order. The strategy is to choose (with certain conditions) random curves C of fixed genus g over suitable fields $\mathbb{F}_q$ and then compute $\mid \mathrm{Pic}_C^0 \mid$.

The main tool is the theory of the local L-series of the Jacobian variety of C determined by the characteristic polynomial of the Frobenius endomorphism.

The fundamental result is due to **Hasse, Deuring** (for $g = 1$) and **Weil.**

A corollary is:

$$||\mathrm{Pic}_C^0| - q^{g_C}| = \mathcal{O}(q^{g_C - 1/2}).$$

This implies: If one could achieve "almost' ' generic security for the DL and aims at the security level 256 bit the size of $g_C \log(q)$ has to be ~ 512 bit.

For $g_C > 3$ and random curves this is currently out of range.

"Fortunately", security analysis shows that for cryptographic purposes this would be of no use either.

The attacks (algorithms for calculating the DL) can be found under the keywords: Tate pairing (F-Rück), Weil-Descent and above all again: Index-Calculus.

Results of Adleman and Huang showed already 1996 that Picard groups of curves of large genus become insecure (subexponential complexity).

Results of Gaudry et al. for small g describe exponential attacks, but are much faster than the generic attacks ($< q^{g/2}$).

This excludes curves of genus > 4 and hyperelliptic curves of genus 4.

But the most powerful result came from **C. Diem**.

Let C be a curve with a plane model C' of degree d (singularities are allowed). (For non-hyperelliptic curves $d = 2g_C - 2$, for hyperelliptic curves $d = 2g + 2$ is possible).

Theorem
Let $d \geq 4$ such that d or $d-1$ is a prime number.
Then the DL in Pic^0_C of curves birationally equivalent to plane curves of degree d can be calculated (except for log factors) in expected time $\mathcal{O}(q^{2-\frac{2}{d-2}})$

For genus 4 and non-hyperelliptic curve C we get $d = 6$ and thus the complexity of DL is lowered, up to log factors, to $\mathcal{O}(q^{3/2})$.

This is significantly smaller than

$$q^2 = |\mathrm{Pic}_C^0|^{1/2},$$

and therefore (together with a result of Gaudry) eliminates curves of genus 4 for DL systems.

For $g_C = 3$ and C is nonhyperelliptic, $d = 4$, and thus the complexity of the DL is restricted by

$$\mathcal{O}(q^{2-\frac{2}{4-2}}) = \mathcal{O}(q),$$

which is too small.

Since there are "many" hyperelliptic curves with $g_C = 3$ whose Jacobian is isogenous to that of a non-hyperelliptic curve, one should avoid curves of genus 3.

Counting Points

The only remaining candidates are curves of genus 1 and 2 over prime fields.

From now on: $g = 1$ and $C = E$.

(More difficult but possible is the construction of cryptographically strong curves of genus 2.)

R. Schoof proves, by using the theorem of Hasse-Weil, the Chinese Remainder Theorem and the evaluation of the operation of the Frobenius endomorphism ϕ_q on small order torsion points:
In principle for all abelian varieties A of fixed dimension, thus also for elliptic curves, point counting algorithms with complexity polynomial in $\log(q)$ for $|A(\mathbb{F}_q)|$ exist.

In practice this is too slow.

Counting is accelerated enormously by the **(SAE)** algorithm (Schoof, Atkin, Elkies).

The key idea is to determine the action of ϕ_q on the isogeny graph of A.

This works quite satisfyingly for $A = J_C$ for $g_C = 2$ and excellently for $g_E = 1$.

The background for this is the theory of moduli schemes of abelian varieties with level structure, for $g = 1$ these are moduli curves, which have a tremendously rich arithmetical and geometrical structure connected to Galois representations via modular forms (Fermat is around the corner).

(SAE) uses the cover $X_1(n) \to X_0(n)$ of degree $\sim n$.

The complexity of point counting is then (with a mild heuristic assumption) reduced to $\mathcal{O}(\log(q)^5)$, and SAE yields cryptographically suitable elliptic curves over prime fields $\mathbb{F}_p$ with $p \sim 512$ bits without relevant problems.

This is the satisfyingly effective and stable **present** state of the art for public key exchange (and for authentication and signature).

But in the **future** there are dark clouds (still quite distant (?)):

quantum computer,

and thus the security in the world of $Q - bits$ has to be analyzed.

Future

Q-bit-complexity

The possibility that quantum computing (i.e., the use of machines based on the manipulation of Q-bits (ternary objects with effects from quantum physics)) may be feasible in the foreseeable future, opens up completely new aspects for the discussion of crypto-primitives. **A consequence:** We have good reasons to assume that the bit-complexity for a family of crypto-primitives behaves exponentially, but it turns out that the Q-bit complexity is subexponential or even polynomial.

The key tool for this is **Quantum Fourier Transform (QFT)**, a linear transform on quantum bits and the quantum analog of the discrete Fourier transform.

With this one obtains (probalistic) estimates of the eigenvalues of unitary operators (cf. **D. Coppersmith (1994). "An approximate Fourier transform useful in quantum factoring". T. R. RC19642, IBM**)

New relationships between objects proposed as crypto-primitives of public-key methods are of particular interest (see **Oded Regev: New lattice based cryptographic constructions, J. ACM 51 (2004).**

For example, the problem of shortest vectors in lattices is related to two problems for dihedral groups we shall state next.

Hidden subgroups: Let $H < G$ and $f : G \to S$ be a black box function (e.g., an oracle) with $f(a) = f(b)$ exactly when $a \sim b$ in G/H.

Determine H!

Related is the

Hidden Shift Problem (HSP):

Let H be a G-set and $f_0 : H \to S$ be a black box function.

For $g_0 \in G$ define $f_{g_0} : H \to S$ by $f_{g_0}(h) := f_0(g_0 \cdot h)$.

Compute g_0.

- *Shor's algorithm* and its extension solve the hidden subgroup problem for abelian groups in polynomial time in the length of the output function.
- *Kuperberg's algorithm* solves HSP in subexponential time if G is abelian.

Consequence of Shor's algorithm:

Discrete logarithms in finite abelian groups and the factorization of natural numbers can be computed by quantum computers in polynomial time.

Key Exchange with Graphs

Abstract Setting

Let $\mathcal{C}, \mathcal{C}^1, \mathcal{C}^2$ be small categories, i.e. objects are sets and morphisms are mappings.

Let $\mathcal{G}, \mathcal{G}^1, \mathcal{G}^2$ be graphs with respective vertices $\mathcal{M} = \{A_i\}, \mathcal{M}^1 = \{A^1_j\}, \mathcal{M}^2 = \{A^2_k\}$ consisting of objects in $\mathcal{C}, \mathcal{C}^1, \mathcal{C}^2$ with (directed) edges $\mathcal{K}, \mathcal{K}^1, \mathcal{K}^2$ consisting of (finitely many) morphisms between vertices in the respective graphs.

Let $F^i : \mathcal{G}^i \to \mathcal{G}$ be mappings of graphs (e.g., inclusions).

Paths are denoted by $s(j,k)$ and $s^i(j,k)$, respectively.

Definition

The tuple $\mathcal{G}, \mathcal{G}^1, \mathcal{G}^2, F^1, F^2$ is of Diffie-Hellman type if there is there is a pair of algorithms $DH = (DH^1, DH^2)$, so that

1. to each path $s^1 := s^1_{0,j_1}$ in $\mathcal{G}^1$ and each endpoint $A^2_{j_2}$ of a path $s^2 := s^2(0, j_2)$, DH^1 computes a path $g_{j_2,k} := DH^1(s^1, A^2_{j_2})$ from $\mathcal{G}$, and
2. to each path $s^2 := s^2_{0,j_2}$ in $\mathcal{G}^2$ and $A^1_{j_1}$, DH^2 computes a path $g_{j_1,k} := DH^2(s^2, A^1_{j_1})$ from $\mathcal{G}$
   such that
3. the endpoint of $g_{j_1,k} \circ F^1(s^1)$ is equal to the endpoint of $g_{j_2,k} \circ F^2(s^2)$.

Diffie-Hellman Type Key Exchange

Two partners P^1 and P^2 want to have a shared secret.

They use $(\mathcal{G}, \mathcal{G}^1, \mathcal{G}^2, F^1, F^2)$ of Diffie-Hellman type with algorithms $DH = (DH^1, DH^2)$.

The private key space of P^i consists of paths $\{s^i_{0,j_i}\}$ o f $\mathcal{G}^i$.

P^i selects s^i.

The public key p_i of P^i is the endpoint $A^i_{j_i}$.

The other partner (and any eavesdropper) can use this information.

Key Exchange

P^1 and P^2 calculate the path

$$DH^1(s^1_{0,j_1}, A^2_{j_2})$$

respectively.

$$DH^2(s^2_{0,j_2}, A^1_{j_1})$$

in $\mathcal{G}$

and have the endpoint of these paths in $\mathcal{G}$ as a common secret. **Remarks**

To make the procedure feasible, objects and the evaluation of morphisms including composition must be fast (depending on the desired security level).

It is often useful to vary the paths, e.g. by splitting them into partial paths, in order to increase the effectiveness. This can also be included in the definition of the graph $\mathcal{G}^i$ (e.g. one can require that the degrees of the morphisms to edges and the lengths of the paths are restricted). This is then public information.

The resulting endpoint after application of DH^i is independent of the choice of the choosen paths.

The security of the key exchange depends on the complexity of the **Diffie-Hellman Computional Problem (CDHP):**

To known endpoints A_i in $\mathcal{M}^i$ of random paths $s^i = s(A_0, A_i)$ in $\mathcal{G}^i$ calculate $A_3 \in \mathcal{M}$ with $A_3 = DH^1(s^1, A_2)$ (without knowing s^1) or analogously $A_3 = DH^2(s^2, A_1)$

Global Diffie-Hellman Graphs

We shall now describe one special case where one can hope to organize a fast key exchange.

1) All vertices in $\mathcal{G}^i$ and $\mathcal{G}$ are contained in a set A.

2) For vertices A_1^i, A_2^i in $\mathcal{G}^i$, the edges are

$$K^i(1,2) \subset \{f^i_{|A_1^i}; f^i \in \mathrm{End}(A) \text{ with } f^i(A_1^i) \subset A_2^i\}.$$

3.)The key space of P^i is a subset $K^i \subset \mathrm{End}(A)$.

4.)**Commutativity (CO)**: K^1 commutates with K^2.

Let s^1, s^2 be paths in $\mathcal{G}^1$ and $\mathcal{G}^2$, respectively, with starting point A_0 and with endpoints A^1 and A^2, respectively, with associated endomorphisms f^1 and f^2 of A, so in particular $f^i(A_0) \subset A_i$.

Define

$DH^1(s^1, A_2)$ as a path in $\mathcal{G}$ with starting point $A_2 = f_2(A_0)$ and ending point $f_1(A_2)$.

$DH^2(s^2, A_1)$ as a path in $\mathcal{G}$ with starting point $A_1 = f_1(A_0)$ and end point $f_2(A_1)$.

Because of (CO), DH^1 and DH^2 satisfy the conditions for Diffie-Hellman graphs with morphisms $F^i =$ inclusion .

Example: G-sets

Let G be a (semi)group and A be a G-set on which G acts simply transitive (for simplicity).

For example, let A be a group and let the G-operation be attached to a principal homogeneous space in $H^1(G, A)$.

(F., STORK workshop, Bruges (Belgium), 26-27 November 2002, J.-M. Couveignes. Hard homogeneous spaces, 1997. Expose)

Let G_1, G_2 be subgroups of G with $G_i \subset Z(G_j)$ for $i \neq j$.

For $g \in G$ define $t_g \in End_{set}(A)$ by $a \mapsto t_g(a) := g \cdot a$.

The graphs $\mathcal{G}, \mathcal{G}^1, \mathcal{G}^2$ have as objects $\{a\} \subset A$ and $\mathrm{Mor}^i(\{a_k\}, \{a_l\}) = \{t_{g_i}, g_i \in G_i; t_{g_i}(a_k) = a_l\}$.

The key space for P^i is G_i.

The centralizer condition implies that the conditions for global Diffie-Hellman graphs are satisfied.

If G is abelian, it follows from **Kuperberg's algorithm** that (CDHP) has at most subexponential complexity.

The system is not necessarily unsafe(cf. DL in $\mathbb{F}_q^*$), possibly one can get a safe system by sufficiently large parameter choice.

Examples based on isogeny classes of elliptic curves

The $\bar{K}$-isomorphism class of E is determined by the absolute invariant $j_E \in K_0$.

Let E, E' be two elliptic curves over K_0.

An isogeny $\eta : E \to E'$ is a non constant homomorphism in $\mathrm{Hom}(E, E')$, the kernel of η is a finite group scheme in E whose order is equal to the degree of η.

Definition

E is ordinary if $\mathrm{End}(E)$ is commutative, else E is *supersingular*.

Theorem

Let E be an ordinary elliptic curve over a field K_0.

- Then $\mathrm{End}(E) = \mathbb{Z}$ (general case) or equal to an order $O_E \subset \mathbb{Q}(\sqrt{-d})$ (CM case).
- Assume that E has CM with O_E.
  Let $\mathcal{S}_E$ be the set of $\mathbb{C}$-isomorphism classes E' with $\mathrm{End}(E') = O_E$.

  Then $\mathcal{S}_E$ is a $\mathrm{Pic}(O_E)$ set.
- Explicitly: For $c \in \mathrm{Pic}(O_E)$, $\mathfrak{A} \in c$ and $\mathbb{C}/O_E = E_0$ we get:
  $c \cdot [E_0]$ is the class of $\mathbb{C}/\mathfrak{A}$.

LiftingTheorem of Deuring

Let E be an ordinary elliptic curve over $\mathbb{F}_q$.
Then, up to $\mathbb{C}$-isomorphisms, there is exactly one elliptic curve $\tilde{E}$ defined over a number field K such that

1. there is a prime divisor $\mathfrak{P}$ of K with

$$\tilde{E} \mod \mathfrak{P} \cong E,$$

   and

2. $\mathrm{End}(E) = \mathrm{End}(\tilde{E}) = O_E$.

Key Exchange à la Couveignes-Stolbunov

Let E_0 be an ordinary elliptic curve over $\mathbb{F}_q$ with $\mathrm{End}(E_0) = O$.

Let S_{E_0} be the set of isomorphism classes of elliptic curves over $\overline{\mathbb{F}_q}$ with endomorphism ring O.

According to Deuring, S_{E_0} is a PHS under $\mathrm{Pic}(O)$.

As described above, we use $(S_{E_0}, [E_0], G = \mathrm{Pic}(O)) = G_1 = G_2)$ for key exchange.

Explicit: Key Exchange

P^1 selects $c^1 \in \mathrm{Pic}(O)$ and publishes the j-invariant j^1 of $c^1 \cdot E_0$.

P^2 lifts the curve belonging to j^1 and applies the ideal class c^2 chosen by it.

The common key of P^1 and P^2 is the j-invariant of the reduction of the constructed elliptic curve.

This is feasible since one can find enough isogenies which are composites of isogenies of small degree (smoothness) and then apply formulas of Vélu.

The **security** depends on the isogeny problem:

Find an isogeny between two isogenous elliptic curves (with known ring of endomorpisms).

(**Question:** Is this equivalent to the HSP?)

The isogeny problem has bit-complexity $\mathcal{O}(q^{1/4+o(1)} \log^2(q) \log(q))$. (Kohel, Galbraith, Hess, Smart et al.) and subexponential Q-bit complexity (Kuperberg, explicitly see Childs-Jao-Soukharev).

Supersingular Elliptic Curves

The disadvantage (besides the subexponential complexity) of the above system is its slowness (in spite of all the tricks), which is due to the structure of the class group of O.

The idea of Castryck-Lange-Martindale-Panny-Renes and of DeFeo and Jao is to use E with $E \times \mathrm{Spec}(\overline{\mathbb{F}_q})$ supersingular.

Supersingular elliptic curves$E$ are isotrivial over $\mathbb{F}_{p^2}$, i.e. $j_E \in \mathbb{F}_{p^2}$.

Take E over $\mathbb{F}_{p^2}$. Then $|E(\mathbb{F}_{p^2})| = (p \pm 1)^2$, and the sign depends on the twist class of E.

$\mathrm{End}_{\overline{\mathbb{F}_p}}(E)$ is a maximal order in the quaternion algebra $\mathbb{Q}_p$ which is unramified outside ∞ and p.

If E is defined over $\mathbb{F}_p$, then $\mathrm{End}_{\mathbb{F}_p}(E)$ is an order O in $\mathbb{Q}(\sqrt{-p})$.

The system of C-L-M-P-R

Let E be an elliptic curve defined over $\mathbb{F}_p$ and supersingular over $\overline{\mathbb{F}_p}$.

The class group of the order $O_E \subset \mathbb{Q}(\sqrt{-p})$ operates on the $\overline{\mathbb{F}_p}$-isogeny classes of E_0.

It is controlled by p.

A clever choice of p allows us to assume that it is generated by powers of elements of order 2 and 3.

The key exchange can use add and double algorithms and is therefore much faster than the key replacement for ordinary elliptic curves. Of course, the security of the above system is again only subexponential (at best).

One hoped to improve this in the next example, which is, as we present it here, not a global Diffie-Hellman graph.

The De Feo-Jao-Plût Key Exchange System

Find and take $p = r^a \cdot s^b \cdot f - 1$ with $p \cong 1 \mod 4$ where r, s are small distict primes, e.g. equal to 2 and 3.

$E_0 : Y^2Z = X^3 + XZ^2$ is a supersingular elliptic curve over $\mathbb{F}_{p^2}$ with $|E_0(\mathbb{F}_{p^2})| = (r^a \cdot s^b \cdot f)^2$.

Fix a level r^a structure (P_1^0, P_2^0) and a level s^b structure (Q_1^0, Q_2^0) of E_0. We define a DH graph structure.

The vertices of the graph $\mathcal{G}$ are $\overline{\mathbb{F}_{p^2}}$-isomorphism classes of supersingular elliptic curves E defined over $\mathbb{F}_{p^2}$ and isogenous to E_0.

The edges of $\mathcal{G}$ are cyclic $\mathbb{F}_{p^2}$- isogenies whose degree divides $r^a \cdot s^b$.

The vertices of the graph $\mathcal{G}^1$ are $\overline{\mathbb{F}_{p^2}}$-isomorphism classes of supersingular elliptic curves E defined over $\mathbb{F}_{p^2}$ *together* with a level s^b structure (Q_1, Q_2). The edges of $\mathcal{G}^1$ are isogenies $\varphi : E \to E'$ with $\deg(\varphi) | r^a$.

The endpoints of the edges are $(E, Q_1, Q_2), (E', \varphi(Q_1), \varphi(Q_2))$.

The graph $\mathcal{G}^2$ is obtained by reversing the roles of r, s and (P_1, P_2), (Q_1, Q_2)

The mappings of $\mathcal{G}^i$ to $\mathcal{G}$ are given by forgetting the level structures. **Key Exchange:**

The partner P^1 chooses in the key space $(\mathbb{Z}/r^a)^2$ the pair (n_1, n_2) and the isogeny

$\eta : E_0 \to E_0/ < n_1P_1^0 + n_2P_2^0 >=: E_1$.

P^1 sends $(E_1, \eta(Q_1^0), \eta(Q_2^0))$.

P^2 selects $(m_1, m_2) \in (\mathbb{Z}/s^b)^2$ and the isogeny $\psi : E_0 \to E_0/ < m_1Q_1^0 + m_2Q_2^0 >=: E_2$.

P^2 sends $(E_2, \psi(P_1^0), \psi(P_2^0))$. **Definition of** DH^i

P^1 computes

$E_3 := E_2/ < n_1\psi(P_1^0) + n_2\psi(P_2^0) >$

and thus obtains a path in $\mathcal{G}$ from E_0 to E_3.

P^2 executes the analoguous computations.

A straightforward check shows: P^1 and P^2 get the same isomorphy class E_3 and can use the j-invariant of E_3 as a key.

A first glance concerning the **security of the key** leads again to the isogeny problem between two curves isogenous to E_0.

For this problem the state of the art is: The best known algorithms have an exponential complexity $p^{1/4}$ (bit computer) or $p^{1/6}$ (quantum computer), and one could be satisfied (NIST round 4).

But to corrupt the key of P^1 it would be enough to determine the kernel of η.

From the beginning it was discussed whether the additional level structures weaken the system.

The problem is: How can one use the information given about the action of η on points of order m of E_0? In fact, during the Conference ANTS XV (2022) two very effective attacks were published.

The first one was published by **Wout Castryck** and **Thomas Decru**[1] and resulted in a probabilistic subexponential algorithm with very convincing numerical examples.

This work uses the special properties of supersingular curves and their ring of endomorphisms.

[1] L. Maino and Ch. Martindale are pursuing related ideas.

Soon after this **Damien Robert** showed:

Theorem (Robert)

Assume that $n < m$ are relatively prime numbers and ℓ the largest prime number dividing m.
Let E_0 and E be two known elliptic curves over a finite field $\mathbb{F}_q$ with an (unknown) cyclic isogeny $\eta : E_0 \to E$ of degree n.

Let $(Q_1, Q_2) \in E_0[m]^2$ be a fixed level-m-structure.
Assume that $(\eta(Q_1), \eta(Q_2)) \in E(\mathbb{F}_q) \times E(\mathbb{F}_q)$, the image of this level-m-structure under η, is known.
Then there is a polynomial time algorithm in $\log(q)$ and ℓ, which calculates the kernel of η.

Remark
No properties of supersingular curves or of quantum computers are used in the proof of Robert's theorem.

Mathematical Background: Mumford's Descent

Let A, B be two principally polarized abelian varieties (e.g. two elliptic curves) over a field K, let n be prime to $\mathrm{Char}(K)$ and let $\eta : A[n] \to B[n]$ be a G_K-isomorphism which is anti-isometric with respect to the Weil pairing on $A \times B$ related to the product polarization.

This yields that $\{(P, \eta(P)); P \in A[n]\} =: \Delta_n)$ is maximally isotropic in $(A \times B)[n]$.

The result of Mumford states that $(A \times B)/\Delta_n$ is a principally polarized abelian variety with respect to the quotient polarization.

Example

Take $A = E, B = E'$ elliptic curves.

Then $(E \times E')/\Delta_n$ is a principally polarized abelian surface with respect to the quotient polarization, i.e. the Jacobian of a curve C_n of genus 2 which is "usually" irreducible.

Remark

This construction was the beginning of a long series of papers of F-Kani and Kani beginning 1991 and still not ending and describing the arithmetic of C_n and corresponding Hurwitz spaces.

The aim was to get information about E by varying E'.

For instance, the ABC-conjecure can be formulated as conjecture about properties of families of curves C on diagonal surfaces.

For this, it was important to determine the instances for which the curve C_n becomes reducible, i.e. isomorphic to two elliptic curves intersecting in one point.

E. Kani succeeded to give an explicit characterization of these special cases.

And exactly this characterization is the core of the two attacks I shall sketch now.

An (easly seen) necessary condition for reducibilityof C_n is that there is an isogeny $\eta : A \to B$.

In his paper in **Crelle 1997** (see also Coll. math. 2016) E. Kani defines special configurations of subgroups of the kernel of η called **Diamonds** and proves that these configurations correspond to uniquely determined anti-isometries ι of degree n (intrinsically given by the diamond) for which $(A \times \iota(A)/\{(P, \iota(P); P \in A[n]\}$ is **reducible**.

The attacks of Castryck-Decru and of Robert create, in an ingenious way, a constellation between abelian varieties (elliptic curves for the Castryck-Decru attack and an eight-dimensional (four-dimensional) product of elliptic curves for the Robert attack) satisfying the conditions of the Diamond theorem of Kani, and exploit this exceptional situation.

The Castryck-Decru Attack (**Wouter Castryck and Thomas Decru: An efficient key recovery attack on SIDH**, COSIC, KU Leuven.

Let $p = 2^a 3^b f - 1$ with $c =: 2^a - 3^b > 0$.

Wanted: The kernel of a cyclic isogeny $\eta : E_0 \to E$ of degree 3^b with known action on $E_0[2^a]$.

The easy case: Assume that 2^a is not too big, so c can be factorized.

Moreover, it is assumed that c is smooth and all prime divisors of c are congruent to 1 mod 4.

We know both E and the level- 2^a structure $(P_1 = \eta(P_1^0), P_2 = \eta(P_2^0)$.

We can quickly compute (see the paper of Castryck-Decru) an arbitrary cyclic isogeny $\gamma : E_0 \to C$

of degree c together with an level 2^a-structure $(P_1^c = \gamma(P_1^0), P_2^c = \gamma(P_2^0))$.

Then there is a test such that probabilistically/heuristically only (very) few cyclic isogenies φ of degree 3^b pass, and among them is one for which $\varphi(P_1^0) = P_1$ and $\varphi(P_2^0) = P_2$.

The Test

Let x be the multiplicative inverse of $3^b \mod 2^a$ and

$$\psi = [-1] \circ \varphi \circ \gamma^{\cdot} C \to E.$$

Check:
If $\varphi(P_i^0) = P_i$ then $[x] \circ \psi_{|C[2^a]}$ is an anti-isometry of level 2^a-structures.

$$\ker(\psi) = H_1 \cdot H_2$$

with uniquely determined cyclic subgroups H_1 of order c and H_2 of order 3^b and $|H_1| + |H_2| = 2^a$.

These are exactly the conditions of Kani for the triplet (ψ, H_1, H_2) being a diamond!

Hence dividing out the graph of $[x] \circ \psi_{|C[2^a]}$ in $C \times E$ yields a reducible abelian surface, and this can be tested by computing a chain of length a of $(2,2)$-isogenies using (one time) gluing and otherwise Richelot's formulas.
For the **general case** one shows that one can proceed inductively and split the search for *eta* in steps fitting into the "easy case".

This induction process is effective and yields subexponential complexity.

For details, we refer to the article by Castryck-Decru.

We remark that for the construction of $\gamma : E_0 \to C$ the knowledge of the ring of endomorphisms of E_0 and hence the supersingularity is important.

The Proof of the Theorem of D. Robert

These complications are avoided by the approach of Robert. In the most challenging part of his algorithm he can rely on results of himself and D. Lubicz concerning the evaluation at arbitrary points of isogenies of (moderately dimensioned) principally polarized abelian varieties with kernels of not too large degree which are maximally isotropic.

As result, he gets the polynomial complexity stated in the theorem above. The Matrix of Robert (cf. also work of E.Kani in Crelle 485 (1997), Collect. math. (2003), (2016)

Let $\eta : E_0 \to E$ is a cyclic isogeny over a field K of degree n, and $m = n + c$ with $c > 0$.

Using (Lagrange, quaternions) that c is the sum of four squares one finds a matrix $M \in M_4(\mathbb{Z})$ with $M^t \cdot M = c \cdot I_4$.

Define $\alpha_{1,1}$ by M operating on E_0^4, $\alpha_{2,2}$ by M^t operating on E^4, and $\alpha_{2,1} = \eta^4 \in Hom(E_0^4, E^4)$.

Let $A = (E_0^4) \times (E^4) = A_0 \times B$ with the product principal polarization.

Define $\alpha \in End(E_0^4 \times E^4)$ by the 2×2-"Matrix"

$\alpha = \begin{pmatrix} \alpha_{1,1} & \alpha_{2,1}^t \\ -\alpha_{2,1} & \alpha_{2,2}^t \end{pmatrix}$ Define "the dual" $r(\alpha) = \begin{pmatrix} \alpha_{1,1}^t & -\alpha_{2,1}^t \\ \alpha_{2,1} & \alpha_{2,2} \end{pmatrix}$ Check: $r(\alpha) \circ \alpha = [m] \cdot id_A$ and

$Ker(\alpha) = r(\alpha(E_0^4[m] \times \{0\}))$.

Conclusion Suppose that M and $\eta_{|E_0[m]}$ are known. Then $\ker(\alpha)$ is known.

Take $K = \mathbb{F}_q$. Check that $ker(\alpha)$ is isotropic. (This follows from Prop.1.1 in Kani:Crelle 1997).

Then, Lubicz and Robert can compute $\alpha(P)$ in polynomial time at every point $P \in A(\mathbb{F}_q)$.

Evaluate α at a base of $0 \times E^4[n]$ to get $\hat{\eta}(E[n])$.

Since η is cyclic, $Ker(\eta) = \hat{\eta}(E[n])$ and so **one can compute the kernel of η.**

Question: How does one has the idea to use the matrix α?

Hint (given by Ernst Kani): The matrix can be viewed as a special case of an isogeny factorization as defined in Crelle 1997 (p.100)

Let me end with the conclusion of Damien Robert in his paper: Breaking the SIDH in polynomial time, ia.cr/2022/1038 (slightly texified):

We have a new toolbox for recovering a cyclic isogeny $\eta : A \to B$ of degree n given by its action on the m-torsion of A as long as $m \geq n$ and m is sufficiently smooth.

This toolbox allows to break SIDH efficiently in all cases. Can it also be used to build new isogeny based cryptosystems?

THANK YOU