

Quantum Equivalence of the DLP and CDHP for Group Actions

Steven Galbraith¹, Lorenz Panny²,
Benjamin Smith³, and Frederik Vercauteren⁴

¹ Mathematics Department, University of Auckland, NZ
s.galbraith@auckland.ac.nz

² Department of Mathematics and Computer Science,
Technische Universiteit Eindhoven, Netherlands

Present affiliation: Institute of Information Science, Academia Sinica, Taipei, Taiwan
lorenz@yx7.cc

³ Inria and École Polytechnique, Institut Polytechnique de Paris, Palaiseau, France
smith@lix.polytechnique.fr

⁴ imec-COSIC, ESAT, KU Leuven, Belgium
frederik.vercauteren@kuleuven.be

Abstract. In this short note we give a polynomial-time quantum reduction from the vectorization problem (DLP) to the parallelization problem (CDHP) for efficiently computable group actions. Combined with the trivial reduction from parallelization to vectorization, we thus prove the quantum equivalence of these problems, which is the post-quantum counterpart to classic results of den Boer and Maurer in the classical Diffie–Hellman setting. In contrast to the classical setting, our reduction holds unconditionally and does not assume knowledge of suitable auxiliary algebraic groups. We discuss the implications of this reduction for isogeny-based cryptosystems including CSIDH.

Keywords: Quantum reduction, group action, discrete-logarithm problem, computational Diffie–Hellman problem.

1 Introduction

In their seminal 1976 paper [8], Diffie and Hellman conjectured that breaking their new key exchange protocol (in the sense of computing the shared secret from

* Author list in alphabetical order; see <https://www.ams.org/profession/leaders/culture/CultureStatement04.pdf>. This work was supported in part by the Ministry for Business, Innovation and Employment in New Zealand project UOAX1933, by the Commission of the European Communities through the Horizon 2020 program under project numbers 643161 (ECRYPT-NET) and 830892 (SPARTA), by the French Agence Nationale de la Recherche through ANR CIAO (ANR-19-CE48-0008), by the Research Council KU Leuven grant C14/18/067, and by CyberSecurity Research Flanders with reference number VR20192203. Date of this document: 2021-06-10.

the public keys) was as hard as computing discrete logarithms. This polynomial-time equivalence was later proven (assuming knowledge of suitable auxiliary algebraic groups of smooth order) for all groups by Maurer [10], based on earlier results of den Boer [7] covering certain special cases.

In this short paper, we prove an unconditional reduction between the analogous problems for group actions in the quantum setting. This result has important implications for the quantum security of the *CSIDH* key-exchange scheme [3].

Cryptographic group actions. In 1997, Couveignes introduced the notion of a *hard homogeneous space* [4], essentially a free and transitive finite abelian group action $*$: $G \times X \rightarrow X$ which is efficiently computable¹ while other computational problems are hard. In Couveignes’ terminology, these are *vectorization* and *parallelization*, named by analogy with the archetypical example of a homogeneous space: a vector space acting on affine space by translations (cf. Figure 1). The vectorization problem is: given x and $g*x$ in X , compute $g \in G$. The parallelization problem is: given x , $g*x$, and $h*x$ in X , compute $gh*x \in X$. The group-exponentiation analogues of these problems are the *discrete logarithm problem* (DLP) and *computational Diffie–Hellman problem* (CDHP).




Figure 1. The vectorization and parallelization problems.

For twenty years, there was little interest in the hard-homogeneous-spaces framework, since all known (conjectural) instantiations were either painfully slow in practice or already captured by the group-exponentiation point of view. However, interest in these one-way group actions has reemerged due to the current focus on post-quantum cryptography, where group-exponentiation Diffie–Hellman is broken in polynomial time by Shor’s algorithm [13], but group *actions* are not. In particular, *CSIDH* is a cryptographic group action that appears to be post-quantum secure and reasonably efficient in many scenarios [3].

DLP–CDHP reductions. Just like in the classical group-exponentiation setting, it is evident that parallelization reduces to vectorization: recover x from $x*g$, then apply x to $y*g$ to obtain $xy*g$. Traditionally, the other direction is much more subtle. The reduction essentially relies on the existence of auxiliary algebraic groups of smooth group order over \mathbb{F}_{q_i} , where the q_i are the prime divisors of the order of the group in which the DLP and CDHP are defined.

¹ See Section 2 for a precise definition.

The first result was given by den Boer [7], who showed the DLP and CDHP to be equivalent in \mathbb{F}_p^\times when p is a prime such that the Euler totient $\varphi(p-1)$ is smooth. The auxiliary groups are simply $\mathbb{F}_{q_i}^\times$ for each prime divisor $q_i \mid p-1$, and the smoothness assumption implies that the DLP in each $\mathbb{F}_{q_i}^\times$ is easy. Maurer [10] generalized this result to arbitrary cyclic groups G , assuming that for each large prime divisor q_i of $|G|$, there exists an efficiently constructible elliptic curve E/\mathbb{F}_{q_i} with smooth group order.

These reductions do not apply in the group-action setting on classical computers [15, §11]. However, we show that there exists a polynomial-time *quantum* reduction from the vectorization to the parallelization problem for group actions, without relying on any extra assumptions. This proves the polynomial-time equivalence of these problems in the quantum setting.

2 Efficient group actions

We now define what it means for a group action $G \times X \rightarrow X$ to be “efficiently computable”. Since our main motivation is CSIDH (where G is an ideal class group and X is a set of elliptic curves), we use the notation $\mathfrak{a}, \mathfrak{b}, \dots$ for elements of the group G , and denote by E an element of the set X .

Definition 1. *Let G be a finite abelian group and X a finite set. We abbreviate “polynomial in $\log(|G|+|X|)$ ” as “polynomial”. A group action $*$: $G \times X \rightarrow X$ is efficiently computable if all elements of G and X have (not necessarily unique) bit representations of polynomial length, a generating set of G of polynomial size is given, and the following tasks can be performed in polynomial time:*

1. Compute the composition $\mathfrak{a}\mathfrak{b} \in G$ of any $\mathfrak{a}, \mathfrak{b} \in G$.
2. Compute the action $\mathfrak{a}*E$ of any $\mathfrak{a} \in G$ on any $E \in X$.
3. Represent elements of X canonically as bit strings.

Vectorization *is*: Given $E \in X$ and $E' \in G * E$, compute any $\mathfrak{a} \in G$ with $E' = \mathfrak{a} * E$.
Parallelization *is*: Given $E \in X$ and $\mathfrak{a} * E, \mathfrak{b} * E \in G * E$, compute $\mathfrak{a}\mathfrak{b} * E \in X$.²

Remark. The notion of a “hard homogeneous space” as defined by Couveignes [4] additionally requires that $*$ is free and transitive, that uniform sampling from G is polynomial-time, and that vectorization and parallelization are hard for $*$. On the other hand, Task 3 is weakened to efficient membership and equality testing.

3 The reduction

Let π be an algorithm that solves the parallelization problem for an efficient group action $G \times X \rightarrow X$. In other words, π takes $\mathfrak{a} * E$ and $\mathfrak{b} * E$ and returns $\mathfrak{a}\mathfrak{b} * E$. We show that oracle access to a quantum circuit that computes π allows one to solve the vectorization problem for $*$: $G \times X \rightarrow X$ in polynomial time.

² The apparent ambiguity in the choice of \mathfrak{a} and \mathfrak{b} lies in the stabilizer subgroup of E , thus cancels out in the result $\mathfrak{a}\mathfrak{b} * E$.

Lemma 1. *Given an element $\mathbf{a} * E \in X$ and access to a parallelization oracle π , one can for any integer $n \geq 0$ compute $\mathbf{a}^n * E$ using $\Theta(\log n)$ queries to π .*

Proof. We perform double-and-add in the “implicit group” [15] of exponents, using the oracle $\pi: (\mathbf{a}^x * E, \mathbf{a}^y * E) \mapsto \mathbf{a}^{x+y} * E$ for addition and doubling. \square

Theorem 1. *Let $*$: $G \times X \rightarrow X$ be an efficiently computable group action. Given quantum access to a perfect parallelization oracle π , one can construct a quantum algorithm for the vectorization problem that runs in polynomial time.*

Proof. We are given an instance $(E, \mathbf{a} * E) \in X^2$ of the vectorization problem.³

From the public description of G , we get a polynomially-sized generating set $\mathbf{g}_1, \dots, \mathbf{g}_r$. For $\underline{x} \in \mathbb{Z}^r$, write $\mathbf{g}^{\underline{x}} = \prod_{i=1}^s \mathbf{g}_i^{x_i}$, and define the map

$$\begin{aligned} h: \mathbb{Z}^r &\longrightarrow X \\ \underline{x} &\longmapsto \mathbf{g}^{\underline{x}} * E. \end{aligned}$$

We apply Boneh and Lipton’s [2] or Kitaev’s [9] higher-dimensional generalisation of Shor’s algorithm [13] to compute the period lattice

$$K = \{ \underline{x} \in \mathbb{Z}^r : \mathbf{g}^{\underline{x}} * E = E \}$$

of the map h in polynomial time. Note that \mathbb{Z}^r / K is isomorphic to $G / \text{Stab}(E)$.

Now, define

$$\begin{aligned} f: \mathbb{Z}^r \times \mathbb{Z} &\longrightarrow X \\ (\underline{x}, y) &\longmapsto \mathbf{g}^{\underline{x}} * (\mathbf{a}^y * E). \end{aligned}$$

Observe that $\mathbf{a}^y * E$ can be computed using Lemma 1: Negative y may be replaced by a positive representative modulo $\det(K)$, which must be a multiple of the order of $\mathbf{a} \cdot \text{Stab}(E)$. Thus, using the efficient algorithm for the group action and the oracle access to π , one can construct a quantum circuit that computes f in polynomial time. The function f is a homomorphism to the implicit group on the orbit of E isomorphic to \mathbb{Z}^r / K , hence defines an instance of the hidden-subgroup problem with respect to its kernel, i.e., the lattice

$$L = \{ (\underline{x}, y) \in \mathbb{Z}^r \times \mathbb{Z} : \mathbf{g}^{\underline{x} + y\mathbf{v}} * E = E \},$$

where $\mathbf{v} \in \mathbb{Z}^r$ is any vector such that $\mathbf{g}^{\mathbf{v}} * E = \mathbf{a} * E$.⁴ This (abelian) hidden-subgroup problem can be solved in polynomial time again using Shor’s algorithm, making use of the efficient circuit to compute f constructed above. Finally, any vector in L of the form $(\underline{x}, 1)$ satisfies $\mathbf{g}^{-\underline{x}} * E = \mathbf{a} * E$, hence yields a solution to the vectorization problem. \square

Remark. If desired, the generating set $\mathbf{g}_1, \dots, \mathbf{g}_r$ can be replaced by a smaller generating set after computing K and before defining f . Moreover, if elements

³ The element \mathbf{a} is only defined up to $\text{Stab}(E)$, but this choice will cancel throughout.

⁴ Note that \mathbf{v} is only defined modulo K , but this does not matter since $L \supseteq K \times \{0\}$.

of G have unique representation, the computation of K can be replaced by a group-structure computation; the benefit is that this is independent of E , hence can be amortized across multiple vectorization instances.

Also note that the computation of K is only necessary to handle negative y when evaluating f ; hence, it seems this step could be omitted by using a variant of Shor’s algorithm that only queries f on the subset $\mathbb{Z}^r \times \mathbb{Z}_{\geq 0}$. The computation of K can also be skipped if the order of G is known a priori, or if the action of inverses can be computed in a different way: For example, in the CSIDH setting, when E is the starting curve chosen in [3], then $\mathfrak{r}^{-1} * E$ can be obtained as the quadratic twist of $\mathfrak{r} * E$.

3.1 Imperfect oracles

It is unclear how to perform the reduction above when π is only guaranteed to succeed with non-negligible probability α , meaning that the probability over all triples $(E, \mathfrak{a} * E, \mathfrak{b} * E) \in X^3$ that the oracle outputs $\mathfrak{a}\mathfrak{b} * E$ is at least α .

In the classical discrete-logarithm setting, it is straightforward to amplify the success probability of CDH oracles using a random self-reduction of problem instances [11, 14]: one computes lists of possible values of g^{ab} by blinding the inputs and unblinding the outputs, and uses majority vote to determine the correct result. Any exponentially small failure probability can be achieved using polynomially many queries [14, §5].

In the group-action setting, however, blinding does not work: The results cited above use a blinding map of the form $g^a \mapsto (g^a)^x g^y = g^{ax+y}$, which relies on the fact that we can multiply two public keys. But the best we can do for a mere group action is to translate the inputs by random elements, i.e., blind as $\mathfrak{a} * E \mapsto \mathfrak{r} * (\mathfrak{a} * E)$ with a random $\mathfrak{r} \in G$, which is insufficient: For example, if \mathcal{A} is a perfect CDH oracle, then the oracle \mathcal{B} that returns the output of \mathcal{A} either unmodified (with probability ϵ), or shifted by a fixed element $\mathfrak{z} \in G$, is entirely unaffected by blinding and hence cannot be amplified using this idea. Thus, we must unfortunately leave the case of imperfect oracles as an open problem.

4 Implications for CSIDH

Let E be an elliptic curve over \mathbb{F}_q with $\text{End}_{\mathbb{F}_q}(E) = \mathcal{O}$ being an order in an imaginary quadratic field. Any invertible \mathcal{O} -ideal \mathfrak{a} gives rise to an isogeny $\varphi_{\mathfrak{a}}: E \rightarrow E'$ with kernel $E[\mathfrak{a}] = \{P \in E(\overline{\mathbb{F}_q}) : \forall \psi \in \mathfrak{a}, \psi(P) = 0\}$. This leads to an action of $\text{cl}(\mathcal{O})$ on a set X of elliptic curves isogenous to E and with the same endomorphism ring as E . Precisely, $\mathfrak{a} * E := E' = E/E[\mathfrak{a}]$. This is the homogeneous space underlying CSIDH [3], the Couveignes and Rostovtsev–Stolbunov cryptosystems [4, 16, 6], and also the SeaSign signature scheme [5].

Public keys are instances $(E, \mathfrak{a} * E)$ of the vectorization problem in this homogeneous space. In CSIDH, the Diffie–Hellman secret shared between Alice and Bob, with public keys $(E, \mathfrak{a} * E)$ and $(E, \mathfrak{b} * E)$, is $\mathfrak{a}\mathfrak{b} * E$. Recovering the shared secret from the public keys is therefore solving a parallelization problem.

Unfortunately, CSIDH is not known to be an efficiently computable group action in general. The standard implementations of CSIDH [3] use secret keys of the form $\mathfrak{a} = \prod_i \mathfrak{l}_i^{e_i}$, where $\underline{e} = (e_1, \dots, e_n) \in \mathbb{Z}^n$ are short exponent vectors and the \mathfrak{l}_i are a fixed set of “small” ideals whose action is efficient. The action of \mathfrak{a} is then evaluated as repeated applications of the \mathfrak{l}_i and their inverses. However, with these implementations, it is no longer efficient to evaluate the action of a *composition* of such “nice” ideals: a sequence of k additions starting from short exponent vectors can result in an exponent vector of 1-norm exponential in k .

If one sets up a sequence of CSIDH instantiations for unbounded security levels, then there is no known polynomial-time method to sample uniformly from the groups G or compute the action in polynomial time. There are two reasons for this. First, one might need relatively large prime ideals \mathfrak{l}_i to generate the class group. Second, and more serious, given a randomly chosen ideal \mathfrak{a} it may be hard to find a short representation of an equivalent ideal of the form $\prod_i \mathfrak{l}_i^{e_i}$. Even when the class group structure is known, finding a short exponent vector \underline{e} requires solving a close(st)-vector problem for the relation lattice $\ker(\mathbb{Z}^n \rightarrow \text{cl}(\mathcal{O}))$. But asymptotically, polynomial-time lattice reduction algorithms cannot guarantee that the output will have norm small enough to ensure that the resulting group action is computable in polynomial time.

For the above reasons, Theorem 1 does not apply directly to the general case of CSIDH or related cryptographic systems. However, this does not mean the result has no practical meaning. For example, since the dimensions n used in CSIDH are rather small (e.g. the CSIDH-512 parameter set from [3] uses $n = 74$), an efficient lattice-reduction algorithm such as BKZ [12] with moderate block size suffices to obtain highly practical results (reducing a random relation lattice of dimension 74 using BKZ with block size 20 yields exponent vectors only 8 times longer than normal CSIDH-512 private keys). As another example, the CSI-FiSh [1] system has a known relation lattice and a relatively efficient group operation, so our theorem shows that the parallelization and vectorization problems are equivalent in a practical sense for this system. Similarly, we expect that in many reasonable cryptographic settings (possibly after some quantum and classical precomputation) our result will provide a meaningful equivalence of the parallelization and vectorization problems.

References

- [1] Ward Beullens, Thorsten Kleinjung, and Frederik Vercauteren. CSI-FiSh: Efficient isogeny based signatures through class group computations. In *ASIACRYPT (1)*, volume 11921 of *Lecture Notes in Computer Science*, pages 227–247. Springer, 2019. <https://ia.cr/2019/498>.
- [2] Dan Boneh and Richard J. Lipton. Quantum cryptanalysis of hidden linear functions (extended abstract). In *CRYPTO*, volume 963 of *Lecture Notes in Computer Science*, pages 424–437. Springer, 1995. <https://crypto.stanford.edu/~dabo/pubs/papers/quantum.pdf>.
- [3] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. CSIDH: An efficient post-quantum commutative group action. In *ASIA-*

- CRYPTO (3)*, volume 11274 of *Lecture Notes in Computer Science*, pages 395–427. Springer, 2018. <https://ia.cr/2018/383>.
- [4] Jean-Marc Couveignes. Hard homogeneous spaces. 2006. IACR Cryptology ePrint Archive 2006/291. <https://ia.cr/2006/291>.
- [5] Luca De Feo and Steven D. Galbraith. SeaSign: Compact isogeny signatures from class group actions. In *EUROCRYPT (3)*, volume 11478 of *Lecture Notes in Computer Science*, pages 759–789. Springer, 2019. <https://ia.cr/2018/824>.
- [6] Luca De Feo, Jean Kieffer, and Benjamin Smith. Towards practical key exchange from ordinary isogeny graphs. In *ASIACRYPT (3)*, volume 11274 of *Lecture Notes in Computer Science*, pages 365–394. Springer, 2018. <https://ia.cr/2018/485>.
- [7] Bert den Boer. Diffie–Hellman is as strong as discrete log for certain primes. In *CRYPTO*, volume 403 of *Lecture Notes in Computer Science*, pages 530–539. Springer, 1988.
- [8] Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Trans. Information Theory*, 22(6):644–654, 1976.
- [9] Alexei Y. Kitaev. Quantum measurements and the abelian stabilizer problem. *Electronic Colloquium on Computational Complexity (ECCC)*, 3(3), 1996. <https://eccc.hpi-web.de/eccc-reports/1996/TR96-003>.
- [10] Ueli M. Maurer. Towards the equivalence of breaking the Diffie–Hellman protocol and computing discrete logarithms. In *CRYPTO*, volume 839 of *Lecture Notes in Computer Science*, pages 271–281. Springer, 1994.
- [11] Ueli M. Maurer and Stefan Wolf. Diffie–Hellman oracles. In *CRYPTO*, volume 1109 of *Lecture Notes in Computer Science*, pages 268–282. Springer, 1996.
- [12] Claus-Peter Schnorr and M. Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Math. Program.*, 66:181–199, 1994.
- [13] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, 1997. <https://arxiv.org/abs/quant-ph/9508027>.
- [14] Victor Shoup. Lower bounds for discrete logarithms and related problems. In *EUROCRYPT*, volume 1233 of *Lecture Notes in Computer Science*, pages 256–266. Springer, 1997.
- [15] Benjamin Smith. Pre- and post-quantum Diffie–Hellman from groups, actions, and isogenies. In *WAIFI*, volume 11321 of *Lecture Notes in Computer Science*, pages 3–40. Springer, 2018. <https://ia.cr/2018/882>.
- [16] Anton Stolbunov. Constructing public-key cryptographic schemes based on class group action on a set of isogenous elliptic curves. *Adv. in Math. of Comm.*, 4(2):215–235, 2010.