J. Cryptol. (2012) 25: 528–555
DOI: 10.1007/s00145-011-9103-4

Journal of
CRYPTOLOGY

Computational Indistinguishability Between Quantum States and Its Cryptographic Application*

Akinori Kawachi

Department of Mathematical and Computing Sciences, Tokyo Institute of Technology, W8-25 Ookayama, Meguro-ku, Tokyo 152-8552, Japan
kawachi@is.titech.ac.jp

Takeshi Koshiba

Division of Mathematics, Electronics and Informatics, Graduate School of Science and Engineering, Saitama University, 255 Shimo-Okubo, Sakura-ku, Saitama 338-8570, Japan
koshiba@mail.saitama-u.ac.jp

Harumichi Nishimura

Department of Mathematics and Information Sciences, Graduate School of Science, Osaka Prefecture University, 1-1 Gakuen-cho, Naka-ku, Sakai 599-8531, Japan
hnishimura@mi.s.osakafu-u.ac.jp

Tomoyuki Yamakami

ERATO-SORST Quantum Computation and Information Project, Japan Science and Technology Agency, Hongo, Bunkyo-ku, Tokyo 113-0033, Japan
TomoyukiYamakami@gmail.com

Communicated by Stefan Wolf

Received 8 May 2006
Online publication 12 April 2011

Abstract. We introduce a computational problem of distinguishing between two specific quantum states as a new cryptographic problem to design a quantum cryptographic scheme that is "secure" against any polynomial-time quantum adversary. Our problem, $\mathrm{QSCD}_{\mathrm{ff}}$, is to distinguish between two types of random coset states with a hidden permutation over the symmetric group of finite degree. This naturally generalizes the commonly-used distinction problem between two probability distributions in computational cryptography. As our major contribution, we show that $\mathrm{QSCD}_{\mathrm{ff}}$ has three properties of cryptographic interest: (i) $\mathrm{QSCD}_{\mathrm{ff}}$ has a trapdoor; (ii) the average-case hardness of $\mathrm{QSCD}_{\mathrm{ff}}$ coincides with its worst-case hardness; and (iii) $\mathrm{QSCD}_{\mathrm{ff}}$ is computationally at least as hard as the graph automorphism problem in the worst case. These cryptographic properties enable us to construct a quantum public-key cryptosystem which is likely to withstand any chosen plaintext attack of a polynomial-time quan-

* The preliminary version [30] appeared in the *Proceedings of EUROCRYPT 2005*, *Lecture Notes in Computer Science*, Vol. 3494, pp. 268–284, Aarhus, Denmark, May 22–26, 2005.

tum adversary. We further discuss a generalization of $\mathrm{QSCD}_{\mathrm{ff}}$, called $\mathrm{QSCD}_{\mathrm{cyc}}$, and introduce a multi-bit encryption scheme that relies on similar cryptographic properties of $\mathrm{QSCD}_{\mathrm{cyc}}$.

Key words. Quantum cryptography, Computational indistinguishability, Trapdoor, Worst-case/average-case equivalence, Graph automorphism problem, Quantum public-key cryptosystem.

1. Introduction

In 1976, Diffie and Hellman [17] first used a computationally intractable problem to design a key exchange protocol. Computational cryptography has since become an important field of extensive study. A number of practical cryptographic systems (e.g., public-key cryptosystems (PKCs), bit commitment schemes (BCSs), pseudorandom generators, and digital signature schemes) have been proposed under popular intractability assumptions, such as the hardness of the integer factorization problem (IFP) and the discrete logarithm problem (DLP), for which no efficient classical algorithm has been found. Using the power of quantum computation, however, we can efficiently solve various number-theoretic problems, including IFP (and thus, the quadratic residuosity problem) [56], DLP (and also the Diffie–Hellman problem) [11,33,56], and the principal ideal problem [24] (see also [16,55]). This indicates that a *quantum adversary* (i.e., an adversary who operates a quantum computer) can easily break any cryptosystems whose security proofs rely on the computational hardness of those problems.

In order to deal with such a powerful quantum adversary, a new area of cryptography, the so-called *quantum cryptography*, has emerged in the past quarter century. In 1984, Bennett and Brassard [8] first proposed a *quantum key distribution scheme*, in which a party can securely send a secret key to another party through a quantum communication channel. Its unconditional security was later proven by Mayers [40] (and more sophisticated proofs were given by, e.g., Shor and Preskill [57] and Renner [51]). Against our early hope, quantum mechanics cannot make all cryptographic schemes information-theoretically secure since, for instance, as Mayers [39] and Lo and Chau [37] independently demonstrated, no quantum BCS can be both unconditionally concealing and binding. Therefore, "computational" approaches are still important and also viable in quantum cryptography. Along this line of study, a number of quantum cryptographic properties have been discussed from complexity-theoretic viewpoints [1,13–15, 18,48].

In fact, a quantum computer is capable of breaking the RSA cryptosystem and many other well-known classical cryptosystems. It is therefore imperative to discover computationally-hard problems from which we can construct a quantum cryptosystem that is secure against any polynomial-time quantum adversary. For instance, the subset sum (knapsack) problem and the shortest vector problem are used as bases of knapsack-based cryptosystems [29,48] as well as lattice-based cryptosystems [4,49,52]. Since we do not know whether these problems withstand any attack by quantum adversaries,

we need to continue searching for better intractable problems that can guard their associated quantum cryptosystems against any computationally-bounded quantum adversary.

This paper naturally generalizes a notion of the computational indistinguishability between two probability distributions [9,20,60] to that between two *quantum states*. In particular, we present a distinction problem, called $\mathrm{QSCD}_{\mathrm{ff}}$ (quantum state computational distinction with fully flipped permutations), between specific ensembles of quantum states. It turns out that $\mathrm{QSCD}_{\mathrm{ff}}$ enjoys useful cryptographic properties as a building block of a secure quantum cryptosystem. Henceforth, $\mathbb{N}$ denotes the set of all non-negative integers.

Definition 1.1. The *advantage* of a polynomial-time quantum algorithm $\mathcal{A}$ that distinguishes between two ensembles $\{\rho_0(l)\}_{l\in\mathbb{N}}$ and $\{\rho_1(l)\}_{l\in\mathbb{N}}$ of quantum states is the function $\delta_{\mathcal{A}}(l)$ defined as:

$$\delta_{\mathcal{A}}(l) = \left|\Pr_{\mathcal{A}}\left[\mathcal{A}\left(\rho_0(l)\right)=1\right] - \Pr_{\mathcal{A}}\left[\mathcal{A}\left(\rho_1(l)\right)=1\right]\right|$$

for two l-qubit quantum states $\rho_0(l)$ and $\rho_1(l)$, where the subscript $\mathcal{A}$ of the probability means that any output of $\mathcal{A}$ is determined by measuring the final state of $\mathcal{A}$ in the standard computational basis. We say that two ensembles $\{\rho_0(l)\}_{l\in\mathbb{N}}$ and $\{\rho_1(l)\}_{l\in\mathbb{N}}$ are *computationally indistinguishable* if the advantage $\delta_{\mathcal{A}}(l)$ is negligible for any polynomial-time quantum algorithm $\mathcal{A}$; namely, for any polynomial p, any polynomial-time quantum algorithm $\mathcal{A}$, and any sufficiently large number l, it holds that $\delta_{\mathcal{A}}(l) < 1/p(l)$. The distinction problem between $\{\rho_0(l)\}_{l\in\mathbb{N}}$ and $\{\rho_1(l)\}_{l\in\mathbb{N}}$ is said to be *solvable with non-negligible advantage* if these ensembles are not computationally indistinguishable, that is, there exist a polynomial-time quantum algorithm $\mathcal{A}$ and a polynomial p such that

$$\left|\Pr_{\mathcal{A}}\left[\mathcal{A}\left(\rho_0(l)\right)=1\right] - \Pr_{\mathcal{A}}\left[\mathcal{A}\left(\rho_1(l)\right)=1\right]\right| > \frac{1}{p(l)}$$

for infinitely many numbers l.

Let $N = \{n \in \mathbb{N} : n \text{ is even and } n/2 \text{ is odd}\} = \{n \in \mathbb{N} : n \equiv 2 \pmod 4\}$. The problem $\mathrm{QSCD}_{\mathrm{ff}}$ asks whether an adversary can distinguish between two sequences of identical copies of $\rho_\pi^+(n)$ and of $\rho_\pi^-(n)$, where n is a length parameter in N and π is unknown to the adversary. For each $n \in N$, let S_n denote the *symmetric group* of degree n and let $\mathcal{K}_n = \{\pi \in S_n : \pi^2 = id \text{ and } \forall i \in \{1,\ldots,n\}[\pi(i) \neq i]\}$, where id stands for the identity permutation. We say a permutation is odd if it can be expressed by an odd number of transpositions, and even otherwise. Denote by sgn the *sign function* of permutations, defined as $\mathrm{sgn}(\pi) = 0$ if π is even and $\mathrm{sgn}(\pi) = 0$ if π is odd. Notice that, for each $n \in N$, $\mathrm{sgn}(\pi) = 1$ for every $\pi \in \mathcal{K}_n$ (i.e., $\pi \in \mathcal{K}_n$ is an odd permutation) since π consists of $n/2$ disjoint transpositions; in other words, it holds that $\pi = (i_1\, i_2)(i_3\, i_4)\cdots(i_{n-1}\, i_n)$ for n distinct numbers $i_1,\ldots,i_n$ in $\{1,\ldots,n\}$. This simple fact will be used for certain properties of $\mathrm{QSCD}_{\mathrm{ff}}$.

Definition 1.2. For each $\pi \in \mathcal{K}_n$, let $\rho_\pi^+(n)$ and $\rho_\pi^-(n)$ be two quantum states defined by

$$\rho_\pi^+(n) = \frac{1}{2n!} \sum_{\sigma \in S_n} \big(|\sigma\rangle + |\sigma\pi\rangle\big)\big(\langle\sigma| + \langle\sigma\pi|\big) \quad \text{and}$$

$$\rho_\pi^-(n) = \frac{1}{2n!} \sum_{\sigma \in S_n} \big(|\sigma\rangle - |\sigma\pi\rangle\big)\big(\langle\sigma| - \langle\sigma\pi|\big).$$

The problem $\mathrm{QSCD}_{\mathrm{ff}}$ is the distinction problem between two quantum states $\rho_\pi^+(n)^{\otimes k(n)}$ and $\rho_\pi^-(n)^{\otimes k(n)}$ for each parameter n in N, where k is a polynomial. For each fixed polynomial k, we use the succinct notation k-$\mathrm{QSCD}_{\mathrm{ff}}$ instead.

To simplify our notation, we often drop the parameter n whenever it is clear from the context. For instance, we write $\rho_\pi^{+\otimes k}$ instead of $\rho_\pi^+(n)^{\otimes k(n)}$. More generally, k-$\mathrm{QSCD}_{\mathrm{ff}}$ can be defined for any integer-valued function k. Note that Definition 1.2 uses the parameter n to express the "length" of the quantum states instead of the parameter l of Definition 1.1. Speaking of polynomial-time indistinguishability, however, there is essentially no difference between n and l because ρ_π^+ and ρ_π^- can be expressed by $O(n \log n)$ qubits and $k(n)$ is a polynomial in n. In this paper, the parameter n serves as a unit of the computational complexity of our target problem and it is often referred to as the *security parameter* in a cryptographic context.

1.1. *Our Contributions*

This paper presents three properties of $\mathrm{QSCD}_{\mathrm{ff}}$ and their direct implications toward building a secure quantum cryptographic scheme. These properties are summarized as follows: (i) $\mathrm{QSCD}_{\mathrm{ff}}$ has a *trapdoor*; namely, we can efficiently distinguish between ρ_π^+ and ρ_π^- if $\pi \in \mathcal{K}_n$ is known; (ii) The average-case hardness of $\mathrm{QSCD}_{\mathrm{ff}}$ over a randomly chosen permutation $\pi \in \mathcal{K}_n$ coincides with its worst-case hardness; (iii) $\mathrm{QSCD}_{\mathrm{ff}}$ is computationally at least as hard in the worst case as the *graph automorphism problem* (GA), where GA is the graph-theoretical problem defined as:

> GRAPH AUTOMORPHISM PROBLEM (GA):
> input: an undirected graph $G = (V, E)$, where V is a set of nodes and E is a set of edges;
> output: YES if G has a non-trivial automorphism, and NO otherwise.

Since there is no known efficient algorithmic solution for GA, the third property suggests that $\mathrm{QSCD}_{\mathrm{ff}}$ should be difficult to solve in polynomial time. We are also able to show, without any assumption, that no time-unbounded quantum algorithm can solve $o(n \log n)$-$\mathrm{QSCD}_{\mathrm{ff}}$. Making use of the aforementioned three cryptographic properties, we can design a computationally-secure quantum PKC whose security relies on the worst-case hardness of GA. The following subsection will discuss in depth numerous advantages of using $\mathrm{QSCD}_{\mathrm{ff}}$ as a basis of secure quantum cryptosystems.

As a further generalization of $\mathrm{QSCD}_{\mathrm{ff}}$, we present another distinction problem $\mathrm{QSCD}_{\mathrm{cyc}}$ which satisfies the following cryptographic properties: (i) it has a trapdoor and

(ii) its average-case hardness coincides with the worst-case hardness. This new problem becomes a basis for another public-key cryptosystem that can encrypt messages longer than those encrypted by the encryption scheme based on $\mathrm{QSCD}_{\mathrm{ff}}$.

1.2. *Comparison Between Our Work and Previous Work*

In a large volume of the existing literature, computational-complexity aspects of quantum states have been spotlighted in connection to quantum cryptography. In the context of quantum zero-knowledge proofs, for instance, the notion of statistical distinguishability between two quantum states was investigated by Watrous [59] and also by Kobayashi [34]. They proved that certain problems of statistical distinction between two quantum states are promise-complete for quantum zero-knowledge proof systems. Concerning the computational complexity of quantum-state generation, Aharonov and Ta-Shma [2] studied its direct connection to quantum adiabatic computing as well as statistical zero-knowledge proofs. In a similar vein, our distinction problem $\mathrm{QSCD}_{\mathrm{ff}}$ is also rooted in computational complexity theory.

In the remaining of this subsection, we briefly discuss various advantages of using $\mathrm{QSCD}_{\mathrm{ff}}$ as a basis of quantum cryptosystems by comparing it with the underlying problems of existing cryptosystems.

Average-Case Hardness Versus Worst-Case Hardness For any given problem, its efficient solvability on average does not, in general, guarantee that the problem should be solved efficiently even in the worst case. Consider the following property of cryptographic problems: the average-case hardness of the problem is "equivalent" to its worst-case hardness under a certain type of polynomial-time reduction. Since the worst-case hardness of the problem is much more desirable, this average-case/worst-case property certainly increases our confidence in the security of the cryptographic scheme. Unfortunately, few cryptographic problems are known to enjoy this property.

In the literature, there are two major categories of worst-case/average-case reductions. The first category involves a *strong reduction*, which transforms an arbitrary instance of length n to a random instance of the same length n or rather length polynomial in n. With this strong reduction, Ajtai [3] found a remarkable connection between average-case hardness and worst-case hardness of certain variants of the so-called shortest vector problem (SVP). He gave an efficient reduction from a problem of approximating the shortest vector of a given n-dimensional lattice in the worst case to another problem of approximating the shortest vector of a random lattice within a larger approximation factor. Later, Micciancio and Regev [41] established a much better average-case/worst-case connection with respect to the approximation of SVP.

Unlike the first one, the second category is represented by a *weak reduction* of Tompa and Woll [58], where the reduction is randomized only over a certain portion of all the instances. A typical example is DLP, which can be randomly reduced to itself by a reduction that maps instances not to all instances of the same length but rather to all instances of the same underlying group. Concerning DLP, it is not known whether an efficient reduction exists from DLP with the worst-case prime to DLP with a random prime. By Shor's algorithm [56], we can efficiently solve DLP and the inverting problem of the RSA function, which have worst-case/average-case reductions of the second category. The graph isomorphism problem (GI) and the aforementioned GA—well-known

graph-theoretical problems—also satisfy weak worst-case/average-case reductions [58] although there is no known cryptosystem whose security relies on their hardness; see [10] and references therein for more information on worst-case/average-case reductions.

In this paper, we show that $\mathrm{QSCD}_{\mathrm{ff}}$ has a worst-case/average-case reduction of the *first category*. Unlike the reduction of DLP, our reduction depends only on the size of each instance. In fact, our distinction problem $\mathrm{QSCD}_{\mathrm{ff}}$ is the *first* cryptographic problem having a worst-case/average-case reduction of the first category; namely, the worst case of the problem can be reduced to the average case of the *same* problem. Our reduction is similar in flavor to the reductions used for the aforementioned lattice problems. In the case of the approximation of SVP, however, an approximation problem of SVP can be reduced randomly only to *another* approximation problem with a worse parameter. Note that, on a quantum computer, no efficient solution is currently known for $\mathrm{QSCD}_{\mathrm{ff}}$.

Computational Hardness of Underlying Computational Problems The hidden subgroup problem (HSP) has played a central role in various discussions on the strengths and limitations of quantum computation. The aforementioned IFP and DLP can be reduced to special cases of HSP on Abelian groups (AHSP). Kitaev [33] showed how to solve AHSP efficiently; in particular, he gave a polynomial-time algorithm that performs the quantum Fourier transformation over Abelian groups, which is a generalization of the quantum Fourier transformation used in, e.g., Shor's algorithm [56]. To solve HSP on non-Abelian groups, a simple application of currently known techniques may not be sufficient despite of the existence of an efficient quantum algorithm for AHSP. Notice that over certain specific non-Abelian groups HSP was already solved in [6,19,22,26, 36,43,50]. Another important variant of HSP is HSP on the dihedral groups (DHSP). Regev [50] demonstrated a quantum reduction from the unique shortest vector problem (uSVP) to a slightly different variant of DHSP, where uSVP can serve as a basis of lattice-based PKCs defined in [4,49]. A subexponential-time quantum algorithm for DHSP was found by Kuperberg [36]. Although these results do not immediately give a desired subexponential-time quantum algorithm for uSVP, it could eventually lead us to design the desired algorithm.

Our problem $\mathrm{QSCD}_{\mathrm{ff}}$ is closely related to another variant: HSP on the *symmetric groups* (SHSP), which appears to be much more difficult to solve than the aforementioned variants of HSP do. Note that no known subexponential-time quantum algorithm exists for SHSP. Recently, Hallgren, Russell, and Ta-Shma [26] introduced a distinction problem, similar to $\mathrm{QSCD}_{\mathrm{ff}}$, between certain two quantum states to discuss the computational intractability of SHSP by a "natural" extension of Shor's algorithm [56]. In this paper, we refer to their distinction problem as DIST. An efficient solution to DIST gives rise to an efficient quantum algorithm for a certain special case of SHSP. To solve DIST, as they showed, we require exponentially many trials of the so-called *weak Fourier sampling* that works on a single copy of the quantum states. In other words, exponentially many copies are needed in total as far as the weak Fourier sampling is used.

This result was improved by Grigni, Schulman, Vazirani, and Vazirani [22], who proved that exponentially many copies are necessary even if we use a powerful method, known as *strong Fourier sampling*, along with a random choice of the bases of the representations of the symmetric group S_n. Concerning the computational hardness of SHSP, Kempe and Shalev [32] further expanded the results of [22,26] with quantum Fourier

sampling methods. Moore, Russell, and Schulman [44], on the contrary, demonstrated that, regardless of the method (such as the above quantum Fourier sampling methods), any time-unbounded quantum algorithm working on a single copy needs $\exp(\Omega(n))$ trials to solve DIST. Even for the case of two copies, Moore and Russell [42] argued that any time-unbounded quantum algorithm that simultaneously works over two copies requires $\exp(\Omega(\sqrt{n}/\log n))$ trials at best. Their results were further improved by Hallgren, Moore, Rötteler, Russell, and Sen [25], who proved that no time-unbounded quantum algorithm solves DIST even if it simultaneously works over $o(n \log n)$ copies. In this paper, we show that the distinction problem DIST is, in fact, polynomial-time reducible to $\mathrm{QSCD}_{\mathrm{ff}}$. This immediately implies, from the above results, that no quantum algorithm solves $\mathrm{QSCD}_{\mathrm{ff}}$ using $o(n \log n)$ copies.

Even by supplying sufficiently many copies to an algorithm, there is no known subexponential-time quantum algorithm that solves $\mathrm{QSCD}_{\mathrm{ff}}$, and thus finding such an algorithm seems a daunting task. This situation indicates that our problem, $\mathrm{QSCD}_{\mathrm{ff}}$, is much more suitable than, for example, uSVP for an underlying intractable problem to build a secure cryptosystem. There is a similarity with the classical case of DLP over different groups; namely, DLP over $\mathbb{Z}_p^*$ (where p is a prime) is classically computable in subexponential time whereas no known classical subexponential-time algorithm exists for DLP over certain groups in elliptic curve cryptography. From this reason, it is generally believed that DLP over such groups is more reliable than DLP over $\mathbb{Z}_p^*$.

We prove that the computational complexity of $\mathrm{QSCD}_{\mathrm{ff}}$ is lower-bounded by that of GA. Well-known upper bounds of GA include $\mathbf{NP} \cap \text{co-}\mathbf{AM}$ [21,54], **SPP** [5], and **UAP** [12]; however, GA is not known to sit in $\mathbf{NP} \cap \text{co-}\mathbf{NP}$. Notice that, since most cryptographic problems fall in $\mathbf{NP} \cap \text{co-}\mathbf{NP}$, very few cryptographic systems are lower-bounded by the worst-case hardness of problems outside of $\mathbf{NP} \cap \text{co-}\mathbf{NP}$.

Quantum Computational Cryptography Apart from PKCs, quantum key distribution gives a foundation to symmetric-key cryptology; for instance, the quantum key distribution scheme in [8] achieves unconditionally secure sharing of secret keys in symmetric-key cryptosystems (SKCs) through an authenticated classical communication channel and an insecure quantum communication channel. Undoubtedly, both SKCs and PKCs have their own advantages and disadvantages. Compared with SKCs, PKCs require fewer secret keys in a large-scale network; however, they often need certain intractability assumptions for their security proofs and are typically vulnerable to, e.g., the man-in-the-middle attack. As an immediate application of $\mathrm{QSCD}_{\mathrm{ff}}$, we propose a new computational quantum PKC whose security relies on the computational hardness of $\mathrm{QSCD}_{\mathrm{ff}}$.

Of many existing PKCs, few make their security proofs solely rely on the *worst-case* hardness of their underlying problems, such as lattice-based PKCs (see, e.g., [52]). A quantum adversary is a powerful foe who can easily break many PKCs whose underlying problems are number-theoretic because these problems can be efficiently solved on a quantum computer. Based on a certain subset of the knapsack problem, Okamoto, Tanaka, and Uchiyama [48] proposed a quantum PKC which withstands certain well-known quantum attacks. Our proposed quantum PKC also seems to fend off a polynomial-time quantum adversary since we can reduce the problem GA to $\mathrm{QSCD}_{\mathrm{ff}}$, where GA is not known to be solved efficiently on a quantum computer.

1.3. *Later Work*

After the publication of the preliminary version [30] of this paper, the notion of quantum-state indistinguishability and its associated quantum encryption schemes have been further studied. Here are some of the recent results related to the topics of this paper. Hayashi, Kawachi, and Kobayashi [27] showed that $\mathrm{QSCD}_{\mathrm{cyc}}$ satisfies the indistinguishability property against time-unbounded quantum algorithms in such a way that $\mathrm{QSCD}_{\mathrm{ff}}$ does. In information-theoretical settings, Nikolopoulos [46] and Nikolopoulos and Ioannou [47] proposed new quantum encryption schemes. Kawachi and Portmann [31] proved that, with respect to the ratio of message length and key size, any quantum encryption scheme has no advantage over a classical one-time pad scheme if we impose certain information-theoretically strong security requirement on the quantum encryption scheme.

2. Cryptographic Properties of $\mathrm{QSCD}_{\mathrm{ff}}$

Through this section, we will show that $\mathrm{QSCD}_{\mathrm{ff}}$ enjoys the following three cryptographically useful properties: (i) a trapdoor, (ii) the equivalence between average-case hardness and worst-case hardness under polynomial-time reductions, and (iii) a reduction from two computationally-hard problems to $\mathrm{QSCD}_{\mathrm{ff}}$. These properties will help us to construct a quantum PKC in Sect. 3. We assume, throughout this paper, the reader's familiarity with the basics of quantum computation [45] and of finite group theory [53].

All the cryptographic properties of $\mathrm{QSCD}_{\mathrm{ff}}$ are consequences of the following characteristics of the set $\mathcal{K}_n$ of the hidden permutations. (i) Each permutation $\pi \in \mathcal{K}_n$ is of order 2. This provides the trapdoor of $\mathrm{QSCD}_{\mathrm{ff}}$. (ii) For any $\pi \in \mathcal{K}_n$, the conjugacy class $\{\tau^{-1}\pi\tau : \tau \in S_n\}$ of π is equal to $\mathcal{K}_n$. This property enables us to prove the equivalence between the worst-case hardness and average-case hardness of $\mathrm{QSCD}_{\mathrm{ff}}$. (iii) The problem GA is (polynomial-time Turing) equivalent to its subproblem with the promise that any given graph has either a unique non-trivial automorphism in $\mathcal{K}_n$ or none at all. This equivalence relation is used to give a complexity-theoretic lower bound of $\mathrm{QSCD}_{\mathrm{ff}}$, that is, the average-case hardness of $\mathrm{QSCD}_{\mathrm{ff}}$ is lower-bounded by the worst-case hardness of GA. To prove those properties, we introduce two new techniques: (i) a variant of the so-called *coset sampling method*, which is widely used in various extensions of Shor's well-known algorithm (see, e.g., [50]), and (ii) a quantum version of the *hybrid argument*, which is a powerful tool for many security reductions used in computational cryptography.

Now, recall the two quantum states $\rho_\pi^+ = \frac{1}{2n!}\sum_{\sigma\in S_n}(|\sigma\rangle + |\sigma\pi\rangle)(\langle\sigma| + \langle\sigma\pi|)$ and $\rho_\pi^- = \frac{1}{2n!}\sum_{\sigma\in S_n}(|\sigma\rangle - |\sigma\pi\rangle)(\langle\sigma| - \langle\sigma\pi|)$ for a permutation $\pi \in \mathcal{K}_n$. For convenience, let $\iota(n)$ (or simply ι) denote the maximally mixed state $\frac{1}{n!}\sum_{\sigma\in S_n}|\sigma\rangle\langle\sigma|$ over S_n, which will appear later.

2.1. *A Trapdoor*

We start by proving that $\mathrm{QSCD}_{\mathrm{ff}}$ has a *trapdoor*. To prove this claim, it suffices to present an efficient distinguishing algorithm between ρ_π^+ and ρ_π^- with an extra knowledge of their hidden permutation $\pi \in \mathcal{K}_n$.

Theorem 2.1 (Distinguishing Algorithm). *There exists a polynomial-time quantum algorithm that, for any security parameter $n \in N$ and for any hidden permutation $\pi \in \mathcal{K}_n$, distinguishes between $\rho_\pi^+(n)$ and $\rho_\pi^-(n)$ using π with probability* 1.

Proof. Fix $n \in N$ arbitrarily. Let χ be any given unknown quantum state which is limited to either ρ_π^+ or ρ_π^-. The desired distinguishing algorithm for χ works as follows:

(D1) Prepare two quantum registers. The first register holds a control bit and the second register holds χ. Apply the Hadamard transformation H to the first register. The state of the system now becomes

$$H|0\rangle\langle 0|H \otimes \chi.$$

(D2) Apply the Controlled-π operator C_π to the both registers, where the operator C_π behaves as $C_\pi|0\rangle|\sigma\rangle = |0\rangle|\sigma\rangle$ and $C_\pi|1\rangle|\sigma\rangle = |1\rangle|\sigma\pi\rangle$ for any given $\sigma \in S_n$. Since $\pi^2 = id$ for every $\pi \in \mathcal{K}_n$, the state of the entire system can be expressed as

$$\frac{1}{n!}\sum_{\sigma\in S_n} |\psi_{\pi,\sigma}^+\rangle\langle\psi_{\pi,\sigma}^+| \quad \text{if } \chi = \rho_\pi^+, \quad \text{and} \quad \frac{1}{n!}\sum_{\sigma\in S_n} |\psi_{\pi,\sigma}^-\rangle\langle\psi_{\pi,\sigma}^-| \quad \text{if } \chi = \rho_\pi^-,$$

where $|\psi_{\pi,\sigma}^+\rangle$ and $|\psi_{\pi,\sigma}^-\rangle$ are defined as

$$\begin{aligned}|\psi_{\pi,\sigma}^\pm\rangle &= C_\pi\left(\frac{1}{2}|0\rangle\big(|\sigma\rangle \pm |\sigma\pi\rangle\big) + \frac{1}{2}|1\rangle\big(|\sigma\rangle \pm |\sigma\pi\rangle\big)\right)\\ &= \frac{1}{2}|0\rangle\big(|\sigma\rangle \pm |\sigma\pi\rangle\big) + \frac{1}{2}|1\rangle\big(|\sigma\pi\rangle \pm |\sigma\rangle\big).\end{aligned}$$

(D3) Apply the Hadamard transformation again to the first register. Since χ is either ρ_π^+ or ρ_π^-, the state of the entire system becomes either

$$\begin{aligned}(H\otimes I)|\psi_{\pi,\sigma}^+\rangle &= \frac{1}{\sqrt{2}}|0\rangle\big(|\sigma\rangle + |\sigma\pi\rangle\big) \quad \text{or}\\ (H\otimes I)|\psi_{\pi,\sigma}^-\rangle &= \frac{1}{\sqrt{2}}|1\rangle\big(|\sigma\rangle - |\sigma\pi\rangle\big),\end{aligned}$$

respectively. Measure the first register in the computational basis. If the measured result is 0, then output YES; otherwise, output NO.

It is clear that the above procedure gives the correct answer with probability 1. □

2.2. *A Reduction from Worst Case to Average Case*

We intend to reduce the worst-case hardness of $\mathrm{QSCD}_{\mathrm{ff}}$ to its average-case hardness. Such a reduction implies that $\mathrm{QSCD}_{\mathrm{ff}}$ with a random permutation π is at least as hard as $\mathrm{QSCD}_{\mathrm{ff}}$ with the fixed permutation π' of the highest complexity. Since the converse reduction is trivial, the average-case hardness of $\mathrm{QSCD}_{\mathrm{ff}}$ is therefore polynomial-time Turing equivalent to its worst-case hardness.

Theorem 2.2. *Let k be any polynomial and let $\mathcal{A}$ be a polynomial-time quantum algorithm that solves k-QSCD$_{\mathrm{ff}}$ with non-negligible advantage for a uniformly random $\pi \in \mathcal{K}_n$; namely, there exists a polynomial p such that, for infinitely many security parameters n in N,*

$$\left| \Pr_{\pi,\mathcal{A}}\left[\mathcal{A}\left(\rho_\pi^+(n)^{\otimes k(n)}\right)=1\right] - \Pr_{\pi,\mathcal{A}}\left[\mathcal{A}\left(\rho_\pi^-(n)^{\otimes k(n)}\right)=1\right]\right| > \frac{1}{p(n)},$$

where π is chosen uniformly at random from $\mathcal{K}_n$. Then, there exists a polynomial-time quantum algorithm $\mathcal{B}$ that solves k-QSCD$_{\mathrm{ff}}$ with non-negligible advantage for any permutation $\pi \in \mathcal{K}_n$.

Proof. Fix an arbitrary parameter $n \in N$ that satisfies the assumption of the theorem. Assume that our input is either $\rho_\pi^+(n)^{\otimes k(n)}$ or $\rho_\pi^-(n)^{\otimes k(n)}$. For each $i \in \{1, 2, \ldots, k(n)\}$, let χ_i be the ith state of the given $k(n)$ states. Clearly, χ_i is either ρ_π^+ or ρ_π^-. From the given average-case algorithm $\mathcal{A}$, we build the desired worst-case algorithm $\mathcal{B}$ in the following way:

(R1) Choose a permutation $\tau \in S_n$ uniformly at random.
(R2) Apply τ to each χ_i, where $i \in \{1, \ldots, k(n)\}$, from the right. If $\chi_i = \rho_\pi^+$, then we obtain the quantum state

$$\chi_i' = \frac{1}{2n!}\sum_{\sigma\in S_n}\left(|\sigma\tau\rangle + \left|\sigma\tau\tau^{-1}\pi\tau\right\rangle\right)\left(\langle\sigma\tau| + \left\langle\sigma\tau\tau^{-1}\pi\tau\right|\right)$$

$$= \frac{1}{2n!}\sum_{\sigma'\in S_n}\left(|\sigma'\rangle + \left|\sigma'\tau^{-1}\pi\tau\right\rangle\right)\left(\langle\sigma'| + \left\langle\sigma'\tau^{-1}\pi\tau\right|\right).$$

When $\chi_i = \rho_\pi^-$, we instead obtain $\chi_i' = \frac{1}{2n!}\sum_{\sigma'\in S_n}(|\sigma'\rangle - |\sigma'\tau^{-1}\pi\tau\rangle)(\langle\sigma'| - \langle\sigma'\tau^{-1}\pi\tau|)$.
(R3) Invoke the average-case quantum algorithm $\mathcal{A}$ on the input $\bigotimes_{i=1}^k \chi_i'$.
(R4) Output the outcome of $\mathcal{A}$.

Let $\pi \in \mathcal{K}_n$. Note that, for each $\tau \in S_n$, $\tau^{-1}\pi\tau$ belongs to $\mathcal{K}_n$. Moreover, for every $\pi' \in \mathcal{K}_n$, there exists a $\tau \in S_n$ satisfying $\tau^{-1}\pi\tau = \pi'$, from which it follows that the conjugacy class $\{\tau^{-1}\pi\tau : \tau \in S_n\}$ of π is equal to $\mathcal{K}_n$. As shown below, the number of all permutations $\tau \in S_n$ for which $\tau^{-1}\pi\tau = \pi'$ is independent of the choice of $\pi' \in \mathcal{K}_n$.

Claim 1. *For any permutations $\pi, \pi', \pi'' \in \mathcal{K}_n$, $|\{\tau \in S_n : \tau^{-1}\pi\tau = \pi'\}| = |\{\tau \in S_n : \tau^{-1}\pi\tau = \pi''\}|$.*

Proof. Define a map $\mu_\tau : \mathcal{K}_n \to \mathcal{K}_n$ as $\mu_\tau(\sigma) = \tau^{-1}\sigma\tau$ and a set $\mathcal{T}_{\pi,\pi'} := \{\mu_\tau : \mu_\tau(\pi) = \pi'\}$. It is obvious that, by defining a group operation "$\cdot$" as $\mu_\tau \cdot \mu_{\tau'}(\cdot) = \mu_\tau(\mu_{\tau'}(\cdot))$, $\mathcal{T}_{\pi,\pi}$ becomes a subgroup of $\mathcal{S}_n := \{\mu_\tau : \tau \in S_n\}$. Therefore, $\mathcal{S}_n$ has a coset decomposition with respect to its subgroup $\mathcal{T}_{\pi,\pi}$ for any $\pi \in \mathcal{K}_n$ and each coset coincides with $\mathcal{T}_{\pi,\pi'}$ for a certain π'. This shows that $|\mathcal{T}_{\pi,\pi'}| = |\mathcal{T}_{\pi,\pi''}|$ for every pair π', π''. Since μ_τ and τ have a one-to-one correspondence, it follows that, for every π', π'', $|\{\tau \in S_n : \tau^{-1}\pi\tau = \pi'\}| = |\{\tau \in S_n : \tau^{-1}\pi\tau = \pi''\}|$. □

The above-mentioned properties imply that $\tau^{-1}\pi\tau$ is indeed uniformly distributed over $\mathcal{K}_n$. Therefore, by feeding the input $\bigotimes_{i=1}^{k}\chi_i'$ to the algorithm $\mathcal{A}$, we can achieve the desired non-negligible advantage of $\mathcal{A}$. This completes the proof. □

2.3. *Computational Hardness*

The third property of $\mathrm{QSCD}_{\mathrm{ff}}$ relates to the computational hardness of $\mathrm{QSCD}_{\mathrm{ff}}$. We want to present two claims that witness its relative hardness against GA. First, we prove that the computational complexity of $\mathrm{QSCD}_{\mathrm{ff}}$ is lower-bounded by that of GA by constructing an efficient reduction from GA to $\mathrm{QSCD}_{\mathrm{ff}}$. Second, we briefly discuss relationships among $\mathrm{QSCD}_{\mathrm{ff}}$, SHSP, and DIST, and we then prove that $\mathrm{QSCD}_{\mathrm{ff}}$ cannot be solved from $o(n \log n)$ copies of input instances.

Now, we prove the first claim concerning the reducibility between GA and $\mathrm{QSCD}_{\mathrm{ff}}$. Our reduction from GA to $\mathrm{QSCD}_{\mathrm{ff}}$ consists of two parts: a reduction from GA to a variant of GA, called $\mathrm{UniqueGA}_{\mathrm{ff}}$, and a reduction from $\mathrm{UniqueGA}_{\mathrm{ff}}$ to $\mathrm{QSCD}_{\mathrm{ff}}$. To describe the desired reduction, we formally introduce $\mathrm{UniqueGA}_{\mathrm{ff}}$. Earlier, Köbler, Schöning, and Torán [35] introduced the following *unique graph automorphism problem* (UniqueGA).

> UNIQUE GRAPH AUTOMORPHISM PROBLEM (UniqueGA):
> input: an undirected graph $G = (V, E)$, where V is a set of nodes and E is a set of edges;
> promise: G has either a unique non-trivial automorphism or no non-trivial automorphism;
> output: YES if G has the non-trivial automorphism, and NO otherwise.

Note that this promise problem UniqueGA is called (1GA, GA) in [35]. The *unique graph automorphism with fully-flipped permutation* ($\mathrm{UniqueGA}_{\mathrm{ff}}$) is a slight modification of UniqueGA. Recall that $N = \{n' \in \mathbb{N} : n' \equiv 2 \pmod 4\}$.

> UNIQUE GRAPH AUTOMORPHISM WITH FULLY-FLIPPED PERMUTATION ($\mathrm{UniqueGA}_{\mathrm{ff}}$):
> input: an undirected graph $G = (V, E)$, where V is a set of nodes and E is a set of edges;
> promise: the number $n = |V|$ of nodes is in N. Moreover, G has either a unique non-trivial automorphism $\pi \in \mathcal{K}_n$ or no non-trivial automorphism;
> output: YES if G has the non-trivial automorphism, and NO otherwise.

Note that every instance G of $\mathrm{UniqueGA}_{\mathrm{ff}}$ is defined only when the number n of nodes belongs to N.

Regarding $\mathrm{UniqueGA}_{\mathrm{ff}}$, we want to prove two helpful lemmas. The first lemma uses a variant of the so-called *coset sampling method*, which has been widely used in many generalizations of Shor's algorithm. Recall that $\iota(n) = \frac{1}{n!}\sum_{\sigma\in S_n} |\sigma\rangle\langle\sigma|$ for each $n \in N$.

Lemma 2.3. *There exists a polynomial-time quantum algorithm that, given an instance G of $\mathrm{UniqueGA}_{\mathrm{ff}}$, generates a quantum state ρ_π^+ if G is a "YES" instance with its unique non-trivial automorphism π, or generates ι if G is a "NO" instance.*

Proof. Let $n \in N$. Given an instance G of UniqueGA$_{\mathrm{ff}}$, we first prepare the quantum state $\frac{1}{\sqrt{n!}} \sum_{\sigma \in S_n} |\sigma\rangle|\sigma(G)\rangle$, where $\sigma(G)$ is the graph resulting from relabeling its nodes according to each permutation σ. By discarding the second register, we can obtain a quantum state χ in the first register. If G is a "YES" instance with the unique non-trivial automorphism π, then this state χ equals ρ_π^+ since $\frac{1}{\sqrt{n!}} \sum_{\sigma} |\sigma\rangle|\sigma(G)\rangle = \frac{1}{\sqrt{n!}} \sum_{\sigma \in S_n/\langle\pi\rangle} (|\sigma\rangle + |\sigma\pi\rangle)|\sigma(G)\rangle$. Otherwise, since $\sigma(G) \neq \sigma'(G)$ for any distinct $\sigma, \sigma' \in S_n$, χ equals $\iota = \frac{1}{n!} \sum_{\sigma \in S_n} |\sigma\rangle\langle\sigma|$. □

The second lemma requires a variant of the coset sampling method as a technical tool. The lemma in essence relies on the fact that the hidden $\pi \in \mathcal{K}_n$ is an odd permutation for each $n \in N$ since, as a special property of $\mathcal{K}_n$, π can be expressed as a product of an odd number of transpositions.

Lemma 2.4. *There exists a polynomial-time quantum algorithm that, given an instance G of UniqueGA$_{\mathrm{ff}}$, generates a quantum state ρ_π^- if G is a "YES" instance with the unique non-trivial automorphism π or generates ι if G is a "NO" instance.*

Proof. Let $n \in N$. Similar to the algorithm given in the proof of Lemma 2.3, we start with the quantum state $\frac{1}{\sqrt{n!}} \sum_{\sigma \in S_n} |\sigma\rangle|\sigma(G)\rangle$ in two registers. Compute the sign of each permutation in the first register and then invert its phase only when the permutation is odd. Consequently, we obtain the quantum state $\frac{1}{\sqrt{n!}} \sum_{\sigma \in S_n} (-1)^{\mathrm{sgn}(\sigma)} |\sigma\rangle|\sigma(G)\rangle$. Recall that $\mathrm{sgn}(\sigma) = 0$ if σ is even, and $\mathrm{sgn}(\sigma) = 1$ otherwise. By discarding the second register, we immediately obtain a certain quantum state, say, χ in the first register. Note that, since π is odd, if σ is odd (even, resp.) then $\sigma\pi$ is even (odd, resp.). Therefore, it follows that $\chi = \rho_\pi^-$ if G is a "YES" instance with the unique non-trivial automorphism π, and $\chi = \iota$ otherwise. □

We are now ready to present a polynomial-time reduction from GA to QSCD$_{\mathrm{ff}}$. This implies that QSCD$_{\mathrm{ff}}$ is computationally at least as hard as GA for infinitely-many input lengths n (and thus in the worst-case).

Theorem 2.5. *If there exist a polynomial k and a polynomial-time quantum algorithm that solves k-QSCD$_{\mathrm{ff}}$ with non-negligible advantage, then there exists a polynomial-time quantum algorithm that solves GA in the worst case for infinitely-many input lengths n.*

Proof. We first show that GA is polynomial-time Turing equivalent to UniqueGA$_{\mathrm{ff}}$. Later, we give a polynomial-time Turing reduction from UniqueGA$_{\mathrm{ff}}$ to QSCD$_{\mathrm{ff}}$. By combining these two reductions, we can reduce GA to QSCD$_{\mathrm{ff}}$. The reduction from GA to UniqueGA$_{\mathrm{ff}}$ we define is similar to the one given by Köbler, Schöning, and Torán [35], who presented a polynomial-time Turing reduction from GA to UniqueGA. Their polynomial-time algorithm for GA makes queries to a given oracle that correctly represents UniqueGA on the promised inputs. This algorithm works correctly because all queries made by the algorithm satisfy the promise of UniqueGA, that is, every query is a graph of even number of nodes with either a unique non-trivial automorphism

without any fixed point or no non-trivial automorphism at all. By a slight modification of their reduction, we can obtain a reduction from GA to UniqueGA$_{\rm ff}$. Furthermore, it is also possible to make our length parameter n satisfy the specific equation $n = 2(2n' + 1)$, where $n' \in \mathbb{N}$. As a result, we obtain the following lemma.

Lemma 2.6. *UniqueGA$_{\rm ff}$ is polynomial-time Turing equivalent to GA.*

In fact, a stronger statement than Lemma 2.6 holds. When a Turing reduction to a promise problem makes only queries that satisfy the promise of the problem, this reduction is called *smart* [23]. The reduction from GA to UniqueGA given by Köbler, Schöning, and Torán [35] is indeed smart, and therefore so is our reduction. For readability, we postpone the proof of Lemma 2.6 until Appendix.

From Lemma 2.6, it suffices to construct a reduction from UniqueGA$_{\rm ff}$ to QSCD$_{\rm ff}$. Assume that there exist two polynomials k and p and also a polynomial-time quantum algorithm $\mathcal{A}$ such that, for infinitely many n's, $\mathcal{A}$ solves k-QSCD$_{\rm ff}$ with advantage $1/p(n)$. Let us fix an arbitrary n for which $\mathcal{A}$ solves k-QSCD$_{\rm ff}$ with advantage $1/p(n)$. On a given instance G of UniqueGA$_{\rm ff}$, we perform the following procedure:

(S1) Generate from G two sequences $S^+ = (\chi^{+\otimes k}, \ldots, \chi^{+\otimes k})$ and $S^- = (\chi^{-\otimes k}, \ldots, \chi^{-\otimes k})$ of $8p^2(n)n$ instances by running the generation algorithms given in Lemmas 2.3 and 2.4, respectively.
(S2) Invoke $\mathcal{A}$ on each component in S^+ and S^- as an input. Let $R^+ = (\mathcal{A}(\chi^{+\otimes k}), \ldots, \mathcal{A}(\chi^{+\otimes k}))$ and $R^- = (\mathcal{A}(\chi^{-\otimes k}), \ldots, \mathcal{A}(\chi^{-\otimes k}))$ be the resulting sequences of $8p^2(n)n$ entries.
(S3) Output YES if the difference ℓ between the number of 1's in R^+ and that in R^- is at least $4p(n)n$; output NO otherwise.

Note that if G is a "YES" instance, then S^+ and S^- should have the form $S^+ = (\rho_\pi^{+\otimes k}, \ldots, \rho_\pi^{+\otimes k})$ and $S^- = (\rho_\pi^{-\otimes k}, \ldots, \rho_\pi^{-\otimes k})$ of $8p^2(n)n$ entries; otherwise, we have $S^+ = S^- = (\iota^{\otimes k}, \ldots, \iota^{\otimes k})$. Therefore, if G is a "YES" instance, the numbers of 1s in R^+ and in R^- are highly likely different.

Finally, we estimate the difference ℓ. Let X^+ and X^- be two random variables respectively expressing the numbers of 1s in R^+ and in R^-. Assume that G is a "YES" instance. Since $\mathcal{A}$ solves k-QSCD$_{\rm ff}$ with advantage $1/p(n)$, we have $|\Pr[\mathcal{A}(\rho_\pi^{+\otimes k}) = 1] - \Pr[\mathcal{A}(\rho_\pi^{-\otimes k}) = 1]| > 1/p(n)$. Next, we want to show that $\Pr[|X^+ - X^-| > 4p(n)n] > 1 - 2e^{-n}$ using the Höffding bounds, which are stated below.

Lemma 2.7 (Höffding [28]). *Let $(X_1, \ldots, X_m)$ be any sequence of independent Bernoulli random variables on $\{0, 1\}$ such that $\Pr[X_i = 1] = p$ for any $i \in \{1, \ldots, m\}$, and let X be a random variable expressing the number of 1s in the sequence, i.e., $X = \sum_{i=1}^m X_i$. Then, for any $0 \le \delta \le 1$, it holds that*

$$\Pr\big[X > (p + \delta)m\big] < e^{-2m\delta^2} \quad \text{and} \quad \Pr\big[X < (p - \delta)m\big] < e^{-2m\delta^2}.$$

For convenience, we define $p_L = \max\{\Pr[\mathcal{A}(\rho_\pi^{+\otimes k}) = 1], \Pr[\mathcal{A}(\rho_\pi^{-\otimes k}) = 1]\}$ and $p_S = \min\{\Pr[\mathcal{A}(\rho_\pi^{+\otimes k}) = 1], \Pr[\mathcal{A}(\rho_\pi^{-\otimes k}) = 1]\}$. From our assumption, we obtain

$p_L - p_S > 1/p(n)$. Note that R^+ and R^- are precisely two sequences of $8p^2(n)n$ independent Bernoulli random variables on $\{0, 1\}$ with probabilities p_L and p_S. We denote by X_L (X_S, resp.) the number of 1s in the sequence associated with p_L (p_S, resp.). The Höffding bounds imply

$$\Pr\bigl[X_L < (p_L - \delta)m\bigr] < e^{-n} \quad \text{and} \quad \Pr\bigl[X_S > (p_S + \delta)m\bigr] < e^{-n},$$

where $m = 8p^2(n)n$ and $\delta = 1/(4p(n))$. Since $p_L - p_S > 1/p(n)$, we obtain $(p_L - p_S - 2\delta)m > 4p(n)n$. From this inequality, it follows that

$$\begin{aligned}\Pr\bigl[\bigl|X^+ - X^-\bigr| > 4p(n)n\bigr] &\geq \Pr\bigl[\bigl|X^+ - X^-\bigr| > (p_L - p_S - 2\delta)m\bigr] \\ &\geq \Pr\bigl[X_L > (p_L - \delta)m \ \wedge \ X_S < (p_S + \delta)m\bigr].\end{aligned}$$

Since X_L and X_S are independent, we obtain a lower bound:

$$\Pr\bigl[X_L > (p_L - \delta)m \ \wedge \ X_S < (p_S + \delta)m\bigr] \geq \bigl(1 - e^{-n}\bigr)^2 > 1 - 2e^{-n},$$

from which we conclude that $\Pr[|X^+ - X^-| > 4p(n)n] > 1 - 2e^{-n}$.

Similarly, when G is a "NO" instance, we have $\Pr[|X^+ - X^-| < 4p(n)n] > 1 - 2e^{-n}$. This guarantees that the above procedure solves $\mathrm{UniqueGA}_{\mathrm{ff}}$ efficiently. □

As noted in Sect. 1, our distinction problem $\mathrm{QSCD}_{\mathrm{ff}}$ has its roots in SHSP. A special case of SHSP is known to be reducible to DIST, which is a problem of distinguishing between $\{\rho_\pi^+(n)\}_{n\in N}$ and $\{\iota(n)\}_{n\in N}$. As Hallgren, Moore, Rötteler, Russell, and Sen [25] demonstrated, solving DIST from $o(n \log n)$ identical copies is impossible even for a time-unbounded quantum algorithm. Now, we show a close relationship between $\mathrm{QSCD}_{\mathrm{ff}}$ and DIST.

Before stating our claim (Theorem 2.9), we present an algorithm that converts ρ_π^+ to ρ_π^- for each fixed $\pi \in \mathcal{K}_n$. This algorithm is a key to the proof of the theorem and further to the construction of a quantum PKC in the subsequent section.

Lemma 2.8 (Conversion Algorithm). *There exists a polynomial-time quantum algorithm that, with certainty, converts $\rho_\pi^+(n)$ into $\rho_\pi^-(n)$ and keeps $\iota(n)$ as it is for any parameter $n \in N$ and any hidden permutation $\pi \in \mathcal{K}_n$.*

Proof. Let $n \in N$ be arbitrary. First, recall the definition of $\mathrm{sgn}(\sigma)$: $\mathrm{sgn}(\sigma) = 0$ if σ is even and $\mathrm{sgn}(\sigma) = 1$ otherwise. Let $\pi \in \mathcal{K}_n$ be any hidden permutation and consider its corresponding quantum state ρ_π^+. On input ρ_π^+, our desired algorithm simply inverts its phase according to the sign of the permutation. This is done by performing the following transformation:

$$|\sigma\rangle + |\sigma\pi\rangle \longmapsto (-1)^{\mathrm{sgn}(\sigma)}|\sigma\rangle + (-1)^{\mathrm{sgn}(\sigma\pi)}|\sigma\pi\rangle.$$

Note that determining the sign of a given permutation takes only time polynomial in n. Since π is odd, $\mathrm{sgn}(\sigma)$ and $\mathrm{sgn}(\sigma\pi)$ are different; thus, the above algorithm obviously converts ρ_π^+ to ρ_π^-. Moreover, the algorithm does not alter the quantum state ι. □

The intractability result of DIST [25] stated above also holds for $\mathrm{QSCD}_{\mathrm{ff}}$. To prove this claim, we want to show in Theorem 2.9 that DIST can be reduced to $\mathrm{QSCD}_{\mathrm{ff}}$ in polynomial time. As a result, no time-unbounded quantum algorithm can solve $\mathrm{QSCD}_{\mathrm{ff}}$ from $o(n \log n)$ copies. The proof of the theorem requires a quantum version of the so-called *hybrid argument* used in computational cryptography.

Theorem 2.9. *Let k be any polynomial. If there exists a quantum algorithm $\mathcal{A}$ such that*

$$\left|\Pr_{\mathcal{A}}\left[\mathcal{A}\left(\rho_\pi^+(n)^{\otimes k(n)}\right)=1\right]-\Pr_{\mathcal{A}}\left[\mathcal{A}\left(\rho_\pi^-(n)^{\otimes k(n)}\right)=1\right]\right|>\varepsilon(n)$$

for any security parameter $n \in N$, then there exists a quantum algorithm $\mathcal{B}$ such that, for each $n \in N$,

$$\left|\Pr_{\mathcal{B}}\left[\mathcal{B}\left(\rho_\pi^+(n)^{\otimes k(n)}\right)=1\right]-\Pr_{\mathcal{B}}\left[\mathcal{B}\left(\iota(n)^{\otimes k(n)}\right)=1\right]\right|>\frac{\varepsilon(n)}{4}.$$

Proof. Fix $n \in N$ arbitrarily, and we hereafter omit this parameter n. Assume that a quantum algorithm $\mathcal{A}$ distinguishes between $\rho_\pi^{+\otimes k}$ and $\rho_\pi^{-\otimes k}$ with advantage at least $\varepsilon(n)$. Let $\mathcal{A}'$ be the algorithm that applies the conversion algorithm of Lemma 2.8 to a given state χ (which is either $\rho_\pi^{+\otimes k}$ or $\iota^{\otimes k}$) and then feeds the resulting state χ' (either $\rho_\pi^{-\otimes k}$ or $\iota^{\otimes k}$) to $\mathcal{A}$. It thus follows that $\mathcal{A}'(\rho_\pi^{+\otimes k})=\mathcal{A}(\rho_\pi^{-\otimes k})$ and $\mathcal{A}'(\iota^{\otimes k})=\mathcal{A}(\iota^{\otimes k})$. By the triangle inequality, we have

$$\left|\Pr_{\mathcal{A}}\left[\mathcal{A}\left(\rho_\pi^{+\otimes k}\right)=1\right]-\Pr_{\mathcal{A}}\left[\mathcal{A}\left(\iota^{\otimes k}\right)=1\right]\right|+\left|\Pr_{\mathcal{A}'}\left[\mathcal{A}'\left(\rho_\pi^{+\otimes k}\right)=1\right]-\Pr_{\mathcal{A}'}\left[\mathcal{A}'\left(\iota^{\otimes k}\right)=1\right]\right|>\varepsilon(n)$$

for any parameter $n \in N$. This inequality leads us to either

$$\left|\Pr_{\mathcal{A}}\left[\mathcal{A}\left(\rho_\pi^{+\otimes k}\right)=1\right]-\Pr_{\mathcal{A}}\left[\mathcal{A}\left(\iota^{\otimes k}\right)=1\right]\right|>\frac{\varepsilon(n)}{2}$$

or

$$\left|\Pr_{\mathcal{A}'}\left[\mathcal{A}'\left(\rho_\pi^{+\otimes k}\right)=1\right]-\Pr_{\mathcal{A}'}\left[\mathcal{A}'\left(\iota^{\otimes k}\right)=1\right]\right|>\frac{\varepsilon(n)}{2}.$$

To complete the proof, we design the desired algorithm $\mathcal{B}$ as follows: first choose either $\mathcal{A}$ or $\mathcal{A}'$ at random and then simulate the chosen algorithm. It is easy to verify that $\mathcal{B}$ distinguishes between $\rho_\pi^{+\otimes k}$ and $\iota^{\otimes k}$ with the advantage of at least $\varepsilon(n)/4$. □

3. An Application to a Quantum Public-Key Cryptosystem

Section 2 has shown the three useful cryptographic properties of $\mathrm{QSCD}_{\mathrm{ff}}$. Founded on these properties, we wish to construct a quantum PKC whose security is guaranteed by the computational hardness of $\mathrm{QSCD}_{\mathrm{ff}}$ (which can be further reduced to the hardness of GA). As the first step, we give an efficient quantum algorithm that generates ρ_π^+ from π.

Lemma 3.1 (ρ_π^+-Generation Algorithm). *There exists a polynomial-time quantum algorithm that, on input $\pi \in \mathcal{K}_n$, generates the quantum state ρ_π^+ with probability* 1.

Proof. The desired generation algorithm, which is given below, uses two registers. Here, we omit the proof of the correctness of the given algorithm because the correctness is obvious from the description of the algorithm.

(G1) Prepare the state $|0\rangle|id\rangle$ in two quantum registers.
(G2) Apply the Hadamard transformation to the first register to obtain the state $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)|id\rangle$.
(G3) Perform the Controlled-π on the both registers, and we then obtain the state $\frac{1}{\sqrt{2}}(|0\rangle|id\rangle + |1\rangle|\pi\rangle)$.
(G4) Subtract 1 from the content of the first register only when the second register contains π. This process gives rise to the state $\frac{1}{\sqrt{2}}(|0\rangle|id\rangle + |0\rangle|\pi\rangle)$.
(G5) Apply a uniformly random permutation σ to the content of the second register from the left. The whole quantum system then becomes $\frac{1}{\sqrt{2}}(|0\rangle|\sigma\rangle + |0\rangle|\sigma\pi\rangle)$.
(G6) Output the content of the second register, which produces the state ρ_π^+ with probability 1.

□

Hereafter, we describe our quantum PKC and then give its security proof. For the security proof, in particular, we need to clarify our model of adversary's attack. Of all attack models discussed in [7], we use a quantum analogue of *the indistinguishability against the chosen plaintext attack* (*IND-CPA*). Our scenario is precisely as follows:

> Suppose that large-scale quantum and classical networks connect a unique network administrator, acting as a trusted third party, and numerous "ordinary" network users, some of who might possibly be malicious against other users. These parties are all capable of running polynomial-time quantum algorithms. In particular, the administrator (say, Charlie) can communicate with each network user via a secure, authenticated communication channel; namely, he can deliver to each individual user a piece of information (both quantum and classical bits) correctly and securely through this channel. It is most likely that a financial reason could force ordinary users to rely on cheap but insecure channels for daily person-to-person communication with other users. From such an insecure channel, a malicious party (say, Eve) might wiretap the communication . To ensure user's secure communication, upon a request from a user (say, Bob) who wants to receive a message from other users, Charlie generates a decryption (or private) key π and sends it through the secure channel to Bob. Charlie also generates an encryption (or public) key ρ_π^+ for anyone who wants to communicate with Bob.
>
> Now, suppose that an honest party, called Alice, wishes to send Bob a classical single-bit message securely. For this purpose, she first requests Charlie for Bob's encryption key ρ_π^+. Using this key, she encrypts her secret message into a quantum state ρ (either ρ_π^+ or ρ_π^-) as a ciphertext and then sends it to Bob through an available insecure quantum channel. To eavesdrop Alice's secret message, Eve intercepts Alice's ciphertext ρ. In addition, since Eve is also a legitimate network user, she can request numerous

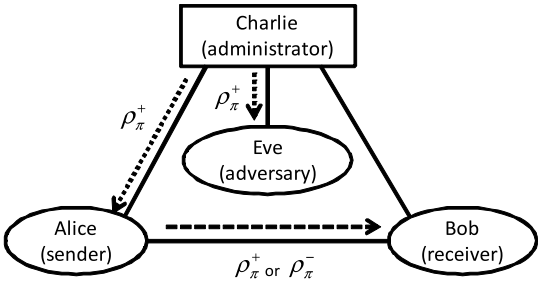

Fig. 1. Our public-key cryptosystem.

copies of the encryption key ρ_π^+ from Charlie (within a polynomial amount of time). Finally, Eve attempts to learn the information involved with Alice's secret message by applying a certain polynomial-time quantum algorithm to the ciphertext ρ as well as a polynomially many copies of the encryption key ρ_π^+ obtained from Charlie as supplemental information.

In the case of classical chosen plaintext attack, all that Eve can collect are Alice's ciphertext and Bob's encryption key. Our scenario is a natural generalization of this classical case because Eve obtains only a quantum state representing Alice's encrypted message and copies of a quantum state serving as an encryption key.

Our scenario demands that the administrator should generate and distribute user's private and public keys. In a practical framework of classical PKCs, such a scenario has been frequently used; for example, a governmental agency may be authorized as a third party to handle those user's keys. Note that Charlie's distribution of decryption keys is done through the secure channel only once at the key setup. With their own single decryption keys, all the users can transmit their messages securely to others a reasonably large number of times, even without any extra secret information shared among them. To the contrary, SKCs require the users to share symmetric secret keys between every pair of them. Thus, even under this scenario, we can enjoy advantages of PKCs over SKCs that stem from the asymmetry of keys in many-to-many communication.

Now, we explain our quantum PKC protocol in detail. In our protocol, Alice transmits a single-bit message to Bob using an $O(n \log n)$-qubit-long encryption key. Our protocol consists of three phases: key setup phase, key transmission phase, and message transmission phase. Figure 1 illustrates our protocol.

The following is the step-by-step description of our quantum PKC protocol.

[Key setup phase]

(A1) Charlie generates Bob's decryption key π uniformly at random from $\mathcal{K}_n$, and then sends it to Bob via a secure and authenticated channel.

[Key transmission phase]

(A2) Alice requests Bob's encryption key from Charlie.
(A3) Using π, Charlie generates a copy of the encryption key ρ_π^+.
(A4) Alice obtains a copy of the encryption key ρ_π^+ from Charlie.

[Message transmission phase]

(A5) Alice encrypts 0 or 1 into ρ_π^+ or ρ_π^-, respectively, and then sends this encrypted message to Bob.
(A6) Bob decrypts Alice's message using the decryption key π.

Step (A1) can be implemented as follows. Recall that $\pi \in \mathcal{K}_n$ consists of $n/2$ disjoint transpositions. We first choose distinct two numbers i_1 and i_2 from $\{1, 2, \ldots, n\}$ uniformly at random, and make a transposition $(i_1\ i_2)$. Next, choosing other distinct two numbers i_3 and i_4 from $\{1, 2, \ldots, n\} \setminus \{i_1, i_2\}$ uniformly at random, we make another transposition $(i_3\ i_4)$. By repeating this process, $n/2$ disjoint transpositions are chosen uniformly at random. From them, define $\pi = (i_1, i_2) \cdots (i_{n/2-1}, i_{n/2})$. Step (A3) is done by the ρ_π^+-generation algorithm of Lemma 3.1. The conversion algorithm of Lemma 2.8 implements Step (A5) since Alice sends Bob either the received state ρ_π^+ or its converted state ρ_π^-. Finally, the distinguishing algorithm of Theorem 2.1 implements Step (A6).

The security proof of our PKC is done by reducing GA to Eve's attacking strategy during the message transmission phase. Our reduction is a simple modification of the reduction given in Theorem 2.5.

Proposition 3.2. *Let $\mathcal{A}$ be any polynomial-time quantum adversary who attacks our quantum PKC during the message transmission phase. Assume that there exist two polynomials $p(n)$ and $l(n)$ satisfying that*

$$\left| \Pr_{\pi,\mathcal{A}}\left[\mathcal{A}\left(\rho_\pi^+, \rho_\pi^{+\otimes l(n)}\right) = 1\right] - \Pr_{\pi,\mathcal{A}}\left[\mathcal{A}\left(\rho_\pi^-, \rho_\pi^{+\otimes l(n)}\right) = 1\right] \right| > \frac{1}{p(n)}$$

for infinitely many parameters $n \in N$. Then, there exists a polynomial-time quantum algorithm that solves GA for infinitely many input sizes n in the worst case with non-negligible probability.

Proof. The proposition immediately follows from the proof of Theorem 2.5 by replacing $\rho_\pi^{+\otimes k}$, $\rho_\pi^{-\otimes k}$, and $\iota^{\otimes k}$ in the proof with $(\rho_\pi^+, \rho_\pi^{+\otimes l(n)})$, $(\rho_\pi^-, \rho_\pi^{+\otimes l(n)})$, and $(\iota, \iota^{\otimes l(n)})$, respectively. □

4. A Generalization of $\mathrm{QSCD}_{\mathrm{ff}}$

In our $\mathrm{QSCD}_{\mathrm{ff}}$-based quantum PKC, Alice encrypts a single-bit message using an $O(n \log n)$-qubit encryption key. We wish to show how to increase the size of Alice's encryption message and construct a multi-bit quantum PKC built upon a generalization of $\mathrm{QSCD}_{\mathrm{ff}}$, called $\mathrm{QSCD}_{\mathrm{cyc}}$ (QSCD with cyclic permutations), which is a distinction problem among *multiple ensembles* of quantum states. Recall that Definition 1.1 has introduced the notion of computational indistinguishability between two ensembles of quantum states. This notion can be naturally generalized as follows to multiple quantum state ensembles.

Definition 4.1. We say that m ensembles $\{\rho_0(l)\}_{l\in\mathbb{N}}, \ldots, \{\rho_{m-1}(l)\}_{l\in\mathbb{N}}$ of quantum states are *computationally indistinguishable* if, for any distinct pair $i, j \in \mathbb{Z}_m$, the advantage of distinguishing between the two ensembles $\{\rho_i(l)\}_{l\in\mathbb{N}}$ and $\{\rho_j(l)\}_{l\in\mathbb{N}}$ is negligible for any polynomial-time quantum algorithm $\mathcal{A}$; namely, for any two ensembles $\{\rho_i(l)\}_{l\in\mathbb{N}}$ and $\{\rho_j(l)\}_{l\in\mathbb{N}}$, any polynomial p, any polynomial-time quantum algorithm $\mathcal{A}$, and any sufficiently large number l, it holds that

$$\left|\Pr_{\mathcal{A}}\left[\mathcal{A}\left(\rho_i(l)\right)=1\right]-\Pr_{\mathcal{A}}\left[\mathcal{A}\left(\rho_j(l)\right)=1\right]\right|<\frac{1}{p(l)}.$$

The distinction problem among the ensembles $\{\rho_0(l)\}_{l\in\mathbb{N}}, \ldots, \{\rho_{m-1}(l)\}_{l\in\mathbb{N}}$ is said to be *solvable with non-negligible advantage* if the ensembles are not computationally indistinguishable, that is, there exist two ensembles $\{\rho_i(l)\}_{l\in\mathbb{N}}$ and $\{\rho_j(l)\}_{l\in\mathbb{N}}$, a polynomial-time quantum algorithm $\mathcal{A}$, and a polynomial p such that

$$\left|\Pr_{\mathcal{A}}\left[\mathcal{A}\left(\rho_i(l)\right)=1\right]-\Pr_{\mathcal{A}}\left[\mathcal{A}\left(\rho_j(l)\right)=1\right]\right|>\frac{1}{p(l)}$$

for infinitely many numbers $l \in \mathbb{N}$.

We wish to define a specific distinction problem, denoted succinctly $\mathrm{QSCD}_{\mathrm{cyc}}$, among m ensembles of quantum states. First, we define a new hidden permutation, which will be encoded into certain quantum states. For any fixed number $n \in \mathbb{N}$, let us assume that $m \geq 2$ and m divides n. The new hidden permutation π consists of disjoint n/m cyclic permutations of length m; in other words, π is of the form

$$\pi = (i_0\, i_1 \cdots i_{m-1}) \cdots (i_{n-m}\, i_{n-m+1} \cdots i_{n-1}),$$

where $i_0, \ldots, i_{n-1} \in \mathbb{Z}_n$ and $i_s \neq i_t$ if $s \neq t$ for any pair (s, t). Such a permutation π has the following two properties: (i) π has no fixed points (i.e., $\pi(i) \neq i$ for any $i \in \mathbb{Z}_n$) and (ii) π is of order m (i.e., $\pi^m = id$). For convenience, we denote by $\mathcal{K}_n^m$ ($\subseteq S_n$) the set of all such permutations.

With a help of the hidden permutation π, we can define the new quantum states $|\Phi_{\pi,s}^{\sigma}\rangle$ as follows. For each $\sigma \in S_n$, $\pi \in \mathcal{K}_n^m$, and $s \in \mathbb{Z}_m$, let

$$\left|\Phi_{\pi,s}^{\sigma}\right\rangle=\frac{1}{\sqrt{m}}\sum_{t=0}^{m-1}\omega_m^{st}\left|\sigma\pi^t\right\rangle,$$

where $\omega_m = e^{2\pi i/m}$. At last, the distinction problem $\mathrm{QSCD}_{\mathrm{cyc}}$ is defined in the following way.

Definition 4.2. The problem $\mathrm{QSCD}_{\mathrm{cyc}}$ is a distinction problem among m ensembles $\{\rho_\pi^{(0)}(n)^{\otimes k(n)}\}_{n\in\mathbb{N}}, \ldots, \{\rho_\pi^{(m-1)}(n)^{\otimes k(n)}\}_{n\in\mathbb{N}}$ of quantum states, where k is an arbitrary polynomial and the notation $\rho_\pi^{(s)}(n)$ denotes the mixed state $\frac{1}{n!}\sum_{\sigma\in S_n}|\Phi_{\pi,s}^{\sigma}\rangle\langle\Phi_{\pi,s}^{\sigma}|$ for each $\pi \in \mathcal{K}_n^m$. When k is fixed, we use the notation k-$\mathrm{QSCD}_{\mathrm{cyc}}$ instead.

Similar to the case of $\mathrm{QSCD}_{\mathrm{ff}}$, we also drop the parameter n wherever possible. Note that $\mathrm{QSCD}_{\mathrm{ff}}$ coincides with $\mathrm{QSCD}_{\mathrm{cyc}}$ with $m = 2$ and n is of the form $2(2n' + 1)$ for a certain number $n' \in \mathbb{N}$.

This new problem $\mathrm{QSCD}_{\mathrm{cyc}}$ also enjoys useful cryptographic properties. We first present a trapdoor of $\mathrm{QSCD}_{\mathrm{cyc}}$. In the case of $\mathrm{QSCD}_{\mathrm{ff}}$, because its trapdoor information π is a permutation of order two, we encode only a single bit into the both quantum states ρ_π^+ and ρ_π^-. On the contrary, since $\mathrm{QSCD}_{\mathrm{cyc}}$ uses a permutation π of order $m \geq 2$, it is possible to encode $\log m$ bits into the m quantum states $\rho_\pi^{(0)}, \ldots, \rho_\pi^{(m-1)}$.

Now, we present a generalized distinguishing algorithm working for $\rho_\pi^{(s)}$'s.

Theorem 4.3 (Generalized Distinguishing Algorithm). *There exists a polynomial-time quantum algorithm that, for each $n \in \mathbb{N}$, $\pi \in \mathcal{K}_n^m$, and $s \in \mathbb{Z}_m$, decrypts $\rho_\pi^{(s)}(n)$ to s with exponentially-small error probability.*

Proof. Let χ be any given quantum state of the form $\rho_\pi^{(s)}$ for a certain hidden permutation $\pi \in \mathcal{K}_n^m$ and also a certain hidden parameter s. Note that χ is a mixture of all pure states $|\Phi_{\pi,s}^\sigma\rangle$ over a randomly chosen $\sigma \in S_n$. It thus suffices to give a polynomial-time quantum algorithm that decrypts $|\Phi_{\pi,s}^\sigma\rangle$ to s for each fixed σ. Such an algorithm can be given by conducting the following *Generalized Controlled-π Test*, which is a straightforward generalization of the distinguishing algorithm given in the proof of Theorem 2.1. To define this test, we first recall the quantum Fourier transformation F_m over $\mathbb{Z}_m$ as well as its inverse F_m^{-1}: for any $x \in \mathbb{Z}_m$,

$$F_m|x\rangle = \frac{1}{\sqrt{m}} \sum_{y \in \mathbb{Z}_m} \omega_m^{xy}|y\rangle \quad \text{and} \quad F_m^{-1}|x\rangle = \frac{1}{\sqrt{m}} \sum_{y \in \mathbb{Z}_m} \omega_m^{-xy}|y\rangle.$$

The Generalized Controlled-π Test is described below.

[Generalized Controlled-π Test]

(D1′) Prepare two quantum registers. The first register holds a control string, initially set to $|0\rangle$, and the second register holds the quantum state $|\Phi_{\pi,s}^\sigma\rangle$. Apply the inverse Fourier transformation F_m^{-1} to the first register. Meanwhile, assume that we can perform the Fourier transformation exactly. The entire system then becomes

$$\frac{1}{\sqrt{m}} \sum_{r=0}^{m-1} |r\rangle \left|\Phi_{\pi,s}^\sigma\right\rangle = \frac{1}{m} \sum_{r,t} \omega_m^{st}|r\rangle \left|\sigma\pi^t\right\rangle.$$

(D2′) Apply π to the content of the second register r times from the right. The state of the entire system evolves into

$$\frac{1}{m} \sum_{r,t} \omega_m^{st}|r\rangle \left|\sigma\pi^{r+t \bmod m}\right\rangle.$$

(D3′) Apply the Fourier transformation F_m to the first register and we then obtain the state

$$\begin{aligned}
&\frac{1}{m}\sum_{r,t}\frac{1}{\sqrt{m}}\sum_{r'=0}^{m-1}\omega_m^{rr'}\left|r'\right\rangle\omega_m^{st}\left|\sigma\pi^{r+t \bmod m}\right\rangle \\
&\quad=\frac{1}{m^{3/2}}\sum_{r,r',t}\omega_m^{st+rr'}\left|r'\right\rangle\left|\sigma\pi^{r+t \bmod m}\right\rangle \\
&\quad=\frac{1}{m^{3/2}}\sum_{r,t}\omega_m^{s(r+t)}|s\rangle\left|\sigma\pi^{r+t \bmod m}\right\rangle \\
&\qquad+\frac{1}{m^{3/2}}\sum_{r,t,r'\neq s}\omega_m^{st+rr'}\left|r'\right\rangle\left|\sigma\pi^{r+t \bmod m}\right\rangle \\
&\quad=\frac{1}{\sqrt{m}}\sum_{u}\omega_m^{su}|s\rangle\left|\sigma\pi^{u}\right\rangle+\frac{1}{m^{3/2}}\sum_{r,u,r'\neq s}\omega_m^{su+r(r'-s)}\left|r'\right\rangle\left|\sigma\pi^{u}\right\rangle \\
&\quad(u:=r+t \bmod m) \\
&\quad=\frac{1}{\sqrt{m}}\sum_{u=0}^{m-1}\omega_m^{su}|s\rangle\left|\sigma\pi^{u}\right\rangle=|s\rangle\left|\Phi_{\pi,s}^{\sigma}\right\rangle \\
&\quad\left(\text{since }\sum_{r}\omega_m^{su+r(r'-s)}=0\text{ for any }u,s,r'(\neq s)\right).
\end{aligned}$$

(D4′) Finally, measure the first register in the computational basis and output the measured result s in $\mathbb{Z}_m$.

The error probability of the above algorithm depends only on the precision of the Fourier transformation over $\mathbb{Z}_m$. As shown in [33], the quantum Fourier transformation can be implemented with exponentially-small error probability by an application of the approximated quantum Fourier transformation. Therefore, the theorem follows. □

Similar to $\mathrm{QSCD}_{\mathrm{ff}}$, the average-case hardness of $\mathrm{QSCD}_{\mathrm{cyc}}$ coincides with its worst-case hardness.

Theorem 4.4. *Let k be any polynomial. Assume that there exists a polynomial-time quantum algorithm $\mathcal{A}$ that solves k-$\mathrm{QSCD}_{\mathrm{cyc}}$ with non-negligible advantage for a uniformly random permutation $\pi\in\mathcal{K}_n^m$; namely, there exist two numbers $s,s'\in\mathbb{Z}_m$ and a polynomial p such that, for infinitely many numbers $n\in\mathbb{N}$,*

$$\left|\Pr_{\pi,\mathcal{A}}\left[\mathcal{A}\left(\rho_\pi^{(s)}(n)^{\otimes k(n)}\right)=1\right]-\Pr_{\pi,\mathcal{A}}\left[\mathcal{A}\left(\rho_\pi^{(s')}(n)^{\otimes k(n)}\right)=1\right]\right|>\frac{1}{p(n)},$$

where π is chosen uniformly at random from $\mathcal{K}_n^m$. Then, there exists a polynomial-time quantum algorithm $\mathcal{B}$ that solves k-$\mathrm{QSCD}_{\mathrm{cyc}}$ with non-negligible advantage.

Proof. This proof follows an argument in the proof of Theorem 2.2. Here, we give only a sketch of our desired algorithm $\mathcal{B}$. Choose a uniformly random permutation

$\tau \in S_n$ and then apply it to $|\Phi^{\sigma}_{\pi,s}\rangle$ from the right. Now, we obtain the state

$$\frac{1}{\sqrt{m}}\sum_{t=0}^{m-1}\omega_m^{st}\big|\sigma\pi^t\tau\big\rangle=\frac{1}{\sqrt{m}}\sum_{t=0}^{m-1}\omega_m^{st}\big|\sigma\tau\tau^{-1}\pi^t\tau\big\rangle=\frac{1}{\sqrt{m}}\sum_{t=0}^{m-1}\omega_m^{st}\big|\sigma\tau\big(\tau^{-1}\pi\tau\big)^t\big\rangle.$$

Note that $\rho^{(s)}_{\tau^{-1}\pi\tau}(n)=\frac{1}{n!}\sum_{\sigma\in S_n}|\Phi^{\sigma\tau}_{\tau^{-1}\pi\tau,s}\rangle\langle\Phi^{\sigma\tau}_{\tau^{-1}\pi\tau,s}|$ is an average-case instance of $\mathrm{QSCD}_{\mathrm{cyc}}$ since $\tau^{-1}\pi\tau$ is distributed uniformly at random over $\mathcal{K}_n^m$. Finally, apply the average-case algorithm $\mathcal{A}$. □

We will exhibit a quantum algorithm that generates the quantum state $\rho_\pi^{(s)}$ efficiently from π and s. This generation algorithm will be used to generate encryption keys in our $\mathrm{QSCD}_{\mathrm{cyc}}$-based multi-bit quantum PKC.

Lemma 4.5 ($\rho_\pi^{(s)}$-Generation Algorithm). *There exists a polynomial-time quantum algorithm that generates $\rho_\pi^{(s)}$ for any $s\in\mathbb{Z}_m$ and any $\pi\in\mathcal{K}_n^m$ with exponentially-small error probability.*

Proof. The desired algorithm is a straightforward generalization of the ρ_π^+-generation algorithm given in the proof of Lemma 3.1. Using the approximated Fourier transformation [33] instead of the Hadamard transformation, we can efficiently approximate from π the Fourier transformation F_π over the cyclic group $\{id,\pi,\pi^2,\ldots,\pi^{m-1}\}$, that is,

$$F_\pi\big|\pi^s\big\rangle=\frac{1}{\sqrt{m}}\sum_{t=0}^{m-1}\omega_m^{st}\big|\pi^t\big\rangle,$$

by employing an argument similar to the proof of Lemma 3.1. Hence, we can perform F_π on $|\pi^s\rangle$ with exponentially-small error probability.

Since the initial state $|\pi^s\rangle$ can be easily generated from π, we immediately obtain an efficient approximation of $F_\pi|\pi^s\rangle$. By applying a uniformly-random permutation $\sigma\in S_n$ to the resulting state from the left, the desired state $\rho_\pi^{(s)}$ can be obtained with exponentially-small error probability. □

Toward the end of this section, we present our multi-bit quantum PKC, based on $\mathrm{QSCD}_{\mathrm{cyc}}$.

[Key setup phase]

(A1′) As Bob's decryption key, Charlie chooses an element π uniformly at random from $\mathcal{K}_n$ and then sends it to Bob via a secure, authenticated channel.

[Key transmission phase]

(A2′) Alice requests Bob's encryption key from Charlie.

(A3′) Charlie generates a copy of the encryption key $(\rho_\pi^{(0)},\ldots,\rho_\pi^{(m-1)})$ from π and sends it to Alice.

(A4′) Alice receives this copy of the encryption key from Charlie.

[Message transmission phase]

(A5′) If her message is $s \in \mathbb{Z}_m$, Alice picks up $\rho_\pi^{(s)}$. She sends it to Bob as a ciphertext.

(A6′) Bob decrypts Alice's message using the decryption key π.

By choosing cycles one by one sequentially, we can perform Step (A1′). The $\rho_\pi^{(s)}$-generation algorithm of Lemma 4.5 immediately implements Step (A3′). Alice can encrypt her message s simply by choosing $\rho_\pi^{(s)}$ out of the series $(\rho_\pi^{(0)}, \ldots, \rho_\pi^{(m-1)})$. Finally, the generalized distinguishing algorithm in Theorem 4.3 achieves Step (A6′).

As the final remark, we refer to a drawback of the above multi-bit encryption scheme. A major drawback is that Charlie should send Alice all the series $(\rho_\pi^{(0)}, \ldots, \rho_\pi^{(m-1)})$ as Bob's encryption key simply because of the lack of a sophisticated converting algorithm among different encryption keys without knowing the hidden decryption key π. This $\mathrm{QSCD}_{\mathrm{cyc}}$-based encryption scheme requires an $O(mn \log n)$-qubit encryption key to encrypt a $\log m$-bit message whereas the $\mathrm{QSCD}_{\mathrm{ff}}$-based encryption scheme needs an $O(n \log n)$-qubit key per a 1-bit message. In a quick comparison, there seems to be no advantage of the $\mathrm{QSCD}_{\mathrm{cyc}}$-based scheme over the $\mathrm{QSCD}_{\mathrm{ff}}$-based scheme in terms of the ratio between message length and encryption key length.

This drawback stems from the conversion algorithm, given in Lemma 2.8, used to swap ρ_π^+ and ρ_π^- in the $\mathrm{QSCD}_{\mathrm{ff}}$-based single-bit encryption scheme. This conversion algorithm utilizes the "parity" of permutations σ and $\sigma\pi$ to invert their phases without using any information on π. More precisely, the algorithm implements the homomorphism f from S_n to $\{+1, -1\}$ ($\cong \mathbb{Z}/2\mathbb{Z}$) satisfying that $f(\sigma) = +1$ (-1, resp.) if σ is even (odd, resp.). Unfortunately, the same algorithm *fails* for $\mathrm{QSCD}_{\mathrm{cyc}}$ because no homomorphism maps S_n to $\{1, \omega_m, \ldots, \omega_m^{m-1}\}$ ($\cong \mathbb{Z}/m\mathbb{Z}$). This is shown as follows. Let us assume, to the contrary, that there exists a homomorphism g mapping S_n to $\{1, \omega_m, \ldots, \omega_m^{m-1}\}$. The *fundamental homomorphism theorem* implies that $S_n / \mathrm{Ker}(g) \cong \mathbb{Z}/m\mathbb{Z}$; namely, there exists an isomorphism from $\sigma\,\mathrm{Ker}(g)$ to $g(\sigma)$ for every $\sigma \in S_n$. Note that $\mathrm{Ker}(g)$ is a normal subgroup in S_n. It is known that such a normal subgroup in S_n equals either the trivial group $\{id\}$ or the alternation group $A_n = \{\sigma \in S_n : \mathrm{sgn}(\sigma) = 0\}$ since A_n is a simple group for $n \geq 5$ (see, e.g., Theorem 3.2.1 in [53]). Apparently, there is neither isomorphism between $\{\sigma A_n : \sigma \in S_n\}$ and $\mathbb{Z}/m\mathbb{Z}$ nor isomorphism between $\{\sigma : \sigma \in S_n\}$ and $\mathbb{Z}/m\mathbb{Z}$ if $n > 4$ and $n \geq m > 2$. This contradicts our assumption on g.

5. Concluding Remarks

We have shown that the computational distinction problem $\mathrm{QSCD}_{\mathrm{ff}}$ satisfies quite useful cryptographic properties which help us to design a quantum PKC whose security is guaranteed by the computational intractability of GA. Although GA is reducible to $\mathrm{QSCD}_{\mathrm{ff}}$ in polynomial time, there seems to be a large gap between the hardness of GA and that of $\mathrm{QSCD}_{\mathrm{ff}}$ because, in the proof of Theorem 2.5, all combinatorial structures of an input graph for GA are completely lost in constructing associated quantum states for $\mathrm{QSCD}_{\mathrm{ff}}$ and, from such states, it is impossible to recover the original graph. It is therefore pressing to find a much better classical problem (for instance, the problems of finding a centralizer or finding a normalizer [38]) that almost matches the compu-

tational hardness of $\mathrm{QSCD}_{\mathrm{ff}}$. Since no fast quantum algorithm is known for $\mathrm{QSCD}_{\mathrm{ff}}$, discovering such a fast algorithm for $\mathrm{QSCD}_{\mathrm{ff}}$ may require new tools and novel proof techniques in quantum complexity theory. Besides our quantum states $\{\rho_\pi^+(n), \rho_\pi^-(n)\}$ used in $\mathrm{QSCD}_{\mathrm{ff}}$, it is imperative to continue searching for other pairs of "simple" quantum states whose computational indistinguishability is helpful to construct a more secure cryptosystem.

Similar to $\mathrm{QSCD}_{\mathrm{ff}}$, $\mathrm{QSCD}_{\mathrm{cyc}}$ also owns useful cryptographic properties for which we have built a multi-bit quantum PKC. Throughout our study, it is not yet clear how difficult $\mathrm{QSCD}_{\mathrm{cyc}}$ is and how secure our multi-bit quantum PKC truly is. If one successfully proves that the worst-case hardness of $\mathrm{QSCD}_{\mathrm{cyc}}$ is lower-bounded by, e.g., the hardness of GA, then our multi-bit quantum PKC might find a more practical use in return.

Acknowledgements

The authors are grateful to Hirotada Kobayashi and Claude Crépeau for fruitful discussions, to John Watrous for useful comments on key ideas, to Donald Beaver, Louis Salvail, and the anonymous reviewers of EUROCRYPT 2005 and Journal of Cryptology for their valuable suggestions. The authors' thanks also go to Cristopher Moore for providing references to a historical account of hidden subgroup problems. This research was partially supported by Grant-in-Aid for Young Scientists (B) No. 17700007 (2005), Grant-in-Aid for Scientific Research on Priority Areas No. 16092206 (2005), No. 18300002 (2006), and No. 21300002 (2009) from the Ministry of Education, Science, Sports and Culture.

Appendix A. A Reduction from GA to $\mathrm{UniqueGA}_{\mathrm{ff}}$

In this appendix, we prove Lemma 2.6, in which $\mathrm{UniqueGA}_{\mathrm{ff}}$ is shown to be polynomial-time Turing equivalent to GA. Earlier, Köbler, Schöning, and Torán [35] established the polynomial-time Turing equivalence between GA and UniqueGA. We first review their reduction and then explain how to modify it to obtain the desired reduction from GA to $\mathrm{UniqueGA}_{\mathrm{ff}}$. Note that the reduction from $\mathrm{UniqueGA}_{\mathrm{ff}}$ to GA is trivial since $\mathrm{UniqueGA}_{\mathrm{ff}}$ is simply a special case of GA.

We begin with explaining our technical tool and notation necessary to describe the reduction of [35]. Their reduction uses a technical tool called a *label* to distinguish each node of a given graph G from the others. Given a graph G, let n be the number of nodes in G. The label j attached to node i consists of two chains: one of which is of length $2n + 3$ connected to node i, and the other is of length j connected to the $n + 2$-nd node of the first chain (see Fig. 2). Note that the total size of the label j is

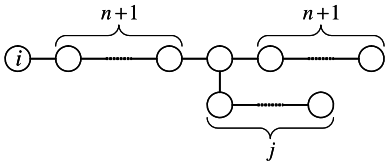

Fig. 2. Label.

$2n + j + 3$. Let $G_{[i]}$ denote the graph obtained from G by attaching the label 1 to the node i. Similarly, $G_{[i_1,\ldots,i_j]}$ is defined as the graph with labels $1,\ldots,j$ respectively attached to nodes $i_1,\ldots,i_j$. Note that any automorphism of $G_{[i]}$ maps the node i into itself and that any label adds no new automorphism into this modified graph. Let $Aut(G)$ be the automorphism group of G and let $Aut(G)_{[1,\ldots,i]}$ be the point-wise stabilizer of $\{1,\ldots,i\}$ in $Aut(G)$, namely, $Aut(G)_{[1,\ldots,i]} = \{\sigma \in Aut(G) : \forall j \in \{1,\ldots,i\}[\sigma(j) = j]\}$.

The following theorem was proven in [35]. For our later reference, we include its proof here.

Theorem A.1 [35, Theorem 1.31]. *GA is polynomial-time Turing reducible to UniqueGA.*

Proof. Let $\mathcal{O}$ be any set that correctly represents UniqueGA on all promised instances. Using $\mathcal{O}$ as an oracle, the following algorithm solves GA in polynomial time. Let G be any given instance of GA.

(U1) Repeat (U2)–(U3) for each i starting with $n-1$ down to 1.
(U2) Repeat (U3) for each j ranging from $i+1$ to n.
(U3) Invoke $\mathcal{O}$ with input graph $G_{[1,\ldots,i-1,i]} \cup G_{[1,\ldots,i-1,j]}$. If the outcome of $\mathcal{O}$ is YES, output YES and halt.
(U4) Output NO.

If G is a "YES" instance, there is at least one non-trivial automorphism. Take the largest number $i \in \{1,\ldots,n\}$ such that there exist a number $j \in \{1,\ldots,n\}$ and a non-trivial automorphism $\pi \in Aut(G)_{[1,\ldots,i-1]}$ for which $\pi(i) = j$ and $i \neq j$. We want to claim that there is exactly one such non-trivial automorphism, i.e., $Aut(G)_{[1,\ldots,i-1]} = \{id, \pi\}$. This is seen as follows. First, note that $Aut(G)_{[1,\ldots,i-1]}$ is expressed as $Aut(G)_{[1,\ldots,i-1]} = \pi_1 Aut(G)_{[1,\ldots,i]} + \cdots + \pi_d Aut(G)_{[1,\ldots,i]}$. For any two distinct cosets $\pi_s Aut(G)_{[1,\ldots,i]}$ and $\pi_t Aut(G)_{[1,\ldots,i]}$ and for any two automorphisms $\sigma \in \pi_s Aut(G)_{[1,\ldots,i]}$ and $\sigma' \in \pi_t Aut(G)_{[1,\ldots,i]}$, it holds that $\sigma(i) \neq \sigma'(i)$. Since $Aut(G)_{[1,\ldots,i]} = \{id\}$ by the definition of i, we obtain $|\pi_k Aut(G)_{[1,\ldots,i]}| = 1$ for any coset $\pi_k Aut(G)_{[1,\ldots,i]}$. Furthermore, there exists the unique coset $\pi Aut(G)_{[1,\ldots,i]}$ satisfying that $\sigma(i) = j$ for any $\sigma \in \pi Aut(G)_{[1,\ldots,i]}$. These facts imply that the non-trivial automorphism π is unique. Note that the unique non-trivial automorphism interchanges two subgraphs $G_{[1,\ldots,i-1,i]}$ and $G_{[1,\ldots,i-1,j]}$. Therefore, the above algorithm successfully outputs YES at Step (U3).

On the contrary, if G is a "NO" instance, then for every distinct i and j, the modified graph has no non-trivial automorphism. Thus, the above algorithm correctly rejects G. □

Finally, we describe the reduction from GA to $\text{UniqueGA}_{\text{ff}}$ by slightly modifying the reduction given in the above proof.

Lemma A.2. *GA is polynomial-time Turing reducible to $\text{UniqueGA}_{\text{ff}}$.*

Proof. Recall the algorithm given in the proof of Theorem A.1. We only need to change the number of nodes to invoke oracle $\text{UniqueGA}_{\text{ff}}$ in (U3). To make such a

change, we first modify the size of each label. Since the number m of all nodes of $G_{[1,\ldots,i-1,i]} \cup G_{[1,\ldots,i-1,j]}$ is even, if there is no k such that $m = 2(2k+1)$, then we add one more node appropriately to the original labels. We then attach our modified labels of length $2n+i+4$ and $2n+j+4$ to the nodes i and j, respectively. Obviously, this modified graph satisfies the promise of $\mathrm{UniqueGA}_{\mathrm{ff}}$. Our algorithm therefore works correctly for any instance of GA. □

References

[1] M. Adcock, R. Cleve, A quantum Goldreich-Levin theorem with cryptographic applications, in *Proceedings of the 19th Annual Symposium on Theoretical Aspects of Computer Science*. LNCS, vol. 2285 (Springer, Berlin, 2002), pp. 323–334

[2] D. Aharonov, A. Ta-Shma, Adiabatic quantum state generation. *SIAM J. Comput.* **37**(1), 47–82 (2007)

[3] M. Ajtai, Generating hard instances of lattice problems (extended abstract), in *Proceedings of the 28th Annual ACM Symposium on Theory of Computing* (1996), pp. 99–108

[4] M. Ajtai, C. Dwork, A public-key cryptosystem with worst-case/average-case equivalence, in *Proceedings of the 29th Annual ACM Symposium on Theory of Computing* (1997), pp. 284–293. See also ECCC TR96-065

[5] V. Arvind, P.P. Kurur, Graph isomorphism is in SPP. *Inf. Comput.* **204**(5), 835–852 (2006)

[6] D. Bacon, A.M. Childs, W. van Dam, From optimal measurement to efficient quantum algorithms for the hidden subgroup problem over semidirect product groups, in *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science* (2005), pp. 469–478

[7] M. Bellare, A. Desai, D. Pointcheval, P. Rogaway, Relations among notions of security for public-key encryption schemes, in *Advances in Cryptology—CRYPTO'98* (Springer, Berlin, 1998), pp. 26–45

[8] C.H. Bennett, G. Brassard, Quantum cryptography: public key distribution and coin tossing, in *Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing* (1984), pp. 175–179

[9] M. Blum, S. Micali, How to generate cryptographically strong sequences of pseudo-random bits. *SIAM J. Comput.* **13**(4), 850–864 (1984)

[10] A. Bogdanov, L. Trevisan, On worst-case to average-case reductions for NP problems. *SIAM J. Comput.* **36**(4), 1119–1159 (2006)

[11] D. Boneh, R.J. Lipton, Quantum cryptanalysis of hidden linear functions (extended abstract), in *Advances in Cryptology—CRYPTO'95*. LNCS, vol. 963 (Springer, Berlin, 1995), pp. 424–437

[12] M. Crâsmaru, C. Glaßer, K.W. Regan, S. Sengupta, A protocol for serializing unique strategies, in *Proceedings of the 29th International Symposium on Mathematical Foundations of Computer Science*. LNCS, vol. 3153 (Springer, Berlin, 2004), pp. 660–672

[13] C. Crépeau, P. Dumais, D. Mayers, L. Salvail, Computational collapse of quantum state with application to oblivious transfer, in *Proceedings of the 1st Theory of Cryptography Conference*. LNCS, vol. 2951 (Springer, Berlin, 2004), pp. 374–393

[14] C. Crépeau, F. Légaré, L. Salvail, How to convert the flavor of a quantum bit commitment, in *Advances in Cryptology—EUROCRYPT'01*. LNCS, vol. 2045 (Springer, Berlin, 2001), pp. 60–77

[15] I. Damgård, S. Fehr, L. Salvail, Zero-knowledge proofs and string commitments withstanding quantum attacks, in *Advances in Cryptology—CRYPTO'04*. LNCS, vol. 3152 (Springer, Berlin, 2004), pp. 254–272

[16] S.-P. Desrosier, De la cryptographie sur les corps quadratiques rels. Master's thesis, Université McGill, Montréal, 2002

[17] W. Diffie, M.E. Hellman, New directions in cryptography. *IEEE Trans. Inf. Theory* **IT-22**(6), 644–654 (1976)

[18] P. Dumais, D. Mayers, L. Salvail, Perfectly concealing quantum bit commitment from any quantum one-way permutation, in *Advances in Cryptology—EUROCRYPT 2000*. LNCS, vol. 1807 (Springer, Berlin, 2000), pp. 300–315

[19] M. Ettinger, P. Høyer, On quantum algorithms for noncommutative hidden subgroups. *Adv. Appl. Math.* **25**, 239–251 (2000)

[20] S. Goldwasser, S. Micali, Probabilistic encryption. *J. Comput. Syst. Sci.* **28**(2), 270–299 (1984)
[21] S. Goldwasser, M. Sipser, Private coins versus public coins in interactive proof system, in *Advances in Computing Research*, ed. by S. Micali. Randomness and Computation, vol. 5 (JAI Press, London, 1989), pp. 73–90
[22] M. Grigni, L.J. Schulman, M. Vazirani, U. Vazirani, Quantum mechanical algorithms for the nonabelian hidden subgroup problem. *Combinatorica* **24**(1), 137–154 (2004)
[23] J. Grollmann, A.L. Selman, Complexity measures for public-key cryptosystems. *SIAM J. Comput.* **17**(2), 309–335 (1988)
[24] S. Hallgren, Polynomial-time quantum algorithms for Pell's equation and the principal ideal problem. *Journal of the ACM* **54**(1) (2007)
[25] S. Hallgren, C. Moore, M. Rötteler, A. Russell, P. Sen, Limitations of quantum coset states for graph isomorphism, in *Proceedings of the 38th ACM Symposium on Theory of Computing* (2006), pp. 604–617. See also quant-ph/0511148 and quant-ph/0511149
[26] S. Hallgren, A. Russell, A. Ta-Shma, The hidden subgroup problem and quantum computation using group representations. *SIAM J. Comput.* **32**(4), 916–934 (2003)
[27] M. Hayashi, A. Kawachi, H. Kobayashi, Quantum measurements for hidden subgroup problems with optimal sample complexity. *Quantum Inf. Comput.* **8**, 345–358 (2008)
[28] W. Höffding, Probability inequalities for sums of bounded random variables. *J. Am. Stat. Assoc.* **58**(301), 13–30 (1963)
[29] R. Impagliazzo, M. Naor, Efficient cryptographic schemes provably as secure as subset sum. *J. Cryptol.* **9**(4), 199–216 (1996)
[30] A. Kawachi, T. Koshiba, H. Nishimura, T. Yamakami, Computational indistinguishability between quantum states and its cryptographic application, in *Advances in Cryptology—EUROCRYPT'05*. LNCS, vol. 3494 (Springer, Berlin, 2005), pp. 268–284
[31] A. Kawachi, C. Portmann, On the power of quantum encryption keys, in *Proceedings of the 2nd International Workshop on Post-Quantum Cryptography*. LNCS, vol. 5299 (Springer, Berlin, 2008), pp. 165–180
[32] J. Kempe, A. Shalev, The hidden subgroup problem and permutation group theory, in *Proceedings of the 16th ACM-SIAM Symposium on Discrete Algorithms* (2005), pp. 1118–1125
[33] A. Kitaev, Quantum measurements and the Abelian stabilizer problem (1995). quant-ph/9511026
[34] H. Kobayashi, Non-interactive quantum perfect and statistical zero-knowledge, in *Proceedings of the 14th Annual International Conference on Algorithms and Computation*. LNCS, vol. 2906 (Springer, Berlin, 2003), pp. 178–188
[35] J. Köbler, U. Schöning, J. Torán, *The Graph Isomorphism Problem: Its Structural Complexity* (Birkhäuser Boston, Cambridge, 1993)
[36] G. Kuperberg, A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. *SIAM J. Comput.* **35**(1), 170–188 (2005)
[37] H.-K. Lo, H.F. Chau, Is quantum bit commitment really possible? *Phys. Rev. Lett.* **78**(17), 3410–3413 (1997)
[38] E.M. Luks, Permutation groups and polynomial-time computation, in *Groups and Computation*, ed. by L. Finklestein, W.M. Kantor. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 5 (Am. Math. Soc., Providence, 1993), pp. 139–175
[39] D. Mayers, Unconditionally secure quantum bit commitment is impossible. *Phys. Rev. Lett.* **78**(17), 3414–3417 (1997)
[40] D. Mayers, Unconditional security in quantum cryptography. *J. ACM* **48**(3), 351–406 (2001)
[41] D. Micciancio, O. Regev, Worst-case to average-case reductions based on Gaussian measure. *SIAM J. Comput.* **37**(1), 267–302 (2007)
[42] C. Moore, A. Russell, The symmetric group defies strong Fourier sampling: Part II (2005). quant-ph/0501066
[43] C. Moore, D. Rockmore, A. Russell, L.J. Schulman, The hidden subgroup problem in affine groups: basis selection in Fourier sampling, in *Proceedings of the 15th ACM–SIAM Symposium on Discrete Algorithms* (2004), pp. 1106–1115
[44] C. Moore, A. Russell, L.J. Schulman, The symmetric group defies strong Fourier sampling. *SIAM J. Comput.* **37**(6), 1842–1864 (2008). See also quant-ph/0501056 and quant-ph/0501066
[45] M.A. Nielsen, I.L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000)

[46] G.M. Nikolopoulos, Applications of single-qubit rotations in quantum public-key cryptography. *Phys. Rev. A* **77**, 032348 (2008)
[47] G.M. Nikolopoulos, L.M. Ioannou, Deterministic quantum-public-key encryption: forward search attack and randomization. *Phys. Rev. A* **79**, 042327 (2009)
[48] T. Okamoto, K. Tanaka, S. Uchiyama, Quantum public-key cryptosystems, in *Advances in Cryptology—CRYPTO 2000*. LNCS, vol. 1880 (Springer, Berlin, 2000), pp. 147–165
[49] O. Regev, New lattice-based cryptographic constructions. *J. ACM* **51**(6), 899–942 (2004)
[50] O. Regev, Quantum computation and lattice problems. *SIAM J. Comput.* **33**(3), 738–760 (2004)
[51] R. Renner, *Security of quantum key distribution*. Ph.D. thesis, ETH Zurich (2005). quant-ph/0512258
[52] O. Regev, On lattices, learning with errors, random linear codes, and cryptography. *J. ACM* **56**(6), 34 (2009)
[53] D. Robinson, *A Course in the Theory of Groups* (Springer, Berlin, 1995)
[54] U. Schöning, Graph isomorphism is in the low hierarchy. *J. Comput. Syst. Sci.* **37**, 312–323 (1988)
[55] A. Schmidt, Quantum algorithms for many-to-one functions to solve the regulator and the principal ideal problem (2009). arXiv:0912.4807
[56] P.W. Shor, Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.* **26**(5), 1484–1509 (1997)
[57] P.W. Shor, J. Preskill, Simple proof of security of the BB84 quantum key distribution protocol. *Phys. Rev. Lett.* **85**, 441–444 (2000)
[58] M. Tompa, H. Woll, Random self-reducibility and zero knowledge interactive proofs of possession of information, in *Proceedings of the 28th IEEE Symposium on Foundations of Computer Science* (1987), pp. 472–482
[59] J. Watrous, Limits on the power of quantum statistical zero-knowledge, in *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science* (2002), pp. 459–468
[60] A.C.-C. Yao, Theory and applications of trapdoor functions (extended abstract), in *Proceedings of the 23rd Annual IEEE Symposium on Foundations of Computer Science* (1982), pp. 80–91