報告番号	甲　第　13789　号

主論文の要旨

論文題目　**Post-Quantum Provable Security in Symmetric-Key Cryptography (共通鍵暗号技術の耐量子証明可能安全性)**

氏　　名　**細山田　光倫**

論文内容の要旨

This paper studies provable security of symmetric-key schemes against adversaries that have quantum computers from both theoretical and practical perspectives.

Provable security is a way to mathematically guarantee the security of a cryptosystem, by showing a theorem that expresses the upper bound on the success probability of an adversary that has specified resources. Most of modern cryptosystems are shown to be secure in the provable security paradigm under some assumptions, e.g., the hardness of certain algebraic problems, or the existence of another secure cryptographic primitive. Sometimes security proofs are provided in an ideal model where the oracle of an ideally random primitive, e.g., a truly random function, is publicly available. Whether a cryptographic primitive can be built from another primitive is a central problem in the theory of cryptology. In addition, if an existing scheme is proven to resist more powerful attacks than previously thought, or if we can prove that a new efficient scheme is secure, the proofs have practical importance. Thus provable security is important both theoretically and practically.

In symmetric cryptology, (tweakable) block ciphers, pseudorandom functions (PRFs), and hash functions play central roles as fundamental underlying primitives to build other cryptosystems such as authenticated encryption schemes. Hence the provable security of such schemes is well-studied.

One of the most important results on provable security is the one on the Luby-Rackoff construction. The Luby-Rackoff construction, or the Feistel construction, is among the most important approaches to construct secure block ciphers from secure pseudorandom

functions (PRFs). The 3-round and 4-round Luby-Rackoff construction are proven to be a pseudorandom permutation (PRP) and a strong PRP, i.e., they are secure against chosen-plaintext attacks (CPAs) and chosen-ciphertext attacks (CCAs), respectively. Another important result on block ciphers is the one by Liskov, Rivest, and Wagner. They showed constructions to convert secure block ciphers into secure tweakable block ciphers, which are called the LRW constructions. As for constructions to convert Merkle-Damgård hash functions into message authentication codes (MACs) or PRFs in a provably secure manner, there has been a long line of research on HMAC and NMAC. They are proven to be secure up to $O(2^{n/2})$ computations when the output length is n bits.

On the other hand, their security has not been studied enough in the setting where adversaries have quantum computers, and many important problems have yet to be solved. On the Luby-Rackoff construction, Kuwakado and Morii showed that a quantum superposed chosen-plaintext attack (qCPA) can distinguish the 3-round construction from a random permutation in polynomial time. In addition, Ito et al. showed a quantum superposed chosen-ciphertext attack (qCCA) that distinguishes the 4-round Luby-Rackoff construction. Since Kuwakado and Morii showed the result, a problem of much interest has been how many rounds are sufficient to achieve provable security against quantum query attacks. Though several years have passed since then, the problem still remains open. Similarly, since Kaplan et al. showed the LRW construction can be broken with a polynomial-time qCPA, it has been open whether there exists a mode of block ciphers to build quantum-secure tweakable block ciphers. For HMAC and NMAC, Song and Yun showed that they are quantum pseudorandom functions (qPRFs) under the standard assumption that the underlying compression function is a qPRF. Their proof guarantees security up to $O(2^{n/5})$ or $O(2^{n/8})$ quantum queries. However, there is a gap between the provable security bound and a simple distinguishing attack that uses $O(2^{n/3})$ quantum queries.

This paper settles these problems. First, we prove that the 4-round Luby-Rackoff construction is secure up to $O(2^{n/6})$ quantum queries, where n is the length of inputs and outputs of the construction. We also prove that the bound is tight by showing an attack that distinguishes the 4-round Luby-Rackoff construction from a random permutation with $O(2^{n/6})$ quantum queries. Our result is the first to demonstrate the tight security of a typical block-cipher construction against quantum query attacks, without any algebraic assumptions.

Second, we show the first design of quantum-secure tweakable block ciphers based on quantum-secure block ciphers, and present a provable security bound. Our construction is simple, and when instantiated with a quantum-secure n-bit block cipher, it is secure against attacks that query arbitrary quantum superpositions of plaintexts and tweaks up to $O(2^{n/6})$ quantum queries.

Third, we close the gap between the security bound and the distinguishing attack of HMAC and NMAC. Specifically, we show that the tight bound of the number of quantum queries to distinguish HMAC or NMAC from a random function is $\Theta(2^{n/3})$ in the quantum random oracle model, where compression functions are modeled as quantum random oracles.

We use an alternative formalization of Zhandry's compressed oracle technique to provide security proofs in the quantum setting. In addition, to show the tight security bound of HMAC and NMAC, we introduce a new proof technique based on the compressed oracle technique, focusing on the symmetry of quantum query records.

Furthermore, we show the classical indifferentiability of the SKINNY-HASH internal function. SKINNY-HASH is a family of function-based sponge hash functions, and it was selected as one of the second round candidates of the NIST lightweight cryptography competition. Its internal function is constructed from the tweakable block cipher SKINNY. The construction of the internal function is very simple and the designers claim n-bit security, where n is the block length of SKINNY. However, a formal security proof of this claim is not given in the original specification of SKINNY-HASH. In this paper, we formally prove that the internal function of SKINNY-HASH has n-bit security, i.e., it is indifferentiable from a random oracle up to $O(2^n)$ queries, substantiating the security claim of the designers. Though the result on the SKINNY-HASH internal function is a classical one, it is unlikely to be broken by quantum attacks. In addition, when post-quantum security of the SKINNY-HASH internal function will be proved, the proof will be based on our classical proof. Thus we believe it will help understanding post-quantum security of hash functions.

The results on the Luby-Rackoff construction and quantum-secure tweakable block ciphers are significant mainly from a theoretical perspective. On the other hand, the results on HMAC and NMAC, and the SKINNY-HASH internal function, are important mainly from a practical perspective.