



Universiteit
Leiden
The Netherlands

Post-quantum security of cryptographic transformations in the random oracle model

Huang, Y.

Citation

Huang, Y. (2026, April 1). *Post-quantum security of cryptographic transformations in the random oracle model*. Retrieved from <https://hdl.handle.net/1887/4300382>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/4300382>

Note: To cite this publication please use the final published version (if applicable).

Summary

A cryptographic scheme is typically designed in multiple steps: one first constructs a simpler (but weaker) scheme, and then later modifies it to a stronger, and more sophisticated one. In some of the steps, it is common to apply a generic transformation that is not tailored to the specific scheme at hand. The purpose of this thesis is to establish rigorous, provable notions of security for various such transformations that are relevant in the scope of *post-quantum cryptography* (PQC). Most of our results are obtained in an idealized setting known as the *random oracle model* (ROM), where cryptographic hash functions (or some of their components) are modelled as a random function — a *random oracle* — that can be queried by the construction as well as its attackers. This is in contrast to the *plain model* that does not involve such an idealization.

In Chapter 3, we study the Fiat-Shamir with aborts (FSwA), and hash-and-sign with retry/aborts (HSwA) design principles. These transformations are popular among post-quantum signature schemes, but their analyses turn out rather tricky. As described in the chapter, all prior security analyses of FSwA share the same subtle but crucial flaw, which also reappears in HSwA as well. This flaw invalidates prior security proofs of all schemes that rely on FSwA, including Dilithium, one of the standards selected by the US National Institute of Standards and Technology (NIST). We re-establish the security of FSwA and HSwA signature schemes, via providing new, fixed security proofs, in a unified framework that covers both.

Our analyses cover both classical and quantum attacks. The technical core here lies in a hybrid sequence that involves *random oracle reprogramming*. Specifically, we argue that, should we reprogram the values of a random oracle in a certain way, and then later undo some of those reprogramming, then in this sequence of modifications to the random oracle, an attacker is unlikely to notice any difference. For an attacker that is given at most q_H quantum queries to the random oracle, and q_S classical queries to the signing oracle, this initially comes with an additive security loss of order $O(q_H \sqrt{q_S \epsilon})$, where ϵ is a scheme-dependent parameter that is close to zero. We then later improve this to $O(q_S \sqrt{q_H \epsilon})$, via a more sophisticated hybrid sequence. This improvement has quantitatively matched state-of-the-art bounds in the better understood case of the Fiat-Shamir transformation (*without* aborts). Although it remains

open whether our bounds are optimal, any further improvements would also need to carry over to the better-understood setting without aborts.

In Chapter 4, we study the BUFF transformation, which is a transformation for signature schemes that aims to provide additional security properties beyond the standard unforgeability. We show that one of these properties — non-resignability (NR) — is more subtle than previously believed. Indeed, the original formal definition of NR in the plain model, put forward by Cremers *et al.*, turns out unachievable, and the same applies to the natural extension of NR in the random oracle model. The unachievability is presented in the form of a simple concrete attack that applies to all “natural” schemes on the table (though leaving a small technical gap for the “un-natural” ones). In particular, it covers all signature schemes that use the BUFF transformation, and so it invalidates prior claimed security of BUFF.

The aforementioned attack, however, does not really threaten the intended applications of NR. Instead, it demonstrates that the original formal definition of NR is inadequate. To recover from this negative state of affairs, we thus go back to the drawing board: proposing a series of new formal definitions, and investigate achievability of these definitions. In the end, we obtain both positive and negative results. On one hand, we show that the BUFF transformation indeed satisfies some meaningful definitions of NR, while on the other hand, whether or not NR is achieved, or achievable at all, still heavily depends on the subtle details of the formal definition.

In Chapter 5, we study a *KEM combiner*. That is, a generic compiler that transforms multiple key-encapsulation mechanisms (KEMs) into a combined KEM that is secure, as long as at least one of the underlying KEMs is secure. Such a combiner would allow one to combine a well-established pre-quantum KEM, say, based on RSA or Diffie-Hellman, to a new post-quantum scheme, and thereby providing a smoother transition for parties to deploy post-quantum cryptographic schemes with less risk.

What we consider, is a particularly efficient construction proposed by Gideon, Heuer, and Poettering that relies on the security of a split-key pseudorandom function (skPRF). Considering its main application in the context of post-quantum cryptography, we prove that the considered skPRF remains secure against quantum attacks, and consequently the corresponding KEM combiner is also secure against quantum attacks.