Tight adaptive reprogramming in the QROM

Alex B. Grilo* Kathrin Hövelmanns† Andreas Hülsing‡ Christian Majenz¶

The *Random oracle model* (ROM) is a fundamental concept in cryptography that has enabled efficient security proofs for crypto systems of practical interest. As the name indicates, a random oracle is a function O drawn uniformly at random from the set of all functions with a desired domain and range, and all parties in a cryptographic protocol (including adversaries) can access it through an interface that on input x outputs $\mathsf{O}(x)$.

The ROM often allows for conceptually simpler and often tighter proofs than the *standard model* (i.e. without oracles / idealized assumptions). Importantly, up to this date, no natural cryptosystem with a security proof in in the ROM has been broken when the random oracle is replaced by a cryptographic hash function.[1]

With the advent of quantum computing, and the possibility of quantum adversaries, the ROM had to be generalized: In this scenario, a quantum adversary interacts with a non-quantum network, meaning that "online" primitives (= functions that can only be computed by the honest participants) stay classical, while the adversary can compute all "offline" primitives (like hash functions) on its own, and hence, in superposition. To account for these stronger capabilities, the quantum-accessible ROM (QROM) was introduced [4]. However, it unfortunately does not come with some advantages of its classical counterpart:

Lack of conceptual simplicity. QROM proofs are extremely complex, since many of the useful properties of the ROM are not known to translate directly to the QROM.

Tightness. Many security proofs that are tight in the ROM are not known to hold in the QROM.

Although it was expected that a proof in the ROM would imply security in the QROM, up to polynomial factors, recent results [20] present schemes which are provably secure in the ROM and insecure in the QROM. As a consequence, presenting QROM proofs is crucial to establish confidence in a post-quantum cryptosystem.

A very useful technique of the (classical) ROM is to *adaptively reprogramm* the oracle: In the security proof, a reduction can simulate the O towards the adversary A. This allows the reduction to pick a random point x and reprogram the value $\mathsf{O}(x)$ to a fresh random value, after A already made queries to O. It is not hard to see that A will likely not notice the change of the value of $\mathsf{O}(x)$, unless A performed essentially as many queries as the size of the domain of O. The key argument here is that if A does not query $\mathsf{O}(x^*)$, it has no information about it.

The ability to query an oracle in superposition renders this formerly simple argument more involved. Intuitively, a query in superposition can be viewed as a query that might contain all input values at once and therefore the very first answer of O can contain information about every value $\mathsf{O}(x)$. It hence was not clear whether it is possible to adaptively reprogram a quantum random oracle without causing a change in the adversary's view.

Previous works used different techniques to replace adaptive reprogramming in the quantum setting, such as variants of Unruh's one-way-to-hiding (O2H) lemma [19, 10, 14], or "history-free" proofs [4, 15, 18, 13]. However, using the O2H lemma leads to non-tight security proofs, damaging the efficiency of candidate replacements for quantum-broken cryptographic protocols that are currently run millions of times per second. History-free proofs do not always exist and if they exist, they are tailor-made proofs that must be examined case by case, being conceptually much more complicated. In addition, they might come at a cost in terms of runtime (as is the case in, e.g., [15]). Hence, in this work we are interested in the fundamental question:

Can we tightly *prove that adaptive reprogramming can also be done in the quantum random oracle model?*

The main contribution of our paper is a strong positive answer for the previous question. More precisely, we prove a lower bound on the quantum query complexity for detecting reprogramming on a random oracle and we show that this lower bound is tight, up to constant multiplicative factors. We remark that with our result, we not only prove the security of cryptographic contructions of practical interest (some of them being in the final round of the NIST standardization competition for post-quantum cryptography [17]), but we also develop technical tools on query complexity that might have other applications in quantum cryptography and algorithms.

*Sorbonne Université † Ruhr-Universität Bochum ‡Eindhoven University of Technology ¶ CWI and QuSoft

[1] There are, however, artificial separation examples [7, 11].

Our contribution. Before we describe our result in more detail, we briefly explain the reprogramming setup. As established in [4, 3], the quantum access to a random oracle $\mathsf{O} : X \to Y$ is modeled via an oracle access to a unitary U_{O}, defined as $|x\rangle_X |y\rangle_Y \mapsto |x\rangle_X |y \oplus \mathsf{O}(x)\rangle_Y$. The adversaries A have access to O via applications of U_{O}, interleaved with arbitrary unitaries (that we keep implicit for now).

While we prove a result for a very general reprogramming setup in our paper [12], we focus on a simpler special case here for ease of presentation. We characterize the distinguishing probability of an adversary A by its ability to distinguish two experiments $\textsc{Repro}_0$ and $\textsc{Repro}_1$. In $\textsc{Repro}_0$, we consider the case where O is reprogrammed at R points, and A has access to the oracle as follows:

1. A makes q queries to the random oracle O.
2. For r from 2 to R:
   (a) The values $\hat{x}_r \sim X$ and $\hat{y}_r \sim Y$ are picked uniformly at random, and we define $\mathsf{O}^{(r)}(x) = \mathsf{O}^{(r-1)}(x)$ for $x \neq \hat{x}_r$ and $\mathsf{O}^{(r)}(\hat{x}_r) = \hat{y}_r$.
   (b) A learns $\hat{x}_r$ and then makes q queries to the $\mathsf{O}^{(r)}$.
3. A measures her internal state and outputs a bit.

$\textsc{Repro}_1$ follows the same structure as $\textsc{Repro}_0$, but the random oracle remains intact, i.e. $\mathsf{O}^{(r)}(x) = \mathsf{O}^{(r-1)}(x)$ for all $x \in X$ and $1 \leq r \leq R$. We are able then to prove the following:

Adaptive reprogramming theorem. *Let X, Y be some finite sets, and let D be any distinguisher, issuing R many reprogramming instructions and q many (quantum) queries to $\mathsf{O} : X \to Y$. Let q denote the number of queries to O that are issued inbetween each potential reprogramming steps. Then*

$$\left|\Pr[\textsc{Repro}_1^{\mathsf{D}} \Rightarrow 1] - \Pr[\textsc{Repro}_0^{\mathsf{D}} \Rightarrow 1]\right| \leq \left(R + \frac{1}{2}\right)\sqrt{\frac{q}{|X|}}.$$

The above theorem constitutes a significant improvement over previous bounds. In [19] and [10], a bound proportional to $\frac{q}{\sqrt{|X|}}$ for the distinguishing advantage in similar settings and in [14], a bound proportional to $\frac{q^2}{|X|}$ is claimed, but that seems to have resulted from a "translation mistake" from [10] and should be similar to the bounds from [19, 10].

As previously mentioned, the theorem that we prove is actually more general than stated, allowing A to have partial control of the distribution from which each $\hat{x}_r$ is picked, with the only requirement that such a distribution must have high-enough entropy (being possibly different at each reprogramming step). In this case, where A has some partial information about the (potentially) reprogrammed values, we obtain the same bound, with $|X|^{-1}$ replaced by the min-entropy of the reprogrammed position.

To prove (the generalization of) the above theorem, we employ Zhandry's superposition oracle technique to translate the distinguishing probability between the games to the task of distinguishing two quantum states. We remark that this reduction from distinguishing oracles to distinguishing quantum states might be of independent interest.

Crucially, we also show that our lower bound is tight, presenting a quantum attack that matches it, up to a constant factor. For simplicity, we restrict our attention to the simple case where O is potentially reprogrammed at a single random position x^* resulting in a new oracle O'. In addition, we set $X = Y = \{0,1\}^n$, again for ease of exposition. Consider an attacker that is allowed q queries to the random oracle, learns a value x^* and peform q extra queries to O/O'.

A classical attack that matches the classical bound for the success probability, $O(q \cdot 2^{-n})$, is the following: pick values $x_1, ..., x_q$ and compute the XOR of the outputs $\mathsf{O}(x_i)$. After the oracle is potentially reprogrammed, the attacker outputs 0 iff the checksum computed before is unchanged. This attack makes $2q$ queries, and runs in time $O(q)$ and space $O(n)$.

In order to match the quantum lower bound in similar time and space complexity, we propose a *quantum checksum* technique, which consists of the classical attack but on a superposition of tuples of inputs: the attacker queries O with the superposition of all possible inputs, and then applies a cyclic permutation σ on the input register.

This process is repeated an additional $q-1$ times, on the same registers. After the potential reprogramming, we repeat the same process, but now applying the permutation σ^{-1} and querying O'. Using techniques from [1], we show how to distinguish the two cases with advantage $\Omega\left(\sqrt{\frac{q}{2^n}}\right)$ in time $\mathsf{poly}(q,n)$. We think that this "quantum checksum" technique might find other applications in the design of quantum algorithms, e.g. for the element distinctness problem.

Applications. We go on to demonstrate the applicability of our tool with three examples.

First, we show a tighter analysis for the hash-and-sign construction that takes a fixed-message-length signature scheme and turns it into a variable-message-length signature scheme by first compressing the message using a hash function. With our technical result, we can tighten a recent security proof [5] for message-compression as described for XMSS [6] in the standard RFC 8391.[2] Our new bound shows that one can use random strings of half the length to randomize the message compression in a provably secure way.

Secondly, we tighten the security of Fiat-Shamir signature schemes, improving quantitatively and qualitatively the analysis from [15].

- Our security proof is conceptually much simpler: while in [15], the random oracle is patched *a priori* with message-dependent randomness using additional machinery like pseudorandom functions, our result follows by a simple (but careful) application of our lemma, i.e. by patching the random oracle *a posteriori*.
- While the technique from [15] causes a quadratic blow-up in the reduction's running time, which renders it non-tight in all practical aspects, we obtain a reduction that is tight with respect to running time.
- While the approach in [15] requires additional assumptions that prevent an application of their analysis to certain concrete constructions like e.g. the alternate NIST candidate Picnic [8], our technique works under very general circumstances, supporting e.g. a security proof for Picnic.

Our third application concerns stronger notions of security of signature schemes. When it comes to real-world implementations, the security of signature schemes should take into account further types of attacks. For instance, an adversary interacting with hardware that realizes a cryptosystem can try to induce a hardware malfunction, also called fault injection, in order to derail the key generation or signing process. Although it is not always straightforward to predict where a triggered malfunction affects the execution, it is well known that even a low-precision malfunction can seriously injure a schemes' security. In the context of the ongoing effort to standardize post-quantum secure primitives [16], it hence made sense to affirm [17] that certain additional security features are desirable. These include, amongst others, resistance against fault attacks and biased randomness generation. With our reprogramming lemma, we generalize the result of [2] from the ROM to the QROM. In doing so, we show for the first time that the *hedged Fiat-Shamir construction* is secure against biased nonces and several types of fault injections, in the QROM. This result can for example be used to argue that alternate NIST candidate Picnic [8] is robust against many types of fault injections, even in the quantum setting.

We would like to remark that the above list of examples is non-exhaustive, and other schemes will also benefit from (variants of) the mentioned security reductions. Schemes that are amenable to our security reductions include e.g. the NIST finalist Dilithium [9], and a number of other post-quantum digital signature schemes.

References

[1] Gorjan Alagic, Christian Majenz, and Alexander Russell. Efficient simulation of random states and random unitaries. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology – EUROCRYPT 2020, Part III*, volume 12107 of *Lecture Notes in Computer Science*, pages 759–787, Zagreb, Croatia, May 10–14, 2020. Springer, Heidelberg, Germany.

[2] Diego F. Aranha, Claudio Orlandi, Akira Takahashi, and Greg Zaverucha. Security of hedged Fiat-Shamir signatures under fault attacks. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology – EUROCRYPT 2020, Part I*, volume 12105 of *Lecture Notes in Computer Science*, pages 644–674, Zagreb, Croatia, May 10–14, 2020. Springer, Heidelberg, Germany.

[2] A *Request For Comments* (RFC) is a publication by Internet Society, the Internet Engineering Task Force or similar bodies, mostly (as a precursor) for standardization. RFC 8391 has been endorsed by NIST in Special Publication 800-208, it can thus be considered standardized.

[3] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. In *39th Annual Symposium on Foundations of Computer Science*, pages 352–361, Palo Alto, CA, USA, November 8–11, 1998. IEEE Computer Society Press.

[4] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In Dong Hoon Lee and Xiaoyun Wang, editors, *Advances in Cryptology – ASIACRYPT 2011*, volume 7073 of *Lecture Notes in Computer Science*, pages 41–69, Seoul, South Korea, December 4–8, 2011. Springer, Heidelberg, Germany.

[5] Joppe W. Bos, Andreas Hülsing, Joost Renes, and Christine van Vredendaal. Rapidly Verifiable XMSS Signatures. Cryptology ePrint Archive, Report 2020/898, 2020. https://eprint.iacr.org/2020/898.

[6] Johannes A. Buchmann, Erik Dahmen, and Andreas Hülsing. XMSS - A practical forward secure signature scheme based on minimal security assumptions. In Bo-Yin Yang, editor, *Post-Quantum Cryptography - 4th International Workshop, PQCrypto 2011*, pages 117–129, Tapei, Taiwan, November 29 – December 2 2011. Springer, Heidelberg, Germany.

[7] Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited. *J. ACM*, 51(4):557–594, July 2004.

[8] Melissa Chase, David Derler, Steven Goldfeder, Claudio Orlandi, Sebastian Ramacher, Christian Rechberger, Daniel Slamanig, and Greg Zaverucha. Post-quantum zero-knowledge and signatures from symmetric-key primitives. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, CCS '17, page 1825–1842, New York, NY, USA, 2017. Association for Computing Machinery.

[9] Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. Crystals-dilithium: A lattice-based digital signature scheme. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2018(1):238–268, Feb. 2018.

[10] Edward Eaton and Fang Song. Making Existential-unforgeable Signatures Strongly Unforgeable in the Quantum Random-oracle Model. In *TQC 2015*, LIPIcs, 2015.

[11] Edward Eaton and Fang Song. A note on the instantiability of the quantum random oracle. In Jintai Ding and Jean-Pierre Tillich, editors, *Post-Quantum Cryptography - 11th International Conference, PQCrypto 2020*, pages 503–523, Paris, France, April 15–17 2020. Springer, Heidelberg, Germany.

[12] Alex B. Grilo, Kathrin Hövelmanns, Andreas Hülsing, and Christian Majenz. Tight adaptive reprogramming in the qrom. Cryptology ePrint Archive, Report 2020/1361, 2020. https://eprint.iacr.org/2020/1361.

[13] Kathrin Hövelmanns, Eike Kiltz, Sven Schäge, and Dominique Unruh. Generic authenticated key exchange in the quantum random oracle model. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020: 23rd International Conference on Theory and Practice of Public Key Cryptography, Part II*, volume 12111 of *Lecture Notes in Computer Science*, pages 389–422, Edinburgh, UK, May 4–7, 2020. Springer, Heidelberg, Germany.

[14] Andreas Hülsing, Joost Rijneveld, and Fang Song. Mitigating multi-target attacks in hash-based signatures. In Chen-Mou Cheng, Kai-Min Chung, Giuseppe Persiano, and Bo-Yin Yang, editors, *PKC 2016: 19th International Conference on Theory and Practice of Public Key Cryptography, Part I*, volume 9614 of *Lecture Notes in Computer Science*, pages 387–416, Taipei, Taiwan, March 6–9, 2016. Springer, Heidelberg, Germany.

[15] Eike Kiltz, Vadim Lyubashevsky, and Christian Schaffner. A concrete treatment of Fiat-Shamir signatures in the quantum random-oracle model. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018, Part III*, volume 10822 of *Lecture Notes in Computer Science*, pages 552–586, Tel Aviv, Israel, April 29 – May 3, 2018. Springer, Heidelberg, Germany.

[16] NIST. National institute for standards and technology. postquantum crypto project, 2017. http://csrc.nist.gov/groups/ST/post-quantum-crypto/.

[17] NIST. Status report on the second round of the nist post-quantum cryptography standardization process. NISTIR 8309, 2020. https://doi.org/10.6028/NIST.IR.8309.

[18] Tsunekazu Saito, Keita Xagawa, and Takashi Yamakawa. Tightly-secure key-encapsulation mechanism in the quantum random oracle model. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018, Part III*, volume 10822 of *Lecture Notes in Computer Science*, pages 520–551, Tel Aviv, Israel, April 29 – May 3, 2018. Springer, Heidelberg, Germany.

[19] Dominique Unruh. Quantum position verification in the random oracle model. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014, Part II*, volume 8617 of *Lecture Notes in Computer Science*, pages 1–18, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.

[20] Takashi Yamakawa and Mark Zhandry. A note on separating classical and quantum random oracles. Cryptology ePrint Archive, Report 2020/787, 2020. https://eprint.iacr.org/2020/787.