

Adaptive versus Static Multi-oracle Algorithms, and Quantum Security of a Split-key PRF

Jelle Don¹, Serge Fehr^{1,2}, and Yu-Hsuan Huang¹

¹Centrum Wiskunde & Informatica, The Netherlands

²Leiden University, The Netherlands

{jelle.don, serge.fehr, yhh}@cwi.nl

Abstract

In the first part of the paper, we show a generic compiler that transforms any oracle algorithm that can query multiple oracles *adaptively*, i.e., can decide on *which* oracle to query at what point dependent on previous oracle responses, into a *static* algorithm that fixes these choices at the beginning of the execution. Compared to naive ways of achieving this, our compiler controls the blow-up in query complexity for each oracle *individually*, and causes a very mild blow-up only.

In the second part of the paper, we use our compiler to show the security of the very efficient hash-based *split-key PRF* proposed by Giacon, Heuer and Poettering (PKC 2018), in the *quantum* random-oracle model. Using a split-key PRF as the key-derivation function gives rise to a secure KEM combiner. Thus, our result shows that the hash-based construction of Giacon *et al.* can be safely used in the context of quantum attacks, for instance to combine a well-established but only classically-secure KEM with a candidate KEM that is believed to be quantum-secure.

Our security proof for the split-key PRF crucially relies on our adaptive-to-static compiler, but we expect our compiler to be useful beyond this particular application. Indeed, we discuss a couple of other, known results from the literature that would have profited from our compiler, in that these works had to go through serious complications in order to deal with adaptivity.

1 Introduction

This paper offers two main contributions. In a first part, we show a generic reduction from adaptive to static multi-oracle algorithms, with a mild increase of the query complexity *for each oracle individually*, and in the second part, exploiting the reduction from the first part, we prove quantum security of the hash-based split-key pseudorandom function (skPRF) proposed in [GHP18]. We now discuss these two contributions in more detail.

Adaptive versus Static Multi-Oracle Algorithms. In certain cryptographic security games, the attacker \mathcal{A} is an *oracle* algorithm that is given query access to *multiple* oracles. This is in particular the case when considering the design of a cryptographic scheme in an idealized setting. Consider for instance the security definitions of public-key encryption and signature schemes in the (quantum) random-oracle model, where the attacker is given oracle access to both: the random-oracle and to a decryption/signing oracle.

By default, such an attacker \mathcal{A} can then choose *adaptively*, i.e., depending on answers to previous queries, at what point to query *which oracle*. This is in contrast to a *static* \mathcal{A} that has a predefined order of when it queries which oracle.¹ In certain cases, proving security for a static attacker is easier than proving security for a full fledged adaptive attacker, or taking care of adaptivity (naively) results in an unnecessary blow-up in the error term (see later).

In this light, it seems to be desirable to have a generic compiler that transforms any adaptive attacker \mathcal{A} into a static attacker $\bar{\mathcal{A}}$ that is equally successful in the attack. And there is actually a

¹In either case, we allow \mathcal{A} to decide adaptively *what input* to query, when having decided (adaptively or statically) on which oracle to query.

simple, naive solution for that. Indeed, let \mathcal{A} be an arbitrary oracle algorithm that makes adaptive queries to n oracles $\mathcal{O}_1, \dots, \mathcal{O}_n$, and consider the static oracle algorithm $\bar{\mathcal{A}}$ defined as follows: $\bar{\mathcal{A}}$ simply runs \mathcal{A} , and at every point in time when \mathcal{A} makes a query to one of $\mathcal{O}_1, \dots, \mathcal{O}_n$ (but due to the adaptivity it will only become clear at the time of the query *which* \mathcal{O}_i is to be queried then), the algorithm $\bar{\mathcal{A}}$ makes n queries, one to every \mathcal{O}_i , and it relays \mathcal{A} 's query to the right oracle, while making dummy queries to the other oracles.

At first glance, this simple solution is not too bad. It certainly transforms any adaptive \mathcal{A} into a static $\bar{\mathcal{A}}$ that will be equally successful, and the blow-up in the total query complexity is a factor n only, which is mild given that the typical case is $n = 2$. However, it turns out that in many situations, considering the blow-up in the total query complexity is not good enough.

For example, consider again the case of an attacker against a public-key encryption scheme in the random-oracle model. In this example, it is typically assumed that \mathcal{A} may make many more queries to the random-oracle than to the decryption oracle, i.e., $q_H \gg q_D$. But then, applying the above simple compiler, $\bar{\mathcal{A}}$ makes the same number of queries to the random-oracle and to the decryption oracle; namely $\bar{q}_H = \bar{q}_D = q_H + q_D$. Furthermore, the actual figure of merit, namely the advantage of an attacker $\bar{\mathcal{A}}$, is typically not (bounded by) a function of the total query complexity, but a function of the two respective query complexities q_H and q_D *individually*. For example, if one can show that the advantage of any *static* attacker $\bar{\mathcal{A}}$ with respective query complexities \bar{q}_H and \bar{q}_D is bounded by, say, $\bar{q}_H \bar{q}_D^2$, then the above compiler gives a bound on the advantage of any *adaptive* attacker \mathcal{A} with respective query complexities q_H and q_D of $q_H^3 + 2q_H^2 q_D + q_H q_D^2$. If $q_H \gg q_D$ then this is significantly worse than $\approx q_H q_D^2$, which one might hope for given the bound for static $\bar{\mathcal{A}}$.

Our first result is a compiler that transforms any *adaptive* oracle algorithm \mathcal{A} that makes at most q_i queries to oracle \mathcal{O}_i for $i = 1, \dots, n$ into a *static* oracle algorithm $\bar{\mathcal{A}}$ that makes at most $\bar{q}_i = nq_i$ queries to oracle \mathcal{O}_i for $i = 1, \dots, n$. Thus, rather than controlling the blow-up in the total number of queries, we can control the blow-up in the number of queries for each oracle *individually*, yet still with the same factor n . Our result applies for *any* vector $\mathbf{q} = (q_1, \dots, q_n) \in \mathbb{N}^n$ and contains no hidden constants. Our compiler naturally depends on \mathbf{q} (or, alternatively, needs \mathbf{q} as input) but otherwise only requires straight-line black-box access to \mathcal{A} , and it preserves efficiency: the run time of $\bar{\mathcal{A}}$ is polynomial in $Q = q_1 + \dots + q_n$, plus the time needed to run \mathcal{A} . Furthermore, the compiler is applicable to any classical or quantum oracle algorithm \mathcal{A} , where in the latter case the queries to the oracles $\mathcal{O}_1, \dots, \mathcal{O}_n$ may be classical or quantum as well; however, the *choice* of the oracle for each query is assumed to be classical (so that individual query complexities are well defined).

In the above made-up example of a public-key encryption scheme with advantage bounded by $\bar{q}_H \bar{q}_D^2$ for any static $\bar{\mathcal{A}}$ with respective query complexities \bar{q}_H and \bar{q}_D , we now get the bound $8q_H q_D^2$ for any adaptive \mathcal{A} with respective query complexities q_H and q_D .

We show the usefulness of our adaptive-to-static compiler by discussing two example results from the literature. One is the security proof by Alkim *et al.* [ABB⁺17] of the qTESLA signature scheme [ABB⁺20] in the quantum random-oracle model; the other is the recent work by Alagic, Bai, Katz and Majenz [ABKM21] on the quantum security of the famous Even-Mansour cipher. In both these works, the adaptivity of the attacker was a serious obstacle and caused a significant overhead and additional complications in the proof. With our results, these complications could have been avoided without sacrificing much in the security loss (as would be the case with using a naive compiler). We also exploit our adaptive-to-static compiler in our second main contribution, discussed below.

Interestingly, all three example applications are in the realm of quantum security (of a classical scheme). This seems to suggest that the kind of adaptivity we consider here is not such much of a hurdle in the case of classical queries. Indeed, in that case, a typical argument works by inspecting the entire query transcript and identifying an event with the property that conditioned on this event, whatever needs to be shown holds *with certainty*, and then it remains to show that this event is very likely to occur. In the case of quantum queries, this kind of reasoning does not apply since one cannot “inspect” the query transcript anymore; instead, one then typically resorts to some sort of hybrid argument where queries are replaced one-by-one, and then adaptivity of the queries may — and sometimes does, as we discuss — form a serious obstacle.

Quantum-security of a Split-key PRF. In the upcoming transition to post-quantum secure cryptographic standards, *combiners* may play an important role. A combiner can be used compile several cryptographic schemes into a new, “combined” scheme, which offers the same (or a similar) functionality, and so that the new scheme is secure as long as *at least one* of the original schemes is secure. For example, combining a well-established but quantum-insecure scheme with a believed-to-be quantum-secure (but less well studied) scheme then offers the best of both worlds: it offers security against quantum attacks, should there really be a quantum computer in the future, but it also offers some protection in case the latter scheme turns out to be insecure (or less secure than expected) even against classical attacks. In other words, using a combiner in this context ensures that we are not making things less secure by trying to aim for quantum security.

In [GHP18], Giacon, Heuer and Poettering showed that any *split-key PRF* (skPRF) gives rise to a secure KEM combiner. In more detail, they show that if a skPRF is used in the (rather) obvious way as a key-derivation function in a KEM combiner, then the resulting combined KEM is IND-CCA secure if at least one of the component KEMs is IND-CCA secure. They also suggest a few candidates for skPRFs. The most efficient of the proposed constructions is a hash-based skPRF, which is proven secure in [GHP18] in the random-oracle model. However, in the context of a quantum attack, which is particularly relevant in the above example application of a combiner, it is crucial to prove security in the *quantum* random-oracle model [BDF⁺11]. Here, we close this gap by proving security of the hash-based skPRF construction proposed by Giacon *et al.* in the quantum random-oracle model.

Our security proof crucially exploits our adaptive-to-static compiler to reduce a general, adaptive attacker/distinguisher to a static one. Namely, in spirit, our security proof is a typical hybrid proof, where we replace, one by one, the queries to the (sk)PRF by queries to a truly random function; however, the crux is that for each hybrid, corresponding to a particular function query that is to be replaced, the closeness of the current to the previous hybrid depends on the number of hash queries *between the current and the previously replaced function query*. In case of an adaptive \mathcal{A} , each such “window” of hash queries between two function queries could be as large as the total number of hash queries in the worst case, giving rise to a huge multiplicative blow-up when using this naive bound. Instead, for a static \mathcal{A} , each such window is bounded by a fixed number, with the sum of these numbers being the total number of hash queries.

By means of our compiler, we can turn the possibly adaptive \mathcal{A} into a static one (almost) for free, and this way avoid an unnecessary blow-up, respectively bypass additional complications that arise by trying to avoid this blow-up by other means.

2 Preliminaries

We consider oracle algorithms $\mathcal{A}^{\mathcal{O}_1, \dots, \mathcal{O}_n}$ that make queries to (possibly unspecified) oracles $\mathcal{O}_1, \dots, \mathcal{O}_n$, see Fig. 1 (left). Sometimes, and in particular when the oracles are not specified, we just write \mathcal{A} and leave it implicit that \mathcal{A} makes oracle calls. We allow \mathcal{A} to be classical or quantum, and in the latter case we may also allow the queries (to some of the oracles) to be quantum; however, the choice of *which* oracle is queried is always classical. For the purpose of our work, we may assume \mathcal{A} to have no input; any potential input could be hardwired into \mathcal{A} . For a vector $\mathbf{q} = (q_1, \dots, q_n) \in \mathbb{N}^n$, we say that \mathcal{A} is a **q-query** oracle algorithm if it makes at most q_i queries to the oracle \mathcal{O}_i .

In general, such an oracle algorithm \mathcal{A} may decide *adaptively* which oracle to query at what step, dependent on previous oracle responses. In contrast to this, a *static* oracle algorithm has an arbitrary but pre-defined order in querying the oracles.

Our goal will be to transform any *adaptive* oracle algorithm \mathcal{A} into a *static* oracle algorithm $\bar{\mathcal{A}}$ that is functionally equivalent, while keeping the blow-up in query complexity for each individual oracle, i.e., the blow-up for each individual q_i , small. By *functionally equivalent* (for certain oracle instantiations) we mean the respective executions of $\mathcal{A}^{\mathcal{O}_1, \dots, \mathcal{O}_n}$ and $\bar{\mathcal{A}}^{\mathcal{O}_1, \dots, \mathcal{O}_n}$ give rise to the same output distribution *for all* (the considered) instantiations $\mathcal{O}_1, \dots, \mathcal{O}_n$ of the oracles $\mathcal{O}_1, \dots, \mathcal{O}_n$. In case of *quantum* oracle algorithms, we require the output state to be the same.

For this purpose, we declare that an *interactive oracle algorithm* \mathcal{B} is an interactive algorithm with two distinct interaction interfaces, one for the interaction with \mathcal{A} (we call this the *simulation interface*), and one for the oracle queries (we call this the *oracle interface*), see Fig. 1 (middle). For any oracle algorithm \mathcal{A} , we then denote by $\mathcal{B}[\mathcal{A}]$ the oracle algorithm that is obtained by

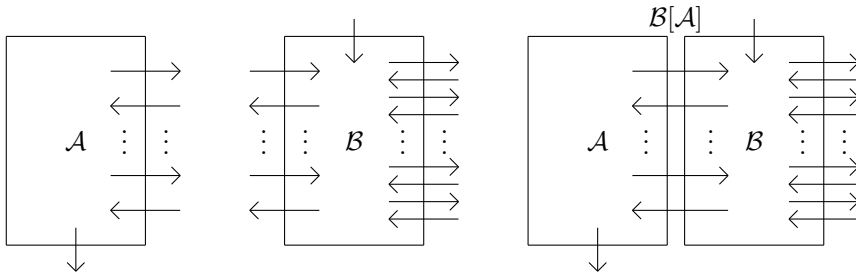


Figure 1: An oracle algorithm \mathcal{A} (left), an interactive oracle algorithm \mathcal{B} (middle), and the oracle algorithm $\mathcal{B}[\mathcal{A}]$ obtained by composing \mathcal{A} and \mathcal{B} (right).

composing \mathcal{A} and \mathcal{B} in the obvious way. In other words, $\mathcal{B}[\mathcal{A}]$ runs \mathcal{A} and answers all of \mathcal{A} 's oracle queries using its simulation interface; furthermore, $\mathcal{B}[\mathcal{A}]$ outputs whatever \mathcal{A} outputs at the end of this run of \mathcal{A} , see Fig. 1 (right).²

In contrast to \mathcal{A} (where, for our purpose, any input could be hardwired), we explicitly allow an interactive oracle algorithm \mathcal{B} to obtain an input. Indeed, our transformation, which turns any adaptive oracle algorithm \mathcal{A} into a static oracle algorithm $\bar{\mathcal{A}}$, needs to “know” \mathbf{q} , i.e., the number of queries \mathcal{A} makes to the different oracles. Thus, this will be provided in the form of an input to \mathcal{B} ; for reasons to be clear, it be provided in unary, i.e., as $1^{\mathbf{q}} := (1^{q_1}, \dots, 1^{q_n})$.

We stress that we do not put any computational restriction on the oracle algorithms \mathcal{A} (beyond bounding the queries to the individual oracles); however, we do want our transformation to preserve efficiency. Therefore, we say that an interactive oracle algorithm \mathcal{B} is *polynomial-time* if the number of local computation steps it performs is bounded to be polynomial in its input size, and where we declare that copying an *incoming* message on the simulation interface to an *outgoing* message on the oracle interface, and vice versa, is unit cost (irrespective of the size of the message). By providing \mathbf{q} in unary, we thus ensure that \mathcal{B} is polynomial-time in $q_1 + \dots + q_n$.

3 A General Adaptive-to-static Reduction for Multi-oracle Algorithms

3.1 Our Result

Let $n \in \mathbb{N}$ be an arbitrary positive integer. We present here a generic adaptiv-to-static compiler \mathcal{B} that, on input a vector $\mathbf{q} \in \mathbb{N}^n$, turns any *adaptive* \mathbf{q} -query oracle algorithm $\mathcal{A}^{\mathcal{O}_1, \dots, \mathcal{O}_n}$ into a *static* $n\mathbf{q}$ -query algorithm.

Theorem 1. *There exists a polynomial-time interactive oracle algorithm \mathcal{B} , such that for any $\mathbf{q} \in \mathbb{N}^n$ and any adaptive \mathbf{q} -query oracle algorithm $\mathcal{A}^{\mathcal{O}_1, \dots, \mathcal{O}_n}$, the oracle algorithm $\mathcal{B}[\mathcal{A}](1^{\mathbf{q}})$ is a static $n\mathbf{q}$ -query oracle algorithm that is functionally equivalent to \mathcal{A} for all stateless instantiations of the oracles $\mathcal{O}_1, \dots, \mathcal{O}_n$.*

Remark 1. As phrased, Theorem 1 applies to oracle algorithms \mathcal{A} that have no input. This is merely for simplicity. In case of an oracle algorithm \mathcal{A} that takes an input, we can simply apply the statement to the algorithm $\mathcal{A}(x)$ that has the input x hardwired, and so argue that Theorem 1 also applies in that case.

Remark 2. $\mathcal{B}[\mathcal{A}]$ is guaranteed to behave the same way as \mathcal{A} for *stateless* (instantiations of the) oracles only. This is because most of the queries that $\mathcal{B}[\mathcal{A}]$ makes are actually dummy queries (i.e., queries on a default input and with the response ignored), which have no effect in case of stateless oracles, but may mess up things in case of stateful oracles. Theorem 1 extends to arbitrary stateful oracles if we allow $\mathcal{B}[\mathcal{A}]$ to *skip* queries instead of making dummy queries (but the skipped queries would still count towards the query complexity).

²Note, we silently assume consistency between \mathcal{A} and \mathcal{B} , i.e. \mathcal{A} should send a message when \mathcal{B} expects one and the format of these messages should match the format of the messages that \mathcal{B} expects (and vice versa), so that the above composition makes sense. Should \mathcal{B} encounter some inconsistency, it will abort.




Figure 2: Constructing the string s by distributing the different symbols evenly within the interval $(0, n]$ (here with $3/q_{\sigma_1} = 2/q_{\sigma_2}$), and then collecting them from left to right.

Given the vector $\mathbf{q} = (q_1, \dots, q_n) \in \mathbb{N}^n$, the core of the problem is to find a *fixed* sequence of \mathcal{O}_i 's in which each individual \mathcal{O}_i occurs at most nq_i times, and so that *every* sequence of \mathcal{O}_i 's that contains each individual \mathcal{O}_i at most q_i times can be embedded into the former. We consider and solve this abstract problem in the following section, and then we wrap up the proof of Theorem 1 in Section 3.3.

3.2 The Technical Core

Let Σ be a non-empty finite set of cardinality n . We refer to Σ as the *alphabet*. As is common, Σ^* denotes the set of finite strings over the alphabet Σ . In other words, the elements of Σ^* are the strings/sequences $s = (s_1, \dots, s_\ell) \in \Sigma^\ell$ with arbitrary $\ell \in \mathbb{N}$ (including $\ell = 0$).

Following standard terminology, for $s = (s_1, \dots, s_\ell)$ and $s' = (s'_1, \dots, s'_m)$ in Σ^* , the *concatenation* of s and s' is the string $s||s' = (s_1, \dots, s_\ell, s'_1, \dots, s'_m)$, and s' is a *subsequence* of s , denoted $s' \sqsubseteq s$ if there exist integers $1 \leq j_1 < \dots < j_m \leq \ell$ with $(s_{j_1}, \dots, s_{j_m}) = (s'_1, \dots, s'_m)$. Such an integer sequence (j_1, \dots, j_m) is then called an *embedding* of s' into s .³

Finally, for a function $q : \Sigma \rightarrow \mathbb{N}, \sigma \mapsto q_\sigma$, we say that $s = (s_1, \dots, s_\ell) \in \Sigma^*$ has *characteristic* (at most) q if $\#\{i \mid s_i = \sigma\} = q_\sigma (\leq q_\sigma)$ for any $\sigma \in \Sigma$.

Lemma 1 (Embedding Lemma). *Let Σ be an alphabet of size n , and let $q : \Sigma \rightarrow \mathbb{N}, \sigma \mapsto q_\sigma$. Then, there exists a string $s \in \Sigma^*$ with characteristic $n \cdot q : \sigma \mapsto n \cdot q_\sigma$ such that any string $s' \in \Sigma^*$ with characteristic at most q is a subsequence of s , i.e., $s' \sqsubseteq s$.*

The idea of the construction of the sequence s is quite simple: First, we evenly distribute $n \cdot q_\sigma$ copies of σ within the interval $(0, n]$ by “attaching” one copy of σ to every point in $(0, n]$ that is an integer multiple of $1/q_\sigma$ (see Fig. 2). Note that it may happen that different symbols are “attached” to the same point. Then, we walk along the interval from 0 and n and, one by one, collect the symbols we encounter in order to build up s' from left to right; in case we encounter a point with multiple symbols “attached” to it, we collect them in an arbitrary order.

It is then not too hard to convince yourself that this s indeed satisfies the claim. Namely, for any $s' = (s'_1, \dots, s'_m)$ as considered, we can again walk along the interval from 0 and n , and we will then encounter all the symbols of s' , one by one: we will encounter the symbol s'_1 within the walk from 0 to $1/q_{s'_1}$, the symbol s'_2 then within the walk from $1/q_{s'_1}$ to $1/q_{s'_1} + 1/q_{s'_2}$, etc.

Putting this idea into a formal proof is somewhat tedious, but in the end not too difficult. In order to formalize things properly, we generalize the standard notion of a sequence $s \in \Sigma^*$ in a way that allows us to talk about “attaching” a symbol to a point on \mathbb{R} , etc., in a rigorous way. Formally, we define a *line sequence* to be an arbitrary finite (possibly empty) subset $S \subseteq \mathbb{R} \times \Sigma$, i.e.,

$$S = \{(t_1, s_1), \dots, (t_\ell, s_\ell)\} \in \mathcal{P}_{<\infty}(\mathbb{R} \times \Sigma),$$

where w.l.o.g. we will always assume that $t_1 \leq \dots \leq t_\ell$. We may think of the symbol s_i to “occur at the time” t_i .⁴ For a subset $T \subset \mathbb{R}$, the set $\mathcal{P}_{<\infty}(T \times \Sigma)$ then obviously denotes the set of line sequences with $t_1, \dots, t_\ell \in T$.

Assuming that the alphabet Σ is equipped with a total order \leq , any line sequence $S = \{(t_1, s_1), \dots, (t_\ell, s_\ell)\}$ is naturally associated with the ordinary sequence

$$\pi(S) := (s_1, \dots, s_\ell) \in \Sigma^*,$$

³We use *string* and *sequence* interchangeably; however, following standard terminology, there is a difference between a *substring* and *subsequence*: namely, a substring is a subsequence that admits an embedding with $j_{i+1} = j_i + 1$.

⁴Note that we allow $t_i = t_j$ for $i \neq j$ while the definition prohibits $(t_i, s_i) = (t_j, s_j)$. If desired, one could allow the latter by letting S be a multi-set, but this is not necessary for us.

which is uniquely determined by the convention $t_1 \leq \dots \leq t_\ell$ and insisting on $s_i \leq s_j$ whenever $t_i = t_j$ for $i < j$.

This projection $\pi : \mathcal{P}_{<\infty}(\mathbb{R} \times \Sigma) \rightarrow \Sigma^*$ preserves the characteristic of the sequence, i.e., if $s = (s_1, \dots, s_\ell) = \pi(S)$ then

$$\#\{t \mid (t, \sigma) \in S\} = \#\{i \mid s_i = \sigma\} \quad (1)$$

for any $\sigma \in \Sigma$. Furthermore, for $T, T' \subset \mathbb{R}$ with $T < T'$ point-wise, and for $S \in \mathcal{P}_{<\infty}(T \times \Sigma)$ and $S' \in \mathcal{P}_{<\infty}(T' \times \Sigma)$, it is easy to see that $\pi(S \cup S') = \pi(S) \parallel \pi(S')$, from which it then follows that for ordinary sequences $s, s' \in \Sigma^*$

$$s \sqsubseteq \pi(S) \wedge s' \sqsubseteq \pi(S') \implies s \parallel s' \sqsubseteq \pi(S) \parallel \pi(S') = \pi(S \cup S'). \quad (2)$$

A final, simple observation, which follows directly from the definitions, is that for $\sigma \in \Sigma$, i.e. a sequence of length $m = 1$, $\sigma \sqsubseteq \pi(S)$ holds if and only if there exists a time $t \in \mathbb{R}$ such that $(t, \sigma) \in S$.

Proof of Lemma 1. For any symbol $\sigma \in \Sigma$ let S_σ be a line sequence

$$S_\sigma := \left\{ \frac{1}{q_\sigma}, \dots, \frac{nq_\sigma}{q_\sigma} \right\} \times \{\sigma\} \in \mathcal{P}_{<\infty}((0, n] \times \Sigma),$$

and set $S := \bigcup_{\sigma \in \Sigma} S_\sigma$. We will show that $s := \pi(S)$ is as claimed.

The claim on the characteristic of s follows from the preservation of the characteristic under π , i.e. (1), and from $\#\{t \mid (t, \sigma) \in S\} = \#S_\sigma = n \cdot q_\sigma$, which holds by construction of S .

Let $s' = (s'_1, \dots, s'_m) \in \Sigma^*$ be arbitrary with characteristic bounded by q . We consider the times $\tau_j := 1/q_{s'_1} + \dots + 1/q_{s'_j}$ for $j \in \{1, \dots, m\}$, and we let T_j be the interval

$$T_j := (\tau_{j-1}, \tau_j] = (\tau_{j-1}, \tau_{j-1} + \frac{1}{q_j}] \subset \mathbb{R},$$

and decompose $S = S_1 \cup \dots \cup S_m$ with $S_j := S \cap (T_j \times \Sigma) \in \mathcal{P}_{<\infty}(T_j \times \Sigma)$. Here, we exploit that

$$\tau_m = \sum_{\sigma \in \Sigma} \frac{\#\{i \mid s'_i = \sigma\}}{q_\sigma} \leq \sum_{\sigma \in \Sigma} \frac{q_\sigma}{q_\sigma} = n,$$

and so the S_j 's indeed cover all of $S \in \mathcal{P}_{<\infty}((0, n] \times \Sigma)$. Given that the interval $T_j \subset (0, n]$ has size $1/q_{s'_j}$, there exists a time $t_j \in T_j \cap \left\{ \frac{1}{q_\sigma}, \dots, \frac{nq_\sigma}{q_\sigma} \right\}$. But then, $(t_j, s'_j) \in S_j$ by construction of S , and therefore $s'_j \sqsubseteq \pi(S_j)$. Finally, since $T_{j-1} < T_j$, property (2) implies that

$$s' = s'_1 \parallel \dots \parallel s'_m \sqsubseteq \pi(S_1 \cup \dots \cup S_m) = s$$

which was to be shown. \square

While Lemma 1 above settles the existence question, the following two observations settle the corresponding efficiency questions. For concreteness, we assume $\Sigma = \{1, \dots, n\}$ below, and thus can identify the function $q : \Sigma \rightarrow \mathbb{N}$, $\sigma \mapsto q_\sigma$ with the vector $\mathbf{q} = (q_1, \dots, q_n)$.

First, we observe that the line sequence S defined in the proof above, as well as its projection $s = \pi(S)$, can be computed in polynomial time in $q_1 + \dots + q_n$; thus, we have the following.

Lemma 2. *There exists a polynomial-time algorithm that, on input $1^{\mathbf{q}}$, computes a string $s \in \Sigma^*$ as specified in the proof of Lemma 1.*

Furthermore, for any $s' \in \Sigma^*$ with characteristic at most q , for which we then know by Lemma 1 that s' can be embedded into s , the following ensures that this embedding can be computed efficiently and *on the fly*.

Lemma 3. *There exists a polynomial-time algorithm \mathcal{E} such that for every string $s \in \Sigma^*$ and every subsequence $s' = (s'_1, \dots, s'_m) \sqsubseteq s$, the following holds. Computing inductively $j_i \leftarrow \mathcal{E}(s, s'_i, j_{i-1})$ for every $i \in [m]$, where $j_0 := 0$, results in an increasing sequence $j_1 < \dots < j_m$ with*

$$s' = (s_{j_1}, \dots, s_{j_m}).$$

The algorithm \mathcal{E} simply follows the obvious greedy strategy: for each s'_i it looks for the next j_i for which $s'_i = s_{j_i}$. More formally:

Proof. The algorithm $\mathcal{E}(s, s'_i, j_{i-1})$ computes

$$j_i := \min \{k \in \mathbb{N} \mid j_{i-1} < k \leq m, s_k = s'_i\} . \quad (3)$$

It can be easily shown that the minimum is well-defined, i.e. taken over a non-empty set for each i by the assumption that s' is a subsequence of s , and thus by construction, every j_i is such that $s'_i = s_{j_i}$ while keeping $j_1 < \dots < j_n$ increasing. This concludes the proof. \square

3.3 Wrapping up the Proof of Theorem 1

The claimed interactive oracle algorithm \mathcal{B} now works in the obvious way. On input \mathbf{q} (provided in unary) and for any \mathcal{A} , $\mathcal{B}[\mathcal{A}]$ will make static oracle queries to $\mathcal{O}_{s_1}, \mathcal{O}_{s_2}, \dots, \mathcal{O}_{s_{nQ}}$, where $s = (s_1, \dots, s_{nQ}) \in \{1, \dots, n\}^*$ is the string promised to exist by Lemma 1, with $Q = q_1 + \dots + q_n$. In more detail, it first computes s using the algorithm from Lemma 2. Then, for the i -th oracle query that \mathcal{B} receives from \mathcal{A} (starting with $i = 1$), and which consists of the identifier $s'_i \in \{1, \dots, n\}$ of which oracle to query now and of the actual input to the oracle $\mathcal{O}_{s'_i}$, the algorithm \mathcal{B} does the following: it computes $j_i \leftarrow \mathcal{E}(s, s'_i, j_{i-1})$ using the algorithm from Lemma 3, makes dummy queries to $\mathcal{O}_{s_{j_{i-1}+1}}, \dots, \mathcal{O}_{s_{j_i-1}}$, and forwards \mathcal{A} 's query input to $\mathcal{O}_{s_{j_i}} = \mathcal{O}_{s'_i}$. The fact that (j_1, \dots, j_Q) computed this way forms an embedding of $s' = (s'_1, \dots, s'_Q)$ into s ensures that \mathcal{B} is able to forward all the queries that \mathcal{A} makes to the right oracle, and so \mathcal{A} will produce its output as in an ordinary run with direct adaptive access to the oracles.

3.4 Applications

To demonstrate the usefulness of our adaptive-to-static compiler, we briefly discuss two results from the literature, where the adaptivity of the attacker was explicitly declared as an obstacle in the security proof, and where dealing with it complicated the proof substantially. In both cases, these complications could be avoided/removed by means of our adaptive-to-static compiler.

3.4.1 Quantum Security of qTESLA.

Our first application is in the context of qTESLA [ABB⁺20], which is a signature scheme that made it into the second round of the NIST post-quantum competition. Its security is based on the Ring-LWE problem, to which the authors of [ABB⁺17] give a reduction in the quantum random-oracle model (QROM).⁵ In the reduction, which starts from the security notion of *Unforgeability under Chosen Message Attack* (UF-CMA), the adversary can query a random-oracle H as well as a signing oracle, where the order of oracle queries may be adaptive.

The reduction strategy of [ABB⁺17] applies only to a static adversary, with a fixed query pattern. Thus, the authors first compile the adaptive into a naive static attacker by letting it do q_H (the number of H -queries of the original adaptive adversary) H -queries between any two signing queries. Leaving it with this would blow up the number of H -queries to $q_S q_H$. In order to avoid that, they give the attacker a “live-switch”, meaning that each query to H may be in superposition of making the query and not making the query, and the total “query magnitude” on actual H -queries is still restricted to q_H . Not so surprising, adding even more “quantumness” to the problem in this way, makes the analysis more complicated (compared to using standard “all-or-nothing” static queries and a standard classical bound on the query complexity), but it allows the authors to avoid the above blow-up in the (classical) query complexity to transpire into the security loss. The overall loss they obtain in the end is $O((q_S q_H^2 + q_S^3 + q_S^2 q_H) \cdot \epsilon)$ for small ϵ determined by the parameters of the scheme.

Since the security reduction in [ABB⁺17] intertwines the adaptive to static hurdle with other aspects of the proof, we cannot simply insert our Theorem 1 and then continue the proof as is. Still, by applying our result, we could obtain a static adversary with almost no cost in the number of H -queries, avoiding the need for the rather complicated “live-switch superposition” attacker, thus simplifying the overall proof significantly. Furthermore, looking ahead at Section 4, our result allows us to obtain the much better $O(\sqrt{q_O q_H^2 \epsilon} + \sqrt{q_O^2 q_H \epsilon})$ loss in a similar context — similar in

⁵We note that some versions of qTESLA have been broken [LS19], but the attack only applies to an optimized variant that was developed for the NIST-competition, and does not apply to the scheme in [ABB⁺17] that we discuss here.

the sense that it also involves two oracles where one reprograms the other at some high-entropy input. The adaptive to static reduction there allows us to apply some additional QROM tools that could potentially also be applied in the setting of qTESLA to improve the bound. However, actually doing this would require us to rewrite the entire proof of [ABB⁺17], which we consider outside the scope of this work.

3.4.2 Quantum Security of the Even-Mansour Cipher.

The recent work [ABKM21] considers the post-quantum security of the Even-Mansour cipher, which turns a public permutation into a keyed permutation, and it proves it secure in the post-quantum setting where an attacker has quantum superposition access to the public permutation, but only classical access to the keyed cipher.

Also here, the fact that the attacker can choose adaptively whether to query the public permutation of the cipher complicates the proof. Indeed, as is explained on page 3 in [ABKM21], this adaptivity issue forces the authors to extend the blinding lemma of Alagic *et al.* to a variant that gives a bound in terms of the *expected* number of queries. While the authors succeed in providing such an extended version of the blinding lemma (Lemma 3 in [ABKM21]), it further increases the complexity of an already involved proof.⁶

Thus, again, our Theorem 1 could be used to simplify the given proof by bypassing the complications that arise due to the attacker choosing adaptively which oracle to query at what point.

4 Quantum Security of a Split-key PRF

4.1 Hybrid Security and skPRFs

A *split-key pseudorandom function* (skPRF), as introduced in [GHP18], is a polynomial-time computable function $F : \mathcal{K}_1 \times \dots \times \mathcal{K}_n \times \mathcal{X} \rightarrow \mathcal{Y}$ that is a pseudorandom function (PRF) in the standard sense *for every* $i \in [n]$ when considered as a keyed function with key space \mathcal{K}_i and message space $\mathcal{K}_1 \times \dots \times \mathcal{K}_{i-1} \times \mathcal{K}_{i+1} \times \dots \times \mathcal{K}_n \times \mathcal{X}$, with the additional restriction that the distinguisher \mathcal{A} (in the standard PRF security definition) must use a fresh $x \in \mathcal{X}$ in every query $(k_1, \dots, k_{i-1}, k_{i+1}, \dots, k_n, x)$.

This restriction on the PRF distinguisher may look artificial, but is motivated by this definition of a skPRF being good enough for the intended purpose of a skPRF, namely to give rise to a *secure KEM combiner*. Indeed, [GHP18] shows that the naturally combined KEM, obtained by concatenating the individual ciphertexts to $C = (c_1, \dots, c_n)$, and combining the individual session keys k_1, \dots, k_n using the above mentioned skPRF as

$$K = F(k_1, \dots, k_n, C),$$

is IND-CCA secure if at least one of the individual KEM's is IND-CCA secure.

The paper [GHP18] also proposes a particularly efficient hash-based construction, given by

$$F(k_1, \dots, k_n, x) := H(g(k_1, \dots, k_n), x) \quad (4)$$

where $g : \mathcal{K}_1 \times \dots \times \mathcal{K}_n \rightarrow \mathcal{W}$ is a polynomial-time mapping with the property that, for some small ϵ ,

$$\Pr_{k_i \leftarrow \mathcal{K}_i} [g(k_1, \dots, k_n) = w] \leq \epsilon, \quad (5)$$

for every $i \in [n]$ and for every $k_1, \dots, k_{i-1}, k_{i+1}, \dots, k_n$ and every w ; furthermore, $H : \mathcal{W} \rightarrow \mathcal{Y}$ is a cryptographic hash function. Simple choices for the function g are $g(k_1, \dots, k_n) = (k_1, \dots, k_n)$ and $g(k_1, \dots, k_n) = k_1 + \dots + k_n$.

It is shown in [GHP18] that this construction is a skPRF when H is modelled as a random-oracle; indeed, it is shown that the distinguishing advantage is upper-bounded by $q_H \epsilon$, where q_H is the number of queries to the random-oracle H .

⁶To be fully precise, Lemma 3 in [ABKM21] also generalizes the original blinding lemma in a different direction by allowing to reprogram to an arbitrary value instead of a uniformly random one; however, this generalization comes for free in that the original proof still applies up to obvious changes, while allowing an expected number of queries, which is needed to deal with the adaptivity issue, requires a new proof.

Given the natural use of combiners in the context of the upcoming transition to post-quantum cryptography, it is natural — and well-motivated — to ask whether F can be proven to be a skPRF in the presence of a *quantum attacker*, i.e., when H is modeled as a *quantum* random-oracle. Below, we answer this in the affirmative.

4.2 Quantum-security of the skPRF

The goal of this section is to show the security of the skPRF (4) in the quantum random-oracle model. In essence, this requires proving that F is a PRF (in the quantum random-oracle model) with respect to *any* of the k_i 's being the key, subject to the restriction of asking a fresh x in each query.

To simplify the notation, we fix the index $i \in [n]$ and simply write k for k_i and x for $(k_1, \dots, k_{i-1}, k_{i+1}, \dots, k_n, x)$, and we abstract away the properties of the function g as follows. We let

$$F(k, x) := H(h(k, x)),$$

where $h : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{W}$ is an arbitrary function with the property that, for some parameter $\epsilon > 0$,

$$\Pr_{k \leftarrow \mathcal{K}} [h(k, x) = w] \leq \epsilon \quad (6)$$

for all $w \in \mathcal{W}$ and $x \in \mathcal{X}$. Furthermore, in the PRF security game, we restrict the attacker/distinguisher \mathcal{A} to queries x with a fresh value of $h(k, x)$, no matter what k is.

More formally, let $\mathcal{A}^{\mathcal{H}, \mathcal{O}}$ be an arbitrary quantum oracle algorithm, making quantum superposition queries to an oracle \mathcal{H} and classical queries to another oracle \mathcal{O} , with the restriction that for every query x to \mathcal{O} it holds that

$$h(\kappa, x) \neq h(\kappa, x'), \quad (7)$$

for any prior query x' to \mathcal{O} and all $\kappa \in \mathcal{K}$. For any such oracle algorithm $\mathcal{A}^{\mathcal{H}, \mathcal{O}}$, we consider the standard PRF security games

$$\text{PR}^1 := \mathcal{A}^{H, F} \quad \text{and} \quad \text{PR}^0 := \mathcal{A}^{H, R},$$

obtained by instantiating \mathcal{H} with a random function H (the random-oracle) in both games, and in one game we instantiate \mathcal{O} with the pseudorandom function F , which we understand to return $F(k, x)$ on query x for a random $k \leftarrow \mathcal{K}$, chosen once and for all queries, and in the other we instantiate \mathcal{O} with a truly random function R instead.

We show that the distinguishing advantage for these two games is bounded as follows.

Theorem 2. *Let $\mathcal{A}^{\mathcal{H}, \mathcal{O}}$ be a $(q_{\mathcal{H}}, q_{\mathcal{O}})$ -query oracle algorithm satisfying (7). Then*

$$|\Pr [1 \leftarrow \text{PR}^1] - \Pr [1 \leftarrow \text{PR}^0]| \leq 4\sqrt{2q_{\mathcal{O}}^2 q_{\mathcal{H}} \epsilon} + 4\sqrt{2q_{\mathcal{H}}^2 q_{\mathcal{O}} \epsilon}.$$

We can now apply Theorem 2 to the function $h(k, x) := (g(k_1, \dots, k_n), \tilde{x})$, where $k := k_i$ and $x := (k_1, \dots, k_{i-1}, k_{i+1}, \dots, k_n, \tilde{x})$. Indeed, the condition (5) on g implies the corresponding condition (7) on h , and the restriction on \tilde{x} being fresh in the original skPRF definition implies the above restriction on $h(k, x)$ being fresh not matter what k is, i.e., (6). Thus, we obtain the following.

Corollary 1. *For any function g satisfying (5) for a given $\epsilon > 0$, the function $F(k_1, \dots, k_n, x) := H(g(k_1, \dots, k_n), x)$ is a skPRF in the quantum random-oracle model with distinguishing advantage at most $4\sqrt{2q_{\mathcal{O}}^2 q_{\mathcal{H}} \epsilon} + 4\sqrt{2q_{\mathcal{H}}^2 q_{\mathcal{O}} \epsilon}$.*

4.3 Proof of Theorem 2

Proof (of Theorem 2). Let $\mathcal{A}^{\mathcal{H}, \mathcal{O}}$ be an oracle algorithm as considered in the previous subsection. Thanks to Theorem 1, taking a factor-2 blow-up in the query complexity into account, we may assume \mathcal{A} to be a *static* $(q_{\mathcal{H}}, q_{\mathcal{O}})$ -query oracle algorithm. It will be convenient to write such a static algorithm as

$$\mathcal{A}^{[\mathcal{H}_0 \circ \mathcal{H}_1 \circ \mathcal{H}_2 \dots \circ \mathcal{H}_{q_{\mathcal{O}}}]},$$

where each block $\mathcal{H}_i = \mathcal{H} \cdots \mathcal{H}$ consists of a (possibly empty) sequence of symbols \mathcal{H} of length $q_i^{\mathcal{H}} = |\mathcal{H}_i|$, and with the understanding that \mathcal{A} first makes $q_0^{\mathcal{H}}$ queries to \mathcal{H} , then a query to \mathcal{O} , then $q_1^{\mathcal{H}}$ queries to \mathcal{H} , etc., where, obviously, $q_0^{\mathcal{H}} + \cdots + q_{q_{\mathcal{O}}}^{\mathcal{H}} = q_{\mathcal{H}}$ then. Instantiating \mathcal{H} with H , and \mathcal{O} with F and R , respectively, we can then write

$$\text{PR}^0 = \mathcal{A}^{[\mathbf{H}_0 R \mathbf{H}_1 \dots R \mathbf{H}_{q_{\mathcal{O}}}]}$$
 and $\text{PR}^1 = \mathcal{A}^{[\mathbf{H}_0 F \mathbf{H}_1 \dots F \mathbf{H}_{q_{\mathcal{O}}}]}$.

For the proof, we introduce certain hybrid games. For this purpose, we introduce the following alternative (*stateful* and R -dependent) instantiation H' of \mathcal{H} . To start with, H' is set to be equal to H , but whenever R is queried on some input x , H' is *reprogrammed* at the point $h(k, x)$ to the value $H'(h(k, x)) := R(x)$. For any i , we now define the two hybrid games

$$\begin{aligned} \text{PR}_i^2 &:= \mathcal{A}^{[\mathbf{H}_0 R \dots R \mathbf{H}_i F \mathbf{H}'_{i+1} F \dots F \mathbf{H}'_{q_{\mathcal{O}}}]}, \\ \widetilde{\text{PR}}_i^2 &:= \mathcal{A}^{[\mathbf{H}_0 R \dots R \mathbf{H}_i R \mathbf{H}'_{i+1} F \dots F \mathbf{H}'_{q_{\mathcal{O}}}]}. \end{aligned}$$

and also spell out

$$\text{PR}_{i+1}^2 = \mathcal{A}^{[\mathbf{H}_0 R \dots R \mathbf{H}_i R \mathbf{H}_{i+1} F \dots F \mathbf{H}'_{q_{\mathcal{O}}}]}$$

to emphasize its relation to $\widetilde{\text{PR}}_i^2$. We note that in all of the above, the first occurrences of \mathcal{H} and \mathcal{O} are instantiated with R and H , respectively, but at some point we switch to R and H' instead.

The extreme cases match up the games we are interested in. Indeed,

$$\text{PR}_0^2 = \mathcal{A}^{[\mathbf{H}_0 F \mathbf{H}'_1 \dots F \mathbf{H}'_{q_{\mathcal{O}}}] = \mathcal{A}^{[\mathbf{H}_0 F \mathbf{H}_1 \dots F \mathbf{H}_{q_{\mathcal{O}}}] = \text{PR}^1,$$

where we exploit that there are no queries to R and thus H' remains equal to H , and, by definition,

$$\text{PR}_{q_{\mathcal{O}}}^2 = \mathcal{A}^{[\mathbf{H}_0 R \mathbf{H}_1 \dots R \mathbf{H}_{q_{\mathcal{O}}}] = \text{PR}^0.$$

Our goal is to prove the closeness of the following games

$$\text{PR}^1 = \text{PR}_0^2 \approx \widetilde{\text{PR}}_0^2 \approx \text{PR}_1^2 \cdots \approx \text{PR}_{q_{\mathcal{O}}-1}^2 \approx \widetilde{\text{PR}}_{q_{\mathcal{O}}-1}^2 \approx \text{PR}_{q_{\mathcal{O}}}^2 = \text{PR}^0.$$

We do this by means of applying Lemma 4 and 5, which we state here and prove further down.

Lemma 4. For each $0 \leq i < q_{\mathcal{O}}$,

$$\left| \Pr [1 \leftarrow \text{PR}_i^2] - \Pr [1 \leftarrow \widetilde{\text{PR}}_i^2] \right| \leq 2 \sqrt{\sum_{1 \leq j \leq i} q_j^{\mathcal{H}} \epsilon}.$$

Lemma 5. For each $0 \leq i < q_{\mathcal{O}}$,

$$\left| \Pr [1 \leftarrow \widetilde{\text{PR}}_i^2] - \Pr [1 \leftarrow \text{PR}_{i+1}^2] \right| \leq 2 q_{i+1}^{\mathcal{H}} \sqrt{q_{\mathcal{O}} \epsilon}.$$

Indeed, by repeated applications of these lemmas, and additionally using that $q_0^{\mathcal{H}} + \cdots + q_i^{\mathcal{H}} \leq q_{\mathcal{H}}$ for all $0 \leq i \leq q_{\mathcal{O}}$, we obtain

$$\begin{aligned} \left| \Pr [1 \leftarrow \text{PR}^1] - \Pr [1 \leftarrow \text{PR}^0] \right| &\leq 2 \sum_{i=0}^{q_{\mathcal{O}}} \sqrt{\sum_{1 \leq j \leq i} q_j^{\mathcal{H}} \epsilon} + 2 \sum_{i=0}^{q_{\mathcal{O}}} q_{i+1}^{\mathcal{H}} \sqrt{q_{\mathcal{O}} \epsilon} \\ &\leq 2 \sqrt{q_{\mathcal{O}}^2 q_{\mathcal{H}} \epsilon} + 2 \sqrt{q_{\mathcal{H}}^2 q_{\mathcal{O}} \epsilon} \end{aligned}$$

which concludes the claim of Theorem 2 when incorporating the factor-2 increase in $q_{\mathcal{H}}$ and $q_{\mathcal{O}}$ due to switching to a static \mathcal{A} . \square

It remains to prove Lemma 4 and 5, which we do below. In both proofs, we use the *gentle measurement lemma* [Wil11, Lemma 9.4.1], which states that if a projective measurement has a very likely outcome then the measurement causes only little disturbance on the state. More formally, for any density operator ρ and any projector P , where $p := \text{tr}(P\rho P)$ then is the probability to

observe the outcome associated with P when measured using the measurement $\{P, \mathbb{I} - P\}$, the trace distance between the original state ρ and the post-measurement state $\rho' := P\rho P/p$ is bounded by $\sqrt{1-p}$. This in turn implies that ρ and ρ' can be distinguished with an advantage $\sqrt{1-p}$ only.

The proof of Lemma 4 additionally makes use of Zhandry's compressed oracle technique [Zha19]. It is out of scope of this work to give a self-contained description of this technique; we refer to the original work [Zha19] instead, or to [CFHL21], which offers an alternative concise description. At the core is the observation that one can *purify* the random choice of the function H and then, by switching to the Fourier basis and doing a suitable measurement, one can check whether a certain input x has been "recorded" in the database (mind though that such a measurement disturbs the state). If the outcome is negative then the oracle is still in a uniform superposition over all possible hash values for x , and as a consequence, when removing the purification by doing a full measurement of H (in the computational basis), $H(x)$ is ensured to be a "fresh" uniformly random value, with no information on $H(x)$ having been leaked in prior queries.

In the proof of Lemma 4, we use this technique to check whether *prior* to the crucial query, which is to F in one and to R in the other game, there was a query to H that would reveal the difference, and we use (6) to argue that it is unlikely that such a query occurred. Since this measurement has a likely outcome, it is also ensured by the gentle measurement lemma that this measurement causes little disturbance.

Proof (of Lemma 4). For convenience, we refer to the *crucial query* as the respective query to F and R that differs between

$$\text{PR}_i^2 = \mathcal{A}^{[\mathbf{H}_0 R \dots R \mathbf{H}_i F \mathbf{H}'_{i+1} F \dots F \mathbf{H}'_{q_{\mathcal{O}}}]}$$
 and $\widetilde{\text{PR}}_i^2 = \mathcal{A}^{[\mathbf{H}_0 R \dots R \mathbf{H}_i R \mathbf{H}'_{i+1} F \dots F \mathbf{H}'_{q_{\mathcal{O}}}]}$.

Furthermore, we let x be the input to that query, and we set $w := h(k, x)$, with k being the key chosen and used by F . Note that up to this very query, the two games are identical. Also, by (7) it is ensured that for any prior query x' to R it holds that $h(k, x') \neq w$.

First, we consider the games \mathbf{G}^1 and $\widetilde{\mathbf{G}}^1$ that work exactly as PR_i^2 and $\widetilde{\text{PR}}_i^2$, respectively, except that, at the beginning of the games we set up the compressed oracle and answer all queries made to H prior to the crucial query using the compressed oracle. Then, once x is received during the crucial query, we do a full measurement of the purified (i.e. uncompressed) oracle in order to obtain the function H , which is then to be used in the remainder of the games. We note that setting up the function H' is then necessarily also deferred to after this measurement, where H' is then set to be equal to H , except that for any prior query x' to R it is reprogrammed to $H'(h(k, x')) := R(x')$. Only once H has been measured and H' set up as above, is the crucial query then actually answered.

It follows from basic properties of the compressed oracle that the respective output distributions of \mathbf{G}^1 and $\widetilde{\mathbf{G}}^1$ match with those of PR_i^2 and $\widetilde{\text{PR}}_i^2$.

Then, we define \mathbf{G}^2 and $\widetilde{\mathbf{G}}^2$ from \mathbf{G}^1 and $\widetilde{\mathbf{G}}^1$, respectively, by introducing one more measurement. Namely, right after x is sent by \mathcal{A} and before H is measured, we measure in the compressed oracle whether the input $w = h(k, x)$ has been recorded in the database, and in case of a positive outcome, the game aborts. By the gentle measurement lemma (and basic properties of the trace distance),

$$|\Pr[1 \leftarrow \mathbf{G}^1] - \Pr[1 \leftarrow \mathbf{G}^2]| \leq \sqrt{\Pr[\mathbf{G}^2 \text{ aborts}]}$$

and similarly for $\widetilde{\mathbf{G}}^1$ and $\widetilde{\mathbf{G}}^2$, where $\widetilde{\mathbf{G}}^2$ aborts with the same probability as \mathbf{G}^2 .

By basic properties, after $t := q_0^{\mathcal{H}} + \dots + q_i^{\mathcal{H}}$ queries to the compressed oracle, no more than t values have been recorded. I.e., if we were to measure, for the sake of the argument, the entire compressed oracle to obtain the full database D , it would hold that $\text{supp}(D) := \{u \mid D(u) \neq \perp\}$ has cardinality at most t . Since k has not been used yet and so is still freshly random (i.e., independent of x and D), the high-entropy condition (6) then ensures that

$$\Pr[\widetilde{\mathbf{G}}^2 \text{ abort}] = \Pr[\mathbf{G}^2 \text{ abort}] = \Pr[w \in \text{supp}(D)] \leq \sum_{j < i} q_j^{\mathcal{H}} \epsilon.$$

It remains to show that \mathbf{G}^2 and $\widetilde{\mathbf{G}}^2$ behave identically conditioned on not aborting. The only difference between the two games is that in \mathbf{G}^2 the crucial query is answered with $y := H(h(k, x)) =$

$H(w)$ and H' is *not* reprogrammed at the point w , while in $\tilde{\mathbf{G}}^2$ the crucial query is answered with $y := R(x)$ and H' is reprogrammed at the point w to $H'(w) := R(x)$. We argue that this difference is not noticeable by \mathcal{A} .

First, we note that y is a fresh random value in both games. In the former game it is because, conditioned on not aborting, the compressed oracle at the register $h(k, x)$ is \perp , and so when uncompressing and measuring to obtain H , the hash value $H(w)$ will be a fresh random value. In the latter game it is because $R(x)$ is a truly random function and, due to (7), x has not been queried to R before.

Second, we observe that $y = H'(w)$ in both games. Indeed, in $\tilde{\mathbf{G}}^2$ this holds by definition; in \mathbf{G}^2 it holds because $H'(w) = H(w)$, which follows from the fact that H' is reprogrammed only at points $w' = h(k, x')$ with x' being a prior query to R , but then (7) ensures that $w' \neq w$.

Thus, in both games, from \mathcal{A} 's perspective, the tuple $(k, y, H', H \setminus w)$ of random variables has the same distribution, where $H \setminus w$ refers to the function (table of) H but with the value at the point w removed. The only difference is that in one game $H'(w) = H(w)$ and in the other not (necessarily). However, the future behavior of \mathcal{A} in both games only depends on $(k, y, H', H \setminus w)$, and thus \mathcal{A} behaves the same way in both games. Here we are exploiting that the future hash queries by \mathcal{A} are to H' (and not to H anymore), and, once more, we are using the restriction (7), here to ensure that for any future F -query x' by \mathcal{A} , it holds that $h(k, x') \neq w$, and thus the response does not depend on $H(w)$. Thus, $H(w)$ does indeed not affect \mathcal{A} 's behavior after the crucial query.

Exploiting that $\text{PR}_i^2 = \mathbf{G}^1 \approx \mathbf{G}^2 = \tilde{\mathbf{G}}^2 \approx \tilde{\mathbf{G}}^1 = \widetilde{\text{PR}}_i^2$, with the approximations bounded as discussed further up, we obtain the claimed closeness claim. This concludes the proof. \square

Proof of Lemma 5. In order to show the closeness between $\widetilde{\text{PR}}_i^2$ and PR_{i+1}^2 , we define the intermediate games

$$\mathbf{G}_{i,j} := \mathcal{A}^{[\mathbf{H}_0 R \dots \mathbf{H}_i R \mathbf{H}'_{i,j} \mathbf{H}_{i,j} F \dots F \mathbf{H}'_{i_0}]}$$

for $0 \leq j \leq m := q_{i+1}^{\mathcal{H}}$, where $\mathbf{H}'_{i,j}$ and $\mathbf{H}_{i,j}$ consists of j and $m - j$ copies of H' and H respectively. Note that for the extreme cases we have

$$\mathbf{G}_{i,0} = \widetilde{\text{PR}}_i^2 \quad \text{and} \quad \mathbf{G}_{i,m} = \text{PR}_{i+1}^2.$$

Thus, it suffices to show closeness between $\mathbf{G}_{i,j}$ and $\mathbf{G}_{i,j+1}$ for any $0 \leq j < m$. Note that they only differ at one query, which is either to H' or to H , which we will refer to as the *crucial query* for convenience. In the remainder, i and j are arbitrary (in the considered ranges) but fixed.

Define the games $\tilde{\mathbf{G}}^1$ and \mathbf{G}^1 from $\mathbf{G}_{i,j}$ and $\mathbf{G}_{i,j+1}$ respectively as follows. Let X be the set of queries x made to R prior to the crucial query, and set $S := \{h(k, x) \mid x \in X\}$. We then measure the crucial query, which may be in a superposition, with the binary measurement that checks whether the crucial query is an element of S , and we abort if this is the case.

In case of a negative outcome, i.e., the crucial query is *not* in S , there is no difference between the reply provided by H and by H' , and thus there is no difference between the two games — and in case of a positive outcome, they both abort. In order to argue that this measurement causes little disturbance, we again use the gentle measurement lemma to argue that

$$|\Pr [1 \leftarrow \mathbf{G}^1] - \Pr [1 \leftarrow \mathbf{G}_{i,j+1}]| \leq \sqrt{\Pr [\mathbf{G}^1 \text{ abort}]},$$

and correspondingly for $\mathbf{G}_{i,j}$ and $\tilde{\mathbf{G}}^1$. So it remains to bound the abort probability. For the purpose of the argument, let us do a full measurement of the query, and let w be the outcome. We note that k has not been used yet, and thus remains a fresh random key, independent of w and X . Thus, using (6),

$$\Pr [\mathbf{G}^1 \text{ abort}] = \Pr [\tilde{\mathbf{G}}^1 \text{ abort}] = \Pr [w \in S] \leq \sum_{x \in X} \Pr [w = h(k, x)] \leq q_{\mathcal{O}} \epsilon.$$

Adding up this error term over the sequence $\mathbf{G}_{i,0} \approx \dots \approx \mathbf{G}_{i,m}$ of approximations, the proof is concluded. \square

Acknowledgments

JD was funded by the ERC-ADG project ALGSTRONGCRYPTO (project number 740972). YHH was funded by the Dutch Research Agenda (NWA) project HAPKIDO (project number NWA.1215.18.002), which is financed by the Dutch Research Council (NWO).

References

- [ABB⁺17] Erdem Alkim, Nina Bindel, Johannes Buchmann, Özgür Dagdelen, Edward Eaton, Gus Gutoski, Juliane Krämer, and Filip Pawlega. Revisiting TESLA in the quantum random oracle model. In Tanja Lange and Tsuyoshi Takagi, editors, *Post-Quantum Cryptography*, pages 143–162. Springer, 2017.
- [ABB⁺20] Erdem Alkim, Paulo S. L. M. Barreto, Nina Bindel, Juliane Krämer, Patrick Longa, and Jefferson E. Ricardini. The lattice-based digital signature scheme qTESLA. In Mauro Conti, Jianying Zhou, Emiliano Casalicchio, and Angelo Spognardi, editors, *Applied Cryptography and Network Security*, pages 441–460. Springer, 2020.
- [ABKM21] Gorjan Alagic, Chen Bai, Jonathan Katz, and Christian Majenz. Post-quantum security of the Even-Mansour cipher. Cryptology ePrint Archive, Report 2021/1601, 2021. <https://ia.cr/2021/1601>.
- [BDF⁺11] 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*, pages 41–69. Springer, 2011.
- [CFHL21] Kai-Min Chung, Serge Fehr, Yu-Hsuan Huang, and Tai-Ning Liao. On the compressed-oracle technique, and post-quantum security of proofs of sequential work. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021*, pages 598–629. Springer, 2021.
- [GHP18] Federico Giacon, Felix Heuer, and Bertram Poettering. KEM combiners. In *IACR International Workshop on Public Key Cryptography*, pages 190–218. Springer, 2018.
- [LS19] Vadim Lyubashevsky and Peter Schwabe. Round 2 official comment: qTESLA. <https://csrc.nist.gov/CSRC/media/Projects/Post-Quantum-Cryptography/documents/round-2/official-comments/qTESLA-round2-official-comment.pdf>, 2019. Accessed: 18-05-2022.
- [Wil11] Mark M Wilde. From classical to quantum Shannon theory. *arXiv preprint arXiv:1106.1445*, 2011.
- [Zha19] Mark Zhandry. How to record quantum queries, and applications to quantum indistinguishability. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019*, pages 239–268. Springer, 2019.