Post-Quantum Security of Tweakable Key-Alternating Feistel Ciphers in the Multi-Key Setting

Rentaro Shiba[1,2][0000−0003−2306−9978] and Tetsu Iwata[1][0000−0002−4729−0979]

[1] Nagoya University, Nagoya, Japan,
shiba.rentaro.k7@s.mail.nagoya-u.ac.jp, tetsu.iwata@nagoya-u.jp
[2] Mitsubishi Electric Corporation, Kamakura, Japan,
Shiba.Rentaro@dc.MitsubishiElectric.co.jp

Abstract. In this paper, we prove the post-quantum security of the Tweakable Key-Alternating Feistel cipher (TKAF) with a public random function in the Q1 model, under the assumption that the adversary is given quantum access to the internal primitive. Specifically, our target is the TKAF studied in the classical setting by Yan *et al.* (ACNS 2020), where the tweak is injected into the round-key XOR via a hash function of ϵ-AXU family. Our proof draws on the post-quantum security proof for the (non-tweakable) key-alternating Feistel ciphers by Basak *et al.* (ASIACRYPT 2025), and adapts it to the tweakable setting and further to the multi-key setting, where an adversary can access multiple classical oracles. As a result, we prove that the 3-round TKAF is post-quantum TPRP-secure and the 4-round TKAF is post-quantum STPRP-secure. Specifically, under the assumption that the adversary is given classical access to ℓ independently specified oracles, at least $\Omega(2^{n/3}/\ell^{2/3})$ classical and quantum queries or $\Omega(\epsilon^{-1/2})$ classical queries are required to break the post-quantum TPRP security of the 3-round TKAF and to break the post-quantum STPRP security of the 4-round TKAF.

Keywords: tweakable block cipher · key-alternating Feistel cipher · provable security · post-quantum security · multi-key security · multi-user security

1 Introduction

In the field of symmetric-key cryptography, post-quantum security analysis has become a major research topic today. Post-quantum security considers adversaries that have the ability of quantum computation in the offline phase. Depending on the allowed form of online access to the construction, we consider two models: the Q1 model, where online queries are restricted to classical, and the Q2 model, where online queries are allowed to be in quantum superposition.

For Feistel networks, which constitute the most fundamental paradigm for building block ciphers, the quantum security analyses have been intensively studied in both the Q1 and Q2 models. Among Feistel networks, the Luby-Rackoff

construction (LR) is a canonical, idealized model in which the round functions are secret random functions. In the Q2 model, Kuwakado and Morii [15] proposed a polynomial-time distinguisher for 3-round LR via Simon's algorithm [21], while qPRP security of 4-round LR is known only under a non-adaptive setting [8] [1].

Although the LR is convenient for studying theoretical aspects of Feistel networks, it is far from practical block-cipher designs. This motivates analyses of more practical Feistel-based constructions such as key-alternating Feistel ciphers (KAFs). In particular, Ito *et al.* [14] showed a polynomial-time quantum distinguisher on 4-round KAF in a model where the adversary has quantum-oracle access to both the encryption and decryption. Moreover, very recently, Basak *et al.* [7] proved PRP security of 3-round KAF and SPRP security of 4-round KAF in the Q1 model by adapting the hybrid argument of Alagic *et al.* [2], a technique also used for other block cipher constructions [3,4].

Motivation. Modern uses of block ciphers often require handling additional public inputs beyond the plaintext, such as nonces, counters, sector numbers, and associated data. This motivates the study of tweakable block ciphers, which extend block ciphers by adding an auxiliary public input called a tweak. Tweakable block ciphers are formalized by Liskov *et al.* [16] and have found broad applications including enciphering modes [9,10] and authenticated encryption schemes [5,18,19].

Given the central role of Feistel networks and the recent progress on the quantum security of KAF in the Q1 model, it is natural to ask whether similar guarantees can be obtained for *tweakable* variants of KAF [22]. In the classical setting, Yan *et al.* [22] studied tweakable KAF by injecting the tweak into the round-key XOR operations via (almost) universal hashing, and proved security bounds for the resulting constructions.

Our contribution. In this paper, we focus on the same tweakable KAF construction for which Yan et al. proved classical security, and we prove the post-quantum security in the Q1 model. In the rest of the paper, we refer to the target construction, the tweakable KAF with a public random function, as TKAF. The details of TKAF are given as follows.

TKAF. Let $f : \{0,1\}^n \to \{0,1\}^n$ be a public random function and $\mathcal{H} = \{H_{k_r} : \mathcal{T} \to \{0,1\}^n \mid k_r \in \mathcal{K}\}$ be ϵ-AXU (The formal definition of ϵ-AXU is given in Section 2.1). Let $(x_{2(r-1)}, x_{2(r-1)+1}) \in \{0,1\}^{2n}$ be an input state for the r-th round, where $x_{2(r-1)}, x_{2(r-1)+1} \in \{0,1\}^n$. Then, the r-th round updates the state by using a function $F^f_{k_r}$ defined as follows.

$$F^f_{k_r} : (x_{2(r-1)}, x_{2(r-1)+1}) \mapsto (x_{2(r-1)+1}, f(x_{2(r-1)+1} \oplus H_{k_r}(t)) \oplus x_{2(r-1)})$$

[1] Hosoyamada and Iwata [13] claimed qPRP security for 4-round LR. Later, Bhaumik *et al.* [8] pointed out the technical issue in the proof and proved qPRP security only for non-adaptive adversaries.

In TKAF, the above operation is sequentially applied up to an arbitrary r. Thus, the r-round TKAF: r-TKAF is defined as

$$r\text{-}\mathsf{TKAF}(x_0, x_1) \coloneqq F^f_{k_r} \circ \cdots \circ F^f_{k_1}(x_0, x_1).$$

We let $r\text{-}\mathsf{TKAF}(x_0, x_1) \coloneqq (x_r, x_{r+1})$. The 4-TKAF is shown in Figure 1. We label the intermediate values as in Figure 1, and we use them in the rest of the paper.

Our proof can be viewed as an extension of the post-quantum security proof of Basak *et al.* for (non-tweakable) KAF with a single public random function. At a high level, Basak *et al.* adapt the proof techniques of [2], which heavily rely on the arbitrary reprogramming lemma and the resampling lemma, to the KAF setting. We generalize their approach to the tweakable setting where each round key is handled by ϵ-AXU hash functions taking a tweak value as input.

Whereas Basak *et al.* work in the single-key setting, we consider the multi-key setting, which generalizes the single-key case, where an adversary is given classical access to multiple keyed instances (resp. tweakable random permutations) in the real world (resp. the ideal world). A systematic post-quantum treatment of such multi-key access has been recently initiated in [20], and it better reflects the practice that symmetric-key primitives are typically deployed under many independent keys.

Let q, p, and ℓ be the number of classical queries, the number of quantum queries, and the number of classical oracles. Then, our contributions are as follows:

- We prove that, for any adversary $\mathcal{A}$ in the Q1 model and the multi-key setting, $\mathsf{Adv}^{\mathsf{pq\text{-}tprp,mk}}_{\text{3-TKAF}}(\mathcal{A}) \leq O(p\sqrt{q/2^n} + \ell q\sqrt{p/2^n} + q^2\ell/2^n + \epsilon q^2)$, where $\mathsf{Adv}^{\mathsf{pq\text{-}tprp,mk}}_{\text{3-TKAF}}(\mathcal{A})$ denotes the post-quantum tweakable pseudorandom permutation (PQ-TPRP) advantage of $\mathcal{A}$.
- We prove that, for any adversary $\mathcal{A}$ in the Q1 model and the multi-key setting, $\mathsf{Adv}^{\mathsf{pq\text{-}stprp,mk}}_{\text{4-TKAF}}(\mathcal{A}) \leq O(p\sqrt{q/2^n} + \ell q\sqrt{p/2^n} + q^2\ell/2^n + \epsilon q^2)$, where $\mathsf{Adv}^{\mathsf{pq\text{-}stprp,mk}}_{\text{4-TKAF}}(\mathcal{A})$ denotes the post-quantum strong tweakable pseudorandom permutation (PQ-STPRP) advantage of $\mathcal{A}$.

Our bounds imply that, at least $\Omega(2^{n/3}/\ell^{2/3})$ classical and quantum queries or $\Omega(\epsilon^{-1/2})$ classical queries are required to break the PQ-TPRP security of the 3-TKAF and to break the PQ-STPRP security of the 4-TKAF. We do not claim that these bounds are tight. In particular, we do not provide attacks showing that the above query complexities are optimal.

Our result also has an implication for KAF with a single public random function. Indeed, TKAF can be viewed as a generalization of KAF, and the security of KAF with a single public random function corresponds to that of our TKAF with an empty tweak. Thus, our result can also be interpreted as showing the post-quantum security of KAF with a single public random function in the multi-key setting. In this sense, our result extends Basak *et al.*'s single-key analysis of KAF to the multi-key setting. Although TKAF here may not directly correspond to widely used concrete block ciphers, the (non-tweakable) KAF is an idealized keyed Feistel construction motivated by concrete Feistel block ciphers such as [1,6,17]. Therefore, its post-quantum multi-key security is useful for un-

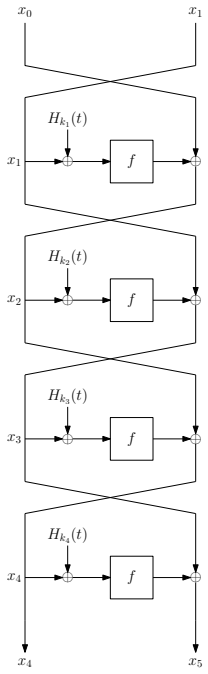


Fig. 1. 4-TKAF

derstanding the idealized behavior of such constructions when many independent keys are used.

Paper organization. Section 2 gives basic notations and definitions, some preliminaries on the quantum security model, and useful lemmas that we use in our proofs. Section 3 provides the PQ-TPRP security proof of the 3-round TKAF and Section 4 provides the PQ-STPRP security proof of the 4-round TKAF.

2 Preliminaries

2.1 Basic Notation and Definition

Throughout the paper, we use the following notation.

- $|S|$: the cardinality of a finite set S.
- $x \leftarrow_{\$} S$: sampling an element x uniformly at random from a finite set S.
- $\mathcal{T}$: a tweak space.
- $\tilde{\mathcal{P}}_{\mathcal{T},n}$: a set of all tweakable permutations on $\{0,1\}^n$ with a tweak space of $\mathcal{T}$.
- $\mathcal{F}_n$: a set of all functions on $\{0,1\}^n$ to $\{0,1\}^n$
- $[n,m]$: the set of integers $\{n, n+1, \ldots, m\}$ (for any integers $n \leq m$).

Uniform AXU Hash Functions. Let $\mathcal{H} = \{H_k : \mathcal{T} \to \{0,1\}^n \mid k \in \mathcal{K}\}$ be a set of hash functions from a set $\mathcal{T}$ to $\{0,1\}^n$ indexed by a key $k \in \mathcal{K}$, where $\mathcal{K}$ is the key space. $\mathcal{H}$ is said to be uniform if, for any $t \in \mathcal{T}$ and $y \in \{0,1\}^n$, it holds that

$$\Pr_{k \leftarrow_{\$} \mathcal{K}}[H_k(t) = y] = 2^{-n}.$$

$\mathcal{H}$ is said to be ϵ-almost XOR-universal (ϵ-AXU) if for all distinct $t_1, t_2 \in \mathcal{T}$ and all $y \in \{0,1\}^n$,

$$\Pr_{k \leftarrow_{\$} \mathcal{K}}[H_k(t_1) \oplus H_k(t_2) = y] \leq \epsilon.$$

In particular, $\mathcal{H}$ is XOR-universal if $\epsilon = 2^{-n}$.

2.2 The Adversary Model

Let $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell$ be ℓ ideal constructions (ideal tweakable permutations), and F be a real world construction. We assume that F^f is constructed from an ideal primitive f (public random function), and it is instantiated by ℓ independently chosen keys $K^1, \ldots, K^\ell$. We specifically denote the real-world construction which is instantiated with f and K^l as $\mathsf{F}^f_{K^l}$ for each $l \in [1, \ell]$. Then, we assume the adversary model where the adversary $\mathcal{A}$ is given classical access to $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell$ (in the ideal world) or $\mathsf{F}^f_{K^1}, \ldots, \mathsf{F}^f_{K^\ell}$ (in the real world), and quantum access to f, *i.e.*, it is given the unitary: $U_f |x\rangle |y\rangle = |x\rangle |y \oplus f(x)\rangle$. If the classical access to $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell$ or $\mathsf{F}^f_{K^1}, \ldots, \mathsf{F}^f_{K^\ell}$ is restricted to the forward direction, we let $\mathcal{A}$'s post-quantum tweakable pseudorandom permutation advantage in the multi-key setting be

$$\mathsf{Adv}^{\mathsf{pq\text{-}tprp,mk}}_{\mathsf{F}^f}(\mathcal{A}) \coloneqq \left| \Pr\left[\mathcal{A}^{\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell, f} = 1\right] - \Pr\left[\mathcal{A}^{\mathsf{F}^f_{K^1}, \ldots, \mathsf{F}^f_{K^\ell}, f} = 1\right] \right|.$$

If the adversary has classical access to $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell$ or $\mathsf{F}^f_{K^1}, \ldots, \mathsf{F}^f_{K^\ell}$ in both directions, we let $\mathcal{A}$'s post-quantum strong tweakable pseudorandom permutation advantage in the multi-key setting be

$$\mathsf{Adv}^{\mathsf{pq\text{-}stprp,mk}}_{\mathsf{F}^f}(\mathcal{A}) \coloneqq \left| \Pr\left[\mathcal{A}^{\tilde{\Pi}_1^{\pm}, \ldots, \tilde{\Pi}_\ell^{\pm}, f} = 1\right] - \Pr\left[\mathcal{A}^{(\mathsf{F}^f_{K^1})^{\pm}, \ldots, (\mathsf{F}^f_{K^\ell})^{\pm}, f} = 1\right] \right|.$$

Adaptivity. In this study, we assume a fully adaptive adversary that can determine both the oracle to be queried and the query input itself based on the transcript so far and its internal state. Following Basak et al.'s proof [7], we use the adaptive-to-static reduction of [11], which states that, for m oracles, regardless of whether they are classical or quantum, an adaptive adversary making q queries to $m\,(\geq 1)$ oracles is functionally equivalent to a static adversary making mq queries. For this reduction, the ℓ classical oracles can be viewed as a single indexed classical oracle, where the oracle index is treated as part of the classical query input, since our model allows the adversary to choose the oracle to be queried in each classical query.

2.3 Useful Lemmas

Arbitrary Reprogramming [2]. Let $\mathcal{D}$ be a distinguisher that runs the following ARP game.

Phase 1: $\mathcal{D}$ specifies a function $F_0 = F : \mathcal{X} \to \mathcal{Y}$ and a randomized algorithm $\mathcal{B}$ which outputs a set $B \subset \mathcal{X} \times \mathcal{Y}$ where each $x \in \mathcal{X}$ is the first element of at most one tuple in B. Let $B_1 = \{x \mid \exists y : (x, y) \in B\}$ and $\varepsilon = \max_{x \in \mathcal{X}}\{\Pr_{B \leftarrow \mathcal{B}}[x \in B_1]\}$.

Phase 2: $\mathcal{B}$ is run to obtain B. Let $F_1 = F^{(B)}$, where $F^{(B)}$ is locally redefined version of F, defined as

$$F^{(B)}(x) \coloneqq \begin{cases} y & \text{if } (x, y) \in B, \\ F(x) & \text{otherwise.} \end{cases}$$

The challenger samples a uniform bit $b \in \{0, 1\}$ and $\mathcal{D}$ is given quantum access to F_b.

Phase 3: $\mathcal{D}$ loses access to F_b, and receives the randomness r used to invoke $\mathcal{B}$. Then $\mathcal{D}$ outputs b', a guess of b.

Lemma 1 (ARP lemma [2]). *For any $\mathcal{D}$ in the ARP game making q queries when its oracle is F_0, it holds that*

$$\mathsf{Adv}^{\mathsf{arp}}(\mathcal{D}) \coloneqq \left|\mathsf{Pr}[\mathcal{D}^{F_b} = 1 \mid b = 1] - \mathsf{Pr}[\mathcal{D}^{F_b} = 1 \mid b = 0]\right| \leq 2q\sqrt{\varepsilon}\,. \tag{1}$$

The proof of Lemma 1 is given in [2, Section 4.1].

Adaptive Resampling [12]. Let $\mathcal{D}$ be a distinguisher that runs the following ARS game.

ARS *game.* Let $F : \mathcal{X} \to \mathcal{Y}$ denote a uniform random function, and set $F_0 = F_1 = F$. Let Φ denote a sampling oracle defined as follows: on the input distribution D_r on $\mathcal{X}$ used in the r-th ($r \in [1, R]$) execution, run the following procedure.

1. sample $x \leftarrow D_r$ and $y \leftarrow_{\$} \mathcal{Y}$
2. set $F_1 \leftarrow F_1^{(x,y)}$, where $F_1^{(x,y)}$ is defined as

$$F_1^{(x,y)}(x^*) \coloneqq \begin{cases} y & \text{if } x^* = x \\ F_1(x^*) & \text{otherwise} \end{cases}$$

3. return x

This algorithm is stateful, since it modifies the definition of F_1. It may also override a previous definition.

The ARS game is a random experiment played by a challenger and a distinguisher $\mathcal{D}$.

Phase 1: The challenger samples a uniform bit $b \in \{0, 1\}$ and $\mathcal{D}$ is given quantum access to F_b and classical access to Φ (at the r-th execution of Φ, the input to Φ is the distribution D_r).

Phase 2: $\mathcal{D}$ loses access to F_b. $\mathcal{D}$ outputs b', a guess of b.

Lemma 2 (ARS lemma [12, Theorem 1]). *In the* ARS *game, let $\mathcal{D}$ make at most R queries to Φ and q quantum queries to F_b.*

Let q_r be the number of quantum queries to F_b issued between the $(r-1)$-st and r-th queries to Φ, and define $\hat{q}_r := \sum_{i\in[0,r-1]} q_i$. (note that $q_0 = 0$). Moreover let $\mu_r := \mathbb{E}[\max_{x\in\mathcal{X}} D_r(x)]$. Then,

$$\begin{aligned}\mathsf{Adv}^{\mathsf{ars}}(\mathcal{D}) &:= \left|\Pr\left[\mathcal{D}^{F_b,\Phi}=1 \mid b=1\right] - \Pr\left[\mathcal{D}^{F_b,\Phi}=1 \mid b=0\right]\right| \\ &\leq \sum_{r\in[1,R]} \left(\sqrt{\hat{q}_r\mu_r} + \frac{1}{2}\hat{q}_r\mu_r\right).\end{aligned}$$

The proof of Lemma 2 is given in [12, Section 5].

2.4 The (Weak) Collisions-in-the-Dark Games

The collisions-in-the-dark (CitD) games are introduced by Basak *et al.* in [7]. In our proof, we use only the weak CitD (WCitD). Since our modified game naturally involves multiple targets arising from multiple independent keys, we refer to our variant as the *multi-target* WCitD (mWCitD) game. Note that mWCitD is a straightforward extension of the original WCitD game. Although the proof is trivial, we show the proof in Appendix A for completeness.

mWCitD *game.* Let $\mathcal{X} = \mathcal{X}_0 \times \mathcal{X}_1$ and $\mathcal{Y}$ be two finite sets, where $\mathcal{X}_1 = \{0,1\}^n$. Let $F : \mathcal{X}_1 \to \mathcal{Y}$ be a random oracle. Fix some relation $\mathbb{R} \in \mathcal{X}^2 \times \mathcal{Y}^2$. The WCitD game is played between a challenger and an adversary $\mathcal{A}$.

Phase 1. The adversary $\mathcal{A}$ makes p quantum queries to the random oracle F and returns a set $\mathcal{X}^* \subseteq \mathcal{X}$ such that $|\mathcal{X}^*| = q$.

Phase 2. The challenger samples ℓ keys $\tilde{k}^1, \ldots, \tilde{k}^\ell \leftarrow_\$ \mathcal{X}_1$. $\mathcal{A}$ wins if the following event holds:

$$\begin{aligned}\mathsf{mWC}_{\mathbb{R}}(\mathcal{X}^*) : \exists x = (x_0, x_1), x' = (x'_0, x'_1) \in \mathcal{X}^*, \exists l \in [1,\ell] \\ \text{such that } (x, x', F(k^l \oplus x_1), F(k^l \oplus x'_1)) \in \mathbb{R}.\end{aligned}$$

Lemma 3 (Multi-Target Weak Collision-in-the-Dark (adapted from [7])). *For any adversary $\mathcal{A}$ in the* mWCitD *game with respect to the relation $\mathbb{R}$, making p queries and returning a set of size q, it holds that*

$$\Pr_{\mathcal{X}^*\leftarrow\mathcal{A}}[\mathsf{mWC}_{\mathbb{R}}(\mathcal{X}^*)] \leq 1.5\ell q\sqrt{\frac{p}{2^n}} + q^2\ell \max_{x,x'\in\mathcal{X}} \Pr_{\substack{r\leftarrow\mathcal{Y}\\ s\leftarrow\mathcal{X}_1}}[(x,x',r,F(s)) \in \mathbb{R}].$$

3 Post-Quantum Security of 3-round TKAF

Theorem 1 (PQ-TPRP security of 3-TKAF). *Let $\mathcal{A}'$ be an adversary who can make q' classical queries to $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell \leftarrow \tilde{\mathcal{P}}_{\mathcal{T},2n}$ or 3-TKAFs instantiated*

as $\mathsf{F}^f_{K^1}, \ldots, \mathsf{F}^f_{K^\ell}$ *where* $K^1 = (k_1^1, k_2^1, k_3^1), \ldots, K^\ell = (k_1^\ell, k_2^\ell, k_3^\ell) \leftarrow \{0,1\}^{3\kappa}$ *and* $f \leftarrow \mathcal{F}_n$*, and* p' *quantum queries to* f*. Then,*

$$\mathsf{Adv}^{\mathsf{pq\text{-}tprp},\mathsf{mk}}_{\mathsf{F}^f}(\mathcal{A}') \leq 4p'\sqrt{\frac{4q'}{2^n}} + 8\ell q'\sqrt{\frac{2p'}{2^n}} + \frac{36{q'}^2\ell}{2^n} + 12\epsilon{q'}^2.$$

Proof. As explained in Section 2.2, we use the adaptive-to-static reduction of [11]. Thus, it suffices to analyze a modified adversary $\mathcal{A}$ that is functionally equivalent to the fully adaptive adversary $\mathcal{A}'$ and makes at most $q = 2q'$ classical queries and $p = 2p'$ quantum queries. Here, the factor 2 is the number of oracle interfaces obtained after regarding the ℓ classical oracles as a single indexed classical oracle, as discussed in Section 2.2.

In the ideal world, the adversary is given classical access to ℓ ideal constructions, which we denote as $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell \leftarrow_\$ \tilde{\mathcal{P}}_{\mathcal{T},2n}$. On the other hand, in the real world, the adversary is given classical access to $\mathsf{F}^f_{K^1}, \ldots, \mathsf{F}^f_{K^\ell}$ where $f \leftarrow_\$ \mathcal{F}_n$ and $K^1, \ldots, K^\ell \leftarrow_\$ \{0,1\}^{r\kappa}$. The adversary is given quantum access to f in both worlds.

As in [20], we define an oracle selection operator as p which selects one from the ℓ independent oracles. The operation p outputs an index in $[1,\ell]$, and we write $\mathsf{p}(i) \in [1,\ell]$ as the index of the oracle queried on the i-th classical query. For example, if ℓ tweakable random permutations in the ideal world (resp. keyed oracles in the real world) are given as $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell$ (resp. $\mathsf{F}^f_{K_1}, \ldots, \mathsf{F}^f_{K_\ell}$), then the i-th query is answered by $\tilde{\Pi}_{\mathsf{p}(i)}$ (resp. $\mathsf{F}^f_{K_{\mathsf{p}(i)}}$).

For each $l \in [1,\ell]$ and $j \in [0,q]$, let j_l denote the number of queries to $\tilde{\Pi}_l$ issued up to the j-th classical query. Clearly, $j_l \leq j$ and $j = \sum_{l\in[1,\ell]} j_l$ (and also $q = \sum_{l\in[1,\ell]} q_l$).

We define a transcript obtained by j_l classical queries to $\tilde{\Pi}_l$, the oracle of the ideal construction indexed by l so far, as

$$\tau^l_{j_l} \coloneqq \left\{(t^{l,1}, (x_0^{l,1}, x_1^{l,1}), (x_3^{l,1}, x_4^{l,1})), \ldots, (t^{l,j_l}, (x_0^{l,j_l}, x_1^{l,j_l}), (x_3^{l,j_l}, x_4^{l,j_l}))\right\},$$

where for all $i \in [1, j_l]$, $t^{l,i}$ is the tweak input, $(x_0^{l,i}, x_1^{l,i})$ is the plaintext, and $(x_3^{l,i}, x_4^{l,i})$ is the ciphertext, of the i-th classical query. We assume that the adversary never issues duplicate classical queries to any fixed oracle. We denote a union of transcripts of all oracles as a multiset $\tau_j \coloneqq \bigcup_{l\in[1,\ell]} \tau^l_{j_l}$.

From the classical transcript obtained from ideal-world oracles, we define the following sets of input and output pairs for the second round and the third round of $\mathsf{F}^f_{K^l}$, for all $l \in [1,\ell]$:

$$\nu^2_{j,K^l} \coloneqq \{(x_2^{l,1} \oplus H_{k_2^l}(t^{l,1}), x_1^{l,1} \oplus x_3^{l,1}), \ldots, (x_2^{l,j_l} \oplus H_{k_2^l}(t^{l,j_l}), x_1^{l,j_l} \oplus x_3^{l,j_l})\}, \quad (2)$$

$$\nu^3_{j,K^l} \coloneqq \{(x_3^{l,1} \oplus H_{k_3^l}(t^{l,1}), x_2^{l,1} \oplus x_4^{l,1}), \ldots, (x_3^{l,j_l} \oplus H_{k_3^l}(t^{l,j_l}), x_2^{l,j_l} \oplus x_4^{l,j_l})\}. \quad (3)$$

We define the reprogramming set, which is a set of all inputs of f to be reprogrammed as a multiset $\nu_{j,K^{[1,\ell]}} \coloneqq \bigcup_{l\in[1,\ell]} \left(\nu^2_{j,K^l} \cup \nu^3_{j,K^l}\right)$. We say that $\nu_{j,K^{[1,\ell]}}$ is valid if

- $x_2^{l,1} \oplus H_{k_2^l}(t^{l,1}), \ldots, x_2^{l,j_l} \oplus H_{k_2^l}(t^{l,j_l})$ for all $l \in [1,\ell]$ are pairwise distinct,
- $x_3^{l,1} \oplus H_{k_3^l}(t^{l,1}), \ldots, x_3^{l,j_l} \oplus H_{k_3^l}(t^{l,j_l})$ for all $l \in [1,\ell]$ are pairwise distinct,
- for all $l, l' \in [1,\ell]$, all distinct $r, r' \in [1,3]$, and all $i \in [1,j_l]$, all $i' \in [1,j_{l'}]$, it holds that $x_r^{l,i} \oplus H_{k_r^l}(t^{l,i}) \neq x_{r'}^{l',i'} \oplus H_{k_{r'}^{l'}}(t^{l',i'})$,

and invalid otherwise. The first and second conditions imply that $\bigcup_{l\in[1,\ell]} \nu^2_{j,K^l}$ and $\bigcup_{l\in[1,\ell]} \nu^3_{j,K^l}$ define functionally consistent input-output pairs for f, respectively. For any valid $\nu_{j,K^{[1,\ell]}}$, we let $f^{\nu_{j,K^{[1,\ell]}}}$ be the locally redefined version defined as

$$f^{\nu_{j,K^{[1,\ell]}}}(x) \coloneqq \begin{cases} y & \text{if } (x,y) \in \nu_{j,K^{[1,\ell]}}, \\ f(x) & \text{otherwise.} \end{cases}$$

We now consider the following experiment.

Experiment H_j. Sample $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell \leftarrow_\$ \tilde{\mathcal{P}}_{\mathcal{T},2n}$, $f \leftarrow_\$ \mathcal{F}_n$, $K^1 = (k_1^1, k_2^1, k_3^1), \ldots, K^\ell = (k_1^\ell, k_2^\ell, k_3^\ell) \leftarrow_\$ \{0,1\}^{3\kappa}$. Then:

1. Run $\mathcal{A}$, answering its j classical queries using $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell$ and quantum queries using f. Abort $\mathcal{A}$ immediately before submitting its $(j+1)$-st classical query.
2. Resume $\mathcal{A}$.
   (a) If $\nu_{j,K^{[1,\ell]}}$ is invalid then:
      - Answer all subsequent classical queries from $\mathcal{A}$ with the error symbol $\perp$.
      - Answer all subsequent quantum queries from $\mathcal{A}$ using f.
   (b) Else:
      - Answer all subsequent classical queries from $\mathcal{A}$ with $\mathsf{F}^f_{K^1}, \ldots, \mathsf{F}^f_{K^\ell}$.
      - Answer all subsequent quantum queries from $\mathcal{A}$ using $f^{\nu_{j,K^{[1,\ell]}}}$.

When $j = 0$, all classical queries are answered using $\mathsf{F}^f_{K^1}, \ldots, \mathsf{F}^f_{K^\ell}$ and all quantum queries are answered using f. Thus, H_0 corresponds to the experiment in the real world. On the other hand, when $j = q$, since all classical queries are answered using $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell$ and all quantum queries are answered using f, H_q corresponds to the experiment in the ideal world.

For simplicity, for all $j \in [0,q]$, we let the event that $\mathcal{A}$ in H_j outputs 1 denote $\mathtt{E}_j$. Then, we have $\mathsf{Adv}^{\mathsf{pq\text{-}tprp,mk}}_{\mathsf{F}^f}(\mathcal{A}) = |\Pr[\mathtt{E}_0] - \Pr[\mathtt{E}_q]|$. We further let the event that the transcript τ_j is invalid be $\mathtt{F}_j$. Then, using the triangle inequality, we have

$$\begin{aligned} |\Pr[\mathtt{E}_0] - \Pr[\mathtt{E}_q]| &= |\Pr[\mathtt{E}_0] - \Pr[\mathtt{E}_q \wedge \neg\mathtt{F}_q] - \Pr[\mathtt{E}_q \wedge \mathtt{F}_q]| \\ &\leq |\Pr[\mathtt{E}_0] - \Pr[\mathtt{E}_q \wedge \neg\mathtt{F}_q]| + \Pr[\mathtt{F}_q] \\ &\leq \left(\sum_{j\in[0,q-1]} |\Pr[\mathtt{E}_j \wedge \neg\mathtt{F}_j] - \Pr[\mathtt{E}_{j+1} \wedge \neg\mathtt{F}_{j+1}]| \right) + \Pr[\mathtt{F}_q]. \quad (4) \end{aligned}$$

Therefore, our aim is to bound $|\Pr[\mathtt{E}_j \wedge \neg\mathtt{F}_j] - \Pr[\mathtt{E}_{j+1} \wedge \neg\mathtt{F}_{j+1}]|$ and $\Pr[\mathtt{F}_q]$. For the latter, we have the following lemma.

Lemma 4.

$$\Pr[\mathtt{F}_q] \le 1.5q\ell\sqrt{\frac{p}{2^n}} + \frac{q^2(\ell+3)}{2^n} + \epsilon q^2 + 2\left(\frac{\epsilon q^2}{2^{\kappa+1}} + \frac{q^2}{2^{n+1}}\right)\left(1-\frac{1}{\ell}\right) + \frac{3\epsilon q^2}{2^\kappa}. \tag{5}$$

Proof. We can divide $\mathtt{F}_q$ into the following three bad events:

- bad_1: There exist $l, l' \in [1,\ell]$ and $j, j' \in [1,q]$ such that $(l,j) \neq (l',j')$ and $x_2^{l,j} \oplus H_{k_2^l}(t^{l,j}) = x_2^{l',j'} \oplus H_{k_2^{l'}}(t^{l',j'})$.
- bad_2: There exist $l, l' \in [1,\ell]$ and $j, j' \in [1,q]$ such that $(l,j) \neq (l',j')$ and $x_3^{l,j} \oplus H_{k_3^l}(t^{l,j}) = x_3^{l',j'} \oplus H_{k_3^{l'}}(t^{l',j'})$.
- bad_3: There exist $l, l' \in [1,\ell]$, $j, j' \in [1,q]$, and distinct $r, r' \in [1,3]$ such that $x_r^{l,j} \oplus H_{k_r^l}(t^{l,j}) = x_{r'}^{l',j'} \oplus H_{k_{r'}^{l'}}(t^{l',j'})$.

Clearly, we can write $\mathtt{F}_q = \mathsf{bad}_1 \cup \mathsf{bad}_2 \cup \mathsf{bad}_3$.

The event bad_2 is the occurrence of a collision on the ciphertext. We can denote $\mathsf{bad}_2 = \mathsf{bad}_{21} \cup \mathsf{bad}_{22}$, where:

- bad_{21}: There exists $l \in [1,\ell]$ and distinct $j, j' \in [1, q_l]$ such that $x_3^{l,j} \oplus H_{k_3^l}(t^{l,j}) = x_3^{l,j'} \oplus H_{k_3^l}(t^{l,j'})$.
- bad_{22}: There exists distinct $l, l' \in [1,\ell]$, $j \in [1,q_l]$ and $j' \in [1,q_{l'}]$ such that $x_3^{l,j} \oplus H_{k_3^l}(t^{l,j}) = x_3^{l',j'} \oplus H_{k_3^{l'}}(t^{l',j'})$.

bad_{21} implies the event of collisions on x_3 (a part of the ciphertext) for each instance. From the AXU property, the probability that $x_3^{l,j} \oplus H_{k_3^l}(t^{l,j}) = x_3^{l,j'} \oplus H_{k_3^l}(t^{l,j'})$ for any $j, j' \in [1, q_l]$ holds under the condition that $t^{l,j} \neq t^{l,j'}$ can be bounded by ϵ. On the other hand, it can be bounded by 2^{-n} under the condition that $t^{l,j} = t^{l,j'}$. Therefore, we have

$$\Pr[\mathsf{bad}_{21}] \le \sum_{l\in[1,\ell]} \frac{q_l(q_l-1)}{2}(\epsilon + 2^{-n}) \le \frac{\epsilon q^2}{2} + \frac{q^2}{2^n}.$$

On the other hand, bad_{22} implies the cross-collisions event across all instances. This event can be further divided into two events and it can be denoted as $\mathsf{bad}_{22} = \mathsf{bad}_{221} \cup \mathsf{bad}_{222}$, where:

- bad_{221}: bad_{22} under the condition that $k_3^l = k_3^{l'}$ for some distinct $l, l' \in [1,\ell]$. (two independently chosen keys collide)
- bad_{222}: bad_{22} under the condition that $k_3^l \neq k_3^{l'}$ for some distinct $l, l' \in [1,\ell]$.

Let $S_{l,j} = x_3^{l,j} \oplus H_{k_3^l}(t^{l,j})$. Then, we can bound the probability of bad_{221} as

$$\begin{aligned}\Pr[\mathsf{bad}_{221}] &\le \sum_{l,l'\in[1,\ell]|l<l'} \sum_{j\le q_l} \sum_{j'\le q_{l'}} \Pr\left[k_3^l = k_3^{l'}\right] \cdot \Pr\left[S_{l,j} = S_{l',j'} \mid k_3^l = k_3^{l'}\right] \\ &\le \frac{1}{2^\kappa}(\epsilon + 2^{-n}) \sum_{l,l'\in[1,\ell]|l<l'} q_l q_{l'}. \end{aligned} \tag{6}$$

Since $q = \sum_{l\in[1,\ell]} q_l$, using the Cauchy-Schwarz inequality, we have

$$\sum_{l,l'\in[1,\ell]|l<l'} q_l q_{l'} = \frac{q^2 - \sum_{l\in[1,\ell]} q_l^2}{2} \leq \frac{q^2}{2}\left(1 - \frac{1}{\ell}\right). \tag{7}$$

Therefore, from Equations (6) and (7), we have

$$\Pr[\mathsf{bad}_{221}] \leq \left(\frac{\epsilon q^2}{2^{\kappa+1}} + \frac{q^2}{2^{\kappa+n+1}}\right)\left(1 - \frac{1}{\ell}\right).$$

On the other hand, assume $k_3^l \neq k_3^{l'}$, $H_{k_3^l}(t^{l,j})$ and $H_{k_3^{l'}}(t^{l',j'})$ are uniform, and then we have

$$\Pr[\mathsf{bad}_{222}] \leq \frac{\sum_{l,l'\in[1,\ell]|l<l'} q_l q_{l'}}{2^n} \leq \frac{q^2}{2^{n+1}}\left(1 - \frac{1}{\ell}\right).$$

Therefore, we have

$$\begin{aligned}\Pr[\mathsf{bad}_2] = \Pr[\mathsf{bad}_{21}] + \Pr[\mathsf{bad}_{22}] &= \Pr[\mathsf{bad}_{21}] + \Pr[\mathsf{bad}_{221}] + \Pr[\mathsf{bad}_{222}] \\ &\leq \frac{\epsilon q^2}{2} + \left(\frac{\epsilon q^2}{2^{\kappa+1}} + \frac{q^2}{2^{n+1}}\right)\left(1 - \frac{1}{\ell}\right).\end{aligned}$$

We bound bad_3 in the same manner as bad_2. The event bad_3 captures collisions between different rounds. There are three relevant pairs of rounds to consider, namely $(1,2)$, $(1,3)$, and $(2,3)$. In contrast to bad_2, bad_3 must also account for collisions occurring within the same instance. Therefore, we have

$$\Pr[\mathsf{bad}_3] \leq 3\left(\frac{\epsilon}{2^\kappa} + \frac{1}{2^n}\right)\sum_{l,l'\in[1,\ell]} q_l q_{l'} \leq \frac{3\epsilon q^2}{2^\kappa} + \frac{3q^2}{2^n}.$$

The event bad_1 is the occurrence of a collision on the internal value. As bad_2, we can denote $\mathsf{bad}_1 = \mathsf{bad}_{11} \cup \mathsf{bad}_{12}$, where:

- bad_{11}: There exists $l \in [1,\ell]$ and distinct $j, j' \in [1, q_l]$ such that $x_2^{l,j} \oplus H_{k_2^l}(t^{l,j}) = x_2^{l,j'} \oplus H_{k_2^l}(t^{l,j'})$.
- bad_{12}: There exists distinct $l, l' \in [1,\ell]$, $j \in [1, q_l]$ and $j' \in [1, q_{l'}]$ such that $x_2^{l,j} \oplus H_{k_2^l}(t^{l,j}) = x_2^{l',j'} \oplus H_{k_2^{l'}}(t^{l',j'})$.

bad_{11} implies a collision on x_2 if $t^{l,j} = t^{l,j'}$ for some $l \in [1,\ell]$. Therefore, this can be bound by using Lemma 3. For adversary inputs $(x_0^{l,j}, x_1^{l,j})$ and $(x_0^{l,j'}, x_1^{l,j'})$, the relation $\mathbb{R}$ is given as

$$\mathbb{R} : x_2^{l,j} = f(x_1^{l,j} \oplus H_{k_1^l}(t^{l,j})) \oplus x_0^{l,j} = f(x_1^{l,j'} \oplus H_{k_1^l}(t^{l,j'})) \oplus x_0^{l,j'} = x_2^{l,j'},$$

for $l \in [1,\ell]$. Since $H_{k_1^l}(\cdot)$ is uniform for $t^{l,j} = t^{l,j'}$ and the output can be used as the uniform key chosen from $\{0,1\}^n$, we can apply Lemma 3. By adding $\epsilon q^2/2$

which is an upper bound on the probability of bad_{11} under the condition that $t^{l,j} \neq t^{l,j'}$, we have

$$\Pr[\mathsf{bad}_{11}] \leq 1.5q_l\ell\sqrt{\frac{p}{2^n}} + q_l^2\ell \max_{x,x' \in \mathcal{X}} \Pr_{\substack{r \leftarrow_{\$} \mathcal{Y} \\ s \leftarrow_{\$} \mathcal{X}_1}} [(x, x', r, f(s)) \in \mathbb{R}] + \frac{\epsilon q^2}{2}$$
$$\leq 1.5q\ell\sqrt{\frac{p}{2^n}} + \frac{q^2\ell}{2^n} + \frac{\epsilon q^2}{2}.$$

On the other hand, the probability of bad_{12} is identical to the probability $\Pr[\mathsf{bad}_{22}]$, due to AXU property of H. Thus, we have

$$\Pr[\mathsf{bad}_{12}] = \Pr[\mathsf{bad}_{22}] \leq \left(\frac{\epsilon q^2}{2^{\kappa+1}} + \frac{q^2}{2^{n+1}}\right)\left(1 - \frac{1}{\ell}\right).$$

Therefore,

$$\Pr[\mathsf{bad}_1] \leq 1.5q\ell\sqrt{\frac{p}{2^n}} + \frac{q^2\ell}{2^n} + \left(\frac{\epsilon q^2}{2^{\kappa+1}} + \frac{q^2}{2^{n+1}}\right)\left(1 - \frac{1}{\ell}\right) + \frac{\epsilon q^2}{2}.$$

□

The remaining part of the proof, *i.e.*, bounding $|\Pr[\mathtt{E}_j \wedge \neg \mathtt{F}_j] - \Pr[\mathtt{E}_{j+1} \wedge \neg \mathtt{F}_{j+1}]|$ for all $j \in [0, q-1]$, relies on the technique proposed in [2]. Now we compare H_j and H_{j+1} under the assumption that the reprogramming set is valid. To make the comparison easy, we introduce the intermediate experiment H'_j, defined as follows.

Experiment H'_j. Sample $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell \leftarrow_{\$} \tilde{\mathcal{P}}_{\mathcal{T},2n}$, $f \leftarrow_{\$} \mathcal{F}_n$, $K^1 = (k_1^1, k_2^1, k_3^1), \ldots, K^\ell = (k_1^\ell, k_2^\ell, k_3^\ell) \leftarrow_{\$} \{0,1\}^{3\kappa}$.

1. Run $\mathcal{A}$, answering its j classical queries using $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell$ and quantum queries using f. Abort $\mathcal{A}$ immediately after submitting the $(j+1)$-st classical query.
2. Resume $\mathcal{A}$.
   (a) If $\nu_{j+1,K^{[1,\ell]}}$ is invalid then:
      – Answer all subsequent classical queries (including the $(j+1)$-st classical query) from $\mathcal{A}$ with the error symbol $\perp$.
      – Answer all subsequent quantum queries from $\mathcal{A}$ using f.
   (b) Else:
      – Answer all subsequent classical queries (including the $(j+1)$-st classical query) from $\mathcal{A}$ with $\mathsf{F}^f_{K^1}, \ldots, \mathsf{F}^f_{K^\ell}$.
      – Answer all subsequent quantum queries from $\mathcal{A}$ using $f^{\nu_{j+1,K^{[1,\ell]}}}$.

Let $\mathtt{E}'_j$ be the event that $\mathcal{A}$ outputs 1 in H'_j, and let $\mathtt{F}'_j$ be the event that the reprogramming set is invalid in the experiment H'_j. Then, we have

$$|\Pr[\mathtt{E}_j \wedge \neg \mathtt{F}_j] - \Pr[\mathtt{E}_{j+1} \wedge \neg \mathtt{F}_{j+1}]| \leq |\Pr[\mathtt{E}_j \wedge \neg \mathtt{F}_j] - \Pr[\mathtt{E}'_j \wedge \neg \mathtt{F}'_j]|$$
$$+ |\Pr[\mathtt{E}'_j \wedge \neg \mathtt{F}'_j] - \Pr[\mathtt{E}_{j+1} \wedge \neg \mathtt{F}_{j+1}]|.$$

Therefore, we can bound $|\Pr[\mathtt{E}_j \wedge \neg \mathtt{F}_j] - \Pr[\mathtt{E}_{j+1} \wedge \neg \mathtt{F}_{j+1}]|$ by comparing H_j and H'_j, and H'_j and H_{j+1} under the assumption that the reprogramming set is valid.

We first compare H'_j and H_{j+1}. For these experiments, we have the following.

Lemma 5. *Let p_{j+1} be the number of quantum queries $\mathcal{A}$ made between the $(j+1)$-st and $(j+2)$nd classical queries. Then, for all $j \in [0, q-1]$,*

$$|\Pr\left[\mathtt{E}'_j \wedge \neg \mathtt{F}'_j\right] - \Pr[\mathtt{E}_{j+1} \wedge \neg \mathtt{F}_{j+1}]| \leq 2p_{j+1}\sqrt{\frac{2(j+1)}{2^n}}.$$

Proof. Under the assumption that the reprogramming sets are valid, experiments H'_j and H_{j+1} can be concisely described by the provided oracle sequences, as follows.

$$\mathsf{H}'_j : \underbrace{f, \tilde{\Pi}_{\mathsf{p}(1)}, f, \ldots, \tilde{\Pi}_{\mathsf{p}(j)}, f}_{j \text{ classical queries}}, \tilde{\Pi}_{\mathsf{p}(j+1)}, f^{\nu_{j+1,K^{[1,\ell]}}}, \underbrace{\mathsf{F}^f_{K^{\mathsf{p}(j+2)}}, f^{\nu_{j+1,K^{[1,\ell]}}}, \ldots, \mathsf{F}^f_{K^{\mathsf{p}(q)}}, f^{\nu_{j+1,K^{[1,\ell]}}}}_{q-j-1 \text{ classical queries}}$$

$$\mathsf{H}_{j+1} : \underbrace{f, \tilde{\Pi}_{\mathsf{p}(1)}, f, \ldots, \tilde{\Pi}_{\mathsf{p}(j)}, f}_{j \text{ classical queries}}, \tilde{\Pi}_{\mathsf{p}(j+1)}, f, \quad \underbrace{\mathsf{F}^f_{K^{\mathsf{p}(j+2)}}, f^{\nu_{j+1,K^{[1,\ell]}}}, \ldots, \mathsf{F}^f_{K^{\mathsf{p}(q)}}, f^{\nu_{j+1,K^{[1,\ell]}}}}_{q-j-1 \text{ classical queries}}$$

The difference between these two experiments reduces to the difference between the primitive, which is quantum accessible at the $(j+1)$-st stage. Since $\mathcal{A}$ can make quantum queries to f or $f^{\nu_{j+1,K^{[1,\ell]}}}$, the locally redefined version of f based on $(j+1)$ classical queries, we construct the ARP distinguisher from $\mathcal{A}$, which proceeds as follows.

Phase 1: $\mathcal{D}$ samples $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell \leftarrow_{\$} \tilde{\mathcal{P}}_{\mathcal{T},2n}$, $f \leftarrow_{\$} \mathcal{F}_n$. It then runs $\mathcal{A}$, answering its classical queries using $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell$ and quantum queries using f. Abort immediately after receiving a response of $\mathcal{A}$'s $(j+1)$-st classical query. Then, $\mathcal{D}$ defines the randomized algorithm $\mathcal{B}$ which proceeds as follows.
  (a) Samples $K^1 = (k^1_1, k^1_2, k^1_3), \ldots, K^\ell = (k^\ell_1, k^\ell_2, k^\ell_3) \leftarrow_{\$} \{0,1\}^{3\kappa}$, and compute the set $\nu_{j+1,K^{[1,\ell]}}$.
  (b) Return the set $\nu_{j+1,K^{[1,\ell]}}$.

Phase 2: $\mathcal{B}$ is run to generate $\nu_{j+1,K^{[1,\ell]}}$. A uniform bit is chosen $b \in \{0,1\}$ and $\mathcal{A}$ is given quantum access to f_b, where $f_0 = f$ and $f_1 = f^{\nu_{j+1,K^{[1,\ell]}}}$.

Phase 3: $\mathcal{D}$ is given $K^1, \ldots, K^\ell$ (or randomness to generate the keys), and resumes $\mathcal{A}$.
  – $\mathcal{D}$ answers all remaining classical queries from $\mathcal{A}$ using $\mathsf{F}^f_{K^1}, \ldots, \mathsf{F}^f_{K^\ell}$.
  – $\mathcal{D}$ answers all remaining quantum queries from $\mathcal{A}$ using $f^{\nu_{j+1,K^{[1,\ell]}}}$.

  Finally, $\mathcal{D}$ outputs whatever $\mathcal{A}$ outputs.

If $b = 0$, $\mathcal{A}$ in $\mathcal{D}$ corresponds to $\mathcal{A}$ in the experiment H_{j+1}. On the other hand, if $b = 1$, $\mathcal{A}$ in $\mathcal{D}$ corresponds to $\mathcal{A}$ in the experiment H'_j. Thus, we have

$$\begin{aligned}|\Pr\left[\mathtt{E}'_j \wedge \neg \mathtt{F}'_j\right] - \Pr[\mathtt{E}_{j+1} \wedge \neg \mathtt{F}_{j+1}]| &= \left|\Pr\left[\mathcal{D}^{f_b} = 1 \mid b = 1\right] - \Pr\left[\mathcal{D}^{f_b} = 1 \mid b = 0\right]\right| \\ &= \mathsf{Adv}^{\mathsf{arp}}(\mathcal{D}) \leq 2p_{j+1}\sqrt{\epsilon}.\end{aligned}$$

Since $|\nu_{j+1,K^{[1,\ell]}}| = 2(j+1)$, ϵ can be bounded as $\epsilon \leq 2(j+1)/2^n$. □

Next we compare H_j and H'_j. For these experiments, we have the following.

Lemma 6. *For all $j \in [0, q]$, we have*

$$\left|\Pr[\mathtt{E}_j \wedge \neg \mathtt{F}_j] - \Pr\left[\mathtt{E}'_j \wedge \neg \mathtt{F}'_j\right]\right| \leq 2\sqrt{\frac{p}{2^{n+1}}} + \frac{p}{2^{n+2}} + \frac{q}{2^{2n}} + \frac{6j}{2^n}.$$

Proof. We further introduce intermediate experiments H^1_j and H^2_j to make the comparison between H'_j and H_j easy. Let $t' = t^{\mathsf{p}(j+1), j_{\mathsf{p}(j+1)}+1}$, $x'_r = x_r^{\mathsf{p}(j+1), j_{\mathsf{p}(j+1)}+1}$ and $k'_r = k_r^{\mathsf{p}(j+1)}$ for all $r \in [0, 3]$. H^1_j proceeds as follows:

Experiment H^1_j. Sample $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell \leftarrow_\$ \tilde{\mathcal{P}}_{\mathcal{T},2n}$, $f \leftarrow_\$ \mathcal{F}_n$, $K^1 = (k^1_1, k^1_2, k^1_3), \ldots, K^\ell = (k^\ell_1, k^\ell_2, k^\ell_3) \leftarrow_\$ \{0,1\}^{3\kappa}$.

1. Run $\mathcal{A}$, answering its j classical queries using $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell$ and quantum queries using f. Abort $\mathcal{A}$ immediately after submitting its $(j+1)$-st classical query without receiving the response.
2. Let $x'_2 = x'_0 \oplus f(x'_1 \oplus H_{k'_1}(t'))$. The challenger interacts with Φ as follows.
   (a) The challenger samples $\hat{k}_2 \leftarrow_\$ \{0,1\}^\kappa$ and defines the distribution D_1, where $D_1(s) = \Pr\left[x'_2 \oplus H_{\hat{k}_2}(t') = s\right]$ for all $s \in \{0,1\}^n$. The challenger runs the resampling operation (queries Φ) with the input distribution D_1 and it outputs s_2. It sets $x'_3 = x'_1 \oplus f(s_2)$.
   (b) The challenger samples $\hat{k}_3 \leftarrow_\$ \{0,1\}^\kappa$ and defines the distribution D_2, where $D_2(s) = \Pr\left[x'_3 \oplus H_{\hat{k}_3}(t') = s\right]$ for all $s \in \{0,1\}^n$. The challenger runs the resampling operation (queries Φ) with the input distribution D_2 and it outputs s_3. It sets $x'_4 = x'_2 \oplus f(s_3)$.

   Let $\nu^{(1)}_{j,K^{[1,\ell]}} = \nu_{j,K^{[1,\ell]}} \cup \nu^{(1)}$ where $\nu^{(1)} = \{s_2, s_3\}$.
3. Resume $\mathcal{A}$.
   (a) If $\nu_{j,K^{[1,\ell]}}$ is invalid then:
      – Answer all subsequent classical queries (including the $(j+1)$-st classical query) from $\mathcal{A}$ with the error symbol $\perp$.
      – Answer all subsequent quantum queries from $\mathcal{A}$ using f.
   (b) Else
      – Answer all subsequent classical queries (including the $(j+1)$-st classical query) from $\mathcal{A}$ using $\mathsf{F}_{K^1}^{f^{\nu^{(1)}_{j,K^{[1,\ell]}}}}, \ldots, \mathsf{F}_{K^\ell}^{f^{\nu^{(1)}_{j,K^{[1,\ell]}}}}$.
      – Answer all subsequent quantum queries from $\mathcal{A}$ using $f^{\nu^{(1)}_{j,K^{[1,\ell]}}}$.

Experiment H^2_j. H^2_j is an identical experiment to H^1_j except that the $(j+1)$-st classical query is answered by $\tilde{\Pi}_{\mathsf{p}(j+1)}$.

Under the assumption that the reprogramming sets are valid, these experiments can be concisely described by the provided oracle sequences, as follows.

$$\mathsf{H}_j : \underbrace{f, \tilde{\Pi}_{\mathsf{p}(1)}, f, \ldots, \tilde{\Pi}_{\mathsf{p}(j)}, f,}_{j \text{ classical queries}} \mathsf{F}^f_{K^{\mathsf{p}(j+1)}}, f^{\nu_{j,K^{[1,\ell]}}}, \underbrace{\mathsf{F}^f_{K^{\mathsf{p}(j+2)}}, f^{\nu_{j,K^{[1,\ell]}}}, \ldots, \mathsf{F}^f_{K^{\mathsf{p}(q)}}, f^{\nu_{j,K^{[1,\ell]}}}}_{q-j-1 \text{ classical queries}}$$

$$\mathsf{H}^1_j : \underbrace{f, \tilde{\Pi}_{\mathsf{p}(1)}, f, \ldots, \tilde{\Pi}_{\mathsf{p}(j)}, f,}_{j \text{ classical queries}} \mathsf{F}^f_{K^{\mathsf{p}(j+1)}}, f^{\nu^{(1)}_{j,K^{[1,\ell]}}}, \underbrace{\mathsf{F}^f_{K^{\mathsf{p}(j+2)}}, f^{\nu^{(1)}_{j,K^{[1,\ell]}}}, \ldots, \mathsf{F}^f_{K^{\mathsf{p}(q)}}, f^{\nu^{(1)}_{j,K^{[1,\ell]}}}}_{q-j-1 \text{ classical queries}}$$

$$\mathsf{H}^2_j : \underbrace{f, \tilde{\Pi}_{\mathsf{p}(1)}, f, \ldots, \tilde{\Pi}_{\mathsf{p}(j)}, f,}_{j \text{ classical queries}} \tilde{\Pi}_{\mathsf{p}(j+1)}, \ f^{\nu^{(1)}_{j,K^{[1,\ell]}}}, \underbrace{\mathsf{F}^f_{K^{\mathsf{p}(j+2)}}, f^{\nu^{(1)}_{j,K^{[1,\ell]}}}, \ldots, \mathsf{F}^f_{K^{\mathsf{p}(q)}}, f^{\nu^{(1)}_{j,K^{[1,\ell]}}}}_{q-j-1 \text{ classical queries}}$$

$$\mathsf{H}'_j : \underbrace{f, \tilde{\Pi}_{\mathsf{p}(1)}, f, \ldots, \tilde{\Pi}_{\mathsf{p}(j)}, f,}_{j \text{ classical queries}} \tilde{\Pi}_{\mathsf{p}(j+1)}, \ f^{\nu_{j+1,K^{[1,\ell]}}}, \underbrace{\mathsf{F}^f_{K^{\mathsf{p}(j+2)}}, f^{\nu_{j+1,K^{[1,\ell]}}}, \ldots, \mathsf{F}^f_{K^{\mathsf{p}(q)}}, f^{\nu_{j+1,K^{[1,\ell]}}}}_{q-j-1 \text{ classical queries}}$$

First, we compare H_j with H^1_j. Under the condition that the reprogramming set is valid, we can construct the following ARS distinguisher from $\mathcal{A}$ between H_j and H^1_j, which proceeds as follows.

Phase 1: $\mathcal{D}$ samples uniform $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell \leftarrow_\$ \tilde{\mathcal{P}}_{\mathcal{T},2n}$, $f \leftarrow_\$ \mathcal{F}_n$, $K^1 = (k^1_1, k^1_2, k^1_3), \ldots, K^\ell = (k^\ell_1, k^\ell_2, k^\ell_3) \leftarrow_\$ \{0,1\}^{3\kappa}$ It then runs $\mathcal{A}$, answering its classical queries using $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell$, and quantum queries using f. Abort immediately after submitting $\mathcal{A}$'s $(j+1)$-st classical query, without receiving the response.

Phase 2: $\mathcal{D}$ sets $x'_2 = x'_0 \oplus f(x'_1 \oplus H_{k'_1}(t'))$. $\mathcal{D}$ runs the following procedure:

(a) $\mathcal{D}$ samples $\hat{k}_2 \leftarrow_\$ \{0,1\}^\kappa$ and defines the distribution D_1 where $D_1(s) = \Pr\left[x'_2 \oplus H_{\hat{k}_2}(t') = s\right]$ for all $s \in \{0,1\}^n$. Then, it queries Φ with D_1 and outputs s_2. It sets $x'_3 = f(s_2) \oplus x'_1$.

(b) $\mathcal{D}$ samples $\hat{k}_3 \leftarrow_\$ \{0,1\}^\kappa$ and defines the distribution D_2 where $D_2(s) = \Pr\left[x'_3 \oplus H_{k'_3}(t') = s\right]$ for all $s \in \{0,1\}^n$. Then, it queries Φ with D_2 and it outputs s_3. It sets $x'_4 = f(s_3) \oplus x'_2$.

(c) $\mathcal{D}$ computes $\nu^{(1)}_{j,K^{[1,\ell]}} = \nu_{j,K^{[1,\ell]}} \cup \nu^{(1)}$, where $\nu^{(1)} = \{s_2, s_3\}$.

Phase 3: A uniform bit $b \in \{0,1\}$ is chosen, and $\mathcal{A}$ is given quantum queries to f_b where $f_0 = f^{\nu_{j,K^{[1,\ell]}}}$ and $f_1 = f^{\nu^{(1)}_{j,K^{[1,\ell]}}}$. Under the assumption $\nu_{j,K^{[1,\ell]}}$ is valid:

- $\mathcal{D}$ answers all remaining classical queries from $\mathcal{A}$ using $\mathsf{F}^{f_b}_{K^1}, \ldots, \mathsf{F}^{f_b}_{K^\ell}$.
- $\mathcal{D}$ answers all remaining quantum queries from $\mathcal{A}$ using f_b.

Finally, $\mathcal{D}$ outputs whatever $\mathcal{A}$ outputs.

Let the event that $\mathcal{A}$ in H^1_j outputs 1 be $\mathtt{E}^1_j$ and the event that the reprogramming set in H^1_j is invalid be $\mathtt{F}^1_j$. Considering the above ARS distinguisher, we have

$$\begin{aligned}
&\left|\Pr[\mathtt{E}_j \wedge \neg \mathtt{F}_j] - \Pr\left[\mathtt{E}^1_j \wedge \neg \mathtt{F}^1_j\right]\right| = \left|\Pr\left[\mathcal{D}^{f_b,\Phi} = 1 \mid b = 1\right] - \Pr\left[\mathcal{D}^{f_b,\Phi} = 1 \mid b = 0\right]\right| \\
&= \mathsf{Adv}^{\mathsf{ars}}(\mathcal{D}) \leq \sum_{r \in \{1,2\}} \left(\sqrt{\hat{p}_r}\mu_r + \frac{1}{2}\hat{p}_r\mu_r\right) \leq 2\sqrt{\frac{p}{2^{n+1}}} + \frac{p}{2^{n+2}}, \qquad (8)
\end{aligned}$$

where $\hat{p}_r = \sum_{i \in [0,r-1]} p_i$ and p_i is the number of queries to f_b that are issued between the $(i-1)$-th and the i-th query to Φ (note that $p_0 = 0$ and $p = \sum_{i \in \{1,2\}} p_i$).

Next, we compare H_j^1 and H_j^2. Let $\mathtt{E}_j^2$ be the event that $\mathcal{A}$ in H_j^2 outputs 1, and $\mathtt{F}_j^2$ be the event that $\nu_{j,K^{[1,\ell]}}$ is invalid. In H_j^1, the $(j+1)$-st classical query is answered by the oracle indexed by $\mathsf{p}(j+1)$ as $x'_3 = f(s_2) \oplus x'_1$, $x'_4 = f(s_3) \oplus x'_2$. On the other hand, in H_j^2, (x'_3, x'_4) is uniformly distributed over $\{0,1\}^{2n}$ given the output of $\tilde{\Pi}_{\mathsf{p}(j+1)}$. Thus, since these experiments are $q_{\mathsf{p}(j+1)}2^{-2n}$-close in statistical distance, we have

$$\left|\Pr\left[\mathtt{E}_j^1 \wedge \neg \mathtt{F}_j^1\right] - \Pr\left[\mathtt{E}_j^2 \wedge \neg \mathtt{F}_j^2\right]\right| \leq \frac{q_{\mathsf{p}(j+1)}}{2^{2n}} \leq \frac{q}{2^{2n}}. \tag{9}$$

Finally, we compare H_j^2 and H'_j assuming that the reprogramming sets are valid in both experiments. In H_j^2, $f^{\nu^{(1)}_{j,K^{[1,\ell]}}}$ is used as the quantum accessible primitive. On the other hand, in H'_j, $f^{\nu_{j+1,K^{[1,\ell]}}}$ is used as the quantum accessible primitive. The bad events in this case are that s_2 or s_3 collide with any value in the reprogramming set generated from classical queries to $\tilde{\Pi}_{\mathsf{p}(j+1)}$ up to the $j_{\mathsf{p}(j+1)}$-th classical query. More precisely, the bad events are as follows.

- bad_4: there exist $i \in [1, j_{\mathsf{p}(j+1)}]$, $r \in [1,3]$ such that $s_2 = x_r^{\mathsf{p}(j+1),i} \oplus H_{k'_r}(t^{\mathsf{p}(j+1),i})$.
- bad_5: there exist $i \in [1, j_{\mathsf{p}(j+1)}]$, $r \in [1,3]$ such that $s_3 = x_r^{\mathsf{p}(j+1),i} \oplus H_{k'_r}(t^{\mathsf{p}(j+1),i})$.

Since s_2 follows the distribution D_1, we have $\Pr[\mathsf{bad}_4] \leq 3j_{\mathsf{p}(j+1)} \max_{x\in\{0,1\}^n} D_1(x) \leq 3j/2^n$. The probabilities of bad_5 can be bounded in the same manner as for bad_4. The only difference is that s_3 is sampled from D_2 rather than from D_1. Therefore, we have $\Pr[\mathsf{bad}_5] \leq 3j_{\mathsf{p}(j+1)} \max_{x\in\{0,1\}^n} D_2(x) \leq 3j/2^n$. Thus,

$$\left|\Pr\left[\mathtt{E}_j^2 \wedge \neg \mathtt{F}_j^2\right] - \Pr\left[\mathtt{E}'_j \wedge \neg \mathtt{F}'_j\right]\right| \leq \Pr[\mathsf{bad}_4] + \Pr[\mathsf{bad}_5] \leq \frac{6j}{2^n}. \tag{10}$$

From Equations (8) to (10), we have Lemma 6. □

From Lemmas 5 and 6, we have

$$\begin{aligned}|\Pr[\mathtt{E}_j \wedge \neg \mathtt{F}_j] - \Pr[\mathtt{E}_{j+1} \wedge \neg \mathtt{F}_{j+1}]| \leq 2p_{j+1}\sqrt{\frac{2(j+1)}{2^n}} + 2\sqrt{\frac{p}{2^{n+1}}} \\ + \frac{p}{2^{n+2}} + \frac{q}{2^{2n}} + \frac{6j}{2^n}.\end{aligned}$$

Further, using Lemma 4,

$$\begin{aligned}&\mathsf{Adv}^{\mathsf{pq\text{-}tprp},\mathsf{mk}}_{\mathsf{F}^f}(\mathcal{A}) = |\Pr[\mathtt{E}_0] - \Pr[\mathtt{E}_q]| \\ &\leq \sum_{j\in[0,q-1]} \left(2p_{j+1}\sqrt{\frac{2(j+1)}{2^n}} + 2\sqrt{\frac{p}{2^{n+1}}} + \frac{p}{2^{n+2}} + \frac{q}{2^{2n}} + \frac{6j}{2^n}\right) \\ &+ 1.5q\ell\sqrt{\frac{p}{2^n}} + \frac{q^2(\ell+3)}{2^n} + \epsilon q^2 + 2\left(\frac{\epsilon q^2}{2^{\kappa+1}} + \frac{q^2}{2^{n+1}}\right)\left(1 - \frac{1}{\ell}\right) + \frac{3\epsilon q^2}{2^\kappa} \\ &\leq 2p\sqrt{\frac{2q}{2^n}} + 4\ell q\sqrt{\frac{p}{2^n}} + \frac{9q^2\ell}{2^n} + 3\epsilon q^2\end{aligned}$$

□

4 Post-Quantum Security of 4-round TKAF

Theorem 2 (PQ-STPRP security of 4-TKAF). *Let $\mathcal{A}'$ be an adversary who can make q' classical queries to $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell \leftarrow_\$ \tilde{\mathcal{P}}_{\mathcal{T},2n}$ or 4-TKAFs instantiated as $\mathsf{F}^f_{K^1}, \ldots, \mathsf{F}^f_{K^\ell}$ where $K^1 = (k^1_1, k^1_2, k^1_3, k^1_4), \ldots, K^\ell = (k^\ell_1, k^\ell_2, k^\ell_3, k^\ell_4) \leftarrow_\$ \{0,1\}^{4\kappa}$ and $f \leftarrow_\$ \mathcal{F}_n$, and p' quantum queries to f. Then,*

$$\mathsf{Adv}^{\mathsf{pq\text{-}stprp,mk}}_{\mathsf{F}^f}(\mathcal{A}') \leq 4p'\sqrt{\frac{4q'}{2^n}} + 10\ell q'\sqrt{\frac{2p'}{2^n}} + \frac{56{q'}^2\ell}{2^n} + 20\epsilon{q'}^2.$$

Proof. The proof is basically the same as the proof of Theorem 1. As in the proof of Theorem 1, we consider the corresponding static adversary $\mathcal{A}$ making at most $q = 2q'$ classical queries and $p = 2p'$ quantum queries. The oracle access given to $\mathcal{A}$ is the same as in the proof of Theorem 1, except that the classical oracles in both the real and ideal worlds provide forward and inverse access to 4-round TKAFs and tweakable random permutations, respectively.

For all $l \in [1, \ell]$, the classical transcript obtained by j_l classical queries to the oracle of the ideal construction indexed by l so far is

$$\tau^l_{j_l} \coloneqq \left\{(\delta^{l,1}, t^{l,1}, (x^{l,1}_0, x^{l,1}_1), (x^{l,1}_4, x^{l,1}_5)), \ldots, (\delta^{l,j_l}, t^{l,j_l}(x^{l,j_l}_0, x^{l,j_l}_1), (x^{l,j_l}_4, x^{l,j_l}_5))\right\}.$$

For all $i \in [1, j_l]$, $\delta^{l,i} \in \{-1, 1\}$ denotes the direction of the i-th classical query to the l-th oracle. $\delta^{l,i} = 1$ indicates a forward query, and $\delta^{l,i} = -1$ indicates an inverse query. Besides, $t^{l,i}$ is the i-th tweak input, $(x^{l,i}_0, x^{l,i}_1)$ is the i-th plaintext, and $(x^{l,i}_4, x^{l,i}_5)$ is the i-th ciphertext. We assume that the adversary never issues duplicate classical queries to any fixed oracle. As per the proof of Theorem 1, we denote a union of transcripts of all oracles as a multiset $\tau_j \coloneqq \bigcup_{l\in[1,\ell]} \tau^l_{j_l}$, where $j = \sum_{l\in[1,\ell]} j_l$. We reuse the set Equations (2) and (3). Then, the reprogramming set $\nu_{j,K^{[1,\ell]}} \coloneqq \bigcup_{l\in[1,\ell]} \left(\nu^2_{j,K^l} \cup \nu^3_{j,K^l}\right)$ is valid if

- $x^{l,1}_2 \oplus H_{k^l_2}(t^{l,1}), \ldots, x^{l,j_l}_2 \oplus H_{k^l_2}(t^{l,j_l})$ for all $l \in [1, \ell]$ are pairwise distinct,
- $x^{l,1}_3 \oplus H_{k^l_3}(t^{l,1}), \ldots, x^{l,j_l}_3 \oplus H_{k^l_3}(t^{l,j_l})$ for all $l \in [1, \ell]$ are pairwise distinct,
- For all $l, l' \in [1, \ell]$, all distinct $r, r' \in [1, 4]$, and all $i \in [1, j_l]$, $i' \in [1, j_{l'}]$, it holds that $x^{l,i}_r \oplus H_{k^l_r}(t^{l,i}) \neq x^{l',i'}_{r'} \oplus H_{k^{l'}_{r'}}(t^{l',i'})$.

Otherwise, $\nu_{j,K^{[1,\ell]}}$ is invalid. As per Section 3, we consider the experiment H_j and H'_j where the 4-TKAF is used as F instead of 3-TKAF, and $K^l = (k^l_1, k^l_2, k^l_3, k^l_4) \in \{0,1\}^{4\kappa}$ is sampled instead of $K^l = (k^l_1, k^l_2, k^l_3) \in \{0,1\}^{3\kappa}$ for all $l \in [1, \ell]$. We prove Theorem 2 by bounding Equation (4) in this setting.

Lemma 7.

$$\Pr[\mathsf{F}_q] \leq 3q\ell\sqrt{\frac{p}{2^n}} + \frac{q^2(\ell+3)}{2^{n-1}} + \left(\frac{\epsilon q^2}{2^\kappa} + \frac{q^2}{2^n}\right)\left(1 - \frac{1}{\ell}\right) + \frac{6\epsilon q^2}{2^\kappa} + \epsilon q^2$$

Proof. F_q can be divided into the following three bad events:

- bad_1: There exist $l, l' \in [1, \ell]$ and $j, j' \in [1, q]$ such that $(l, j) \neq (l', j')$ and $x_2^{l,j} \oplus H_{k_2^l}(t^{l,j}) = x_2^{l',j'} \oplus H_{k_2^{l'}}(t^{l',j'})$.
- bad_2: There exist $l, l' \in [1, \ell]$ and $j, j' \in [1, q]$ such that $(l, j) \neq (l', j')$ and $x_3^{l,j} \oplus H_{k_3^l}(t^{l,j}) = x_3^{l',j'} \oplus H_{k_3^{l'}}(t^{l',j'})$.
- bad_3: There exist $l, l' \in [1, \ell]$, $j, j' \in [1, q]$, and distinct $r, r' \in [1, 4]$ such that $x_r^{l,j} \oplus H_{k_r^l}(t^{l,j}) = x_{r'}^{l',j'} \oplus H_{k_{r'}^{l'}}(t^{l',j'})$.

We can write $\mathsf{F}_q = \mathsf{bad}_1 \cup \mathsf{bad}_2 \cup \mathsf{bad}_3$. Observe that bad_1 in the Lemma 4 for Theorem 1 are essentially the same type of event as bad_1 and bad_2 in this proof. In all cases, the event captures a collision among internal intermediate values that are not directly observable by the adversary. Consequently, the probability of bad_1 and bad_2 are the same as bad_1 in Lemma 4, and thus we have

$$\Pr[\mathsf{bad}_1] = \Pr[\mathsf{bad}_2] \leq 1.5q\ell\sqrt{\frac{p}{2^n}} + \frac{q^2\ell}{2^n} + \left(\frac{\epsilon q^2}{2^{\kappa+1}} + \frac{q^2}{2^{n+1}}\right)\left(1 - \frac{1}{\ell}\right) + \frac{\epsilon q^2}{2}.$$

The event bad_3 can also be analyzed in the same manner as bad_3 in Lemma 5. The only difference is that, in the 4-round construction, there are more pairs of distinct rounds that may collide. Accounting for these additional round pairs yields $\Pr[\mathsf{bad}_3] \leq \frac{6\epsilon q^2}{2^\kappa} + \frac{6q^2}{2^n}$. □

The remaining part, $|\Pr[\mathtt{E}_j \wedge \neg \mathtt{F}_j] - \Pr[\mathtt{E}_{j+1} \wedge \neg \mathtt{F}_{j+1}]|$ for all $j \in [0, q-1]$ can be bounded as the same manner as the proof in Section 3.

Lemma 8. *For all $j \in [0, q-1]$,*

$$\left|\Pr\left[\mathtt{E}'_j \wedge \neg \mathtt{F}'_j\right] - \Pr[\mathtt{E}_{j+1} \wedge \neg \mathtt{F}_{j+1}]\right| \leq 2p_{j+1}\sqrt{\frac{2(j+1)}{2^n}}.$$

Proof. The proof is the same as the proof of Lemma 5 except that 4-round TKAF is used as F^f, and $K^l = (k_1^l, k_2^l, k_3^l, k_4^l) \in \{0,1\}^{4\kappa}$ is sampled instead of $K^l = (k_1^l, k_2^l, k_3^l) \in \{0,1\}^{3\kappa}$ for all $l \in [1, \ell]$, in H'_j and H_{j+1}. Since this change does not affect the distinguishability of the ARP distinguisher defined in the proof of Lemma 5, we have the above bound for the modified ARP distinguisher. □

Lemma 9. *For all $j \in [0, q]$,*

$$\left|\Pr[\mathtt{E}_j \wedge \neg \mathtt{F}_j] - \Pr\left[\mathtt{E}'_j \wedge \neg \mathtt{F}'_j\right]\right| \leq 2\sqrt{\frac{p}{2^{n+1}}} + \frac{p}{2^{n+2}} + \frac{q}{2^{2n}} + \frac{8j}{2^n}.$$

Proof. This part is slightly different from the proof of Lemma 6 since we consider the adversary's inverse classical queries. We first introduce intermediate experiments H_j^1 and H_j^2. Let $\delta' = \delta^{\mathsf{p}(j+1), j_{\mathsf{p}(j+1)}+1}$, $t' = t^{\mathsf{p}(j+1), j_{\mathsf{p}(j+1)}+1}$, $x'_r = x_r^{\mathsf{p}(j+1), j_{\mathsf{p}(j+1)}+1}$ and $k'_r = k_r^{\mathsf{p}(j+1)}$ for all $r \in [0, 4]$. H_j^1 proceeds as follows.

Experiment H^1_j. Sample $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell \leftarrow_\$ \tilde{\mathcal{P}}_{\mathcal{T},2n}$, $f \leftarrow_\$ \mathcal{F}_n$, $K^1 = (k^1_1, k^1_2, k^1_3, k^1_4), \ldots, K^\ell = (k^\ell_1, k^\ell_2, k^\ell_3, k^\ell_4) \leftarrow_\$ \{0,1\}^{4\kappa}$.

1. Run $\mathcal{A}$, answering its j classical queries using $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell$ and quantum queries using f. Abort $\mathcal{A}$ immediately after submitting its $(j+1)$-st classical query without receiving the response.
2. The challenger runs the following procedure:
   - if $\delta' = 1$ then:
     (a) It sets $x'_2 = x'_0 \oplus f(x'_1 \oplus H_{k'_1}(t'))$.
     (b) $\mathcal{D}$ samples $\hat{k}_2 \leftarrow_\$ \{0,1\}^\kappa$ and defines the distribution D_1 where $D_1(s) = \Pr\left[x'_2 \oplus H_{\hat{k}_2}(t') = s\right]$ for all $s \in \{0,1\}^n$. Then, it queries Φ with D_1 and outputs s_2. It sets $x'_3 = x'_1 \oplus f(s_2)$.
     (c) $\mathcal{D}$ samples $\hat{k}_3 \leftarrow_\$ \{0,1\}^\kappa$ and defines the distribution D_2 where $D_2(s) = \Pr\left[x'_3 \oplus H_{\hat{k}_3}(t') = s\right]$ for all $s \in \{0,1\}^n$. Then, it queries Φ with D_2 and outputs s_3. It sets $x'_4 = x'_2 \oplus f(s_3)$ and $x'_5 = x'_3 \oplus f(x'_4 \oplus H_{k'_4}(t'))$.
   - if $\delta' = -1$ then:
     (a) It sets $x'_3 = x'_5 \oplus f(x'_4 \oplus H_{k'_4}(t'))$.
     (b) $\mathcal{D}$ samples $\hat{k}_3 \leftarrow_\$ \{0,1\}^\kappa$ and defines the distribution D_2 where $D_2(s) = \Pr\left[x'_3 \oplus H_{\hat{k}_3}(t') = s\right]$ for all $s \in \{0,1\}^n$. Then, it queries Φ with D_2 and outputs s_3. It sets $x'_2 = x'_4 \oplus f(s_3)$.
     (c) $\mathcal{D}$ samples $\hat{k}_2 \leftarrow_\$ \{0,1\}^\kappa$ and defines the distribution D_1 where $D_1(s) = \Pr\left[x'_2 \oplus H_{\hat{k}_2}(t') = s\right]$ for all $s \in \{0,1\}^n$. Then, it queries Φ with D_1 and outputs s_2. It sets $x'_1 = x'_3 \oplus f(s_2)$ and $x'_0 = x'_2 \oplus f(x'_1 \oplus H_{k'_1}(t'))$.

   Let $\nu^{(1)}_{j,K^{[1,\ell]}} = \nu_{j,K^{[1,\ell]}} \cup \nu^{(1)}$ where $\nu^{(1)} = \{s_2, s_3\}$.
3. Resume $\mathcal{A}$.
   (a) If $\nu_{j,K^{[1,\ell]}}$ is invalid then:
     - Answer all subsequent classical queries (including the $(j+1)$-st classical query) from $\mathcal{A}$ with the error symbol $\perp$.
     - Answer all subsequent quantum queries from $\mathcal{A}$ using f.
   (b) Else
     - Answer all subsequent classical queries (including the $(j+1)$-st classical query) from $\mathcal{A}$ using $\mathsf{F}_{K^1}^{f^{\nu^{(1)}_{j,K^{[1,\ell]}}}}, \ldots, \mathsf{F}_{K^\ell}^{f^{\nu^{(1)}_{j,K^{[1,\ell]}}}}$.
     - Answer all subsequent quantum queries from $\mathcal{A}$ using $f^{\nu^{(1)}_{j,K^{[1,\ell]}}}$

Experiment H^2_j. H^2_j is an identical experiment to H^1_j except that the $(j+1)$-st classical query is answered by $\tilde{\Pi}_{\mathsf{p}(j+1)}$.

We can construct an ARS distinguisher from $\mathcal{A}$ between H'_j and H^1_j, which proceeds as follows.

Phase 1: $\mathcal{D}$ samples uniform $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell \leftarrow_\$ \tilde{\mathcal{P}}_{\mathcal{T},2n}$, $f \leftarrow_\$ \mathcal{F}_n$, $K^1 = (k^1_1, k^1_2, k^1_3, k^1_4), \ldots, K^\ell = (k^\ell_1, k^\ell_2, k^\ell_3, k^\ell_4) \leftarrow_\$ \{0,1\}^{4\kappa}$. It then runs $\mathcal{A}$, answering its classical queries

using $\tilde{\Pi}_1, \ldots, \tilde{\Pi}_\ell$, and quantum queries using f. Abort immediately after submitting $\mathcal{A}$'s $(j+1)$-st classical query, without receiving the response.

Phase 2: $\mathcal{D}$ runs the following procedure:

(a) if $\delta' = 1$ then:
   i. $\mathcal{D}$ sets $x'_2 = x'_0 \oplus f(x'_1 \oplus H_{k'_1}(t'))$.
   ii. $\mathcal{D}$ samples $\hat{k}_2 \leftarrow_{\$} \{0,1\}^\kappa$ and defines the distribution D_1 where $D_1(s) = \Pr\left[x'_2 \oplus H_{\hat{k}_2}(t') = s\right]$ for all $s \in \{0,1\}^n$. Then, it queries Φ with D_1 and it outputs s_2. It sets $x'_3 = x'_1 \oplus f(s_2)$.
   iii. $\mathcal{D}$ samples $\hat{k}_3 \leftarrow_{\$} \{0,1\}^\kappa$ and defines the distribution D_2 where $D_2(s) = \Pr\left[x'_3 \oplus H_{\hat{k}_3}(t') = s\right]$ for all $s \in \{0,1\}^n$. Then, it queries Φ with D_2 and it outputs s_3. It sets $x'_4 = x'_2 \oplus f(s_3)$ and $x'_5 = x'_3 \oplus f(x'_4 \oplus H_{k'_4}(t'))$.

(b) if $\delta' = -1$ then:
   i. $\mathcal{D}$ sets $x'_3 = x'_5 \oplus f(x'_4 \oplus H_{k'_4}(t'))$.
   ii. $\mathcal{D}$ samples $\hat{k}_3 \leftarrow_{\$} \{0,1\}^\kappa$ and defines the distribution D_2 where $D_2(s) = \Pr\left[x'_3 \oplus H_{\hat{k}_3}(t') = s\right]$ for all $s \in \{0,1\}^n$. Then, it queries Φ with D_2 and it outputs s_3. It sets $x'_2 = x'_4 \oplus f(s_3)$.
   iii. $\mathcal{D}$ samples $\hat{k}_2 \leftarrow_{\$} \{0,1\}^\kappa$ and defines the distribution D_1 where $D_1(s) = \Pr\left[x'_2 \oplus H_{\hat{k}_2}(t') = s\right]$ for all $s \in \{0,1\}^n$. Then, it queries Φ with D_1 and it outputs s_2. It sets $x'_1 = x'_3 \oplus f(s_2)$ and $x'_0 = x'_2 \oplus f(x'_1 \oplus H_{k'_1}(t'))$.

(c) $\mathcal{D}$ computes $\nu^{(1)}_{j,K^{[1,\ell]}} = \nu_{j,K^{[1,\ell]}} \cup \nu^{(1)}$, where $\nu^{(1)} = \{s_2, s_3\}$.

Phase 3: A uniform bit $b \in \{0,1\}$ is chosen, and $\mathcal{A}$ is given quantum queries f_b, where $f_0 = f^{\nu_{j,K^{[1,\ell]}}}$ and $f_1 = f^{\nu^{(1)}_{j,K^{[1,\ell]}}}$. Under the assumption that $\nu_{j,K^{[1,\ell]}}$ is valid:
- $\mathcal{D}$ answers all remaining classical queries from $\mathcal{A}$ using $\mathsf{F}^{f_b}_{K^1}, \ldots, \mathsf{F}^{f_b}_{K^\ell}$.
- $\mathcal{D}$ answers all remaining quantum queries from $\mathcal{A}$ using f_b.

By the same argument as in Equation (8), we have

$$\left|\Pr[\mathtt{E}_j \wedge \neg \mathtt{F}_j] - \Pr\left[\mathtt{E}^1_j \wedge \neg \mathtt{F}^1_j\right]\right| = \mathsf{Adv}^{\mathsf{ars}}(\mathcal{D}) \leq 2\sqrt{\frac{p}{2^{n+1}}} + \frac{p}{2^{n+2}}. \tag{11}$$

Next, we compare H^1_j and H^2_j. Let $\mathtt{E}^2_j$ be the event that $\mathcal{A}$ in H^2_j outputs 1, and $\mathtt{F}^2_j$ be the event that $\nu_{j,K^{[1,\ell]}}$ is invalid. In H^1_j, the $(j+1)$-st classical query is answered by the oracle indexed by $\mathsf{p}(j+1)$ as $x'_4 = f(s_3) \oplus x'_2$, $x'_5 = f(x'_4 \oplus H_{k'_4}(t')) \oplus x'_1 \oplus f(s_2)$ if $\delta' = 1$, and $x'_0 = f(x'_1 \oplus H_{k'_1}(t')) \oplus x'_4 \oplus f(s_3)$, $x'_1 = f(s_2) \oplus x'_3$ if $\delta' = -1$. In both cases, the outputs (x'_4, x'_5) and (x'_0, x'_1) are uniformly distributed over $\{0,1\}^{2n}$. On the other hand, in H^2_j, the response of $(j+1)$-st classical query is answered by $\tilde{\Pi}_{\mathsf{p}(j+1)}$. Therefore,

$$\left|\Pr\left[\mathtt{E}^1_j \wedge \neg \mathtt{F}^1_j\right] - \Pr\left[\mathtt{E}^2_j \wedge \neg \mathtt{F}^2_j\right]\right| \leq \frac{q_{\mathsf{p}(j+1)}}{2^{2n}} \leq \frac{q}{2^{2n}}. \tag{12}$$

Finally, we compare H_j^2 and H'_j. The bad events that make these experiments differ are that s_2 or s_3 collide with any value in the reprogramming set generated from classical queries to $\tilde{\Pi}_{\mathsf{p}(j+1)}$ up to the $j_{\mathsf{p}(j+1)}$-th classical query.

These experiments differ if at least one of the following bad events occurs.

- bad_4: there exist $i \in [1, j_{\mathsf{p}(j+1)}]$, $r \in [1,4]$ such that $s_2 = x_r^{\mathsf{p}(j+1),i} \oplus H_{k_r^{\mathsf{p}(j+1)}}(t^{\mathsf{p}(j+1),i})$.
- bad_5: there exist $i \in [1, j_{\mathsf{p}(j+1)}]$, $r \in [1,4]$ such that $s_3 = x_r^{\mathsf{p}(j+1),i} \oplus H_{k_r^{\mathsf{p}(j+1)}}(t^{\mathsf{p}(j+1),i})$.

Since s_2 and s_3 follow the distributions D_1 and D_2, respectively, and $\max_{x\in\{0,1\}^n} D_1(x) = \max_{x\in\{0,1\}^n} D_2(x) = 2^{-n}$, we have

$$\left|\Pr\left[\mathsf{E}_j^2 \wedge \neg \mathsf{F}_j^2\right] - \Pr\left[\mathsf{E}'_j \wedge \neg \mathsf{F}'_j\right]\right| \leq \Pr[\mathsf{bad}_4] + \Pr[\mathsf{bad}_5] \leq \frac{8j}{2^n}. \tag{13}$$

From Equations (11) to (13), we have Lemma 9.

□

Then, we have

$$\begin{aligned}
&\mathsf{Adv}_{\mathsf{F}f}^{\mathsf{pq\text{-}stprp,mk}}(\mathcal{A}) = |\Pr[\mathtt{E}_0] - \Pr[\mathtt{E}_q]| \\
&\leq \sum_{j\in[0,q-1]} \left(2p_{j+1}\sqrt{\frac{2(j+1)}{2^n}} + 2\sqrt{\frac{p}{2^{n+1}}} + \frac{p}{2^{n+2}} + \frac{q}{2^{2n}} + \frac{8j}{2^n}. \right) \\
&+ 3q\ell\sqrt{\frac{p}{2^n}} + \frac{q^2(\ell+3)}{2^{n-1}} + \left(\frac{\epsilon q^2}{2^\kappa} + \frac{q^2}{2^n}\right)\left(1 - \frac{1}{\ell}\right) + \frac{6\epsilon q^2}{2^\kappa} + \epsilon q^2 \\
&\leq 2p\sqrt{\frac{2q}{2^n}} + 5\ell q\sqrt{\frac{p}{2^n}} + \frac{14q^2\ell}{2^n} + 5\epsilon q^2.
\end{aligned}$$

□

5 Conclusion

In this paper, we studied the post-quantum security of TKAF in the Q1 model and the multi-key setting. We proved that 3-TKAF achieves PQ-TPRP security and that 4-TKAF achieves PQ-STPRP security, when the tweak is injected via an ϵ-AXU hash function. Our proof extends the hybrid-based analysis of KAF by Basak *et al.* to the tweakable and multi-key setting by analyzing cross-collision events.

Several possible directions for future work can be considered. One direction is to improve the quantitative bounds obtained in this paper. It would also be interesting to study variants of TKAF based on public random permutations rather than public random functions. Another direction is to consider different ways of generating or injecting round keys, including other key schedules. Finally, analyzing higher-round variants, such as constructions with more than five rounds, may lead to stronger post-quantum security.

References

1. Adams, C.: The CAST-128 Encryption Algorithm. RFC 2144 (1997), available at `https://datatracker.ietf.org/doc/html/rfc2144`
2. Alagic, G., Bai, C., Katz, J., Majenz, C.: Post-quantum security of the Even-Mansour cipher. In: EUROCRYPT (3). Lecture Notes in Computer Science, vol. 13277, pp. 458–487. Springer (2022)
3. Alagic, G., Bai, C., Katz, J., Majenz, C., Struck, P.: Post-quantum security of tweakable Even-Mansour, and applications. In: EUROCRYPT (1). Lecture Notes in Computer Science, vol. 14651, pp. 310–338. Springer (2024)
4. Alagic, G., Bai, C., Majenz, C., Shi, K.: Post-quantum security of block cipher constructions. In: Selected Areas in Cryptography (2026), `https://sacworkshop.org/SAC26/preproceedings/SAC2026_1-final18.pdf`, to Appear
5. Andreeva, E., Bogdanov, A., Luykx, A., Mennink, B., Tischhauser, E., Yasuda, K.: Parallelizable and authenticated online ciphers. In: ASIACRYPT (1). Lecture Notes in Computer Science, vol. 8269, pp. 424–443. Springer (2013)
6. Aoki, K., Ichikawa, T., Kanda, M., Matsui, M., Moriai, S., Nakajima, J., Tokita, T.: Camellia: A 128-bit block cipher suitable for multiple platforms - design and analysis. In: Selected Areas in Cryptography. Lecture Notes in Computer Science, vol. 2012, pp. 39–56. Springer (2000)
7. Basak, J., Bhaumik, R., Chauhan, A.K., Jejurikar, R., Jha, A., Roy, A., Schrottenloher, A., Talnikar, S.: Post-quantum security of key-alternating Feistel ciphers. In: ASIACRYPT (1). Lecture Notes in Computer Science, vol. 16245, pp. 446–478. Springer (2025)
8. Bhaumik, R., Cogliati, B., Ethan, J., Jha, A.: Mind the bad norms - revisiting compressed oracle-based quantum indistinguishability proofs. In: ASIACRYPT (9). Lecture Notes in Computer Science, vol. 15492, pp. 215–247. Springer (2024)
9. Bhaumik, R., List, E., Nandi, M.: ZCZ - achieving n-bit SPRP security with a minimal number of tweakable-block-cipher calls. In: ASIACRYPT (1). Lecture Notes in Computer Science, vol. 11272, pp. 336–366. Springer (2018)
10. Datta, N., Dutta, A., Ghosh, S., List, E., Nandi, H.: HCTR+: an optimally secure tbc-based accordion mode. IACR Trans. Symmetric Cryptol. **2025**(3), 183–229 (2025)
11. Don, J., Fehr, S., Huang, Y.: Adaptive versus static multi-oracle algorithms, and quantum security of a split-key PRF. In: TCC (1). pp. 33–51. Lecture Notes in Computer Science, Springer (2022)
12. Grilo, A.B., Hövelmanns, K., Hülsing, A., Majenz, C.: Tight adaptive reprogramming in the QROM. In: ASIACRYPT (1). Lecture Notes in Computer Science, vol. 13090, pp. 637–667. Springer (2021)
13. Hosoyamada, A., Iwata, T.: 4-round Luby-Rackoff construction is a qprp. In: ASIACRYPT (1). Lecture Notes in Computer Science, vol. 11921, pp. 145–174. Springer (2019)
14. Ito, G., Hosoyamada, A., Matsumoto, R., Sasaki, Y., Iwata, T.: Quantum chosen-ciphertext attacks against Feistel ciphers. In: CT-RSA. Lecture Notes in Computer Science, vol. 11405, pp. 391–411. Springer (2019)
15. Kuwakado, H., Morii, M.: Quantum distinguisher between the 3-round Feistel cipher and the random permutation. In: ISIT. pp. 2682–2685. IEEE (2010)
16. Liskov, M.D., Rivest, R.L., Wagner, D.A.: Tweakable block ciphers. In: CRYPTO. Lecture Notes in Computer Science, vol. 2442, pp. 31–46. Springer (2002)

17. National Institute of Standards and Technology: Data Encryption Standard (DES). Federal Information Processing Standards Publication 46-3 (1999), withdrawn May 19, 2005. Available at `https://csrc.nist.gov/pubs/fips/46-3/final`
18. Peyrin, T., Seurin, Y.: Counter-in-tweak: Authenticated encryption modes for tweakable block ciphers. In: CRYPTO (1). Lecture Notes in Computer Science, vol. 9814, pp. 33–63. Springer (2016)
19. Rogaway, P.: Efficient instantiations of tweakable blockciphers and refinements to modes OCB and PMAC. In: ASIACRYPT. Lecture Notes in Computer Science, vol. 3329, pp. 16–31. Springer (2004)
20. Shiba, R., Iwata, T.: Multi-key security in the quantum world: Revisiting tweakable Even-Mansour and FX. IACR Trans. Symmetric Cryptol. **2026**(2) (2026)
21. Simon, D.R.: On the power of quantum computation. In: FOCS. pp. 116–123. IEEE Computer Society (1994)
22. Yan, H., Wang, L., Shen, Y., Lai, X.: Tweaking key-alternating Feistel block ciphers. In: ACNS (1). Lecture Notes in Computer Science, vol. 12146, pp. 69–88. Springer (2020)

A Multi-Target Variants of Collisions-in-the-Dark

In our proof of TKAF, we use the multi-target variant of weak CitD. Although the extension is trivial, we give the proof in this section for completeness. As the original CitD, the multi-target variant of the weak CitD variant is a union bound of the strong variant. Therefore, in this section, we first introduce the multi-target strong CitD (mSCitD) and then we prove the multi-target weak CitD (mWCitD)

A.1 Multi-Target Strong Collisions-in-the-Dark

mSCitD *game.* Let $\mathcal{X} = \mathcal{X}_0 \times \mathcal{X}_1$ and $\mathcal{Y}$ be two finite sets, where $\mathcal{X}_1 = \{0,1\}^n$. Let $F : \mathcal{X}_1 \to \mathcal{Y}$ be a random oracle. Fix some relation $\mathbb{R} \in \mathcal{X}^2 \times \mathcal{Y}^2$. The mSCitD game is played between a challenger and an adversary $\mathcal{A}$.

Phase 1. The adversary $\mathcal{A}$ makes p quantum queries to the random oracle F and returns a tuple $(\mathcal{X}^*, x^*)$, where $\mathcal{X}^* \subseteq \mathcal{X}$, $|\mathcal{X}^*| = q$, and $x^* = (x_0^*, x_1^*) \in \mathcal{X}$.

Phase 2. The challenger samples ℓ keys $\tilde{k}^1, \ldots, \tilde{k}^\ell \leftarrow \mathcal{X}_1$. $\mathcal{A}$ wins if the following event holds:

$$\mathsf{mSC}_{\mathbb{R}}(\mathcal{X}^*, x^*) : \ \exists\, x = (x_0, x_1) \in \mathcal{X}^*,\ \exists\, l \in [1, \ell] \ \text{ such that } (x, x^*, F(\tilde{k}^l \oplus x_1),\ F(\tilde{k}^l \oplus x_1^*)) \in \mathbb{R}.$$

Lemma 10 (Multi-Target Strong Collisions-in-the-Dark (adapted from [7])). *For any p quantum query adversary $\mathcal{A}$ in the* mSCitD *game with respect to the relation $\mathbb{R}$, it holds that*

$$\Pr_{(\mathcal{X}^*, x^*)}[\mathsf{mSC}_{\mathbb{R}}(X^*, x^*)] \leq 1.5\ell\sqrt{\frac{p}{2^n}} + q\ell \max_{x, x^* \in \mathcal{X}_1} \Pr_{\substack{r \leftarrow_\$ \mathcal{Y} \\ s \leftarrow_\$ \mathcal{X}_1}}[(x, x^*, r, F(s)) \in \mathbb{R}].$$

Proof. We construct the ARS distinguisher $\mathcal{D}$ from $\mathcal{A}$ that is given quantum access to $F : \mathcal{X}_1 \to \mathcal{Y}$ and classical access to the sampling oracle Φ. The ARS distinguisher we consider here is slightly different from the experiment in Lemma 2. Here, the resampling oracle Φ is invoked with an empty input ($\perp$), *i.e.*, we do not specify an input distribution to Φ. Therefore, the relevant distinguishing advantage is captured by [12, Proposition 1], and it is bounded as

$$\mathsf{Adv}^{\mathsf{ars}}(\mathcal{D}) = \left|\Pr\left[\mathcal{D}^{F_b,\Phi} = 1 \mid b = 1\right] - \Pr\left[\mathcal{D}^{F_b,\Phi} = 1 \mid b = 0\right]\right| \leq 1.5R\sqrt{\frac{p}{2^n}}, \tag{14}$$

where p is the number of quantum queries and R is the number of classical queries to Φ.

$\mathcal{D}$ proceeds as follows.

1. Run $\mathcal{A}$ which makes p quantum queries to F.
2. Once $\mathcal{A}$ has made all its p quantum queries, it measures a register and returns the pair $(\mathcal{X}^*, x^*)$ where $\mathcal{X}^*$ has size q and $x^* = (x_0^*, x_1^*)$.
3. $\mathcal{D}$ makes ℓ classical queries to Φ and gets uniformly random values $s_1, \ldots, s_\ell$.
4. $\mathcal{D}$ gets $\tilde{k}_1 = x_1^* \oplus s_1, \ldots, \tilde{k}_\ell = x_1^* \oplus s_\ell$.
5. Now if $\mathsf{mSC}_{\mathbb{R}}(\mathcal{X}^*, x^*)$ holds then $\mathcal{D}$ returns 0, and 1 otherwise.

When $\mathcal{D}$ is interacting with F_0, it perfectly simulates the mSCitD game for $\mathcal{A}$. Thus,

$$\Pr\left[\mathcal{D}^{F_b,\Phi} = 1 \mid b = 0\right] = \Pr_{(\mathcal{X}^*,x^*)\leftarrow\mathcal{A}}[\mathsf{mSC}_{\mathbb{R}}(\mathcal{X}^*, x^*)]. \tag{15}$$

On the other hand, when $\mathcal{D}$ is interacting with F_1, $F(s_1) = r_1, \ldots, F(s_\ell) = r_\ell$ are uniformly distributed and independent of each other and all other variables. Thus,

$$\begin{aligned}
\Pr\left[\mathcal{D}^{F_b,\Phi} = 1 \mid b = 1\right] &\leq \sum_{l\in[1,\ell]} \Pr_{\substack{r_l\leftarrow_{\$}Y \\ s_l\leftarrow_{\$}\mathcal{X}_1}}[\exists x \in \mathcal{X}^* : (x, x^*, r_l, F(s_l \oplus x_1^* \oplus x_1)) \in \mathbb{R}] \\
&\leq \ell \sum_{x\in\mathcal{X}^*} \Pr_{\substack{r\leftarrow_{\$}\mathcal{Y} \\ s\leftarrow_{\$}\mathcal{X}_1}}[(x, x^*, r, F(s \oplus x_1^* \oplus x_1)) \in \mathbb{R}] \\
&\leq q\ell \max_{x,x^*\in\mathcal{X}_1} \Pr_{\substack{r\leftarrow_{\$}\mathcal{Y} \\ s\leftarrow_{\$}\mathcal{X}_1}}[(x, x^*, r, F(s)) \in \mathbb{R}]
\end{aligned} \tag{16}$$

By subtracting Equation (15) from Equation (16) and using Equation (14) ([12, Proposition 1]), we have

$$\Pr_{(\mathcal{X}^*,x^*)}[\mathsf{mSC}_{\mathbb{R}}(X^*, x^*)] \leq 1.5\ell\sqrt{\frac{p}{2^n}} + q\ell \max_{x,x^*\in\mathcal{X}_1} \Pr_{\substack{r\leftarrow_{\$}\mathcal{Y} \\ s\leftarrow_{\$}\mathcal{X}_1}}[(x, x^*, r, F(s)) \in \mathbb{R}].$$

A.2 Multi-Target Weak Collisions-in-the-Dark

Proof (Proof of Lemma 3). Let $(x^1, \ldots, x^q)$ be a fixed ordering of $\mathcal{X}^*$. For all $j \in \{1, \ldots, q-1\}$, let $\mathcal{X}_j^* = \{x^1, \ldots, x^j\}$, and define

$$\begin{aligned}\mathtt{E}_j : \text{there exist } x \in \mathcal{X}_j^* \text{ and } l \in [1, \ell] \text{ such that}\\ (x, x^{j+1}, F(k^l \oplus x_1), F(k^l \oplus x_1^{j+1})) \in \mathbb{R}\end{aligned}$$

Clearly, $\mathsf{mWC}_{\mathbb{R}}(\mathcal{X}^*) = \mathtt{E}_1 \vee \cdots \vee \mathtt{E}_{q-1}$. Therefore, taking the union bound, we have

$$\Pr_{\mathcal{X}^* \leftarrow \mathcal{A}}[\mathsf{mWC}_{\mathbb{R}}(\mathcal{X}^*)] = \Pr[\mathtt{E}_1 \vee \cdots \vee \mathtt{E}_{q-1}] \leq \sum_{j \in [1, q-1]} \Pr[\mathtt{E}_j].$$

For any fixed j, we can construct the mSCitD distinguisher $\mathcal{D}_j$ whose advantage corresponds to $\Pr[\mathtt{E}_j]$. Then,

$$\Pr_{\mathcal{X}^* \leftarrow \mathcal{A}}[\mathsf{mSC}_{\mathbb{R}}(\mathcal{X}^*, x^{j+1})] = \Pr[\mathtt{E}_j].$$

From Lemma 10,

$$\Pr[\mathtt{E}_j] \leq 1.5\ell\sqrt{\frac{p}{2^n}} + j\ell \max_{x, x^* \in \mathcal{X}_1} \Pr_{\substack{r \leftarrow_\$ \mathcal{Y} \\ s \leftarrow_\$ \mathcal{X}_1}}[(x, x^*, r, F(s)) \in \mathbb{R}].$$

Taking the union bound,

$$\begin{aligned}\sum_{j \in [1, q-1]} \Pr[\mathtt{E}_j] &\leq \sum_{j \in [1, q-1]} \left(1.5\ell\sqrt{\frac{p}{2^n}} + j\ell \max_{x, x^* \in \mathcal{X}_1} \Pr_{\substack{r \leftarrow_\$ \mathcal{Y} \\ s \leftarrow_\$ \mathcal{X}_1}}[(x, x^*, r, F(s)) \in \mathbb{R}] \right) \\ &\leq 1.5q\ell\sqrt{\frac{p}{2^n}} + q^2\ell \max_{x, x^* \in \mathcal{X}_1} \Pr_{\substack{r \leftarrow_\$ \mathcal{Y} \\ s \leftarrow_\$ \mathcal{X}_1}}[(x, x^*, r, F(s)) \in \mathbb{R}].\end{aligned}$$