

Watermarking PRFs against Quantum Adversaries

Fuyuki Kitagawa¹

Ryo Nishimaki¹

¹ NTT Corporation, Tokyo, Japan
{fuyuki.kitagawa.yh,ryo.nishimaki.zk}@hco.ntt.co.jp

Abstract

We initiate the study of software watermarking against quantum adversaries. A quantum adversary generates a *quantum state* as a pirate software that potentially removes an embedded message from a *classical* marked software. Extracting an embedded message from quantum pirate software is difficult since measurement could irreversibly alter the quantum state. In software watermarking against classical adversaries, a message extraction algorithm crucially uses the (input-output) behavior of a classical pirate software to extract an embedded message. Even if we instantiate existing watermarking PRFs with quantum-safe building blocks, it is not clear whether they are secure against quantum adversaries due to the quantum-specific property above. Thus, we need entirely new techniques to achieve software watermarking against quantum adversaries.

In this work, we define secure watermarking PRFs for quantum adversaries (unremovability against quantum adversaries). We also present two watermarking PRFs as follows.

- We construct a privately extractable watermarking PRF against quantum adversaries from the quantum hardness of the learning with errors (LWE) problem. The marking and extraction algorithms use a public parameter and a private extraction key, respectively. The watermarking PRF is unremovable even if adversaries have (the public parameter and) access to the extraction oracle, which returns a result of extraction for a queried quantum circuit.
- We construct a publicly extractable watermarking PRF against quantum adversaries from indistinguishability obfuscation (IO) and the quantum hardness of the LWE problem. The marking and extraction algorithms use a public parameter and a public extraction key, respectively. The watermarking PRF is unremovable even if adversaries have the extraction key (and the public parameter).

We develop a quantum extraction technique to extract information (a classical string) from a quantum state without destroying the state too much. We also introduce the notion of extraction-less watermarking PRFs as a crucial building block to achieve the results above by combining the tool with our quantum extraction technique.

Keywords: watermarking, pseudorandom function, post-quantum cryptography

Contents

1	Introduction	1
1.1	Background	1
1.2	Our Result	2
1.3	Technical Overview	3
1.4	More on Related Work	8
2	Preliminaries	10
2.1	Quantum Information	10
2.2	Measurement Implementation	11
2.3	Cryptographic Tools	15
3	Definition of Quantum Watermarking	19
3.1	Syntax and Pseudorandomness	19
3.2	Unremovability against Quantum Adversaries	21
4	Definition of Extraction-Less Watermarking	23
4.1	Syntax and Pseudorandomness	23
4.2	Simulatability for Mark-Dependent Distributions (SIM-MDD Security)	23
5	Watermarking PRF from Extraction-Less Watermarking PRF	25
6	Extraction-Less Watermarking PRF from LWE	28
7	Extraction-Less Watermarking PRF with Public Simulation from IO	34
7.1	Puncturable Encryption, Revisited	34
7.2	Construction of Extraction-less Watermarking PRF with Public Simulation	35
8	Putting Pieces Altogether	43
A	Achieving QSIM-MDD from SIM-MDD	48
A.1	Semi-Classical One-Way to Hiding (O2H) Lemma	48
A.2	Proof	48
B	Puncturable Encryption with Strong Ciphertext Pseudorandomness	50
B.1	Tools for PE	50
B.2	PE Scheme Description	51
B.3	PE Security Proof	52
B.4	Original Ciphertext Pseudorandomness of PE	60

1 Introduction

1.1 Background

Software watermarking is a cryptographic primitive that achieves a digital analog of watermarking. A marking algorithm of software watermarking can embed an arbitrary message (bit string) into a computer software modeled as a circuit. A marked software almost preserves the functionality of the original software. An extraction algorithm of software watermarking can extract the embedded message from a marked software. Secure software watermarking should guarantee that no adversary can remove the embedded message without significantly destroying the functionality of the original software (called unremovability).

Barak, Goldreich, Impagliazzo, Rudich, Sahai, Vadhan, and Yang [BGI⁺12] initiate the study of software watermarking and present the first definition of cryptographically secure software watermarking. Hopper, Molnar, and Wagner [HMW07] also study the definition of cryptographically secure watermarking for perceptual objects. However, both works do not present a secure concrete scheme. A few works study secure constructions of watermarking for cryptographic primitives [NSS99, YF11, Nis13, Nis19], but they consider only restricted removal strategies. Cohen, Holmgren, Nishimaki, Wichs, and Vaikuntanathan [CHN⁺18] present stronger definitions for software watermarking and the first secure watermarking schemes for cryptographic primitives *against arbitrary removal strategies*. After the celebrated work, watermarking for cryptographic primitives have been extensively studied [BLW17, KW21, QWZ18, KW19, YAL⁺19, GKM⁺19, YAYX20, Nis20].

Primary applications of watermarking are identifying ownership of objects and tracing users that distribute illegal copies. Watermarking for cryptographic primitives also has another exciting application. Aaronson, Liu, Liu, Zhandry, and Zhang [ALL⁺21] and Kitagawa, Nishimaki, and Yamakawa [KNY21] concurrently and independently find that we can construct secure software leasing schemes by combining watermarking with quantum cryptography.¹ Secure software leasing [AL21] is a quantum cryptographic primitive that prevents users from generating authenticated pirated copies of leased software.² Since watermarking has such an exciting application in quantum cryptography and quantum computers might be an imminent threat to cryptography due to rapid progress in research on quantum computing, it is natural and fascinating to study secure software watermarking in the quantum setting.

In quantum cryptography, building blocks must be quantum-safe such as lattice-based cryptography [Reg09]. However, even if we replace building blocks of existing cryptographic primitives/protocols with quantum-safe ones, we do not necessarily obtain quantum-safe cryptographic primitives/protocols [BDF⁺11, ARU14]. We sometimes need new proof techniques which are different from classical ones due to quantum specific properties such as no-cloning and superposition access [Wat09, Zha12b, Zha12a, Unr12, Zha19, CMSZ21]. Even worse, we must consider entirely different security models in some settings. Zhandry [Zha20] studies traitor tracing [CFN94] in the quantum setting as such an example. In quantum traitor tracing, an adversary can output a *quantum state* as a pirate decoder. Zhandry shows that we need new techniques for achieving quantum traitor tracing because running a quantum pirate decoder to extract information may irreversibly alter the state due to measurement.

Zhandry [Zha20] refers to software watermarking as a cryptographic primitive that has a similar issue to quantum traitor tracing. However, his work focuses only on traitor tracing and does not study software watermarking against quantum adversaries. If we use software watermarking in the quantum setting, an adversary can output a *quantum state* as a pirate circuit where an embedded message might be removed. However, previous works consider a setting where an adversary outputs a *classical* pirate circuit. It is not clear whether watermarking schemes based on quantum-safe cryptography are secure against quantum

¹Precisely speaking, Aaronson et al. achieve copy-detection schemes [ALL⁺21], which are essentially the same as secure software leasing schemes.

²Leased software must be a quantum state since classical bit strings can be easily copied.

adversaries because we need an entirely new extraction algorithm to extract an embedded message from a *quantum* pirate circuit. Thus, the main question in this study is:

Can we achieve secure watermarking for cryptographic primitives against quantum adversaries?

We affirmatively answer this question in this work.

1.2 Our Result

Our main contributions are two-fold. One is the definitional work. We define watermarking for pseudorandom functions (PRFs) against quantum adversaries, where adversaries output a quantum state as a pirate circuit that distinguishes a PRF from a random function.³ The other one is constructing the first secure watermarking PRFs against quantum adversaries. We present two watermarking PRFs as follows.

- We construct a privately extractable watermarking PRF against quantum adversaries from the quantum hardness of the learning with errors (LWE) problem. This watermarking PRF is secure in the presence of the extraction oracle and supports public marking. That is, the marking and extraction algorithms use a public parameter and secret extraction key, respectively. The watermarking PRF is unremovable even if adversaries have access to the extraction oracle, which returns a result of extraction for a queried quantum circuit.
- We construct a publicly extractable watermarking PRF against quantum adversaries from indistinguishability obfuscation (IO) and the quantum hardness of the LWE problem. This watermarking PRF also supports public marking. That is, the marking and extraction algorithms use a public parameter and a public extraction key, respectively. The watermarking PRF is unremovable (we do not need to consider the mark and extraction oracles since it supports public marking and public extraction).

The former and latter PRFs satisfy weak pseudorandomness and standard (strong) pseudorandomness even against a watermarking authority, respectively.

We develop a quantum extraction algorithm to achieve the results above. Zhandry [Zha20] presents a useful technique for extracting information from quantum states without destroying them too much. However, we cannot simply apply his technique to the watermarking setting. Embedded information (arbitrary string) is chosen from an exponentially large set in the watermarking setting. On the other hand, in the traitor tracing setting, we embed a user index, which could be chosen from a polynomially large set, in a decryption key. Zhandry’s technique is tailored to traitor tracing based on private linear broadcast encryption (PLBE) [BSW06] where user information is chosen from a polynomially large set with linear structure. Thus, we extend Zhandry’s technique [Zha20] to extract information chosen from an exponentially large set. We also introduce the notion of extraction-less watermarking as a crucial tool to achieve watermarking against quantum adversaries. This tool is a suitable building block for our quantum extraction technique in our watermarking extraction algorithm. These are our technical contributions. See Section 1.3 for the detail.

Although this paper focuses on watermarking PRFs against quantum adversaries, it is easy to extend our definitions to watermarking public-key encryption (PKE) against quantum adversaries. In particular, our construction technique easily yields watermarking PKE (where a decryption circuit is marked) schemes. We will provide the detail of them in a future version.

We also focus on watermarking PRFs with public marking in this paper. However, we can easily convert our PRFs into ones with private marking. See Remark 3.4 for the detail.

³This definitional choice comes from the definition of traceable PRFs [GKWW21]. See Sections 1.3 and 1.4 for the detail.

1.3 Technical Overview

Syntax of watermarking PRF. We first review the syntax of watermarking PRF used in this work. A watermarking PRF scheme consists of five algorithms (Setup , Gen , Eval , Mark , $\mathcal{E}\text{xtract}$).⁴ Setup outputs a public parameter pp and an extraction key xk . Gen is given pp and outputs a PRF key prfk and a public tag τ . Eval is the PRF evaluation algorithm that takes as an input prfk and x in the domain and outputs y . By using Mark , we can generate a marked evaluation circuit that has embedded message $m \in \{0, 1\}^{\ell_m}$ and can be used to evaluate $\text{Eval}(\text{prfk}, x')$ for almost all x' . Finally, $\mathcal{E}\text{xtract}$ is the extraction algorithm supposed to extract the embedded message from a pirated quantum evaluation circuit generated from the marked evaluation circuit. By default, in this work, we consider the public marking setting, where anyone can execute Mark . Thus, Mark takes pp as an input. On the other hand, we consider both the private extraction and the public extraction settings. Thus, the extraction key xk used by $\mathcal{E}\text{xtract}$ is kept secret by an authority in the private extraction setting and made public in the public extraction setting.

In this work, we allow $\mathcal{E}\text{xtract}$ to take the public tag τ generated with the original PRF key corresponding to the pirate circuit. In reality, we execute $\mathcal{E}\text{xtract}$ for a software when a user claims that the software is illegally generated by using her/his PRF key. Thus, it is natural to expect we can use a user’s public tag for extraction. Moreover, pirate circuits are distinguishers, not predictors in this work. As discussed by Goyal et al. [GKWW21], security against pirate distinguishers is much preferable compared to security against pirate predictors considered in many previous works on watermarking. In this case, it seems that such additional information fed to $\mathcal{E}\text{xtract}$ is unavoidable. For a more detailed discussion on the syntax, see the discussion in Section 3.1.

It is also natural to focus on distinguishers breaking weak pseudorandomness of PRFs when we consider pirate distinguishers instead of pirate predictors. Goyal et al. [GKWW21] already discussed this point. Thus, we focus on watermarking weak PRF in this work.

Definition of unremovability against quantum adversaries. We say that a watermarking PRF scheme satisfies unremovability if given a marked evaluation circuit \tilde{C} that has an embedded message m , any adversary cannot generate a circuit such that it is a “good enough circuit”, but the extraction algorithm fails to output m . In this work, we basically follow the notion of “good enough circuit” defined by Goyal et al. [GKWW21] as stated above. Let D be the following distribution for a PRF $\text{Eval}(\text{prfk}, \cdot) : \text{Dom} \rightarrow \text{Ran}$.

D : Generate $b \leftarrow \{0, 1\}$, $x \leftarrow \text{Dom}$, and $y_0 \leftarrow \text{Ran}$. Compute $y_1 \leftarrow \text{Eval}(\text{prfk}, x)$. Output (b, x, y_b) .

A circuit is defined as good enough circuit with respect to $\text{Eval}(\text{prfk}, \cdot)$ if given (x, y_b) output by D , it can correctly guess b with probability significantly greater than $1/2$. In other words, a circuit is defined as good enough if the circuit breaks weak PRF security.

Below, for a distribution D' whose output is of the form (b, x, y) , let $\mathcal{M}_{D'} = (\mathcal{M}_{D',0}, \mathcal{M}_{D',1})$ be binary positive operator valued measures (POVMs) that represents generating random (b, x, y) from D' and testing if a quantum circuit can guess b from (x, y) . Then, for a quantum state $|\psi\rangle$, the overall distinguishing advantage of it for the above distribution D is $\langle \psi | \mathcal{M}_{D,0} | \psi \rangle$. Thus, a natural adaptation of the above notion of goodness for quantum circuits might be to define a quantum state $|\psi\rangle$ as good if $\langle \psi | \mathcal{M}_{D,0} | \psi \rangle$ is significantly greater than $1/2$. However, this notion of goodness for quantum circuits is not really meaningful. The biggest issue is that it does not consider the stateful nature of quantum programs.

This issue was previously addressed by Zhandry [Zha20] in the context of traitor tracing against quantum adversaries. In the context of classical traitor tracing or watermarking, we can assume that a pirate circuit is stateless, or can be rewound to its original state. This assumption is reasonable. If we have the software description of the pirate circuit, such a rewinding is trivial. Even if we have a hardware box in which a pirate circuit is built, it seems that such a rewinding is possible by hard reboot or cutting

⁴In this paper, standard math font stands for classical algorithms, and calligraphic font stands for quantum algorithms.

power. On the other hand, in the context of quantum watermarking, we have to consider that a pirate circuit is inherently stateful since it is described as a quantum state. Operations to a quantum state can alter the state, and in general, it is impossible to rewind the state into its original state. Regarding the definition of good quantum circuits above, if we can somehow compute the average success probability $\langle \psi | \mathcal{M}_{D,0} | \psi \rangle$ of the quantum state $|\psi\rangle$, the process can change or destroy the quantum state $|\psi\rangle$. Namely, even if we once confirm that the quantum state $|\psi\rangle$ is good by computing $\langle \psi | \mathcal{M}_{D,0} | \psi \rangle$, we cannot know the success probability of the quantum state even right after the computation. Clearly, the above notion of goodness is not the right notion, and we need one that captures the stateful nature of quantum programs.

In the work on traitor tracing against quantum adversaries, Zhandry [Zha20] proposed a notion of goodness for quantum programs that solves the above issue. We adopt it. For the above POVMs \mathcal{M}_D , let \mathcal{M}'_D be the projective measurement $\{P_p\}_{p \in [0,1]}$ that projects a state onto the eigenspaces of $\mathcal{M}_{D,0}$, where each p is an eigenvalue of $\mathcal{M}_{D,0}$. \mathcal{M}'_D is called projective implementation of \mathcal{M}_D and denoted as $\text{ProjImp}(\mathcal{M}_D)$. Zhandry showed that the following process has the same output distribution as \mathcal{M}_D :

1. Apply the projective measurement $\mathcal{M}'_D = \text{ProjImp}(\mathcal{M}_D)$ and obtain p .
2. Output 0 with probability p and output 1 with probability $1 - p$.

Intuitively, \mathcal{M}'_D project a state to an eigenvector of $\mathcal{M}_{D,0}$ with eigenvalue p , which can be seen as a quantum state with success probability p . Using \mathcal{M}'_D , Zhandry defined that a quantum circuit is Live if the outcome of the measurement \mathcal{M}'_D is significantly greater than $1/2$. The notion of Live is a natural extension of the classical goodness since it collapses to the classical goodness for a classical decoder. Moreover, we can ensure that a quantum state that is tested as Live still has a high success probability. On the other hand, the above notion of goodness cannot say anything about the post-tested quantum state's success probability even if the test is passed. In this work, we use the notion of Live quantum circuits as the notion of good quantum circuits.

Difficulty of quantum watermarking PRF. From the above discussion, our goal is to construct a watermarking PRF scheme that guarantees that we can extract the embedded message correctly if a pirated quantum circuit is Live. In watermarking PRF schemes, we usually extract an embedded message by applying several tests on success probability to a pirate circuit. When a pirate circuit is a quantum state, the set of tests that we can apply is highly limited compared to a classical circuit due to the stateful nature of quantum states.

One set of tests we can apply without destroying the quantum state is $\text{ProjImp}(\mathcal{M}_{D'})$ for distributions D' that are indistinguishable from D from the view of the pirate circuit.⁵ We denote this set as $\{\text{ProjImp}(\mathcal{M}_{D'}) \mid D' \stackrel{\epsilon}{\approx} D\}$. Zhandry showed that if distributions D_1 and D_2 are indistinguishable, the outcome of $\text{ProjImp}(\mathcal{M}_{D_1})$ is close to that of $\text{ProjImp}(\mathcal{M}_{D_2})$. By combining this property with the projective property of projective implementations, as long as the initial quantum state is Live and we apply only tests contained in $\{\text{ProjImp}(\mathcal{M}_{D'}) \mid D' \stackrel{\epsilon}{\approx} D\}$, the quantum state remains Live. On the other hand, if we apply a test outside of $\{\text{ProjImp}(\mathcal{M}_{D'}) \mid D' \stackrel{\epsilon}{\approx} D\}$, the quantum state might be irreversibly altered. This fact is a problem since the set $\{\text{ProjImp}(\mathcal{M}_{D'}) \mid D' \stackrel{\epsilon}{\approx} D\}$ only is not sufficient to implement the existing widely used construction method for watermarking PRF schemes.

To see this, we briefly review the method. In watermarking PRF schemes, the number of possible embedded messages is super-polynomial, and thus we basically need to extract an embedded message in a bit-by-bit manner. In the method, such a bit-by-bit extraction is done as follows. For every $i \in [\ell_m]$, we define two distributions $S_{i,0}$ and $S_{i,1}$ whose output is of the form (b, x, y) as D above. Then, we design a marked circuit with embedded message $m \in \{0, 1\}^{\ell_m}$ so that it can be used to guess b from (x, y) with probability significantly greater than $1/2$ only for $S_{i,0}$ (resp. $S_{i,1}$) if $m[i] = 0$ (resp. $m[i] = 1$). The

⁵In the actual extraction process, we use an approximation of projective implementation introduced by Zhandry [Zha20] since applying a projective implementation is inefficient. In this overview, we ignore this issue for simplicity.

extraction algorithm can extract i -th bit of the message $m[i]$ by checking for which distributions of $S_{i,0}$ and $S_{i,1}$ a pirate circuit has a high distinguishing advantage.

As stated above, we cannot use this standard method to extract a message from quantum pirate circuits. The reason is that $S_{i,0}$ and $S_{i,1}$ are typically distinguishable. This implies that at least either one of $\text{ProjImp}(\mathcal{M}_{S_{i,0}})$ or $\text{ProjImp}(\mathcal{M}_{S_{i,1}})$ is not contained in $\{\text{ProjImp}(\mathcal{M}_{D'}) \mid D' \stackrel{c}{\approx} D\}$. Since the test outside of $\{\text{ProjImp}(\mathcal{M}_{D'}) \mid D' \stackrel{c}{\approx} D\}$ might destroy the quantum state, we might not be able to perform the process for all i , and fail to extract the entire bits of the embedded message.

It seems that to perform the bit-by-bit extraction for a quantum state, we need to extend the set of applicable tests and come up with a new extraction method.

Our solution: Use of reverse projective property. We find that as another applicable set of tests, we have $\text{ProjImp}(\mathcal{M}_{D'})$ for distributions D' that are indistinguishable from D^{rev} , where D^{rev} is the following distribution.

D^{rev} : Generate $b \leftarrow \{0,1\}$, $x \leftarrow \text{Dom}$, and $y_0 \leftarrow \text{Ran}$. Compute $y_1 \leftarrow \text{Eval}(\text{prfk}, x)$. Output $(1 \oplus b, x, y_b)$.

We denote the set as $\{\text{ProjImp}(\mathcal{M}_{D'}) \mid D' \stackrel{c}{\approx} D^{\text{rev}}\}$. D^{rev} is the distribution the first bit of whose output is flipped from that of D . Then, $\mathcal{M}_{D^{\text{rev}}}$ can be seen as POVMs that represents generating random (b, x, y_b) from D and testing if a quantum circuit *cannot* guess b from (x, y_b) . Thus, we see that $\mathcal{M}_{D^{\text{rev}}} = (\mathbf{M}_{D,1}, \mathbf{M}_{D,0})$. Recall that $\mathcal{M}_D = (\mathbf{M}_{D,0}, \mathbf{M}_{D,1})$.

Let $D_1 \in \{\text{ProjImp}(\mathcal{M}_{D'}) \mid D' \stackrel{c}{\approx} D\}$ and D_1^{rev} be the distribution that generates $(b, x, y) \leftarrow D_1$ and outputs $(1 \oplus b, x, y)$. D_1^{rev} is a distribution contained in $\{\text{ProjImp}(\mathcal{M}_{D'}) \mid D' \stackrel{c}{\approx} D^{\text{rev}}\}$. Similarly to the relation between D and D^{rev} , if $\mathcal{M}_{D_1} = (\mathbf{M}_{D_1,0}, \mathbf{M}_{D_1,1})$, we have $\mathcal{M}_{D_1^{\text{rev}}} = (\mathbf{M}_{D_1^{\text{rev}},1}, \mathbf{M}_{D_1^{\text{rev}},0})$. Since $\mathbf{M}_{D_1,0} + \mathbf{M}_{D_1,1} = \mathbf{I}$, $\mathbf{M}_{D_1,0}$ and $\mathbf{M}_{D_1,1}$ share the same set of eigenvectors, and if a vector is an eigenvector of $\mathbf{M}_{D_1,0}$ with eigenvalue p , then it is also an eigenvector of $\mathbf{M}_{D_1,1}$ with eigenvalue $1 - p$. Thus, if apply $\text{ProjImp}(\mathcal{M}_{D_1})$ and $\text{ProjImp}(\mathcal{M}_{D_1^{\text{rev}}})$ successively to a quantum state and obtain the outcomes \tilde{p}_1 and \tilde{p}'_1 , it holds that $\tilde{p}'_1 = 1 - \tilde{p}_1$. We call this property the reverse projective property of the projective implementation.

Combining projective and reverse projective properties and the outcome closeness for indistinguishable distributions of the projective implementation, we see that the following key fact holds.

Key fact: As long as the initial quantum state is Live and we apply tests contained in $\{\text{ProjImp}(\mathcal{M}_{D'}) \mid D' \stackrel{c}{\approx} D\}$ or $\{\text{ProjImp}(\mathcal{M}_{D'}) \mid D' \stackrel{c}{\approx} D^{\text{rev}}\}$, the quantum state remains Live. Moreover, if the outcome of applying $\text{ProjImp}(\mathcal{M}_D)$ to the initial state is p , we get the outcome close to p every time we apply a test in $\{\text{ProjImp}(\mathcal{M}_{D'}) \mid D' \stackrel{c}{\approx} D\}$, and we get the outcome close to $1 - p$ every time we apply a test in $\{\text{ProjImp}(\mathcal{M}_{D'}) \mid D' \stackrel{c}{\approx} D^{\text{rev}}\}$.

In this work, we perform bit-by-bit extraction of embedded messages by using the above key fact of the projective implementation. To this end, we introduce the new notion of extraction-less watermarking PRF as an intermediate primitive.

Via extraction-less watermarking PRF. An extraction-less watermarking PRF scheme has almost the same syntax as a watermarking PRF scheme, except that it does not have an extraction algorithm $\mathcal{E}xtract$ and instead has a simulation algorithm Sim . Sim is given the extraction key xk , the public tag τ , and an index $i \in [\ell_m]$, and outputs a tuple of the form (γ, x, y) . Sim simulates outputs of D or D^{rev} for a pirate circuit depending on the message embedded to the marked circuit corresponding to the pirate circuit. More concretely, we require that from the view of the pirate circuit generated from a marked circuit with embedded message $m \in \{0,1\}^{\ell_m}$, outputs of Sim are indistinguishable from those of D if $m[i] = 0$ and

are indistinguishable from those of D^{rev} if $m[i] = 1$ for every $i \in [\ell_m]$. We call this security notion simulatability for mark-dependent distributions (SIM-MDD security).

By using an extraction-less watermarking PRF scheme ELWMPRF, we construct a watermarking PRF scheme WMPRF against quantum adversaries as follows. We use Setup, Gen, Eval, Mark of ELWMPRF as Setup, Gen, Eval, Mark of WMPRF, respectively. We explain how to construct the extraction algorithm $\mathcal{E}xtract$ of WMPRF using Sim of ELWMPRF. For every $i \in [\ell_m]$, we define $D_{\tau,i}$ as the distribution that outputs randomly generated $(\gamma, x, y) \leftarrow \text{Sim}(xk, \tau, i)$. Given xk , τ , and a quantum state $|\psi\rangle$, $\mathcal{E}xtract$ extracts the embedded message in the bit-by-bit manner by repeating the following process for every $i \in [\ell_m]$.

- Apply $\text{ProjImp}(\mathcal{M}_{D_{\tau,i}})$ to $|\psi_{i-1}\rangle$ and obtain the outcome \tilde{p}_i , where $|\psi_0\rangle = |\psi\rangle$ and $|\psi_{i-1}\rangle$ is the state after the $(i-1)$ -th loop for every $i \in [\ell_m]$.
- Set $m'_i = 0$ if $\tilde{p}_i > 1/2$ and otherwise $m'_i = 1$.

The extracted message is set to $m'_1 \parallel \dots \parallel m'_{\ell_m}$.

We show that the above construction satisfies unremovability. Suppose an adversary is given marked circuit $\tilde{C} \leftarrow \text{Mark}(pp, prfk, m)$ and generates a quantum state $|\psi\rangle$, where $(pp, xk) \leftarrow \text{Setup}(1^\lambda)$ and $(prfk, \tau) \leftarrow \text{Gen}(pp)$. Suppose also that $|\psi\rangle$ is Live. This assumption means that the outcome p of applying $\text{ProjImp}(\mathcal{M}_D)$ to $|\psi\rangle$ is $1/2 + \epsilon$, where ϵ is an inverse polynomial. For every $i \in [\ell_m]$, from the SIM-MDD security of ELWMPRF, $D_{\tau,i}$ is indistinguishable from D if $m[i] = 0$ and is indistinguishable from D^{rev} if $m[i] = 1$. This means that $D_{\tau,i} \in \{\text{ProjImp}(\mathcal{M}_{D'}) \mid D' \stackrel{c}{\approx} D\}$ if $m[i] = 0$ and $D_{\tau,i} \in \{\text{ProjImp}(\mathcal{M}_{D'}) \mid D' \stackrel{c}{\approx} D^{\text{rev}}\}$ if $m[i] = 1$. Then, from the above key fact of the projective implementation, it holds that \tilde{p}_i is close to $1/2 + \epsilon > 1/2$ if $m[i] = 0$ and is close to $1/2 - \epsilon < 1/2$ if $m[i] = 1$. Therefore, we see that $\mathcal{E}xtract$ correctly extract m from $|\psi\rangle$. This means that WMPRF satisfies unremovability.

The above definition, construction, and security analysis are simplified and ignore many subtleties. The most significant point is that we use approximated projective implementations introduced by Zhandry [Zha20] instead of projective implementations in the actual construction since applying a projective implementation is an inefficient process. Moreover, though the outcomes of (approximate) projective implementations for indistinguishable distributions are close, in the actual analysis, we have to take into account that the outcomes gradually change every time we apply an (approximate) projective implementation. These issues can be solved by doing careful parameter settings.

Comparison with the work by Zhandry [Zha20]. Some readers familiar with Zhandry's work [Zha20] might think that our technique contradicts the lesson from Zhandry's work since it essentially says that once we find a large gap in success probabilities, the tested quantum pirate circuit might self-destruct. However, this is not the case. What Zhandry's work really showed is the following. Once a quantum pirate circuit itself detects that there is a large gap in success probabilities, it might self-destruct. Even if an extractor finds a large gap in success probabilities, if the tested quantum pirate circuit itself cannot detect the large gap, the pirate circuit cannot self-destruct. In Zhandry's work, whenever an extractor finds a large gap, the tested pirate circuit also detects the large gap. In our work, the tested pirate circuit cannot detect a large gap throughout the extraction process while an extractor can find it.

The reason why a pirate circuit cannot detect a large gap in our scheme even if an extractor can find it is as follows. Recall that in the above extraction process of our scheme based on an extraction-less watermarking PRF scheme, we apply $\text{ProjImp}(\mathcal{M}_{D_{\tau,i}})$ to the tested pirate circuit for every $i \in [\ell_m]$. Each $D_{\tau,i}$ outputs a tuple of the form (b, x, y) and is indistinguishable from D or D^{rev} depending on the embedded message. In the process, we apply $\text{ProjImp}(\mathcal{M}_{D_{\tau,i}})$ for every $i \in [\ell_m]$, and we get the success probability p if $D_{\tau,i}$ is indistinguishable from D and we get $1 - p$ if $D_{\tau,i}$ is indistinguishable from D^{rev} . The tested pirate circuit needs to know which of D or D^{rev} is indistinguishable from the distribution $D_{\tau,i}$

behind the projective implementation to know which of p or $1 - p$ is the result of an application of a projective implementation. However, this is impossible. The tested pirate circuit receives only (x, y) part of $D_{\tau,i}$'s output and not b part. (Recall that the task of the pirate circuit is to guess b from (x, y) .) The only difference between D and D^{rev} is that the first-bit b is flipped. Thus, if the b part is dropped, $D_{\tau,i}$ is, in fact, indistinguishable from both D and D^{rev} . As a result, the pirate program cannot know which of p or $1 - p$ is the result of an application of a projective implementation. In other words, the pirate circuit cannot detect a large gap in our extraction process.

Instantiating extraction-less watermarking PRF. In the rest of this overview, we will explain how to realize extraction-less watermarking PRF.

We consider the following two settings similar to the ordinary watermarking PRF. Recall that we consider the public marking setting by default.

Private-simulatable: In this setting, the extraction key xk fed into Sim is kept secret. We require that SIM-MDD security hold under the existence of the simulation oracle that is given a public tag τ' and an index $i' \in [\ell_m]$ and returns $\text{Sim}(xk, \tau', i')$. An extraction-less watermarking PRF scheme in this setting yields a watermarking PRF scheme against quantum adversaries in private-extractable setting where unremovability holds for adversaries who can access the extraction oracle.

Public-simulatable: In this setting, the extraction key xk is publicly available. An extraction-less watermarking PRF scheme in this setting yields a watermarking PRF scheme against quantum adversaries in the public-extractable setting.

We provide a construction in the first setting using private constrained PRF based on the hardness of the LWE assumption. Also, we provide a construction in the second setting based on IO and the hardness of the LWE assumption.

To give a high-level idea behind the above constructions, in this overview, we show how to construct a public-simulatable extraction-less watermarking PRF in the token-based setting [CHN⁺18]. In the token-based setting, we treat a marked circuit $\tilde{C} \leftarrow \text{Mark}(\text{pp}, \text{prfk}, m)$ as a tamper-proof hardware token that an adversary can only access in a black-box way.

Before showing the actual construction, we explain the high-level idea. Recall that SIM-MDD security requires that an adversary \mathcal{A} who is given $\tilde{C} \leftarrow \text{Mark}(\text{pp}, \text{prfk}, m)$ cannot distinguish $(\gamma^*, x^*, y^*) \leftarrow \text{Sim}(xk, \tau, i^*)$ from an output of D if $m[i^*] = 0$ and from that of D^{rev} if $m[i^*] = 1$. This is the same as requiring that \mathcal{A} cannot distinguish $(\gamma^*, x^*, y^*) \leftarrow \text{Sim}(xk, \tau, i^*)$ from that of the following distribution D_{real,i^*} . We can check that D_{real,i^*} is identical with D if $m[i^*] = 0$ and with D^{rev} if $m[i^*] = 1$.

D_{real,i^*} : Generate $\gamma \leftarrow \{0, 1\}$ and $x \leftarrow \text{Dom}$. Then, if $\gamma = m[i^*]$, generate $y \leftarrow \text{Ran}$, and otherwise, compute $y \leftarrow \text{Eval}(\text{prfk}, x)$. Output (γ, x, y) .

Essentially, the only attack that \mathcal{A} can perform is to feed x^* contained in the given tuple (γ^*, x^*, y^*) to \tilde{C} and compares the result $\tilde{C}(x^*)$ with y^* , if we ensure that γ^*, x^* are pseudorandom. In order to make the construction immune to this attack, letting $\tilde{C} \leftarrow \text{Mark}(\text{pp}, \text{prfk}, m)$ and $(\gamma^*, x^*, y^*) \leftarrow \text{Sim}(xk, \tau, i^*)$, we have to design Sim and \tilde{C} so that

- If $\gamma = m[i^*]$, $\tilde{C}(x^*)$ outputs a value different from y^* .
- If $\gamma \neq m[i^*]$, $\tilde{C}(x^*)$ outputs y^* .

We achieve these conditions as follows. First, we set (γ^*, x^*, y^*) output by $\text{Sim}(xk, \tau, i^*)$ so that γ^* and y^* is random values and x^* is an encryption of $y^* \| i^* \| \gamma^*$ by a public-key encryption scheme with pseudorandom ciphertext property, where the encryption key pk is included in τ . Then, we set \tilde{C} as a token such that it has the message m and the decryption key sk corresponding to pk hardwired, and it outputs y^* if the input is decryptable and $\gamma^* \neq m[i^*]$ holds for the decrypted $y^* \| i^* \| \gamma^*$, and otherwise behaves as $\text{Eval}(\text{prfk}, \cdot)$. The actual construction is as follows.

Let PRF be a PRF family consisting of functions $\{F_{\text{prfk}}(\cdot) : \{0,1\}^n \rightarrow \{0,1\}^\lambda | \text{prfk}\}$, where λ is the security parameter and n is sufficiently large. Let PKE = (KG, E, D) be a CCA secure public-key encryption scheme satisfying pseudorandom ciphertext property. Using these ingredients, We construct an extraction-less watermarking PRF scheme ELWMPRF = (Setup, Gen, Eval, Mark, Sim) as follows.

Setup(1^λ): In this construction, $\text{pp} := \perp$ and $\text{xk} := \perp$.

Gen(pp): It generates a fresh PRF key prfk of PRF and a key pair $(\text{pk}, \text{sk}) \leftarrow \text{KG}(1^\lambda)$. The PRF key is (prfk, sk) and the corresponding public tag is pk .

Eval($(\text{prfk}, \text{sk}), x$): It simply outputs $F_{\text{prfk}}(x)$.

Mark(pp, $(\text{prfk}, \text{sk}), m$): It generates the following taken $\tilde{C}[\text{prfk}, \text{sk}, m]$.

Hard-Coded Constants: $\text{prfk}, \text{sk}, m$.

Input: $x \in \{0,1\}^n$.

1. Try to decrypt $y \| i \| \gamma \leftarrow D(\text{sk}, x)$ with $y \in \{0,1\}^\lambda$, $i \in [\ell_m]$, and $\gamma \in \{0,1\}$.
2. If decryption succeeds, output y if $\gamma \neq m[i]$ and $F_{\text{prfk}}(x)$ otherwise.
3. Otherwise, output $F_{\text{prfk}}(x)$.

Sim(xk, τ, i): It first generates $\gamma \leftarrow \{0,1\}$ and $y \leftarrow \{0,1\}^\lambda$. Then, it parses $\tau := \text{pk}$ and generates $x \leftarrow E(\text{pk}, y \| i \| \gamma)$. Finally, it outputs (γ, x, y) .

We check that ELWMPRF satisfies SIM-MDD security. For simplicity, we fix the message $m \in [\ell_m]$ embedded into the challenge PRF key. Then, for any adversary \mathcal{A} and $i^* \in [\ell_m]$, SIM-MDD security requires that given $\tilde{C}[\text{prfk}, \text{sk}, m] \leftarrow \text{Mark}(\text{mk}, \text{prfk}, m)$ and $\tau = \text{pk}$, \mathcal{A} cannot distinguish $(\gamma^*, x^* = E(\text{pk}, y^* \| i^* \| \gamma^*), y^*) \leftarrow \text{Sim}(\text{xk}, \tau, i^*)$ from an output of D if $m[i^*] = 0$ and is indistinguishable from D^{rev} if $m[i^*] = 1$.

We consider the case of $m[i^*] = 0$. We can finish the security analysis by considering the following sequence of mutually indistinguishable hybrid games, where \mathcal{A} is given $(\gamma^*, x^* = E(\text{pk}, y^* \| i^* \| \gamma^*), y^*) \leftarrow \text{Sim}(\text{xk}, \tau, i^*)$ in the first game, and on the other hand, is given $(\gamma^*, x^*, y^*) \leftarrow D$ in the last game. We first change the game so that x^* is generated as a uniformly random value instead of $x^* \leftarrow E(\text{pk}, y^* \| i^* \| \gamma^*)$ by using the pseudorandom ciphertext property under CCA of PKE. This is possible since the CCA oracle can simulate access to the marked token $\tilde{C}[\text{prfk}, \text{sk}, m]$ by \mathcal{A} . Then, we further change the security game so that if $\gamma^* = 1$, y^* is generated as $F_{\text{prfk}}(x^*)$ instead of a uniformly random value by using the pseudorandomness of PRF. Note that if $\gamma^* = 0$, y^* remains uniformly at random. We see that if $\gamma^* = 1$, the token $\tilde{C}[\text{prfk}, \text{sk}, m]$ never evaluate $F_{\text{prfk}}(x^*)$ since $m[i^*] \neq \gamma^*$. Thus, this change is possible. We see that now the distribution of (γ^*, x^*, y^*) is exactly the same as that output by D . Similarly, in the case of $m[i^*] = 1$, we can show that an output of $\text{Sim}(\text{xk}, \tau, i^*)$ is indistinguishable from that output by D^{rev} . The only difference is that in the final step, we change the security game so that y^* is generated as $F_{\text{prfk}}(x^*)$ if $\gamma^* = 0$.

In the actual public-simulatable construction, we implement this idea using iO and puncturable encryption [CHN⁺18] instead of token and CCA secure public-key encryption. Also, in the actual secret-simulatable construction, we basically follow the same idea using private constrained PRF and secret-key encryption.

1.4 More on Related Work

Watermarking against classical adversaries. Cohen et al. [CHN⁺18] present a publicly extractable watermarking PRF from IO and injective OWFs. It is unremovable against adversaries who can access the mark oracle only before a target marked circuit is given. The mark oracle returns a marked circuit for a queried *arbitrary* polynomial-size circuit. Suppose we additionally assume the hardness of the

decisional Diffie-Hellman or LWE problem. In that case, their watermarking PRF is unremovable against adversaries that can access the mark oracle before and after a target marked circuit is given. However, adversaries can query a valid PRF key to the mark oracle in that case. They also present definitions and constructions of watermarking for public-key cryptographic primitives.

Boneh, Lewi, and Wu [BLW17] present a privately extractable watermarking PRF from privately programmable PRFs, which are variants of private constrained PRFs [BLW17, CC17]. It is unremovable in the presence of the mark oracle. However, it is not secure in the presence of the extraction oracle and does not support public marking. They instantiate a privately programmable PRF with IO and OWFs, but later, Peikert and Shiehian [PS18] instantiate it with the LWE assumption.

Kim and Wu [KW21] (KW17), Quach, Wichs, and Zirdelis [QWZ18] (QWZ), and Kim and Wu [KW19] (KW19) present privately extractable watermarking PRFs from the LWE assumption. They are secure in the presence of the mark oracle. KW17 construction is not secure in the presence of the extraction oracle and does not support public marking. QWZ construction is unremovable in the presence of the extraction oracle and supports public marking. However, it does not have pseudorandomness against an authority that generates a marking and extraction key. KW19 construction is unremovable in the presence of the extraction oracle and has some restricted pseudorandomness against an authority (see the reference [KW19] for the detail). However, it does not support public marking.⁶

Yang et al. [YAL⁺19] present a collusion-resistant watermarking PRF from IO and the LWE assumption. Collusion-resistant watermarking means unremovability holds even if adversaries receive multiple marked circuits with different embedded messages generated from one target circuit.

Goyal, Kim, Manohar, Waters, and Wu [GKM⁺19] improve the definitions of watermarking for public-key cryptographic primitives and present constructions. In particular, they introduce collusion-resistant watermarking and more realistic attack strategies for public-key cryptographic primitives. Nishimaki [Nis20] present a general method for equipping many existing public-key cryptographic schemes with the watermarking functionality.

Goyal, Kim, Waters, and Wu [GKWW21] introduce the notion of traceable PRFs, where we can identify a user that creates a pirate copy of her/his authenticated PRF. The difference between traceable PRF and (collusion-resistant) watermarking PRF is that there is only one target original PRF and multiple authenticated copies of it with different identities in traceable PRF. In (collusion-resistant) watermarking PRF, we consider many different PRF keys. In addition, Goyal et al. introduce a refined attack model. Adversaries in previous watermarking PRF definitions output a pirate PRF circuit that correctly computes the original PRF values for $1/2 + \epsilon$ fraction of inputs. However, adversaries in traceable PRFs output a pirate circuit that *distinguishes* whether an input pair consists of a random input and a real PRF value or a random input and output value. This definition captures wide range of attacks. For example, it captures adversaries who create a pirate PRF circuit that can compute the first quarter bits of the original PRF output. Such an attack is not considered in previous watermarking PRFs. We adopt the refined attack model in our definitions.

Learning information from adversarial entities in the quantum setting. Zhandry [Zha20] introduces the definition of secure traitor tracing against quantum adversaries. In traitor tracing, each legitimate user receives a secret key that can decrypt broadcasted ciphertexts and where identity information is embedded. An adversary outputs a pirate decoder that can distinguish whether an input is a ciphertext of m_0 or m_1 where m_0 and m_1 are adversarially chosen plaintexts. A tracing algorithm must identify a malicious user's identity such that its secret decryption key is embedded in the pirate decoder. Thus, we need to extract information from adversarially generated objects. Such a situation also appears in security proofs of interactive proof systems [Wat09, Unr12, ARU14, CMSZ21] (but not in real cryptographic algorithms) since we rewind a verifier.

⁶Their construction supports public marking in the random oracle model.

Zhandry presents how to estimate the success decryption probability of a *quantum* pirate decoder without destroying the decoding (distinguishing) capability. He achieves a quantum tracing algorithm that extracts a malicious user identity by combining the probability approximation technique above with PLBE [BSW06]. However, his technique is limited to the setting where user identity spaces are only polynomially large while there are several traitor tracing schemes with exponentially large identity spaces [NWZ16, GKW19]. As observed in previous works [GKM⁺19, Nis20, GKWW21], traitor tracing and watermarking have similarities since an adversary outputs a pirate circuit in the watermarking setting and an extraction algorithm tries to retrieve information from it. However, a notable difference is that we must consider exponentially large message spaces by default in the (message-embedding) watermarking setting.

Application of (classical) watermarking. As we explained above, Aaronson et al. [ALL⁺21], and Kitagawa et al. [KNY21] achieve secure software leasing schemes by using watermarking. A leased software consists of a quantum state and watermarked circuit. Although they use watermarking schemes in the quantum setting, it is sufficient for their purpose to use secure watermarking against adversaries that output a *classical* pirate circuit. This is because a returned software is verified by a checking algorithm and must have a specific format in secure software leasing.⁷ That is, a returned software is rejected if it does not have a classical circuit part that can be tested by an extraction algorithm of the building block watermarking.

2 Preliminaries

Notations and conventions. In this paper, standard math or sans serif font stands for classical algorithms (e.g., C or Gen) and classical variables (e.g., x or pk). Calligraphic font stands for quantum algorithms (e.g., \mathcal{G}) and calligraphic font and/or the bracket notation for (mixed) quantum states (e.g., q or $|\psi\rangle$). For strings x and y , $x\|y$ denotes the concatenation of x and y . Let $[\ell]$ denote the set of integers $\{1, \dots, \ell\}$, λ denote a security parameter, and $y := z$ denote that y is set, defined, or substituted by z .

In this paper, for a finite set X and a distribution D , $x \leftarrow X$ denotes selecting an element from X uniformly at random, $x \leftarrow D$ denotes sampling an element x according to D . Let $y \leftarrow A(x)$ and $y \leftarrow \mathcal{A}(\chi)$ denote assigning to y the output of a probabilistic or deterministic algorithm A and a quantum algorithm \mathcal{A} on an input x and χ , respectively. When we explicitly show that A uses randomness r , we write $y \leftarrow A(x; r)$. PPT and QPT algorithms stand for probabilistic polynomial time algorithms and polynomial time quantum algorithms, respectively. Let negl denote a negligible function.

2.1 Quantum Information

Let \mathcal{H} be a finite-dimensional complex Hilbert space. A (pure) quantum state is a vector $|\psi\rangle \in \mathcal{H}$. Let $\mathcal{S}(\mathcal{H})$ be the space of Hermitian operators on \mathcal{H} . A density matrix is a Hermitian operator $X \in \mathcal{S}(\mathcal{H})$ with $\text{Tr}(X) = 1$, which is a probabilistic mixture of pure states. A quantum state over $\mathcal{H} = \mathbb{C}^2$ is called qubit, which can be represented by the linear combination of the standard basis $\{|0\rangle, |1\rangle\}$. More generally, a quantum system over $(\mathbb{C}^2)^{\otimes n}$ is called an n -qubit quantum system for $n \in \mathbb{N} \setminus \{0\}$.

A Hilbert space is divided into registers $\mathcal{H} = \mathcal{H}^{R_1} \otimes \mathcal{H}^{R_2} \otimes \dots \otimes \mathcal{H}^{R_n}$. We sometimes write \mathcal{X}^{R_i} to emphasize that the operator \mathcal{X} acts on register \mathcal{H}^{R_i} .⁸ When we apply \mathcal{X}^{R_1} to registers \mathcal{H}^{R_1} and \mathcal{H}^{R_2} , \mathcal{X}^{R_1} is identified with $\mathcal{X}^{R_1} \otimes I^{R_2}$.

A unitary operation is represented by a complex matrix \mathbf{U} such that $\mathbf{U}\mathbf{U}^\dagger = \mathbf{I}$. The operation \mathbf{U} transforms $|\psi\rangle$ and X into $\mathbf{U}|\psi\rangle$ and $\mathbf{U}X\mathbf{U}^\dagger$, respectively. A projector \mathbf{P} is a Hermitian operator ($\mathbf{P}^\dagger = \mathbf{P}$) such that $\mathbf{P}^2 = \mathbf{P}$.

⁷A valid software must run on a legitimate platform. For example, a video game title of Xbox must run on Xbox.

⁸The superscript parts are gray colored.

For a quantum state X over two registers \mathcal{H}^{R_1} and \mathcal{H}^{R_2} , we denote the state in \mathcal{H}^{R_1} as $X[R_1]$, where $X[R_1] = \text{Tr}_2[X]$ is a partial trace of X (trace out R_2).

Given a function $F : X \rightarrow Y$, a quantum-accessible oracle O of F is modeled by a unitary transformation U_F operating on two registers \mathcal{H}^{in} and \mathcal{H}^{out} , in which $|x\rangle|y\rangle$ is mapped to $|x\rangle|y \oplus F(x)\rangle$, where \oplus denotes XOR group operation on Y . We write $\mathcal{A}^{(O)}$ to denote that the algorithm \mathcal{A} 's oracle O is a quantum-accessible oracle.

Definition 2.1 (Quantum Program with Classical Inputs and Outputs [ALL⁺21]). A quantum program with classical inputs is a pair of quantum state q and unitaries $\{U_x\}_{x \in [N]}$ where $[N]$ is the domain, such that the state of the program evaluated on input x is equal to $U_x q U_x^\dagger$. We measure the first register of $U_x q U_x^\dagger$ to obtain an output. We say that $\{U_x\}_{x \in [N]}$ has a compact classical description \mathbf{U} when applying U_x can be efficiently computed given \mathbf{U} and x .

Definition 2.2 (Positive Operator-Valued Measure). Let \mathcal{I} be a finite index set. A positive operator valued measure (POVM) \mathcal{M} is a collection $\{M_i\}_{i \in \mathcal{I}}$ of Hermitian positive semi-definite matrices M_i such that $\sum_{i \in \mathcal{I}} M_i = \mathbf{I}$. When we apply POVM \mathcal{M} to a quantum state X , the measurement outcome is i with probability $p_i = \text{Tr}(X M_i)$. We denote by $\mathcal{M}(|\psi\rangle)$ the distribution obtained by applying \mathcal{M} to $|\psi\rangle$.

Definition 2.3 (Quantum Measurement). A quantum measurement \mathcal{E} is a collection $\{E_i\}_{i \in \mathcal{I}}$ of matrices E_i such that $\sum_{i \in \mathcal{I}} E_i^\dagger E_i = \mathbf{I}$. When we apply \mathcal{E} to a quantum state X , the measurement outcome is i with probability $p_i = \text{Tr}(X E_i^\dagger E_i)$. Conditioned on the outcome being i , the post-measurement state is $E_i X E_i^\dagger / p_i$.

We can construct a POVM \mathcal{M} from any quantum measurement \mathcal{E} by setting $M_i := E_i^\dagger E_i$. We say that \mathcal{E} is an implementation of \mathcal{M} . The implementation of a POVM may not be unique.

Definition 2.4 (Projective Measurement/POVM). A quantum measurement $\mathcal{E} = \{E_i\}_{i \in \mathcal{I}}$ is projective if for all $i \in \mathcal{I}$, E_i is a projector. This implies that $E_i E_j = \mathbf{0}$ for distinct $i, j \in \mathcal{I}$. In particular, two-outcome projective measurement is called a binary projective measurement, and is written as $\mathcal{E} = (\mathbf{P}, \mathbf{I} - \mathbf{P})$, where \mathbf{P} is associated with the outcome 1, and $\mathbf{I} - \mathbf{P}$ with the outcome 0. Similarly, a POVM \mathcal{M} is projective if for all $i \in \mathcal{I}$, M_i is a projector. This also implies that $M_i M_j = \mathbf{0}$ for distinct $i, j \in \mathcal{I}$.

Definition 2.5 (Controlled Projection). Let $\mathcal{P} = \{M_i\}_{i \in \mathcal{I}}$ be a collection of projective measurement over a Hilbert space \mathcal{H} , where $M_i = (\Pi_i, \mathbf{I} - \Pi_i)$ for $i \in \mathcal{I}$. Let D be a distribution whose randomness space is \mathcal{R} . The controlled projection $\text{CProj}_{\mathcal{P}, D} = (\text{CProj}_{\mathcal{P}, D}^1, \text{CProj}_{\mathcal{P}, D}^0)$ is defined as follows.⁹

$$\text{CProj}_{\mathcal{P}, D}^1 := \sum_{r \in \mathcal{R}} |r\rangle \langle r| \otimes \Pi_{D(r)}, \quad \text{CProj}_{\mathcal{P}, D}^0 := \sum_{r \in \mathcal{R}} |r\rangle \langle r| \otimes (\mathbf{I} - \Pi_{D(r)})$$

2.2 Measurement Implementation

Definition 2.6 (Projective Implementation). Let:

- $\mathcal{P} = (\mathbf{P}, \mathbf{I} - \mathbf{P})$ be a binary outcome POVM
- D be a finite set of distributions over outcomes $\{0, 1\}$
- $\mathcal{E} = \{E_D\}_{D \in \mathcal{D}}$ be a projective measurement with index set \mathcal{D} .

We define the following measurement.

1. Measure under the projective measurement \mathcal{E} and obtain a distribution D over $\{0, 1\}$.
2. Output a bit sampled from the distribution D .

⁹We use superscript b to denote that it is associated with the outcome b here.

We say this measurement is a projective implementation of \mathcal{P} , denoted by $\text{ProjImp}(\mathcal{P})$ if it is equivalent to \mathcal{P} .

Theorem 2.7 ([Zha20, Lemma 1]). Any binary outcome POVM $\mathcal{P} = (\mathbf{P}, \mathbf{I} - \mathbf{P})$ has a projective implementation $\text{ProjImp}(\mathcal{P})$.

Definition 2.8 (Shift Distance). For two distributions D_0, D_1 , the shift distance with parameter ϵ , denoted by $\Delta_{\text{Shift}}^\epsilon(D_0, D_1)$, is the smallest quantity δ such that for all $x \in \mathbb{R}$:

$$\begin{aligned} \Pr[D_0 \leq x] &\leq \Pr[D_1 \leq x + \epsilon] + \delta, & \Pr[D_0 \geq x] &\leq \Pr[D_1 \geq x - \epsilon] + \delta, \\ \Pr[D_1 \leq x] &\leq \Pr[D_0 \leq x + \epsilon] + \delta, & \Pr[D_1 \geq x] &\leq \Pr[D_0 \geq x - \epsilon] + \delta. \end{aligned}$$

For two real-valued measurements \mathcal{M} and \mathcal{N} over the same quantum system, the shift distance between \mathcal{M} and \mathcal{N} with parameter ϵ is

$$\Delta_{\text{Shift}}^\epsilon(\mathcal{M}, \mathcal{N}) := \sup_{|\psi\rangle} \Delta_{\text{Shift}}^\epsilon(\mathcal{M}(|\psi\rangle), \mathcal{N}(|\psi\rangle)).$$

Definition 2.9 ((ϵ, δ) -Almost Projective [Zha20]). A real-valued quantum measurement $\mathcal{M} = \{\mathbf{M}_i\}_{i \in \mathcal{I}}$ is (ϵ, δ) -almost projective if the following holds. For any quantum state $|\psi\rangle$, we apply \mathcal{M} twice in a row to $|\psi\rangle$ and obtain measurement outcomes x and y , respectively. Then, $\Pr[|x - y| \leq \epsilon] \geq 1 - \delta$.

Theorem 2.10 ([Zha20, Theorem 2]). Let D be any probability distribution and \mathcal{P} be a collection of projective measurements. For any $0 < \epsilon, \delta < 1$, there exists an algorithm of measurement $\mathcal{API}_{\mathcal{P}, D}^{\epsilon, \delta}$ that satisfies the following.

- $\Delta_{\text{Shift}}^\epsilon(\mathcal{API}_{\mathcal{P}, D}^{\epsilon, \delta}, \text{ProjImp}(\mathcal{P}_D)) \leq \delta$.
- $\mathcal{API}_{\mathcal{P}, D}^{\epsilon, \delta}$ is (ϵ, δ) -almost projective.
- The expected running time of $\mathcal{API}_{\mathcal{P}, D}^{\epsilon, \delta}$ is $T_{\mathcal{P}, D} \cdot \text{poly}(1/\epsilon, \log(1/\delta))$ where $T_{\mathcal{P}, D}$ is the combined running time of D , the procedure mapping $i \rightarrow (\mathbf{P}_i, \mathbf{I} - \mathbf{P}_i)$, and the running time of measurement $(\mathbf{P}_i, \mathbf{I} - \mathbf{P}_i)$.

Theorem 2.11 ([Zha20, Corollary 1]). Let q be an efficiently constructible, potentially mixed state, and D_0, D_1 efficiently samplable distributions. If D_0 and D_1 are computationally indistinguishable, for any inverse polynomial ϵ and any function δ , we have $\Delta_{\text{Shift}}^{3\epsilon}(\mathcal{API}_{\mathcal{P}, D_0}^{\epsilon, \delta}, \mathcal{API}_{\mathcal{P}, D_1}^{\epsilon, \delta}) \leq 2\delta + \text{negl}(\lambda)$.

Note that the indistinguishability of D_0 and D_1 needs to hold against distinguishers who can construct q in the theorem above. However, this fact is not explicitly stated in [Zha20]. We need to care about this condition if we need secret information to construct q , and the secret information is also needed to sample an output from D_0 or D_1 . We handle such a situation when analyzing the unremovability of our privately extractable watermarking PRF. In that situation, we need a secret extraction key to construct q and sample an output from D_0 and D_1 .

We also define the notion of the reverse almost projective property of API.

Definition 2.12 ((ϵ, δ) -Reverse Almost Projective). Let $\mathcal{P} = \{(\Pi_i, \mathbf{I} - \Pi_i)\}_i$ be a collection of binary outcome projective measurements. Let D be a distribution. We also let $\mathcal{P}^{\text{rev}} = \{(\mathbf{I} - \Pi_i, \Pi_i)\}_i$. We say \mathcal{API} is (ϵ, δ) -reverse almost projective if the following holds. For any quantum state $|\psi\rangle$, we apply $\mathcal{API}_{\mathcal{P}, D}^{\epsilon, \delta}$ and $\mathcal{API}_{\mathcal{P}^{\text{rev}}, D}^{\epsilon, \delta}$ in a row to $|\psi\rangle$ and obtain measurement outcomes x and y , respectively. Then, $\Pr[|(1 - x) - y| \leq \epsilon] \geq 1 - \delta$.

Algorithm $\mathcal{API}_{\mathcal{P},D}^{\epsilon,\delta}$

Parameter: Collection of projective measurement \mathcal{P} , distribution D , real values ϵ, δ .

Input: A quantum state $|\psi\rangle$.

1. Initialize a state $|\mathbb{1}_{\mathcal{R}}\rangle |\psi\rangle$.
2. Initialize a classical list $L := (1)$.
3. Repeat the following $T := \left\lceil \frac{\ln 4/\delta}{\epsilon^2} \right\rceil$.
 - (a) Apply $\text{CProj}_{\mathcal{P},D}$ to register $\mathcal{H}^{\text{R}} \otimes \mathcal{H}^{\text{H}}$. Let b_{2i-1} be the measurement outcome and set $L := (L, b_{2i-1})$.
 - (b) Apply $\text{IsU}_{\mathcal{R}}$ to register \mathcal{H}^{R} . Let b_{2i} be the measurement outcome and set $L := (L, b_{2i})$.
4. Let t be the number of index i such that $b_{i-1} = b_i$ in the list $L = (0, b_1, \dots, b_{2T})$, and $\tilde{p} := t/2T$.
5. If $b_{2T} = 0$, repeat the loop again until $b_{2i} = 1$.
6. Discard \mathcal{H}^{R} register, and output \tilde{p} .

Figure 1: The description of \mathcal{API} .

We show that the measurement algorithm $\mathcal{API}_{\mathcal{P},D}^{\epsilon,\delta}$ in Theorem 2.10 also satisfies Definition 2.12. First, we describe the detail of $\mathcal{API}_{\mathcal{P},D}^{\epsilon,\delta}$ in Figure 1. \mathcal{API} uses an ancilla register \mathcal{H}^{R} besides the original Hilbert space \mathcal{H}^{H} . Let \mathcal{R} be the randomness space of distribution D . We define $\text{IsU}_{\mathcal{R}} := (|\mathbb{1}_{\mathcal{R}}\rangle \langle \mathbb{1}_{\mathcal{R}}|, \mathbf{I} - |\mathbb{1}_{\mathcal{R}}\rangle \langle \mathbb{1}_{\mathcal{R}}|)$ where

$$|\mathbb{1}_{\mathcal{R}}\rangle := \frac{1}{\sqrt{|\mathcal{R}|}} \sum_{r \in \mathcal{R}} |r\rangle.$$

We use the following lemma to analyze $\mathcal{API}_{\mathcal{P},D}^{\epsilon,\delta}$.

Lemma 2.13 ([Jor75]). *For any two Hermitian projectors Π_v and Π_w on a Hilbert space \mathcal{H} , there exists an orthogonal decomposition of \mathcal{H} into one-dimensional and two-dimensional subspaces (the Jordan subspaces) that are invariant under both Π_v and Π_w . Moreover:*

- in each one-dimensional space, Π_v and Π_w act as identity or rank-zero projectors; and
- in each two-dimensional subspace S_j , Π_v and Π_w are rank-one projectors: there exists $|v_j\rangle, |w_j\rangle \in S_j$ such that Π_v projects onto $|v_j\rangle$ and Π_w projects onto $|w_j\rangle$.

For each two-dimensional subspace S_j , we call $p_j := |\langle v_j | w_j \rangle|^2$ the eigenvalue of the j -th subspace. It is easy to see that $|v_j\rangle$ is an eigenvector of the Hermitian matrix $\Pi_v \Pi_w \Pi_v$ with eigenvalue p_j .

As previous works observed [MW05, Reg, CMSZ21], we obtain the following by Lemma 2.13. There exists orthogonal vectors $|v_j\rangle, |v_j^\perp\rangle$ that span S_j , such that $\Pi_v |v_j\rangle = |v_j\rangle$ and $\Pi_v |v_j^\perp\rangle = 0$. Similarly, $\Pi_w |w_j\rangle = |w_j\rangle$ and $\Pi_w |w_j^\perp\rangle = 0$. By setting appropriate phases, we have

$$\begin{aligned} |w_j\rangle &= \sqrt{p_j} |v_j\rangle + \sqrt{1-p_j} |v_j^\perp\rangle, & |w_j^\perp\rangle &= \sqrt{1-p_j} |v_j\rangle - \sqrt{p_j} |v_j^\perp\rangle, \\ |v_j\rangle &= \sqrt{p_j} |w_j\rangle + \sqrt{1-p_j} |w_j^\perp\rangle, & |v_j^\perp\rangle &= \sqrt{1-p_j} |w_j\rangle - \sqrt{p_j} |w_j^\perp\rangle, \end{aligned}$$

where $|v_j^\perp\rangle, |w_j^\perp\rangle \in S_j$ such that $\langle v_j | v_j^\perp \rangle = \langle w_j | w_j^\perp \rangle = 0$. We also have

$$\Pi_v |w_j\rangle = \sqrt{p_j} |v_j\rangle, \quad \Pi_w |v_j\rangle = \sqrt{p_j} |w_j\rangle.$$




Figure 2: Solid lines denote that the measurement outcome is 1. Dashed line denote that the measurement outcome is 0. Double lines denote we apply $\text{CProj}_{\mathcal{P},D} = (\text{CProj}_{\mathcal{P},D}^1, \text{CProj}_{\mathcal{P},D}^0)$. Single lines denote we apply $\text{IsU}_{\mathcal{R}} = (|\mathbb{1}_{\mathcal{R}}\rangle\langle\mathbb{1}_{\mathcal{R}}|, \mathbf{I} - |\mathbb{1}_{\mathcal{R}}\rangle\langle\mathbb{1}_{\mathcal{R}}|)$.




Figure 3: Solid lines denote that the measurement outcome is 1. Dashed line denote that the measurement outcome is 0. Double lines denote we apply $\text{CProj}_{\mathcal{P},D}^{\text{rev}} = (\text{CProj}_{\mathcal{P},D}^{1,\text{rev}}, \text{CProj}_{\mathcal{P},D}^{0,\text{rev}}) = (\text{CProj}_{\mathcal{P},D}^0, \text{CProj}_{\mathcal{P},D}^1)$. Single lines denote we apply $\text{IsU}_{\mathcal{R}} = (|\mathbb{1}_{\mathcal{R}}\rangle\langle\mathbb{1}_{\mathcal{R}}|, \mathbf{I} - |\mathbb{1}_{\mathcal{R}}\rangle\langle\mathbb{1}_{\mathcal{R}}|)$.

Theorem 2.14. $\mathcal{A}PI_{\mathcal{P},D}^{\epsilon,\delta}$ in Figure 1 is (ϵ, δ) -reverse almost projective.

Proof of Theorem 2.14. To analyze $\mathcal{A}PI_{\mathcal{P},D}^{\epsilon,\delta}$, we set $\Pi_w := \text{CProj}_{\mathcal{P},D}^1$, $\Pi_v := |\mathbb{1}_{\mathcal{R}}\rangle\langle\mathbb{1}_{\mathcal{R}}| \otimes \mathbf{I}$, and apply Lemma 2.13. Then, we have the following relationships:

$$\begin{aligned}
\text{CProj}_{\mathcal{P},D}^1 |v_j\rangle &= \sqrt{p_j} |w_j\rangle, & (|\mathbb{1}_{\mathcal{R}}\rangle\langle\mathbb{1}_{\mathcal{R}}|) |w_j\rangle &= \sqrt{p_j} |v_j\rangle \\
\text{CProj}_{\mathcal{P},D}^0 |v_j\rangle &= \sqrt{1-p_j} |w_j^\perp\rangle, & (\mathbf{I} - |\mathbb{1}_{\mathcal{R}}\rangle\langle\mathbb{1}_{\mathcal{R}}|) |w_j\rangle &= \sqrt{1-p_j} |v_j^\perp\rangle \\
\text{CProj}_{\mathcal{P},D}^1 |v_j^\perp\rangle &= \sqrt{1-p_j} |w_j\rangle, & (|\mathbb{1}_{\mathcal{R}}\rangle\langle\mathbb{1}_{\mathcal{R}}|) |w_j^\perp\rangle &= \sqrt{1-p_j} |v_j\rangle \\
\text{CProj}_{\mathcal{P},D}^0 |v_j^\perp\rangle &= -\sqrt{p_j} |w_j^\perp\rangle, & (\mathbf{I} - |\mathbb{1}_{\mathcal{R}}\rangle\langle\mathbb{1}_{\mathcal{R}}|) |w_j^\perp\rangle &= -\sqrt{p_j} |v_j^\perp\rangle
\end{aligned}$$

where $|w_j\rangle$ and $|v_j\rangle$ are decompositions of Π_w and Π_v , and $p_j = |\langle v_j | w_j \rangle|^2$.

Suppose we apply $\mathcal{A}PI_{\mathcal{P},D}^{\epsilon,\delta}$ to a $|\psi\rangle$ on \mathcal{H}^H . We can write the initial state in Figure 1 as $|\mathbb{1}_{\mathcal{R}}\rangle|\psi\rangle = \sum_j \alpha_j |v_j\rangle + \alpha_j^\perp |v_j^\perp\rangle$ since $\{|v_j\rangle, |v_j^\perp\rangle\}_j$ is a basis. We also have that $|\mathbb{1}_{\mathcal{R}}\rangle\langle\mathbb{1}_{\mathcal{R}}| \otimes \mathbf{I} = \sum_j |v_j\rangle\langle v_j|$ since Π_v projects onto $|v_j\rangle$ in each decomposed subspace S_j . It is easy to see that $(|\mathbb{1}_{\mathcal{R}}\rangle\langle\mathbb{1}_{\mathcal{R}}| \otimes \mathbf{I}) |\mathbb{1}_{\mathcal{R}}\rangle|\psi\rangle = |\mathbb{1}_{\mathcal{R}}\rangle|\psi\rangle$. Thus, for all j , $\alpha_j^\perp = 0$. Therefore, for any $|\psi\rangle$, we can write $|\mathbb{1}_{\mathcal{R}}\rangle|\psi\rangle = \sum_j \alpha_j |v_j\rangle$. As we see in Figure 1, when we run $\mathcal{A}PI_{\mathcal{P},D}^{\epsilon,\delta}(|\psi\rangle)$, the initial state is $|\mathbb{1}_{\mathcal{R}}\rangle|\psi\rangle = \sum_j \alpha_j |v_j\rangle$ and we apply $\text{CProj}_{\mathcal{P},D}$ and $\text{IsU}_{\mathcal{R}}$ alternately. Therefore, the quantum state $|v_j\rangle$ in each decomposed subspace S_j changes as in Figure 2 when we run $\mathcal{A}PI_{\mathcal{P},D}^{\epsilon,\delta}(|\psi\rangle)$.

Next, suppose we apply $\mathcal{A}PI_{\mathcal{P}^{\text{rev}},D}^{\epsilon,\delta}$ to the quantum state immediately after applying $\mathcal{A}PI_{\mathcal{P},D}^{\epsilon,\delta}$ to $|\psi\rangle$. $\mathcal{A}PI_{\mathcal{P},D}^{\epsilon,\delta}$ ensures that the final measurement is $\text{IsU}_{\mathcal{R}}$ and its result is 1. This means that the state going into the main loop of $\mathcal{A}PI_{\mathcal{P}^{\text{rev}},D}^{\epsilon,\delta}$ (the third item in Figure 1) is identical to the state before \mathcal{H}^R is discarded

at the application of $\mathcal{API}_{\mathcal{P},D}^{\epsilon,\delta}$ to $|\psi\rangle$. By the definition of \mathcal{P}^{rev} , we have $\text{CProj}_{\mathcal{P},D}^b = \text{CProj}_{\mathcal{P}^{\text{rev}},D}^{1-b}$ for $b \in \{0,1\}$. Thus, the quantum state $|v_j\rangle$ in each decomposed subspace S_j changes as in Figure 3 when we apply $\mathcal{API}_{\mathcal{P}^{\text{rev}},D}^{\epsilon,\delta}$.

From the above discussions, we can view a successive execution of $\mathcal{API}_{\mathcal{P},D}^{\epsilon,\delta}$ and $\mathcal{API}_{\mathcal{P}^{\text{rev}},D}^{\epsilon,\delta}$ to $|\psi\rangle$ as the following single experiment.

- Sample p_j from $\{p_j\}_j$ with the probability α_j^2 .
- Flip $2T$ biased random coins whose probability of outputting 1 is p_j .
- Flip an even number of additional random coins until 0 is found.
- Flip $2T$ biased random coins whose probability of outputting 1 is $1 - p_j$.
- Let K be the overall list of coin flips.

Let \tilde{p}_x and \tilde{p}_y be the outcome of $\mathcal{API}_{\mathcal{P},D}^{\epsilon,\delta}$ and $\mathcal{API}_{\mathcal{P}^{\text{rev}},D}^{\epsilon,\delta}$, respectively. \tilde{p}_x is the fraction of 1's in the first $2T$ bits of K . Also, \tilde{p}_y is the fraction of 1's in the last $2T$ bits of K . Then, we have

$$\Pr[|\tilde{p}_x - p_j| \geq \epsilon/2] \leq \delta/2 \quad (1)$$

$$\Pr[|\tilde{p}_y - (1 - p_j)| \geq \epsilon/2] \leq \delta/2. \quad (2)$$

It is easy to see that Equation (2) is equivalent to $\Pr[|(1 - \tilde{p}_y) - p_j| \geq \epsilon/2] \leq \delta/2$ due to $|a| = |-a|$. Therefore, by combining it with Equation (1), we obtain

$$\Pr[|\tilde{p}_x - (1 - \tilde{p}_y)| \geq \epsilon] \leq \delta.$$

This completes the proof. ■

2.3 Cryptographic Tools

Definition 2.15 (Learning with Errors). Let $n, m, q \in \mathbb{N}$ be integer functions of the security parameter λ . Let $\chi = \chi(\lambda)$ be an error distribution over \mathbb{Z} . The LWE problem $\text{LWE}_{n,m,q,\chi}$ is to distinguish the following two distributions.

$$D_0 := \{(A, \mathbf{s}^\top A + \mathbf{e}) \mid A \leftarrow \mathbb{Z}_q^{n \times m}, \mathbf{s} \leftarrow \mathbb{Z}_q^n, \mathbf{e} \leftarrow \chi^m\} \text{ and } D_1 := \{(A, \mathbf{u}) \mid A \leftarrow \mathbb{Z}_q^{n \times m}, \mathbf{u} \leftarrow \mathbb{Z}_q^m\}.$$

When we say we assume the quantum hardness of the LWE problem or the QLWE assumption holds, we assume that for any QPT adversary \mathcal{A} , it holds that

$$|\Pr[\mathcal{A}(D_0) = 1] - \Pr[\mathcal{A}(D_1) = 1]| \leq \text{negl}(\lambda).$$

Definition 2.16 (Pseudorandom Generator). A pseudorandom generator (PRG) $\text{PRG} : \{0,1\}^\lambda \rightarrow \{0,1\}^{\lambda+\ell(\lambda)}$ with stretch $\ell(\lambda)$ (ℓ is some polynomial function) is a polynomial-time computable function that satisfies the following. For any QPT adversary \mathcal{A} , it holds that

$$\left| \Pr[\mathcal{A}(\text{PRG}(s)) = 1 \mid s \leftarrow \mathcal{U}_\lambda] - \Pr[\mathcal{A}(r) \mid r \leftarrow \mathcal{U}_{\lambda+\ell(\lambda)}] \right| \leq \text{negl}(\lambda),$$

where \mathcal{U}_m denotes the uniform distribution over $\{0,1\}^m$.

Theorem 2.17 ([HILL99]). If there exists a OWF, there exists a PRG.

Definition 2.18 (Quantum-Accessible Pseudo-Random Function). Let $\{\text{PRF}_K : \{0, 1\}^{\ell_1} \rightarrow \{0, 1\}^{\ell_2} \mid K \in \{0, 1\}^\lambda\}$ be a family of polynomially computable functions, where ℓ_1 and ℓ_2 are some polynomials of λ . We say that PRF is a quantum-accessible pseudo-random function (QPRF) family if for any QPT adversary \mathcal{A} , it holds that

$$\text{Adv}_{\mathcal{A}}^{\text{prf}}(\lambda) = \left| \Pr \left[\mathcal{A}^{\text{PRF}_K(\cdot)}(1^\lambda) = 1 \mid K \leftarrow \{0, 1\}^\lambda \right] - \Pr \left[\mathcal{A}^{\text{R}(\cdot)}(1^\lambda) = 1 \mid R \leftarrow \mathcal{U} \right] \right| \leq \text{negl}(\lambda),$$

where \mathcal{U} is the set of all functions from $\{0, 1\}^{\ell_1}$ to $\{0, 1\}^{\ell_2}$.

Theorem 2.19 ([Zha12a]). If there exists a OWF, there exists a QPRF.

Definition 2.20 (Puncturable PRF). A puncturable PRF (PPRF) is a tuple of algorithms $\text{PPRF} = (\text{PRF.Gen}, F, \text{Puncture})$ where $\{F_K : \{0, 1\}^{\ell_1} \rightarrow \{0, 1\}^{\ell_2} \mid K \in \{0, 1\}^\lambda\}$ is a PRF family and satisfies the following two conditions. Note that ℓ_1 and ℓ_2 are polynomials of λ .

Punctured correctness: For any polynomial-size set $S \subseteq \{0, 1\}^{\ell_1}$ and any $x \in \{0, 1\}^{\ell_1} \setminus S$, it holds that

$$\Pr \left[F_K(x) = F_{K_{\neq S}}(x) \mid K \leftarrow \text{PRF.Gen}(1^\lambda), K_{\neq S} \leftarrow \text{Puncture}(K, S) \right] = 1.$$

Pseudorandom at punctured point: For any polynomial-size set $S \subseteq \{0, 1\}^{\ell_1}$ and any QPT distinguisher \mathcal{A} , it holds that

$$\left| \Pr \left[\mathcal{A}(F_{K_{\neq S}}, \{F_K(x_i)\}_{x_i \in S}) = 1 \right] - \Pr \left[\mathcal{A}(F_{K_{\neq S}}, (\mathcal{U}_{\ell_2})^{|S|}) = 1 \right] \right| \leq \text{negl}(\lambda),$$

where $K \leftarrow \text{PRF.Gen}(1^\lambda)$, $K_{\neq S} \leftarrow \text{Puncture}(K, S)$ and \mathcal{U}_{ℓ_2} denotes the uniform distribution over $\{0, 1\}^{\ell_2}$.

If $S = \{x^*\}$ (i.e., puncturing a single point), we simply write $F_{\neq x^*}(\cdot)$ instead of $F_{K_{\neq S}}(\cdot)$.

It is easy to see that the Goldwasser-Goldreich-Micali tree-based construction of PRFs (GGM PRF) [GGM86] from one-way function yield puncturable PRFs where the size of the punctured key grows polynomially with the size of the set S being punctured [BW13, BGI14, KPTZ13]. Thus, we have:

Theorem 2.21 ([GGM86, BW13, BGI14, KPTZ13]). If OWFs exist, then for any polynomials $\ell_1(\lambda)$ and $\ell_2(\lambda)$, there exists a PPRF that maps ℓ_1 -bits to ℓ_2 -bits.

Definition 2.22 (SKE). An SKE scheme with plaintext space $\mathcal{P} = \{\mathcal{P}_\lambda\}_{\lambda \in \mathbb{N}}$ and ciphertext space $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$, where $\mathcal{C}_\lambda \subseteq \{0, 1\}^{\ell_{\text{ct}}}$ for some $\ell_{\text{ct}} = \ell_{\text{ct}}(\lambda)$, is a tuple of three algorithms.

$\text{Gen}(1^\lambda) \rightarrow k$: The key generation algorithm takes as input the security parameter λ , and outputs an encryption key k .

$\text{Enc}(k, m) \rightarrow \text{ct}$: The encryption algorithm takes as input k and a plaintext $m \in \mathcal{P}_\lambda$, and outputs a ciphertext $\text{ct} \in \mathcal{C}_\lambda$.

$\text{Dec}(k, \text{ct}) \rightarrow m'$: The decryption algorithm takes as input k and $\text{ct} \in \mathcal{C}_\lambda$, and outputs a plaintext $m' \in \mathcal{P}_\lambda \cup \{\perp\}$.

Correctness: An SKE scheme is correct if for all $\lambda \in \mathbb{N}$ and $m \in \mathcal{P}_\lambda$,

$$\Pr \left[\text{Dec}(k, \text{ct}) = m \mid k \leftarrow \text{Gen}(1^\lambda, 1^\kappa), \text{ct} \leftarrow \text{Enc}(k, m) \right] = 1.$$

Sparseness: In this work, we also require that most strings are not valid ciphertexts under a randomly generated key of an SKE scheme:

$$\Pr \left[\text{Dec}(k, c) \neq \perp \mid k \leftarrow \text{Gen}(1^\lambda), c \leftarrow \{0, 1\}^{\ell_{\text{ct}}} \right] \leq \text{negl}(\lambda).$$

Definition 2.23 (Ciphertext Pseudorandomness for SKE). An SKE scheme satisfies ciphertext pseudorandomness if for any (stateful) QPT \mathcal{A} , it holds that

$$2 \left| \Pr \left[\mathcal{A}^{\text{Enc}(k, \cdot)}(\text{ct}_b) = b \mid \begin{array}{l} 1^\kappa \leftarrow \mathcal{A}(1^\lambda), k \leftarrow \text{Gen}(1^\lambda, 1^\kappa), \\ m \leftarrow \mathcal{A}^{\text{Enc}(k, \cdot)}, b \leftarrow \{0, 1\}, \\ \text{ct}_0 \leftarrow \text{Enc}(k, m), \text{ct}_1 \leftarrow \{0, 1\}^{\ell_{\text{ct}}} \end{array} \right] - \frac{1}{2} \right| \leq \text{negl}(\lambda).$$

Theorem 2.24. If OWFs exist, there exists an SKE scheme with sparseness and ciphertext pseudorandomness.

The well-known PRF-based SKE satisfies ciphertext pseudorandomness. However, we need padding for sparseness. That is, a ciphertext is $(r, \text{PRF}_k(r) \oplus 0^{\ell - \ell_\mu} \| m)$ where $r \in \{0, 1\}^n$ is randomness of encryption, k is a PRF key, $\text{PRF} : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$ is a PRF, and $|m| = \ell_\mu$. We check that the first $\ell - \ell_\mu$ bits of $m' = \text{Dec}(k, (c_1, c_2))$ equals to $0^{\ell - \ell_\mu}$. If ℓ is sufficiently long, the scheme has sparseness.

Definition 2.25 (Constrained PRF (Syntax)). A constrained PRF (CPRF) with domain Dom , range Ran , and constraint family $\mathcal{F} = \{\mathcal{F}_{\lambda, \kappa}\}_{\lambda, \kappa \in \mathbb{N}}$ where $\mathcal{F}_{\lambda, \kappa} = \{f : \text{Dom} \rightarrow \{0, 1\}\}$ is a tuple of four algorithms.

$\text{Setup}(1^\lambda, 1^\kappa) \rightarrow \text{msk}$: The setup algorithm takes as input the security parameter λ and a constraint-family parameter κ , and outputs a master PRF key msk .

$\text{Constrain}(\text{msk}, f) \rightarrow \text{sk}_f$: The constrain algorithm takes as input λ and a constraint $f \in \mathcal{F}_{\lambda, \kappa}$, and outputs a constrained key sk_f .

$\text{Eval}(\text{msk}, x) \rightarrow y$: The evaluation algorithm takes as input msk and an input $x \in \text{Dom}$, and outputs a value $y \in \text{Ran}$.

$\text{CEval}(\text{sk}_f, x) \rightarrow y$: The constrained evaluation algorithm takes as input sk_f and $x \in \text{Dom}$, and outputs a value $y \in \text{Ran}$.

Definition 2.26 (Security for CPRF). A private CPRF should satisfy correctness, pseudorandomness, and privacy.

Correctness: A CPRF is correct if for any (stateful) QPT adversary \mathcal{A} , it holds that

$$\Pr \left[\begin{array}{l} \text{Eval}(\text{msk}, x) \neq \text{CEval}(\text{sk}_f, x) \\ \wedge x \in \text{Dom} \wedge f(x) = 0 \end{array} \mid \begin{array}{l} (1^\kappa, f) \leftarrow \mathcal{A}(1^\lambda), \\ \text{msk} \leftarrow \text{Setup}(1^\lambda, 1^\kappa), \\ \text{sk}_f \leftarrow \text{Constrain}(\text{msk}, f), \\ x \leftarrow \mathcal{A}^{\text{Eval}(\text{msk}, \cdot)}(\text{sk}_f) \end{array} \right] \leq \text{negl}(\lambda).$$

Selective single-key pseudorandomness: A CPRF is selectively single-key pseudorandom if for any (stateful) QPT adversary \mathcal{A} , it holds that

$$2 \left| \Pr \left[\begin{array}{l} \mathcal{A}^{\text{Eval}(\text{msk}, \cdot)}(y_b) = b \\ \wedge x \notin \mathcal{Q}_e \\ \wedge f(x) \neq 0 \end{array} \mid \begin{array}{l} (1^\kappa, f) \leftarrow \mathcal{A}(1^\lambda), \\ \text{msk} \leftarrow \text{Setup}(1^\lambda, 1^\kappa), \\ \text{sk}_f \leftarrow \text{Constrain}(\text{msk}, f) \\ x \leftarrow \mathcal{A}^{\text{Eval}(\text{msk}, \cdot)}(\text{sk}_f) \\ y_0 := \text{Eval}(\text{msk}, x), y_1 \leftarrow \text{Ran}, \\ b \leftarrow \{0, 1\} \end{array} \right] - \frac{1}{2} \right| \leq \text{negl}(\lambda),$$

where \mathcal{Q}_e is the sets of queries to $\text{Eval}(\text{msk}, \cdot)$.

Selective single-key privacy: A CPRF is selectively single-key private if for any (stateful) QPT adversary \mathcal{A} , there exists a stateful PPT simulator $\text{Sim} = (\text{Sim}_1, \text{Sim}_2)$ that satisfying that

$$2 \left| \Pr \left[\mathcal{A}^{\mathcal{O}_b(\cdot)}(\text{sk}_b) = b \mid \begin{array}{l} (1^k, f) \leftarrow \mathcal{A}(1^\lambda), \\ \text{msk} \leftarrow \text{Setup}(1^\lambda, 1^k), b \leftarrow \{0, 1\}, \\ \text{sk}_0 \leftarrow \text{Constrain}(\text{msk}, f), \\ (\text{st}_{\text{Sim}}, \text{sk}_1) \leftarrow \text{Sim}_1(1^k, 1^\lambda) \end{array} \right] - \frac{1}{2} \right| \leq \text{negl}(\lambda),$$

where $\mathcal{O}_0(\cdot) := \text{Eval}(\text{msk}, \cdot)$ and $\mathcal{O}_1(\cdot) := \text{Sim}_2(\text{st}_{\text{Sim}}, \cdot, f(\cdot))$.

We say that a CPRF is a selectively single-key private CPRF if it satisfies correctness, selective single-key pseudorandomness, and selective single-key privacy.

Theorem 2.27 ([BTWV17, PS18]). If the QLWE assumption holds, there exists a selectively single-key private CPRF for polynomial-size classical circuits.

Definition 2.28 (PKE). A PKE with plaintext space $\mathcal{P} = \{\mathcal{P}_\lambda\}_{\lambda \in \mathbb{N}}$, ciphertext space $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ is a tuple of three algorithms.

$\text{Gen}(1^\lambda) \rightarrow (\text{pk}, \text{sk})$: The key generation algorithm takes as input the security parameter λ and outputs a key pair (pk, sk) .

$\text{Enc}(\text{pk}, m) \rightarrow \text{ct}$: The encryption algorithm takes as input pk , a plaintext $m \in \mathcal{P}$, and outputs a ciphertext $\text{ct} \in \mathcal{C}$.

$\text{Dec}(\text{sk}, \text{ct}) \rightarrow m' / \perp$: The decryption algorithm takes as input sk and $\text{ct} \in \mathcal{C}$, and outputs a plaintext $m' \in \mathcal{P}$ or \perp .

Correctness: A PKE scheme is correct if for all $\lambda \in \mathbb{N}$ and $m \in \mathcal{P}_\lambda$, it holds that

$$\Pr \left[\text{Dec}(\text{sk}, \text{ct}) = m \mid (\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda), \text{ct} \leftarrow \text{Enc}(\text{pk}, m) \right] = 1.$$

Definition 2.29 (CCA Security for PKE). A PKE scheme is CCA secure if for any (stateful) QPT adversary \mathcal{A} , it holds that

$$2 \left| \Pr \left[\begin{array}{l} b^* = b \\ \wedge \text{ct}_b \notin \mathcal{Q} \end{array} \mid \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda), \\ (m_0, m_1) \leftarrow \mathcal{A}^{\text{Dec}(\text{sk}, \cdot)}(1^\lambda, \text{pk}), \\ b \leftarrow \{0, 1\}, \text{ct}_b \leftarrow \text{Enc}(\text{pk}, m_b), \\ b' \leftarrow \mathcal{A}^{\text{Dec}(\text{sk}, \cdot)}(1^\lambda, \text{ct}_b) \end{array} \right] - \frac{1}{2} \right| \leq \text{negl}(\lambda),$$

where \mathcal{Q} is the set of queries to $\text{Dec}(\text{sk}, \cdot)$ after \mathcal{A} is given ct_b .

Theorem 2.30 ([Pei09]). If the QLWE assumption holds, there exists a PKE scheme that satisfies CCA security.

Definition 2.31 (Indistinguishability Obfuscator [BGI⁺12]). A PPT algorithm $i\mathcal{O}$ is a secure IO for a classical circuit class $\{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ if it satisfies the following two conditions.

Functionality: For any security parameter $\lambda \in \mathbb{N}$, circuit $C \in \mathcal{C}_\lambda$, and input x , we have that

$$\Pr [C'(x) = C(x) \mid C' \leftarrow i\mathcal{O}(C)] = 1.$$

Indistinguishability: For any PPT Samp and QPT distinguisher \mathcal{D} , the following holds:

If $\Pr [\forall x, C_0(x) = C_1(x) \mid (C_0, C_1, \text{aux}) \leftarrow \text{Samp}(1^\lambda)] > 1 - \text{negl}(\lambda)$, then we have

$$\text{Adv}_{i\mathcal{O}, \mathcal{D}}^{\text{io}}(\lambda) := \left| \Pr \left[\mathcal{D}(i\mathcal{O}(C_0), \text{aux}) = 1 \mid (C_0, C_1, \text{aux}) \leftarrow \text{Samp}(1^\lambda) \right] - \Pr \left[\mathcal{D}(i\mathcal{O}(C_1), \text{aux}) = 1 \mid (C_0, C_1, \text{aux}) \leftarrow \text{Samp}(1^\lambda) \right] \right| \leq \text{negl}(\lambda).$$

There are a few candidates of secure IO for polynomial-size classical circuits against quantum adversaries [BGMZ18, CHVW19, AP20, DQV⁺21]. In some candidates [WW21, GP21], the assumptions behind the constructions were found to be false [HJL21].

3 Definition of Quantum Watermarking

We introduce definitions for watermarking PRFs against quantum adversaries in this section.

3.1 Syntax and Pseudorandomness

Definition 3.1 (Watermarking PRF). A watermarking PRF WMPRF for the message space $\mathcal{M} := \{0, 1\}^{\ell_m}$ with domain Dom and range Ran is a tuple of five algorithms $(\text{Setup}, \text{Gen}, \text{Eval}, \text{Mark}, \text{Extract})$.

$\text{Setup}(1^\lambda) \rightarrow (\text{pp}, \text{xk})$: The setup algorithm takes as input the security parameter and outputs a public parameter pp and an extraction key xk .

$\text{Gen}(\text{pp}) \rightarrow (\text{prfk}, \tau)$: The key generation algorithm takes as input the public parameter pp and outputs a PRF key prfk and a public tag τ .

$\text{Eval}(\text{prfk}, x) \rightarrow y$: The evaluation algorithm takes as input a PRF key prfk and an input $x \in \text{Dom}$ and outputs $y \in \text{Ran}$.

$\text{Mark}(\text{pp}, \text{prfk}, m) \rightarrow \tilde{C}$: The mark algorithm takes as input the public parameter pp , a PRF key prfk , and a message $m \in \{0, 1\}^{\ell_m}$, and outputs a marked evaluation circuit \tilde{C} .

$\text{Extract}(\text{xk}, \tau, C', \epsilon) \rightarrow m'$: The extraction algorithm takes as input an extraction key xk , a tag τ , a quantum circuit with classical inputs and outputs $C' = (q, \mathbf{U})$, and a parameter ϵ , and outputs m' where $m' \in \{0, 1\}^{\ell_m} \cup \{\text{unmarked}\}$.

Evaluation Correctness: For any message $m \in \{0, 1\}^{\ell_m}$, it holds that

$$\Pr \left[\tilde{C}(x) = \text{Eval}(\text{prfk}, x) \mid \begin{array}{l} (\text{pp}, \text{xk}) \leftarrow \text{Setup}(1^\lambda) \\ (\text{prfk}, \tau) \leftarrow \text{Gen}(\text{pp}) \\ \tilde{C} \leftarrow \text{Mark}(\text{pp}, \text{prfk}, m) \\ x \leftarrow \text{Dom} \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

Remark 3.2 (On extraction correctness). Usually, a watermarking PRF scheme is required to satisfy extraction correctness that ensures that we can correctly extract the embedded mark from an honestly marked circuit. However, as observed by Quach et al. [QWZ18], if we require the extraction correctness to hold for a randomly chosen PRF key, it is implied by unremovability defined below. Note that the unremovability defined below considers a distinguisher as a pirate circuit. However, it implies the extraction correctness since we can easily transform an honestly marked circuit into a successful distinguisher. Thus, we do not explicitly require a watermarking PRF scheme to satisfy extraction correctness in this work.

Remark 3.3 (On public marking). We consider only watermarking PRFs with public marking as in Definition 3.1 since we can achieve public marking by default. The reason is as follows. Suppose that we generate pp , xk , and a marking key mk at the setup. When we generate a PRF key and a public tag at Gen , we can first generate $(\text{pp}', \text{xk}', \text{mk}') \leftarrow \text{Setup}(1^\lambda)$ from scratch (ignoring the original $(\text{pp}, \text{xk}, \text{mk})$) and set a PRF key $\widehat{\text{prfk}} := (\text{prfk}', \text{mk}')$ and a public tag $\widehat{\tau} := (\text{pp}', \text{xk}', \tau')$ where $(\text{prfk}', \tau') \leftarrow \text{Gen}(\text{pp}')$. That is, anyone can generate a marked circuit from $\widehat{\text{prfk}} = (\text{prfk}', \text{mk}')$ by $\text{Mark}(\text{mk}', \text{prfk}', m)$. Therefore, we consider public marking by default in our model.

Remark 3.4 (On private marking). We might prefer private marking in some settings since we might want to prevent adversaries from forging a watermarked PRF. We can convert watermarking PRFs in Definition 3.1 into ones with private marking by using signatures. Below, we assume that a PRF key prfk includes its public tag τ since it does not harm security. At the setup phase, we also generate a signature key pair $(\text{vk}, \text{sk}) \leftarrow \text{SIG.Gen}(1^\lambda)$ and set a mark key $\text{mk}' := (\text{vk}, \text{sk})$ and an extraction key $\text{xk}' := (\text{xk}, \text{vk})$. To embed a message m into prfk , we generate a signature $\sigma \leftarrow \text{SIG.Sign}(\text{sk}, \tau \| m)$ and generate $\tilde{C} \leftarrow \text{Mark}(\text{pp}, \text{prfk}, m \| \sigma)$. To extract a message, we run $m' \leftarrow \text{Extract}(\text{xk}', \tau, \tilde{C}, \epsilon)$, parse $m' = m \| \sigma$, and run $\text{SIG.Vrfy}(\text{vk}, \tau \| m, \sigma)$. If the verification result is \top , we output m . This conversion is the same as what Goyal et al. [GKM⁺19] proposed. Adversaries cannot forge a signature for $\tau^* \| m^* \neq \tau \| m$ by the unforgeability of SIG. Intuitively, if an adversary can forge a watermarked PRF whose functionality is different from those of watermarked PRFs given from a mark oracle, $\tau^* \neq \tau$ should hold since public tags are related to PRF keys. This breaks the unforgeability of SIG. Thus, we expect that adversaries cannot break the unforgeability of watermarking. However, we do not formally define watermarking unforgeability against quantum adversaries since it is not a scope of this work. We leave it as future work.

Discussion on syntax. Definition 3.1 is a natural quantum variant of classical watermarking PRFs except that the key generation algorithm outputs a public tag τ , and the extraction algorithm uses it. Such a public tag is not used in previous works on watermarking PRFs [CHN⁺18, KW21, QWZ18, KW19, YAL⁺19]. A public tag should not harm watermarking PRF security. We justify using τ as follows.

First, we need to obtain many pairs of input and output to extract an embedded message from a marked PRF in almost all known (classical) watermarking constructions [CHN⁺18, BLW17, KW21, QWZ18, KW19, YAL⁺19, GKM⁺19, Nis20]. This is because we must check whether a tested PRF circuit outputs particular values for particular inputs which *depends on the target PRF* (such particular inputs are known as marked points). Suppose marked points are fixed and do not depend on a PRF that will be marked. In that case, an adversary can easily remove an embedded message by destroying functionalities at the fixed marked points that could be revealed via a (non-target) marked PRF that an adversary generated. Recall that we consider the public marking setting. The attack was already observed by Cohen et al. [CHN⁺18].

Second, we consider a stronger adversary model than that in most previous works as the definition of traceable PRFs by Goyal et al. [GKWW21]. An adversary outputs a distinguisher-based pirate circuit in our security definition rather than a pirate circuit that computes an entire output of a PRF. This is a refined and realistic model, as Goyal et al. [GKWW21] argued (and we explain in Section 1.4). In this model, we cannot obtain a valid input-output pair from a pirate circuit anymore. Such a pair is typical information related to a target PRF. Goyal et al. resolve this issue by introducing a tracing key that is generated from a target PRF. Note that parameters of watermarking (pp and xk) should *not* be generated from a PRF since we consider many different PRF keys in the watermarking PRF setting.

Thus, if we would like to achieve an extraction algorithm and the stronger security notion simultaneously, an extraction algorithm should somehow take information related to a target PRF as input to correctly extract an embedded message. In the weaker adversary model, an extraction algorithm can easily obtain many valid input and output pairs by running a tested circuit many times. However, in the stronger distinguisher-based pirate circuit model, a pirate circuit outputs a single decision bit.

To resolve this issue, we introduce public tags. We think it is natural to have information related to the original PRF key in an extraction algorithm. In reality, we check a circuit when a user claims that her/his PRF key (PRF evaluation circuit) is illegally used. Thus, it is natural to expect we can use a user's public tag for extraction. This setting resembles watermarking for public-key cryptographic primitives, where a user public key is available in an extraction algorithm. In addition, public tags do not harm PRF security in our constructions. It is unclear whether we can achieve unremovability in the stronger distinguisher-based model without any syntax change (even in the classical setting).¹⁰

¹⁰Even if we consider the weaker adversary model, the same issue appears in the quantum setting in the end. If we run a

Extended pseudorandomness. We consider extended weak pseudorandomness, where weak pseudorandomness holds even if the adversary generates pp . This notion is the counterpart of extended pseudorandomness by Quach et al. [QWZ18], where pseudorandomness holds in the presence of the extraction oracle. However, our pseudorandomness holds even against an authority unlike extended pseudorandomness by Quach et al. since we allow adversaries to generate a public parameter.

Definition 3.5 (Extended Weak Pseudorandomness against Authority). *To define extended weak pseudorandomness for watermarking PRFs, we define the game $\text{Exp}_{\mathcal{A}, \text{WMPRF}}^{\text{ext-wprf}}(\lambda)$ as follows.*

1. \mathcal{A} first sends pp to the challenger.
2. The challenger generates $(\text{prfk}, \tau) \leftarrow \text{Gen}(\text{pp})$ and sends τ to \mathcal{A} .
3. The challenger chooses $\text{coin} \leftarrow \{0, 1\}$. \mathcal{A} can access to the following oracles.

O_{wprf} : When this is invoked (no input), it returns (a, b) where $a \leftarrow \text{Dom}$ and $b := \text{Eval}(\text{prfk}, a)$.

O_{chall} : When this is invoked (no input), it returns:

- (a, b) where $a \leftarrow \text{Dom}$ and $b := \text{Eval}(\text{prfk}, a)$ if $\text{coin} = 0$,
- (a, b) where $a \leftarrow \text{Dom}$ and $b \leftarrow \text{Ran}$ if $\text{coin} = 1$.

This oracle is invoked only once.

4. When \mathcal{A} terminates with output coin' , the challenger outputs 1 if $\text{coin} = \text{coin}'$ and 0 otherwise.

We say that WMPRF is extended weak pseudorandom if for every QPT \mathcal{A} , we have

$$\text{Adv}_{\mathcal{A}, \text{WMPRF}}^{\text{ext-wprf}}(\lambda) = 2 \left| \Pr \left[\text{Exp}_{\mathcal{A}, \text{WMPRF}}^{\text{ext-wprf}}(\lambda) = 1 \right] - \frac{1}{2} \right| = \text{negl}(\lambda).$$

3.2 Unremovability against Quantum Adversaries

We define unremovability for watermarking PRFs against quantum adversaries.

Definition 3.6 (Unremovability for private extraction). *We consider the public marking and secret extraction setting here. Let $\epsilon \geq 0$. We define the game $\text{Exp}_{\mathcal{A}, \text{WMPRF}}^{\text{nmv}}(\lambda, \epsilon)$ as follows.*

1. The challenger generates $(\text{pp}, \text{xk}) \leftarrow \text{Setup}(1^\lambda)$ and gives pp to the adversary \mathcal{A} . \mathcal{A} send $\text{m} \in \{0, 1\}^{\ell_{\text{m}}}$ to the challenger. The challenger generates $(\text{prfk}, \tau) \leftarrow \text{Gen}(\text{pp})$, computes $\tilde{\text{C}} \leftarrow \text{Mark}(\text{pp}, \text{prfk}, \text{m})$, and sends τ and $\tilde{\text{C}}$ to \mathcal{A} .
2. \mathcal{A} can access to the following oracle.

O_{ext} : On input τ' and a quantum circuit \mathcal{C} , it returns $\text{Extract}(\text{xk}, \mathcal{C}, \tau', \epsilon)$.

3. Finally, the adversary outputs a “pirate” quantum circuit $\mathcal{C}_{\text{sk}} = (q, \mathbf{U})$, where \mathcal{C}_{sk} is a quantum program with classical inputs and outputs whose first register (i.e., output register) is \mathbb{C}^2 and \mathbf{U} is a compact classical description of $\{\mathbf{U}_{x,y}\}_{x \in \text{Dom}, y \in \text{Ran}}$.

Let D be the following distribution.

D : Generate $b \leftarrow \{0, 1\}$, $x \leftarrow \text{Dom}$, and $y_0 \leftarrow \text{Ran}$. Compute $y_1 \leftarrow \text{Eval}(\text{prfk}, x)$. Output (b, x, y_b) .

quantum circuit for an input and measure the output, the measurement could irreversibly alter the quantum state and we lost the functionality of the original quantum state. That is, there is no guarantee that we can correctly check whether a tested quantum circuit is marked or not *after* we obtain a single valid pair of input and output by running the circuit. However, as we explained above, we want to obtain information related to a target PRF for extraction. Thus, we need a public tag in the syntax in either case.

We also let $\mathcal{P} = (\mathbf{P}_{b,x,y}, \mathbf{Q}_{b,x,y})_{b,x,y}$ be a collection of binary outcome projective measurements, where

$$\mathbf{P}_{b,x,y} = \mathbf{U}_{x,y}^\dagger |b\rangle \langle b| \mathbf{U}_{x,y} \quad \text{and} \quad \mathbf{Q}_{b,x,y} = \mathbf{I} - \mathbf{P}_{b,x,y}.$$

Moreover, we let $\mathcal{M}_D = (\mathbf{P}_D, \mathbf{Q}_D)$ be binary outcome POVMs, where

$$\mathbf{P}_D = \sum_{r \in \mathcal{R}} \frac{1}{|\mathcal{R}|} \mathbf{P}_{D(r)} \quad \text{and} \quad \mathbf{Q}_D = \mathbf{I} - \mathbf{P}_D.$$

Live: When applying the measurement $\text{ProjImp}(\mathcal{M}_D)$ to q , we obtain a value p such that $p \geq \frac{1}{2} + \epsilon$.

GoodExt: When Computing $m' \leftarrow \text{Extract}(xk, C_{\mathfrak{G}}, \tau, \epsilon)$, it holds that $m' \neq \text{unmarked}$.

BadExt: When Computing $m' \leftarrow \text{Extract}(xk, C_{\mathfrak{G}}, \tau, \epsilon)$, it holds that $m' \notin \{m, \text{unmarked}\}$.

We say that WMPRF satisfies unremovability if for every $\epsilon > 0$ and QPT \mathcal{A} , we have

$$\Pr[\text{BadExt}] \leq \text{negl}(\lambda) \quad \text{and} \quad \Pr[\text{GoodExt}] \geq \Pr[\text{Live}] - \text{negl}(\lambda).$$

Intuitively, $(\mathbf{P}_{b,x,y}, \mathbf{Q}_{b,x,y})$ is a projective measurement that feeds (x, y) to $C_{\mathfrak{G}}$ and checks whether the outcome is b or not (and then uncomputes). Then, \mathcal{M}_D can be seen as POVMs that results in 0 with the probability that $C_{\mathfrak{G}}$ can correctly guess b from (x, y_b) for (b, x, y_b) generated randomly from D .

Remark 3.7 (On attack model). We check whether $C_{\mathfrak{G}}$ correctly distinguishes a real PRF value from a random value or not by applying $\text{ProjImp}(\mathcal{M}_D)$ to q . This attack model follows the refined and more realistic attack model by Goyal et al. [GKWW21]. The adversary outputs a pirate circuit that computes an entire PRF value in all previous works except their work.

The distinguisher-based pirate circuit model is compatible with the (quantum) pirate decoder model of traitor tracing. Thus, our attack model also follows the attack model of quantum traitor tracing (the black box projection model) by Zhandry [Zha20, Section 4.2].¹¹

As in the traitor tracing setting [Zha20], $\text{ProjImp}(\mathcal{M}_D)$ is inefficient in general. We can handle this issue as Zhandry did. We will use an approximate version of $\text{ProjImp}(\mathcal{M}_D)$ to achieve an efficient reduction. In addition, we cannot apply both $\text{ProjImp}(\mathcal{M}_D)$ and Extract to $C_{\mathfrak{G}}$ simultaneously. However, the condition $\Pr[\text{GoodExt}] \geq \Pr[\text{Live}] - \text{negl}(\lambda)$ claims that an embedded mark cannot be removed as long as the pirate circuit is alive. This fits the spirit of watermarking. See Zhandry's paper [Zha20, Section 4] for more discussion on the models.

Remark 3.8 (On selective message). As we see in Definition 3.6, we consider the selective setting for private extraction case, where \mathcal{A} must send the target message m to the challenger before \mathcal{A} accesses to the oracle O_{ext} and after pp is given. This is the same setting as that by Quach et al. [QWZ18]. We can consider the fully adaptive setting, where \mathcal{A} can send the target message m after it accesses to the oracle O_{ext} , as Kim and Wu [KW19]. However, our privately extractable watermarking PRF satisfies only selective security. Thus, we write only the selective variant for the private extraction case.

Definition 3.9 (Unremovability for Public Extraction). This is the same as Definition 3.6 except we use the game $\text{Exp}_{\mathcal{A}, \text{WMPRF}}^{\text{pub-ext-nrmv}}(\lambda, \epsilon)$ defined in the same way as $\text{Exp}_{\mathcal{A}, \text{WMPRF}}^{\text{nrmv}}(\lambda, \epsilon)$ except the following differences.

- In item 1, \mathcal{A} is given xk together with pp .
- Item 2 is removed.

¹¹In the watermarking setting, an extraction algorithm can take the description of a pirate circuit as input (corresponding to the software decoder model [Zha20, Section 4.2]), unlike the black-box tracing model of traitor tracing. However, we use a pirate circuit in the black box way for our extraction algorithms. Thus, we follow the black box projection model by Zhandry [Zha20].

4 Definition of Extraction-Less Watermarking

We introduce the notion of extraction-less watermarking PRF as an intermediate primitive towards watermarking PRFs secure against quantum adversaries.

4.1 Syntax and Pseudorandomness

Definition 4.1 (Extraction-Less Watermarking PRF). *An extraction-less watermarking PRF WMPRF for the message space $\{0,1\}^{\ell_m}$ with domain Dom and range Ran is a tuple of five algorithms (Setup, Gen, Eval, Mark, Sim), where the first four algorithms have the same input/output behavior as those defined in Definition 3.1 and Sim has the following input/output behavior.*

$\text{Sim}(xk, \tau, i) \rightarrow (\gamma, x, y)$: *The simulation algorithm Sim takes as input the extraction key xk , a tag τ , and an index i , and outputs a tuple (γ, x, y) .*

Evaluation Correctness: *It is defined in exactly the same way as the evaluation correctness for watermarking PRF defined in Definition 3.1.*

Extended pseudorandomness. Extended pseudorandomness for extraction-less watermarking PRF is defined in exactly the same way as that for watermarking PRF, that is Definition 3.5.

4.2 Simulatability for Mark-Dependent Distributions (SIM-MDD Security)

We introduce the security notion for extraction-less watermarking PRF that we call simulatability for mark-dependent distributions. Let D and D^{rev} be the following distributions.

D : Generate $b \leftarrow \{0,1\}$, $x \leftarrow \text{Dom}$, and $y_0 \leftarrow \text{Ran}$. Compute $y_1 \leftarrow \text{Eval}(\text{prfk}, x)$. Output (b, x, y_b) .

D^{rev} : Generate $(b, x, y) \leftarrow D$. Output $(1 \oplus b, x, y)$.

Namely, D is the distribution that outputs a random value if the first bit $b = 0$ and a PRF evaluation if the first bit $b = 1$, and D^{rev} is its opposite (i.e., a PRF evaluation if $b = 0$ and a random value if $b = 1$). SIM-MDD security is a security notion that guarantees that an adversary given $\tilde{C} \leftarrow \text{Mark}(\text{mk}, \text{prfk}, m)$ cannot distinguish an output of $\text{Sim}(xk, \tau, i)$ from that of D if $m[i] = 0$ and from that of D^{rev} if $m[i] = 1$.

Definition 4.2 (SIM-MDD Security with Private Simulation). *To define SIM-MDD security with private simulation, we define the game $\text{Expt}_{i^*, \mathcal{A}, \text{WMPRF}}^{\text{sim-mdd}}(\lambda)$ as follows, where $i^* \in [\ell_m]$.*

1. *The challenger generates $(\text{pp}, xk) \leftarrow \text{Setup}(1^\lambda)$ and sends pp to \mathcal{A} . \mathcal{A} sends $m \in \{0,1\}^{\ell_m}$ to the challenger. The challenger generates $(\text{prfk}, \tau) \leftarrow \text{Gen}(\text{pp})$ and computes $\tilde{C} \leftarrow \text{Mark}(\text{mk}, \text{prfk}, m)$. The challenger sends τ and \tilde{C} to \mathcal{A} .*

2. *\mathcal{A} can access to the following oracle.*

O_{sim} : *On input τ' and $i' \in [\ell_m]$, it returns $\text{Sim}(xk, \tau', i')$.*

3. *Let D_{real, i^*} be the following distribution. Note that D_{real, i^*} is identical with D if $m[i^*] = 0$ and with D^{rev} if $m[i^*] = 1$.*

D_{real, i^*} : *Generate $\gamma \leftarrow \{0,1\}$ and $x \leftarrow \text{Dom}$. Then, if $\gamma = m[i^*]$, generate $y \leftarrow \text{Ran}$, and otherwise, compute $y \leftarrow \text{Eval}(\text{prfk}, x)$. Output (γ, x, y) .*

The challenger generates $\text{coin} \leftarrow \{0,1\}$. If $\text{coin} = 0$, the challenger samples $(\gamma, x, y) \leftarrow D_{\text{real}, i^}$. If $\text{coin} = 1$, the challenger generates $(\gamma, x, y) \leftarrow \text{Sim}(xk, \tau, i^*)$. The challenger sends (γ, x, y) to \mathcal{A} .*

4. When \mathcal{A} terminates with output coin' , the challenger outputs 1 if $\text{coin} = \text{coin}'$ and 0 otherwise.

Note that \mathcal{A} is not allowed to access to O_{sim} after \mathcal{A} is given (γ, x, y) .

We say that WMPRF is SIM-MDD secure if for every $i^* \in [\ell_m]$ and QPT \mathcal{A} , we have

$$\text{Adv}_{i^*, \mathcal{A}, \text{WMPRF}}^{\text{sim-mdd}}(\lambda) = 2 \left| \Pr \left[\text{Exp}_{i^*, \mathcal{A}, \text{WMPRF}}^{\text{sim-mdd}}(\lambda) = 1 \right] - \frac{1}{2} \right| = \text{negl}(\lambda).$$

We consider the selective setting above as unremovability for private extraction in Definition 3.6 since we use SIM-MDD security with private simulation to achieve unremovability for private simulation.

Remark 4.3 (On multi challenge security). We can prove that the above definition implies the multi-challenge variant where polynomially many outputs of $\text{Sim}(xk, \tau, i^*)$ are required to be indistinguishable from those of D_{real, i^*} . This is done by hybrid arguments where outputs of $\text{Sim}(xk, \tau, i^*)$ are simulated using O_{sim} and those of D_{real, i^*} are simulated using \tilde{C} . To apply Theorem 2.11, we need the multi challenge variant. However, we consider the single challenge variant due to the implication above. A similar remark is applied to the variants of SIM-MDD security introduced below.

SIM-MDD security with private simulation under the \mathcal{API} oracle. Let the \mathcal{API} oracle be an oracle that is given $(\epsilon, \delta, \tau', i')$ and a quantum state q , and returns the result of $\mathcal{API}_{\mathcal{P}, D_{\tau', i'}}^{\epsilon, \delta}(q)$ and the post measurement state, where \mathcal{P} is defined in the same way as that in Definition 3.6 and $D_{\tau', i'}$ be the distribution that outputs randomly generated $(\gamma, x, y) \leftarrow \text{Sim}(xk, \tau', i')$. The \mathcal{API} oracle cannot be simulated using the simulation oracle O_{sim} since we need superposition of outputs of Sim to compute $\mathcal{API}_{\mathcal{P}, D_{\tau', i'}}^{\epsilon, \delta}(q)$. When constructing watermarking PRFs with private simulation from extraction-less watermarking PRFs, the underlying extraction-less watermarking PRF scheme needs to satisfy SIM-MDD security with private simulation under the \mathcal{API} oracle that we call QSIM-MDD security with private simulation. The reason is as follows. In the security analysis of the construction, the indistinguishability guarantee provided by SIM-MDD security needs to hold for an adversary against the resulting watermarking scheme who can access the extraction oracle. This means that it also needs to hold for an adversary who can access the \mathcal{API} oracle since \mathcal{API} is repeatedly invoked in the extraction algorithm of the resulting scheme.

Fortunately, as we will see, we can generically convert an extraction-less watermarking PRF scheme satisfying SIM-MDD security with private simulation into one satisfying QSIM-MDD security with private simulation, using QPRFs. Thus, when realizing an extraction-less watermarking PRF scheme as an intermediate step towards privately extractable watermarking PRFs, we can concentrate on realizing one satisfying SIM-MDD security with private simulation.

Remark 4.4. There is a similar issue in the traitor tracing setting. If PLBE is a secret-key based one, we need a counterpart of QSIM-MDD in secret-key based PLBE to achieve traitor tracing with a secret tracing algorithm against quantum adversaries by using Zhandry's framework [Zha20]. Note that Zhandry focuses on public-key based PLBE in his work [Zha20].

Definition 4.5 (QSIM-MDD Security with Private Simulation). Let $D_{\tau, i}$ be a distribution defined as follows.

$D_{\tau, i}$: Output $(\gamma, x, y) \leftarrow \text{Sim}(xk, \tau, i)$.

Then, we define the game $\text{Exp}_{i^*, \mathcal{A}, \text{WMPRF}}^{\text{q-sim-mdd}}(\lambda)$ in the same way as $\text{Exp}_{i^*, \mathcal{A}, \text{WMPRF}}^{\text{sim-mdd}}(\lambda)$ except that in addition to O_{sim} , \mathcal{A} can access to the following oracle in the step 2.

O_{api} : On input $(\epsilon, \delta, \tau', i')$ and a quantum state q , it returns the result of $\mathcal{API}_{\mathcal{P}, D_{\tau', i'}}^{\epsilon, \delta}(q)$ and the post measurement state, where \mathcal{P} is defined in the same way as that in Definition 3.6.

We say that WMPRF is QSIM-MDD secure with private simulation if for every $i^* \in [\ell_m]$ and QPT \mathcal{A} , we have

$$\text{Adv}_{i^*, \mathcal{A}, \text{WMPRF}}^{\text{q-sim-mdd}}(\lambda) = 2 \left| \Pr \left[\text{Exp}_{i^*, \mathcal{A}, \text{WMPRF}}^{\text{q-sim-mdd}}(\lambda) = 1 \right] - \frac{1}{2} \right| = \text{negl}(\lambda).$$

We have the following theorem.

Theorem 4.6. *Assume there exists an extraction-less watermarking PRF scheme satisfying SIM-MDD security with private simulation and a QPRF. Then, there exists an extraction-less watermarking PRF scheme satisfying QSIM-MDD security with private simulation.*

We prove this theorem in Appendix A.

Definition 4.7 (SIM-MDD Security with Public Simulation). *We define the game $\text{Exp}_{i^*, \mathcal{A}, \text{WMPRF}}^{\text{sim-mdd-pub}}(\lambda)$ in the same way as $\text{Exp}_{i^*, \mathcal{A}, \text{WMPRF}}^{\text{sim-mdd}}(\lambda)$ except the following differences, where $i^* \in [\ell_m]$.*

- In item 1, \mathcal{A} is given xk together with pp .
- Item 2 is removed.

We say that WMPRF satisfies SIM-MDD security with public simulation if for every $i^* \in [\ell_m]$ and QPT \mathcal{A} , we have

$$\text{Adv}_{i^*, \mathcal{A}, \text{WMPRF}}^{\text{sim-mdd-pub}}(\lambda) = 2 \left| \Pr \left[\text{Exp}_{i^*, \mathcal{A}, \text{WMPRF}}^{\text{sim-mdd-pub}}(\lambda) = 1 \right] - \frac{1}{2} \right| = \text{negl}(\lambda).$$

5 Watermarking PRF from Extraction-Less Watermarking PRF

We show how to construct watermarking PRF secure against quantum adversaries from extraction-less watermarking PRF.

Let ELWMPRF = (Setup, Gen, Eval, Mark, Sim) be an extraction-less watermarking PRF scheme whose message space is $\{0, 1\}^{\ell_m+1}$. We construct a watermarking PRF scheme WMPRF = (WM.Setup, WM.Gen, WM.Eval, WM.Mark, $\mathcal{E}_{\text{extract}}$) whose message space is $\{0, 1\}^{\ell_m}$ as follows. We use Setup, Gen, and Eval as WM.Setup, WM.Gen, and WM.Eval, respectively. Thus, the domain and range of WMPRF are the same as those of ELWMPRF. Also, we construct WM.Mark and $\mathcal{E}_{\text{extract}}$ as follows.

WM.Mark(pp, prfk, m):

- Output $\tilde{C} \leftarrow \text{Mark}(pp, \text{prfk}, m \| 0)$.

$\mathcal{E}_{\text{extract}}(xk, C, \tau, \epsilon)$:

- Let $\epsilon' = \epsilon/4(\ell_m + 1)$ and $\delta' = 2^{-\lambda}$.
- Parse $(q, \mathbf{U}) \leftarrow C$.
- Let \mathcal{P} be defined in the same way as that in Definition 3.6 and $D_{\tau, i}$ be the following distribution for every $i \in [\ell_m + 1]$.

$D_{\tau, i}$: Output $(\gamma, x, y) \leftarrow \text{Sim}(xk, \tau, i)$.

- Compute $\tilde{p}_{\ell_m+1} \leftarrow \mathcal{A}PI_{\mathcal{P}, D_{\tau, \ell_m+1}}^{\epsilon', \delta'}(q)$. If $\tilde{p}_{\ell_m+1} < \frac{1}{2} + \epsilon - 4\epsilon'$, return unmarked. Otherwise, letting q_0 be the post-measurement state, go to the next step.
- For all $i \in [\ell_m]$, do the following.

1. Compute $\tilde{p}_i \leftarrow \mathcal{A}PI_{\mathcal{P}, D_{\tau, i}}^{\epsilon', \delta'}(q_{i-1})$. Let q_i be the post-measurement state.

2. If $\tilde{p}_i > \frac{1}{2} + \epsilon - 4(i+1)\epsilon'$, set $m'_i = 0$. If $\tilde{p}_i < \frac{1}{2} - \epsilon + 4(i+1)\epsilon'$, set $m'_i = 1$. Otherwise, exit the loop and output $m' = 0^{\ell_m}$.
- Output $m' = m'_1 \parallel \dots \parallel m'_{\ell_m}$.

We have the following theorems.

Theorem 5.1. *If ELWMPRF satisfies extended weak pseudorandomness against authority, then so does WMPRF.*

Theorem 5.2. *If ELWMPRF is an extraction-less watermarking PRF that satisfies QSIM-MDD security, WMPRF is a privately extractable watermarking PRF.*

Theorem 5.3. *If ELWMPRF is an extraction-less watermarking PRF that satisfies SIM-MDD security with public simulation, WMPRF is a publicly extractable watermarking PRF.*

It is clear that Theorem 5.1 holds since the evaluation algorithm of WMPRF is the same as that of ELWMPRF and extended weak pseudorandomness is insensitive to how the marking and extraction algorithms are defined. Thus, we omit a formal proof.

The proofs of Theorems 5.2 and 5.3 are almost the same. Thus, we only provide the proof for the former, and omit the proof for the latter.

Proof of Theorem 5.2. Let $\epsilon > 0$. Let \mathcal{A} be a QPT adversary attacking the unremovability of WMPRF. The description of $\text{Expt}_{\mathcal{A}, \text{WMPRF}}^{\text{rmv}}(\lambda, \epsilon)$ is as follows.

1. The challenger generates $(\text{pp}, \text{xk}) \leftarrow \text{Setup}(1^\lambda)$ and gives pp to the adversary \mathcal{A} . \mathcal{A} sends $m \in \{0, 1\}^{\ell_m}$ to the challenger. The challenger generates $(\text{prfk}, \tau) \leftarrow \text{Gen}(\text{pp})$, computes $\tilde{C} \leftarrow \text{Mark}(\text{pp}, \text{prfk}, m \parallel 0)$, and sends \tilde{C} to \mathcal{A} .
2. \mathcal{A} can access to the following oracle.
 O_{ext} : On input τ' and a quantum circuit C , it returns $\mathcal{E}\text{xtract}(\text{xk}, C, \tau', \epsilon)$.
3. Finally, the adversary outputs a quantum circuit $C_{\mathcal{A}} = (q, \mathbf{U})$.

We define D , \mathcal{P} , \mathcal{M}_D , and the three events Live, GoodExt, and BadExt in the same way as Definition 3.6.

The proof of $\Pr[\text{GoodExt}] \geq \Pr[\text{Live}] - \text{negl}(\lambda)$. $\mathcal{E}\text{xtract}$ outputs unmarked if and only if $\tilde{p}_{\ell+1} < \frac{1}{2} + \epsilon - 4\epsilon'$, that is we have $\Pr[\text{GoodExt}] = \Pr[\tilde{p}_{\ell+1} \geq \frac{1}{2} + \epsilon - 4\epsilon']$. Let p the probability obtained by applying $\text{ProjImp}(\mathcal{M}_D)$ to q . Then, we have $\Pr[\text{Live}] = \Pr[p \geq \frac{1}{2} + \epsilon]$. Let \tilde{p} be the outcome obtained if we apply $\mathcal{A}\text{PI}_{\mathcal{P}, D}^{\epsilon', \delta'}$ to q . From the property of $\mathcal{A}\text{PI}$, we have

$$\Pr[\text{Live}] = \Pr\left[p \geq \frac{1}{2} + \epsilon\right] \leq \Pr\left[\tilde{p} \geq \frac{1}{2} + \epsilon - \epsilon'\right] + \text{negl}(\lambda).$$

D and D_{τ, ℓ_m+1} are computationally indistinguishable from the QSIM-MDD security of ELWMPRF since outputs of $\text{Sim}(\text{xk}, \tau, i)$ is indistinguishable from those of D if $m[i] = 0$. This indistinguishability holds even under the existence of O_{api} . Then, from Theorem 2.11, we have

$$\Pr\left[\tilde{p} \geq \frac{1}{2} + \epsilon - \epsilon'\right] \leq \Pr\left[\tilde{p}_{\ell+1} \geq \frac{1}{2} + \epsilon - 4\epsilon'\right] + \text{negl}(\lambda) = \Pr[\text{GoodExt}] + \text{negl}(\lambda).$$

By combining the above two equations, we obtain $\Pr[\text{GoodExt}] \geq \Pr[\text{Live}] - \text{negl}(\lambda)$.

The reason D and $D_{\tau, \ell+1}$ need to be computationally indistinguishable under the existence of O_{api} to apply Theorem 2.11 is as follows. In this application of Theorem 2.11, the quantum state appeared in the statement of it is set as q contained in the quantum circuit C output by \mathcal{A} . Then, Theorem 2.11 (implicitly) requires that D and $D_{\tau, \ell+1}$ be indistinguishable for distinguishers who can construct q . To construct q , we need to execute \mathcal{A} who can access to O_{ext} in which \mathcal{API} is repeatedly executed. This is the reason D and $D_{\tau, \ell+1}$ need to be indistinguishable under the existence of O_{api} .

The proof of $\Pr[\text{BadExt}] \leq \text{negl}(\lambda)$. We define the event BadExt_i as follows for every $i \in [\ell_m]$.

BadExt _{i} : When Running $\mathcal{Extract}(\text{xk}, C_{\mathfrak{g}}, \tau^*, \epsilon)$, the following conditions hold.

- $\tilde{p}_{\ell+1} \geq \frac{1}{2} + \epsilon - 4\epsilon'$ holds.
- $m'_j = m_j$ holds for every $j \in [i-1]$.
- $\mathcal{Extract}$ exits the i -th loop or $m'_i \neq m_i$ holds.

Then, we have $\Pr[\text{BadExt}] \leq \sum_{i \in [\ell]} \Pr[\text{BadExt}_i]$. Below, we estimate $\Pr[\text{BadExt}_i]$.

We first consider the case of $m_{i-1} = 0$ and $m_i = 0$. Assume $m'_{i-1} = m_{i-1} = 0$ holds. Then, we have $\tilde{p}_{i-1} > \frac{1}{2} + \epsilon - 4i\epsilon'$. Let $\tilde{p}'_{i-1} \leftarrow \mathcal{API}_{\mathcal{P}, D_{\tau, i-1}}^{\epsilon', \delta'}(q_{i-1})$. From, the almost-projective property of \mathcal{API} , we have

$$\Pr \left[\tilde{p}'_{i-1} > \frac{1}{2} + \epsilon - 4i\epsilon' - \epsilon' \right] \geq 1 - \delta'.$$

When $m_{i-1} = 0$ and $m_i = 0$, $D_{\tau, i-1}$ and $D_{\tau, i}$ are computationally indistinguishable since both of them are computationally indistinguishable from D by the QSIM-MDD security of ELWMPRF. This indistinguishability holds under the existence of O_{api} . Thus, from Theorem 2.11, we have

$$1 - \delta' \leq \Pr \left[\tilde{p}'_{i-1} > \frac{1}{2} + \epsilon - (4i+1)\epsilon' \right] \leq \Pr \left[\tilde{p}_i > \frac{1}{2} + \epsilon - 4(i+1)\epsilon' \right] + \text{negl}(\lambda).$$

This means that $\Pr[\text{BadExt}_i] = \text{negl}(\lambda)$ in this case. Note that the reason the indistinguishability of $D_{\tau, i-1}$ and $D_{\tau, i}$ needs to hold under O_{api} is that Theorem 2.11 requires it hold for distinguishers who can construct q_{i-1} .

Next, we consider the case of $m_{i-1} = 0$ and $m_i = 1$. Assume $m'_{i-1} = m_{i-1} = 0$ holds. Then, we have $\tilde{p}_{i-1} > \frac{1}{2} + \epsilon - 4i\epsilon'$. We then define an additional distribution $D_{\tau, i}^{\text{rev}}$ as follows.

$D_{\tau, i}^{\text{rev}}$: Generate $(\gamma, x, y) \leftarrow \text{Sim}(\text{xk}, \tau, i)$. Output $(1 \oplus \gamma, x, y)$.

That is, the first bit of the output is flipped from $D_{\tau, i}$. Then, for any random coin r , we have $(\mathbf{P}_{D_{\tau, i}^{\text{rev}}(r)}, \mathbf{Q}_{D_{\tau, i}^{\text{rev}}(r)}) = (\mathbf{Q}_{D_{\tau, i}(r)}, \mathbf{P}_{D_{\tau, i}(r)})$. This is because we have $\mathbf{Q}_{b, x, y} = \mathbf{I} - \mathbf{P}_{b, x, y} = \mathbf{P}_{1 \oplus b, x, y}$ for any tuple (b, x, y) . Therefore, $\mathcal{API}_{\mathcal{P}, D_{\tau, i}^{\text{rev}}}^{\epsilon', \delta'}$ is exactly the same process as $\mathcal{API}_{\mathcal{P}^{\text{rev}}, D_{\tau, i-1}}^{\epsilon', \delta'}$. Let $\tilde{p}'_{i-1} \leftarrow \mathcal{API}_{\mathcal{P}, D_{\tau, i-1}^{\text{rev}}}^{\epsilon', \delta'}(q_{i-1})$. From, the reverse-almost-projective property of \mathcal{API} , we have

$$\Pr \left[\tilde{p}'_{i-1} < \frac{1}{2} - \epsilon + 4i\epsilon' + \epsilon' \right] \geq 1 - \delta'.$$

When $m_{i-1} = 0$ and $m_i = 1$, $D_{\tau, i-1}^{\text{rev}}$ and $D_{\tau, i}$ are computationally indistinguishable since both of them are computationally indistinguishable from the following distribution D^{rev} by the QSIM-MDD security of ELWMPRF.

D^{rev} : Generate $(\gamma, x, y) \leftarrow D$. Output $(1 \oplus \gamma, x, y)$.

This indistinguishability holds under the existence of O_{api} . Thus, from Theorem 2.11, we have

$$1 - \delta' \leq \Pr \left[\tilde{p}'_{i-1} < \frac{1}{2} - \epsilon + (4i+1)\epsilon' \right] \leq \Pr \left[\tilde{p}_i < \frac{1}{2} - \epsilon + 4(i+1)\epsilon' \right] + \text{negl}(\lambda).$$

This means that $\Pr[\text{BadExt}_i] = \text{negl}(\lambda)$ also in this case. Note that the reason the indistinguishability of $D_{\tau, i-1}^{\text{rev}}$ and $D_{\tau, i}$ needs to hold under O_{api} is that Theorem 2.11 requires it hold for distinguishers who can construct q_{i-1} .

Similarly, we can prove that $\Pr[\text{BadExt}_i] = \text{negl}(\lambda)$ holds in the case of $(m_{i-1}, m_i) = (1, 0)$ and $(m_{i-1}, m_i) = (1, 1)$.

Overall, we see that $\Pr[\text{BadExt}] = \text{negl}(\lambda)$ holds in all cases. ■

6 Extraction-Less Watermarking PRF from LWE

We present an extraction-less watermarking PRF, denoted by PRF_{cprf} , whose message space is $\{0, 1\}^{\ell_m}$ with domain $\{0, 1\}^n$ and range $\{0, 1\}^m$. We use the following tools, which can be instantiated with the QLWE assumption (See Theorems 2.24, 2.27 and 2.30):

- Private CPRF $\text{CPRF} = (\text{CPRF.Setup}, \text{CPRF.Eval}, \text{CPRF.Constrain}, \text{CPRF.CEval})$. For ease of notation, we denote CPRF evaluation circuit $\text{CPRF.Eval}(\text{msk}, \cdot)$ and constrained evaluation circuits $\text{CPRF.CEval}(\text{sk}_f, \cdot)$ by $G : \{0, 1\}^n \rightarrow \{0, 1\}^m$ and $G_{\notin \mathcal{V}} : \{0, 1\}^n \rightarrow \{0, 1\}^m$, respectively, where $x \in \mathcal{V}$ iff $f(x) = 1$.
- SKE scheme $\text{SKE} = (\text{SKE.Gen}, \text{SKE.Enc}, \text{SKE.Dec})$. The plaintext space and ciphertext space of SKE are $\{0, 1\}^{\ell_{\text{ske}}}$ and $\{0, 1\}^n$, respectively, where $\ell_{\text{ske}} = \log \ell_m + 1$.
- PKE scheme $\text{PKE} = (\text{Gen}, \text{Enc}, \text{Dec})$. The plaintext space of PKE is $\{0, 1\}^{2\lambda}$.

Construction overview. We already explained the high-level idea for how to realize extraction-less watermarking PRFs in Section 1.3. However, the construction of PRF_{cprf} requires some additional efforts. Thus, before providing the actual construction, we provide a high-level overview of PRF_{cprf} .

Recall that letting $\tilde{C} \leftarrow \text{Mark}(\text{pp}, \text{prfk}, m)$ and $(\gamma^*, x^*, y^*) \leftarrow \text{Sim}(\text{xk}, \tau, i^*)$, we have to design Sim and \tilde{C} so that

- If $\gamma = m[i^*]$, $\tilde{C}(x^*)$ outputs a value different from y^* .
- If $\gamma \neq m[i^*]$, $\tilde{C}(x^*)$ outputs y^* .

In the token-based construction idea, we achieve these conditions by setting x^* as an encryption of $y^* \| i^* \| \gamma^*$ and designing \tilde{C} as a token such that it outputs y^* if the input is decryptable and $\gamma^* \neq m[i^*]$ holds for the decrypted value $y^* \| i^* \| \gamma^*$, and otherwise behaves as the original evaluation circuit. However, in PRF_{cprf} , we use a constrained evaluation circuit of CPRF as \tilde{C} , and thus we cannot program output values for specific inputs. Intuitively, it seems that Sim needs to use the original PRF key prfk to achieve the above two conditions.

To solve the issue, we adopt the idea used by Quach et al. [QWZ18]. In PRF_{cprf} , the setup algorithm Setup generates $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$ of PKE, and sets $\text{pp} = \text{pk}$ and $\text{xk} = \text{sk}$. Then, the PRF key generation algorithm is given pk , generates $G \leftarrow \text{CPRF.Setup}(1^\lambda, 1^\kappa)$ along with $\text{ske.k} \leftarrow \text{SKE.Gen}(1^\lambda)$, and sets the public tag τ as an encryption of $(G, \text{ske.k})$ under pk . The evaluation algorithm of PRF_{cprf} is simply that of CPRF.

Now, we explain how to design Sim and $\tilde{C} \leftarrow \text{Mark}(\text{pp}, \text{prfk}, m)$ to satisfy the above two conditions. Given $\text{xk} = \text{sk}$, $\tau = \text{Enc}(\text{pk}, \text{prfk})$ and i , Sim is able to extract $\text{prfk} = (G, \text{ske.k})$. Then, Sim generates $\gamma \leftarrow \{0, 1\}$ and sets $x \leftarrow \text{SKE.Enc}(\text{ske.k}, i \| \gamma)$ and $y \leftarrow G(x)$. We set \tilde{C} as a constrained version of G for a circuit D that outputs 1 if the input x is decryptable by ske.k and $\gamma = m[i]$ holds for decrypted value

$i||\gamma$, and otherwise outputs 0. For an input x , the constrained version of G outputs the correct output $G(x)$ if and only if $D(x) = 0$. We can check that PRF_{cprf} satisfies the above two conditions.

The above construction does not satisfy extended weak pseudorandomness against authority since the authority can extract the original CPRF key G by $xk = sk$. However, this problem can be fixed by constraining G . We see that Sim needs to evaluate G for valid ciphertexts of SKE. Thus, to implement the above mechanism, it is sufficient to set the public tag τ as an encryption of ske.k and a constrained version of G for a circuit D_{auth} that output 0 if and only if the input is decryptable by ske.k . Then, the authority can only extract such a constrained key. By requiring sparseness for SKE, the constrained key cannot be used to break the pseudorandomness of PRF_{cprf} for random inputs. This means that PRF_{cprf} satisfies extended weak pseudorandomness against an authority. Note that we only need a single-key CPRF for PRF_{cprf} since either a user or the authority (not both) is a malicious entity in security games.

The description of PRF_{cprf} is as follows.

$\text{Setup}(1^\lambda)$:

- Generate $(pk, sk) \leftarrow \text{Gen}(1^\lambda)$.
- Output $(pp, xk) := (pk, sk)$.

$\text{Gen}(pp)$:

- Parse $pp = pk$.
- Generate $G \leftarrow \text{CPRF.Setup}(1^\lambda, 1^\kappa)$. In our construction, κ is the size of circuit $D[\text{ske.k}, m]$ described in Figure 5, which depends on ℓ_m (and λ).
- Generate $\text{ske.k} \leftarrow \text{SKE.Gen}(1^\lambda)$.
- Construct a circuit $D_{\text{auth}}[\text{ske.k}]$ described in Figure 4.
- Compute $G_{\notin \mathcal{V}_{\text{auth}}} := \text{CPRF.Constrain}(G, D_{\text{auth}}[\text{ske.k}])$, where $\mathcal{V}_{\text{auth}} \subset \{0, 1\}^n$ is a set such that $x \in \mathcal{V}_{\text{auth}}$ iff $D_{\text{auth}}[\text{ske.k}](x) = 1$.
- Output $\text{prfk} := (G, \text{ske.k})$ and $\tau \leftarrow \text{Enc}(pk, (G_{\notin \mathcal{V}_{\text{auth}}}, \text{ske.k}))$.

$\text{Eval}(\text{prfk}, x \in \{0, 1\}^n)$: Recall that G is a keyed CPRF evaluation circuit.

- Parse $\text{prfk} = (G, \text{ske.k})$.
- Output $y := G(x)$.

$\text{Mark}(pp, \text{prfk}, m)$:

- Parse $pp = pk$ and $\text{prfk} = (G, \text{ske.k})$.
- Construct a circuit $D[\text{ske.k}, m]$ described in Figure 5.
- Compute $G_{\notin \mathcal{V}} \leftarrow \text{CPRF.Constrain}(G, D[\text{ske.k}, m])$, where $\mathcal{V} \subset \{0, 1\}^n$ is a set such that $x \in \mathcal{V}$ iff $D[\text{ske.k}, m](x) = 1$.
- Output $\tilde{C} = G_{\notin \mathcal{V}}$.

$\text{Sim}(xk, \tau, i)$:

- Parse $xk = sk$.
- Compute $(G_{\notin \mathcal{V}_{\text{auth}}}, \text{ske.k}) \leftarrow \text{Dec}(sk, \tau)$.
- Choose $\gamma \leftarrow \{0, 1\}$.
- Compute $x \leftarrow \text{SKE.Enc}(\text{ske.k}, i||\gamma)$ and $y \leftarrow G_{\notin \mathcal{V}_{\text{auth}}}(x)$.
- Output (γ, x, y) .

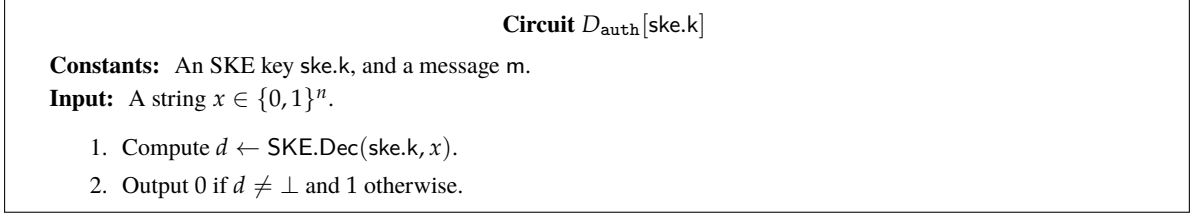


Figure 4: The description of D_{auth}

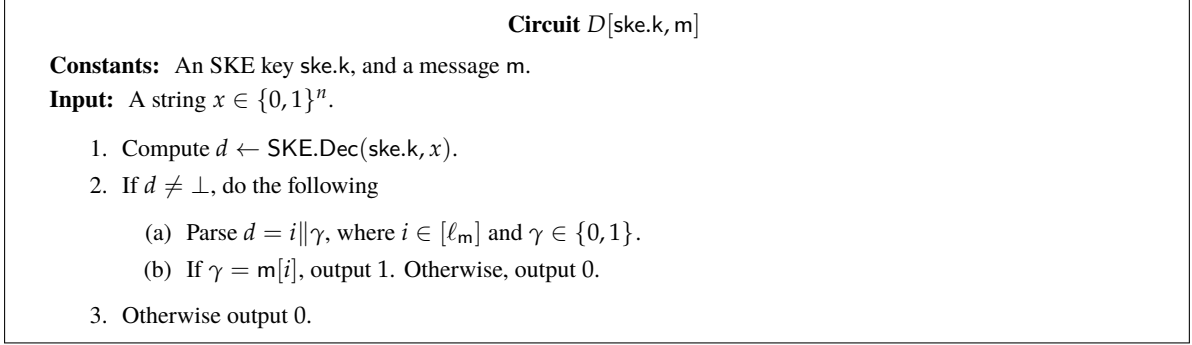


Figure 5: The description of D

The evaluation correctness of PRF_{cprf} follows from the sparseness of SKE and the correctness of CPRF. For the security of PRF_{cprf} , we have the following theorems.

Theorem 6.1. *SKE is a secure SKE scheme with pseudorandom ciphertext, CPRF is a selectively single-key private CPRF, PKE is a CCA secure PKE scheme, then PRF_{cprf} is an extraction-less watermarking PRF satisfying SIM-MDD security.*

Theorem 6.2. *If CPRF is a selective single-key private CPRF, PRF_{cprf} satisfies extended weak pseudo-randomness.*

SIM-MDD security. First, we prove the SIM-MDD security of PRF_{cprf} .

Proof of Theorem 6.1. We define a sequence of hybrid games to prove the theorem.

Hyb₀: This is the same as the case $\text{coin} = 1$ in $\text{Exp}_{i^*, \mathcal{A}, \text{PRF}_{\text{cprf}}}^{\text{sim-mdd}}(\lambda)$. In this game, \mathcal{A} is given $\tau \leftarrow \text{Enc}(\text{pk}, (G_{\notin \mathcal{V}_{\text{auth}}}, \text{ske.k}))$ and $G_{\notin \mathcal{V}} \leftarrow \text{CPRF.Constrain}(G, D[\text{ske.k}, m])$ as a public tag and a marked circuit. After τ and $G_{\notin \mathcal{V}}$ are given, \mathcal{A} can access to O_{sim} . Finally, after finishing the access to O_{sim} , \mathcal{A} is given (γ^*, x^*, y^*) as the challenge tuple and outputs $\text{coin}' \in \{0, 1\}$, where $\gamma^* \leftarrow \{0, 1\}$, $x^* \leftarrow \text{SKE.Enc}(\text{ske.k}, i^* \parallel \gamma^*)$, and $y^* \leftarrow G_{\notin \mathcal{V}_{\text{auth}}}(x^*)$.

Hyb₁: This is the same as Hyb₀ except for the following two changes. First, G is used instead of $G_{\notin \mathcal{V}_{\text{auth}}}$ when generating the challenge tuple (γ^*, x^*, y^*) . Second, we change the behavior of O_{sim} as follows. When \mathcal{A} sends τ' and i' to O_{sim} , if $\tau' = \tau$, O_{sim} performs the remaining procedures by using $(G, \text{ske.k})$ (without decrypting $\tau' = \tau$).

Hyb₂: This is the same as Hyb₁ except that we use $\tau \leftarrow \text{Enc}(\text{pk}, 0^{2\lambda})$ instead of $\tau \leftarrow \text{Enc}(\text{pk}, (G_{\notin \mathcal{V}_{\text{auth}}}, \text{ske.k}))$.

Hyb₃: This is the same as Hyb₂ except that if $m[i^*] = \gamma^*$, we use $y^* \leftarrow \{0, 1\}^m$ instead of $y^* \leftarrow G(x^*)$.

Hyb₄: This is the same as Hyb₃ except that we use a simulated $(\text{st}_{\text{Sim}}, \widehat{G}) \leftarrow \text{CPRF.Sim}_1(1^\kappa, 1^\lambda)$ instead of $G_{\notin \mathcal{V}} \leftarrow \text{CPRF.Constrain}(G, D[\text{ske.k}, m])$ for the challenge marked circuit. Also, if $m[i^*] \neq \gamma^*$,

the challenger computes $y^* \leftarrow \text{Sim}_2(\text{st}_{\text{sim}}, x^*, 0)$. In addition, we also change the behavior of O_{sim} as follows. Given τ' and i' , if $\tau' \neq \tau$, O_{sim} answers in the same way as Hyb_3 . Otherwise, it returns (γ, x, y) , where $\gamma \leftarrow \{0, 1\}$, $x \leftarrow \text{SKE.Enc}(\text{ske.k}, i' \| \gamma)$, and $y \leftarrow \text{Sim}_2(\text{st}_{\text{sim}}, x, 1)$ if $m[i'] = \gamma$ and $y \leftarrow \text{Sim}_2(\text{st}_{\text{sim}}, x, 0)$ otherwise.

Hyb₅: This is the same as Hyb_4 except that we use $x^* \leftarrow \{0, 1\}^n$ instead of $x^* \leftarrow \text{SKE.Enc}(\text{ske.k}, i^* \| \gamma^*)$.

Hyb₆: We undo the change at Hyb_4 .

Hyb₇: We undo the change at Hyb_2 .

Hyb₈: We undo the change at Hyb_1 . This is the same as the case $\text{coin} = 0$ in $\text{Exp}_{i^*, \mathcal{A}, \text{PRF}_{\text{cprf}}}^{\text{sim-mdd}}(\lambda)$.

Proposition 6.3. *If CPRF, SKE, and PKE are correct, it holds that $|\Pr[\text{Hyb}_0 = 1] - \Pr[\text{Hyb}_1 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Proposition 6.3. For the first change, $x^* \notin \mathcal{V}_{\text{auth}}$ holds since $\perp \neq \text{SKE.Dec}(\text{ske.k}, x^*)$ from the correctness of SKE. Then, from the correctness of CPRF, we have $G_{\notin \mathcal{V}_{\text{auth}}}(x^*) = G(x^*)$, and thus the first change does not affect the view of \mathcal{A} . For the second change, from the correctness of PKE, $(G_{\notin \mathcal{V}_{\text{auth}}}, \text{ske.k}) \leftarrow \text{Dec}(\text{sk}, \tau')$ if $\tau' = \tau$. Then, similarly to the first change, we can see that the second change does not affect the view of \mathcal{A} . ■

Proposition 6.4. *If PKE is CCA secure, it holds that $|\Pr[\text{Hyb}_1 = 1] - \Pr[\text{Hyb}_2 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Proposition 6.4. We construct an algorithm \mathcal{B} that breaks CCA security of PKE by using \mathcal{A} .

B: \mathcal{B} is given pk from the challenger. \mathcal{B} generates $\text{ske.k} \leftarrow \text{SKE.Gen}(1^\lambda)$, and $G \leftarrow \text{CPRF.Setup}(1^\lambda)$, and sets $\text{pp} := \text{pk}$. \mathcal{B} sends pp to \mathcal{A} and obtain m from \mathcal{A} . \mathcal{B} then generate $G_{\notin \mathcal{V}_{\text{auth}}} := \text{CPRF.Constrain}(G, D_{\text{auth}}[\text{ske.k}])$, sets $(m_0, m_1) := ((G_{\notin \mathcal{V}_{\text{auth}}}, \text{ske.k}), 0^{2\lambda})$ as the challenge plaintext of the CCA game and receives τ from its challenger. \mathcal{B} also constructs $D[\text{ske.k}, m]$, generates $G_{\notin \mathcal{V}} \leftarrow \text{CPRF.Constrain}(G, D[\text{ske.k}, m])$, and sends τ and $G_{\notin \mathcal{V}}$ to \mathcal{A} as the challenge public tag and marked circuit.

O_{sim} : When \mathcal{A} sends τ' and i' to O_{sim} , \mathcal{B} simulates the answer by using G , ske.k , and the decryption oracle $\text{Dec}(\text{sk}, \cdot)$.

After finishing \mathcal{A} 's oracle access to O_{sim} , \mathcal{B} chooses $\gamma^* \leftarrow \{0, 1\}$, generates $x^* \leftarrow \text{SKE.Enc}(\text{ske.k}, i^* \| \gamma^*)$ and $y^* \leftarrow G_{\notin \mathcal{V}_{\text{auth}}}(x^*) = G(x^*)$, and sends (γ^*, x^*, y^*) to \mathcal{A} . Note that $G_{\notin \mathcal{V}_{\text{auth}}}(x^*) = G(x^*)$ holds since $\perp \neq \text{SKE.Dec}(\text{ske.k}, x^*)$ and thus $x^* \notin \mathcal{V}_{\text{auth}}$.

Finally, when \mathcal{A} terminates with output coin' , \mathcal{B} outputs coin' and terminates.

\mathcal{B} perfectly simulates if Hyb_1 if $\tau \leftarrow \text{Enc}(\text{pk}, (G_{\notin \mathcal{V}_{\text{auth}}}, \text{ske.k}))$, and Hyb_2 if $\tau \leftarrow \text{Enc}(\text{pk}, 0^{2\lambda})$. This completes the proof. ■

Proposition 6.5. *If CPRF satisfies selective pseudorandomness, it holds that*

$$|\Pr[\text{Hyb}_2 = 1] - \Pr[\text{Hyb}_3 = 1]| \leq \text{negl}(\lambda).$$

Proof of Proposition 6.5. We use selective single-key pseudorandomness of G . We construct an algorithm \mathcal{B} that breaks the selective single-key pseudorandomness of G by using \mathcal{A} .

B: \mathcal{B} generates $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$, $\text{ske.k} \leftarrow \text{SKE.Gen}(1^\lambda)$, and $\tau \leftarrow \text{Enc}(\text{pk}, 0^{2\lambda})$, sets and sends $\text{pp} := \text{pk}$ to \mathcal{A} , and obtains m from \mathcal{A} . \mathcal{B} constructs $D[\text{ske.k}, m]$, sends $D[\text{ske.k}, m]$ to its challenger, and receives $G_{\notin \mathcal{V}}$. \mathcal{B} sends τ and $G_{\notin \mathcal{V}}$ to \mathcal{A} as the challenge public tag and marked circuit.

O_{sim} : When \mathcal{A} sends τ' and $i' \in [\ell_m]$ to O_{sim} , if $\tau' \neq \tau$, \mathcal{B} computes $(G', \text{ske.k}') \leftarrow \text{Dec}(\text{sk}, \tau')$, and computes and returns the answer (γ, x, y) by using $(G', \text{ske.k}')$. If $\tau' = \tau$, \mathcal{B} returns (γ, x, y) computed as follows. \mathcal{B} chooses $\gamma \leftarrow \{0, 1\}$, and generates $x \leftarrow \text{SKE.Enc}(\text{ske.k}, i' \| \gamma)$. \mathcal{B} finally sends x to its PRF evaluation oracle and receives $y \leftarrow G(x)$.

After finishing \mathcal{A} 's oracle access to O_{sim} , \mathcal{B} sends (γ^*, x^*, y^*) computed as follows to \mathcal{A} . \mathcal{B} first chooses $\gamma^* \leftarrow \{0, 1\}$ and generates $x^* \leftarrow \text{SKE.Enc}(\text{ske.k}, i^* \| \gamma^*)$. If $m[i^*] = \gamma^*$, \mathcal{B} sends x^* to its challenge oracle and receives y^* . If $m[i^*] \neq \gamma^*$, \mathcal{B} sends x^* to its PRF evaluation oracle and receives y^* .

Finally, when \mathcal{A} terminates with output coin' , \mathcal{B} outputs coin' and terminates.

\mathcal{B} perfectly simulates Hyb_2 if the challenge oracle returns $y^* = G(x^*)$, and Hyb_3 if it returns $y^* \leftarrow \{0, 1\}^m$. Note that in these games, $x^* \leftarrow \text{SKE.Enc}(\text{ske.k}, i^* \| \gamma^*)$, and thus if $m[i^*] = \gamma^*$, we have $x^* \in \mathcal{V}(D[\text{ske.k}, m](x^*) = 1)$. This completes the proof. ■

Proposition 6.6. *If CPRF satisfies selective single-key privacy, it holds that*

$$|\Pr[\text{Hyb}_3 = 1] - \Pr[\text{Hyb}_4 = 1]| \leq \text{negl}(\lambda).$$

Proof of Proposition 6.6. We use selective single-key privacy of G . We construct an algorithm \mathcal{B} that breaks the selective privacy of G by using \mathcal{A} .

\mathcal{B} : \mathcal{B} generates $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$, $\text{ske.k} \leftarrow \text{SKE.Gen}(1^\lambda)$, and $\tau \leftarrow \text{Enc}(\text{pk}, 0^{2\lambda})$, sends $\text{pp} := \text{pk}$ to \mathcal{A} , and obtains m from \mathcal{A} . \mathcal{B} constructs $D[\text{ske.k}, m]$, sends $D[\text{ske.k}, m]$ to its challenger, and receives G^* . \mathcal{B} sends τ and G^* to \mathcal{A} as the challenge public tag and marked circuit.

O_{sim} : When \mathcal{A} sends τ' and $i' \in [\ell_m]$ to O_{sim} , if $\tau' \neq \tau$, \mathcal{B} computes $(G', \text{ske.k}') \leftarrow \text{Dec}(\text{sk}, \tau')$ and returns the answer (γ, x, y) computed by using $(G', \text{ske.k}')$. If $\tau' = \tau$, \mathcal{B} returns the answer (γ, x, y) computed as follows. \mathcal{B} chooses $\gamma \leftarrow \{0, 1\}$, and generates $x \leftarrow \text{SKE.Enc}(\text{ske.k}, i' \| \gamma)$. \mathcal{B} sends x to its oracle and receives y .

After finishing \mathcal{A} 's oracle access to O_{sim} , \mathcal{B} sends (γ^*, x^*, y^*) computed as follows to \mathcal{A} . \mathcal{B} chooses $\gamma^* \leftarrow \{0, 1\}$ and generates $x^* \leftarrow \text{SKE.Enc}(\text{ske.k}, i^* \| \gamma^*)$. If $m[i^*] = \gamma^*$, \mathcal{B} chooses $y^* \leftarrow \{0, 1\}^m$. If $m[i^*] \neq \gamma^*$, \mathcal{B} sends x^* to its oracle and receives y^* .

Finally, when \mathcal{A} terminates with output coin' , \mathcal{B} outputs coin' and terminates.

\mathcal{B} perfectly simulates Hyb_3 if $G^* = \text{CPRF.Constrain}(G, D[\text{ske.k}, m])$ and \mathcal{B} has access to $G(\cdot)$, and Hyb_4 if $G^* = \text{Sim}_1(1^\kappa, 1^\lambda)$ and \mathcal{B} has access to $\text{Sim}_2(\text{st}_{\text{Sim}}, \cdot, D[\text{ske.k}, m](\cdot))$. This completes the proof. ■

Proposition 6.7. *If SKE satisfies ciphertext pseudorandomness, it holds that*

$$|\Pr[\text{Hyb}_4 = 1] - \Pr[\text{Hyb}_5 = 1]| \leq \text{negl}(\lambda).$$

Proof of Proposition 6.7. We construct an algorithm \mathcal{B} that breaks the ciphertext pseudorandomness of SKE by using \mathcal{A} .

\mathcal{B} : \mathcal{B} generates $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$ and sends $\text{pp} := \text{pk}$ to \mathcal{A} , and obtains m from \mathcal{A} . \mathcal{B} then generates $\tau \leftarrow \text{Enc}(\text{pk}, 0^{2\lambda})$ and $(\text{st}_{\text{Sim}}, \widehat{G}) \leftarrow \text{Sim}_1(1^\kappa, 1^\lambda)$, and sends τ and \widehat{G} to \mathcal{A} as the challenge public tag and marked circuit.

O_{sim} : When \mathcal{A} sends τ' and $i' \in [\ell_m]$ to O_{sim} , if $\tau' \neq \tau$, \mathcal{B} computes $(G', \text{ske.k}') \leftarrow \text{Dec}(\text{sk}, \tau)$ and returns the answer (γ, x, y) computed by using $(G', \text{ske.k}')$. If $\tau' = \tau$, \mathcal{B} returns the answer (γ, x, y) computed as follows. \mathcal{B} chooses $\gamma \leftarrow \{0, 1\}$, sends $i' \parallel \gamma$ to its encryption oracle, and receives $x \leftarrow \text{SKE.Enc}(\text{ske.k}, i' \parallel \gamma)$. \mathcal{B} computes $y \leftarrow \text{Sim}_2(\text{st}_{\text{sim}}, x, 1)$ if $\gamma = m[i']$ and $y \leftarrow \text{Sim}_2(\text{st}_{\text{sim}}, x, 0)$ otherwise.

After finishing \mathcal{A} 's oracle access to O_{sim} , \mathcal{B} sends (γ^*, x^*, y^*) computed as follows to \mathcal{A} . \mathcal{B} chooses $\gamma^* \leftarrow \{0, 1\}$, sends $i^* \parallel \gamma^*$ to its challenger as the challenge plaintext, and receives x^* . \mathcal{B} generates $y^* \leftarrow \{0, 1\}^m$ if $m[i^*] = \gamma^*$ and $y^* \leftarrow \text{Sim}_2(\text{st}_{\text{sim}}, x, 0)$ otherwise.

Finally, when \mathcal{A} terminates with output coin' , \mathcal{B} outputs coin' and terminates.

\mathcal{B} perfectly simulates Hyb_4 if $x^* \leftarrow \text{SKE.Enc}(\text{ske.k}, i^* \parallel \gamma^*)$, and Hyb_5 if $x^* \leftarrow \{0, 1\}^n$. This completes the proof. ■

Proposition 6.8. *If CPRF satisfies selective single-key privacy, it holds that*

$$|\Pr[\text{Hyb}_5 = 1] - \Pr[\text{Hyb}_6 = 1]| \leq \text{negl}(\lambda).$$

Proof of Proposition 6.8. This proof is almost the same as that of Proposition 6.6. ■

Proposition 6.9. *If PKE is CCA secure, it holds that $|\Pr[\text{Hyb}_6 = 1] - \Pr[\text{Hyb}_7 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Proposition 6.9. This proof is almost the same as that of Proposition 6.4. ■

Proposition 6.10. *If CPRF, SKE, and PKE are correct, it holds that $|\Pr[\text{Hyb}_7 = 1] - \Pr[\text{Hyb}_8 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Proposition 6.10. This proof is almost the same as that of Proposition 6.3. ■

By Propositions 6.3 to 6.10, we complete the proof of Theorem 6.1. ■

Extended weak pseudorandomness. Next, we prove the extended pseudorandomness of PRF_{cprf} .

Proof of Theorem 6.2. Let \mathcal{A} be an adversary attacking the extended weak pseudorandomness of PRF_{cprf} . We construct \mathcal{B} that attacks the selective single-key pseudorandomness of CPRF.

\mathcal{B} : Given pp from \mathcal{A} , \mathcal{B} first generates $\text{ske.k} \leftarrow \text{SKE.Gen}(1^\lambda)$, sends $D_{\text{auth}}[\text{ske.k}]$ to its challenger, and obtains G^* . \mathcal{B} sets $\text{pp} := \text{pk}$, generates $\tau \leftarrow \text{Enc}(\text{pk}, (G^*, \text{ske.k}))$, and sends it to \mathcal{A} . \mathcal{B} answers \mathcal{A} 's queries as follows.

O_{wprf} : When this is invoked (no input), \mathcal{B} generates $a \leftarrow \{0, 1\}^n$, sends it to its evaluation oracle, and obtains b . Then, \mathcal{B} returns (a, b) to \mathcal{A} .

O_{chall} : When this is invoked (no input), \mathcal{B} generates $a^* \leftarrow \{0, 1\}^n$, outputs a^* as its challenge input, and obtains b^* . \mathcal{B} returns (a^*, b^*) to \mathcal{A} . Note that this oracle is invoked only once.

When \mathcal{A} terminates with output b' , \mathcal{B} outputs b' and terminates.

Due to the sparseness of SKE, without negligible probability, we have $\text{SKE.Dec}(\text{ske.k}, a^*) = \perp$ thus $D_{\text{auth}}[\text{ske.k}](a^*) = 1$, and a generated when answering to a query to O_{wprf} is different from a^* . Therefore, without negligible probability, \mathcal{B} is a valid adversary against the selective single-key pseudorandomness of CPRF. When \mathcal{B} is valid, we see that the advantage of \mathcal{B} is the same as that of \mathcal{A} . This completes the proof. ■

7 Extraction-Less Watermarking PRF with Public Simulation from IO

We construct an extraction-less watermarking PRF satisfying SIM-MDD security with public simulation. We first introduce a tool.

7.1 Puncturable Encryption, Revisited

Cohen et al. [CHN⁺18] introduced the notion of puncturable encryption (PE). They used a PE scheme as a crucial building block to construct a publicly extractable watermarking PRF against classical adversaries. We also use a PE scheme to construct an extraction-less watermarking PRF with public simulation (against quantum adversaries). However, we find that the original PE definition is not sufficient for proving unremovability (and our purpose) since there is a subtle issue in the security proof by Cohen et al. [CHN⁺18]. However, we can fix the issue since their PE scheme satisfies a stronger security notion than what they proved. Thus, we introduce a stronger security notion for PE in this section.

The syntax of PE is almost the same as that of the original PE.

Definition 7.1 (Puncturable Encryption (Syntax)). A puncturable encryption (PE) scheme PE for a plaintext space $\mathcal{P} = \{0, 1\}^{\ell_p}$ is a triple of PPT algorithms $(\text{Gen}, \text{Puncture}, \text{Enc})$ and a deterministic algorithm Dec. The ciphertext space will be $\{0, 1\}^{\ell_{ct}}$ where $\ell_{ct} = \text{poly}(\lambda, \ell_p)$.

$\text{Gen}(1^\lambda) \rightarrow (\text{ek}, \text{dk})$: The key generation algorithm takes as input the security parameter 1^λ and outputs an encryption key ek and a decryption key dk.

$\text{Puncture}(\text{dk}, \{c^*\}) \rightarrow \text{dk}_{\neq c^*}$: The puncturing algorithm takes as input dk and a string $c^* \in \{0, 1\}^{\ell_{ct}}$, and outputs a ‘‘punctured’’ decryption key $\text{dk}_{\neq c^*}$.

$\text{Enc}(\text{ek}, m) \rightarrow c$: The encryption algorithm takes as input ek and a plaintext $m \in \{0, 1\}^{\ell_p}$, and outputs a ciphertext c in $\{0, 1\}^{\ell_{ct}}$.

$\text{Dec}(\text{dk}', c') \rightarrow m' \text{ or } \perp$: The decryption algorithm takes a possibly punctured decryption key dk' and a string $c' \in \{0, 1\}^{\ell_{ct}}$. It outputs a plaintext m' or the special symbol \perp .

There are four security requirements on PE. Three of those are the same as those in the original PE security. The difference is ciphertext pseudorandomness.

Definition 7.2 (Puncturable Encryption Security). A PE scheme $\text{PE} = (\text{Gen}, \text{Puncture}, \text{Enc}, \text{Dec})$ with plaintext space $\mathcal{P} = \{0, 1\}^{\ell_p}$ and ciphertext space $\mathcal{C} = \{0, 1\}^{\ell_{ct}}$ is required to satisfy the following properties.

Correctness: We require that for all plaintext $m \in \mathcal{P}$ and $(\text{ek}, \text{dk}) \leftarrow \text{Gen}(1^\lambda)$, it holds that $\text{Dec}(\text{dk}, \text{Enc}(\text{ek}, m)) = m$.

Punctured Correctness: We require the same to hold for punctured keys. For all possible keys $(\text{ek}, \text{dk}) \leftarrow \text{Gen}(1^\lambda)$, all string $c^* \in \mathcal{C}$, all punctured keys $\text{dk}_{\neq c^*} \leftarrow \text{Puncture}(\text{dk}, \{c^*\})$, and all potential ciphertexts $c \in \mathcal{C} \setminus \{c^*\}$:

$$\text{Dec}(\text{dk}, c) = \text{Dec}(\text{dk}_{\neq c^*}, c).$$

Sparseness: We also require that most strings are not valid ciphertexts:

$$\Pr \left[\text{Dec}(\text{dk}, c) \neq \perp \mid (\text{ek}, \text{dk}) \leftarrow \text{Gen}(1^\lambda), c \leftarrow \{0, 1\}^{\ell_{ct}} \right] \leq \text{negl}(\lambda).$$

Ciphertext Pseudorandomness: We require that PE has strong ciphertext pseudorandomness defined in Definition 7.3.

Definition 7.3 (Strong Ciphertext Pseudorandomness). We define the following experiment $\text{Exp}_{\mathcal{A}}^{\text{s-cpr}}(\lambda)$.

1. \mathcal{A} sends a message $m^* \in \mathcal{P} = \{0, 1\}^{\ell_p}$ to the challenger.
2. The challenger does the following:
 - Generate $(ek, dk) \leftarrow \text{Gen}(1^\lambda)$
 - Compute a ciphertext $c^* \leftarrow \text{Enc}(ek, m^*)$.
 - Choose $r^* \leftarrow \mathcal{C} = \{0, 1\}^{\ell_{ct}}$.
 - Choose $\text{coin} \leftarrow \{0, 1\}$ and set $x_0 := c^*$ and $x_1 := r^*$.
 - Generate a punctured key $dk_{\neq x_{\text{coin}}} \leftarrow \text{Puncture}(dk, \{x_{\text{coin}}\})$
 - Send $(x_{\text{coin}}, ek, dk_{\neq x_{\text{coin}}})$ to \mathcal{A} :
3. \mathcal{A} outputs coin^* and the experiment outputs 1 if $\text{coin} = \text{coin}^*$; otherwise 0.

We say that PE has strong ciphertext pseudorandomness if for every QPT adversary \mathcal{A} , it holds that

$$\text{Adv}_{\mathcal{A}}^{\text{s-cpr}}(\lambda) := 2 \left| \Pr[\text{Exp}_{\mathcal{A}}^{\text{s-cpr}}(\lambda) = 1] - \frac{1}{2} \right| \leq \text{negl}(\lambda).$$

Remark 7.4 (Difference from the original PE). In the original PE definition, Puncture takes two strings $\{c_0, c_1\} \subset \{0, 1\}^{\ell_{ct}}$ and outputs a punctured decryption key $dk_{\notin \{c_0, c_1\}}$ and punctured correctness is accordingly defined.

In the original ciphertext pseudorandomness (described in Appendix B.4), a punctured decryption key is punctured at both c^* and r^* . That is, the information about m^* remains in $dk_{\notin \{c^*, r^*\}}$ for $\text{coin} \in \{0, 1\}$. This is an issue for our purpose (and the proof by Cohen et al. [CHN⁺18]). Thus, we introduce the strong ciphertext pseudorandomness, where the information about m^* disappears in the case $\text{coin} = 1$ since the punctured decryption key is $dk_{\neq r^*}$ when $\text{coin} = 1$.

In fact, the PE scheme PE by Cohen et al. [CHN⁺18] satisfies strong ciphertext pseudorandomness (and thus, we can also fix the issue in the proof by Cohen et al.¹²).

Theorem 7.5. *If there exists secure IO for circuits and the QLWE assumption holds, there exists secure PE that satisfies strong ciphertext pseudorandomness.*

We prove this theorem in Appendix B.

7.2 Construction of Extraction-less Watermarking PRF with Public Simulation

We describe our extraction-less watermarking PRF PRF_{io} for message space $\{0, 1\}^{\ell_m}$ with domain $\{0, 1\}^{\ell_{in}}$ and range $\{0, 1\}^{\ell_{out}}$ below. We use the following tools:

- PPRF $\text{PRF} = \text{PRF}(\text{Gen}, \text{Eval}, \text{Puncture})$. We denote a PRF evaluation circuit $\text{PRF.Eval}_{\text{prfk}}(\cdot)$ by $F : \{0, 1\}^{\ell_{in}} \rightarrow \{0, 1\}^{\ell_{out}}$, a PRF evaluation circuit with punctured key $\text{PRF.Eval}_{\text{prfk}_{\neq x}}(\cdot)$ by $F_{\neq x}$ (that is, we omit prfk and simply write $F(\cdot)$ instead of $F_{\text{prfk}}(\cdot)$) for ease of notations.
- PE scheme $\text{PE} = \text{PE}(\text{Gen}, \text{Puncture}, \text{Enc}, \text{Dec})$. The plaintext and ciphertext space of PE are $\{0, 1\}^{\ell_{pt}}$ and $\{0, 1\}^{\ell_{ct}}$, respectively, where $\ell_{pt} = \ell + \log \ell_m + 1$ and $\ell_{in} := \ell_{ct}$ ($\ell_{ct} = \text{poly}(\ell, \log \ell_m)$).
- Indistinguishability obfuscator $i\mathcal{O}$.
- PRG $\text{PRG} : \{0, 1\}^{\ell} \rightarrow \{0, 1\}^{\ell_{out}}$.

¹²See Appendix B.4 for the detail of the issue.

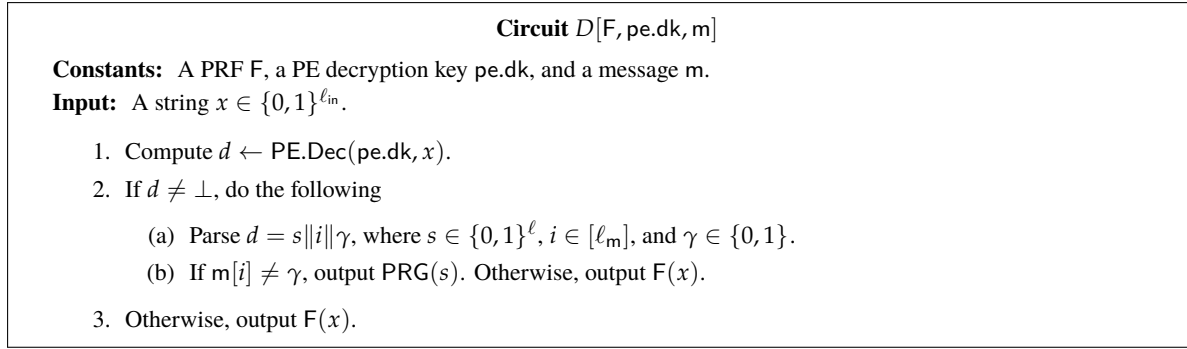


Figure 6: The description of D

$\text{Setup}(1^\lambda)$:

- Output $(\text{pp}, \text{xk}) := (\perp, \perp)$.

$\text{Gen}(\text{pp})$:

- Parse $\text{pp} = \perp$.
- Compute $F \leftarrow \text{PRF.Gen}(1^\lambda)$.
- Generate $(\text{pe.ek}, \text{pe.dk}) \leftarrow \text{PE.Gen}(1^\lambda)$.
- Output $\text{prfk} := (F, \text{pe.dk})$ and $\tau := \text{pe.ek}$.

$\text{Eval}(\text{prfk}, x \in \{0, 1\}^{\ell_{\text{in}}})$:

- Parse $\text{prfk} = (F, \text{pe.dk})$.
- Compute and output $y \leftarrow F(x)$.

$\text{Mark}(\text{pp}, \text{prfk}, m \in \{0, 1\}^{\ell_m})$:

- Parse $\text{pp} = \perp$ and $\text{prfk} = (F, \text{pe.dk})$.
- Construct a circuit $D[F, \text{pe.dk}, m]$ described in Figure 6.
- Compute and output $\tilde{C} := i\mathcal{O}(D[F, \text{pe.dk}, m])$.

$\text{Sim}(\text{xk}, \tau, i)$:

- Parse $\text{xk} = \perp$ and $\tau = \text{pe.ek}$.
- Choose $\gamma \leftarrow \{0, 1\}$ and $s \leftarrow \{0, 1\}^{\ell}$.
- Compute $y := \text{PRG}(s)$.
- Compute $x \leftarrow \text{PE.Enc}(\text{pe.ek}, s || i || \gamma)$.
- Output (γ, x, y)

The size of the circuit D is appropriately padded to be the maximum size of all modified circuits, which will appear in the security proof.

The evaluation correctness of PRF_{io} immediately follows from the sparseness of PE and the functionality of $i\mathcal{O}$.¹³ PRF_{io} trivially satisfies pseudorandomness (against an authority) since Setup

¹³In fact, PRF_{io} satisfies a stronger evaluation correctness than one written in Definition 4.1. The evaluation correctness holds even for any PRF key prfk and input $x \in \text{Dom}$ like the statistical correctness by Cohen et al. [CHN⁺18].

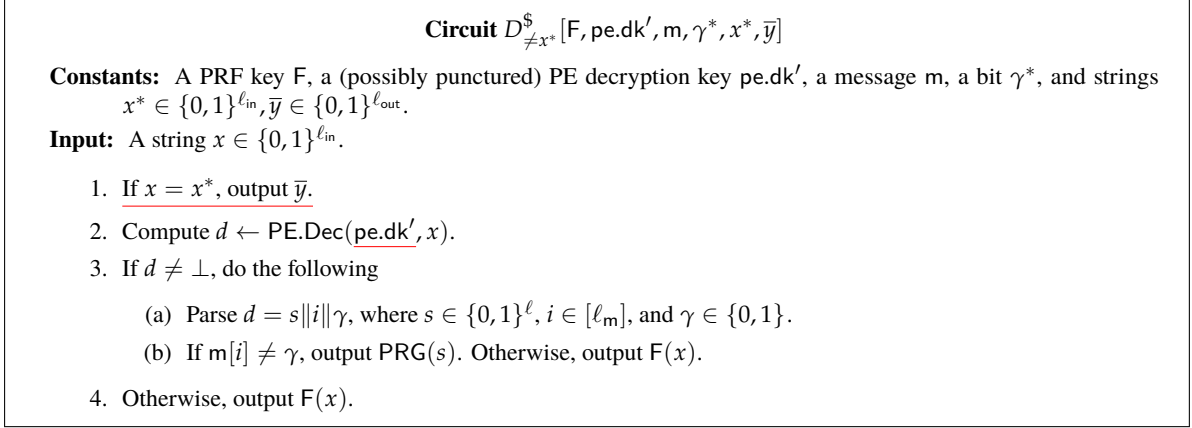


Figure 7: The description of $D_{\neq x^*}^{\$}$ (for $\gamma^* = m[i^*]$)

outputs nothing, τ is a public key pe.ek , and Eval is independent of $(\text{pe.ek}, \text{pe.dk})$ (pe.dk is not used in Eval). Moreover, we have the following theorem.

Theorem 7.6. *If PRF is a secure PPRF, PRG is a secure PRG, PE is a secure PE with strong ciphertext pseudorandomness, and $i\mathcal{O}$ is a secure IO, then $\text{PRF}_{i\mathcal{O}}$ is an extraction-less watermarking PRF satisfying SIM-MDD security with public simulation.*

Proof of Theorem 7.6. We define a sequence of hybrid games. We sometimes omit hard-coded values when we write some circuits. For example, we simply write D instead of $D[F, \text{pe.dk}, m]$ when hard-coded values $(F, \text{pe.dk}, m)$ are not important in arguments or clear from the context.

Hyb₀: This is the same as the case $b = 1$ in $\text{Exp}_{i^*, \mathcal{A}, \text{PRF}_{i\mathcal{O}}}^{\text{pub-sim-mdd}}(\lambda)$. In this game, \mathcal{A} is given $\tau = \text{pe.ek}$ and $\tilde{C} = i\mathcal{O}(D[F, \text{pe.dk}, m])$ as a public tag and a marked circuit, where $(\text{pe.ek}, \text{pe.dk}) \leftarrow \text{PE.Gen}(1^\lambda)$, $(F, \text{pe.dk})$ is the target PRF key, and m is the target message from \mathcal{A} . Also, \mathcal{A} is given $(\gamma^*, x^*, y^*) = (\gamma_1, x_1, y_1)$ as the challenge tuple, where $\gamma_1 \leftarrow \{0, 1\}$, $x_1 \leftarrow \text{PE.Enc}(\text{pe.ek}, s^* \| i^* \| \gamma_1)$, $y_1 := \text{PRG}(s^*)$, and $s^* \leftarrow \{0, 1\}^{\ell}$.

Case $\gamma^* = m[i^*]$: We consider two cases separately hereafter. First, we consider the case where $\gamma^* = m[i^*]$. We denote these hybrid games by Hyb_k^- . Note that we can choose γ^* at any time and hard-code it into \tilde{D} in the proof since it is a uniformly random bit in all hybrid games.

Hyb₁⁻: This is the same as Hyb_0 except that if $\gamma_1 = m[i^*]$, we use $\tilde{D} \leftarrow i\mathcal{O}(D_{\neq x_1}^{\$}[F, \text{pe.dk}_{\neq x_1}, m, \gamma_1, x_1, \bar{y}])$, where $D_{\neq x^*}^{\$}$ is described in Figure 7 and $\bar{y} := F(x_1)$. We use a punctured decryption key $\text{pe.dk}_{\neq x_1}$ instead of pe.dk . However, we do *not* use a punctured key for F .

Hyb₂⁻: This is the same as Hyb_1^- except that we generate

- $x_0 \leftarrow \{0, 1\}^{\ell_{\text{in}}}$,
- $\tilde{D} \leftarrow i\mathcal{O}(D_{\neq x_0}^{\$}[F, \text{pe.dk}_{\neq x_0}, m, \gamma_1, x_0, \bar{y}])$.

That is, we replace $x^* = x_1$ and $\text{pe.dk}' = \text{pe.dk}_{\neq x_1}$ with $x^* = x_0$ and $\text{pe.dk}' = \text{pe.dk}_{\neq x_0}$, respectively.

We also rename $\gamma_1 \leftarrow \{0, 1\}$ into $\gamma_0 \leftarrow \{0, 1\}$ (these distributions are the same).

Hyb₃⁻: This is the same as Hyb_2^- except that we use $y_0 \leftarrow \{0, 1\}^{\ell_{\text{out}}}$ instead of $y_1 := \text{PRG}(s^*)$.

We describe the high-level overview of hybrid games for $\gamma^* = m[i^*]$ in Figures 9 and 10.

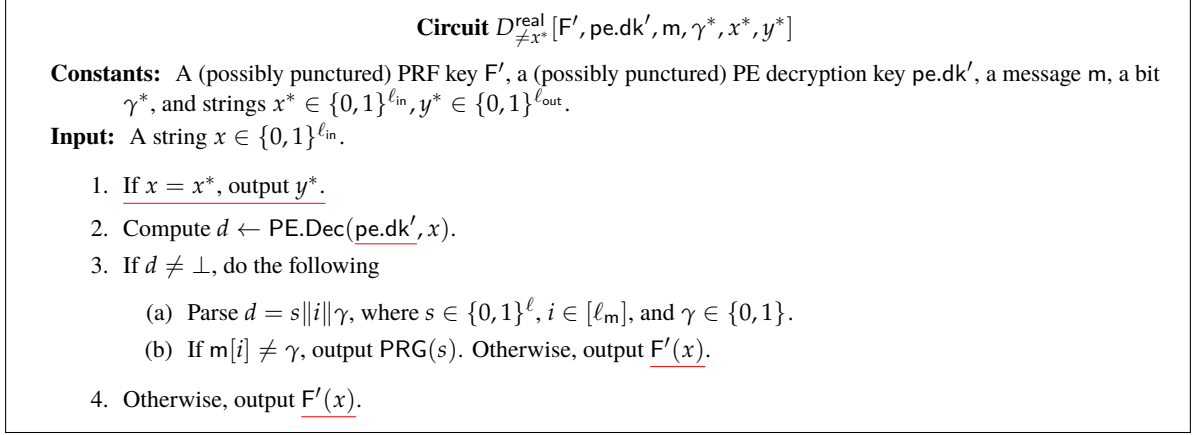


Figure 8: The description of $D_{\neq x^*}^{\text{real}}$ (for $\gamma^* = 1 - m[i^*]$)

Case $\gamma^* \neq m[i^*]$: Next, we consider the case where $\gamma^* \neq m[i^*]$. We denote these hybrid games by Hyb_k^{\neq} .

Hyb_1^{\neq} : This is the same as Hyb_0 except that if $\gamma_1 \neq m[i^*]$, we generate $\tilde{D} \leftarrow i\mathcal{O}(D_{\neq x_1}^{\text{real}}[F, \text{pe.dk}_{\neq x_1}, m, \gamma_1, x_1, y_1])$, where $D_{\neq x^*}^{\text{real}}$ is described in Figure 8 and $y_1 := \text{PRG}(s^*)$. We use a punctured decryption key $\text{pe.dk}_{\neq x_1}$ instead of pe.dk . However, we do *not* use a puncture key for F at this point.

Hyb_2^{\neq} : This is the same as Hyb_1^{\neq} except that

- $x_0 \leftarrow \{0, 1\}^{\ell_{\text{in}}}$,
- $\tilde{D} \leftarrow i\mathcal{O}(D_{\neq x_0}^{\text{real}}[F, \text{pe.dk}_{\neq x_0}, m, \gamma_1, x_0, y_1])$.

That is, we replace $x^* = x_1$ and $\text{pe.dk}' = \text{pe.dk}_{\neq x_1}$ with $x^* = x_0$ and $\text{pe.dk}' = \text{pe.dk}_{\neq x_0}$, respectively. We also rename $\gamma_1 \leftarrow \{0, 1\}$ into $\gamma_0 \leftarrow \{0, 1\}$ (these distributions are the same).

Hyb_3^{\neq} : This is the same as Hyb_2^{\neq} except that we use $y_0 \leftarrow \{0, 1\}^{\ell_{\text{out}}}$ instead of $y_1 := \text{PRG}(s^*)$.

Hyb_4^{\neq} : This is the same as Hyb_3^{\neq} except that we use $F_{\neq x_0}$ instead of F .

Hyb_5^{\neq} : This is the same as Hyb_4^{\neq} except that we use $y_0 := F(x_0)$ instead of $y_0 \leftarrow \{0, 1\}^{\ell_{\text{out}}}$.

Hyb_6^{\neq} : This is the same as Hyb_5^{\neq} except that we use F instead of $F_{\neq x_0}$.

We describe the high-level overview of hybrid games for $\gamma^* \neq m[i^*]$ in Figures 9 and 11.

End of case analysis: The two case analyses end. Remaining transitions are the reverse of transitions from Hyb_0 to Hyb_1^{\neq} or Hyb_1^{\neq} .

Hyb_4^{\neq} and Hyb_7^{\neq} : These are the same as Hyb_3^{\neq} and Hyb_6^{\neq} , respectively except that

- if $\gamma_0 = m[i^*]$, we use $\tilde{D} \leftarrow i\mathcal{O}(D[F, \text{pe.dk}, m])$ instead of $\tilde{D} \leftarrow i\mathcal{O}(D_{\neq x_0}^{\$}[F, \text{pe.dk}_{\neq x_0}, m, \gamma_0, x_0, \bar{y}])$, where $D_{\neq x^*}^{\$}$ is described in Figure 7 and $\bar{y} := F(x_0)$.
- if $\gamma_0 \neq m[i^*]$, we use $\tilde{D} \leftarrow i\mathcal{O}(D[F, \text{pe.dk}, m])$ instead of $\tilde{D} \leftarrow i\mathcal{O}(D_{\neq x_0}^{\text{real}}[F, \text{pe.dk}_{\neq x_0}, m, \gamma_0, x_0, y_0])$, where $D_{\neq x^*}^{\text{real}}$ is described in Figure 8 and $y_0 := F(x_0)$.

The each last hybrid is the same as the case $b = 0$ in $\text{Exp}_{i^*, \mathcal{A}, \text{PRF}_{\text{io}}}^{\text{pub-sim-mdd}}(\lambda)$. That is, \mathcal{A} is given $\tilde{C} = i\mathcal{O}(D[F, \text{pe.dk}, m])$ and $(\gamma_0, x_0, y_0) \leftarrow D_{\text{real}, i^*}$. Recall that $x_0 \leftarrow \{0, 1\}^{\ell_{\text{in}}}$ and

- $y_0 \leftarrow \{0, 1\}^{\ell_{\text{out}}}$ if $\gamma_0 = m[i^*]$ (see Hyb_3^-),
- $y_0 := F(x_0)$ if $\gamma_0 \neq m[i^*]$ (see Hyb_5^\neq).

		O_{chall}		security
		$\gamma_1 = m[i^*]$	$\gamma_1 \neq m[i^*]$	
Hyb_0	(F, pe.dk)	$i\mathcal{O}(D)$	$i\mathcal{O}(D)$	
Hyb_1^-	(F, <u>pe.dk$_{\neq x_1}$</u>)	<u>$i\mathcal{O}(D_{\neq x_1}^\\$)$</u>	N/A	IO & PE p-Cor.
Hyb_1^\neq	(F, <u>pe.dk$_{\neq x_1}$</u>)	N/A	<u>$i\mathcal{O}(D_{\neq x_1}^{\text{real}})$</u>	IO & PE p-Cor.

Figure 9: High-level overview of hybrid games from Hyb_0 to Hyb_1^- and Hyb_1^\neq . Note that in these hybrid games, $(\gamma_1, x_1, y_1) \leftarrow \text{Sim}(xk, \tau, i^*)$ and $x^* = x_1$. We use $\text{pe.dk}_{\neq x_1}$ in $D_{\neq x_1}^\$$ and $D_{\neq x_1}^{\text{real}}$, but F is not punctured yet. In “security” column, PE p-Cor. means PE punctured correctness.

$\gamma^* = m[i]$						
	$x^* (x_1/x_0)$	$y^* (y_1/y_0)$	\bar{y}	O_{chall}	prfk	security
Hyb_1^-	$\text{PE.Enc}(p_{i^*})$	$\text{PRG}(s^*)$	$F(x_1)$	$i\mathcal{O}(D_{\neq x_1}^\$ [x_1 \mapsto \bar{y}])$	(F, <u>pe.dk$_{\neq x_1}$</u>)	
Hyb_2^-	<u>$x_0 \leftarrow \\$</u>	$\text{PRG}(s^*)$	$F(x_0)$	$i\mathcal{O}(D_{\neq x_0}^\$ [x_0 \mapsto \bar{y}])$	(F, <u>pe.dk$_{\neq x_0}$</u>)	S-CPR
Hyb_3^-	$\$$	<u>$y_0 \leftarrow \\$</u>	$F(x_0)$	$i\mathcal{O}(D_{\neq x_0}^\$ [x_0 \mapsto \bar{y}])$	(F, <u>pe.dk$_{\neq x_0}$</u>)	PRG
Hyb_4^-	$\$$	$\$$	$F(x_0)$	<u>$i\mathcal{O}(D)$</u>	(F, <u>pe.dk</u>)	IO & PE p-Cor.

Figure 10: High-level overview of hybrid games from Hyb_1^- to Hyb_4^- . Here, $p_{i^*} := (s^* || i^* || \gamma_1)$ and $x_1 \leftarrow \text{PE.Enc}(\text{pe.ek}, p_{i^*})$. Note that \bar{y} is an output of $D_{\neq x^*}^\$$ for input x^* for $\gamma_1 = m[i^*]$ case. $D_{\neq x^*}^\$ [x^* \mapsto \bar{y}]$ means $D_{\neq x^*}^\$ (x^*)$ outputs the hard-coded value \bar{y} . In “security” column, S-CPR means the Strong Ciphertext PseudoRandomness of PE.

We prove the lemma by proving the following propositions.

The case $\gamma^* = m[i]$. We first prove propositions for the case $\gamma^* = m[i]$.

Proposition 7.7. *If $i\mathcal{O}$ is a secure IO and PE satisfies punctured correctness, it holds that*

$$|\Pr[\text{Hyb}_0 = 1] - \Pr[\text{Hyb}_1^- = 1]| \leq \text{negl}(\lambda).$$

Proof of Proposition 7.7. The difference between the two games is that $D_{\neq x_1}^\$ [F, \text{pe.dk}_{\neq x_1}, m, \gamma, x_1, \bar{y}]$ is used for O_{chall} instead of $D[F, \text{pe.dk}, m]$ in the case where $\gamma_1 = m[i^*]$. These two circuits are the same except that

- for input x_1 , $D_{\neq x_1}^\$$ directly outputs \bar{y} ,

due to the punctured correctness of PE. Thus, if the following hold, $D_{\neq x_1}^\$$ and D are functionally equivalent:

- $D(x_1)$ outputs $\bar{y} = F(x_1)$ when $\gamma_1 = m[i^*]$.

This holds since $x_1 \leftarrow \text{PE.Enc}(\text{pe.ek}, s^* || i^* || \gamma_1)$ and $D(x_1)$ runs the item (b) in Figure 6, but $\gamma_1 \neq m[i^*]$ does not hold in this case.

Thus, $D_{\neq x_1}^\$$ and D are functionally equivalent and the proposition holds due to IO security. ■

$\gamma^* \neq m[i]$					
	$x^* (x_1/x_0)$	$y^* (y_1/y_0)$	O_{chall}	prfk	security
Hyb ₁ [≠]	PE.Enc(p_{i^*})	PRG(s^*)	$i\mathcal{O}(D_{\neq x_1}^{\text{real}}[x_1 \mapsto y_1])$	(F, pe.dk _{≠x₁})	
Hyb ₂ [≠]	$x_0 \leftarrow \$$	PRG(s^*)	$i\mathcal{O}(D_{\neq x_0}^{\text{real}}[x_0 \mapsto y_1])$	(F, pe.dk _{≠x₀})	S-CPR
Hyb ₃ [≠]	$\$$	$y_0 \leftarrow \$$	$i\mathcal{O}(D_{\neq x_0}^{\text{real}}[x_0 \mapsto y_0])$	(F, pe.dk _{≠x₀})	PRG
Hyb ₄ [≠]	$\$$	$\$$	$i\mathcal{O}(D_{\neq x_0}^{\text{real}}[x_0 \mapsto y_0])$	(F _{≠x₀} , pe.dk _{≠x₀})	IO & PPRF p-Cor.
Hyb ₅ [≠]	$\$$	$F(x_0)$	$i\mathcal{O}(D_{\neq x_0}^{\text{real}}[x_0 \mapsto y_0])$	(F _{≠x₀} , pe.dk _{≠x₀})	PPRF
Hyb ₆ [≠]	$\$$	$F(x_0)$	$i\mathcal{O}(D_{\neq x_0}^{\text{real}}[x_0 \mapsto y_0])$	(F, pe.dk _{≠x₀})	IO & PPRF p-Cor.
Hyb ₇ [≠]	$\$$	$F(x_0)$	$i\mathcal{O}(D)$	(F, pe.dk)	IO & PE p-Cor.

Figure 11: High-level overview of hybrid games from Hyb₁[≠] to Hyb₇[≠]. Here, $p_{i^*} := (s^* \| i^* \| \gamma_1)$ and $x_1 \leftarrow \text{PE.Enc}(\text{pe.ek}, p_{i^*})$. $D_{\neq x^*}^{\text{real}}[x^* \mapsto y^*]$ means $D_{\neq x^*}^{\text{real}}(x^*)$ outputs the hard-coded value y^* . In “security” column, S-CPR and PPRF p-Cor. mean the Strong Ciphertext Pseudorandomness of PE and punctured correctness of PPRF, respectively.

Proposition 7.8. *If PE satisfies strong ciphertext pseudorandomness, it holds that*

$$|\Pr[\text{Hyb}_1^- = 1] - \Pr[\text{Hyb}_2^- = 1]| \leq \text{negl}(\lambda).$$

Proof of Proposition 7.8. We construct an algorithm \mathcal{B} for strong ciphertext pseudorandomness by using \mathcal{A} . \mathcal{B} generates $F \leftarrow \text{PRF.Gen}(1^\lambda)$, and chooses $s^* \leftarrow \{0, 1\}^\ell$ and $\gamma_1 \leftarrow \{0, 1\}$. \mathcal{B} sends $s^* \| i^* \| \gamma_1$ to the challenger. The challenger returns $(x^*, \text{pe.ek}, \text{pe.dk}_{\neq x^*})$ to \mathcal{B} .

Then, \mathcal{B} passes $\text{pp} := \perp$ and $\text{xk} := \perp$ to \mathcal{A} . \mathcal{B} also computes $\bar{y} := F(x^*)$.

Challenge: When \mathcal{A} sends a challenge query m , \mathcal{B} does the following

- Construct $D_{\neq x^*}^\$[F, \text{pe.dk}_{\neq x^*}, m, \gamma_1, x^*, \bar{y}]$ as described in Figure 7.
- Return $\tilde{C} := i\mathcal{O}(D_{\neq x^*}^\$[F, \text{pe.dk}_{\neq x^*}, m, \gamma_1, x^*, \bar{y}])$ and $\tau := \text{pe.ek}$ to \mathcal{A} .

After finishing \mathcal{A} 's challenge query, \mathcal{B} computes $y^* := \text{PRG}(s^*)$ and sends (γ_1, x^*, y^*) to \mathcal{A} . Finally, when \mathcal{A} terminates with output coin' , \mathcal{B} outputs coin' and terminates. \mathcal{B} perfectly simulates

- Hyb₁⁻ if $x^* \leftarrow \text{PE.Enc}(\text{pe.ek}, s^* \| i^* \| \gamma_1)$,
- Hyb₂⁻ if $x^* \leftarrow \{0, 1\}^{\ell_{\text{in}}}$.

Thus, we see that the proposition holds. ■

Proposition 7.9. *If PRG is a secure PRG, it holds that $|\Pr[\text{Hyb}_2^- = 1] - \Pr[\text{Hyb}_3^- = 1]| \leq \text{negl}(\lambda)$.*

Proof of Proposition 7.9. The difference between the two games is that y^* in the target triple (γ_0, x_0, y^*) is $\text{PRG}(s^*)$ or random in the case where $\gamma_0 = m[i]$. Recall that we rename $\gamma_1 \leftarrow \{0, 1\}$ to $\gamma_0 \leftarrow \{0, 1\}$. Note that we randomly choose $x_0 \leftarrow \{0, 1\}^{\ell_{\text{in}}}$ and use F and $\text{pe.dk}_{\neq x_0}$ in these games. Thus, we can apply pseudorandomness of PRG since the value s^* is never used anywhere else. ■

The case where $\gamma^* \neq m[i]$. Next, we prove propositions for the case where $\gamma^* \neq m[i]$.

Proposition 7.10. *If $i\mathcal{O}$ is a secure IO and PE satisfies punctured correctness, it holds that*

$$|\Pr[\text{Hyb}_0 = 1] - \Pr[\text{Hyb}_1^\neq = 1]| \leq \text{negl}(\lambda).$$

Proof of Proposition 7.10. The difference between the two games is that $D_{\neq x_1}^{\text{real}}[F, \text{pe.dk}_{\neq x_1}, m, \gamma, x_1, y_1]$ is used for the challenge query instead of $D[F, \text{pe.dk}, m]$ in the case where $\gamma_1 \neq m[i^*]$. These two circuits are the same except that

- for input x_1 , $D_{\neq x_1}^{\text{real}}$ directly outputs the hard-wired value $y_1 = \text{PRG}(s^*)$,

due to the punctured correctness of PE. Thus, if the following hold, $D_{\neq x_1}^{\text{real}}$ and D are functionally equivalent:

- $D(x_1)$ outputs $y_1 = \text{PRG}(s^*)$ when $\gamma_1 \neq m[i^*]$,

This holds since $x_1 \leftarrow \text{PE.Enc}(\text{pe.ek}, s^* \| i^* \| \gamma_1)$, $D(x_1)$ runs the item (b) in Figure 6, and $\gamma_1 \neq m[i^*]$ holds in this case. Thus, $D_{\neq x_1}^{\text{real}}$ and D are functionally equivalent and the proposition holds due to IO security. ■

Proposition 7.11. *If PE satisfies strong pseudorandom ciphertext, it holds that*

$$\left| \Pr[\text{Hyb}_1^{\neq} = 1] - \Pr[\text{Hyb}_2^{\neq} = 1] \right| \leq \text{negl}(\lambda).$$

Proof of Proposition 7.11. We construct an algorithm \mathcal{B} for strong ciphertext pseudorandomness by using a distinguisher \mathcal{A} . \mathcal{B} generates $F \leftarrow \text{PRF.Gen}(1^\lambda)$ and chooses $s^* \leftarrow \{0, 1\}^\ell$ and $\gamma_1 \leftarrow \{0, 1\}$. \mathcal{B} sends $s^* \| i^* \| \gamma_1$ to the challenger. The challenger returns $(x^*, \text{pe.ek}, \text{pe.dk}_{\neq x^*})$ to \mathcal{B} .

Then, \mathcal{B} passes $\text{pp} := \perp$ and $\text{xk} := \perp$ to \mathcal{A} . \mathcal{B} also computes $y^* := \text{PRG}(s^*)$.

Challenge: When \mathcal{A} sends a challenge query m , \mathcal{B} does the following

- Construct $D_{\neq x^*}^{\text{real}}[F, \text{pe.dk}_{\neq x^*}, m, \gamma_1, x^*, y^*]$ where $y^* := \text{PRG}(s^*)$ as described in Figure 8.
- Return $\tilde{C} := i\mathcal{O}(D_{\neq x^*}^{\text{real}})$ and $\tau := \text{pe.ek}$ to \mathcal{A} .

After finishing \mathcal{A} 's challenge, \mathcal{B} sends (γ_1, x^*, y^*) to \mathcal{A} . Finally, when \mathcal{A} terminates with output coin' , \mathcal{B} outputs coin' and terminates.

\mathcal{B} perfectly simulates

- Hyb_1^{\neq} if $x^* \leftarrow \text{PE.Enc}(\text{pe.ek}, s^* \| i^* \| \gamma_1)$,
- Hyb_2^{\neq} if $x^* \leftarrow \{0, 1\}^{\ell_{\text{in}}}$.

Thus, we see that the proposition holds. ■

Proposition 7.12. *If PRG is a secure PRG, it holds that $\left| \Pr[\text{Hyb}_2^{\neq} = 1] - \Pr[\text{Hyb}_3^{\neq} = 1] \right| \leq \text{negl}(\lambda)$.*

Proof of Proposition 7.12. The difference between the two games is that y^* in the target triple (γ_0, x_0, y^*) is $\text{PRG}(s^*)$ or random in the case where $\gamma_0 \neq m[i]$. Recall that we rename $\gamma_1 \leftarrow \{0, 1\}$ to $\gamma_0 \leftarrow \{0, 1\}$. Note that we randomly choose $x_0 \leftarrow \{0, 1\}^{\ell_{\text{in}}}$ and use F and $\text{pe.dk}_{\neq x_0}$ in these games. Thus, we can apply pseudorandomness of PRG since the value s^* is never used anywhere else. ■

Proposition 7.13. *If $i\mathcal{O}$ is a secure IO and F satisfies punctured correctness, it holds that*

$$\left| \Pr[\text{Hyb}_3^{\neq} = 1] - \Pr[\text{Hyb}_4^{\neq} = 1] \right| \leq \text{negl}(\lambda).$$

Proof of Proposition 7.13. The difference between the two games is that $D_{\neq x_0}^{\text{real}}[F_{\neq x_0}, \text{pe.dk}_{\neq x_0}, m, \gamma_0, x_0, y_0]$ is used for the challenge query instead of $D_{\neq x_0}^{\text{real}}[F, \text{pe.dk}_{\neq x_0}, m, \gamma_0, x_0, y_0]$ in the case where $\gamma_0 \neq m[i^*]$. These two circuits are the same except that we use $F_{\neq x_0}$ instead of F . Those two circuits above are functionally equivalent since $F_{\neq x_0}(\cdot)$ is functionally equivalent to F except for x_0 and both

$D_{\neq x_0}^{\text{real}}[F_{\neq x_0}, \text{pe.dk}_{\neq x_0}, m, \gamma_0, x_0, y_0](x_0)$ and $D_{\neq x_0}^{\text{real}}[F, \text{pe.dk}_{\neq x_0}, m, \gamma_0, x_0, y_0](x_0)$ directly outputs y_0 by the description of $D_{\neq x_0}^{\text{real}}$. Note that $D_{\neq x_0}^{\text{real}}$ does not have any “if branch” condition that uses F or $F_{\neq x_0}$. Thus, the proposition holds due to IO security. ■

Proposition 7.14. *If F satisfies punctured pseudorandomness, it holds that*

$$\left| \Pr[\text{Hyb}_4^{\neq} = 1] - \Pr[\text{Hyb}_5^{\neq} = 1] \right| \leq \text{negl}(\lambda).$$

Proof of Proposition 7.14. We construct an algorithm \mathcal{B} that breaks the pseudorandomness at punctured points of F by using \mathcal{A} .

\mathcal{B} generates $(\text{pe.ek}, \text{pe.dk}) \leftarrow \text{PE.Gen}(1^\lambda)$, chooses $x_0 \leftarrow \{0, 1\}^{\ell_{\text{in}}}$ and $\gamma_0 \leftarrow \{0, 1\}$, sends x_0 as the challenge to its challenger of F , and receives $F_{\neq x_0}$ and y^* . Here x_0 does not rely on m , so we can generate x_0 before m is fixed. \mathcal{B} sends $\text{pp} := \perp$ and $\text{xk} := \perp$ to \mathcal{A} . \mathcal{B} also computes $\text{pe.dk}_{\neq x_0} \leftarrow \text{PE.Puncture}(\text{pe.dk}, x_0)$.

Challenge: For query m , \mathcal{B} can simulate the target marked circuit $\tilde{C} = i\mathcal{O}(D_{\neq x_0}^{\text{real}}[F_{\neq x_0}, \text{pe.dk}_{\neq x_0}, m, \gamma_0, x_0, y^*])$ by using $\text{pe.dk}_{\neq x_0}$, $F_{\neq x_0}$, y^* , and the public tag $\tau = \text{pe.ek}$.

After finishing \mathcal{A} 's challenge query, \mathcal{B} sends (γ_0, x_0, y^*) to \mathcal{A} . Finally, when \mathcal{A} terminates with output coin' , \mathcal{B} outputs coin' and terminates.

\mathcal{B} perfectly simulates

- Hyb_4^{\neq} if $y^* \leftarrow \{0, 1\}^{\ell_{\text{out}}}$,
- Hyb_5^{\neq} if $y^* := F(x_0)$.

The punctured pseudorandomness of F immediately implies this proposition. ■

Proposition 7.15. *If $i\mathcal{O}$ is a secure IO and F satisfies punctured correctness, it holds that*

$$\left| \Pr[\text{Hyb}_5^{\neq} = 1] - \Pr[\text{Hyb}_6^{\neq} = 1] \right| \leq \text{negl}(\lambda).$$

Proof of Proposition 7.15. The difference between the two games is that $D_{\neq x_0}^{\text{real}}[F, \text{pe.dk}_{\neq x_0}, m, \gamma_0, x_0, y_0]$ is used for the challenge query instead of $D_{\neq x_0}^{\text{real}}[F_{\neq x_0}, \text{pe.dk}_{\neq x_0}, m, \gamma_0, x_0, y_0]$ in the case where $\gamma_0 \neq m[i^*]$. These two circuits are the same except that we use F instead of $F_{\neq x_0}$. This proof is the same as that of Proposition 7.14 (in a reverse manner). Thus, we omit it. ■

End of case analyses. We complete the two case analyses.

Proposition 7.16. *If $i\mathcal{O}$ is a secure IO, F satisfies punctured correctness, and PE satisfies punctured correctness, it holds that $|\Pr[\text{Hyb}_3^{\neq} = 1] - \Pr[\text{Hyb}_4^{\neq} = 1]| \leq \text{negl}(\lambda)$ and $|\Pr[\text{Hyb}_6^{\neq} = 1] - \Pr[\text{Hyb}_7^{\neq} = 1]| \leq \text{negl}(\lambda)$.*

Proof. This proof is the same as that of Proposition 7.7 and Proposition 7.10, respectively (in a reverse manner). Thus, we omit them. ■

We complete the proof of Theorem 7.6. ■

8 Putting Pieces Altogether

Privately extractable watermarking PRF. We summarize how to obtain our privately extractable watermarking PRF.

By Theorems 2.24, 2.27, 2.30, 6.1 and 6.2, we obtain an extraction-less watermarking with private simulation from the QLWE assumption. By combining this with Theorems 4.6 and 5.2, we obtain the following theorem.

Theorem 8.1. *If the QLWE assumption holds, there exists a privately extractable watermarking PRF.*

Publicly extractable watermarking PRF. We summarize how to obtain our publicly extractable watermarking PRF.

By Theorems 2.17, 2.21, 7.5 and 7.6, we obtain an extraction-less watermarking with public simulation from IO and the QLWE assumption since OWFs can be instantiated with the QLWE assumption. By combining this with Theorem 5.3, we obtain a publicly extractable watermarking PRF from IO and the QLWE assumption. Thus, we obtain the following theorem.

Theorem 8.2. *If there exists a secure IO and the QLWE assumption holds, there exists a publicly extractable watermarking PRF.*

References

- [AHU19] Andris Ambainis, Mike Hamburg, and Dominique Unruh. Quantum security proofs using semi-classical oracles. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 269–295. Springer, Heidelberg, August 2019. (Cited on page 48.)
- [AKPW13] Joël Alwen, Stephan Krenn, Krzysztof Pietrzak, and Daniel Wichs. Learning with rounding, revisited - new reduction, properties and applications. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 57–74. Springer, Heidelberg, August 2013. (Cited on page 51.)
- [AL21] Prabhanjan Ananth and Rolando L. La Placa. Secure software leasing. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part II*, volume 12697 of *LNCS*, pages 501–530. Springer, Heidelberg, October 2021. (Cited on page 1.)
- [ALL⁺21] Scott Aaronson, Jiahui Liu, Qipeng Liu, Mark Zhandry, and Ruizhe Zhang. New approaches for quantum copy-protection. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 526–555, Virtual Event, August 2021. Springer, Heidelberg. (Cited on page 1, 10, 11.)
- [AP20] Shweta Agrawal and Alice Pellet-Mary. Indistinguishability obfuscation without maps: Attacks and fixes for noisy linear FE. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 110–140. Springer, Heidelberg, May 2020. (Cited on page 19.)
- [ARU14] Andris Ambainis, Ansis Rosmanis, and Dominique Unruh. Quantum attacks on classical proof systems: The hardness of quantum rewinding. In *55th FOCS*, pages 474–483. IEEE Computer Society Press, October 2014. (Cited on page 1, 9.)

- [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, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 41–69. Springer, Heidelberg, December 2011. (Cited on page 1.)
- [BGI⁺12] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. *Journal of the ACM*, 59(2):6:1–6:48, 2012. (Cited on page 1, 18.)
- [BGI14] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In Hugo Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 501–519. Springer, Heidelberg, March 2014. (Cited on page 16.)
- [BGMZ18] James Bartusek, Jiaxin Guan, Fermi Ma, and Mark Zhandry. Return of GGH15: Provable security against zeroizing attacks. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 544–574. Springer, Heidelberg, November 2018. (Cited on page 19.)
- [BHH⁺19] Nina Bindel, Mike Hamburg, Kathrin Hövelmanns, Andreas Hülsing, and Edoardo Persichetti. Tighter proofs of CCA security in the quantum random oracle model. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019, Part II*, volume 11892 of *LNCS*, pages 61–90. Springer, Heidelberg, December 2019. (Cited on page 48.)
- [BLW17] Dan Boneh, Kevin Lewi, and David J. Wu. Constraining pseudorandom functions privately. In Serge Fehr, editor, *PKC 2017, Part II*, volume 10175 of *LNCS*, pages 494–524. Springer, Heidelberg, March 2017. (Cited on page 1, 9, 20.)
- [BSW06] Dan Boneh, Amit Sahai, and Brent Waters. Fully collusion resistant traitor tracing with short ciphertexts and private keys. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 573–592. Springer, Heidelberg, May / June 2006. (Cited on page 2, 10.)
- [BTVW17] Zvika Brakerski, Rotem Tsabary, Vinod Vaikuntanathan, and Hoeteck Wee. Private constrained PRFs (and more) from LWE. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 264–302. Springer, Heidelberg, November 2017. (Cited on page 18.)
- [BW13] Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 280–300. Springer, Heidelberg, December 2013. (Cited on page 16.)
- [CC17] Ran Canetti and Yilei Chen. Constraint-hiding constrained PRFs for NC^1 from LWE. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part I*, volume 10210 of *LNCS*, pages 446–476. Springer, Heidelberg, April / May 2017. (Cited on page 9.)
- [CFN94] Benny Chor, Amos Fiat, and Moni Naor. Tracing traitors. In Yvo Desmedt, editor, *CRYPTO'94*, volume 839 of *LNCS*, pages 257–270. Springer, Heidelberg, August 1994. (Cited on page 1.)
- [CHN⁺18] Aloni Cohen, Justin Holmgren, Ryo Nishimaki, Vinod Vaikuntanathan, and Daniel Wichs. Watermarking cryptographic capabilities. *SIAM Journal on Computing*, 47(6):2157–2202, 2018. (Cited on page 1, 7, 8, 20, 34, 35, 36, 51, 52, 60, 61.)

- [CHVW19] Yilei Chen, Minki Hhan, Vinod Vaikuntanathan, and Hoeteck Wee. Matrix PRFs: Constructions, attacks, and applications to obfuscation. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019, Part I*, volume 11891 of *LNCS*, pages 55–80. Springer, Heidelberg, December 2019. (Cited on page 19.)
- [CMSZ21] Alessandro Chiesa, Fermi Ma, Nicholas Spooner, and Mark Zhandry. Post-quantum succinct arguments: Breaking the quantum rewinding barrier. In Nisheeth Vishnoi, editor, *FOCS 2021 (to appear)*. IEEE, 2021. (Cited on page 1, 9, 13.)
- [DQV⁺21] Lalita Devadas, Willy Quach, Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. Succinct lwe sampling, random polynomials and obfuscation. In Kobbi Nissim and Brent Waters, editors, *TCC 2021*, *LNCS*. Springer, 2021. (Cited on page 19.)
- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, 1986. (Cited on page 16.)
- [GKM⁺19] Rishab Goyal, Sam Kim, Nathan Manohar, Brent Waters, and David J. Wu. Watermarking public-key cryptographic primitives. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 367–398. Springer, Heidelberg, August 2019. (Cited on page 1, 9, 10, 20.)
- [GKW19] Rishab Goyal, Venkata Koppula, and Brent Waters. New approaches to traitor tracing with embedded identities. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019, Part II*, volume 11892 of *LNCS*, pages 149–179. Springer, Heidelberg, December 2019. (Cited on page 10.)
- [GKWW21] Rishab Goyal, Sam Kim, Brent Waters, and David J. Wu. Beyond software watermarking: Traitor-tracing for pseudorandom functions. In Mehdi Tibouchi and Huaxiong Wang, editors, *Asiacrypt 2021 (to appear)*, Lecture Notes in Computer Science. Springer, 2021. (Cited on page 2, 3, 9, 10, 20, 22.)
- [GP21] Romain Gay and Rafael Pass. Indistinguishability obfuscation from circular security. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 736–749. ACM, 2021. (Cited on page 19.)
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999. (Cited on page 15.)
- [HJL21] Samuel B. Hopkins, Aayush Jain, and Huijia Lin. Counterexamples to new circular security assumptions underlying iO. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part II*, volume 12826 of *LNCS*, pages 673–700, Virtual Event, August 2021. Springer, Heidelberg. (Cited on page 19.)
- [HMW07] Nicholas Hopper, David Molnar, and David Wagner. From weak to strong watermarking. In Salil P. Vadhan, editor, *TCC 2007*, volume 4392 of *LNCS*, pages 362–382. Springer, Heidelberg, February 2007. (Cited on page 1.)
- [Jor75] Camille Jordan. Essai sur la géométrie à n dimensions. *Bulletin de la Société Mathématique de France*, 3:103–174, 1875. (Cited on page 13.)

- [KNY21] Fuyuki Kitagawa, Ryo Nishimaki, and Takashi Yamakawa. Secure software leasing from standard assumptions. In Kobbi Nissim and Brent Waters, editors, *TCC 2021*, LNCS. Springer, 2021. (Cited on page 1, 10.)
- [KPTZ13] Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 2013*, pages 669–684. ACM Press, November 2013. (Cited on page 16.)
- [KW19] Sam Kim and David J. Wu. Watermarking PRFs from lattices: Stronger security via extractable PRFs. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 335–366. Springer, Heidelberg, August 2019. (Cited on page 1, 9, 20, 22.)
- [KW21] Sam Kim and David J. Wu. Watermarking cryptographic functionalities from standard lattice assumptions. *J. Cryptol.*, 34(3):28, 2021. (Cited on page 1, 9, 20.)
- [MW05] Chris Marriott and John Watrous. Quantum arthur-merlin games. *Comput. Complex.*, 14(2):122–152, 2005. (Cited on page 13.)
- [Nao91] Moni Naor. Bit commitment using pseudorandomness. *Journal of Cryptology*, 4(2):151–158, January 1991. (Cited on page 51.)
- [Nis13] Ryo Nishimaki. How to watermark cryptographic functions. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 111–125. Springer, Heidelberg, May 2013. (Cited on page 1.)
- [Nis19] Ryo Nishimaki. How to watermark cryptographic functions by bilinear maps. *IEICE Transactions*, 102-A(1):99–113, 2019. (Cited on page 1.)
- [Nis20] Ryo Nishimaki. Equipping public-key cryptographic primitives with watermarking (or: A hole is to watermark). In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part I*, volume 12550 of *LNCS*, pages 179–209. Springer, Heidelberg, November 2020. (Cited on page 1, 9, 10, 20.)
- [NSS99] David Naccache, Adi Shamir, and Julien P. Stern. How to copyright a function? In Hideki Imai and Yuliang Zheng, editors, *PKC'99*, volume 1560 of *LNCS*, pages 188–196. Springer, Heidelberg, March 1999. (Cited on page 1.)
- [NWZ16] Ryo Nishimaki, Daniel Wichs, and Mark Zhandry. Anonymous traitor tracing: How to embed arbitrary information in a key. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 388–419. Springer, Heidelberg, May 2016. (Cited on page 10.)
- [Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 333–342. ACM Press, May / June 2009. (Cited on page 18.)
- [PS18] Chris Peikert and Sina Shiehian. Privately constraining and programming PRFs, the LWE way. In Michel Abdalla and Ricardo Dahab, editors, *PKC 2018, Part II*, volume 10770 of *LNCS*, pages 675–701. Springer, Heidelberg, March 2018. (Cited on page 9, 18.)
- [PW11] Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. *SIAM Journal on Computing*, 40(6):1803–1844, 2011. (Cited on page 51.)

- [QWZ18] Willy Quach, Daniel Wichs, and Giorgos Zirdelis. Watermarking PRFs under standard assumptions: Public marking and security with extraction queries. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 669–698. Springer, Heidelberg, November 2018. (Cited on page 1, 9, 19, 20, 21, 22, 28.)
- [Reg] Oded Regev. Witness-preserving amplification of qma (lecture notes). https://cims.nyu.edu/~regev/teaching/quantum_fall_2005/ln/qma.pdf. (Cited on page 13.)
- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM*, 56(6):34:1–34:40, 2009. (Cited on page 1.)
- [SW21] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: Deniable encryption, and more. *SIAM J. Comput.*, 50(3):857–908, 2021. (Cited on page 50.)
- [Unr12] Dominique Unruh. Quantum proofs of knowledge. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 135–152. Springer, Heidelberg, April 2012. (Cited on page 1, 9.)
- [Wat09] John Watrous. Zero-knowledge against quantum attacks. *SIAM J. Comput.*, 39(1):25–58, 2009. (Cited on page 1, 9.)
- [WW21] Hoeteck Wee and Daniel Wichs. Candidate obfuscation via oblivious LWE sampling. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part III*, volume 12698 of *LNCS*, pages 127–156. Springer, Heidelberg, October 2021. (Cited on page 19.)
- [YAL⁺19] Rupeng Yang, Man Ho Au, Junzuo Lai, Qiuliang Xu, and Zuoxia Yu. Collusion resistant watermarking schemes for cryptographic functionalities. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part I*, volume 11921 of *LNCS*, pages 371–398. Springer, Heidelberg, December 2019. (Cited on page 1, 9, 20.)
- [YAYX20] Rupeng Yang, Man Ho Au, Zuoxia Yu, and Qiuliang Xu. Collusion resistant watermarkable PRFs from standard assumptions. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 590–620. Springer, Heidelberg, August 2020. (Cited on page 1.)
- [YF11] Maki Yoshida and Toru Fujiwara. Toward digital watermarking for cryptographic data. *IEICE Transactions*, 94-A(1):270–272, 2011. (Cited on page 1.)
- [Zha12a] Mark Zhandry. How to construct quantum random functions. In *53rd FOCS*, pages 679–687. IEEE Computer Society Press, October 2012. (Cited on page 1, 16, 50.)
- [Zha12b] Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 758–775. Springer, Heidelberg, August 2012. (Cited on page 1, 50.)
- [Zha19] Mark Zhandry. How to record quantum queries, and applications to quantum indistinguishability. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 239–268. Springer, Heidelberg, August 2019. (Cited on page 1.)
- [Zha20] Mark Zhandry. Schrödinger’s pirate: How to trace a quantum decoder. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part III*, volume 12552 of *LNCS*, pages 61–91. Springer, Heidelberg, November 2020. (Cited on page 1, 2, 3, 4, 6, 9, 12, 22, 24.)

A Achieving QSIM-MDD from SIM-MDD

We prove Theorem 4.6, that is, we show that we can transform extraction-less watermarking PRF satisfying SIM-MDD security with private simulation into one satisfying QSIM-MDD security with private simulation, by using a QPRF. Before the proof, we introduce semi-classical one-way to hiding (O2H) lemma.

A.1 Semi-Classical One-Way to Hiding (O2H) Lemma

We recall a few lemmas.

Definition A.1 (Punctured oracle). *Let $F : X \rightarrow Y$ be any function, and $S \subset X$ be a set. The oracle $F \setminus S$ (“ F punctured by S ”) takes as input a value $x \in X$. It first computes whether $x \in S$ into an auxiliary register and measures it. Then it computes $F(x)$ and returns the result. Let Find be the event that any of the measurements returns 1.*

Lemma A.2 (Semi-classical O2H [AHU19, Theorem 1]). *Let $G, H : X \rightarrow Y$ be random functions, z be a random value, and $S \subseteq X$ be a random set such that $G(x) = H(x)$ for every $x \notin S$. The tuple (G, H, S, z) may have arbitrary joint distribution. Furthermore, let \mathcal{A} be a quantum oracle algorithm. Let Ev be any classical event. Then we have*

$$\left| \Pr[\text{Ev} : \mathcal{A}^{|H\rangle}(z)] - \Pr[\text{Ev} : \mathcal{A}^{|G\rangle}(z)] \right| \leq 2\sqrt{(q+1) \cdot \Pr[\text{Find} : \mathcal{A}^{|H \setminus S\rangle}(z)]} .$$

Lemma A.3 (Search in semi-classical oracle [AHU19, Theorem 2]). *Let $H : X \rightarrow Y$ be a random function, let z be a random value, and let $S \subset X$ be a random set. (H, S, z) may have arbitrary joint distribution. Let \mathcal{A} be a quantum oracle algorithm. If for each $x \in X$, $\Pr[x \in S] \leq \epsilon$ (conditioned on H and z), then we have*

$$\Pr[\text{Find} : \mathcal{A}^{|H \setminus S\rangle}(z)] \leq 4q\epsilon ,$$

where q is the number of queries to H by \mathcal{A} .

Note that the above lemma is originally introduced in [AHU19], but we use a variant that is closer to Lemma 4 in [BHH⁺19].

A.2 Proof

Construction. We start with the construction. Let $\text{ELWMPRF} = (\text{Setup}, \text{Gen}, \text{Eval}, \text{Mark}, \text{Sim})$ be an extraction-less watermarking PRF scheme satisfying SIM-MDD security with private simulation. We also let the message space of ELWMPRF is $\{0, 1\}^{\ell_m}$. Let PRF be a QPRF with domain $\{0, 1\}^\lambda$ and range \mathcal{R}_{Sim} , which is the randomness space of Sim . We construct an extraction-less watermarking PRF scheme $\text{QELWMPRF} = (\text{QEL.Setup}, \text{QEL.Gen}, \text{QEL.Eval}, \text{QEL.Mark}, \text{QEL.Sim})$ satisfying QSIM-MDD security with private simulation as follows. We use Gen , Eval , and Mark as QEL.Gen , QEL.Eval , and QEL.Mark , respectively. The domain and range of QELWMPRF are the same as those of ELWMPRF . The mark space of QELWMPRF is $\{0, 1\}^{\ell_m}$. Also, we construct QEL.Setup and QEL.Sim as follows.

$\text{QEL.Setup}(1^\lambda)$:

- Generate $(\text{pp}, \text{xk}) \leftarrow \text{Setup}(1^\lambda)$.
- Genrate $K \leftarrow \{0, 1\}^\lambda$.
- Outputs $(\text{pp}, \text{qxk} := (\text{xk}, K))$.

QEL.Sim(qxk, $\tau, i; r$):

- Parse $(xk, K) \leftarrow \text{qxk}$.
- Output $(\gamma, x, y) \leftarrow \text{Sim}(xk, \tau, i; \text{PRF}_K(r))$.

Security analysis. Let $i^* \in [\ell_m]$ and \mathcal{A} be any QPT adversary for QSIM-MDD security with private simulation making total q queries to O_{sim} and O_{api} . We prove that for any polynomial w , it holds that $\text{Adv}_{i^*, \mathcal{A}, \text{QELWMPRF}}^{\text{q-sim-mdd}}(\lambda) \leq 1/w$. We prove it using hybrid games. Let SUC_X be the event that the final output is 1 in Game X. We define a distribution $D_{\tau', i'}$ as

$D_{\tau', i'}$: Output $(\gamma, x, y) \leftarrow \text{Sim}(xk, \tau', i')$.

Game 1: This is $\text{Expt}_{i^*, \mathcal{A}, \text{QELWMPRF}}^{\text{q-sim-mdd}}(\lambda)$. Thus, $\text{Adv}_{i^*, \mathcal{A}, \text{QELWMPRF}}^{\text{q-sim-mdd}}(\lambda) = 2|\Pr[\text{SUC}_1] - 1/2|$.

1. The challenger generates $(\text{pp}, xk) \leftarrow \text{Setup}(1^\lambda)$ and $K \leftarrow \{0, 1\}^\lambda$, and gives pp to \mathcal{A} . \mathcal{A} send $m \in \{0, 1\}^{\ell_m}$ to the challenger. The challenger generates $(\tau, \text{prfk}) \leftarrow \text{Gen}(\text{pp})$, computes $\tilde{C} \leftarrow \text{Mark}(\text{pp}, \text{prfk}, m)$, and sends \tilde{C} to \mathcal{A} .
2. \mathcal{A} can access to the following oracles.
 - O_{sim} : On input τ' and i' , it returns $\text{Sim}(xk, \tau', i'; \text{PRF}_K(r))$, where $r \leftarrow \{0, 1\}^\lambda$.
 - O_{api} : On input $(\epsilon, \delta, \tau', i')$ and a quantum state q , it returns the result of $\mathcal{A}PI_{\mathcal{P}, D_{\tau', i'}^{\text{PRF}}}^{\epsilon, \delta}(q)$ and the post measurement state, where $D_{\tau', i'}^{\text{PRF}} = D_{\tau', i'}(\text{PRF}_K(\cdot))$.
3. The challenger generates coin $\leftarrow \{0, 1\}$. If coin = 0, the challenger samples $(\gamma, x, y) \leftarrow D_{\text{real}, i^*}$. If coin = 1, the challenger generates $(\gamma, x, y) \leftarrow \text{Sim}(xk, \tau, i^*; \text{PRF}_K(r^*))$, where, $r^* \leftarrow \{0, 1\}^\lambda$. The challenger sends (γ, x, y) to \mathcal{A} .
4. When \mathcal{A} terminates with output coin', the challenger outputs 1 if coin = coin' and 0 otherwise.

Game 2: This game is the same as Game 1 except that PRF_K is replaced with a quantum-accessible random function R .

We have $|\Pr[\text{SUC}_1] - \Pr[\text{SUC}_2]| = \text{negl}(\lambda)$ from the the security of PRF.

Game 3: This game is the same as Game 2 except that R is replaced with

$$V(r) = \begin{cases} v^* & (\text{if } r = r^*) \\ R(r) & (\text{otherwise}), \end{cases}$$

where $v^* \leftarrow \mathcal{R}_{\text{Sim}}$.

We have $|\Pr[\text{SUC}_2] - \Pr[\text{SUC}_3]| = 0$.

Game 4: This game is the same as Game 3 except the followings. When \mathcal{A} makes a query τ' and i' to O_{sim} , $\text{Sim}(xk, \tau', i'; R(r))$ is returned instead of $\text{Sim}(xk, \tau', i'; V(r))$. Also, when \mathcal{A} makes a query $(\epsilon, \delta, \tau', i')$ to O_{api} , $\mathcal{A}PI_{\mathcal{P}, D_{\tau', i'}^R}^{\epsilon, \delta}(q)$ is performed instead of $\mathcal{A}PI_{\mathcal{P}, D_{\tau', i'}^V}^{\epsilon, \delta}(q)$, where $D_{\tau', i'}^R = D_{\tau', i'}(R(\cdot))$ and $D_{\tau', i'}^V = D_{\tau', i'}(V(\cdot))$.

By this change, V is now used only for generating the challenge tuple $(\gamma, x, y) \leftarrow \text{Sim}(xk, \tau, i^*; V(r^*)) = \text{Sim}(xk, \tau, i^*; v^*)$.

We have $|\Pr[\text{SUC}_3] - \Pr[\text{SUC}_4]| = O(\sqrt{\frac{q^2}{2^\lambda}})$ from Lemma A.2 and Lemma A.3.

Game 5: This game is the same as Game 4 except that R is replaced with $G \circ F$, where $F : \{0, 1\}^\lambda \rightarrow [s]$ and $G : [s] \rightarrow \mathcal{R}_{\text{Sim}}$ are random functions and s is a polynomial of λ specified later.

Theorem A.4 (Small Range Distribution [Zha12a]). *For any QPT adversary \mathcal{B} making q quantum queries to R or $G \circ F$, we have $\left| \Pr[\mathcal{B}^{(R)}(1^\lambda) = 1] - \Pr[\mathcal{B}^{(G \circ F)}(1^\lambda) = 1] \right| \leq O(q^3/s)$.*

By the above theorem, we have $|\Pr[\text{SUC}_4] - \Pr[\text{SUC}_5]| = O(q^3/s)$.

We can simulate F using a $2q$ -wise independent function E by the following theorem.

Theorem A.5 ([Zha12b]). *For any QPT adversary \mathcal{B} making q quantum queries to F or E , we have $\Pr[\mathcal{B}^{(F)}(1^\lambda) = 1] = \Pr[\mathcal{B}^{(E)}(1^\lambda) = 1]$.*

We can efficiently simulate $\mathcal{A}PI_{\mathcal{P}, D_{\tau', i'}}^{\epsilon, \delta}$ in Game 5 using s samples from $D_{\tau', i'}$ since $D_{\tau', i'}(G(\cdot))$ can be interpreted as a mapping for s samples from $D_{\tau', i'}$. Then, from the SIM-MDD security with private simulation of ELWMPRF, we have $|\Pr[\text{SUC}_5] - 1/2| = \text{negl}(\lambda)$. From the above, we also have $\text{Adv}_{i^*, \mathcal{A}, \text{QELWMPRF}}^{\text{q-sim-mdd}}(\lambda) \leq O(q^3/s) + 2\gamma$ for some negligible function γ . Thus, by setting $s = O(q^3 \cdot w^2)$, we obtain $\text{Adv}_{i^*, \mathcal{A}, \text{QELWMPRF}}^{\text{q-sim-mdd}}(\lambda) \leq 1/w$.

Since w is any polynomial, this means that $\text{Adv}_{i^*, \mathcal{A}, \text{QELWMPRF}}^{\text{q-sim-mdd}}(\lambda) = \text{negl}(\lambda)$.

Remark A.6. It is easy to see that the extended weak pseudorandomness of ELWMPRF is preserved after we apply the transformation above since the evaluation algorithm is the same as that of ELWMPRF and extended weak pseudorandomness holds against adversaries that generate pp. Thus, we omit a formal proof.

B Puncturable Encryption with Strong Ciphertext Pseudorandomness

We prove Theorem 7.5 in this section.

B.1 Tools for PE

Definition B.1 (Statistically Injective PPRF). *If a PPRF family $\mathcal{F} = \{F_K : \{0, 1\}^{\ell_1(\lambda)} \rightarrow \{0, 1\}^{\ell_2(\lambda)} \mid K \in \{0, 1\}^\lambda\}$ satisfies the following, we call it a statistically injective PPRF family with failure probability $\epsilon(\cdot)$. With probability $1 - \epsilon(\lambda)$ over the random choice of $K \leftarrow \text{PRF.Gen}(1^\lambda)$, for all $x, x' \in \{0, 1\}^{\ell_1(\lambda)}$, if $x \neq x'$, then $F_K(x) \neq F_K(x')$. If $\epsilon(\cdot)$ is not specified, it is a negligible function.*

Sahai and Waters show that we can convert any PPRF into a statistically injective PPRF [SW21].

Theorem B.2 ([SW21]). *If OWFs exist, then for all efficiently computable functions $n(\lambda)$, $m(\lambda)$, and $e(\lambda)$ such that $m(\lambda) \geq 2n(\lambda) + e(\lambda)$, there exists a statistically injective PPRF family with failure probability $2^{-e(\lambda)}$ that maps $n(\lambda)$ bits to $m(\lambda)$ bits.*

Definition B.3. *An injective bit-commitment with setup consists of PPT algorithms (Gen, Com) .*

$\text{Gen}(1^\lambda)$: *The key generation algorithm takes as input the security parameter 1^λ and outputs a commitment key ck .*

$\text{Com}_{\text{ck}}(b)$: *The commitment algorithm takes as input ck and a bit b and outputs a commitment com .*

These satisfy the following properties.

Computationally Hiding: *For any QPT \mathcal{A} , it holds that*

$$\left| \Pr[\mathcal{A}(\text{Com}_{\text{ck}}(0)) = 1 \mid \text{ck} \leftarrow \text{Gen}(1^\lambda)] - \Pr[\mathcal{A}(\text{Com}_{\text{ck}}(1)) = 1 \mid \text{ck} \leftarrow \text{Gen}(1^\lambda)] \right| \leq \text{negl}(\lambda).$$

Statistically Binding: *It holds that*

$$\Pr \left[\text{com}_0 = \text{com}_1 \left| \begin{array}{l} \text{ck} \leftarrow \text{Gen}(1^\lambda) \\ \text{com}_0 \leftarrow \text{Com}_{\text{ck}}(0) \\ \text{com}_1 \leftarrow \text{Com}_{\text{ck}}(1) \end{array} \right. \right] \leq \text{negl}(\lambda).$$

Injective: *For every security parameter λ , there is a bound ℓ_r on the number of random bits used by Com such that if $\text{ck} \leftarrow \text{Gen}(1^\lambda)$, $\text{Com}_{\text{ck}}(\cdot; \cdot)$ is an injective function on $\{0, 1\} \times \{0, 1\}^{\ell_r}$ except negligible probability.*

Theorem B.4. *If the QLWE assumption holds, there exists a secure injective bit-commitment with setup.*

This theorem follows from the following theorems.

Theorem B.5 ([Nao91]). *If there exists (injective) OWFs, there exists (injective) bit-commitment.*

Theorem B.6 ([PW11, AKPW13, Adapted]). *If the QLWE assumption holds, there exists a secure injective OWF with evaluation key generation algorithms.*

Remark B.7. The injective OWFs achieved in Theorem B.6 needs evaluation key generation algorithms unlike the standard definition of OWFs. However, OWFs with evaluation key generation algorithms are sufficient for proving Theorem B.4 by using Theorem B.5 since we use *commitment key generation algorithm* Gen (i.e., setup) in Definition B.3. Note that there is no post-quantum secure injective OWF *without evaluation key generation algorithm* so far.

B.2 PE Scheme Description

We review the puncturable encryption scheme by Cohen et al. [CHN⁺18]. We can see Theorem 7.5 holds by inspecting their PE scheme. The scheme utilizes the following ingredients and the length n of ciphertexts is 12 times the length ℓ of plaintexts:

- A length-doubling PRG : $\{0, 1\}^\ell \rightarrow \{0, 1\}^{2\ell}$
- An injective PPRFs (See Definition B.1) $F : \{0, 1\}^{3\ell} \rightarrow \{0, 1\}^{9\ell}$.
- A PPRF $G : \{0, 1\}^{9\ell} \rightarrow \{0, 1\}^\ell$.
- An injective bit-commitment with setup (Com.Gen, Com) using randomness in $\{0, 1\}^{9\ell}$. We only use this in our security proof.

Scheme. The scheme PE by Cohen et al. [CHN⁺18] is as follows.

Gen(1^λ): Sample functions F and G , generates pe.ek as the obfuscated circuit $i\mathcal{O}(E)$ where E is described in Figure 12, and returns $(\text{pe.ek}, \text{pe.dk}) := (i\mathcal{O}(E), D)$, where pe.dk is the (un-obfuscated) program D in Figure 13.

Puncture($\text{pe.dk}, c^*$): Output $\text{pe.dk}_{\neq c^*}$, where $\text{pe.dk}_{\neq c^*}$ is the obfuscated circuits $i\mathcal{O}(D_{\neq c^*})$ where $D_{\neq c^*}$ is described in Figure 14, that is, $\text{pe.dk}_{\neq c^*} := i\mathcal{O}(D_{\neq c^*})$.

Enc($\text{pe.ek}, m$): Take $m \in \{0, 1\}^\ell$, sample $s \leftarrow \{0, 1\}^\ell$, and outputs $c \leftarrow \text{pe.ek}(m, s)$.

Dec($\text{pe.dk}, c$): Take $c \in \{0, 1\}^{12\ell}$ and returns $m := \text{pe.dk}(c)$.

The size of the circuits is appropriately padded to be the maximum size of all modified circuits, which will appear in the security proof.

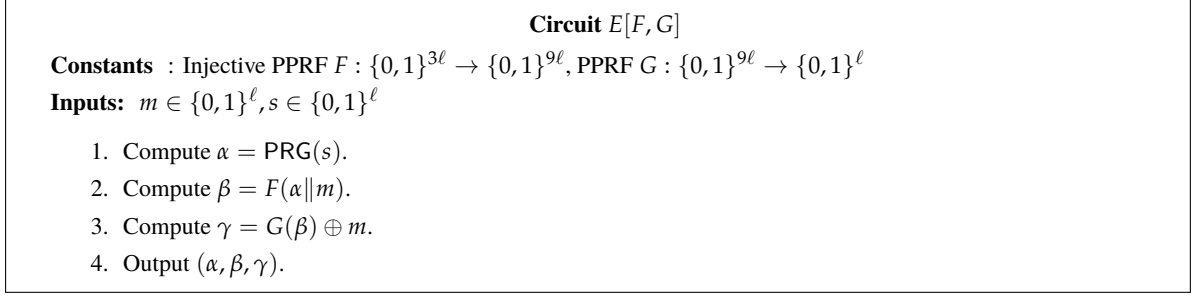


Figure 12: Description of encryption circuit E




Figure 13: Description of decryption circuit D

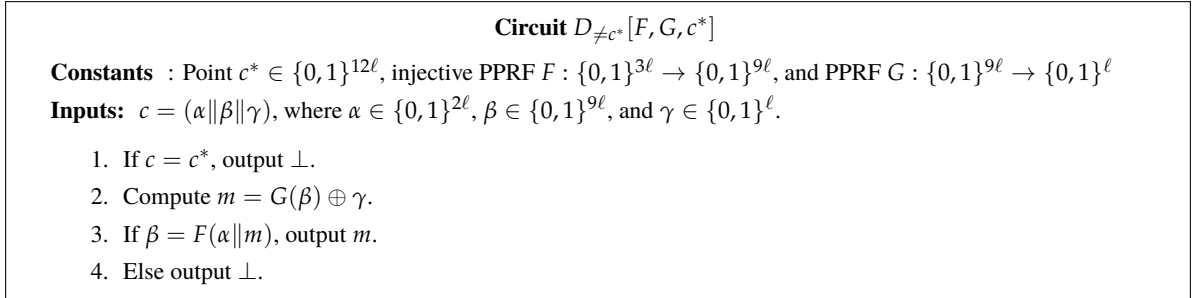


Figure 14: Description of punctured decryption circuit $D_{\neq c^*}$ at c^*

B.3 PE Security Proof

Cohen et al. [CHN⁺18] proved correctness, punctured correctness, and sparseness of PE above by using secure PRG PRG, secure injective PPRF F , secure PPRF G , and secure IO $i\mathcal{O}$. Thus, we complete the proof of Theorem 7.5 by combining Theorems B.2 and B.4, and Theorem B.8 below, which we prove in this section.

Theorem B.8. *If PRG is a secure PRG, F is a secure injective PPRF, G is a secure PPRF, Com is a secure injective bit-commitment with setup, and $i\mathcal{O}$ is a secure IO, then PE is a secure PE that satisfies strong ciphertext pseudorandomness.*

Proof of Theorem B.8. To prove $x_0 := c^* \leftarrow \text{Enc}(\text{pe.ek}, m^*)$ is indistinguishable from $x_1 := r^* \leftarrow \{0, 1\}^\ell$, we define a sequence hybrid games.

Real: This is the same as the real game with $b = 0$. That is, for queried m^* the challenger does the following.

1. Choose an injective PPRF $F : \{0, 1\}^{3\ell} \rightarrow \{0, 1\}^{9\ell}$ and PPRF $G : \{0, 1\}^{9\ell} \rightarrow \{0, 1\}^\ell$.

2. Choose $s \leftarrow \{0, 1\}^\ell$ and compute $\alpha_0 := \text{PRG}(s)$, $\beta_0 := F(\alpha_0 \| m^*)$, and $\gamma_0 := G(\beta_0) \oplus m^*$.
3. Set $x_0 := \alpha_0 \| \beta_0 \| \gamma_0$ and computes $\text{pe.ek} := i\mathcal{O}(E)$ and $\text{pe.dk}_{\neq x_0} := i\mathcal{O}(D_{\neq x_0})$.
4. Send $(x_0, \text{pe.ek}, \text{pe.dk}_{\neq x_0})$ to the adversary.

Hyb₁: This is the same as Hyb₀(0) except that α_0 is uniformly random.

Hyb₂: This is the same as Hyb₁ except that we use punctured $F_{\neq \alpha_0 \| m^*}$ and modified circuits $E_{\neq \alpha_0 \| m^*}$ and $D_{\neq \alpha_0 \| m^*}^2$ described in Figures 16 and 17. Intuitively, these modified circuits are punctured at input $\alpha_0 \| m^*$ and use exceptional handling for this input.

Hyb₃: This is the same as Hyb₂ except that $\beta_0 \leftarrow \{0, 1\}^{9\ell}$.

Hyb₄: This is the same as Hyb₃ except that we use punctured $G_{\neq \beta_0}$ and modified circuits $E_{\neq \alpha_0 \| m^*, \neq \beta_0}$ and $D_{\neq \alpha_0 \| m^*, \neq \beta_0}^4$ described in Figures 18 and 19. Intuitively, these modified circuits are punctured at input β_0 and use $F_{\neq \alpha_0 \| m^*}$ and exceptional handling for β_0 .

Hyb₅ = Rand₂: This is the same as Hyb₄ except that γ_0 is uniformly random. Now, $\alpha_0, \beta_0, \gamma_0$ are uniformly random and we rewrite them into $\alpha_1, \beta_1, \gamma_1$, respectively. For ease of notation, we also denote this game by Rand₂.

Rand₁: This is the same as Hyb₅ = Rand₂ except that we use un-punctured G , circuit $E_{\neq \alpha_0 \| m^*, \neq \beta_0}$ reverts to $E_{\neq \alpha_1 \| m^*}$ described in Figure 16, and we change circuit $D_{\neq \alpha_1 \| m^*, \neq \beta_1}^4$ into $D_{\neq \alpha_1 \| m^*}^r$ described in Figure 21.

Rand: This is the same as the real game with $b = 1$. That is, for queried m^* the challenger does the following.

1. Choose an injective PPRF $F : \{0, 1\}^{3\ell} \rightarrow \{0, 1\}^{9\ell}$ and PPRF $G : \{0, 1\}^{9\ell} \rightarrow \{0, 1\}^\ell$.
2. Choose $\alpha_1 \leftarrow \{0, 1\}^{2\ell}$, $\beta_1 \leftarrow \{0, 1\}^{9\ell}$, and $\gamma_1 \leftarrow \{0, 1\}^\ell$.
3. Set $x_1 := \alpha_1 \| \beta_1 \| \gamma_1$ and computes $\text{pe.ek} := i\mathcal{O}(E)$ and $\text{pe.dk}_{\neq x_1} := i\mathcal{O}(D_{\neq x_1})$.
4. Send $(x_1, \text{pe.ek}, \text{pe.dk}_{\neq x_1})$ to the adversary.

We described the overview of these hybrid games in Figure 15. If we prove these hybrid games are

	α^*	β^*	γ^*	$\text{pe.ek} := i\mathcal{O}(\cdot)$	$\text{pe.dk} := i\mathcal{O}(\cdot)$
Real	PRG(s)	$F(\alpha_0 \ m^*)$	$G(\beta_0) \oplus m^*$	E	$D_{\neq x_0}$
Hyb ₁	\$	$F(\alpha_0 \ m^*)$	$G(\beta_0) \oplus m^*$	E	$D_{\neq x_0}$
Hyb ₂	\$	$F(\alpha_0 \ m^*)$	$G(\beta_0) \oplus m^*$	$E_{\neq \alpha_0 \ m^*}$	$D_{\neq \alpha_0 \ m^*}^2 [F_{\neq \alpha_0 \ m^*}]$
Hyb ₃	\$	\$	$G(\beta_0) \oplus m^*$	$E_{\neq \alpha_0 \ m^*}$	$D_{\neq \alpha_0 \ m^*}^2 [F_{\neq \alpha_0 \ m^*}]$
Hyb ₄	\$	\$	$G(\beta_0) \oplus m^*$	$E_{\neq \alpha_0 \ m^*, \neq \beta_0}$	$D_{\neq \alpha_0 \ m^*, \neq \beta_0}^4 [F_{\neq \alpha_0 \ m^*}, G_{\neq \beta_0}]$
Hyb ₅	\$	\$	\$	$E_{\neq \alpha_1 \ m^*, \neq \beta_1}$	$D_{\neq \alpha_1 \ m^*, \neq \beta_1}^4 [F_{\neq \alpha_1 \ m^*}, G_{\neq \beta_1}]$
Rand ₁	\$	\$	\$	$E_{\neq \alpha_1 \ m^*}$	$D_{\neq \alpha_1 \ m^*}^r [F_{\neq \alpha_1 \ m^*}]$
Rand	\$	\$	\$	E	$D_{\neq x_1}$

Figure 15: High-level overview of hybrid games from Real to Rand. Recall that Hyb₅ = Rand₂. Transitions from Rand₂ to Rand are basically the reverse transitions from Hyb₀ to Hyb₄, but there are subtle differences.

indistinguishable, we complete the proof of Theorem B.8. ■

We prove that those hybrid games in Figure 15 are indistinguishable by Lemmata B.9, B.10, B.14, B.15, B.19, B.20 and B.25.




Figure 16: Description of encryption circuit $E_{\neq \alpha^* \| m^*}$

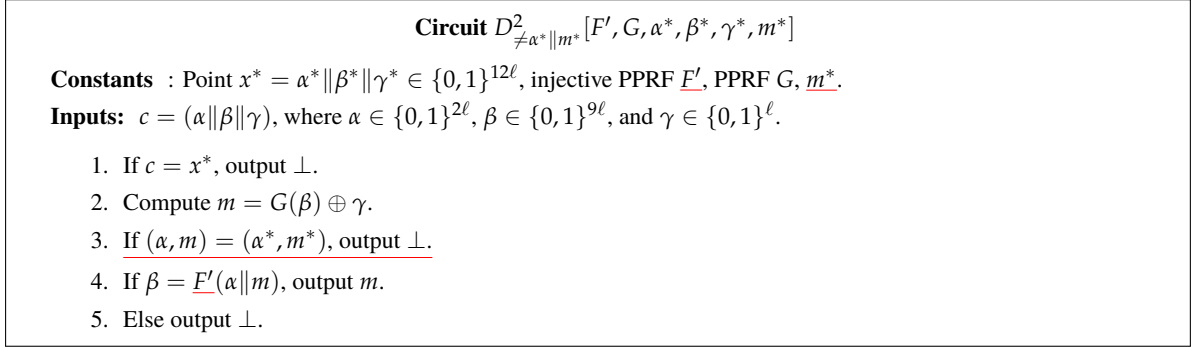


Figure 17: Description of punctured decryption circuit $D_{\neq \alpha^* \| m^*}^2$

From Real to Hyb₅. We first move from Real to Hyb₅.

Lemma B.9. *If PRG is a secure PRG, it holds that $|\Pr[\text{Hyb}_0(0) = 1] - \Pr[\text{Hyb}_1 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Lemma B.9. The randomness s for encryption is never used anywhere except $\alpha_0 := \text{PRG}(s)$. We can apply the PRG security and immediately obtain the lemma. ■

Lemma B.10. *If $i\mathcal{O}$ is a secure IO and F is a secure injective PPRF, it holds that*

$$|\Pr[\text{Hyb}_1 = 1] - \Pr[\text{Hyb}_2 = 1]| \leq \text{negl}(\lambda).$$

Proof of Lemma B.10. We change E and $D_{\neq x_0}$ into $E_{\neq \alpha_0 \| m^*}$ and $D_{\neq \alpha_0 \| m^*}^2$, respectively.

We define a sequence of sub-hybrid games.

Hyb₁¹: This is the same as Hyb₁ except that we generate $F_{\neq \alpha_0 \| m^*}$ and set $F' := F_{\neq \alpha_0 \| m^*}$ and $\text{pe.ek} := i\mathcal{O}(E_{\neq \alpha_0 \| m^*})$ described in Figure 16.

Hyb₁²: This is the same as Hyb₁¹ except that we set $\text{pe.dk}_{\neq x_0} := i\mathcal{O}(D_{\neq \alpha_0 \| m^*}^2[F, G, \alpha_0, \beta_0, \gamma_0, m^*])$ described in Figure 17. That is, we still use F , but modify the circuit.

Proposition B.11. *If $i\mathcal{O}$ is a secure IO, it holds that $|\Pr[\text{Hyb}_1 = 1] - \Pr[\text{Hyb}_1^1 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Proposition B.11. In these games, value $\alpha_0 \leftarrow \{0, 1\}^{2\ell}$ is not in the image of PRG except with negligible probability. The only difference between the two games is that $F_{\neq \alpha_0 \| m^*}$ is used in Hyb₁¹. Thus, E and $E_{\neq \alpha_0 \| m^*}$ are functionally equivalent except with negligible probability. We can obtain the proposition by applying the IO security. ■

Circuit $E_{\neq \alpha^* \| m^*, \neq \beta^*}[F', G']$

Constants : Injective PPRF F' , PPRF G'

Inputs: $m \in \{0, 1\}^\ell, s \in \{0, 1\}^\ell$

1. Compute $\alpha = \text{PRG}(s)$.
2. Compute $\beta = F'(\alpha \| m)$.
3. Compute $\gamma = G'(\beta) \oplus m$.
4. Output (α, β, γ) .

Figure 18: Description of encryption circuit $E_{\neq \alpha^* \| m^*, \neq \beta^*}$

Proposition B.12. *If $i\mathcal{O}$ is a secure IO and F is injective, it holds that $|\Pr[\text{Hyb}_1^1 = 1] - \Pr[\text{Hyb}_1^2 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Proposition B.12. We analyze the case where $(\alpha, m) = (\alpha_0, m^*)$ since it is the only difference between $D_{\neq x_0}$ and $D_{\neq \alpha_0 \| m^*}^2$.

- If $c = x_0$, $D_{\neq \alpha_0 \| m^*}^2$ outputs \perp by the first line of the description. Thus, the output of $D_{\neq \alpha_0 \| m^*}^2(x_0)$ is the same as that of $D_{\neq x_0}(x_0)$.
- If $c \neq x_0$, it holds $(\beta_0, \gamma_0) \neq (\beta, \gamma)$ in this case. However, it should be $\beta_0 = \beta$ due to the injectivity of F and $\beta_0 = F(\alpha_0 \| m^*)$. Thus, both $D_{\neq x_0}(c)$ and $D_{\neq \alpha_0 \| m^*}^2(c)$ output \perp in this case ($D_{\neq x_0}(c)$ outputs \perp at the first line).

Therefore, $D_{\neq x_0}$ and $D_{\neq \alpha_0 \| m^*}^2$ are functionally equivalent. We can obtain the proposition by applying the IO security. ■

Proposition B.13. *If $i\mathcal{O}$ is a secure IO, it holds that $|\Pr[\text{Hyb}_1^2 = 1] - \Pr[\text{Hyb}_2 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Proposition B.13. Due to the exceptional handling in the third item of $D_{\neq \alpha_0 \| m^*}^2$, $F(\alpha \| m)$ is never computed for input (α_0, m^*) . Thus, even if we use $F_{\neq \alpha_0 \| m^*}$ instead of F , $D_{\neq \alpha_0 \| m^*}^2[F]$ and $D_{\neq \alpha_0 \| m^*}^2[F_{\neq \alpha_0 \| m^*}]$ are functionally equivalent. We can obtain the proposition by the IO security. ■

We complete the proof of Lemma B.10. ■

Lemma B.14. *If F is a secure injective PPRF, it holds that $|\Pr[\text{Hyb}_2 = 1] - \Pr[\text{Hyb}_3 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Lemma B.14. The difference between these two games is that β_0 is $F(\alpha_0 \| m^*)$ or random. We can immediately obtain the lemma by applying punctured pseudorandomness of F since we use $F_{\neq \alpha_0 \| m^*}$ in these games. ■

Lemma B.15. *If $i\mathcal{O}$ is a secure IO and F is a secure injective PPRF, it holds that*

$$|\Pr[\text{Hyb}_3 = 1] - \Pr[\text{Hyb}_4 = 1]| \leq \text{negl}(\lambda).$$

Proof of Lemma B.15. We change $E_{\neq \alpha^* \| m^*}$ and $D_{\neq \alpha^* \| m^*}^2$ into $E_{\neq \alpha^* \| m^*, \neq \beta^*}$ and $D_{\neq \alpha^* \| m^*, \neq \beta^*}^4$, respectively.

We define a sequence of sub-hybrid games.

Hyb₃¹: This is the same as **Hyb₃** except that we use punctured $G_{\neq \beta_0}$ and set $\text{pe.ek} := i\mathcal{O}(E_{\neq \alpha_0 \| m^*, \neq \beta_0}[F_{\neq \alpha_0 \| m^*}, G_{\neq \beta_0}])$.

Hyb₃²: This is the same as **Hyb₃¹** except that we still use G but set $\text{pe.dk}_{\neq c_0} := i\mathcal{O}(D_{\neq \alpha_0 \| m^*, \neq \beta_0}^4[F_{\neq \alpha_0 \| m^*}, G])$.

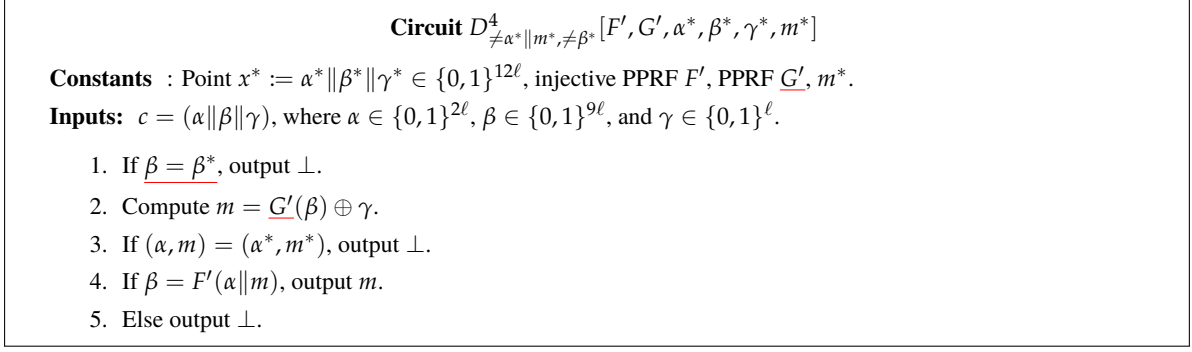


Figure 19: Description of punctured decryption circuit $D_{\neq \alpha^* \| m^*, \neq \beta^*}^4$

Proposition B.16. *If $i\mathcal{O}$ is a secure IO, it holds that $|\Pr[\text{Hyb}_3 = 1] - \Pr[\text{Hyb}_3^1 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Proposition B.16. In these games $\beta_0 \leftarrow \{0, 1\}^{9\ell}$ is uniformly random. By the sparsity of F , β_0 is not in the image of F except with negligible probability. Thus, $E_{\neq \alpha_0 \| m^*}$ and $E_{\neq \alpha_0 \| m^*, \neq \beta_0}$ are functionally equivalent except with negligible probability. We obtain the proposition by the IO security. ■

Proposition B.17. *If $i\mathcal{O}$ is a secure IO, it holds that $|\Pr[\text{Hyb}_3^1 = 1] - \Pr[\text{Hyb}_3^2 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Proposition B.17. The difference between $D_{\neq \alpha_0 \| m^*}^2$ and $D_{\neq \alpha_0 \| m^*, \neq \beta_0}^4$ is that we replace “If $c = x_0$, outputs \perp .” with “If $\beta = \beta_0$, outputs \perp .”. In these games, $\beta_0 \leftarrow \{0, 1\}^{9\ell}$ is not in the image of F except with negligible probability. Recall that $c = x_0$ means $c = \alpha_0 \| \beta_0 \| \gamma_0$. Thus, those two circuits may differ when $\beta = \beta_0$ but $(\alpha, \gamma) \neq (\alpha_0, \gamma_0)$. However, it does not happen $\beta = F'(\alpha \| (G(\beta) \oplus \gamma))$ in this case due to the injectivity of F . Thus, $D_{\neq \alpha_0 \| m^*}^2$ and $D_{\neq \alpha_0 \| m^*, \neq \beta_0}^4$ are functionally equivalent and we obtain the proposition by applying the IO security. ■

Proposition B.18. *If $i\mathcal{O}$ is a secure IO, it holds that $|\Pr[\text{Hyb}_3^2 = 1] - \Pr[\text{Hyb}_4 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Proposition B.18. The difference between these two games that we use $D_{\neq \alpha_0 \| m^*, \neq \beta_0}^4 [F_{\neq \alpha_0 \| m^*}, G_{\neq \beta_0}]$ instead of $D_{\neq \alpha_0 \| m^*, \neq \beta_0}^4 [F_{\neq \alpha_0 \| m^*}, G]$. However, $G_{\neq \beta_0}(\beta_0)$ is never computed by the first item of $D_{\neq \alpha_0 \| m^*, \neq \beta_0}^4$. We obtain the proposition by the IO security. ■

We complete the proof of Lemma B.15. ■

Lemma B.19. *If G is a secure PPRF, it holds that $|\Pr[\text{Hyb}_4 = 1] - \Pr[\text{Hyb}_5 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Lemma B.19. The difference between these two games is that γ_0 is $G(\beta_0)$ or random. We can immediately obtain the lemma by applying punctured pseudorandomness of G since we use $G_{\neq \beta_0}$ in these games. ■

In Hyb_5 , α_0 , β_0 , and γ_0 are uniformly random strings as α_1 , β_1 , and γ_1 .

From Rand to Hyb_5 . We leap to Rand and move from Rand to $\text{Rand}_2 = \text{Hyb}_5$ instead of directly moving from $\text{Hyb}_5 = \text{Rand}_2$ to Rand since $\text{Real} \approx \text{Hyb}_5$ and $\text{Rand}_2 \approx \text{Rand}$ is almost symmetric (but not perfectly symmetric).

Lemma B.20. *If $i\mathcal{O}$ is a secure IO and F is a secure injective PPRF, it holds that*

$$|\Pr[\text{Rand} = 1] - \Pr[\text{Rand}_1 = 1]| \leq \text{negl}(\lambda).$$

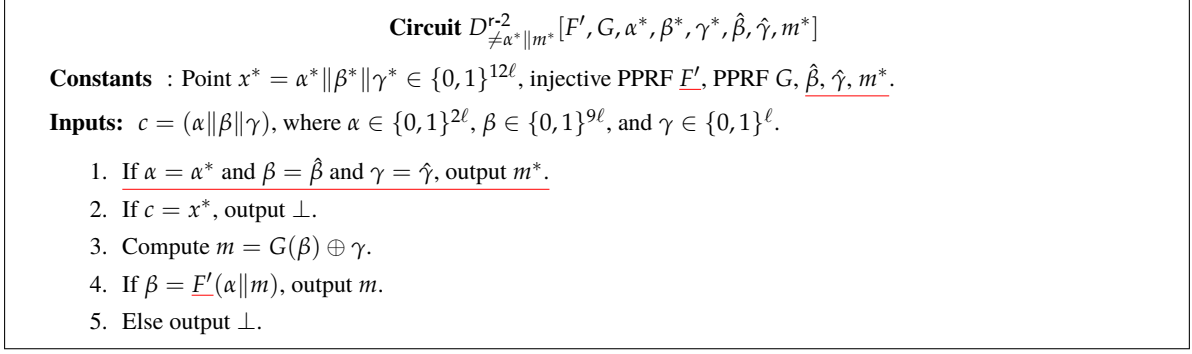


Figure 20: Description of punctured decryption circuit $D_{\neq \alpha^* \| m^*}^{r-2}$

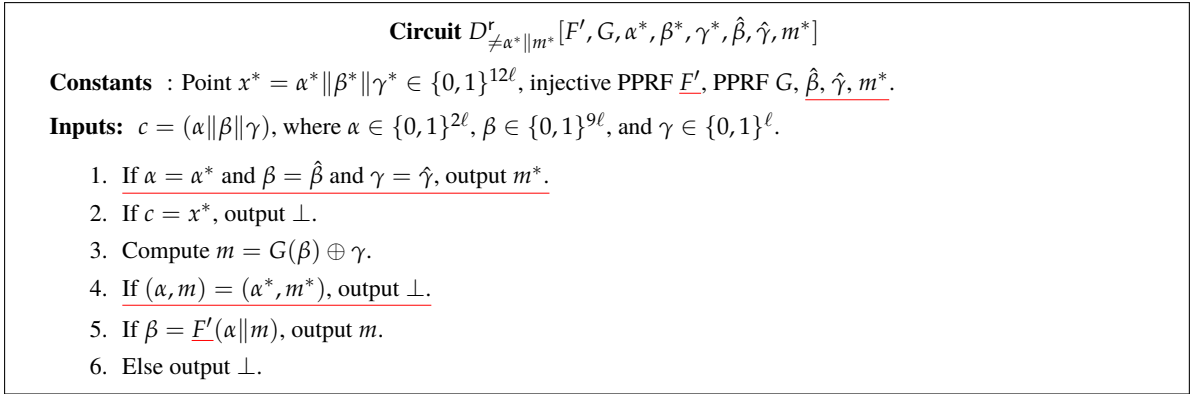


Figure 21: Description of punctured decryption circuit $D_{\neq \alpha^* \| m^*}^r$

Proof of Lemma B.20. We change E and $D_{\neq x_1}$ into $E_{\neq \alpha_1 \| m^*}$ and $D_{\neq \alpha_1 \| m^*}^r$, respectively.

We define a sequence of sub-hybrid games.

rHyb¹: This is the same as Rand except that we generate $F_{\neq \alpha_1 \| m^*}$ and set $F' := F_{\neq \alpha_1 \| m^*}$ and $\text{pe.ek} := i\mathcal{O}(E_{\neq \alpha_1 \| m^*})$ described in Figure 16.

rHyb²: This is the same as rHyb¹ except that we set $\text{pe.dk}_{\neq x_1} := i\mathcal{O}(D_{\neq \alpha_1 \| m^*}^{r-2}[F, G])$ described in Figure 20. That is, we still use F , but the modified circuit that outputs m^* for input $\alpha_1 \| \hat{\beta} \| \hat{\gamma}$, where $\hat{\beta} := F(\alpha_1 \| m^*)$ and $\hat{\gamma} := G(\hat{\beta}) \oplus m^*$.

rHyb³: This is the same as rHyb² except that we set $\text{pe.dk}_{\neq x_1} := i\mathcal{O}(D_{\neq \alpha_1 \| m^*}^r[F, G])$ described in Figure 21. That is, we still use F , but the modified circuit outputs \perp for an input such that $(\alpha, m) = (\alpha_1, m^*)$.

Proposition B.21. *If $i\mathcal{O}$ is a secure IO, it holds that $\left| \Pr[\text{Rand} = 1] - \Pr[\text{rHyb}^1 = 1] \right| \leq \text{negl}(\lambda)$.*

Proof of Proposition B.21. In these games, value $\alpha_1 \leftarrow \{0, 1\}^{2\ell}$ is not in the image of PRG except with negligible probability. Thus, E and $E_{\neq \alpha_1 \| m^*}$ are functionally equivalent except with negligible probability. We can obtain the proposition by applying the IO security. ■

Proposition B.22. *If $i\mathcal{O}$ is a secure IO and F is injective, it holds that $\left| \Pr[\text{rHyb}^1 = 1] - \Pr[\text{rHyb}^2 = 1] \right| \leq \text{negl}(\lambda)$.*

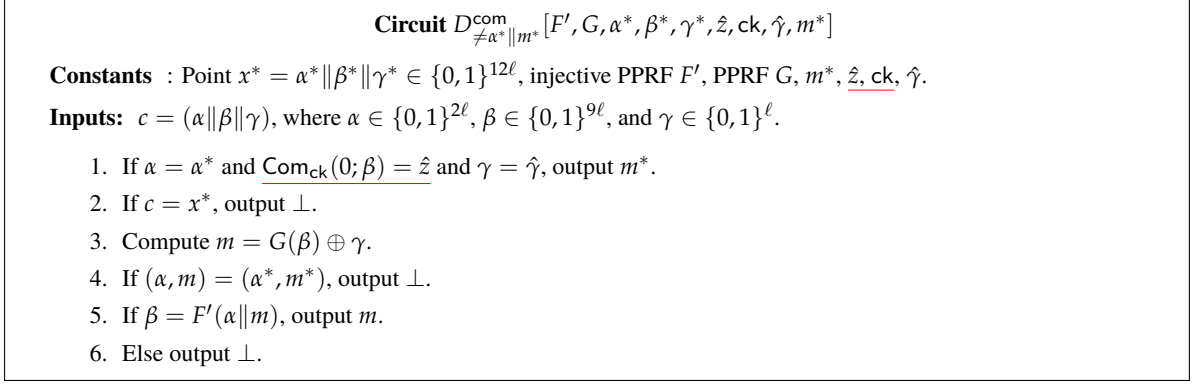


Figure 22: Description of punctured decryption circuit $D_{\neq \alpha^* \| m^*}^{\text{com}}$

Proof of Proposition B.22. The difference between $D_{\neq x_1}$ and $D_{\neq \alpha_1 \| m^*}^{r-2}$ is “If $\alpha = \alpha^*$ and $\beta = \hat{\beta}$ and $\gamma = \hat{\gamma}$, output m^* .”. Although $\alpha_1 \| \hat{\beta} \| \hat{\gamma}$ is a valid encryption, $\hat{\beta} = F(\alpha_1 \| m^*)$ is not equal to β_1 except with negligible probability since β_1 is uniformly random. Similarly, $\hat{\gamma}$ is not equal to γ_1 except with negligible probability. Thus, $D_{\neq x_1}(\alpha_1 \| \hat{\beta} \| \hat{\gamma})$ outputs m^* . That is, $D_{\neq x_1}$ and $D_{\neq \alpha_1 \| m^*}^{r-2}$ are functionally equivalent. We can obtain the proposition by applying the IO security. ■

Proposition B.23. *If $i\mathcal{O}$ is a secure IO and F is injective, it holds that $|\Pr[\text{rHyb}^2 = 1] - \Pr[\text{rHyb}^3 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Proposition B.23. We analyze the case where $(\alpha, m) = (\alpha_1, m^*)$. We can reach the fourth line of $D_{\neq \alpha_1 \| m^*}^r$ if $c \neq x_1$. If $c \neq x_1$ and $(\alpha, m) = (\alpha_1, m^*)$, it holds that $(\beta, \gamma) \neq (\beta_1, \gamma_1)$. However, it should be $\beta_1 = \beta$ in this case due to the injectivity of F . That is, if $D_{\neq \alpha_1 \| m^*}^r(c)$ outputs \perp at the fourth line, $D_{\neq \alpha_1 \| m^*}^{r-2}(c)$ also outputs \perp at the second line. Therefore, $D_{\neq \alpha_1 \| m^*}^{r-2}$ and $D_{\neq \alpha_1 \| m^*}^r$ are functionally equivalent. We can obtain the proposition by applying the IO security. ■

Proposition B.24. *If $i\mathcal{O}$ is a secure IO, it holds that $|\Pr[\text{rHyb}^3 = 1] - \Pr[\text{Rand}_1 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Proposition B.24. Due to the exceptional handling in the fourth line of $D_{\neq \alpha_1 \| m^*}^r$, $F(\alpha \| m)$ is never computed for input (α_1, m^*) . Thus, even if we use $F_{\neq \alpha_1 \| m^*}$ instead of F , $D_{\neq \alpha_1 \| m^*}^r[F]$ and $D_{\neq \alpha_1 \| m^*}^r[F_{\neq \alpha_1 \| m^*}]$ are functionally equivalent. We can obtain the proposition by the IO security. ■

We complete the proof of Lemma B.20. ■

Lemma B.25. *If $i\mathcal{O}$ is a secure IO, F is a secure injective PPRF, and $(\text{Com.Gen}, \text{Com})$ is a secure injective bit-commitment with setup, it holds that $|\Pr[\text{Rand}_1 = 1] - \Pr[\text{Rand}_2 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Lemma B.25. We change $E_{\neq \alpha^* \| m^*}$ and $D_{\neq \alpha^* \| m^*}^r$ into $E_{\neq \alpha^* \| m^*, \neq \beta^*}$ and $D_{\neq \alpha^* \| m^*, \neq \beta^*}^4$, respectively.

We define a sequence of sub-hybrid games.

rHyb_1^1 : This is the same as Rand_1 except that we use $\hat{\beta} \leftarrow \{0, 1\}^{9\ell}$ instead of $F(\alpha_1 \| m^*)$.

rHyb_1^2 : This is the same as rHyb_1^1 except that we use $D_{\neq \alpha_1 \| m^*}^{\text{com}}$ described in Figure 22, where $\text{ck} \leftarrow \text{Com.Gen}(1^\lambda)$ and $\hat{z} = \text{Com}_{\text{ck}}(0; \hat{\beta})$ are hardwired, instead of $D_{\neq \alpha_1 \| m^*}^r$.

rHyb_1^3 : This is the same as rHyb_1^2 except that we hard-code $\hat{z} = \text{Com}_{\text{ck}}(1; \hat{\beta})$ into $D_{\neq \alpha_1 \| m^*}^{\text{com}}$ instead of $\text{Com}_{\text{ck}}(0; \hat{\beta})$.

rHyb_1^4 : This is the same as rHyb_1^3 except that we use $D_{\neq \alpha_1 \| m^*}^F$ described in Figure 23

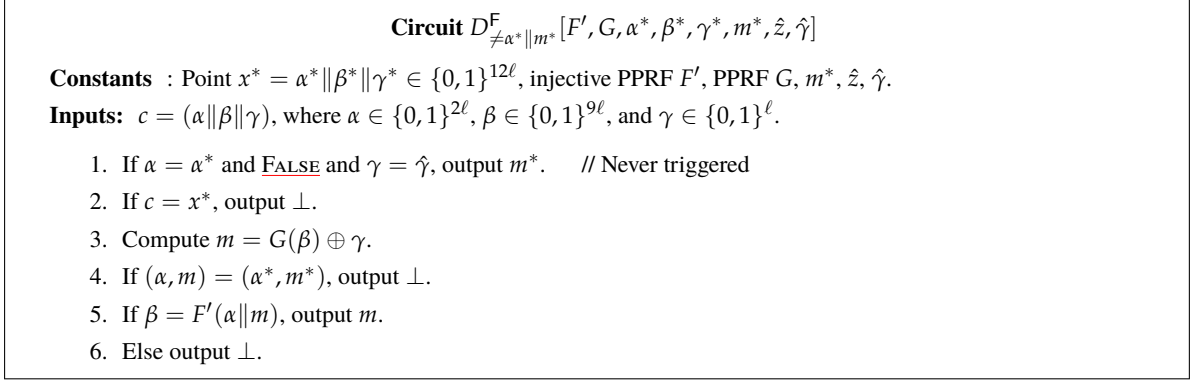


Figure 23: Description of punctured decryption circuit $D_{\neq \alpha^* \| m^*}^F$

rHyb_1^5 : This is the same as rHyb_1^4 except that we use punctured $G_{\neq \beta_1}$ and set $\text{pe.ek} := i\mathcal{O}(E_{\neq \alpha_1 \| m^*, \neq \beta_1}[F_{\neq \alpha_1 \| m^*}, G_{\neq \beta_1}])$.

rHyb_1^6 : This is the same as rHyb_1^5 except that we still use G but set $\text{pe.dk}_{\neq c_1} := i\mathcal{O}(D_{\neq \alpha_1 \| m^*, \neq \beta_1}^4[F_{\neq \alpha_1 \| m^*}, G])$.

Proposition B.26. *If F is a secure PPRF, it holds that $|\Pr[\text{Rand}_1 = 1] - \Pr[\text{rHyb}_1^1 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Proposition B.26. In these games, we use $F_{\neq \alpha_1 \| m^*}$ in $E_{\neq \alpha_1 \| m^*}$ and $D_{\neq \alpha_1 \| m^*}^r$. Thus, we can apply the punctured pseudorandomness and immediately obtain the proposition. ■

Proposition B.27. *If $i\mathcal{O}$ is a secure IO and Com_{ck} is injective, it holds that*

$$|\Pr[\text{rHyb}_1^1 = 1] - \Pr[\text{rHyb}_1^2 = 1]| \leq \text{negl}(\lambda).$$

Proof of Proposition B.27. The difference between $D_{\neq \alpha_1 \| m^*}^{\text{com}}$ and $D_{\neq \alpha_1 \| m^*}^r$ is whether we use “ $\text{Com}_{\text{ck}}(0; \beta) = \hat{z}$ ” or “ $\beta = \hat{\beta}$ ”, where $\hat{z} = \text{Com}_{\text{ck}}(0; \hat{\beta})$ and $\text{ck} \leftarrow \text{Com.Gen}(1^\lambda)$. Since Com is injective, these two conditions are equivalent. Therefore, those two circuits are functionally equivalent. We obtain the proposition by applying the IO security. ■

Proposition B.28. *If $(\text{Com.Gen}, \text{Com})$ is computationally hiding, it holds that*

$$|\Pr[\text{rHyb}_1^2 = 1] - \Pr[\text{rHyb}_1^3 = 1]| \leq \text{negl}(\lambda).$$

Proof of Proposition B.28. The only difference between these two games is that $\hat{z} = \text{Com}_{\text{ck}}(0; \hat{\beta})$ or $\hat{z} = \text{Com}_{\text{ck}}(1; \hat{\beta})$. Note that $\hat{\beta}$ is never used anywhere else. We can obtain the proposition by the hiding property of Com . ■

Proposition B.29. *If $i\mathcal{O}$ is a secure IO and $(\text{Com.Gen}, \text{Com})$ is statistically binding, it holds that $|\Pr[\text{rHyb}_1^3 = 1] - \Pr[\text{rHyb}_1^4 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Proposition B.29. The difference between $D_{\neq \alpha_1 \| m^*}^F$ and $D_{\neq \alpha_1 \| m^*}^{\text{com}}$ is that the first line of $D_{\neq \alpha_1 \| m^*}^F$ is never executed. However, $\hat{z} = \text{Com}_{\text{ck}}(1; \hat{\beta})$ is hardcoded in $D_{\neq \alpha_1 \| m^*}^{\text{com}}$. Thus, the first line of $D_{\neq \alpha_1 \| m^*}^{\text{com}}$, in particular, condition “ $\text{Com}_{\text{ck}}(0; \beta) = \hat{z} = \text{Com}_{\text{ck}}(1; \hat{\beta})$ ” is also never true except negligible probability due to the statistical binding property of Com . That is, these two circuits are functionally equivalent except negligible probability. We obtain the proposition by applying the IO security. ■

Proposition B.30. *If $i\mathcal{O}$ is a secure IO, it holds that $|\Pr[\text{rHyb}_1^4 = 1] - \Pr[\text{rHyb}_1^5 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Proposition B.30. In these games $\beta_1 \leftarrow \{0, 1\}^{9\ell}$ is uniformly random. By the sparsity of F , β_1 is not in the image of F except with negligible probability. Thus, $E_{\neq \alpha_1 \| m^*}$ and $E_{\neq \alpha_1 \| m^*, \neq \beta_1}$ are functionally equivalent except with negligible probability. We obtain the proposition by the IO security. ■

Proposition B.31. *If $i\mathcal{O}$ is a secure IO, it holds that $|\Pr[\text{rHyb}_1^5 = 1] - \Pr[\text{rHyb}_1^6 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Proposition B.31. The difference between $D_{\neq \alpha_1 \| m^*}^F$ in Figure 23 and $D_{\neq \alpha_1 \| m^*, \neq \beta_1}^4$ in Figure 19 is that we replace “If $c = x_1$, outputs \perp .” with “If $\beta = \beta_1$, outputs \perp .” since the first line of $D_{\neq \alpha_1 \| m^*}^F$ is never triggered. In these games, $\beta_1 \leftarrow \{0, 1\}^{9\ell}$ is not in the image of F except with negligible probability. Recall that $c = x_1$ means $c = \alpha_1 \| \beta_1 \| \gamma_1$. Thus, those two circuits may differ when $\beta = \beta_1$ but $(\alpha, \gamma) \neq (\alpha_1, \gamma_1)$. However, it does not happen $\beta = F'(\alpha \| (G(\beta) \oplus \gamma))$ in this case due to the injectivity of F . Thus, $D_{\neq \alpha_1 \| m^*}^F$ and $D_{\neq \alpha_1 \| m^*, \neq \beta_1}^4$ are functionally equivalent and we obtain the proposition by applying the IO security. ■

Proposition B.32. *If $i\mathcal{O}$ is a secure IO, it holds that $|\Pr[\text{rHyb}_1^6 = 1] - \Pr[\text{Rand}_2 = 1]| \leq \text{negl}(\lambda)$.*

Proof of Proposition B.32. The difference between these two games that we use $D_{\neq \alpha_1 \| m^*, \neq \beta_1}^4[F_{\neq \alpha_1 \| m^*}, G_{\neq \beta_1}]$ instead of $D_{\neq \alpha_1 \| m^*, \neq \beta_1}^4[F_{\neq \alpha_1 \| m^*}, G]$. However, $G_{\neq \beta_1}(\beta_1)$ is never computed by the first line of $D_{\neq \alpha_1 \| m^*, \neq \beta_1}^4$. We obtain the proposition by the IO security. ■

We complete the proof of Lemma B.25. ■

B.4 Original Ciphertext Pseudorandomness of PE

We describe the original ciphertext pseudorandomness of PE defined by Cohen et al. [CHN⁺18] in this section for reference.

Definition B.33 (Ciphertext Pseudorandomness). *We define the following experiment $\text{Expt}_{\mathcal{A}}^{\text{cpr}}(\lambda)$ for PE.*

1. \mathcal{A} sends a message $m^* \in \{0, 1\}^{\ell_p}$ to the challenger.
2. The challenger does the following:
 - Generate $(\text{ek}, \text{dk}) \leftarrow \text{Gen}(1^\lambda)$
 - Compute encryption $c^* \leftarrow \text{Enc}(\text{ek}, m^*)$.
 - Choose $r^* \leftarrow \{0, 1\}^{\ell_{\text{ct}}}$.
 - Generate the punctured key $\text{dk}_{\notin \{c^*, r^*\}} \leftarrow \text{Puncture}(\text{dk}, \{c^*, r^*\})$
 - Choose $\text{coin} \leftarrow \{0, 1\}$ and sends the following to \mathcal{A} :

$$(c^*, r^*, \text{ek}, \text{dk}_{\notin \{c^*, r^*\}}) \text{ if } \text{coin} = 0$$

$$(r^*, c^*, \text{ek}, \text{dk}_{\notin \{c^*, r^*\}}) \text{ if } \text{coin} = 1$$

3. \mathcal{A} outputs coin^* and the experiment outputs 1 if $\text{coin} = \text{coin}^*$; otherwise 0.

We say that PE has ciphertext pseudorandomness if for every QPT adversary \mathcal{A} , it holds that

$$\text{Adv}_{\mathcal{A}}^{\text{cpr}}(\lambda) := 2 \cdot \Pr[\text{Expt}_{\mathcal{A}}^{\text{cpr}}(\lambda) = 1] - 1 \leq \text{negl}(\lambda).$$

Issue in the proof by Cohen et al. In the watermarking PRF by Cohen et al. [CHN⁺18], we use $x_0 \leftarrow \text{PE.Enc}(\text{pe.ek}, a\|b\|c\|i)$ to extract an embedded message. They replace $x_0 \leftarrow \text{PE.Enc}(\text{pe.ek}, a\|b\|c\|i)$ with $x_1 \leftarrow \{0, 1\}^{\ell_{\text{ct}}}$ in their proof of unremovability [CHN⁺18, Lemma 6.7]. Then, they use PRG security [CHN⁺18, Lemma 6.8] to replace $\text{PRG}(c)$ with a uniformly random string since the information about c disappears from the PE ciphertext. However, there is a subtle issue here. The information about c remains in the punctured decryption key $\text{dk}_{\neq\{x_0, x_1\}} \leftarrow \text{Puncture}(\text{pe.dk}, \{x_0, x_1\})$, which is punctured both at x_0 and x_1 , since they use ciphertext pseudorandomness in Definition B.33 and need to use the punctured decryption key. Thus, we cannot apply PRG security even after we apply the ciphertext pseudorandomness in Definition B.33. This is the reason why we introduce the strong ciphertext pseudorandomness in Definition 7.3.