

Composability in Watermarking Schemes

Jiahui Liu¹ and Mark Zhandry²

¹Massachusetts Institute of Technology, jiahuiliu.crypto@gmail.com

²NTT Research, mark.zhandry@ntt-research.com

Abstract

Software watermarking allows for embedding a mark into a piece of code, such that any attempt to remove the mark will render the code useless. Provably secure watermarking schemes currently seems limited to programs computing various cryptographic operations, such as evaluating pseudorandom functions (PRFs), signing messages, or decrypting ciphertexts (the latter often going by the name “traitor tracing”). Moreover, each of these watermarking schemes has an ad-hoc construction of its own.

We observe, however, that many cryptographic objects are used as building blocks in larger protocols. We ask: just as we can compose building blocks to obtain larger protocols, can we compose watermarking schemes for the building blocks to obtain watermarking schemes for the larger protocols? We give an affirmative answer to this question, by precisely formulating a set of requirements that allow for composing watermarking schemes. We use our formulation to derive a number of applications.

1 Introduction

Watermarking is an old idea, which aims to embed a mark in some object, such that any attempt to remove the mark destroys the object. In software watermarking, this means embedding a mark into program code, such that any attempt to remove the code will make the code useless. Such watermarking aims to deter piracy by identifying the source of pirated software. Recently, software watermarking has become an active area of research within cryptography, with numerous positive results for watermarking cryptographic functionalities, such as trapdoor functions [Nis13], pseudorandom functions [CHN⁺16, KW17, GKWW21], decryption [CFNP00] (under the name “traitor tracing”), and more [GKM⁺19].

In this work, we initiate the study of *composing* watermarked functionalities. That is, if a cryptographic primitive A (or perhaps several primitives) is used to build a primitive B , can we use a watermarking scheme for A to realize a watermarking scheme for B ? Our aim is to show when such a composition is possible, based on properties of the construction and security proof for B using A .

1.1 Motivation

Abstractions are central to cryptography, as they allow for decomposing various tasks into smaller building blocks, which can then be instantiated independently. The literature is full of results

that show how to generically realize one abstraction assuming solutions to one or more input abstractions.

Given that cryptographic primitives are often composed, and that most watermarking schemes with provable security are for cryptographic primitives, an interesting question is whether watermarking schemes can be composed. To the best of our knowledge, this question has not been previously asked. In contrast, existing watermarking schemes are each developed “from scratch.” Even if the underlying techniques are similar, watermarking schemes for different primitives must go through separate constructions and security proofs. This can be a time-consuming process.

What does it mean to watermark a PRF? As a running example throughout this introduction, we will use pseudorandom functions (PRFs), which are one of the main workhorses in symmetric key cryptography, and are used as a central component in many higher-level protocols. In [CHN⁺16], the authors show how to watermark the evaluation procedure for a certain class of PRFs. We point out, however, that PRFs are typically not considered a cryptographic end goal, but rather a tool used to build other cryptographic notions. So *what, then, is the utility of watermarking a PRF?*

Another fundamental question is the following: for important reasons that we will not get into here, the watermarking guarantee proved by [CHN⁺16] (and all subsequent work on watermarking PRFs) is weaker than one may expect. Namely, they show that it is impossible to remove the mark without causing the program to fail on *random* inputs. Certain PRFs called “weak PRFs” are only guaranteed secure when the adversary sees evaluations on random inputs, and so for this reason may authors (e.g. [GKWW21, MW22, KN22]) refer to such a scheme as watermarking “weak PRFs.” Weak PRFs can be used as building blocks for many applications, though not as many as ordinary PRFs. In the context of using PRFs as a building block, what is implication of watermarking a *weak* PRF?

Composition of Watermarking Schemes. Our thesis is that watermarking, at least in many cases, should be defined and executed in such a way as to be composable, allowing watermarking schemes for building blocks to generically imply watermarking schemes for higher-level protocols. Watermarking a weak PRF, for example, should generically enable watermarking for many applications of weak PRFs, such as CPA-secure symmetric encryption. Naturally, a more ambitious goal is: watermarking a message-authentication code scheme or a digital signatures scheme, when composed with a watermarkable PRF scheme, should lead to a watermarkable CCA-secure SKE scheme.

Not all compositions support watermarking. Certainly, not all composition results from cryptography can be used to compose watermarking schemes. For example, pseudorandom functions can be built from any pseudorandom generator (PRG) [GGM86]. But this does not seem to yield a viable path toward constructing a watermarking scheme for PRFs. After all, what would it even mean to watermarking a PRG, given that the evaluation algorithm for the PRG is public?

1.2 Overview of Our Results

Defining Watermarking Security. Most cryptographic primitives are defined by an interactive game between an adversary A and challenger. Given a security game, we can translate the game

into a watermarking definition, as follows. The attacker gets a watermarked secret key, and then tries to produce a program A . We say that A is “good” if it can win the security game with non-negligible advantage. We then require the existence of a tracing algorithm, which can extract the mark from any good program A produced by the adversary. In the case of watermarking public key decryption functionalities, our notion corresponds exactly to existing notions of traitor tracing, since the security game is non-interactive. However, for other primitives, our definition is potentially stronger than existing notions: existing notions only ask for mark extraction for non-interactive programs A that can evaluate some function, whereas we must extract from any A which wins a security experiment, which is potentially interactive. The strengthened definitions will be crucial for our composition theorem, which we now describe.

Composition Theorem. Our first result is a composition theorem, which gives conditions under which a construction of a target primitive P from input primitives P_1, \dots, P_k can be turned into a compiler for watermarking schemes.

Theorem 1.1 (Main Theorem (Informal)). *If the construction of a target primitive P from input primitives P_1, \dots, P_k satisfies some given conditions and the watermarking schemes for P_1, \dots, P_k satisfy the above watermarking security definition, then we can compose the construction to watermarking schemes for P_1, \dots, P_k , black-boxly into a watermarking scheme for P that satisfies the above watermarking definition.*

Our conditions apply to a large class of constructions. They are, very roughly, as follows:

- The construction of P from P_1, \dots, P_k is black-box. Moreover, the secret key sk for P is $sk = (sk_1, \dots, sk_k)$ where sk_i is the secret key for P_i .
- The security proof for P turns an adversary A for P into adversaries A_1, \dots, A_k for P_1, \dots, P_k , respectively, such that if A has non-negligible advantage, so does *at least* one of the P_i . Typical proofs in cryptography utilizing hybrid proofs will usually have this form.
- Moreover, we require a property of the security proof, which we call a “watermarking-compatible reduction”. This is a rather technical definition, but roughly we allow the security games for P, P_1, \dots, P_k to consist of two phases, where the adversary’s winning condition is only dependent on the second phase. We require that the reduction respects the phases, in the sense that the second stages of A_1, \dots, A_k only depend on the second stage of A .
- Depending on the construction and reduction, *not all* input primitives P_1, \dots, P_k have to be watermarkable to give a watermarking construction for the target primitive P . That is, we can compose the “plain” constructions of these input primitives with the watermarkable versions of the other input primitives to give a watermarkable construction for P .
- If all input primitives (that need to be watermarkable) have collusion-resistant watermarking security, the resulting watermarkable P also satisfies collusion-resistant security. If all input primitives can be traced using only a public key, the same is true of P .

Many reductions in the literature are watermarking-compatible. For example, consider constructing CPA-secure symmetric encryption from weak PRFs. Here, the first stage for CPA-security consists of all queries occurring prior to the challenge query, and the second stage consists of the challenge query and all subsequent queries. The winning condition does not depend on any first stage CPA queries (or second stage). The first and second stages for the weak PRF are just two

rounds of queries to the weak PRF oracle, and here again the win condition does not depend on the actually queries. Thus, this reduction is watermarking compatible.

A non-example would be many proofs involving signatures as the target primitive P . The issue is that, the winning condition for a signature scheme security game is that the adversary produces a “new” signature on a message that was not seen in a previous query. But checking this win condition requires knowing all the queries. So if there is a first phase where the attacker can make signature queries, then the win condition is not solely dependent on the second stage. If we let the winning condition depend only on the second stage, the adversary can trivially win by querying a signature in the first stage and give it to the second stage adversary as the final output.

Applications. We demonstrate that many well-known cryptographic constructions have watermarking-compatible reductions, and therefore we can compose the watermarkable constructions of the input primitives to obtain watermarkable constructions of the target primitive. Within the scope of this work, we give the following examples:

- Two most simple examples are:
 - Watermarkable CPA-secure secret-key encryption scheme from watermarkable weak PRF
 - Watermarkable CCA2-secure secret-key encryption scheme from watermarkable weak PRF and watermarkable MACs.
- Some more advanced examples are:
 - Watermarkable CCA2-secure public-key encryption from watermarkable selectively secure identity based encryption and strong one-time signatures, which can in turn be based on LWE. Here, the strong one-time signature we need is a “plain scheme” which does not need to be watermarkable.
 - Watermarkable CCA2-secure PKE from watermarkable CPA-secure PKE and NIZK (without watermarking), which can in turn be based on LWE too.
 - Watermarkable weak pseudorandom permutation from watermarkable weak PRF.
 - Watermarkable CCA2-secure hybrid encryption scheme from watermarkable CCA2-secure PKE and CCA2-secure SKE (without watermarking).

Additionally, we show that all the input primitives (weak PRF, CPA-secure PKE, selective IBE, signatures, etc.) in the above examples have constructions that satisfy our security definition for watermarking. We can obtain these constructions by using or modifying the existing watermarking schemes in [GKM⁺19], [GKWW21, MW22]. Therefore, the above composed schemes all have concrete constructions.

We also briefly discuss how the functional encryption construction in [GKP⁺13] is also watermarking compatible. We cannot possibly elaborate all the concrete examples of watermarking compatible reductions within the scope of this work, but given our generic framework of composition, one can easily verify whether a construction is watermarking compatible by looking into its security proof.

1.3 Other Related Work

[Nis20] presents a general framework for constructing watermarking schemes for any primitive which admits a certain “all-but-one” reduction. The work and ours focus on different aspects of

watermarking: theirs is focused on constructing watermarking schemes in the first place, whereas our composition theorem shows how to combine watermarking schemes. We also note that the framework in [Nis20] seems very much tied to the collusion-free setting, whereas ours is much more general, and can accommodate collusions if the input tracing mechanisms are collusion resistant.

1.4 Technical Overview

Watermarking We first briefly recall the definition of watermarking a cryptographic primitive: a watermarkable cryptographic scheme, apart from the usual evaluation algorithms (such as key generation, encrypt, decrypt or sign and verify), additionally has a Mark algorithm and an Extract algorithm (as well as the corresponding marking and extraction keys). The Mark algorithm allows one to embed a mark into the secret key used to evaluate the cryptographic functionality; the Extract key allows one to extract a mark from an allegedly marked key. The watermarking security, usually referred to as "unremovability", states that given a marked secret key with an adversarially requested mark τ , the adversary should not be able to produce a circuit that has the same functionality as the secret key such that the above mark τ cannot be extracted from this adversarial circuit.

As explored in this work and some previous works ([GKM⁺19, GKWW21]), one important definitional aspect in the above security lies in what we mean by the "adversary's output circuit has the same functionality as the original secret key". We will elaborate in the following sections of the overview.

Watermarking-Compatible Reduction Before we go into how we compose watermarking schemes, we expand more on the type of reductions that allow us to build a watermarking composition upon. We call these reductions watermarking-compatible. We will then give a concrete example to help comprehension.

Consider a black-box construction of a target primitive P from input primitives P_1, \dots, P_k and consider a reduction algorithm \mathcal{B} from security of P to the security of P_i : we let both the adversary \mathcal{A} and the reduction \mathcal{B} be divided into two stages. In the first stage, stage-1 adversary \mathcal{A}_1 , for the security game of P , receives some public parameters from stage-1 \mathcal{B}_1 and makes some queries; \mathcal{B}_1 answers these queries by making queries to the oracle provided by the challenger of the security game for P_i .

Entering the second stage, as one can expect, \mathcal{A}_1 can give an arbitrary state it to stage-2 adversary \mathcal{A}_2 . However, what \mathcal{B}_1 can give to the second stage reduction \mathcal{B}_2 is more restricted: \mathcal{B}_1 can give all the public parameters to the second stage but *none of the queries* made by \mathcal{A}_1 , to \mathcal{B}_2 . \mathcal{B}_2 continues to simulate the query stage to answer \mathcal{A} 's queries. In particular, \mathcal{B} records \mathcal{A}_2 's queries. Note that the challenge phase of the security game where \mathcal{A}_2 (and resp. \mathcal{B}) receives its challenge from the challenger always happens in stage 2¹.

In the final output phase, \mathcal{B}_2 's answer to the challenger in game P_i will be dependent on (some of) the following pieces of information: challenger's challenge input to \mathcal{B}_2 , \mathcal{A}_2 's queries made during stage 2, \mathcal{B}_2 's randomness used to prepare the challenge input for \mathcal{A}_2 and \mathcal{A}_2 's final output.

¹But as we will see in some concrete examples, \mathcal{A} can commit to some "challenge messages" in stage 1, which will be taken into stage-2 as an "auxiliary input" and later used by the challenger to prepare the challenge.

We give some intuition on why we need B to be "oblivious" about \mathcal{A} 's queries in stage 1 of the above reduction. Looking forward, in the actual watermarking (unremovability) security game, we can replace "answering \mathcal{A} 's queries in stage 1" with a giving out a watermarked secret key to \mathcal{A} , where B will have no idea what inputs \mathcal{A} has evaluated on using the key. We therefore model the reduction as the above to capture this scenario.

Example: CCA2-Secure Secret-Key Encryption To make the above abstract description concrete, we take an example of the reduction from CCA2-security to weak PRF and MAC.

We briefly recall the textbook construction: the scheme's secret key consists of the PRF's secret key sk_1 and MAC's secret key sk_2 . The encryption algorithm, on input message m , computes ciphertext $ct = (r, ct' = \text{PRF.Eval}(sk_1, r) \oplus m, \text{sig})$, where $\text{sig} \leftarrow \text{MAC.Sign}(sk_2, (r, ct'))$. The decryption algorithm will first verify the signature sig on (r, ct') , if yes continue to decrypt using sk_1 , else abort.

It's not hard to see how the CCA2 security game can fit into the format of a two-stage game we have depicted above: stage 1 consists simply of letting the adversary making queries to the encryption and decryption oracles. Entering stage 2, the adversary \mathcal{A} is allowed to make more queries and then submits the challenge plaintexts; then \mathcal{A} receives challenge ciphertexts from the challenger. \mathcal{A} continues to make more admissible queries and finally outputs its answer.

By viewing the security game in two stages, we will recall the security proof for CCA2 security. The proof can go through a case-by-case analysis:

1. In case 1, in stage 2 of the CCA2 security game, there is a decryption query from \mathcal{A} of the form $ct = (ct_0, \sigma)$, such that it has never been the output of an encryption query. Also, the decryption oracle did not output \perp on this query.
2. In case 2, there is no such decryption query in stage 2.

In the first case, we can do a reduction to unforgeability of MAC:

- \mathcal{A}_{MAC} samples its own PRF secret key sk_1 . For stage 1: For any encryption query from stage-1 adversary, \mathcal{A}^1 , stage-1 reduction $\mathcal{A}_{\text{MAC}}^1$ will simulate the response as follows: it will compute the ct_0 part of ciphertext and then query the signing oracle $\text{MAC.Sign}(sk_2, \cdot)$ in the security game G_{MAC} . For any decryption query, $\mathcal{A}_{\text{MAC}}^1$ will query the verification oracle $\text{MAC.Verify}(sk_2, \cdot)$ first and decrypt those whose response is not \perp .
- After entering stage 2, for any encryption or decryption query from \mathcal{A}^2 , $\mathcal{A}_{\text{MAC}}^2$ will still simulate the response as the above. Recall that $\mathcal{A}_{\text{MAC}}^2$ will not get to see any queries made in stage 1 and $\mathcal{A}_{\text{MAC}}^2$ will record all queries from \mathcal{A}_2 .
- In the challenge phase, \mathcal{A}^2 sends in challenge messages (m_0, m_1) ; $\mathcal{A}_{\text{MAC}}^2$ flips a coin $b \leftarrow \{0, 1\}$, prepares and sends the signed encryption ct^* of m_b to \mathcal{A}^2 . $\mathcal{A}_{\text{MAC}}^2$ continues to simulate the encryption and decryption oracles for \mathcal{A}^2 , on queries $ct \neq ct^*$.
- In the end, $\mathcal{A}_{\text{MAC}}^2$ will look up \mathcal{A}^2 's queries: find a decryption query $ct = (ct', \sigma)$ such that it has never been the output of an encryption query and the decryption oracle did not output \perp on this query. $\mathcal{A}_{\text{MAC}}^2$ output (ct', σ) as its forgery.

In the second case, we consider a reduction \mathcal{A}_{PRF} to the weak pseudorandomness of the PRF. In the following reduction, for the sake of simplicity, let's consider a variant of the weak PRF game: the adversary \mathcal{A}_{PRF} is given adaptive query access to an oracle $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ that computes a PRF function; there will be a challenge phase where a challenge input r^* is sampled at

random; \mathcal{A}_{PRF} will receive one of $(r^*, F(r^*))$ or $(r^*, y \leftarrow \{0, 1\}^m)$ at random and try to tell which one it receives ².

- In stage 1: \mathcal{A}_{PRF} samples its own MAC key sk_2 . For any encryption query from \mathcal{A}^1 , $\mathcal{A}_{\text{PRF}}^1$ will simulate the response as follows: it will query the PRF oracle $F(\cdot) := \text{PRF.Eval}(sk_1, \cdot)$ on a fresh r of its own choice. After obtaining $(r, F(r))$, \mathcal{A}_{PRF} signs the message $(r, F(r) \oplus m)$ using sk_2 and sends the entire ciphertext to \mathcal{A} . Similarly, for decryption queries, \mathcal{A}_{PRF} first verifies the signature and then queries $F(\cdot)$ oracle to decrypt.
- After entering stage 2, for any encryption query from \mathcal{A}^2 , $\mathcal{A}_{\text{PRF}}^2$ will still simulate the response as the above. Recall that $\mathcal{A}_{\text{PRF}}^2$ will not get to see any queries made in stage 1.
- In the challenge phase, $\mathcal{A}_{\text{PRF}}^2$ will receive a challenge input-output pair (r^*, y^*) from the weak PRF challenger: $r^* \leftarrow \{0, 1\}^\ell, y^* \in \{0, 1\}^\ell$ where y^* is either $\text{PRF.Eval}(sk_1, r^*)$ or uniformly random.
- \mathcal{A}^2 sends in challenge messages, (m_0, m_1) ; $\mathcal{A}_{\text{PRF}}^2$ flips a coin $b \leftarrow \{0, 1\}$ and sends $ct = (r^*, y^* \oplus m_b, \sigma^*)$ to \mathcal{A}^2 .
- If \mathcal{A}^2 guesses the correct b , then $\mathcal{A}_{\text{PRF}}^2$ output 0, for " y^* is $\text{PRF.Eval}(sk_1, r^*)$ ", else $\mathcal{A}_{\text{PRF}}^2$ output 1, for " y^* is uniformly random".

Note that in the above reductions, both \mathcal{A}_{PRF} and \mathcal{A}_{MAC} do not need the stage-1 queries from \mathcal{A} to help them win their own games.

Definition and Composition Framework of Watermarking Now we relate the above reduction to the watermarking goal: in our watermarking (i.e. unremovability) security game for primitive P , the watermarking adversary \mathcal{A} acts like a "stage-1" adversary in the original security game for the underlying primitive P . Then \mathcal{A} produces a program C —this signifies the end of "stage 1" of the security game.

A notable feature of our extraction algorithm is the following. The adversarial program C produced by \mathcal{A} will be treated as a stage-2 adversary by the extraction algorithm: the extraction algorithm will try to extract a watermark by having only black-box access to C . Since C is supposed to function like a stage-2 adversary that wins the security game of P , the extraction procedure ensures C operates properly by simulating the stage-2 security game for C .

Finally, \mathcal{A} wins if C is a "good" program in winning the corresponding security game for P , but no valid watermarks can be extracted from C .

Suppose a target primitive P can be built from input primitives P_1, \dots, P_k via watermarking-compatible reductions, and each watermarking scheme for P_i satisfies our definition, then we can compose the watermarkable versions of P_1, \dots, P_k to give a watermarkable P .

The composition construction works roughly as follows:

1. A watermarkable P 's key generation algorithm is similar to the unwatermarkable (plain) construction of P : by generating all keys of P_1, \dots, P_k and concatenate them, for secret/public keys respectively. The extra step is that now we also concatenate the marking keys and extraction keys generated from the watermarkable P_1, \dots, P_k .
2. How P evaluates is exactly the same as in the plain construction of P from P_1, \dots, P_k : by running the P_1, \dots, P_k algorithms as black-box subroutines.

²Watermarkable weak PRF with this type of security can be constructed in [GKM⁺19].

3. To watermark P , we simply watermark the secret keys of the P_1, \dots, P_k respectively. The watermarked key for P is then just the concatenation of the watermarked keys for P_1, \dots, P_k . Since the construction is black-box and assumes the key is just a tuple of keys for P_1, \dots, P_k , the evaluation with watermarked keys is still the same, running the evaluation algorithm with black-box from subroutines P_1, \dots, P_k , except now with the marked keys respectively. We obtain correctness (functionality-preserving property), following from the correctness of the plain construction and functionality preserving properties of watermarkable P_1, \dots, P_k .
4. In order to extract a mark, we turn any pirated algorithm A for P into pirated algorithm A_1, \dots, A_k for P_1, \dots, P_k . We do this by applying the security reduction for P to the pirated algorithm A , and let A_1, \dots, A_k be the adversaries produced by the reduction. We then interpret A_1, \dots, A_k as pirated programs for P_1, \dots, P_k , and attempt to extract the mark from each of them. Note that any mark extraction algorithm is supposed to use only the input-output behavior of the pirate program.

The guarantee of the reduction is that one of these pirated programs A_i must actually be a "good" program in the security game for the corresponding primitive P_i , which means that mark extraction must succeed for that A_i . The security proof therefore shows that we are guaranteed to extract some mark ³.

Formalizing this composition takes some care, as we need to ensure that the mark extraction algorithm has the ability to actually transform A into each of the A_i . There are a couple issues with getting this to work:

- In a reduction, it is typically assumed that, when the reduction is attacking P_i , it has access to the keys for $P_j, j \neq i$. This access is often used to simulate the view of A when constructing A_i . In our watermarking construction, we therefore need to ensure that the tracing algorithm has this information.
- Security proofs often work via hybrid arguments that change various terms. Importantly, the hybrid arguments get to simulate the *entire* game. For us, however, we only get to simulate the second stage of the game; the first stage has actually already been fixed by the time the adversary produces its pirated program A . So we have to ensure that the tracing algorithm can carry out the reduction to create a program A_i , even though it cannot simulate the entire security experiment from the very beginning.

Our definition and proof take care of these issues and more, to give a composition theorem that applies any time our criteria are met.

Note that the restriction to games where the winning condition is independent of the first stage seems necessary. This is because the adversary actually gets in its possession a watermarked key, which lets it compute the various cryptographic functions of the key. For example, in the case of digital signatures, the adversary can actually compute signatures for itself. In the case of encryption, the adversary can encrypt and decrypt messages. Moreover, there is no way to track what the adversary is doing, since it is all happening internally to the adversary rather than being explicitly queried (suppose for example, that the adversary signs a few messages, but has its entire state encrypted under a fully homomorphic encryption (FHE) scheme so that the

³The Extract algorithm assumes the (input) adversarial circuit to be possibly stateful and interactive, but does not require any input circuit to be stateful and interactive. For example, an honestly generated watermarked circuit is not stateful or interactive in our construction.

signature and message being signed are all “under the hood” of the FHE scheme). We model this ability by having a first stage of the security experiment where the adversary can freely query the functionalities, but then do not have the win condition depend on those queries.

Watermarking Composition Example: CCA2-Secure SKE Now we demonstrate how the above abstract watermarking composition works by using the concrete example of a CCA2-secure SKE built from weak PRF and MAC again.

Consider having input constructions as watermarkable weak PRF and watermarkable MAC which satisfies our syntax and unremovability security requirements, a watermarking CCA-secure scheme consists of several algorithms: WMS_{Setup}, Enc, Dec, Mark, Extract. The WMS_{Setup} generates the secret key, marking key and extraction key by concatenating the corresponding keys from the PRF and MAC. The marking algorithm is on input $sk = (sk_{\text{PRF}}, sk_{\text{MAC}})$ and watermarking message τ , output $\tilde{sk} = (\tilde{sk}_{\text{PRF}} \leftarrow \text{PRF.Mark}(sk_{\text{PRF}}, \tau), \tilde{sk}_{\text{MAC}} \leftarrow \text{PRF.Mark}(sk_{\text{MAC}}, \tau))$.

Finally, the Extract algorithm is slightly more complex: on input extraction key $xk = (xk_{\text{MAC}}, xk_{\text{PRF}})$ and a circuit C , do *both* of the following:

1. Create a circuit C_{PRF} :

- Simulate the stage-2 CCA2-security for C (as described in our previous paragraph on watermarking-compatible reduction from CCA2 security to PRF). Note that just as a real reduction C_{PRF} is only hardcoded with the MAC extraction key xk_{MAC} (unless the underlying watermarkable PRF scheme has a public extraction key), which enables it to simulate the signing and verification of MAC; to compute the part that involves the PRF evaluation oracle, C_{PRF} need to make external queries to some "challenger" to get answers.

How C_{PRF} uses C to get its final output is exactly the same as in the watermarking-compatible reduction from CC SKE to weak PRF we described.

After C_{PRF} is created, the extraction algorithm uses the watermarkable PRF’s extraction algorithm to extract a mark from C_{PRF} : $\tau/\perp \leftarrow \text{PRF.Extract}(xk_{\text{PRF}}, C_{\text{PRF}})$. Note that the "external queries" made by the circuit will be answered now by the $\text{PRF.Extract}(xk_{\text{PRF}}, \cdot)$ algorithm, because it treats C_{PRF} as a stage-2 adversary in the weak pseudorandomness security game and uses the extraction key xk_{PRF} to simulate the game.

2. Create a circuit C_{MAC} :

- Simulate the stage-2 CCA2-security for C (as described in our previous paragraph on watermarking-compatible reduction from CCA2 security to MAC). Note that just as a real reduction C_{MAC} is only hard-coded with the PRF extraction key xk_{PRF} , which enables it to simulate the oracles of PRF; to compute the part that involves the MAC signing/verifying oracles, C_{MAC} need to make external queries to some "challenger" to get answers.

How C_{MAC} uses C to get its final output is exactly the same as in the watermarking-compatible reduction from CCA SKE to MAC we described.

After C_{MAC} is created, the extraction algorithm will use the watermarkable MAC’s extraction algorithm to extract a mark from C_{MAC} : $\tau/\perp \leftarrow \text{MAC.Extract}(xk_{\text{MAC}}, C_{\text{PRF}})$. Note that the "external queries" made by the circuit will be answered by the algorithm $\text{MAC.Extract}(xk_{\text{MAC}}, \cdot)$ algorithm, because it treats C_{MAC} as a stage-2 adversary in the MAC security game and uses the extraction key xk_{MAC} to simulate the game.

We say that the adversary \mathcal{A} wins the watermarkable CCA2-secure SKE's unremovability game if no (previously queried) watermark can be extracted from either of the above circuits $C_{\text{PRF}}, C_{\text{MAC}}$ created during the extraction algorithm and C is a "good" program in winning the CCA2 security game. The intuition for the security proof is relatively straightforward: if C is a "good" program for winning CCA2-security game, then at least one of C_{PRF} and C_{MAC} must be a "good" program for winning its corresponding security game of weak PRF or MAC, by the property of the reduction.

For instance, imagine a reduction algorithm \mathcal{B}_{PRF} to the unremovability of watermarkable PRF: \mathcal{B}_{PRF} will simulate the marking query stage for \mathcal{A} by making marking queries to the marking oracle of the weak PRF challenger, along with the MAC key it sampled on its own. Note that the evaluation queries to the PRF evaluation oracle can now be replaced by having access to a marked PRF key. After \mathcal{A} outputs circuit C , \mathcal{B}_{PRF} simply creates the same circuit C'_{PRF} as the C_{PRF} created in the above extraction algorithm. If the C_{PRF} circuit created in the extraction is "good" at winning the weak pseudorandomness game, then so is C'_{PRF} . Since we cannot extract a watermark from C_{PRF} , neither can we extract from C'_{PRF} because they have the same input-output behavior when using the same C as its subroutine black-boxly. The same argument applies to the reduction to unremovability of watermarkable MAC.

More Advanced Watermarking Constructions with "Unwatermarkable" Building Blocks The watermarking reduction in the above example of CCA2-secure SKE is relatively straightforward to acquire. Some similar construction examples include a watermarkable weak PRP from watermarkable weak PRF (see Appendix B).

However, there exist various constructions that look markedly distinct from the above example of CCA2-secure SKE and may be much more involved.

More importantly, some constructions have building blocks which are cryptographic primitives that don't make sense to watermark, for instance, a non-interactive zero knowledge proof scheme. Nevertheless, we show that a large class of them can be watermarking-compatible and in many scenarios, we simply do *not* need to watermark all the building blocks.

We briefly discuss two examples on a high level:

- CCA-secure PKE from IBE and One-Time Signatures** The first example is the CCA2-secure PKE from selectively secure IBE and one-time signatures by [BCHK07]. In the original scheme, the encryption samples a fresh signature signing/verification key pair (sk, vk) , compute an IBE ciphertext by using the vk as the identity, then signs this ciphertext using signing key sk , and finally outputs the IBE ciphertext, the signature and vk . To decrypt, one first verifies the signature under vk , if valid then derives an identity-embedded decryption key from the IBE master secret key using vk as the identity and decrypt the ciphertext. Note that in this scheme, the keys for one-time signatures are only sampled "online" upon every encryption algorithm and how to watermark such keys is not well-defined. However, as it turns out—we don't have to watermark the signing key in our composed construction. A watermarkable CCA2-secure PKE using the above construction only has to watermark the IBE secret key. The security proof would show that a successful unremovability adversary can help us either break the unremovability security of the watermark IBE scheme or break the strong one-time unforgeability security of the signature scheme. In other words,

leveraging the "plain" security of the one-time signature scheme suffices for the composed watermarking scheme.

In more detail, we can show the following: the pirated algorithm C produced by an adversary is a "good" program for breaking the CCA-security and meanwhile no watermark can be extracted from C . However, it does not help us win the selective IBE CPA-security game. Then, simply following the hybrid argument in the security analysis of the CCA-secure PKE itself, we observe that C must be breaking the security of the one-time signature scheme. Thereby, we can use C black-boxly to create a reduction for the one-time signature scheme.

- **CCA-secure PKE from CPA-secure PKE and NIZK** A second example is the [NY90] CCA-secure PKE scheme built from CPA-secure PKE and a non-interactive zero-knowledge proof scheme (NIZK). The encryption algorithm encrypts a message twice under two different public keys and uses a NIZK proof to prove that they encrypt the same message. The watermarkable version of the above construction does not watermark the NIZK scheme, but only the two decryption keys of the PKE scheme. How the NIZK scheme is used in the watermarking scheme is less obvious to see: intuitively, the extraction algorithm will try to create two circuits that break the underlying CPA security game of the PKE, with the corresponding keys. If we look carefully into the proof for the [NY90] scheme, to make this reduction go through, one has to first work in a hybrid game where the real NIZK proof in the challenge ciphertext is replaced with a simulated proof. Our extraction algorithm thereby uses such a simulated proof when interacting with the input circuit C . The security of NIZK helps us say that if C is "good" in the original CCA-security game, so will C be good in the game simulated by the extraction algorithm.

To summarize, we only have to watermark the secret keys of input primitives P_i which have their keys generated in the main Key Generation algorithm of the target primitive P and have their secret keys used during the evaluation algorithm of P . If an input primitive has its keys sampled "freshly" during every invocation of P 's evaluation algorithms or sampled in key generation, but not used in P 's evaluation algorithm (e.g. only used in the security proof instead), then we don't need a watermarkable version of P_i to build a watermarkable P . Leveraging P_i 's original security property suffices.

More advanced examples include the functional encryption from attribute-based encryption, FHE and garbled circuits in [GKP⁺13], where we only have to watermark the ABE scheme; a hybrid CCA-secure encryption scheme from CCA-secure PKE and CCA secure SKE, where we only have to watermark the PKE scheme. We refer the readers to the construction and discussions of these examples in Appendix A, 6, B,E, F.

2 Organizations

In Section 3.1 we will first present a generic definition for a definition of a cryptographic primitive, i.e. its algorithms, correctness and security properties. Then we will describe what is a watermarking-compatible construction for a cryptographic primitive black-boxly from some input primitives in

In Section 4, we will first revisit the definition of a cryptographic construction/algorithm with watermarking security. Next, we present and prove the properties of our watermarking composition scheme.

In the rest of the sections, we will elaborate on some examples of cryptographic constructions that are watermarking-compatible and their corresponding watermarking schemes.

3 Definitions: General Cryptographic Primitive and Watermarking-Compatible Constructions

3.1 General Cryptographic Primitive

In this section, we present syntax and definitions for a general cryptographic primitive. The notations and definitions formalized in this section will assist the demonstration of our generic watermarking framework.

General Cryptographic Primitive Syntax A cryptographic primitive $P = (\text{KeyGen}, \text{SecEval}, \text{PubEval})$ consists of the following algorithms:

- $\text{KeyGen}(1^\lambda) \rightarrow (\text{sk}, \text{pk})$: is a (randomized) algorithm that takes a security parameter λ and interacts with an adversary \mathcal{A} : $(\text{sk}, \text{pk}) \leftarrow (\mathcal{A} \Leftrightarrow \text{KeyGen}(1^\lambda))$ where sk is some secret information unknown to \mathcal{A} , and \mathcal{A} can get some public information pk from the interaction.
- $\text{SecEval}(\text{sk}, \text{pk}, x \in \mathcal{X}) \rightarrow y \in \mathcal{Y}_s$: a secret-evaluation algorithm that takes in the secret information sk , public information pk and some input x from input space \mathcal{X} , and output a value $y \in \mathcal{Y}_s$.
- $\text{PubEval}(\text{pk}, x \in \mathcal{X}) \rightarrow y \in \mathcal{Y}_p$: a public-evaluation algorithm that takes in the public information pk and some input x from input space \mathcal{X} , and output a value $y \in \mathcal{Y}_p$.

More generally, the KeyGen algorithm can generate several secret keys $\text{sk}_1, \dots, \text{sk}_\ell$ for some polynomial ℓ . There will be ℓ different secret algorithms $\{\text{SecEval}_i(\text{sk}_i, \text{pk}, x \in \mathcal{X}_{s,i})\}_{i \in [\ell]}$ and (some constant) m different public algorithms $\{\text{PubEval}_j(\text{pk}, x \in \mathcal{X}_{p,j})\}_{j \in [m]}$. For example, an identity-based encryption scheme can have two decryption algorithms, one using the master secret key, the other with a secret key embedded with an identity.

However, without loss of generality, we will make two simplifications:

- All algorithms have their input space padded to be the same length as \mathcal{X} ;
- We view all secret algorithms $\{\text{SecEval}_i(\text{sk}_i, \text{pk}, x \in \mathcal{X}_{s,i})\}_{i \in [\ell]}$ as one algorithm $\text{SecEval}(\text{sk}, \text{pk}, \cdot)$ that will take in an index $i \in [\ell]$ to decide which mode to use, similarly for public algorithms. But we will assume such an index specification to be implicit and omit the use of indices in the algorithm to avoid an overflow of letters.

Occasionally, we denote all the algorithms $(\{\text{SecEval}_i(\text{sk}_i, \text{pk}, x \in \mathcal{X}_{s,i})\}_{i \in [\ell]}, \{\text{PubEval}_j(\text{pk}, x \in \mathcal{X}_{p,j})\}_{j \in [m]})$ as a combined functionality hardcoded with the corresponding keys $\text{Eval}(\text{pk}, \text{sk})$. We call it Eval for short.

Correctness for Predicate F_R Before going into the correctness property, we first define a notion important to our generic definition:

Definition 3.1 (Predicate). A predicate $F(C, x, z_1, \dots, z_k, r)$ is a binary outcome function that runs a program C on a some input x to get output y , and outputs 0/1 depending on whether $(x, y, z_1, \dots, z_k, r) \in R$ for some binary relation defined by R . The randomness of input x , program C both depend on randomness r . z_1, \dots, z_k are auxiliary inputs that specify the relation.

A correctness property of a primitive P with respect to predicate F_R says that

Definition 3.2 (Correctness for Predicate F_R). P satisfies correctness if there exists some function $\epsilon = \epsilon(\lambda) \in [0, 1]$ so that for all $\lambda \in \mathbb{N}, x \in \mathcal{X}$ ⁴:

$$\Pr_{r \leftarrow D_r, (\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(1^\lambda)} [F_R(\text{Eval}, \text{pk}, \text{sk}, x, r) = 1] \geq 1 - \epsilon.$$

The randomness used in checking the predicate is sampled from a distribution D_r . We can simply take D_r to be the uniform distribution, which can be mapped to any distribution we need when computing the predicate.

All ϵ within the scope of this work is negligible in λ .

Remark 3.3. To further explain the above correctness property, we consider the following (abstract) example. Given $\text{Eval} = (\{\text{SecEval}_i(\text{sk}_i, \text{pk}, x \in \mathcal{X})\}_{i \in [\ell]}, \{\text{PubEval}_j(\text{pk}, x \in \mathcal{X})\}_{j \in [m]})$, F_R samples input $x \in \mathcal{X}$ using the first part of r ; then it runs algorithm (supposing randomized, using part 2 of string r) $\text{SecEval}_1(\text{sk}_1, \text{pk}, x)$ to give some outcome y_1 ; then runs (supposing deterministic) $\text{PubEval}_1(\text{pk}_1, y_1)$ to give outcome y_2 ; check if $(x, y_1, y_2) \in R$; output 1 if yes, 0 otherwise.

A concrete example is a public key encryption scheme: the predicate is encrypting a message x using r and then decrypting it to check if one can recover the original message.

Game-based Security A security property of the cryptographic primitive P is described by a interactive procedure G_P between a challenger and adversary \mathcal{A} .

G_P will involve KeyGen , SecEval , PubEval as its subroutines. G_P outputs a bit 1 if the game is not aborted and a certain condition has been met at the end of the game; else it outputs 0.

The security of a primitive P with respect to G_P says: For all $\lambda \in \mathbb{N}$, there exists some function $\eta = \eta(\lambda)$ such that for all admissible \mathcal{A} , there exists a function $\text{negl}(\lambda)$:

$$\Pr[G_P(1^\lambda, \mathcal{A}) = 1] \leq \eta + \text{negl}(\lambda)$$

where η is the trivial probability for any admissible \mathcal{A} to make G_P output 1 and the probability is over the randomness used in G_P .

2-Stage Game-based Security To be compatible with the context of watermarking, we will view the game G_P as a 2-stage security game. 2-stage game-based security is a central notion that connects a primitive's "plain security" to its watermarking security.

The KeyGen procedure and a first part of the interactions between the challenger and \mathcal{A} are taken out of the original game and executed as a first stage G_P^1 . Then, G_P^2 denotes the rest of the game and will take in parameters (sk, pk) generated in stage 1 as well as some auxiliary input.

Definition 3.4 (2-Stage Game-based Security). A security property of the cryptographic primitive P is described by a stage-1 (possibly interactive) key generation procedure a challenger and adversary \mathcal{A}_1 followed by a fixed-parameter game G_P^2 between a challenger and adversary \mathcal{A}_2 .

We denote $G_P^1(1^\lambda, \mathcal{A}_1)$ as the first stage adversary \mathcal{A}_1 interacting a challenger which runs the $\text{KeyGen}(1^\lambda)$, SecEval , PubEval algorithm; together they output a key pair (sk, pk) and some auxiliary parameters aux

⁴For a cryptographic primitive P , there can be many different correct properties, each defined with respect to a different predicate.

which will be later used in the game G_P . \mathcal{A}_1 only gets pk , aux but may make arbitrary polynomial number of admissible queries to oracles provided by the challenger during the interaction.

The stage-2 game challenger G_P^2 is parametrized by inputs sk , pk , aux generated during stage 1 and will involve SecEval , PubEval as its subroutines. Stage-2 adversary \mathcal{A}_2 gets an arbitrary polynomial size state st from stage 1 \mathcal{A}_1 . G_P^2 outputs a bit 1 if the game is not aborted and a certain condition has been met at the end of the game; else it outputs 0.

The security of a primitive P with respect to G_P says: there exists some function $\eta = \eta(\lambda)$ such that for all admissible $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ and any non-negligible $\gamma = \gamma(\lambda)$, there exists a function $\text{negl}(\lambda)$ for all $\lambda \in \mathbb{N}$:

$$\Pr[\mathcal{A}_2(\text{st}) \text{ is } \gamma\text{-good in } G_P^2(\text{sk}, \text{pk}, \text{aux}, \cdot) : \{(\text{sk}, \text{pk}), \text{aux}, \text{st}\} \leftarrow G_P^1(\mathcal{A}_1, 1^\lambda)] \leq \text{negl}(\lambda)$$

where \mathcal{A}_2 is said to be γ -good if:

$$\Pr[G_P^2(\text{sk}, \text{pk}, \text{aux}, \mathcal{A}_2) = 1] \geq \eta + \gamma$$

where η is the trivial probability for any admissible \mathcal{A} to make G_P output 1 and the probability is over the randomness used in KeyGen and by G_P .

Remark 3.5. While some readers may find the introduction of parameter γ in the above 2-stage game confounding, we would like to make a note that the above two definitions are essentially equivalent.

An alternative way of stating the 2-stage game security would be: there exists some function $\eta = \eta(\lambda)$ such that for all admissible $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, there exists two negligible functions $\text{negl}_1(\lambda)$, $\text{negl}_2(\lambda)$ such that for all $\lambda \in \mathbb{N}$:

$$\Pr \left[\Pr[G_P^2(\text{sk}, \text{pk}, \text{aux}, \mathcal{A}_2(\text{st})) = 1] \leq \eta + \text{negl}_1(\lambda) : \{(\text{sk}, \text{pk}), \text{aux}, \text{st}\} \leftarrow G_P^1(\mathcal{A}_1, 1^\lambda) \right] \geq 1 - \text{negl}_2(\lambda)$$

Remark 3.6 (Division into a 2-stage Game). How we divide a security game into two stages depends on the application and suppose that the game G_P has different stages by its definition, the way we divide stage-1 and 2 is usually not the same as in the original definition of G_P . We will see examples later.

In most settings, stage-1 game G_P^1 only involves running the key generation KeyGen and letting \mathcal{A}_1 make some queries. In a few special settings, \mathcal{A}_1 needs to commit to some challenge messages which will be put into aux and be part of the input to stage 2. Examples include the challenge attribute/identity in an attribute/identity-based encryption, because the queries made in the first stage need to go through the "admissibility" check that depends on \mathcal{A} 's choice of challenge messages.

Remark 3.7. In the rest of this work, it is usually clear from the context which stage of G_P we are referring to because G_P^2 will take in parameters sk , pk , aux while the entire game G_P takes in only security parameter 1^λ . Plus G_P^1 is seldomly used explicitly in our language. Occasionally we will omit the superscript and slightly abuse the notation to denote G_P^2 as $G_P(\text{sk}, \text{pk}, \text{aux}, \cdot)$

For our convenience in later notations, we also give the following simple definition:

Definition 3.8 (Stage-2 Game View). The stage-2 G_P^2 can also output a view $\text{view}(\text{sk}, \text{pk}, \text{aux}, \mathcal{A}_2)$ that is the transcript of interaction of \mathcal{A}_2 with the challenger in G_P^2 , including the final output.

Remark 3.9. We make a few notes on the scope of security games we consider in this work:

1. We mainly focus on game-based (and falsifiable) security notions within the scope of this work. We need the challenger in the security to be efficient for the watermarking construction to be efficient.
2. All the admissible adversary need to be PPT. We do not consider security models other than polynomially bounded in time/circuit size (such as bounded storage).

3.2 Watermarking-Compatible Construction of Cryptographic Primitive

In this section, we characterize what type of black-box cryptographic constructions are watermarking compatible. As we will see in later sections, watermarking compatible constructions allow us to give a watermarking scheme for the target primitive constructed, black-boxly from the watermarking schemes of the building blocks.

Outline and Intuition Let us denote P as the target, or outcome primitive built in a construction. Let $P_i, i \in [k]$ denote each underlying primitive we use to build P .

Intuitively, one expects to watermark all underlying building blocks to ensure watermarking security of the target primitive P . As discussed in our technical overview section "More advanced watermarking constructions with unwatermarked building blocks", many constructions possess building blocks of primitives that do *not* need to be watermarkable to ensure the watermarkability of the final target primitive.

Here, we discuss the matter in more details: we need the partitioning for primitives in $\{P_i\}_{i \in [k]}$ into S and \bar{S} . These primitives in these two sets will play different roles when we construct the watermarking scheme for the target primitive P : the primitives in set S need to be watermarkable and those in set \bar{S} do not need to be watermarkable. The reasoning is that the primitives in set S will have their secret keys generated during the Key Generation of the target primitive P , but primitives in set \bar{S} will only have their secret keys generated freshly during every run of SecEval or PubEval algorithm of the target primitive P .

A simple example is CCA2-secure PKE scheme based on IBE and one-time signatures, which we discussed in the technical overview. During the encryption algorithm, one generates fresh one-time signature keys and the message is encrypted with the one-time signature's verification key as the identity. Therefore, the one-time signature in the above construction belongs to set \bar{S} because its keys are only generated during the encryption algorithm.

In the watermarking construction for P , the KeyGen, PubEval, SecEval algorithms will follow from the plain construction for P from primitives in set S and \bar{S} . Therefore, we cannot watermark the keys for the primitives in set \bar{S} because their keys are not generated when we give out the watermarked key (which only contains keys in set S) to a user. We therefore distinguish these two sets in our presentation for watermarking-compatible construction in this section and watermarkable implementation in Section 4.2.

Moreover, looking forward: we will observe that we indeed do *not* need watermarking security (i.e. unremovability) for the primitives in set \bar{S} , to achieve watermarking security for the target primitive P . We only need to rely on their "plain security" (e.g. unforgeability for a signature scheme, IND-CPA-security for an encryption scheme).

Notations Now we give formalization for the above outline.

Suppose a cryptographic primitive P is constructed black-boxly from primitives P_1, P_2, \dots, P_k in the following way, where $P_i, P_j, i \neq j$ are allowed to be the same primitive.

- Let S be some fixed subset of $[k]$ defined in the construction. S specifies two ways of using primitive P_i . The major difference is: for primitives $P_i, i \in S$, the algorithm $\text{KeyGen}(1^\lambda)$ will compute $P_i.\text{KeyGen}(1^\lambda)$ and the secret keys $P_i.\text{sk}$ generated will be used in the secret evaluation algorithm $w\text{SecEval}$.

Without loss of generality, we let the first $|\mathcal{S}|$ number of P_i 's be those corresponding to the set \mathcal{S} .

- For primitives $P_i, i \notin \mathcal{S}$, the algorithm $\text{KeyGen}(1^\lambda)$ will not compute $P_i.\text{KeyGen}(1^\lambda)$. They will either (1) have their keys generated freshly upon every run of SecEval or PubEval (2) will have a reference string (which can be viewed as a public key) generated during $\text{KeyGen}(1^\lambda)$, but have no secret keys or their secret keys will not be used in the SecEval algorithm of the target primitive P .
- For the sake of formality, we make a further division of the set $\bar{\mathcal{S}}$ into $\mathcal{T}_S, \mathcal{T}_P, \mathcal{T}_K$. They perform slightly different functionalities in the construction.
 - Let $\mathcal{T}_S \in [k]$ denote the set of indices i where P_i 's keys will be generated during the secret evaluation algorithm SecEval ;
 - Let $\mathcal{T}_P \in [k]$ denote the set of indices i where P_i 's keys will be generated during the public evaluation algorithm PubEval ;
 - Let \mathcal{T}_K denote the set of indices i where P_i have no secret keys/secret keys are not used in SecEval and will have a public reference string (which can be viewed as a public key) generated during $\text{KeyGen}(1^\lambda)$.

⁵.

Watermarking-Compatible Construction Syntax

$\text{KeyGen}(1^\lambda) \rightarrow (\text{sk}, \text{pk})$:

1. compute $(\text{sk}_i, \text{pk}_i) \leftarrow P_i.\text{KeyGen}(1^\lambda)$ for all $i \in \mathcal{S}$; compute $(\text{pk}_i, \text{td}_i) \leftarrow P_i.\text{KeyGen}(1^\lambda)$ for all $i \in \mathcal{T}_K$.
2. output $\text{sk} = (\{\text{sk}_i\}_{i \in \mathcal{S}})$; $\text{pk} = (\{\text{pk}_i\}_{i \in \mathcal{S} \cup \mathcal{T}_K})$.

$\text{SecEval}(\text{sk}, \text{pk}, x \in \mathcal{X}) \rightarrow y \in \mathcal{Y}_s$: is an algorithm that

1. uses $P_i.\text{SecEval}(\text{sk}_i, \text{pk}_i, \cdot), i \in \mathcal{S}$ as subroutines.
2. uses $P_i.\text{PubEval}(\text{pk}_i, \cdot), i \in \mathcal{S} \cup \mathcal{T}_K$ as subroutines.
3. computes $(\text{sk}_j, \text{pk}_j) \leftarrow P_j.\text{KeyGen}(1^\lambda)$ for some $j \in \mathcal{T}_S$ (can include these pk_j generated as part of the output).

$\text{PubEval}(\text{pk}, x \in \mathcal{X}) \rightarrow y \in \mathcal{Y}_p$: is an algorithm that

1. uses $P_i.\text{PubEval}(\text{pk}_i, \cdot), i \in \mathcal{S} \cup \mathcal{T}_K$ as subroutines.
2. computes $(\text{sk}_j, \text{pk}_j) \leftarrow P_j.\text{KeyGen}(1^\lambda)$ for all $j \in \mathcal{T}_P$ (may include these $\{\text{pk}_j\}_j$ generated as part of the output).

We say the construction is watermarking compatible if the above construction of P satisfies:

1. **Correctness of Construction** : the above construction of P satisfies a correctness property defined by predicate F_R . Moreover, the proof of this correctness property follows from correctness properties of P_1, \dots, P_k for predicates F_{R_i} respectively.

⁵As we will see, primitives in \mathcal{T}_k will have their secret keys (called trapdoors td here) used only in the security proofs.

Looking forward: the sets $\mathcal{T}_S, \mathcal{T}_P$ will play the same role in the watermarking reduction, and the set \mathcal{T}_K will incur a slightly different argument but will essentially play the same role as the other primitives in set $\bar{\mathcal{S}}$. That is, only their plain security is required to achieve the watermarking security of the target primitive P .

2. **Security of Construction:** The above construction satisfies security defined with game G_P . The security can be proved from the security properties of P_1, \dots, P_k with respect to some game G_{P_1}, \dots, G_{P_k} .
3. **Watermarking Compatible Reduction** The above security proof is of the following format, which we call *Watermarking-Compatible Reduction*, shown below.

Among the above 3 properties, the first two are natural properties that come with any black-box cryptographic constructions with provable security and correctness. We will focus on discussing property 3.

Watermarking Compatible Reductions On a high-level, a watermarking-compatible reduction is essentially a natural security reduction, except with the following feature: we view both the the security game G_P for the target security primitive P and the security game G_{P_i} for the input primitive P_i as two-stage games defined in Definition 3.4. The supposed adversary \mathcal{A} for G_P will just be any usual PPT adversary. But we restrict the reduction in stage 2 to be "oblivious" about the queries made by \mathcal{A} in stage 1: that is, it cannot pass on any queries made by \mathcal{A} to the stage 2 reduction but we should nevertheless make the reduction go through successfully.

We divide our watermarking compatible reductions into two types. Type 1 reduction captures the reduction from breaking the security of P to breaking the security of a primitive $P_i, i \in \mathcal{S}$, i.e. the primitives that have their secret keys generated in the KeyGen algorithm and used in the SecEval algorithm of P .

Type 2 reduction accordingly captures the reduction from breaking the security of P to breaking the security of a primitive $P_i, i \in \bar{\mathcal{S}}$.

The main difference is in how they are used in proving the unremovability security in the composed watermarking scheme. In the actual watermarking unremovability, type 1 reduction is supposed to be oblivious about the queries made by the adversary and will eventually lead to breaking the unremovability of the underlying primitive; type 2 reduction operates similarly but will eventually lead to breaking the plain security of the underlying primitive.

Watermarking-Compatible Reduction Type 1 First we consider reductions for primitives P_i for $i \in \mathcal{S}$.

1. Consider a reduction to P_i for the j -th $P_i \in \mathcal{S}$; let $\mathcal{A}_i = (\mathcal{A}_i^1, \mathcal{A}_i^2)$ be the two-stage adversary for the 2-stage game of primitive P_i defined in Definition 3.4.
2. \mathcal{A}_i receives public key pk_i
3. \mathcal{A}_i prepares $(\text{sk}_j, \text{pk}_j) \leftarrow \text{KeyGen}_j(1^\lambda)$ for all $j \neq i, j \in \mathcal{S}$ and \mathcal{A}_i sends all $(\{\text{pk}_j\}_{j \in \mathcal{S}})$ to \mathcal{A}_1 .
4. \mathcal{A}_i simulates the security game G_P for P with \mathcal{A} , while being the adversary in game G_{P_i} .
 - In stage 1, \mathcal{A}_i^1 provides the oracles needed for game G_P^1 and answer \mathcal{A}^1 's queries as follows:
 - For any (admissible) queries in the game G_P^1 , if sk_i is required to compute the answer for the query, \mathcal{A}_i can use the oracles provided in game $G_{P_i}^1$. To answer the entire query, \mathcal{A}_i may finish the rest of the computation using $\{\text{sk}_j\}_{j \neq i}$.
 - If only $\{\text{sk}_j\}_{j \neq i}$ are needed, \mathcal{A}_i answers the queries by itself.
 - During the interaction of stage-1 game, \mathcal{A}^1 and \mathcal{A}_i^1 generates an auxiliary information aux which is public to both of them.

- Upon entering stage 2, \mathcal{A}^1 generates an arbitrarily polynomial-size state st and gives it to \mathcal{A}^2 .
 - \mathcal{A}_i also enters its second stage \mathcal{A}_i^2 . \mathcal{A}_i^2 also receives all of $(\{\text{sk}_j\}_{j \neq i}, \{\text{pk}\}_{i \in \mathcal{S}}, \text{aux})$ from \mathcal{A}^1 . But \mathcal{A}_i^2 does *not* obtain any of \mathcal{A}^1 's queries in stage 1.
 - \mathcal{A}_i^2 simulates the stage-2 game G_P^2 for \mathcal{A}^2 , using the above strategy as \mathcal{A}_i^1 uses. Note that the oracle operations done by \mathcal{A}_i^2 and the conditions on admissible queries *cannot* be dependent on of \mathcal{A}^1 's queries in stage 1, because \mathcal{A}_i^2 cannot see these queries.
 - \mathcal{A}_i^2 also records all of \mathcal{A}^2 's queries.
 - In the challenge phase of G_P^2 , \mathcal{A}_i^2 receives a challenge input inp_i from the challenger.
 - In G_P 's challenge phase, \mathcal{A}_i^2 samples some randomness r and prepares a challenge input inp for \mathcal{A} using r and inp_i . \mathcal{A}_i^2 sends inp_i to \mathcal{A}_2 .⁶
 - \mathcal{A}_i^2 continues to simulate the oracles needed in G_P^2 game for \mathcal{A}^2 if there are further query stages.
5. \mathcal{A}_i computes an efficiently computable function $f_i(\text{out}, r, \mathcal{Q}, \text{inp}_i)$ on input of \mathcal{A}^2 's final answer out , the randomness used to prepare \mathcal{A}^2 's challenge input, \mathcal{A}^2 's queries \mathcal{Q} and \mathcal{A}_i^1 's challenge input inp_i , and gives it to the challenger of primitive P_i .

Remark 3.10. *In the actual watermarking unremovability reduction, \mathcal{A}_i^1 will receive a watermarked secret key $\tilde{\text{sk}}_i$ and simulate the queries required using the oracles in game $G_{P_i}^1$ with $\tilde{\text{sk}}_i$ instead.*

More generically, we can also model the queries made by \mathcal{A}^1 in stage 1, related to the keys in the set \mathcal{S} , $\{\text{sk}_i\}_{i \in \mathcal{S}}$ as some arbitrary polynomial size (admissible) leakage on these keys. In the plain security reduction, such leakage correspond to \mathcal{A}^1 's adaptive queries made to the oracles provided. In the watermarking unremovability game, it corresponds to admissible marking queries where we give out marked secret keys.

Watermarking-Compatible Reduction Type 2 Watermarking-compatible reduction type 2 operates essentially the same as the above reduction for P_i but for some $i \notin \mathcal{S}$.

The main difference is that in the actual watermarking unremovability security game, in order to simulate oracle queries for G_P^1 , the stage 1 reduction \mathcal{A}_i^1 will actually need oracles in the game $G_{P_i}^1$. Naturally, this is because the primitives in set \mathcal{S} are not watermarked in the watermarking composition and the corresponding reduction will not get a watermarked key from the challenger, but only the oracles provided in the plain security game of G_{P_i} . Even though this also means that the reduction \mathcal{A}_i^1 gets to observe \mathcal{A}^1 's queries (using the oracles of the game G_{P_i}) in the actual watermarking security game, and can make use of them in stage 2, \mathcal{A}_i^1 does not need to make any of such queries to the challenger since the keys of any primitive in set \mathcal{S} will only be sampled online by \mathcal{A}_i^1 itself. Meanwhile, \mathcal{A}_i^2 should still not inherit any queries related to the "watermarked" primitives in set \mathcal{S} since in the real watermarking unremovability game, \mathcal{A}_i^2 is not supposed to see them.

1. Consider doing reduction to P_i for some $P_i, i \notin \mathcal{S}$; let $\mathcal{A}_i = (\mathcal{A}_i^1, \mathcal{A}_i^2)$ be the two-stage adversary for primitive P_i , where the stages are defined by Definition 3.4.
2. \mathcal{A}_i receives public key pk_i from the challenger.
3. \mathcal{A}_i prepares $(\text{sk}_\ell, \text{pk}_\ell) \leftarrow P_\ell.\text{KeyGen}(1^\lambda)$ for all $\ell \in \mathcal{S}$. \mathcal{A}_i sends all $(\{\text{pk}_\ell\})$ to \mathcal{A}_1 .
4. \mathcal{A}_i simulates the security game G_P for P with \mathcal{A} :

⁶Note that G_P 's challenge phase may be before, after or concurrent with G_{P_i} 's challenge phase depending on the reduction. Similarly, the challenge \mathcal{A}^2 receives may be dependent on inp .

- In stage-1 game G_P^1 , for any (admissible) queries if $sk_\ell, \ell \in \mathcal{S}$ is required to compute the answer for the query, \mathcal{A}_i can answer the query since it possesses the keys $\{sk_\ell\}_{\ell \in \mathcal{S}}$. The secret keys of primitives $\{P_j\}_{j \notin \mathcal{S}}$ are all sampled freshly upon every run of $wSecEval(sk, pk, \cdot)$ (or sampled freshly in $wPubEval(pk, \cdot)$), which \mathcal{A}_i^1 can do it on its own.
 - During the interaction of stage-1 game, \mathcal{A}^1 and \mathcal{A}_i^1 generates an auxiliary information aux which is public to both of them.
 - Entering stage 2, \mathcal{A}^1 generates an arbitrarily polynomial-size state st and gives it to \mathcal{A}^2 . \mathcal{A}_i also enters its second stage \mathcal{A}_i^2 . \mathcal{A}_i^2 also receives all of $(\{sk_\ell\}_{\ell \in \mathcal{S}}, \{pk_\ell\}_{\ell \in \mathcal{S}}, aux)$ but does *not* obtain any of \mathcal{A}^1 's queries involving the use of $\{sk_\ell\}_{\ell \notin \mathcal{S}}$.
 - \mathcal{A}_i^2 continues to provide the oracles needed for game G_P^2 and answer $\mathcal{A}^2(st)$'s queries using the above strategy as \mathcal{A}_i^1 uses.
 - \mathcal{A}_i^2 also records all of \mathcal{A}^2 's queries.
 - In the challenge phase of G_{P_i} , \mathcal{A}_i^2 receives a challenge input inp_i from the challenger.
 - In G_P^2 's challenge phase, \mathcal{A}_i^2 samples some randomness r and prepares a challenge input inp for \mathcal{A} using r and inp_i . \mathcal{A}_i^2 sends inp_i to \mathcal{A}_2 .
 - \mathcal{A}_i^2 continues to simulate the oracles needed in G_P^2 game for \mathcal{A}^2 if there are further query stages
5. \mathcal{A}_i computes an efficiently computable function $f_i(out, r, \mathcal{Q}, inp_i)$ on input of \mathcal{A}^2 's final answer out , its queries \mathcal{Q} and \mathcal{A}_i 's challenge input inp_i and secret randomness r , and gives it to the challenger of primitive P_i .

Properties of Watermarking-Compatible Reductions We further give some properties of watermarking-compatible reduction. These properties come with any natural black-box security roof with hybrid argument and reductions. We present them here for the sake of convenience and use them as facts later.

Fact 3.11 (Reduction Property 1). *A watermarking-compatible reduction guarantees that: if there exists some \mathcal{A} such that the advantage of \mathcal{A}^2 winning the game G_P is non-negligible, i.e. $\Pr[G_P(1^\lambda, \mathcal{A}) = 1] \geq \eta + \gamma$ where η is the trivial winning probability and γ is non-negligible in λ , then there exists some $i \in [k]$ such that \mathcal{A}_i , using the above reduction strategy, wins G_{P_i} with some non-negligible advantage γ_i .*

Remark 3.12. *The above property follows naturally from any hybrid argument of proof.*

If using the language of the 2-stage security game: If there exists some adversary $\mathcal{A} = (\mathcal{A}^1, \mathcal{A}^2)$ such that for some non-negligible ϵ and some non-negligible γ , we have:

$$\Pr \left[\Pr[G_P^2(sk, pk, aux, \mathcal{A}_2(st)) = 1] \geq \eta + \gamma : \{(sk, pk), aux, st\} \leftarrow G_P^1(\mathcal{A}_1, 1^\lambda) \right] \geq \epsilon$$

Then there exists some $i \in [k]$ such that $\mathcal{A}_i = (\mathcal{A}_i^1, \mathcal{A}_i^2)$, using the above reduction strategy, such that for some non-negligible γ_i, ϵ_i , we have:

$$\Pr \left[\Pr[G_{P_i}^2(sk_i, pk_i, aux, \mathcal{A}_i^2(st)) = 1] \geq \eta + \gamma_i : \{(sk_i, pk_i), aux, st\} \leftarrow G_{P_i}^1(\mathcal{A}_i^1, 1^\lambda) \right] \geq \epsilon_i.$$

Fact 3.13 (Reduction Property 2). *For all $i \in \mathcal{S}$, once give the secret key $P_i.sk$ in the clear, there exists a PPT algorithm T_i that wins the security game G_P with probability 1.*

Additionally, such a T_i can be used black-boxly by the watermarking-compatible reduction algorithm \mathcal{A}_i to win the security game G_{P_i} with noticeable probability.

Reduction Function Recall that in the end of the reduction, \mathcal{A}_i computes an efficiently computable function $f_i(\text{out}, r, \mathcal{Q}, \text{inp}_i)$ on input of \mathcal{A}^2 's final answer out , the randomness used to prepare \mathcal{A}^2 's challenge input, \mathcal{A}^2 's queries \mathcal{Q} and \mathcal{A}_i 's challenge input inp_i , and gives it to the challenger of primitive P_i .

For convenience, we will refer to the function f_i used by the reduction as *reduction function*.

Remark 3.14 (Reduction Function). *We refer to the above function f_i used by the reduction as reduction function.*

Remark 3.15 (Special Primitives in \mathcal{T}_K). *The role of the primitives in the set \mathcal{T}_K may be confusing to the readers at this point. We make some further explanation. As we go into later sections and go into examples (Appendix A), its role will become clear.*

More specifically, \mathcal{T}_k only contains the simulation-based primitives P_ℓ where

1. *there exists an efficient simulator algorithm such that the output of $P_\ell.\text{PubEval}(P_\ell.\text{pk}, \cdot)$ is indistinguishable from the output of the simulator $\text{Sim}(P_\ell.\text{td}, \cdot)$. Their secret key (trapdoor td) will only come up in the security proof for P .*
2. *In the security proof for P , the rest of the hybrid arguments and reduction happen in a hybrid world where we have already invoked the above security for P_ℓ .*

The one example primitive in the set \mathcal{T}_K we use within this work is a NIZK scheme with a trapdoor. Its "keys" (the common reference string CRS, and trapdoor td) are generated in the main key generation algorithm, but the trapdoor is not used in any of SecEval , PubEval , and only in the security proof.

4 Watermarking Composition Framework

4.1 Definition: Watermarkable Implementation of a Cryptographic Primitive

In this section, we will define what we call a "watermarkable implementation" of a cryptographic primitive P . We distinguish it from the usual naming ("watermarkable P") because of some differences in syntax and definition. But we will sometimes use watermarkable P in our work for convenience. Please note that all "watermarkable P" used in the technical part of this work refers to a watermarkable implementation of a cryptographic primitive P defined below.

On a high-level, a watermarkable implementation of P differs from most existing watermarking definition in two aspects:

1. **Unremovability security:** A pirate circuit produced by the adversary in the unremovability security game is considered to function successfully as long as it can break the security game G_P .
2. **Extraction algorithm syntax:** the extraction key used to extract a watermark is able to simulate the game G_P for the underlying "plain security" of P .

In more detail, a watermarkable implementation of a cryptographic primitive has the following syntax and properties.

Watermarkable Primitive Syntax A watermarkable primitive WP for a cryptographic primitive P with a security game G_P consists of the following algorithms:

$\text{WMSSetup}(1^\lambda) \rightarrow (\text{sk}, \text{pk}, \text{mk}, \text{xk})$: on security parameter, outputs a secret key sk , public key pk , marking key mk , extraction key xk .
 $\text{wSecEval}(\text{sk}, \text{pk}, x)$: takes in the secret information sk , public information pk and some input x from input space \mathcal{X} , and output a value $y \in \mathcal{Y}_s$.
 $\text{wPubEval}(\text{pk}, x \in \mathcal{X})$: takes in the public information pk and some input x from input space \mathcal{X} , and output a value $y \in \mathcal{Y}_p$.
 $\text{Mark}(\text{mk}, \text{sk}, \tau \in \mathcal{M}_{\text{Mark}}) \rightarrow \text{sk}_\tau$: takes in marking key mk , a secret key sk and a message $\tau \in \mathcal{M}_{\text{Mark}}$; output a marked key sk_τ ⁷.
 $\text{Extract}(\text{xk}, \text{pk}, \text{aux}, C) \rightarrow \tau \in \mathcal{M}_{\text{Mark}} / \vec{\tau} \in \mathcal{M}_{\text{Mark}}^k / \perp$: on input extraction key sk , public key pk , auxiliary information aux and circuit C , outputs a mark message $\tau \in \mathcal{M}_\tau$ or a tuple of marked messages $\vec{\tau} \in \mathcal{M}_\tau^k$ where k is a constant/small polynomial parameter related to the concrete watermarking construction.

Remark 4.1. *The Extract algorithm is allowed to output a tuple of marks when working on adversarial programs, when the watermarkable implementation is one that comes from composing underlying watermarking implementations. More details to be discussed later in Section 4.2.*

Remark 4.2. *The extraction algorithm takes in the public key pk (which is not a limitation because it's public) and an auxiliary input aux . As we will see in the concrete examples, depending on the security game of the primitives we watermark, we may need Extract to take in some aux or may not.*

When it is clear from the context that there is no public key or auxiliary information for Extract to take, we will omit them from the input parameters for notation cleanness.

Remark 4.3. *There can be several different (constant number of) wSecEval and wPubEval . As aforementioned, we view all secret algorithms as one algorithm $\text{SecEval}(\text{sk}, \text{pk}, \cdot)$ that will take in an index $i \in [\ell]$ to decide which mode to use, similarly for public algorithms. We thus also use sk to denote $(\{\text{sk}_i\}_{i \in [\ell]})$.*

Let us denote $\text{Eval} = (\text{wSecEval}, \text{wPubEval})$.

Correctness The construction should satisfy "unmarked" correctness for unmarked keys: a watermarkable implementation of P is said to be correct with respect to predicate F_R if there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}, x \in \mathcal{X}$:

$$\Pr[F_R(\text{Eval}, \text{pk}, \text{sk}, x, r) = 1 : (\text{sk}, \text{pk}) \leftarrow \text{WMSSetup}(1^\lambda), r \leftarrow D_r] \geq 1 - \text{negl}(\lambda).$$

where the probability is taken over randomness used in WMSSetup and D_r . Recall that randomness r is used in checking whether the predicate $F_R(\text{Eval}, \text{pk}, \text{sk}, x, r)$ is satisfied Definition 3.1 and D_r can be simply taken to be the uniform distribution.

Functionality Preserving vs. Exact Functionality Preserving We present both definitions of functionality-preserving for the sake of comprehensiveness because they have both appeared in watermarking literatures. Before [GKM⁺19], the watermarking literature mainly used exact functionality preserving only. While it is not particularly vital to the contributions in this work, we

⁷In definitions in the watermarking literature, the output of Mark is the program $\text{wSecEval}(\text{sk}_\tau, \text{pk}, \cdot)$. Since conventionally by Kerckhoff's principle, the evaluation algorithm itself is public, giving out only the marked secret key is equivalent to giving out the program, we give out the key for convenience that come up later. There are watermarking schemes where the evaluation algorithm may be different when running on a marked key; in that case we can consider wSecEval to have two modes, one for unmarked keys one for marked keys.

would like to mention that exact functionality preserving is a stronger notion. For some of the underlying building blocks we use, only constructions under the relaxed functionality-preserving definition are known, for example the watermarkable signature scheme in [GKM⁺19].

Functionality Preserving The functionality-preserving property says: if for a predicate F_R , the underlying primitive P satisfies the correctness in Definition 3.2, then there exists a negligible function $\text{negl}(\lambda)$ such that for all $\lambda \in \mathbb{N}, x \in \mathcal{X}$:

$$\Pr \left[F_R(\text{Eval}, \text{pk}, \tilde{\text{sk}}, x, r) = 1 : \begin{array}{l} (\text{sk}, \text{pk}, \text{xk}, \text{mk}) \leftarrow \text{WMSetup}(1^\lambda) \\ \tilde{\text{sk}} \leftarrow \text{Mark}(\text{mk}, \text{sk}, \tau), r \leftarrow D_r \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

Exact Functionality Preserving The exact functionality preserving is a stronger property that: there exists a negligible function $\text{negl}(\lambda)$ such that for all $\lambda \in \mathbb{N}, x \in \mathcal{X}, \tau \in \mathcal{M}_\tau$:

$$\Pr \left[\text{wSecEval}(\tilde{\text{sk}}, \text{pk}, x) = \text{SecEval}(\text{sk}, \text{pk}, x) : \begin{array}{l} (\text{sk}, \text{pk}, \text{xk}, \text{mk}) \leftarrow \text{WMSetup}(1^\lambda) \\ \tilde{\text{sk}} \leftarrow \text{Mark}(\text{mk}, \text{sk}, \tau) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

Correctness of Extraction There exists a publicly known efficient algorithm T and a negligible function $\text{negl}(\lambda)$ such that for all $\lambda \in \mathbb{N}$, for all $\tau \in \mathcal{M}_\tau$:

$$\Pr \left[\text{Extract}(\text{xk}, T(\text{sk}_\tau, \cdot)) = \tau : \begin{array}{l} (\text{sk}, \text{pk}, \text{xk}, \text{mk}) \leftarrow \text{WMSetup}(1^\lambda) \\ \tilde{\text{sk}} \leftarrow \text{Mark}(\text{mk}, \text{sk}, \tau) \end{array} \right] \geq 1 - \text{negl}(\lambda)$$

Security with Security Game G_P The watermarkable primitive WP satisfies the same security defined by the security game G_P for P , except that the subroutines used in G_P , KeyGen, SecEval, PubEval are replaced with WMSetup, wSecEval, wPubEval respectively. The Extract, Mark algorithms and keys are ignored in the context of these games.

Watermarking Security Compatible with A Security Game Now we present the unremovability security of the watermarking implementation for primitive P .

Informally, the security guarantees that: any PPT adversary \mathcal{A} , when given the watermarked key(s), generates a pirate circuit C^* ; if C^* can win the security game G_P with some non-negligible advantage, then we must be able to extract a (previously queried) watermark from C^* .

Definition 4.4 (γ -Unremovability with Game G_P). *We say a watermarkable implementation of P is γ -unremovable if:*

For every stateful γ -unremovable admissible PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, the following holds:

$$\Pr \left[\text{Extract}(\text{xk}, \text{pk}, \text{aux}, C^*) \notin \mathcal{Q} : \begin{array}{l} (\text{sk}, \text{pk}, \text{mk}, \text{xk}) \leftarrow \text{WMSetup}(1^\lambda) \\ (\text{aux}, C^*) \leftarrow \mathcal{A}^{\text{Mark}(\text{mk}, \text{sk}, \cdot)}(1^\lambda, \text{pk}) \end{array} \right] \leq \text{negl}(\lambda)$$

\mathcal{Q} denotes the set of f marks queried by \mathcal{A} to the marking oracle $\text{Mark}(\text{mk}, \text{sk}, \cdot)$, G_P^1 is the stage 1 of a security game G_P , as defined in Definition 3.4.

We call a PPT adversary \mathcal{A} as γ -unremovable admissible if \mathcal{A} 's output C^* is an admissible adversary in the stage-2 security game G_P^2 and is γ -good, where γ -good is defined as:

$$\Pr[G_P^2(\text{sk}, \text{pk}, \text{aux}, C^*) = 1] \geq \eta + \gamma$$

Here, η is the trivial success probability for any admissible adversary in G_P .

The randomness over testing whether C^* is γ -good is the randomness used answering any oracle queries from C^* and preparing the challenge for C^* , in the stage-2 security game G_P^2 .

To match our syntax, we denote $\text{Extract}(\text{xk}, \text{pk}, \text{aux}, C^*) \notin \mathcal{Q}$ to mean that:

1. If $\text{Extract}(\text{xk}, \text{pk}, \text{aux}, C^*)$ outputs a single mark $\tau \in \mathcal{M}_\tau / \tau = \perp$, it is considered to be not in the query set \mathcal{Q} if and only if $\tau \notin \mathcal{Q}$.
2. If $\text{Extract}(\text{xk}, \text{pk}, \text{aux}, C^*)$ outputs a tuple of marks $\vec{\tau} = (\tau_1, \dots, \tau_k)$, where each $\tau_i \in \mathcal{M}_\tau / \tau_i = \perp$, it is considered not in the query set \mathcal{Q} if and only if for all i , $\tau_i \notin \mathcal{Q}$.

Remark 4.5. The Extract algorithm will only output a tuple of marks when used on adversarial circuits C . Because of composition, the adversary might be able to acquire a composed watermarked circuit where each "component" has a different watermark, but this does not formulate an attack as long as we can still extract any valid (previously queried) watermarks from the circuit. We will discuss this aspect in more details in Section 4.2.

Remark 4.6 (Discussions on the γ -Unremovability Definition). We give the above definition with a parameter γ and it helps illustrate how this definition corresponds to the 2-stage security game in Definition 3.4: \mathcal{A} corresponds to the "stage-1" adversary and the circuit C^* corresponds to stage-2 adversary.

If we want a watermarking scheme's γ to be any non-negligible function (the strongest notion here) in terms of the security parameter, we can simply require it to satisfy γ -unremovability for any non-negligible γ (defined below as "strong unremovability").

In our paper, all our underlying watermarking building blocks used for composition satisfy γ -unremovability for any non-negligible gamma. So they are not sensitive to γ . Our statements about the watermarking compiler also only consider this case (see Lemma 4.21).

Besides, we also inherit this γ -removability definition partially from [GKM⁺19]. In some other watermarking literatures, it may be also meaningful to consider a weaker security where gamma is not necessarily negligible.

Definition 4.7 (Strong Unremovability). We say a watermarking implementation of P satisfies strong unremovability if it satisfies γ -unremovability for any non-negligible γ .

Remark 4.8 (Statefulness of the Adversarial Program). Since the adversarial circuit in our security game is assumed to be interactive and stateful, one may ask how stateful it can be.

This issue with stateful pirates is almost always present in the traitor tracing and watermarking literature. For this reason, the vast majority of works operate in a stateless program model. This requires an assumption that the pirate program can be reset by the tracer to its initial conditions, which makes sense if, say, the program is given as software. In our work, we allow the pirate program to be stateful, in the sense that we allow it to play an interactive security game. However, we always assume between runs of the game that the program is reset to its initial conditions. This notion of statefulness generalizes the existing watermarking/traitor tracing definitions from non-interactive games to interactive games, while also avoiding the issue of self-destruction. Like in the existing models, our model makes sense if, say, the program is given as software.

We leave more remark on the unremovability definition at the end of this section Remark 4.12.

Security Game Simulation Property of the Extraction Key We additionally require the extraction key `Extract` to be able to simulate the (stage-2) security game G_P^2 for the primitive P to be watermarked.

Consider the security game G_P for the underlying primitive P : $G_P = (G_P^1, G_P^2)$ is an interactive game between a challenger and admissible $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$. Recall that in the stage-2 Definition 3.4 G_P^2 can output a view $\text{view}(\text{sk}, \text{pk}, \text{aux}, \mathcal{A}_2)$ (Definition 3.8).

Definition 4.9 (Extraction key simulation property). *The extraction key simulation property of the extraction key says that: given $(\text{sk}, \text{pk}, \text{aux}) \leftarrow G_P^1(\mathcal{A}_1, 1^\lambda)$ where we use $\text{WMSetup}(1^\lambda)$ to obtain $(\text{sk}, \text{pk}, \text{xk})$ in G_P^1 , there exists a PPT simulator $\text{Sim}(\text{xk}, \text{pk}, \text{aux}, \mathcal{A}_2)$ (where Sim interacts with \mathcal{A}_2 black-boxly) that outputs a simulated view $\text{view}_{\text{Sim}}(\text{xk}, \text{pk}, \text{aux}, \mathcal{A}_2)$, such that for any $\lambda \in \mathbb{N}$, for any admissible \mathcal{A} in G_P , the distributions $\text{view}_{G_P^2} = \{\text{view}(\text{sk}, \text{pk}, \text{aux}, \mathcal{A}_2)\}$ and $\text{view}_{\text{Sim}} = \{\text{view}_{\text{Sim}}(\text{xk}, \text{pk}, \text{aux}, \mathcal{A}_2)\}$ are perfectly/statistically/computationally indistinguishable.*

Extraction Syntax: Simulation of Security Game G_P inside `Extract` Following the above property, we require that the `Extract` algorithm in a watermarkable implementation of a primitive P follows a specific format: it simulates the stage-2 security game G_P^2 for any input circuit, while trying to extract a mark.

Extract Algorithm

On input $(\text{xk}, \text{pk}, \text{aux})$ and a program C :

- Run an algorithm E where E uses the following subroutine (for possibly $\text{poly}(\lambda)$) many times):
 - Use $(\text{xk}, \text{pk}, \text{aux})$ to simulate the stage-2 game $G_P^2(\text{sk}, \text{pk}, \text{aux}, C)$ for primitive P , running C black-boxly by treating C as the stage-2 adversary \mathcal{A}_2 defined in Definition 3.4.
 - If C is non-interactive (i.e. C does not make any queries or respond to interaction in the challenge phase of G_P^2), then E samples challenge inputs on its own and runs C on them.
- E can take any of C 's outputs, including C 's queries made during the above simulated game, as inputs to compute the extraction algorithm's final output.
- Output a watermark $\tau \in \mathcal{M}_\tau$ or \perp

Examples of Extraction Algorithm that Simulates a Security Game It is not hard to create an extraction algorithm with the above syntax and simulation capability. In fact, some existing works have already built watermarking schemes that have `Extract` algorithm with such a format.

- To enable the extraction key to simulate the security game, a naive solution in the private extraction case is to simply let xk contain the secret key sk . One example is the extraction algorithm in a watermarkable signature scheme (see Appendix D). To answer the signature queries for the pirate program, the extraction key is the signing key.

- In some other scenarios, we don't need the secret key to simulate the security game and one can even have public extraction.

For example, a CPA secure PKE scheme. Another example is when we can sample from the input-output space of the evaluation oracles without having the actual key: in the weak PRF setting (with non-adaptive queries), we can simply answer its queries by sampling. In [GKWW21], one can use indistinguishability obfuscation to build such a sampler, so that we have watermarkable weak PRF with public extraction.

Remark 4.10 (Extraction of watermarks from honestly generated circuit). *Note that even though our extraction algorithm simulates a possibly interactive game for any input program, it does not require or assume the input program to be interactive or even stateful. As we have described, when the input circuit is not interactive, then the extraction algorithm simply samples challenge inputs and run C on them to get C 's outputs.*

For example, we can also correctly extract from an honestly watermarked circuit which is neither stateful nor interactive.

Remark 4.11 (Computational Simulation). *In most of our applications, the extraction key can simulate the security game's view with perfect indistinguishability. But whether the simulation quality is perfect/statistical/computational does not affect our watermarking composition theorem and applications, as long as the difference is negligible.*

We will encounter computational simulation only when there are primitives in the set \mathcal{T}_K : the extraction algorithm will simulate a "hybrid version" of the security game, after invoking the security of the primitive in \mathcal{T}_K , instead of the original security game. See more discussions Remark 4.16.

The one example provided in our paper is the use of NIZK proof with trapdoors in Appendix A.

Meaningfulness The meaningfulness property says that an honestly generated key should not have watermarks.

A watermarking implementation of primitive P scheme satisfies the meaningfulness property if there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$:

$$\Pr [\text{Extract}(xk, sk) \neq \perp : (sk, pk, mk, xk) \leftarrow \text{WMSetup}(1^\lambda)] \leq \text{negl}(\lambda)$$

Private Extraction vs. Public Extraction A watermarkable implementation of a cryptographic primitive is a scheme with public-extraction if the extraction key xk can be made public while the above properties still hold. Otherwise it has private extraction.

Single-key vs. Collusion Resistant A watermarkable implementation of a cryptographic primitive is collusion resistant if in the unremovability security game, \mathcal{A} is allowed to query the marking oracle $\text{Mark}(mk, sk, \cdot)$ for arbitrarily many times and q -bounded collusion resistance if it is allowed to query for a-priori bounded polynomial q -times. the scheme is single-key secure if $q = 1$.

Remark 4.12 (Defining stages in the unremovability game). *As discussed briefly in Definition 3.4: an example of subtlety when defining the staged unremovability game is: in the original security game for P , \mathcal{A} needs to commit to some challenge messages that will be used by the challenger to prepare the final challenge for \mathcal{A} . The question is whether we make \mathcal{A} commit to such challenge messages in stage 1, or let the stage-2 adversary, i.e. the circuit C output by \mathcal{A} choose their challenge messages later in stage 2.*

In summary, if admissibility of queries made in stage 1 will be dependent on the challenge, then we must make \mathcal{A} commit to its own part on the challenge (as some auxiliary information aux) before it produces C , and we run C with these prefixed aux in stage-2 game $G_P^2(\text{pk}, \text{xk}, \text{aux}, C)$.

Let's take a "flexible" example, where we can let the stage-2 adversary \mathcal{A}_2 do such commitment, instead of forcing it to commit in stage 1. Accordingly, we can let the watermarking adversary's output program C selects its own commitment when we run C : for example, in CPA/CCA encryption games, we can allow \mathcal{A} to submit its challenge messages (m_0, m_1) either together/before it outputs C or output a C that will choose its own (m_0, m_1) . This flexibility comes from the fact that whether the queries made by \mathcal{A} and C are valid does not depend on the choice of (m_0, m_1) .

A non-flexible example would be the challenge attribute/identity in ABE/IBE. We must let \mathcal{A} choose its challenge attribute/identity before/at the same time with outputting C . C is simulated in a second stage security game with this pre-selected attribute/identity as its input. Otherwise \mathcal{A} can easily cheat because the extraction algorithm(stage-2 game) does not get to check if the previous key generation queries made by \mathcal{A} are valid with respect to challenge attribute/identity.

4.2 Watermarking Composition: Target Primitive from Input Primitives

In this section, we show that if primitive P is built from P_1, \dots, P_k where the security can be shown via a watermarking-compatible reduction, we can construct a watermarkable implementation of P , named WP to satisfy definition in 4.1, from existing watermarkable implementations of P_1, \dots, P_k called $\text{WP}_1, \text{WP}_2, \dots, \text{WP}_k$, satisfying definition in 4.1. We will still refer to P as the target primitive and P_1, \dots, P_k as input primitives.

Outline and Intuition On a high level, the watermarking scheme of the target primitive P simply follows the construction of plain P from the plain underlying building blocks in terms of evaluation algorithms $\text{SecEval}, \text{PubEval}$. Correctness and functionality-preserving compose in a relatively natural way.

To mark a key of P , the marking algorithm concatenates the marked keys of all underlying P_i that need to be marked. To extract a mark, we attempt to run the extraction algorithm of all underlying P_i on the input circuit. If any valid mark is extracted, then the circuit is considered marked.

In particular, the extraction algorithm will treat the circuit as an adversary in the stage-2 security game of P and turns it black-boxly into a reduction for the security game of each underlying P_i , one by one and then run the underlying extraction algorithm of P_i on it.

In order to remove a mark from a P 's key, the adversary \mathcal{A} must remove all marks from each P_i 's key. Meanwhile, the pirate program still needs to win the security game of P , then we must be able to use to break the unremovability of at least one underlying P_i . In more generic scenarios, the pirate program made by \mathcal{A} may not break any unremovability security of watermarkable building blocks, but get around the task of removing marks by breaking the security of some unwatermarked building blocks. By similar means, we can use the pirate program to build our reduction to the security of these unwatermarked building blocks. By the properties of the watermarking-compatible reduction (Section 3.2) that P 's construction satisfies, the above analysis is exhaustive.

Construction of WP Similar to the description of construction in Section 3.2, we recall the following notations:

- Let $\mathcal{S} \subset [k]$ be a fixed set used in P 's construction from P_1, \dots, P_k , where the primitives P_i with $i \in \mathcal{S}$'s keys will be generated in the KeyGen algorithm of P . Without loss of generality, we let the first $|\mathcal{S}|$ number of P_i 's be those corresponding to the set \mathcal{S} .
- Let $\mathcal{T}_S \in [k]$ denote the set of indices i where P_i 's keys will be generated during SecEval; Let $\mathcal{T}_P \in [k]$ denote the set of indices i where P_i 's keys will be generated during PubEval. Let \mathcal{T}_K denote the set of indices i where P_i 's secret key (we call trapdoor td) will be generated during KeyGen(1^λ) but will not be used in P 's algorithms.
- Without loss of generality, we sometimes assume a numbering on the primitives so that the first $|\mathcal{S}|$ primitives are in the set \mathcal{S} .

WMS $\text{Setup}(1^\lambda) \rightarrow (\text{sk}, \text{pk}, \text{xk}, \text{mk}) :$

1. compute $(\text{sk}_i, \text{pk}_i, \text{xk}_i, \text{mk}_i) \leftarrow \text{WP}_i.\text{WMS}\text{Setup}(1^\lambda)$ for all $i \in \mathcal{S}$.
2. compute $(\text{pk}_\ell, \text{td}_\ell) \leftarrow P_\ell.\text{KeyGen}(1^\lambda)$ for all $\ell \in \mathcal{T}_K$;
3. output $\text{sk} = (\text{sk}_{j_1}, \dots, \text{sk}_{|\mathcal{S}|}); \text{pk} = (\text{pk}_1, \dots, \text{pk}_{|\mathcal{S}|}, \{\text{pk}_\ell\}_{\ell \in \mathcal{T}_K});$
 $\text{xk} = (\text{xk}_1, \dots, \text{xk}_{|\mathcal{S}|}, \{\text{td}_\ell\}_{\ell \in \mathcal{T}_K}); \text{mk} = (\text{mk}_1, \dots, \text{mk}_{|\mathcal{S}|})$

wSecEval(sk, pk, x):

1. parse input $\text{sk} = (\text{sk}_1, \dots, \text{sk}_{|\mathcal{S}|}); \text{pk} = (\text{pk}_1, \dots, \text{pk}_{|\mathcal{S}|}, \{\text{pk}_\ell\}_{\ell \in \mathcal{T}_K})$.
2. wSecEval($\text{sk}, \text{pk}, \cdot$) is the same algorithm as $P.\text{SecEval}(\text{sk}, \text{pk}, \cdot)$ in the construction of P from P_1, \dots, P_k , except that $P_i.\text{SecEval}(\text{sk}_i, \cdot)$ is replaced with $\text{WP}_i.\text{wSecEval}(\text{sk}_i, \cdot)$ for $i \in \mathcal{S}$. Overall, wSecEval($\text{sk}, \text{pk}, \cdot$) is an algorithm that:
 - (a) uses $\text{WP}_i.\text{wSecEval}(\text{sk}_i, \text{pk}_i, \cdot), i \in \mathcal{S}$ as subroutines.
 - (b) uses $\text{WP}_i.\text{wPubEval}(\text{pk}_i, \cdot), i \in \mathcal{S} \cup \mathcal{T}_K$ as subroutines.
 - (c) computes $(\text{sk}_j, \text{pk}_j) \leftarrow P_j.\text{KeyGen}(1^\lambda)$ for some $j \in \mathcal{T}_S$ (and may include these pk_j generated as part of the output).

wPubEval(pk, x)

1. parse input $\text{pk} = (\text{pk}_1, \dots, \text{pk}_{|\mathcal{S}|}, \{\text{pk}_\ell\}_{\ell \in \mathcal{T}_K})$.
2. wPubEval(pk, \cdot) is the same algorithm as $P.\text{PubEval}(\text{pk}, \cdot)$ in the construction of P from P_1, \dots, P_k , except that $P_i.\text{PubEval}(\text{pk}_i, \cdot)$ is replaced with $\text{WP}_i.\text{wPubEval}(\text{pk}_i, \cdot)$ for $i \in \mathcal{S} \cup \mathcal{T}_K$. Overall, wPubEval($\text{sk}, \text{pk}, \cdot$) is an algorithm that:
 - (a) uses $\text{WP}_i.\text{wPubEval}(\text{pk}_i, \cdot), i \in \mathcal{S}$ as subroutines.
 - (b) uses $\text{WP}_i.\text{wPubEval}(\text{pk}_i, \cdot), i \in \mathcal{S}$ as subroutines.
 - (c) computes $(\text{sk}_j, \text{pk}_j) \leftarrow P_j.\text{KeyGen}(1^\lambda)$ for some $j \in \mathcal{T}_P$ (and may include these pk_j generated as part of the output).

Mark($\text{mk}, \text{sk}, \tau \in \mathcal{M}_{\text{Mark}}$) $\rightarrow \text{sk}_\tau$

1. parse $\text{mk} = (\text{mk}_1, \dots, \text{mk}_{|\mathcal{S}|}); \text{sk} = (\text{sk}_1, \dots, \text{sk}_{|\mathcal{S}|}); \text{pk} = (\text{pk}_1, \dots, \text{pk}_{|\mathcal{S}|}, \{\text{pk}_\ell\}_{\ell \in \mathcal{T}_K})$.
2. Compute $\text{sk}_{i,\tau} \leftarrow \text{WP}_i.\text{Mark}(\text{mk}_i, \text{sk}_i, \tau)$ for all $i \in \mathcal{S}$.
3. output $\tilde{\text{sk}} = (\text{sk}_{1,\tau}, \dots, \text{sk}_{|\mathcal{S}|,\tau})$

Extract($\text{xk}, \text{pk}, \text{aux}, C$) $\rightarrow \tau \in \mathcal{M}_{\text{Mark}} / \vec{\tau} \in \mathcal{M}_{\text{Mark}}^{|\mathcal{S}|} / \perp :$

1. parse $\text{xk} = (\text{xk}_1, \dots, \text{xk}_{|\mathcal{S}|}, \{\text{td}_\ell\}_{\ell \in \mathcal{T}_K}); \text{aux} = (\text{aux}_1, \dots, \text{aux}_{|\mathcal{S}|}); \text{pk} = (\text{pk}_1, \dots, \text{pk}_{|\mathcal{S}|}, \{\text{pk}_\ell\}_{\ell \in \mathcal{T}_K})$.
2. Initialize an empty set $\vec{\tau}$.
3. For each $i \in \mathcal{S}$:
 - (a) prepare the following circuit C_i using black-box access to C .

- i. C_i uses $(\{xk_j\}_{j \in \mathcal{S}, j \neq i}, \{td_\ell\}_{\ell \in \mathcal{T}_K})$ and external queries to simulate security game $G_P^2(\text{sk}, \text{pk}, \text{aux}, C)$ for C ⁸:
 - For any (admissible) queries in the stage-2 game $G_P^2(\text{sk}, \text{pk}, \text{aux}, C)$, if the oracles provided in security game $G_{P_i}(\text{sk}_i, \text{pk}_i, \text{aux}_i, \cdot)$ is required to compute the answer for the query, C_i can will make a query to an external challenger. To answer the entire query, C_i may finish the rest of the computation using $(\{xk_j\}_{j \neq i}, \{td_\ell\}_{\ell \in \mathcal{T}_K})$.
 - If only $\{xk_j\}_{j \neq i}$ are needed to answer a query, C_i answers it by itself.
 - ii. C_i records queries from C into a set \mathcal{Q}_C .
 - iii. C_i receives its challenge input inp_i from the interaction with an external challenger, samples a random string r , and prepares a challenge input inp for C using inp_i and r .
 - iv. If there is a query phase after the challenge phase, C_i continues to simulate the oracles required using $\{xk_j\}_{j \neq i}$ and external queries to a challenger.
 - v. After C makes its final output out , C_i computes the function $f_i(\text{out}, r, \mathcal{Q}_C, \text{inp})$ where f_i is the reduction function (Remark 3.14) used in watermarking-compatible reduction of P to P_i . Output the result of this computation.
- (b) compute $\tau_i / \perp \leftarrow \text{WP}_i.\text{Extract}(xk_i, C_i)$.
- (c) add τ_i (or \perp) to the tuple $\vec{\tau}$, and go to step 2 with $i := i + 1$.

4. Output

- $\vec{\tau} = (\tau_1, \dots, \tau_{|\mathcal{S}|})$ if $\exists i, j$ where $\tau_i \neq \tau_j$;
- else if $\tau_i = \tau_j = \tau$ (or \perp) for all i, j , output τ (or \perp resp.).

Remark 4.13. We also give a more systematic description of *Extract* is using the language of watermarking-compatible reduction:

The *Extract* algorithm:

1. On input circuit C and $xk, \text{pk}, \text{aux}$; for $i \in \mathcal{S}$:

- (a) Create program C_i (C_i has black-box access to C). C_i treats C as the stage-2 adversary \mathcal{A}^2 in the stage-2 game $G_P^2(\text{sk}, \text{pk}, \text{aux}, \mathcal{A}^2)$.
- (b) C acts as a stage-2 reduction \mathcal{A}_i^2 in the watermarking-compatible reduction from primitive P to P_i , with the difference that C uses $\{xk_j\}_{j \neq i}$ (instead of $\{\text{sk}_j\}_{j \neq i}$) and external queries to challenger to simulate the game.
- (c) compute $\tau / \perp \leftarrow \text{WP}_i.\text{Extract}(xk_i, C_i)$; add it to the extracted mark tuple.

2. Output the tuple of marks extracted.

Remark 4.14. As discussed previously, for the convenience of discussions in some later parts, our *Mark* algorithm outputs a marked key instead of circuits since the *wSecEval* algorithm itself is public.

Remark 4.15 (Outputting a tuple of marks). The *Extract* algorithm will only output a tuple of marks (with $|\mathcal{S}|$ number of them at most) when used on adversarial circuits C . For honestly generated circuit C , the extraction will always output only one marked message.

⁸ C_i can also have xk_i if the watermarkable implementation WP_i has a public extraction key. In this case, C_i also does not need to make external queries.

First this is clearly not a limitation because $|\mathcal{S}|$ is the number of input primitives and usually a small constant/polynomial. Second, the need for the extraction algorithm to output a tuple of marks is essential for the composability of watermarking schemes, otherwise we cannot capture the following trivial "attack": the adversary receives an honestly marked output primitive key, which consists of k marked secret keys of the input primitives. It can simply destroy all the keys except one and this leftover marked key suffices for breaking the security of P . Such an "attack" should not be considered as a successful attack because we can still extract a valid watermark from the entire program. Therefore, we need the extraction to output a tuple of marks so that we can check if the adversary has succeeded at removing/replacing marks from each underlying functionality.

Remark 4.16. As previously discussed in Remark 4.11, when we have a primitive in the set \mathcal{T}_K , the Extract algorithm needs to use the trapdoor(s) td of the primitives in set \mathcal{T}_K to simulate a hybrid version of the security game G_P . This simulation is supposedly computationally indistinguishable from the original game, by the security of the primitive(s) in \mathcal{T}_K .

In more detail, \mathcal{T}_k only contains primitives P_ℓ where

1. There exists an efficient simulator algorithm such that the output of $P_\ell.\text{PubEval}(P_\ell.\text{pk}, \cdot)$ is indistinguishable from the output of the simulator $\text{Sim}(P_\ell.\text{td}, \cdot)$. Their secret key (trapdoor td) will only come up in the security proof for P .
2. In the security proof for P , the rest of the hybrid arguments and reduction happen in a hybrid game where we have already invoked the above security for P_ℓ .

Due to the above reasons, the Extract algorithm has to simulate hybrid version of game G_P where the security of $P_\ell \in \mathcal{T}_K$ is already invoked, in order for the reduction (and thus extraction) for the rest of the primitives to go through.

Essentially, we simply need the plain security of primitives in the set \mathcal{T}_K to realize the watermarking security of P , like other primitives in $\bar{\mathcal{S}}$.

Such examples are rare. The only example in this work is a NIZK proof system with a trapdoor used to simulate proofs (Appendix A).

We now state our main theorem:

Theorem 4.17 (Watermarking Composition). *Suppose the input primitives in the set \mathcal{S} , $\{P_i\}_{i \in \mathcal{S}}$ have watermarkable implementations $\{\text{WP}_i\}_{i \in \mathcal{S}}$ which satisfy the definitions in 4.1, the primitives outside the set \mathcal{S} have secure constructions (defined by their corresponding correctness and security in Section 3.1), and the construction of P is watermarking-compatible by the definitions in 3.2, then the output watermarkable implementation WP for primitive P will satisfy the properties in Definition 4.1.*

To prove Theorem 4.17, we show that all the properties from Section 4.1 hold.

Functionality-Preserving

Lemma 4.18 (Functionality-Preserving and Exact Functionality-Preserving). *Suppose the above construction of P from P_1, \dots, P_k satisfies correctness of construction, and constructions $\text{WP}_1, \dots, \text{WP}_{|\mathcal{S}|}$ satisfy the functionality-preserving properties in Definition 4.1, then WP satisfies correctness defined by predicate F_R .*

If all $\text{WP}_1, \dots, \text{WP}_{|\mathcal{S}|}$ satisfy the exact functionality-preserving properties, then WP satisfies exact functionality-preserving.

Proof. Since WP_1, \dots, WP_k satisfy the functionality-preserving properties in Definition 4.1, the correctness properties of P_1, \dots, P_k are preserved by their watermarkable implementations. Since P 's correctness follows from the correctness properties of P_1, \dots, P_k and moreover, WP 's evaluation algorithms are constructed the same way as they are in P (even though the $WMS\text{Setup}(1^\lambda)$ algorithm differs from $\text{KeyGen}(1^\lambda)$, we can ignore the watermarking-related components xk, mk in this context). The constructions for primitives $P_i, i \notin \mathcal{S}$ also do not affect functionality preserving since we use their plain, unwatermarked keys during evaluations. WP 's correctness property thereby follows.

It is easy to see that if all WP_i satisfies exact functionality preserving, then WP satisfies exact functionality preserving. \square

Correctness of Extraction Our correctness property is a computational correctness that relies on the unremovability of the input watermarkable implementations. Since our Extract algorithm interacts with any input circuit black-boxly by simulating the security game G_P , the only way to extract a watermark is through the circuit's input-output behavior in this game. We use the guarantee that "if the circuit can win the game, then we can extract a watermark" to show correctness on a circuit embedded with an honestly watermarked key (which can win the game with probability 1 naturally). Also note that we do not need our honestly watermarked circuit to be stateful/interactive as in Remark 4.10.

Lemma 4.19 (Correctness of Extraction). *Suppose the above construction of P from P_1, \dots, P_k has a watermarking-compatible reduction, and all watermarkable implementation of input primitives in set \mathcal{S} , $WP_1, \dots, WP_{|\mathcal{S}|}$ satisfy strong unremovability and the extraction syntax, then the watermarkable implementation of target primitive WP satisfies correctness of extraction.*

Proof. Given an honestly marked key $sk_\tau = (sk_{1,\tau}, \dots, sk_{|\mathcal{S}|,\tau}) \leftarrow WP.\text{Mark}(mk, sk)$: we can observe that there exists a public algorithm T that once given sk_τ , can win the game $G_P(sk, pk, \cdot)$ with probability $(1 - \text{negl}(\lambda))$. We can sample a random $i \leftarrow \mathcal{S}$ and design a circuit T to work as follows: T on any challenge input, uses a strategy associated with the key sk_i to compute the challenge given in game G_P^2 .

Take the concrete example of the watermarkable implementation WP for a CCA2 encryption scheme, used in our technical overview section (Section 1.4). The watermarked key is a concatenation of a watermarked decryption key/weak PRF key $sk_{1,\tau}$ and a watermarked MAC signing key $sk_{2,\tau}$. Then the algorithm T can choose the weak PRF key as its functionality: on any challenge ciphertext, T simply decrypts the ciphertext using the watermarked key $sk_{1,\tau}$ and outputs the correct answer.

We then look into what happens when we compute $\text{Extract}(xk, C)$: C is treated as a stage-2 adversary in the watermarking-compatible reduction from the security of P to the security of P_i and C_i with black box access to C is a stage-2 reduction algorithm. C , for each $i \in \mathcal{S}$. We input $T(sk_\tau, \cdot)$ into the extraction algorithm. By the properties of watermarking-compatible reduction, such a $T(sk_\tau, \cdot)$ will be used to build C_i that is γ_i -good for some noticeable γ_i for each $i \in \mathcal{S}$. When it comes to the index i that we choose to be our functionality in T , then the created circuit C_i using T will naturally help win the security game $G_{P_i}^2$.

If we view the honest user as an adversary that makes one marking query on some message τ , then by the unremovability security of WP_i , we must have $\Pr[WP_i.\text{Extract}(xk_i, pk_i, aux, C_i = \tau) \geq 1 - \text{negl}(\lambda)] \geq 1 - \text{negl}(\lambda)$ for all i . Therefore $\Pr[WP.\text{Extract}(xk, pk, aux, (sk_\tau, \cdot)) = \tau] \geq 1 - \text{negl}(\lambda)$. Since extracting one valid watermark suffices, we can conclude the correctness of extraction.

Going back to the CCA2 encryption example: the circuit T can answer any challenge by decrypting using the PRF key; therefore, when it is used by the circuit C_1 created in Extract algorithm, its output helps C_1 distinguish between pseudorandom input and real random input. By the unremovability security of the watermarkable weak PRF, we must be able to extract the queried watermark τ from C_1 . □

Security with G_P We show that the (plain) security of P with game G_P holds.

Lemma 4.20 (Security with G_P). *Suppose the above construction of P from P_1, \dots, P_k has a provable black-box security, and watermarkable implementations $WP_1, \dots, WP_{|S|}$ as well as the (unwatermarked) constructions for $\{P_i\}_{i \notin S}$ satisfy the security with the security defined by game $G_{P_i}, i \in [k]$ respectively, then WP satisfies security defined by game G_P .*

Proof. Since WP_1, \dots, WP_k satisfy the security property with game G_{P_i} in Definition 4.1, the security properties of P_1, \dots, P_k are preserved by their watermarkable implementations. P 's security with respect to G_P follows from the security properties of P_1, \dots, P_k with respect to $\{G_{P_i}\}_{i \in [k]}$. Moreover, WP 's evaluation algorithms are constructed the same way as they are in P and even though the $WMSetup(1^\lambda)$ algorithm differs from $KeyGen(1^\lambda)$, we can ignore the watermarking-related components xk, mk in this context. Therefore, WP 's security property follows from the security of WP_1, \dots, WP_k as P 's security property follows from the security of P_1, \dots, P_k . □

Unremovability We now show that the construction satisfies unremovability if all building blocks satisfy unremovability (or plain security, for unwatermarked building blocks) and the properties on extraction key/algorithm defined in section 4.1.

Lemma 4.21 (Unremovability). *Suppose the above construction of P from P_1, \dots, P_k has a watermarking-compatible reduction, the watermarkable implementations $WP_1, \dots, WP_{|S|}$ satisfy strong unremovability (Definition 4.7), the constructions for $\{P_j\}_{j \notin S}$ satisfy security defined by their game $\{G_{P_j}\}_{j \notin S}$ respectively, and the extraction algorithms of $WP_1, \dots, WP_{|S|}$ satisfy the extraction syntax in Section 4.1, then the watermarkable implementation WP satisfies strong unremovability.*

Proof. If there exists an adversary that breaks γ -unremovability with game G_P as defined in Definition 4.4, then we have $\Pr[\text{Extract}(xk, C) \notin Q \wedge C \text{ is } \gamma\text{-good}] \geq \epsilon$ for some non-negligible ϵ where C is γ -good in the sense that $\Pr[G_P^2(\text{sk}, \text{pk}, \text{aux}, C) = 1] \geq \eta + \gamma$. Here the overall winning probability of \mathcal{A} is taken over the randomness used in $KeyGen$ and the unremovability game; the probability for $G_P^2(\text{sk}, \text{pk}, \text{aux}, C) = 1$ is taken over the randomness used in the stage-2 game G_P^2 .

By the design of our Extract algorithm, it must be the case that, for all $i \in S, WP_i. \Pr[\text{Extract}(xk_i, C_i) \notin Q] \geq \epsilon$, where C_i is the circuit made black-boxly from C in WP .Extract algorithm for index i .

Now we divide our analysis into cases. First, we consider that there is no input primitive in the set \mathcal{T}_k , i.e. no trapdoor td used in the extraction procedure and the extraction key can perfectly simulate the security game for G_P .

For any \mathcal{A} producing a circuit C such that $\Pr[\text{Extract}(xk, \text{pk}, \text{aux}, C) \notin Q \wedge C \text{ is } \gamma\text{-good}] \geq \epsilon$ (where the probability is taken over the randomness in $KeyGen$ and randomness used throughout the game), one of the following cases must hold:

- **Case 1: For some non-negligible probability ϵ_i , during the execution of $\text{Extract}(\text{xk}, \text{pk}, \text{aux}, C)$, there exists some $i^* \in \mathcal{S}$ such that C_{i^*} is a γ_{i^*} -good program in $G_{P_i}^2(\text{sk}_i, \text{pk}_i, \text{aux}_i, C_i)$ for some non-negligible γ_{i^*} .**

Then, there exists a reduction to break the γ_{i^*} -unremovability of WP_{i^*} for some non-negligible γ_{i^*} , using \mathcal{A} . Let \mathcal{A}_{i^*} be the reduction and adversary in γ_{i^*} -unremovability game of WP_{i^*} .

- \mathcal{A}_{i^*} receives pk_{i^*} from the challenger and samples $\{\text{sk}_i, \text{pk}_i, \text{xk}_i, \text{mk}_i\}_{i \in \mathcal{S}, i \neq i^*}$ from $\text{WP}_{i^*}.\text{WMSetup}$. It gives $\{\text{pk}_i\}_{i \in \mathcal{S}, i \neq i^*}$ to \mathcal{A} .
- When \mathcal{A} makes a marking query on some symbol τ , \mathcal{A}_{i^*} responds as follows:
 - * compute $\text{sk}_{i,\tau} \leftarrow \text{Mark}(\text{mk}_i, \text{sk}_i, \tau)$ for $i \in \mathcal{S}, i \neq i^*$.
 - * query the WP_{i^*} challenger on oracle $\text{Mark}(\text{mk}_{i^*}, \text{sk}_{i^*}, \tau)$ to obtain $\text{sk}_{i^*,\tau}$.
 - * Let $\text{sk}_\tau = \{\text{sk}_{i,\tau}\}_{i \in \mathcal{S}}$; send sk_τ to \mathcal{A} .
- Next, \mathcal{A} outputs circuit polynomial-size C .
- \mathcal{A}_{i^*} creates the following circuit C'_{i^*} using C as a black-box:
 - * C'_{i^*} is hardcoded with $\{\text{sk}_i\}_{i \in \mathcal{S}, i \neq i^*}, \{\text{pk}_i\}_{i \in \mathcal{S}}$.
 - * When C makes a query to oracles in game G_P , C'_{i^*} responds as follows: if sk_{i^*} is required to answer the query, C'_{i^*} queries the external challenger on the oracles provided in game $G_{P_{i^*}}$; otherwise, C'_{i^*} answers the query using $\{\text{sk}_i\}_{i \in \mathcal{S}, i \neq i^*}$.
 - * C'_{i^*} records all queries from C , denoted as \mathcal{Q}_C .
 - * In the challenge phase, C'_{i^*} receives challenge input inp_i from the interaction with an external challenger, samples a random string r , and prepares a challenge input inp for C using inp_i and r .
 - * When C outputs its final answer out , C'_{i^*} computes the function $f_{i^*}(\text{out}, \mathcal{Q}_C, \text{inp})$ according to the reduction function f_{i^*} in the watermarking-compatible reduction from P to P_{i^*} .
- In short, C is treated as a stage adversary \mathcal{A}_2 in security game G_P ; circuit C_{i^*} is a stage-2 reduction algorithm in the watermarking-compatible reduction from target primitive P to input primitive P_{i^*} , with black box access to C .

By the design of the WP.Extract algorithm, it using extraction keys $\{\text{xk}_i\}_{i \in \mathcal{S}}$, will treat the input circuit as an adversary in stage-2 game G_P^2 and simulate G_P^2 for it.

Given any adversarial circuit C that is a γ -good adversary in the security game G_P , following our extraction algorithm syntax requirement, the extraction algorithm running at step $i = i^*$ will create a circuit C_{i^*} that has exactly the same functionality as the circuit C'_{i^*} created in the above reduction⁹. Therefore, if C_{i^*} created by Extract is a γ_{i^*} -good in breaking $G_{P_{i^*}}$, then so is C'_{i^*} created by the above procedure; since we have $\text{WP}_{i^*}.\text{Pr}[\text{Extract}(\text{xk}_{i^*}, C_{i^*}) \notin Q] \geq \epsilon$ and the set Q is the same for \mathcal{A} and \mathcal{A}_{i^*} , then we also have $\text{WP}_i.\text{Pr}[\text{Extract}(\text{xk}_i, C'_{i^*}) \notin Q] \geq \epsilon$.

- **Case 2: With overwhelming probability, during the execution of $\text{Extract}(\text{xk}, \text{pk}, \text{aux}, C)$, there exists no $i \in \mathcal{S}$ such that C_i is a γ_i -good adversary in $G_{P_i}^2(\text{sk}_i, \text{pk}_i, \text{aux}_i, C_i)$ for some non-negligible γ_i .**

In this case, it must be that there exists some $i^* \notin \mathcal{S}$ such that we can use the unremovability adversary \mathcal{A} for WP to break the *security* of P_{i^*} , i.e. winning game $G_{P_i}, i \notin \mathcal{S}$ with some non-negligible advantage.

⁹Note that the external queries made by C_{i^*} will be answered by $\text{WP}_{i^*}.\text{Extract}(\text{xk}_{i^*}, C_{i^*})$ during its execution, where by our requirement on the extraction syntax, $\text{WP}_{i^*}.\text{Extract}(\text{xk}_{i^*}, C_{i^*})$ will simulate the game $G_{P_{i^*}}$ for C_{i^*} , answering its queries.

By the properties of watermarking-compatible reduction, if (with non-negligible probability), the circuit C produced by \mathcal{A} is a γ -good adversary in G_P , then there must exist some $i \in [k]$ such that using the watermarking-compatible reduction, one can create a γ_i -good adversary for some non-negligible γ_i . By the above analysis, C_i supposedly performs exactly the reduction for P_i and if there is no good C_i for $i \in \mathcal{S}$, then C can only be used to build a reduction to break the security of some $P_i, i \notin \mathcal{S}$.

The reduction is as follows. Let \mathcal{A}_{i^*} be the reduction and adversary in the security game $G_{P_{i^*}}$.

- \mathcal{A}_{i^*} receives pk_{i^*} from the challenger and samples $\{\text{sk}_i, \text{pk}_i, \text{xk}_i, \text{mk}_i\}_{i \in \mathcal{S}}$ from $\text{WP}_{i^*}.\text{WMSSetup}$. It gives $\{\text{pk}_i\}_{i \in \mathcal{S}}$ to \mathcal{A} .
 - When \mathcal{A} makes a marking query on some symbol τ , \mathcal{A}_{i^*} answers the query on its own because it has all the $\{\text{sk}_i, \text{mk}_i\}_{i \in \mathcal{S}}$.
 - Besides answering marking queries, \mathcal{A}_{i^*} will treat \mathcal{A} as a stage-1 adversary in the watermarking-compatible reduction Type 2 (Section 3.2): when \mathcal{A} makes a query to oracles in game G_P , \mathcal{A}_{i^*} responds as follows: if sk_{i^*} is required to answer the query, \mathcal{A}_{i^*} queries the challenger in $G_{P_{i^*}}$ security game; otherwise, \mathcal{A}_{i^*} answers the query using $\{\text{sk}_i\}_{i \in \mathcal{S}}$.
- For the secret keys of $\{\text{WP}_i\}_{i \notin \mathcal{S}, i \neq i^*}$, they are all sampled freshly upon every query of $\text{wSecEval}(\text{sk}, \text{pk}, \cdot)$ (or sampled freshly in $\text{wPubEval}(\text{pk}, \cdot)$ which \mathcal{A} can do it on its own).
- Next, \mathcal{A} outputs circuit polynomial-size C .
 - \mathcal{A}_{i^*} creates the following circuit C^* using C as a black-box:
 - * C^* is hardcoded with $\{\text{sk}_i, \text{pk}_i\}_{i \in \mathcal{S}}$.
 - * When C makes a query to oracles in game G_P , C^* responds as follows: if sk_{i^*} is required to answer the query, C^* queries the external challenger on the oracles provided in game $G_{P_{i^*}}$; otherwise, C^* answers the query using $\{\text{sk}_i\}_{i \in \mathcal{S}}$.
 - * C^* records all queries from C , denoted as \mathcal{Q}_C .
 - * In the challenge phase, C^* receives challenge input inp_i from the interaction with an external challenger, samples a random string r , and prepares a challenge input inp for C using inp_i and r .
 - * When C outputs its final answer out , C^* computes the function $f_{i^*}(\text{out}, \mathcal{Q}_C, \text{inp})$ according to the reduction function f_{i^*} in the watermarking-compatible reduction from P to P_{i^*} .
 - In short, C is treated as a stage-2 adversary \mathcal{A}_2 in security game G_P ; circuit C^* is a stage-2 reduction algorithm in the watermarking-compatible reduction Type 2 from target primitive P to input primitive P_{i^*} , with black box access to C . By the property of the reduction, C^* should win the security game $G_{P_{i^*}}$ with non-negligible advantage.

Finally, we discuss the rare case when the extraction key simulates the security game for C statistically/computationally close to a real game.

In the setting where $(\{\text{xk}_i\}_{i \in \mathcal{S}}, \{\text{td}_\ell\}_{\ell \in \mathcal{T}_K})$ can only simulate G_P^2 for C statistically close by some negligible distance, then if C is a γ -good circuit in a real $G_P^2(\text{sk}, \text{pk}, \text{aux}, \cdot)$ game, then C is a $(\gamma - \text{negl}(\lambda))$ -good circuit in the game simulated by the Extract algorithm and thus by the reduction. Afterwards, we can go back to use C in the above Case 1 and Case 2.

Coming to the computational setting, we will see that they will also be transformed into Case 1 or Case 2.

Note that within the scope of this work, this case corresponds precisely to having primitives

in the set \mathcal{T}_k and the computational property comes from primitives in the set \mathcal{T}_k . Recall that \mathcal{T}_k only contains the primitives P_ℓ where there exists a simulator algorithm such that the output of $P_\ell.\text{PubEval}(P_\ell.\text{pk}, \cdot)$ is indistinguishable from the output of the simulator $\text{Sim}(P_\ell.\text{td}, \cdot)$ and in the security proof for P , the rest of the hybrid arguments and reduction happen in a hybrid world where we have already invoked the security for P_ℓ . (see discussions in Remark 3.15, Definition 4.9 and Remark 4.11, Remark 4.16).

We provide one such example where we need to use a NIZK proof and provide simulated proofs to the adversarial circuits in the extraction algorithm. See the example in Appendix A for details.

First consider the case where we only have one primitive P_ℓ in set \mathcal{T}_K in our construction (this is the case in our example Appendix A). By our premise, using $(\{xk_i\}_{i \in S}, \text{td}_\ell)$ where $\text{td}_\ell \leftarrow P_\ell.\text{KeyGen}(1^\lambda)$, one can simulate G_P^2 for C computationally close by some negligible amount. If with some non-negligible probability ϵ , we have C is a γ -good circuit in a real $G_P^2(\text{sk}, \text{pk}, \text{aux}, \cdot)$ game. Then with $(\epsilon - \text{negl}(\lambda))$ probability, C is a $(\gamma - \text{negl}(\lambda))$ -good circuit in the game simulated by the Extract algorithm and the reductions, otherwise there exists a PPT distinguisher that can use C to distinguish between the real game and simulated one to break the security of a primitive in set \mathcal{T}_k : a reduction \mathcal{B} can samples all $\text{sk}, \text{pk}, xk, \text{mk}$ on its own and receive the public parameters from the challenger of some G_{P_ℓ} . The rest of simulation is the same as in Case 2 of the above analysis. At the end, \mathcal{B} tests if the circuit output by \mathcal{A} is a γ -good one, if yes, output guess "real view; else output guess "simulated view".

If there are more than one $P_\ell \in \mathcal{T}_K$, we can then use a hybrid to apply the above to each of them, following the hybrid order in the security proof for the plain security of P .

Therefore, by the security of P_ℓ , we have that C is a $(\gamma - \text{negl}(\lambda))$ -good circuit in the game simulated by the Extract algorithm. Since the drop in C 's advantage is negligible, we can go back to Case 1 and Case 2. □

Additional Properties Finally, we discuss some additional properties realized by the watermarking composition.

Private vs. Public Extraction Firstly, we have a relatively straightforward observation that if all underlying constructions satisfy public extraction, then the target construction also does.

Lemma 4.22 (Public Extraction). *Suppose a watermakable implementation WP is constructed from watermarkable implementations $\text{WP}_1, \dots, \text{WP}_{|S|}$, if $\text{WP}_1, \dots, \text{WP}_{|S|}$ all have public extraction and there exists no primitive in the set \mathcal{T}_K , then WP has public extraction.*

Proof. It is easy to observe that if all WP_i have public extraction, then all xk_i can be made public and the extraction algorithm of WP can be made public. During the extraction procedure, since the keys in the primitives in the set $\mathcal{T}_S, \mathcal{T}_P$ are sampled online, even a public procedure can sample them. But if we have a primitive in the set \mathcal{T}_K , then we need a trapdoor during extraction, so we have to rule out this case. □

Meaningfulness The meaningfulness of WP follows from the meaningfulness of each WP_i construction naturally.

Collusion Resistance Finally, we show that collusion resistance also composes. Informally, suppose a watermarkable implementation WP is constructed from watermarkable implementations $WP_1, \dots, WP_{|\mathcal{S}|}$, if all of $WP_1, \dots, WP_{|\mathcal{S}|}$ are collusion resistant, then WP is collusion resistant. If at least one of them is q -bounded collusion resistant, then WP is q -collusion resistant q (by the smallest $q_i, i \in |\mathcal{S}|$).

Lemma 4.23 (Collusion Resistance). *Suppose the above construction of P from P_1, \dots, P_k has a watermarking-compatible reduction, the watermarkable implementations $WP_1, \dots, WP_{|\mathcal{S}|}$ all satisfy strong unremovability (Definition 4.7) with q -collusion-resistance, the constructions for $\{P_j\}_{j \notin \mathcal{S}}$ satisfy security defined by their game $\{G_{P_j}\}_{j \notin \mathcal{S}}$ respectively, and the extraction algorithms of $WP_1, \dots, WP_{|\mathcal{S}|}$ satisfy the extraction syntax in Section 4.1, then the watermarkable implementation WP satisfies strong unremovability with q -collusion resistance.*

Proof. Since each marking query of WP requires querying each underlying marking oracle WP_i . Mark once on the same marking message, the number of queries allowed for the reduction from WP's unremovability is bounded by the input primitive WP_i 's query bound. If all underlying schemes are q -collusion resistant for some q , then all reductions can query q times and the unremovability proof for Lemma 4.21 stays the same. WP is also q -collusion resistant. \square

Extraction Simulation Property and Extraction Syntax Having this additional property in the composed watermarkable implementation scheme allows us to further compose watermarking schemes from input schemes that come from composition themselves.

It is easy to see that the ability of the extraction algorithm to simulate the security game also composes.

We can observe that if all watermarkable implementations of input primitives satisfy the extraction syntax (Section 4.1, and thus key simulation property Definition 4.9), then the combined keys ($\{xk_i\}_{i \in \mathcal{S}}, \{pk_i\}_{i \in \mathcal{S}}$) can be used to simulate the game G_P because all oracles used in G_P can all be built from oracles used in $\{G_{P_i}\}_{i \in \mathcal{S}}$ according to the watermarking-compatible reduction.

Since all the WP_i .Extract(xk_i, pk_i, aux_i, \cdot) algorithm perfectly (statistically/computationally, resp.) simulates the security game for G_{P_i} stage 2 (as a subroutine of extraction), when each C_i created from input program C is executed inside the procedure WP_i .Extract(xk_i, pk_i, aux_i, C_i): C_i will simulate the game G_P stage 2 for C using external oracle queries when interacting with the algorithm WP_i .Extract(xk_i, pk_i, aux_i, C_i) and the other extraction keys $\{xk_i\}_{i \in \mathcal{S}}$. The entire Extract algorithm consists of subroutines where one simulates G_P perfectly (statistically/computationally, resp.) stage 2 for C when running C .

When we have input primitives in \mathcal{T}_k , as we have discussed in several remarks: in the extract trapdoor(s) of primitives in \mathcal{T}_k to simulate part of the security game, which results in a computationally close simulation from the original game. One concrete example of such primitive in \mathcal{T}_k is NIZK, which we will use in Appendix A and Appendix G.

5 Two Simple Examples: Watermarkable CPA and CCA-secure SKE

In this section, we give two relatively simple examples through the following two steps:

- First we give their "plain" construction and show that their reduction is watermarking-compatible.

- Then give their watermarking implementation of the target primitive in the watermarking-compatible reduction language.

For the remaining examples we show in this work, we will integrate the above two discussions for brevity.

5.1 Watermarkable CPA-secure SKE from Watermarkable weak PRF

In this section, we show a simple example of obtaining a watermarkable implementation of a CPA secure secret-key encryption scheme from a watermarkable implementation of weak PRF. The security of watermarkable weak PRF satisfying our requirements in Section 4.1 can be built from LWE (with private extraction) or iO (with public extraction) [GKWW21]. The following example was also discussed in [GKWW21].

5.2 Definition: Watermarkable Implementation of PRF

We give the functionality-preserving unremovability security definition for PRF and CPA-secure SKE. The other definitions (pseudorandomness, CPA-security, correctness, functionality preserving) follow naturally from the correctness and security definitions of PRF and SKE respectively.

A watermarkable implementation of a PRF scheme consists of the following algorithms (WMS_{Setup}, Eval, Mark, and satisfies the following properties:

WMS_{Setup}(1^λ) \rightarrow (sk, mk, xk): on input security parameter, outputs a secret key sk, marking key mk, extraction key xk.

Eval(sk, $x \in \{0, 1\}^\ell$) $\rightarrow y \in \{0, 1\}^\ell$: a deterministic evaluation algorithm that on input sk, x , outputs y .

Mark(mk, sk, τ): on marking key mk, secret key sk, a message τ , output a marked key sk_τ .

Extract(xk, C): on extraction key sk and program C , output a mark $\tau \in \mathcal{M}_\tau$ or \perp .

Functionality-Preserving A watermarkable PRF is functionality preserving if then there exists a negligible function $\text{negl}(\lambda)$ such that for all $\lambda \in \mathbb{N}$, $x \in \mathcal{X}$, $\tau \in \mathcal{M}_\tau$:

$$\Pr \left[\text{Eval}(\tilde{\text{sk}}, x) = \text{Eval}(\text{sk}, x) : \begin{array}{l} (\text{sk}, \text{pk}, \text{xk}, \text{mk}) \leftarrow \text{WMS}_{\text{Setup}}(1^\lambda) \\ \text{sk} \leftarrow \text{Mark}(\text{mk}, \text{sk}, \tau) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

Weak Pseudorandomness A watermarkable PRF satisfies weak pseudorandomness if for all PPT \mathcal{A} , there exists a negligible function $\text{negl}(\lambda)$ such that for all $\lambda \in \mathbb{N}$:

$$\Pr \left[\mathcal{A}^{\mathcal{O}(\text{sk}, \cdot)}(x, y_b) = b : x \leftarrow \{0, 1\}^\ell, y_0 = \text{Eval}(\text{sk}, x), y_1 \leftarrow \{0, 1\}^\ell, b \leftarrow \{0, 1\} \right] \leq \frac{1}{2} + \text{negl}(\lambda).$$

where the oracle $\mathcal{O}(\text{sk}, \cdot)$ samples a uniformly random $x \leftarrow \{0, 1\}^\ell$ upon every query and outputs $(x, \text{Eval}(\text{sk}, \cdot))$.

We can also consider the equivalent notion where \mathcal{A} is given an oracle $\mathcal{O}(\text{sk}, \cdot)$ which either computes PRF.Eval(sk, \cdot) on random inputs or computes a real random function that samples random input together with a random output. There will be no challenge input and \mathcal{A} outputs a bit indicating which oracle it is given. Similarly for the security below.

Variation of Weak Pseudorandomness Game To be better compatible with our unremovability definition, we consider the following variation of the above weak PRF game, described in a 2-stage fashion: in stage 1, \mathcal{A}^1 is allowed to query *adaptively* $\mathcal{O}(\text{sk}, \cdot)$ on any input of its own choice. Entering stage 2, it can only query $\mathcal{O}(\text{sk}, \cdot)$ in a way that each input is sampled at random. Then \mathcal{A}^2 is fed with challenge input (x, y_b) where $x \leftarrow \{0, 1\}^\ell, y_0 = \text{Eval}(\text{sk}, x), y_1 \leftarrow \{0, 1\}^\ell, b \leftarrow \{0, 1\}$ and needs to output the correct b .

γ -Unremovability of weak PRF The γ -Unremovability for a watermarkable implementation of a weak PRF scheme says, for all PPT admissible stateful adversary \mathcal{A} , there exists a negligible function $\text{negl}(\lambda)$ such that for all $\lambda \in \mathbb{N}$:

$$\Pr \left[\text{Extract}(\text{xk}, C) \notin \mathcal{Q} \wedge C \text{ is } \gamma\text{-good} : \begin{array}{l} (\text{sk}, \text{xk}, \text{mk}) \leftarrow \text{WMSetup}(1^\lambda) \\ C \leftarrow \mathcal{A}^{\text{Mark}(\text{mk}, \text{sk}, \cdot)}(1^\lambda) \end{array} \right] \leq \text{negl}(\lambda).$$

$\mathcal{O}(\text{sk}, \cdot)$ samples a uniformly random $x \leftarrow \{0, 1\}^\ell$ upon every query and outputs $(x, \text{Eval}(\text{sk}, \cdot))$.

\mathcal{Q} is the set of marks queried by \mathcal{A} and C is a PPT admissible, stateful γ -good adversary in the security game $G_{\text{PRF}}(\text{sk}, \cdot)$ if:

$$\Pr \left[C^{\mathcal{O}(\text{sk}, \cdot)}(x, y_b) = b : x \leftarrow \{0, 1\}^\ell, y_0 = \text{Eval}(\text{sk}, x), y_1 \leftarrow \{0, 1\}^\ell, b \leftarrow \{0, 1\} \right] \geq \frac{1}{2} + \gamma.$$

Remark 5.1. In the above setting, the oracle $\mathcal{O}(\cdot)$ will sample a random r and output $r, \mathcal{O}(r)$.

We can consider an even slightly stronger notion of weak PRF: in the "fully adaptive-query" setting, \mathcal{A}, C are allowed to make any query to the real Eval oracle; only the challenge input x is sampled uniformly at random. Then C is asked to distinguish between $y_0 = \text{Eval}(\text{sk}, x), y_1 \leftarrow \{0, 1\}^\ell$.

Note that [GKWW21] consider the first notion so that they can realize public extraction only by letting the extraction key sample input-output pairs. This notion also suffices to prove watermarkable CPA security. Both can be constructed in [GKWW21].

5.3 Definition: Watermarkable Implementation of CPA-secure SKE

A watermarkable implementation of a SKE scheme consists of the following algorithms:

$\text{WMSetup}(1^\lambda) \rightarrow (\text{sk}, \text{mk}, \text{xk})$: on input security parameter, outputs a secret key sk , marking key mk , extraction key xk .

$\text{Enc}(\text{sk}, m) \rightarrow \text{ct}$: on input secret key sk and message m , outputs a ciphertext ct .

$\text{Dec}(\text{sk}, \text{ct}) \rightarrow m$: on input secret key sk and ciphertext ct outputs a message m .

$\text{Mark}(\text{mk}, \text{sk}, \tau)$: on marking key mk , secret key sk , a message τ , output a marked key sk_τ .

$\text{Extract}(\text{xk}, C)$: on extraction key xk and program C , output a mark $\tau \in \mathcal{M}_\tau$ or \perp .

The correctness property is the same as the correctness property of a plain SKE scheme and we combine it with the functionality preserving property it due to it being standard.

Correctness and Functionality Preserving A watermarkable CPA-secure SKE is functionality preserving if then there exists a negligible function $\text{negl}(\lambda)$ such that for all $\lambda \in \mathbb{N}, m \in \mathcal{M}, \tau \in \mathcal{M}_\tau$:

$$\Pr \left[\text{Dec}(\tilde{\text{sk}}, \text{ct}') = m \wedge \text{Dec}(\text{sk}, \text{ct}) = m : \begin{array}{l} (\text{sk}, \text{pk}, \text{xk}, \text{mk}) \leftarrow \text{WMSetup}(1^\lambda) \\ \tilde{\text{sk}} \leftarrow \text{Mark}(\text{mk}, \text{sk}, \tau) \\ \text{ct} \leftarrow \text{Enc}(\text{sk}, m), \text{ct}' \leftarrow \text{Enc}(\text{sk}_\tau, m) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

CPA-security A watermarkable CPA secure SKE scheme satisfies CPA security if for all PPT admissible stateful adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, there exists a negligible function $\text{negl}(\lambda)$ such that for all $\lambda \in \mathbb{N}$:

$$\Pr \left[\mathcal{A}_2^{\text{Enc}(\text{sk}, \cdot)}(\text{ct}_b) = b : \begin{array}{l} (m_0, m_1) \leftarrow \mathcal{A}_1^{\text{Enc}(\text{sk}, \cdot)} \\ \text{ct}_b \leftarrow \text{Enc}(\text{sk}, m_b), b \leftarrow \{0, 1\} \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda).$$

γ -Unremovability of CPA-secure SKE The γ -Unremovability for a watermarkable CPA secure SKE scheme says, for all PPT admissible stateful adversary \mathcal{A} , there exists a negligible function $\text{negl}(\lambda)$ such that for all $\lambda \in \mathbb{N}$:

$$\Pr \left[\text{Extract}(\text{xk}, C) \notin \mathcal{Q} \wedge C \text{ is } \gamma\text{-good} : \begin{array}{l} (\text{sk}, \text{xk}, \text{mk}) \leftarrow \text{WMSetup}(1^\lambda) \\ C \leftarrow \mathcal{A}^{\text{Mark}(\text{mk}, \text{sk}, \cdot)}(1^\lambda) \end{array} \right] \leq \text{negl}(\lambda).$$

where \mathcal{Q} is the set of marks queried by \mathcal{A} and $C = (C_1, C_2)$ is said to be a (PPT admissible, stateful) γ -good program in the security game $G_{CCA}^2(\text{sk}, \cdot)$, more specifically:

$$\Pr \left[C_2^{\text{Enc}(\text{sk}, \cdot)}(\text{ct}_b) = b : \begin{array}{l} (m_0, m_1) \leftarrow C_1^{\text{Enc}(\text{sk}, \cdot)} \\ \text{ct}_b \leftarrow \text{Enc}(\text{sk}, m_b), b \leftarrow \{0, 1\} \end{array} \right] \geq \frac{1}{2} + \gamma.$$

Remark 5.2. The above definition is equivalent to letting A output (m_0, m_1) and making (m_0, m_1) as the auxiliary input aux for G_{CCA} and Extract . Because we can view any distribution over (m_0, m_1) used by C as a convex combination of different message pairs $\{(m_0, m_1)_i\}_i$ and its corresponding strategy such that the overall winning probability is $1/2 + \gamma$. Thus, we can let \mathcal{A} pick the message-pair and corresponding strategy with the largest winning probability and hardcode them into C instead. The other way of implication is easy to see.

This also applies to the CCA security game.

5.4 Watermarking-Compatible Reduction

We show that the textbook construction of CPA encryption from weak PRF is watermarking-compatible. Given a PRF scheme that satisfies weak PRF security $\text{PRF} = (\text{PRF.KeyGen}, \text{PRF.Eval})$, where input and output lengths are ℓ and key length is n .

The syntax of PRF and CPA-security is standard and we refer them to the watermarkable definitions Section 5.3, Section 5.3, by replace WMSetup with KeyGen and removing Extract , Mark algorithms.

Plain Construction of CPA-secure Encryption from weak PRF

- $\text{KeyGen}(1^\lambda) : \text{sk} \leftarrow \text{PRF.KeyGen}(\lambda)$
- $\text{Enc}(\text{sk}, m \in \{0, 1\}^\ell) \rightarrow \text{ct} : \text{samples } r \leftarrow \{0, 1\}^\ell; \text{ output } \text{ct} \leftarrow (r, \text{PRF.Eval}(\text{sk}, r) \oplus m).$
- $\text{Dec}(\text{sk}, \text{ct}) : \text{parse } \text{ct} := (r, \text{ct}'); \text{ compute } m := \text{PRF.Eval}(\text{sk}, r) \oplus \text{ct}'.$

Watermarking-Compatible Reduction We briefly recall the textbook reduction for the above construction, in the language of watermarking-compatible.

Claim 5.3. *The CPA-secure SKE to weak PRF pseudorandomness reduction is a watermarking compatible reduction.*

PRF pseudorandomness adversary \mathcal{A}_{PRF} simulates CPA game for adversary \mathcal{A} as follows:

- There is no public key and public evaluation algorithm PubEval in the scheme, so \mathcal{A}_{PRF} and \mathcal{A} will not receive information regarding public key pk.
- First consider stage 1: For any encryption query from \mathcal{A}^1 , $\mathcal{A}_{\text{PRF}}^1$ will simulate the response as follows: it will query the oracle $\mathcal{O}(\text{sk}, \cdot)$ in the security game G_{PRF} , where $\mathcal{O}(\cdot)$ outputs $\text{PRF.Eval}(\text{sk}, x)$ on any query x .
- After entering stage 2, for any encryption query from \mathcal{A}^2 , $\mathcal{A}_{\text{PRF}}^2$ will still simulate the response as the above. Recall that $\mathcal{A}_{\text{PRF}}^2$ will not get to see any queries made in stage 1.
- Note that $\mathcal{A}_{\text{PRF}}^2$'s access to $\mathcal{O}(\cdot)$ is changed to non-adaptive. $\mathcal{O}(\cdot)$ performs the following: upon every query, sample a fresh r , output $\text{PRF.Eval}(\text{sk}, r)$.
- In the challenge phase, $\mathcal{A}_{\text{PRF}}^2$ will use the challenge input from the weak PRF challenger: $r^* \leftarrow \{0, 1\}^\ell$, $y^* \in \{0, 1\}^\ell$ where y^* is either $\text{PRF.Eval}(\text{sk}, r^*)$ or uniformly random.
- \mathcal{A}^2 sends in challenge messages (m_0, m_1) ; $\mathcal{A}_{\text{PRF}}^2$ flips a coin $b \leftarrow \{0, 1\}$ and sends $\text{ct} = (r^*, y^* \oplus m_b)$ to \mathcal{A}^2 .
- If \mathcal{A}^2 guesses the correct b , then $\mathcal{A}_{\text{PRF}}^2$ output 0, for " y^* is $\text{PRF.Eval}(\text{sk}, r^*)$ ", else $\mathcal{A}_{\text{PRF}}^2$ output 1, for " y^* is uniformly random".

In the above reduction, $\mathcal{A}_{\text{PRF}}^2$'s reduction function $f(\text{out}, b, \mathcal{Q}, \text{inp})$ is: ignore \mathcal{Q}, inp , check if \mathcal{A}^2 's output out is equal to the randomness b used in preparing the challenge input, if yes output 0 else 1.

We denote the security game for weak PRF as G_{PRF} and for CPA encryption as G_{CPA} ; we also denote the advantage of \mathcal{A} in G_{CPA} as Adv_{CPA} and advantage of \mathcal{A}_{PRF} as Adv_{PRF} . We can observe that $\text{Adv}_{\text{PRF}} \geq \text{Adv}_{\text{CPA}}/2 - \frac{q}{2^\ell}$ where q is the number of queries made by \mathcal{A} .

5.4.1 Security of Watermarkable Implementation for CPA-secure SKE

Theorem 5.4. *Assuming that LWE is secure, then there exists secure single-key watermarkable implementation of CPA-secure SKE with private tracing.*

Assuming that there exists indistinguishability obfuscation, then there exists secure collusion-resistant watermarkable implementation of CPA-secure SKE with public tracing.

Theorem 5.5. *[[GKWW21, MW22]] Assuming that LWE is secure, then there exists secure collusion-resistant watermarkable implementation of weak PRF with private extraction.*

Assuming that there exists indistinguishability obfuscation, then there exists secure collusion-resistant watermarkable implementation of weak PRF with public extraction.

Note that both schemes above satisfy the properties we require in Section 4.1, including: extraction key of the scheme can be used to simulate the weak PRF pseudorandomness security game *perfectly*; the extraction procedure satisfies the extraction syntax¹⁰.

Remark 5.6. *In [GKWW21, MW22], the watermarkable implementation of weak PRF is called traceable PRF. Their exact extraction procedure does not satisfy our requirement that the adversarial circuit C can*

¹⁰For convenience, in the rest of this work, when we cite an existing watermarking scheme as a "watermarking implementation of a primitive", then the scheme satisfies all properties defined in Section 4.1.

make $\text{PRF.Eval}(\text{sk}, \cdot)$ queries during the extraction/tracing procedure. But it can be easily modified to satisfy this property: their tracing key has the capability to sample from the input-output space of the PRF, and thus can answer weak PRF queries.

Lemma 5.7. *Assuming that there exists collusion-resistant (resp. single-key) secure watermarkable implementation of weak PRF with public (resp. private) extraction, then there exists secure collusion-resistant (resp. single-key) watermarkable implementation of CPA-secure SKE with public (resp. private) extraction.*

To prove Lemma 5.7, we present the construction as below.

Watermarking Implementation of CPA-secure Encryption Given a watermarkable implementation of weak PRF (wPRF.WMSetup , wPRF.Eval , wPRF.Mark , wPRF.Extract), we can construct a watermarkable implementation of CPA-secure SKE as follows:

$\text{WMSetup}(1^\lambda) \rightarrow (\text{sk}, \text{xk}, \text{mk})$: compute $(\text{sk}, \text{xk}, \text{mk}) \leftarrow \text{wPRF.WMSetup}(1^\lambda)$
 $\text{Enc}(\text{sk}, m \in \{0, 1\}^\ell)$: samples $r \leftarrow \{0, 1\}^\ell$; output $\text{ct} \leftarrow (r, \text{wPRF.Eval}(\text{sk}, r) \oplus m)$.
 $\text{Dec}(\text{sk}, \text{ct})$: parse $\text{ct} := (r, \text{ct}')$; compute $m := \text{wPRF.Eval}(\text{sk}, r) \oplus \text{ct}'$.
 $\text{Mark}(\text{mk}, \text{sk}, \tau \in \mathcal{M}_{\text{Mark}})$:

1. parse $\text{mk} = \text{wPRF.mk}$
2. output $\text{sk}_\tau \leftarrow \text{wPRF.Mark}(\text{mk}, \text{sk}, \tau)$;

$\text{Extract}(\text{xk}, C)$:

1. On input circuit C and $\text{xk} = \text{wPRF.xk}$;
2. Treat C as the stage-2 adversary \mathcal{A}_2 in the watermarking-compatible reduction from CPA-security to weak PRF security to create circuit C_{PRF} (C_{PRF} has black-box access to C) as a stage-2 reduction $\mathcal{A}_{\text{PRF}}^2$ algorithm from CPA-security to the weak PRF security.
3. compute $\tau/\perp \leftarrow \text{wPRF.Extract}(\text{xk}, C_{\text{PRF}})$; if the mark extracted is not \perp , abort and output the mark.

Proof. We describe the reduction in the watermarking-compatible reduction language.

Suppose there exists a PPT adversary $\mathcal{A}_{\text{Mark}}$ for the γ -unremovability security of the above watermarkable CPA encryption, we can consider the adversary $\mathcal{A}_{\text{Mark}}$ to be the stage-1 adversary \mathcal{A}^1 in the CPA-security game: $\mathcal{A}_{\text{Mark}}$ gets to query the marking oracle $\text{Mark}(\text{mk}, \text{sk}, \cdot)$. The reduction \mathcal{B} which is an adversary in the strong unremovability security game of the watermarkable implementation of weak PRF, can simulate both the answers to the marking queries by querying the PRF's marking oracle.

Then $\mathcal{A}_{\text{Mark}}$ outputs a circuit C ; during the $\text{Extract}(\text{xk}, \cdot)$ algorithm. C will be treated as a stage-2 adversary in the CPA-security game; with black-box access to C , circuit C_{PRF} will act like a stage-2 reduction from CPA-security to weak PRF security. Since $\text{wPRF.Extract}(\text{wPRF.sk}, \cdot)$ algorithm will perfectly simulate the weak PRF security game on its input circuit, C_{PRF} will be able to answer C 's queries and use C 's output to finish the reduction.

The reduction \mathcal{B} will create a circuit C' with black-box use of C the same way as the above C_{PRF}^C and output C' in the $\gamma/2$ -unremovability security game of the watermarkable implementation of weak PRF.

By our assumption, C should satisfy that $\Pr[G_{\text{CPA}}(\text{sk}, C) = 1] \geq \frac{1}{2} + \gamma$ and $\Pr[\text{Extract}(\text{xk}, C) \notin \mathcal{Q}] \geq \epsilon$ for some non-negligible ϵ . By the design of our Extract algorithm, it must be the case that

$\Pr[\text{wPRF.Extract}(\text{wPRF.xk}, C_{\text{PRF}}) \notin \mathcal{Q}] \geq \epsilon$. Since the circuit C' created by \mathcal{B} is exactly the same as C_{PRF} and the set of marking queries \mathcal{Q} made by \mathcal{B} to the marking oracle is the set as $\mathcal{A}_{\text{Mark}}$'s, we must have that $\Pr[\text{wPRF.Extract}(\text{wPRF.xk}, C') \notin \mathcal{Q}] \geq \epsilon$.

Meanwhile, by the property of the reduction, we have that

$$\Pr[G_{\text{PRF}}(\text{wPRF.sk}, C') = 1] \geq \frac{1}{2} + \gamma/3$$

. Therefore, \mathcal{B} breaks $\gamma/3$ -unremovability for watermarkable implementation of weak PRF for some non-negligible γ . □

5.5 Watermarkable CCA-secure Encryption from Watermarkable MAC and PRF

In this section, we show an example of obtaining a watermarkable implementation of a CCA secure secret-key encryption scheme from a watermarkable implementation of weak PRF and a watermarkable implementation of MAC(or signature).

The security of watermarkable implementation of signatures (and therefore MAC) satisfying our requirements in Section 4.1 can be built from a modification of the watermarkable signature scheme in [GKM⁺19].

5.5.1 Watermarking-Compatible Reduction of CCA Security to MAC and weak PRF

We first give the "plain" construction from weak PRF and MAC to CCA-secure SKE and show that the reduction is watermarking-compatible.

Given a PRF scheme that satisfies weak PRF security $\text{PRF} = (\text{PRF.KeyGen}, \text{PRF.Eval})$, where input and output lengths are ℓ and key length is n and a MAC scheme $\text{MAC} = (\text{MAC.KeyGen}, \text{MAC.Sign}, \text{MAC.Verify})$.

Construction

- $\text{KeyGen}(1^\lambda) : \text{PRF.sk} \leftarrow \text{PRF.KeyGen}(\lambda), \text{MAC.sk} \leftarrow \text{MAC.KeyGen}(\lambda); \text{output sk} = (\text{PRF.sk}, \text{MAC.sk})$.
- $\text{Enc}(\text{sk}, m \in \{0, 1\}^\ell) \rightarrow \text{ct} :$
 - parse $\text{sk} = (\text{PRF.sk}, \text{MAC.sk})$;
 - samples $r \leftarrow \{0, 1\}^\ell$;
 - compute $\text{ct}_0 \leftarrow (r, \text{PRF.Eval}(\text{sk}, r) \oplus m)$;
 - compute $\sigma = \text{MAC.Sign}(\text{MAC.sk}, \text{ct}_0)$
 - output $\text{ct} = (\text{ct}_0, \sigma)$,
- $\text{Dec}(\text{sk}, \text{ct}) :$
 - parse $\text{ct} = (\text{ct}_0, \sigma)$ and $\text{sk} = (\text{PRF.sk}, \text{MAC.sk})$;
 - If $\text{MAC.Verify}(\text{MAC.sk}, \text{ct}_0, \sigma) = 1$ continue, else abort and output \perp .
 - parse $\text{ct}_0 := (r, \text{ct}')$; compute $m := \text{PRF.Eval}(\text{PRF.sk}, r) \oplus \text{ct}'$.
 - output m .

Claim 5.8. *The CCA-secure SKE to weak PRF pseudorandomness's reduction and the CCA-secure SKE to EUF-CMA-security of MACs reduction are both watermarking compatible reductions.*

Watermarking-Compatible Reduction We describe the security reduction of the above construction in the language of watermarking-compatible.

First consider a reduction for MAC security, \mathcal{A}_{MAC} simulates CPA game for adversary \mathcal{A} as follows:

- There is no public key and public evaluation algorithm PubEval in the scheme, so \mathcal{A}_{PRF} and \mathcal{A} will not receive information regarding public key pk.
- \mathcal{A}_{MAC} samples its own PRF secret key PRF.sk.
- For stage 1: For any encryption query from \mathcal{A}^1 , $\mathcal{A}_{\text{MAC}}^1$ will simulate the response as follows: it will compute the ct_0 part of ciphertext and then query the signing oracle $\text{MAC.Sign}(\text{MAC.sk}, \cdot)$ in the security game G_{MAC} . For any decryption query, \mathcal{A}_{MAC} will query the verification oracle $\text{MAC.Verify}(\text{MAC.sk}, \cdot)$.
- After entering stage 2, for any encryption or decryption query from \mathcal{A}^2 , $\mathcal{A}_{\text{MAC}}^2$ will still simulate the response as the above. Recall that $\mathcal{A}_{\text{MAC}}^2$ will not get to see any queries made in stage 1 and $\mathcal{A}_{\text{MAC}}^2$ will record all queries from \mathcal{A}_2 .
- In the challenge phase, \mathcal{A}^2 sends in challenge messages (m_0, m_1) ; $\mathcal{A}_{\text{MAC}}^2$ flips a coin $b \leftarrow \{0, 1\}$ and sends the encryption of m_b to \mathcal{A}^2 .
- $\mathcal{A}_{\text{MAC}}^2$ continues to simulate the encryption and decryption oracles for \mathcal{A}^2 .
- In the end, $\mathcal{A}_{\text{MAC}}^2$ will look up \mathcal{A}^2 's queries: find a decryption query $\text{ct} = (\text{ct}_0, \sigma)$ such that it has never been the output of an encryption query and the decryption oracle did not output \perp on this query. $\mathcal{A}_{\text{MAC}}^2$ output (ct_0, σ) as its forgery.

If such a decryption query exists in the end, then \mathcal{A}_{MAC} break the EUF-CMA-security of MAC (Otherwise, we will do a reduction to weak PRF).

In the above reduction, $\mathcal{A}_{\text{MAC}}^2$'s reduction function $f(\text{out}, r, \mathcal{Q}, \text{inp})$ is: ignore out, inp, r , check if \mathcal{A}^2 's queries \mathcal{Q} contain a decryption query where the input of this query is not from the output of an encryption query; if found, output this query.

Since such a ciphertext pair was never output by the encryption oracle, then it means that the reduction \mathcal{A}_{MAC} never queried the corresponding message-signature out of the decryption oracle. Thus \mathcal{A}_{MAC} can use it to break EUF-CMA-security of MAC.

If no such a decryption query exists, we can build a reduction to break the weak pseudorandomness of PRF. The reduction to weak PRF security will be similar to Section 5.1, except that the reduction will sample its own MAC keys to simulate the oracles. Since we have discussed this example in sufficient detail in the technical overview, we omit repeating the details here.

5.5.2 Definition: Watermarkable Implementation of CCA-secure SKE

We give the CCA2 security and unremovability security definition for CCA2-secure SKE. The other definitions (correctness, functionality preserving) follow naturally from the correctness and previously discussed CPA secure SKE.

CCA2-security A watermarkable CPA secure SKE scheme satisfies CPA security if for all PPT admissible stateful adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, there exists a negligible function $\text{negl}(\lambda)$ such that for all $\lambda \in \mathbb{N}$:

$$\Pr \left[\mathcal{A}_2^{\text{Enc}(\text{sk}, \cdot), \text{Dec}(\text{sk}, \cdot)}(\text{ct}_b) = b : \begin{array}{l} (m_0, m_1) \leftarrow \mathcal{A}_1^{\text{Enc}(\text{sk}, \cdot), \text{Dec}(\text{sk}, \cdot)} \\ \text{ct}_b \leftarrow \text{Enc}(\text{sk}, m_b), b \leftarrow \{0, 1\} \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda).$$

where \mathcal{A}_2 can only make decryption queries on $ct \neq ct_b$.

γ -Unremovability of CCA2-secure SKE The γ -Unremovability for a watermarkable CCA2 secure SKE scheme says, for all PPT admissible stateful adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, there exists a negligible function $\text{negl}(\lambda)$ such that for all $\lambda \in \mathbb{N}$:

$$\Pr \left[\text{Extract}(xk, C) \notin \mathcal{Q} \wedge C \text{ is } \gamma\text{-good} : \begin{array}{l} (sk, xk, mk) \leftarrow \text{WMSetup}(1^\lambda) \\ C \leftarrow \mathcal{A}^{\text{Mark}(mk, sk, \cdot)}(1^\lambda) \end{array} \right] \leq \text{negl}(\lambda).$$

where \mathcal{Q} is the set of marks queried by \mathcal{A} and $C = (C_1, C_2)$ is a PPT admissible, stateful γ -good adversary in the security game $G_{CCA}(sk, \cdot)$, more specifically:

$$\Pr \left[C_2^{\text{Enc}(sk, \cdot), \text{Dec}(sk, \cdot)}(ct_b) = b : \begin{array}{l} (m_0, m_1) \leftarrow C_1^{\text{Enc}(sk, \cdot), \text{Dec}(sk, \cdot)} \\ ct_b \leftarrow \text{Enc}(sk, m_b), b \leftarrow \{0, 1\} \end{array} \right] \geq \frac{1}{2} + \gamma.$$

C is admissible if it only make queries to the oracle on $ct \neq ct_b$.

We refer the watermarkable MAC (and signatures) definition to Appendix D.2.

5.5.3 Security of Watermarkable Implementation for CCA-secure SKE

Theorem 5.9. *Assuming LWE, then there exists secure bounded collusion-resistant watermarkable implementation of CCA-secure SKE with private extraction.*

Assuming that there exists indistinguishability obfuscation, then there exists secure bounded collusion-resistant watermarkable implementation of CCA-secure SKE with private tracing.

The above theorem is based on Theorem 5.5 and the following:

Theorem 5.10. *[Revision of [GKM⁺19]] Assuming that there exists secure OWFs, then there exists secure watermarkable implementation of digital signatures (and MAC) with private extraction and bounded collusion resistance.*

Remark 5.11. *The construction for the watermarkable digital signature scheme in the original [GKM⁺19] does not exactly satisfy our requirements in Section 4.1. But we will show that a minor modification of the scheme will suffice.*

Lemma 5.12. *Assuming that there exists collusion-resistant (resp. single-key) secure watermarkable implementation of weak PRF and MAC, then there exists secure collusion-resistant (resp. single-key) watermarkable implementation of CCA-secure SKE with private extraction.*

Remark 5.13. *[GKWW21, MW22] can give us a watermarkable weak PRF scheme that suffices for our use, though their definition of weak PRF allows only non-adaptive queries of random r and its evaluation. We need the weak PRF to allow adaptive queries to be used in CCA construction. But their private-key watermarkable PRF scheme already satisfies this stronger definition, by having the PRF evaluation key included in the tracing key.*

To prove Lemma 5.12, we present the construction as below.

Watermarking Implementation of CCA-secure Encryption Given a watermarkable implementation of weak PRF ($w\text{PRF.WMSetup}$, $w\text{PRF.Eval}$, $w\text{PRF.Mark}$, $w\text{PRF.Extract}$) and one of MAC ($w\text{MAC.WMSetup}$, $w\text{MAC.Sign}$, $w\text{MAC.Verify}$, $w\text{MAC.Mark}$, $w\text{MAC.Extract}$), we can construct a watermarkable implementation of CCA-secure SKE as follows:

$\text{WMSetup}(1^\lambda) \rightarrow (\text{sk}, \text{xk}, \text{mk})$:

1. compute $(w\text{PRF.sk}, w\text{PRF.xk}, w\text{PRF.mk}) \leftarrow w\text{PRF.WMSetup}(1^\lambda)$; compute $(w\text{MAC.sk}, w\text{MAC.xk}, \text{MAC.mk}) \leftarrow w\text{MAC.WMSetup}(1^\lambda)$
2. output $\text{sk} = (w\text{PRF.sk}, w\text{MAC.sk})$; $\text{mk} = (w\text{PRF.mk}, w\text{MAC.mk})$; $\text{xk} = (w\text{PRF.xk}, w\text{MAC.xk})$

$\text{Enc}(\text{sk}, m \in \{0, 1\}^\ell)$:

1. parse $\text{sk} = (w\text{PRF.sk}, w\text{MAC.sk})$;
2. samples $r \leftarrow \{0, 1\}^\ell$;
3. compute $\text{ct}_0 \leftarrow (r, w\text{PRF.Eval}(w\text{PRF.sk}, r) \oplus m)$;
4. compute $\sigma = w\text{MAC.Sign}(w\text{MAC.sk}, \text{ct}_0)$
5. output $\text{ct} = (\text{ct}_0, \sigma)$,

$\text{Dec}(\text{sk}, \text{ct})$:

1. parse $\text{ct} = (\text{ct}_0, \sigma)$ and $\text{sk} = (w\text{PRF.sk}, w\text{MAC.sk})$;
2. If $w\text{MAC.Verify}(w\text{MAC.sk}, \text{ct}_0, \sigma) = 1$ continue, else abort and output \perp .
3. parse $\text{ct}_0 := (r, \text{ct}')$; compute $m := w\text{PRF.Eval}(w\text{PRF.sk}, r) \oplus \text{ct}'$.
4. output m .

$\text{Mark}(\text{mk}, \text{sk}, \tau \in \mathcal{M}_{\text{Mark}})$:

1. parse $\text{mk} = (w\text{PRF.mk}, w\text{MAC.mk})$
2. compute $\text{sk}_{\text{PRF}, \tau} \leftarrow w\text{PRF.Mark}(w\text{PRF.mk}, w\text{PRF.sk}, \tau)$;
 $\text{sk}_{\text{MAC}, \tau} \leftarrow w\text{MAC.Mark}(w\text{MAC.mk}, w\text{MAC.sk}, \tau)$;
3. output $(\text{sk}_{\text{PRF}, \tau}, \text{sk}_{\text{MAC}, \tau})$

$\text{Extract}(\text{xk}, C)$:

1. On input circuit C and $\text{xk} = (w\text{PRF.xk}, w\text{MAC.xk})$;
2. Initialize empty tuple $\vec{\tau}$.
3. Let $P_1 := \text{PRF}$ and $P_2 := \text{MAC}$; For $i = 1, 2$:
 - (a) Treat C as the stage-2 adversary \mathcal{A}_2 in the watermarking-compatible reduction from CCA-security to P_i security; create a circuit $C_{P_i}^C$ (i.e. C_{P_i} has black-box access to C) as a stage-2 reduction algorithm \mathcal{A}_i^2 from CCA-security game to P_i 's security game.
 - (b) compute $\tau/\perp \leftarrow wP_i.\text{Extract}(wP_i.\text{xk}, C_{P_i}^C)$; if the mark extracted is not \perp , add it to $\vec{\tau}$.
4. output $\vec{\tau}$

Now we prove the unremovability security of the watermarkable implementation of the CCA secure encryption scheme.

Proof. We describe the reduction in the watermarking-compatible reduction language.

Suppose there exists a PPT adversary $\mathcal{A}_{\text{Mark}}$ for the γ -unremovability security of the above watermarkable CCA encryption, we can consider the adversary $\mathcal{A}_{\text{Mark}}$ to be the stage-1 adversary

\mathcal{A}^1 in the CCA-security game: $\mathcal{A}_{\text{Mark}}$ gets to query the marking oracle $\text{Mark}(\text{mk}, \text{sk}, \cdot)$, which is a leakage on the secret key $\{\text{wPRF.sk}, \text{wMAC.sk}\}$.

Let us denote $P_1 = \text{PRF}$; $P_2 = \text{wMAC}$. The reduction \mathcal{B}_i which is an adversary in the $\gamma/3$ -unremovability security game of the watermarkable implementation of primitive P_i , can simulate both the answers to the marking queries by querying the challenger.

Then $\mathcal{A}_{\text{Mark}}$ outputs a circuit C ; during the $\text{Extract}(\text{xk}, \cdot)$ algorithm. C will be treated as a stage-2 adversary in the CCA-security game; with black-box access to C . circuit C_{P_i} will act like a stage-2 reduction from CCA-security to P_i 's security. Since $\text{wP}_i.\text{Extract}(\text{wP}_i.\text{xk}, \cdot)$ algorithm will perfectly simulate the security game G_{P_i} on its input circuit, C_{P_i} will be able to answer C 's queries and use C 's output to finish the reduction.

Since we are assuming that $\Pr[\text{Extract}(\text{xk}, C) \notin \mathcal{Q}] \geq \epsilon$, for some non-negligible ϵ . By the design of our Extract algorithm, it must be that for both $i \in \{1, 2\}$, $\Pr[\text{wP}_i.\text{Extract}(\text{wP}_i.\text{xk}, C_i^C) \notin \mathcal{Q}] \geq \epsilon$.

We now consider two cases. For any $\mathcal{A}^{\text{Enc}(\text{sk}, \cdot), \text{Dec}(\text{sk}, \cdot), \text{Mark}(\text{mk}, \text{sk}, \cdot)}(\text{pk})$ producing a γ -good C such that $\Pr[\text{Extract}(\text{xk}, \text{pk}, C) \notin \mathcal{Q}] \geq \epsilon$, one of the following cases must hold:

- Case 1:** In case 1, with some non-negligible probability, C_1 is γ_1 -good in the game $G_{\text{PRF}}(\text{wPRF.sk}, \cdot)$ for some non-negligible γ_1 . In other words, during the execution of $\text{Extract}(\text{xk}, C)$, C makes decryption queries on ciphertexts that all come from the encryption queries.

The reduction \mathcal{B}_1 will create a circuit C'_1 with black-box use of C the same way as the circuit $C_{P_1}^C$ created in the extraction algorithm and output C'_1 in the γ_1 -unremovability security game of the watermarkable implementation of weak PRF.

Since the circuit C' created by \mathcal{B}_1 is exactly the same as C_{PRF}^C and the set of marking queries \mathcal{Q} made by \mathcal{B} to the marking oracle is the set as $\mathcal{A}_{\text{Mark}}$'s, we must have that $\Pr[\text{wPRF.Extract}(\text{wPRF.xk}, C'_1) \notin \mathcal{Q}] \geq \epsilon$.

Meanwhile, by the property of the reduction, we have that $\Pr[G_{\text{PRF}}(\text{wPRF.sk}, C'_1) = 1] \geq \frac{1}{2} + \gamma_1$ for some non-negligible γ_1 . Therefore, \mathcal{B} breaks γ_1 -unremovability for watermarkable implementation of weak PRF.
- Case 2:** In case 2, with some non-negligible probability, C_2 is γ_2 -good in the game $G_{\text{MAC}}(\text{wMAC.sk}, \cdot)$ for some non-negligible γ_2 .

In other words, during the execution of $\text{Extract}(\text{xk}, C)$, C makes decryption queries on ciphertexts that not only come from the encryption queries.

The reduction \mathcal{B}_2 will create a circuit C'_2 with black-box use of C the same way as the circuit C_{MAC}^C created in the extraction algorithm and output C'_2 in the γ_2 -unremovability security game of the watermarkable implementation of weak PRF.

Since the circuit C' created by \mathcal{B}_2 is exactly the same as C_{MAC} and the set of marking queries \mathcal{Q} made by \mathcal{B}_2 to the marking oracle is the set as $\mathcal{A}_{\text{Mark}}$'s, we must have that $\Pr[\text{wMAC.Extract}(\text{wMAC.xk}, C'_2) \notin \mathcal{Q}] \geq \epsilon$.

Meanwhile, by the property of the reduction, we have that $\Pr[G_{\text{MAC}}(\text{wMAC.sk}, C'_2) = 1] \geq \frac{1}{2} + \gamma_2$ for some non-negligible γ_2 . Therefore, \mathcal{B}_2 breaks γ_2 -unremovability for watermarkable implementation of MAC.

By the property of the watermarking-compatible reduction (shown in Section 5.5.1), one of the above two cases must hold. \square

6 Watermarkable CCA-secure PKE from Watermarkable Identity-based Encryption and Strong One-Time Signature

In this section, we give an example of how to achieve watermarkable implementation of CCA-secure PKE, where type-2 reduction is used. The watermarkable CCA-secure PKE scheme is based on the construction from selectively secure IBE and strong one-time signatures ([BCHK07])

6.1 Definition: Watermarkable CCA2-secure PKE

A watermarkable implementation of a SKE scheme consists of the following algorithms:

WMS $\text{Setup}(1^\lambda) \rightarrow (\text{sk}, \text{pk}, \text{mk}, \text{xk})$: on input security parameter, outputs a secret key sk , public key pk , marking key mk , extraction key xk .
 $\text{Enc}(\text{pk}, m) \rightarrow \text{ct}$: on input public key pk and message m , outputs a ciphertext ct .
 $\text{Dec}(\text{sk}, \text{ct}) \rightarrow m$: on input secret key sk and ciphertext ct outputs a message m .
 $\text{Mark}(\text{mk}, \text{sk}, \tau)$: on marking key mk , secret key sk , a message τ , output a marked key sk_τ .
 $\text{Extract}(\text{xk}, \text{pk}, C)$: on extraction key xk , public key pk and program C , output a mark $\tau \in \mathcal{M}_\tau$ or \perp .

The correctness property is the same as the correctness property of a plain PKE scheme and we combine it with the functionality preserving property it due to it being standard.

Correctness and Functionality Preserving A watermarkable PKE is functionality preserving if then there exists a negligible function $\text{negl}(\lambda)$ such that for all $\lambda \in \mathbb{N}, m \in \mathcal{M}, \tau \in \mathcal{M}_\tau$:

$$\Pr \left[\text{Dec}(\tilde{\text{sk}}, \text{ct}) = m \wedge \text{Dec}(\text{sk}, \text{ct}) = m : \begin{array}{l} (\text{sk}, \text{pk}, \text{xk}, \text{mk}) \leftarrow \text{WMS\text{Setup}}(1^\lambda) \\ \tilde{\text{sk}} \leftarrow \text{Mark}(\text{mk}, \text{sk}, \tau) \\ \text{ct} \leftarrow \text{Enc}(\text{pk}, m) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

CCA2-security A watermarkable PKE scheme satisfies CCA2 security if for all PPT admissible stateful adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, there exists a negligible function $\text{negl}(\lambda)$ such that for all $\lambda \in \mathbb{N}$:

$$\Pr \left[\mathcal{A}_2^{\text{Dec}(\text{sk}, \cdot)}(\text{st}, \text{ct}_b) = b : \begin{array}{l} (m_0, m_1, \text{st}) \leftarrow \mathcal{A}_1^{\text{Dec}(\text{sk}, \cdot)}(\text{pk}) \\ \text{ct}_b \leftarrow \text{Enc}(\text{pk}, m_b), b \leftarrow \{0, 1\} \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda).$$

\mathcal{A} is admissible if it only make queries to the oracle on $\text{ct} \neq \text{ct}_b$.

γ -Unremovability of CCA2-secure PKE The γ -Unremovability for a watermarkable CCA2 secure PKE scheme says, for all PPT admissible stateful adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, there exists a negligible function $\text{negl}(\lambda)$ such that for all $\lambda \in \mathbb{N}$:

$$\Pr \left[\text{Extract}(\text{xk}, \text{pk}, C) \notin \mathcal{Q} \wedge C \text{ is } \gamma\text{-good} : \begin{array}{l} (\text{sk}, \text{pk}, \text{xk}, \text{mk}) \leftarrow \text{WMS\text{Setup}}(1^\lambda) \\ C \leftarrow \mathcal{A}^{\text{Mark}(\text{mk}, \text{sk}, \cdot)}(1^\lambda, \text{pk}) \end{array} \right] \leq \text{negl}(\lambda).$$

where \mathcal{Q} is the set of marks queried by \mathcal{A} and $C = (C_1, C_2)$ is a PPT admissible, stateful γ -good adversary in the security game $G_{CCA}(\text{sk}, \text{pk}, \cdot)$, more specifically:

$$\Pr \left[C_2^{\text{Dec}(\text{sk}, \cdot)}(\text{ct}_b) = b : \begin{array}{l} (m_0, m_1) \leftarrow C_1^{\text{Dec}(\text{sk}, \cdot)}(\text{pk}, 1^\lambda) \\ \text{ct}_b \leftarrow \text{Enc}(\text{pk}, m_b), b \leftarrow \{0, 1\} \end{array} \right] \geq \frac{1}{2} + \gamma.$$

C is admissible if it only make queries to the oracle on $\text{ct} \neq \text{ct}_b$.

6.1.1 Preliminaries

Definition: Watermarkable Implementation of Identity-Based Encryption For clarity, we first give the definition for watermarkable implementation of a selectively secure identity-based encryption wIBE where one can mark the master secret key, though it is easy to see that they match the definition in Section 4.1. It consists of the following algorithms.

- WMSSetup(1^λ) \rightarrow (msk, mpk, mk, xk): on input security parameter, outputs a master secret key msk, master public key mpk, marking key mk, extraction key xk.
- KeyGen(msk, id) \rightarrow (sk_{id}): on input mster secret key, and id $\in \mathcal{ID}$, output a key sk_{id}.
- Enc(mpkm, m, id) \rightarrow ct_{id}: on input public key pk, message m and id $\in \mathcal{ID}$, outputs a ciphertext ct.
- Dec(sk_{id}, ct) \rightarrow m : on input secret key sk_{id} and ciphertext ct outputs a message m .
- Mark(mk, msk, τ) \rightarrow msk _{τ} : on marking key mk, secret key msk, a message τ , output a marked key msk _{τ} .
- Extract(xk, mpk, aux, C) \rightarrow τ/\perp : on extraction key sk, public key mpk auxiliary input aux and program C , output a mark $\tau \in \mathcal{M}_\tau$ or \perp .

Note that for our applications, we only consider marking a master secret key. But a scheme where one marks an ID-embedded key can also be applicable in certain compositions. Both types of watermarkable IBE can be consructed from modifications of the watermarkable ABE scheme in [GKM⁺19].

Correctness The functionality-preserving property says, for all $m \in \mathcal{M}$, id $\in \mathcal{ID}$, there exists a negligible function $\text{negl}(\lambda)$ such that:

$$\Pr \left[\begin{array}{l} (\text{Dec}(\text{sk}_{\text{id}}, \text{ct}) = m) \wedge \\ (\text{Dec}(\text{msk}, \text{ct}) = m) \end{array} : \begin{array}{l} (\text{msk}, \text{mpk}, \text{xk}, \text{mk}) \leftarrow \text{WMSSetup}(1^\lambda) \\ \text{sk}_{\text{id}} \leftarrow \text{KeyGen}(\text{msk}, \text{id}) \\ \text{ct} \leftarrow \text{Enc}(\text{mpk}, \text{id}, m) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

Functionality-Preserving The functionality-preserving property says, for all $m \in \mathcal{M}$, id $\in \mathcal{ID}$ and all $\tau \in \mathcal{M}_\tau$, there exists a negligible function $\text{negl}(\lambda)$ such that:

$$\Pr \left[\begin{array}{l} (\text{Dec}(\text{sk}_{\text{id}, \tau}, \text{ct}) = m) \wedge \\ (\text{Dec}(\text{msk}_\tau, \text{ct}) = m) \end{array} : \begin{array}{l} (\text{msk}, \text{mpk}, \text{xk}, \text{mk}) \leftarrow \text{WMSSetup}(1^\lambda) \\ \text{msk}_\tau \leftarrow \text{Mark}(\text{mk}, \text{sk}, \tau) \\ \text{sk}_{\text{id}, \tau} \leftarrow \text{KeyGen}(\text{msk}_\tau, \text{id}) \\ \text{ct} \leftarrow \text{Enc}(\text{mpk}, \text{id}, m) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

Selective-ID Security The selective-ID CPA security property says, for all $m \in \mathcal{M}$, $\text{id} \in ID$ and all $\tau \in \mathcal{M}_\tau$, and for all PPT admissible stateful adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$, there exists a negligible function $\text{negl}(\lambda)$ such that:

$$\Pr \left[\mathcal{A}_3^{\text{KeyGen}(\text{msk}, \cdot)}(\text{ct}_b) = b : \begin{array}{l} \text{id}^* \leftarrow \mathcal{A}_1(1^\lambda) \\ (\text{msk}, \text{mpk}, \text{xk}, \text{mk}) \leftarrow \text{WMSetup}(1^\lambda) \\ (m_0, m_1) \leftarrow \mathcal{A}_2^{\text{KeyGen}(\text{msk}, \cdot)}(\text{mpk}) \\ \text{ct} \leftarrow \text{Enc}(\text{mpk}, \text{id}^*, m_b), b \leftarrow \{0, 1\} \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda).$$

where $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ is admissible if $\mathcal{A}_2, \mathcal{A}_3$ only makes queries to the oracle $\text{KeyGen}(\text{msk}, \cdot)$ on $\text{id} \neq \text{id}^*$.

Additionally, we denote the above security game after $\text{id}^*, \text{msk}, \text{mpk}$ are fixed as $G_{\text{IBE}}(\text{msk}, \text{mpk}, \text{id}^*, \cdot)$.

Note that the stages in the above definition do not exactly match the stage in the watermarking-compatible reduction, as shown below.

γ -Unremovability The γ -Unremovability for a watermarkable selective-ID secure IBE scheme says, for all $m \in \mathcal{M}$, $\text{id} \in ID$ and all $\tau \in \mathcal{M}_\tau$, and for all PPT admissible stateful adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, there exists a negligible function $\text{negl}(\lambda)$ such that:

$$\Pr \left[\begin{array}{l} \text{Extract}(\text{xk}, \text{mpk}, \text{aux} = \text{id}^*, C) \notin \mathcal{Q} \\ \wedge C \text{ is } \gamma\text{-good} \end{array} : \begin{array}{l} \text{id}^* \leftarrow \mathcal{A}_1(1^\lambda) \\ (\text{msk}, \text{mpk}, \text{xk}, \text{mk}) \leftarrow \text{WMSetup}(1^\lambda) \\ C \leftarrow \mathcal{A}_2^{\text{Mark}(\text{mk}, \text{msk}, \cdot)}(\text{mpk}) \end{array} \right] \leq \text{negl}(\lambda).$$

where \mathcal{Q} is the set of marks queried by \mathcal{A}_2 and $C = (C_1, C_2)$ is a PPT admissible, stateful γ -good adversary in the security game $G_{\text{IBE}}(\text{msk}, \text{mpk}, \text{id}^*, \cdot)$, more specifically:

$$\Pr \left[C_2^{\text{KeyGen}(\text{msk}, \cdot)}(\text{ct}_b) = b : \begin{array}{l} (m_0, m_1) \leftarrow C_1^{\text{KeyGen}(\text{msk}, \cdot)}(\text{mpk}) \\ \text{ct}_b \leftarrow \text{Enc}(\text{mpk}, \text{id}^*, m_b), b \leftarrow \{0, 1\} \end{array} \right] \geq \frac{1}{2} + \gamma.$$

C is admissible if and only if C only make queries to the oracle $\text{KeyGen}(\text{msk}, \cdot)$ on $\text{id} \neq \text{id}^*$.

6.2 Watermarkable CCA-secure PKE Construction

In this section, we prove Theorem A.1 with an alternative approach, based on [BCHK07].

Theorem 6.1. *Assuming the hardness of LWE, there exists secure watermarkable implementation of identity-based encryption with selective-ID security, with private extraction and collusion-resistant security.*

We prove this above theorem in Appendix C by making a modification of [GKM⁺19].

Lemma 6.2. *Assuming there exists secure watermarkable implementation of (private-extraction, collusion resistant) identity-based encryption with selective-ID security and strongly unforgeable one-time signatures, then there exists secure (private-extraction, collusion resistant) watermarkable implementation of CCA-secure PKE*

Now we describe the construction for watermarkable implementation of CCA-secure PKE from a watermarkable implementation of selective-ID secure IBE scheme $\text{wIBE} = (\text{WMSetup}, \text{KeyGen}, \text{Enc}, \text{Dec}, \text{Mark}, \text{Extract})$ and a one-time strongly unforgeable signature scheme $\text{OTS} = (\text{KeyGen}, \text{Sign}, \text{Verify})$ (without watermarking).

$\text{WMSSetup}(1^\lambda) \rightarrow (\text{msk}, \text{mpk}, \text{mk}, \text{xk})$: compute $(\text{wIBE.msk}, \text{wIBE.mpk}, \text{wIBE.mk}, \text{wIBE.xk}) \leftarrow \text{wIBE.WMSSetup}(1^\lambda)$;
output $\text{sk} = \text{wIBE.msk}$; $\text{pk} = \text{wIBE.mpk}$; $\text{mk} = \text{wIBE.mk}$; $\text{xk} = (\text{wIBE.xk}, \text{wIBE.mpk})$.

$\text{Enc}(\text{pk}, m) \rightarrow \text{ct}$:

1. compute $(\text{OTS.vk}, \text{OTS.sk}) \leftarrow \text{OTS.KeyGen}(1^\lambda)$;
2. compute $\text{ct}_{\text{vk}} \leftarrow \text{wIBE.Enc}(\text{pk}, \text{id} = \text{vk}, m)$;
3. compute $\sigma \leftarrow \text{OTS.Sign}(\text{OTS.sk}, \text{ct}_{\text{vk}})$;
4. output $\text{ct} = (\text{ct}_{\text{vk}}, \text{sig}, \text{vk})$.

$\text{Dec}(\text{sk}, \text{ct}) \rightarrow m$:

1. parse $\text{ct} := (\text{ct}_{\text{vk}}, \text{sig}, \text{vk})$, $\text{sk} = \text{wIBE.msk}$;
2. if $\text{OTS.Verify}(\text{vk}, (\text{ct}_{\text{vk}}, \text{sig})) = 1$, continue; else abort and output \perp .
3. compute $\text{sk}_{\text{vk}} \leftarrow \text{wIBE.KeyGen}(\text{wIBE.msk}, \text{id} = \text{vk})$;
4. Output $m \leftarrow \text{wIBE.Dec}(\text{sk}_{\text{vk}}, \text{ct}_{\text{vk}})$.

$\text{Mark}(\text{mk}, \text{sk}, \tau) \rightarrow \text{sk}_\tau$: compute $\text{sk}_\tau \leftarrow \text{wIBE.Mark}(\text{mk}, \text{sk} = \text{wIBE.msk}, \tau)$.

$\text{Extract}(\text{xk}, \text{aux} = \perp, C) \rightarrow \tau / \perp$:

1. parse $\text{xk} := \text{wIBE.xk}, \text{wIBE.mpk}$. Samples $(\text{OTS.vk}^*, \text{OTS.sk}^*) \leftarrow \text{OTS.KeyGen}(1^\lambda)$.
2. Create a circuit C_{IBE} that has black-box access to C .
 - C_{IBE} is hardcoded with challenge $\text{id}^* = \text{OTS.vk}^*$ and OTS.sk^* , wIBE.mpk
 - C_{IBE} simulates the CCA PKE security game for C by simulating $\text{Dec}(\text{sk}, \cdot)$ oracles as follows:
 - (a) when C makes a query $\text{ct} = (\text{ct}_{\text{vk}}, \text{sig}, \text{vk})$; first check if $\text{OTS.Verify}(\text{vk}, (\text{ct}_{\text{vk}}, \text{sig})) = 1$, if yes continue; else output \perp .
 - (b) make an external query to a $\text{wIBE.KeyGen}(\text{msk}, \cdot)$ oracle for $\text{sk}_{\text{vk}} \leftarrow \text{wIBE.KeyGen}(\text{msk}, \text{vk})$. Output $m \leftarrow \text{wIBE.Dec}(\text{sk}_{\text{vk}}, \text{ct}_{\text{vk}})$.
 - In the challenge phase, C outputs (m_0, m_1) ; then C_{IBE} submits (m_0, m_1) to the external challenger and receives challenger ciphertext $\text{ct}_{b, \text{vk}^*} = \text{Enc}(\text{mpk}, \text{OTS.vk}^*, m_b)$, $b \leftarrow \{0, 1\}$;
 - feed C with $\text{ct}_b^* = (\text{ct}_{b, \text{vk}^*}, \text{sig}^* = \text{OTS.Sign}(\text{OTS.sk}^*, \text{ct}_{b, \text{vk}^*}), \text{OTS.vk}^*)$.
 - Continue to simulate $\text{Dec}(\text{sk}, \cdot)$ oracle for C and only answer queries $\text{ct} \neq \text{ct}_b^*$; finally C outputs a bit b' and C_{IBE} outputs the same.
3. Output $\tau / \perp \leftarrow \text{wIBE.Extract}(\text{wIBE.xk}, \text{aux} = \text{id}^*, C_{\text{IBE}})$;

Proof. We devide our analysis into cases.

Suppose there exists adversary \mathcal{A} that breaks the γ -unremovability of watermarkable CCA-secure PKE for some non-negligible γ , i.e. $\mathcal{A}^{\text{Dec}(\text{sk}, \cdot), \text{Mark}(\text{mk}, \text{sk}, \cdot)}(\text{pk})$ produces some program C such that $\Pr[G_{\text{CCA}}(\text{sk}, \text{pk}, C) = 1] \geq \frac{1}{2} + \gamma$ (for G_{CCA} see Section 5.5.2) and $\Pr[\text{Extract}(\text{xk}, C) \in \mathcal{Q}] \geq \epsilon$ for some non-negligible ϵ .

For any $(\text{sk}, \text{pk}, \text{xk}, \text{mk})$ generated by WMSSetup and any $\mathcal{A}^{\text{Mark}(\text{mk}, \text{sk}, \cdot)}(\text{pk})$ producing a γ -good C such that $\Pr[\text{Extract}(\text{xk}, \text{pk}, C) \notin \mathcal{Q}] \geq \epsilon$, one of the following cases must hold:

- **Case 1: during the execution of $\text{Extract}(\text{xk}, C)$, the program C produced by the adversary \mathcal{A} will only make queries that have the format $\text{ct} = (\text{ct}_{\text{vk}}, \text{vk}, \text{sig})$, where $\text{vk} \neq \text{vk}^*$, vk^* is the OTS's verification key part in the challenge ciphertext $\text{ct}_b^* = (\text{ct}_{b, \text{vk}^*}, \text{sig}^*, \text{OTS.vk}^*)$.**

In this case, we show that we can use \mathcal{A} to break the γ -unremovability of the watermarkable IBE.

The reduction $\mathcal{B}_{\text{wIBE}}$ interacts with \mathcal{A} as follows:

- \mathcal{B} samples $(vk^*, sk^*) \leftarrow \text{OTS.KeyGen}(1^\lambda)$ and submits vk^* as the challenge id^* to the challenger;
- $\mathcal{B}_{\text{wIBE}}$ receives wIBE.mpk from the challenger and gives it as the public key to \mathcal{A} . For \mathcal{A} 's marking queries, \mathcal{B} queries the marking oracle $\text{wIBE.Mark}(\text{wIBE.mk}, \text{wIBE.msk}, \cdot)$;

Next, \mathcal{A} outputs a program C . \mathcal{B} creates a circuit C_{wIBE} with black-box access to C . C is treated as a stage-2 adversary in the reduction from CCA-secure PKE to selective-ID secure IBE:

- C_{wIBE} is hardcoded with the challenge $id^* = vk^*$ and its corresponding signing key sk^* and wIBE.mpk . Note that C_{wIBE} does not receive any queries \mathcal{B} has answered/queried when answering \mathcal{A} 's queries in the previous stage.
- C_{wIBE} simulates answers to C 's decryption queries: when C makes a query $ct = (ct_{vk}, \text{sig}, vk)$; first check if $\text{OTS.Verify}(vk, (ct_{vk}, \text{sig})) = 1$, if yes continue; else output \perp . Make an external query to the $\text{wIBE.KeyGen}(\text{msk}, \cdot)$ oracle for $sk_{vk} \leftarrow \text{wIBE.KeyGen}(\text{msk}, vk)$. Output $m \leftarrow \text{wIBE.Dec}(sk_{vk}, ct_{vk})$. Since all the vk queried are not equal to vk^* , these queries are all valid.
- In the challenge phase, C outputs (m_0, m_1) ; then C_{wIBE} submits (m_0, m_1) to the external challenger and receives challenger ciphertext $ct_{b, vk^*} = \text{Enc}(\text{mpk}, vk^*, m_b)$, $b \leftarrow \{0, 1\}$;
- feed C with $ct_b^* = (ct_{b, vk^*}, \text{sig}^* = \text{OTS.Sign}(sk^*, ct_{b, vk^*}), vk^*)$.
- Continue to simulate $\text{Dec}(sk, \cdot)$ oracle for C and only answer queries $ct \neq ct_b^*$; finally C outputs a bit b' and C_{wIBE} outputs the same.

In the execution of the extraction $\text{Extract}(xk, C)$, C is used in the program C_{IBE} as a black box and by our design, it must hold that $\Pr[\text{wIBE.Extract}(xk, \text{mpk}, C_{\text{IBE}}) \in \mathcal{Q} \wedge C \text{ is } \gamma\text{-good}] \geq \epsilon$ for some non-negligible ϵ . Note that C_{IBE} created by $\text{Extract}(xk, C)$ differs from C_{wIBE} created by \mathcal{B} in that: $\text{Extract}(xk, C)$ samples a fresh $(vk^*, sk^*) \leftarrow \text{OTS.KeyGen}(1^\lambda)$; C_{wIBE} is hardcoded with a (vk^*, sk^*) sampled by \mathcal{B} previously. However, the challenge id 's in both settings are sampled freshly at random, independent of the $\text{wIBE.KeyGen}(\text{wIBE.msk}, \cdot)$ (i.e. $\text{Dec}(sk, \cdot)$) queries made in the stage of \mathcal{B} interacting with \mathcal{A} before \mathcal{A} outputs C (i.e. stage-1 in the watermarking-compatible reduction). Therefore, the output of $C_{\text{wIBE}}, C_{\text{IBE}}$ are the same unless in the real CCA PKE security game, \mathcal{A} has made a query on some ciphertext $ct = (ct_{vk^*} = \text{wIBE.Enc}(\text{mpk}, vk^*, m_b), vk^*, \text{sig})$ for the same vk^* that is later sampled by $\text{Extract}(xk, C)$ and for one of the challenge messages $m_b \in \{m_0, m_1\}$. But this happens with negligible probability since vk^* is sampled at random.

Since we have $\Pr[\text{wIBE.Extract}(xk, \text{aux} = vk^*, C_{\text{IBE}}) \in \mathcal{Q}] \geq \epsilon$ and given our assumption in case 1, we also have $\Pr[G_{\text{IBE}}(\text{wIBE.mpk}, \text{wIBE.msk}, id^* = vk^*, C_{\text{IBE}}) = 1] \geq 1/2 + \gamma$, we should deduce that the program C_{wIBE} produced by the reduction $\mathcal{B}_{\text{wIBE}}$ from C performs the same with all but negligible difference: with probability $\epsilon - \text{negl}(\lambda)$, we still have $\Pr[\text{wIBE.Extract}(xk, \text{aux} = id^* C_{\text{wIBE}}) \in \mathcal{Q}]$ and $\Pr[G_{\text{IBE}}(\text{wIBE.mpk}, \text{wIBE.msk}, \text{aux} = id^*, C_{\text{wIBE}}) = 1] \geq 1/2 + \gamma$.

- **Case 2: during the execution of $\text{Extract}(xk, C)$, the program C produced by the adversary \mathcal{A} will make at least one query that have the format $ct = (ct_{vk}, \text{sig}, vk)$, where $vk = \text{OTS.vk}^*$, OTS.vk^* is the OTS's verification key part in the challenge ciphertext ct^* ; further more, the query $ct = (ct_{vk}, \text{sig}, vk)$ satisfies $\text{OTS.Verify}(\text{OTS.vk}^*, ct_{vk}, \text{sig}) = 1$ and $ct^* \neq ct$**
In this case, we show a Type-2 watermarking compatible reduction to the strongly unforgeability security of the one-time signature scheme.

In stage-1 of the watermarking-compatible reduction, $\mathcal{B}_{\text{OTS}}^1$ operates as follows:

- $\mathcal{B}_{\text{OTS}}^1$ receives vk^* from the challenger. It samples $w\text{BE.xk}, msk, mpk, mk$ on its own. $\mathcal{B}_{\text{OTS}}^1$ can thus answer all the marking queries from \mathcal{A} .

After \mathcal{A} outputs program C , since we are given that $\Pr[G_{\text{CCA}}(\text{sk}, \text{pk}, C) = 1] \geq 1/2 + \gamma$ and C makes decryption queries on valid ciphertexts with the vk^* . \mathcal{B}_{OTS} can do the following.

Entering stage-2 of the reduction, $\mathcal{B}_{\text{OTS}}^2$ receives all stateful information from $\mathcal{B}_{\text{OTS}}^1$ except for the queries made by \mathcal{A} .

$\mathcal{B}_{\text{OTS}}^2$ runs C and simulates the CCA-security game, $\text{Dec}(\text{sk}, \cdot)$ oracles for C as follows:

- For decryption queries of the format $ct = (ct_0, vk, sig)$, first checks if $\text{OTS.Verify}(vk, (ct_{vk}, sig)) = 1$, if yes continue; else output \perp .

Then decrypt ct directly using $sk_{vk} \leftarrow \text{KeyGen}(msk, vk)$ and $m \leftarrow \text{Dec}(sk_{vk}, ct_{vk})$.

In the challenge phase, C submits 2 messages (m_0, m_1) and $\mathcal{B}_{\text{OTS}}^2$ outputs challenge ciphertext $ct^* = (ct_{vk^*, b} = w\text{BE.Enc}(mpk, vk^*, m_b \leftarrow (m_0, m_1)), vk^*, sig^*)$ by making one query to the signing oracle $\text{OTS.Sign}(sk^*, \cdot)$ on message $ct_{vk^*, b}$.

Then the query phase continues and C is only allowed to query on $ct \neq ct^*$. Since in the end, there will be one query ct such that $ct = (ct', vk^*, sig')$ where $ct \neq ct^*$ but $\text{OTS.Verify}(vk^*, ct', sig') = 1$. Thus $\mathcal{B}_{\text{OTS}}^2$ can output (ct', sig') as a forgery.

□

7 References

- [BCHK07] Dan Boneh, Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. *SIAM Journal on Computing*, 36(5):1301–1328, 2007.
- [CFNP00] B. Chor, A. Fiat, M. Naor, and B. Pinkas. Tracing traitors. *IEEE Transactions on Information Theory*, 46(3):893–910, 2000.
- [CHN⁺16] Aloni Cohen, Justin Holmgren, Ryo Nishimaki, Vinod Vaikuntanathan, and Daniel Wichs. Watermarking cryptographic capabilities. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing, STOC '16*, page 1115–1127, New York, NY, USA, 2016. Association for Computing Machinery.
- [DDN91] Danny Dolev, Cynthia Dwork, and Moni Naor. Non-malleable cryptography. In *Proceedings of the twenty-third annual ACM symposium on Theory of computing*, pages 542–552, 1991.
- [GGH⁺16] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. *SIAM Journal on Computing*, 45(3):882–929, 2016.
- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, aug 1986.
- [GKM⁺19] Rishab Goyal, Sam Kim, Nathan Manohar, Brent Waters, and David J Wu. Watermarking public-key cryptographic primitives. In *Advances in Cryptology–CRYPTO*

2019: *39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2019, Proceedings, Part III 39*, pages 367–398. Springer, 2019.

- [GKP⁺13] Shafi Goldwasser, Yael Kalai, Raluca Ada Popa, Vinod Vaikuntanathan, and Nickolai Zeldovich. Reusable garbled circuits and succinct functional encryption. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 555–564, 2013.
- [GKWW21] Rishab Goyal, Sam Kim, Brent Waters, and David J Wu. Beyond software watermarking: traitor-tracing for pseudorandom functions. In *Advances in Cryptology–ASIACRYPT 2021: 27th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 6–10, 2021, Proceedings, Part III 27*, pages 250–280. Springer, 2021.
- [KN22] Fuyuki Kitagawa and Ryo Nishimaki. Watermarking prfs against quantum adversaries. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology – EUROCRYPT 2022*, pages 488–518, Cham, 2022. Springer International Publishing.
- [KW17] Sam Kim and David J. Wu. Watermarking cryptographic functionalities from standard lattice assumptions. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology – CRYPTO 2017*, pages 503–536, Cham, 2017. Springer International Publishing.
- [MW22] Sarasij Maitra and David J. Wu. Traceable prfs: Full collusion resistance and active security. In Goichiro Hanaoka, Junji Shikata, and Yohei Watanabe, editors, *Public-Key Cryptography – PKC 2022*, pages 439–469, Cham, 2022. Springer International Publishing.
- [Nis13] Ryo Nishimaki. How to watermark cryptographic functions. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology – EUROCRYPT 2013*, pages 111–125, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [Nis20] Ryo Nishimaki. Equipping public-key cryptographic primitives with watermarking (or: A hole is to watermark). In Rafael Pass and Krzysztof Pietrzak, editors, *Theory of Cryptography*, pages 179–209, Cham, 2020. Springer International Publishing.
- [NY90] Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, pages 427–437, 1990.
- [PS19] Chris Peikert and Sina Shiehian. Noninteractive zero knowledge for np from (plain) learning with errors. In *Annual International Cryptology Conference*, pages 89–114. Springer, 2019.
- [Wat24] Brent Waters. A new approach for non-interactive zero-knowledge from learning with errors. *Cryptology ePrint Archive, Paper 2024/340*, 2024. <https://eprint.iacr.org/2024/340>.

A Watermarkable CCA-secure PKE from Watermarkable CPA-secure PKE and Statistically Simulation-Sound NIZK

In this section we present a watermarkable CCA2-secure PKE, based on a modification of the [NY90] scheme.

A.1 Preliminaries: Statistically Simulation Sound NIZK Proof for NP

A statistically simulation sound NIZK proof for language $L \in \text{NP}$ with relation R_L consists of the following efficient algorithms.

- $\text{NIZK.Setup}(1^\lambda) \rightarrow (\text{CRS}, \text{td})$: On input the security parameter 1^λ , the setup returns a common reference string CRS and a trapdoor td.
- $\text{NIZK.Prove}(\text{CRS}, w, x) \rightarrow \pi$: On input a common reference string CRS, a classical witness w , and a statement x , the proving algorithm returns a proof π .
- $\text{NIZK.Ver}(\text{CRS}, \pi, x) \rightarrow 0/1$: On input a common reference string CRS, a proof π , and a statement x , the verification algorithm returns a bit in $\{0, 1\}$.

A SSS-NIZK for NP scheme should satisfy the following properties:

Correctness A NIZK proof (NIZK.Setup , NIZK.Prove , NIZK.Ver) is correct if there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, all $x \in L$, and all $w \in \mathcal{R}_L(x)$ it holds that

$$\Pr[\text{NIZK.Ver}(\text{CRS}, \text{NIZK.Prove}(\text{CRS}, w, x), x) = 1] = 1 - \text{negl}(\lambda)$$

where $\text{CRS} \leftarrow \text{NIZK.Setup}(1^\lambda)$.

Statistical Soundness A NIZK proof (NIZK.Setup , NIZK.Prove , NIZK.Ver) is computationally sound if there exist a negligible function $\text{negl}(\cdot)$ such that for all unbounded adversaries \mathcal{A} and all $x \notin L$, it holds that:

$$\Pr[\text{NIZK.Ver}(\text{CRS}, \pi \leftarrow \mathcal{A}(\text{CRS}, x), x) = 1] = \text{negl}(\lambda)$$

where $\text{CRS} \leftarrow \text{NIZK.Setup}(1^\lambda)$.

Computational Zero Knowledge A NIZK proof (NIZK.Setup , NIZK.Prove , NIZK.Ver) is computationally zero-knowledge if there exists a PPT simulator Sim such that for all non-uniform PPT adversaries, all statements $x \in L$ and all witnesses $w \in \mathcal{R}_L(x)$, it holds that $\text{Sim}(1^\lambda, \text{CRS}, \text{td}, x) \approx_c \text{NIZK.Prove}(\text{CRS}, w, x)$ where $\text{CRS} \leftarrow \text{NIZK.Setup}(1^\lambda)$.

Statistically Simulation Soundness A NIZK proof satisfies statistically simulation soundness if it is infeasible to convince an honest verifier of a false statement even when the adversary itself is provided with a simulated proof.

Formally, for all statements x and all (even unbounded) adversaries $\mathcal{A} = (\text{Sim}_1, \text{Sim}_2)$, there exists a negligible function $\text{negl}(\cdot)$:

$$\Pr \left[(\text{CRS}, \text{td}) \leftarrow \text{Sim}_1(1^\lambda, x), \pi \leftarrow \text{Sim}_2(x, \text{CRS}, \text{td}) : \begin{array}{c} \exists(x', \pi') \wedge x' \neq x \wedge x' \in L_{no} \\ \text{and} \\ V(\text{CRS}, x', \pi') = 1 \end{array} \right] \leq \text{negl}(\cdot)$$

where td is a trapdoor generated by the simulator along with the simulated CRS.

Instantiation A SSS-NIZK proof system can be realized from a standard NIZK proof system [GGH⁺16], which can be built on the hardness of LWE [PS19, Wat24].

A.2 Construction and Security

Theorem A.1. *Assuming the hardness of LWE, there exists secure watermarkable implementation of CCA-secure PKE with private extraction and collusion resistance.*

Lemma A.2. *Assuming the security of watermarkable implementation of a CPA-secure PKE $wPKE = (WMS\text{Setup}, \text{Enc}, \text{Dec}, \text{Mark}, \text{Extract})$ and statistically simulation sound NIZK scheme $NIZK = (\text{Setup}, \text{Prove}, \text{Verify})$ (without watermarking), there exists a secure watermarkable implementation of CCA(2)-secure PKE.*

The two building blocks: watermarkable implementation of a CPA-secure PKE $wPKE = (WMS\text{Setup}, \text{Enc}, \text{Dec}, \text{Mark}, \text{Extract})$ and a SSS-NIZK scheme $NIZK = (\text{Setup}, \text{Prove}, \text{Verify})$ can both be obtained from LWE.

Construction Given a watermarkable implementation of a CPA-secure PKE $wPKE = (WMS\text{Setup}, \text{Enc}, \text{Dec}, \text{Mark}, \text{Extract})$ and a SSS-NIZK scheme $NIZK = (\text{Setup}, \text{Prove}, \text{Verify})$, the construction of watermarkable implementation of CCA-secure PKE is as follows.

$\text{KeyGen}(1^\lambda, 1^n)$: Compute $(wPKE.pk_1, wPKE.sk_1, wPKE.xk_1, wPKE.mk_1) \leftarrow wPKE.WMS\text{Setup}(\lambda)$;
 $(wPKE.pk_2, wPKE.sk_2, wPKE.xk_2, wPKE.mk_2) \leftarrow wPKE.WMS\text{Setup}(\lambda)$;
 Compute $(\text{CRS}, \text{td}) \leftarrow NIZK.\text{Setup}(1^\lambda)$; Output $\text{sk} = (wPKE.sk_i)_{i \in \{1,2\}}$; $\text{pk} = (\{wPKE.pk_i\}_{i \in \{1,2\}}, \text{CRS})$; $\text{xk} = (\{wPKE.xk_i\}_{i \in \{1,2\}}, \text{td})$; $\text{mk} = \{wPKE.mk_i\}_{i \in \{1,2\}}$

$\text{Enc}(\text{pk}, m)$:

- parse $\text{pk} := \text{pk}_1, \text{pk}_2, \text{CRS}$;
- compute $\text{ct}_1 \leftarrow wPKE.\text{Enc}(\text{pk}_1, m)$; $\text{ct}_2 \leftarrow wPKE.\text{Enc}(\text{pk}_2, m)$;
- compute $\pi \leftarrow NIZK.\text{Prove}(\text{CRS}, (\text{ct}_1, \text{ct}_2), (r_1, r_2, m))$ for the following statement:
 \exists witness (r_1, r_2, m) such that $\text{ct}_i = wPKE.\text{Enc}(\text{pk}_i, m; r_i)$ for all $i \in [2]$, where r_i is the randomness used in encryption.
- Output $\text{ct} = (\{\text{ct}_i\}_{i \in [2]}, \pi)$.

$\text{Dec}(\text{sk}, \text{ct})$:

- parse $\text{ct} = (\text{ct}_1, \text{ct}_2, \pi)$; $\text{sk} = (\text{sk}_1, \text{sk}_2)$;
- if $NIZK.\text{Verify}(\text{CRS}, \pi, (\text{ct}_1, \text{ct}_2)) = 1$, continue; else abort and output \perp .
- compute $m \leftarrow wPKE.\text{Dec}(\text{sk}_1, \text{ct}_1)$; output m' .

$\text{Mark}(\text{mk}, \text{sk}, \tau)$: parse $\text{mk} = (\text{mk}_1, \text{mk}_2)$; $\text{sk} = (\text{sk}_1, \text{sk}_2)$; output $(\text{sk}_{\tau,1} \leftarrow wPKE.\text{Mark}(\text{sk}_1, \tau); \text{sk}_{\tau,2} \leftarrow wPKE.\text{Mark}(\text{sk}_2, \tau))$;

$\text{Extract}(\text{xk}, \text{pk}, \text{aux}, C)$:

- parse $\text{xk} := (\text{xk}_1, \text{xk}_2, NIZK.\text{td})$; $\text{pk} := (\text{pk}_1, \text{pk}_2, \text{CRS})$. Initialize an empty tuple \vec{r} ;
- For $i = 1, 2$: Create the following circuit C_i with black-box access to C :
 - C_i is hardcoded with $(\text{pk}_1, \text{pk}_2, \text{xk}_{j \neq i}, \text{td}, \text{CRS})$ and simulates the CCA – PKE game for C as follows:
 - For C 's decryption queries $\text{ct} = (\text{ct}_1, \text{ct}_2, \pi)$:
 - * First check if $NIZK.\text{Verify}(\text{CRS}, \pi) = 1$, if 0 output \perp ; else continue;

- * By the extraction key simulation property, since $xk_{j \neq i}$ can be used to simulate the oracles used in $G_{CPA}(pk_j, sk_j, \cdot)$, C_i can simulate the oracle $wPKE.Dec(sk_j, \cdot)$ and thus decrypt ct_j in the ciphertext ct to output m .
- C submits challenge messages (m_0, m_1) ; C_i submits (m_0, m_1) to the external challenger; C_i receives challenge ciphertext $ct_i^* = wPKE.Enc(pk_i, m_b), b \leftarrow \{0, 1\}$ from the challenger.
 C_i prepares the following ciphertext: compute $ct_j \leftarrow wPKE.Enc(pk_j, m_{b_j}), b_j \leftarrow \{0, 1\}$;
compute $\hat{\pi} \leftarrow \text{Sim}(td, CRS, (ct_1^*, ct_2^*))$ where Sim is the simulator algorithm for NIZK. Then it sends $ct^* = (ct_1^*, ct_2^*, \hat{\pi})$ to C .
- C continues to simulate the decryption oracle as above to decrypt only valid ciphertexts $ct \neq ct^*$.
- In the end, C_i outputs the same as C outputs.
- Add $\tau/\perp \leftarrow wPKE.Extract(wPKE.xk_i, wPKE.pk_i, aux = \perp, C_i)$ to $\vec{\tau}$.
- Output $\vec{\tau}$.

Proof. We make a few claims for the proof of unremovability.

First, we consider a hybrid game where stage-1 (before the unremovability adversary \mathcal{A} produces program C) is the same, but the stage-2 game for CCA – PKE where the program C plays in is different. We call this stage-2 game $G_{CCA, H_1}(sk, pk, \cdot)$, defined as follows:

- On input $(sk, pk) = (pk_1, pk_2, sk_1, sk_2, td, CRS)$; the challenger plays the CCA – PKE game for C as follows:
- For C 's decryption queries $ct = (ct_1, ct_2, \pi)$:
 - First check if $NIZK.Verify(CRS, \pi) = 1$, if 0 output \perp ; else continue; output $m' \leftarrow wPKE.Dec(sk_1, ct_1)$.
- C submits challenge messages (m_0, m_1)
The challenger prepares the following ciphertext: compute $ct_i^* \leftarrow wPKE.Enc(pk_i, m_b)$ for all $i = 1, 2, b \leftarrow \{0, 1\}$;
compute $\hat{\pi} \leftarrow \text{Sim}(td, CRS, (ct_1^*, ct_2^*))$ where Sim is the simulator algorithm for NIZK.
Then it sends $ct^* = (ct_1^*, ct_2^*, \hat{\pi})$ to \mathcal{A} and C outputs a guess b' .
- continues to simulate the decryption oracle as above to decrypt only valid ciphertexts $ct \neq ct^*$.

Claim A.3. *Assuming the computational zero-knowledge property of NIZK, for any admissible PPT C , if C is a γ -good program in the original game $G_{CCA-PKE}(sk, pk, \cdot)$, then C is a $(\gamma - \text{negl}(\lambda))$ -good program in the above $G_{CCA, H_1}(sk, pk, \cdot)$ defined.*

Proof. We refer the definition of $G_{CCA-PKE}(sk, pk, C)$ to Section 6.1.

Let view_{CCA, H_1} (see Definition 3.8 for definition) be the transcript output by the above game.

It is to see that the only difference in $G_{CCA, H_1}(sk, pk, C)$ and $G_{CCA-PKE}(sk, pk, C)$ is in the generation of the proof π : $G_{CCA-PKE}(sk, pk, C)$ the proof is generated honestly using NIZK. Prove and in $G_{CCA, H_1}(sk, pk, C)$ it is generated by $\text{Sim}(td, \cdot)$. By the computational zero knowledge property of NIZK, $\text{view}_{G_{CCA-PKE}}$ and view_{CCA, H_1} are computationally indistinguishable. Otherwise we can build a distinguisher, given NIZK.CRS, that samples $WMS\text{Setup}$ on its own and use C to break the computational zero knowledge property.

Therefore, the output distributions of any admissible PPT \mathcal{A} in $G_{CCA,H_1}(\text{sk}, \text{pk}, C)$ and $G_{CCA-PKE}(\text{sk}, \text{pk}, C)$ must be computationally indistinguishable. Thus, any γ -good C in $G_{CCA,H_1}(\text{sk}, \text{pk}, \cdot)$ must be $(\gamma - \text{negl}(\lambda))$ -good in $G_{CCA-PKE}(\text{sk}, \text{pk}, \cdot)$ \square

Next, we consider a next stage-2 hybrid game $G_{CCA,H_2}(\text{sk}, \text{pk}, \cdot)$ (note that the stage before \mathcal{A} outputs C is still the same:

- On input $(\text{sk}, \text{pk}) = (\text{pk}_1, \text{pk}_2, \text{sk}_1, \text{sk}_2, \text{td}, \text{CRS})$;
- For C 's decryption queries $\text{ct} = (\text{ct}_1, \text{ct}_2, \pi)$:
 - First check if $\text{NIZK.Verify}(\text{CRS}, \pi) = 1$, if 0 output \perp ; else continue; output $m' \leftarrow \text{wPKE.Dec}(\text{sk}_1, \text{ct}_1)$.
- C submits challenge messages (m_0, m_1)
The challenger prepares the following ciphertext: compute $\text{ct}_1^* \leftarrow \text{wPKE.Enc}(\text{pk}_1, m_{b_1})$, $b_1 \leftarrow \{0, 1\}$; compute $\text{ct}_2^* \leftarrow \text{wPKE.Enc}(\text{pk}_2, m_{b_2})$ $b_2 \leftarrow \{0, 1\}$
compute $\hat{\pi} \leftarrow \text{Sim}(\text{td}, \text{CRS}, (\text{ct}_1^*, \text{ct}_2^*))$ where Sim is the simulator algorithm for NIZK.
Then it sends $\text{ct}^* = (\text{ct}_1^*, \text{ct}_2^*, \hat{\pi})$ to \mathcal{A} and C outputs a guess b' .
- continues to simulate the decryption oracle as above to decrypt only valid ciphertexts $\text{ct} \neq \text{ct}^*$.

Claim A.4. For any admissible PPT C , if C is a γ -good program in the game $G_{CCA,H_1}(\text{sk}, \text{pk}, \cdot)$, then C is a $\gamma/2$ -good program in the above $G_{CCA,H_2}(\text{sk}, \text{pk}, \cdot)$ defined.

Proof. The only difference between $G_{CCA,H_1}(\text{sk}, \text{pk}, \cdot)$ and $G_{CCA,H_2}(\text{sk}, \text{pk}, \cdot)$ is that the bits b_i used to decide which of m_0, m_1 to encrypt for the challenge ciphertext is the same for $G_{CCA,H_1}(\text{sk}, \text{pk}, \cdot)$ and independent for $G_{CCA,H_2}(\text{sk}, \text{pk}, \cdot)$. When $b_1 = b_2$, the two games are the same. Therefore, we have that C is a $\gamma/2$ -good program in the above $G_{CCA,H_2}(\text{sk}, \text{pk}, \cdot)$ given that C is a γ -good program in the above $G_{CCA,H_1}(\text{sk}, \text{pk}, \cdot)$. \square

Next we show that for any γ -unremovability adversary \mathcal{A} where the output program C is γ -good in $G_{CCA,H_2}(\text{sk}, \text{pk}, \cdot)$, \mathcal{A} can be used to build a reduction to γ -unremovability of CPA-security of for wPKE.

Claim A.5. For any γ -unremovability adversary \mathcal{A} where the output program C is γ -good in $G_{CCA-PKE}(\text{sk}, \text{pk}, \cdot)$, \mathcal{A} can be used to build a reduction to $(\gamma/2 - \text{negl}(\lambda))$ -unremovability of CPA-security of for wPKE.

Proof. By Claim A.3 and Claim A.4, we know that if there is a γ -unremovability adversary \mathcal{A} for Section 6.1, then if we put its output circuit C into our $\text{Extract}(\text{xk}, \text{pk}, \text{aux}, C)$ procedure, the game simulated by Extract , when $i = 2$ is exactly the same as $G_{CCA,H_2}(\text{sk}, \text{pk}, \cdot)$, i.e. the NIZK proof in the challenge ciphertext is generated by simulator and the bits for encryption are independent.

Thus, for any C produced by a winning \mathcal{A} , if C is γ -good by definition in $G_{CCA-PKE}(\text{sk}, \text{pk}, \cdot)$, C will be a $(\gamma - \text{negl}(\lambda))/2$ -good program during the execution of $\text{Extract}(\text{xk}, \text{pk}, C)$ for $i = 2$.

Also, by the statistical soundness and statistical simulation soundness of NIZK, no adversary \mathcal{A} and its output program C will make a query on a ct that contains a false proof π' , that will pass the verification, even after C has seen the simulated proof $\hat{\pi}$, except with negligible probability.

Therefore we can say that with overwhelming probability, all queries made in the unremovability game (including execution of Extract contain valid proofs.

Now we can create a reduction \mathcal{B}_2 to break the $\gamma/2 - \text{negl}(\lambda)$ -unremovability of CPA-secure PKE (of sk_2). \mathcal{B}_2 simulates the marking oracle for \mathcal{A} by querying the marking oracle for wPKE.sk_1 . \mathcal{B}_2 also samples wPKE.sk_1 on its own.

After \mathcal{A} outputs C , \mathcal{B}_2 created program C'_2 that works exactly as the circuit C_2 in our $\text{Extract}(\text{xk}, \text{pk}, C)$ algorithm when $i = 2$, except that C'_2 can hardcode the secret key sk_1 sampled by \mathcal{B}_2 to simulate the decryption oracle.

Also, since we have that $\Pr[\text{Extract}(\text{xk}, \text{pk}, C) \notin \mathcal{Q}] \geq \epsilon$ for some non-negligible ϵ , we must have that $\Pr[\text{wPKE.Extract}(\text{xk}_2, \text{pk}_2, C'_2) \notin \mathcal{Q}] \geq \epsilon$ for some non-negligible ϵ by the design of our Extract algorithm. □

We next analyze a symmetric case, consider the following game $G_{\text{CCA}, H_3}(\text{sk}, \text{pk}, \cdot)$.

First, before \mathcal{A} outputs program C , we switch to answering \mathcal{A} 's decryption queries using $\text{wPKE.Dec}(\text{sk}_2, \cdot)$.

Note that the underlined differences are between *game* $G_{\text{CCA}, H_1}(\text{sk}, \text{pk}, \cdot)$ and *game* $G_{\text{CCA}, H_3}(\text{sk}, \text{pk}, \cdot)$.

- On input $(\text{sk}, \text{pk}) = (\text{pk}_1, \text{pk}_2, \text{sk}_1, \text{sk}_2, \text{td}, \text{CRS})$;
- For C 's decryption queries $\text{ct} = (\text{ct}_1, \text{ct}_2, \pi)$:
 - First check if $\text{NIZK.Verify}(\text{CRS}, \pi) = 1$, if 0 output \perp ; else continue; output $m' \leftarrow \text{wPKE.Dec}(\text{sk}_2, \text{ct}_1)$.
- C submits challenge messages (m_0, m_1)
The challenger prepares the following ciphertext: compute $\text{ct}_1^* \leftarrow \text{wPKE.Enc}(\text{pk}_1, m_{b_1})$,
 $b_1 \leftarrow \{0, 1\}$; compute $\text{ct}_2^* \leftarrow \text{wPKE.Enc}(\text{pk}_2, m_{b_2})$ $b_2 \leftarrow \{0, 1\}$
compute $\hat{\pi} \leftarrow \text{Sim}(\text{td}, \text{CRS}, (\text{ct}_1^*, \text{ct}_2^*))$ where Sim is the simulator algorithm for NIZK.
Then it sends $\text{ct}^* = (\text{ct}_1^*, \text{ct}_2^*, \hat{\pi})$ to \mathcal{A} and C outputs a guess b' .
- continues to simulate the decryption oracle as above to decrypt only valid ciphertexts $\text{ct} \neq \text{ct}^*$.

Note that the only difference is that we now answer decryption queries of C using sk_2 . This results in no difference on the adversary's view because by the statistical soundness of NIZK, all ciphertexts submitted to the decryption oracle have their ct_1, ct_2 encrypt the same message or will result in \perp replies.

By the exact (but symmetric) analysis as Claim A.5, we can do a reduction to break the $(\gamma/2 - \text{negl}(\lambda))$ -unremovability of wPKE where the reduction is challenged with CPA-security game under $(\text{pk}_1, \text{sk}_1, \text{mk}_1, \text{xk}_1)$ and samples $(\text{pk}_2, \text{sk}_2, \text{mk}_2, \text{xk}_2)$ on its own. □

B Watermarkable Weak PRP from Watermarkable weak PRF

In this section, we present a watermarkable weak PRP built using the two-round Feistel network from a watermarkable weak PRF.

A weak PRP of input-output space $\{0, 1\}^\ell$ consists of algorithms:

$\text{KeyGen}(1^\lambda) \rightarrow \text{sk}$: a randomized algorithm that generates a secret key sk .

$\text{Eval}(\text{sk}, x \in \{0, 1\}^\ell) \rightarrow y \in \{0, 1\}^\ell$: a deterministic evaluation algorithm that on input sk, x , outputs y .

The other syntax and definitions (such as correctness, functionality preserving) are the same as watermarkable weak PRF and we omit them here.

γ -Unremovability of weak PRP The γ -Unremovability for a watermarkable implementation of a weak PRF scheme says, for all PPT admissible stateful adversary \mathcal{A} , there exists a negligible function $\text{negl}(\lambda)$ such that for all $\lambda \in \mathbb{N}$:

$$\Pr \left[\text{Extract}(\text{xk}, C) \notin \mathcal{Q} \wedge C \text{ is } \gamma\text{-good} : \begin{array}{l} (\text{sk}, \text{xk}, \text{mk}) \leftarrow \text{WMSetup}(1^\lambda) \\ C \leftarrow \mathcal{A}^{\text{Mark}(\text{mk}, \text{sk}, \cdot)}(1^\lambda) \end{array} \right] \leq \text{negl}(\lambda).$$

\mathcal{Q} is the set of marks queried by \mathcal{A} and C is said to be a PPT admissible, stateful γ -good circuit if:

$$\Pr \left[C^{\text{Eval}(\text{sk}, \cdot)}(x, y_b) = b : x \leftarrow \{0, 1\}^\ell, y_0 = \text{Eval}(\text{sk}, x), y_1 \leftarrow \{0, 1\}^\ell, b \leftarrow \{0, 1\} \right] \geq \frac{1}{2} + \gamma.$$

$\text{Eval}(\cdot)$ samples a uniformly random $x \leftarrow \{0, 1\}^\ell$ upon every query and outputs $(x, \text{Eval}(\text{sk}, x))$.

B.1 Construction and Security

Given a watermarkable implementation of a weak PRF $\text{wPRF} = (\text{WMSetup}, \text{Eval}, \text{Mark}, \text{Extract})$ where the input and output space of Eval is $\{0, 1\}^{\ell/2}$, we give the construction based on a 2-round Feistel network as below:

$\text{WMSetup}(1^\lambda)$: compute $(\text{sk}_1, \text{xk}_1, \text{mk}_1) \leftarrow \text{wPRF.WMSetup}(1^\lambda)$; $(\text{sk}_2, \text{xk}_2, \text{mk}_2) \leftarrow \text{wPRF.WMSetup}(1^\lambda)$;
output $\text{sk} = (\text{sk}_1, \text{sk}_2)$; $\text{xk} = (\text{xk}_1, \text{xk}_2)$; $\text{mk} = (\text{mk}_1, \text{mk}_2)$.

$\text{Eval}(\text{sk}, x \in \{0, 1\}^\ell)$:

1. parse $\text{sk} = (\text{sk}_1, \text{sk}_2)$; parse $x = x_0 \| x_1$ where $x_0, x_1 \in \{0, 1\}^{\ell/2}$.
2. for $i = 1, 2$:
 $x_{i+1} = \text{wPRF.Eval}(\text{sk}_i, x_i) \oplus x_{i-1}$
3. output $x_2 \| x_3$

$\text{Mark}(\text{mk}, \text{sk}, \tau)$: parse $\text{sk} = (\text{sk}_1, \text{sk}_2)$; $\text{mk} = (\text{mk}_1, \text{mk}_2)$;

Output $(\tilde{\text{sk}}_1 \leftarrow \text{wPRF.Mark}(\text{sk}_1, \tau), \tilde{\text{sk}}_2 \leftarrow \text{wPRF.Mark}(\text{sk}_2, \tau))$.

$\text{Extract}(\text{xk}, C)$:

1. parse $\text{xk} = (\text{xk}_1, \text{xk}_2)$;
2. For $i = 1$: create the following circuit C_1
 - (a) C_1 is hardcoded with xk_2 has black-box access to C . C_1 answers C 's queries using xk_2 and external queries:
 - i. Upon C 's query: query external oracle $\mathcal{O}(\text{sk}_1, \cdot)$ which will return $(x_1 \leftarrow \{0, 1\}^{\ell/2}, \text{wPRF.Eval}(\text{sk}_1, x_1))$;
 - ii. sample $(x_2 \leftarrow \{0, 1\}^{\ell/2}, \text{wPRF.Eval}(\text{sk}_2, x_2))$ using xk_2 . Let $x_0 := \text{wPRF.Eval}(\text{sk}_1, x_1) \oplus x_2$ and $x_3 = \text{wPRF.Eval}(\text{sk}_2, x_2) \oplus x_1$.
 - iii. output $(x_0 \| x_1, x_3 \| x_2)$.

- (b) In the challenge phase, receive the challenge (x_1^*, y_1^*) from the external challenger; compute challenge input for C : $(x_0^* \| x_1^*, x_3^* \| x_2^*)$ the same way as above queries.
 - (c) C outputs a guess b' and C_1 outputs the same.
3. compute $\tau_1 / \perp \leftarrow \text{wPRF.Extract}(xk_1, C_1)$;
 4. For $i = 2$: create the following circuit C_2
 - (a) C_2 is hardcoded with xk_1 has black-box access to C . C_2 answers C 's queries using xk_1 and external queries:
 - i. Upon C 's query: query external oracle $\mathcal{O}(\text{sk}_2, \cdot)$ which will return $(x_2 \leftarrow \{0, 1\}^{\ell/2}, \text{wPRF.Eval}(\text{sk}_2, x_2))$;
 - ii. sample $(x_1 \leftarrow \{0, 1\}^{\ell/2}, \text{wPRF.Eval}(\text{sk}_1, x_1))$ using xk_1 . Let $x_0 := \text{wPRF.Eval}(\text{sk}_1, x_1) \oplus x_2$ and $x_3 = \text{wPRF.Eval}(\text{sk}_2, x_2) \oplus x_1$.
 - iii. output $(x_0 \| x_1, x_3 \| x_2)$.
 - (b) In the challenge phase, receive the challenge (x_2^*, y_2^*) from the external challenger; compute challenge input for C : $(x_0^* \| x_1^*, x_3^* \| x_2^*)$ the same way as above queries.
 - (c) C outputs a guess b' and C_2 outputs the same.
 5. compute $\tau_2 / \perp \leftarrow \text{wPRF.Extract}(xk_2, C_2)$;
 6. output (τ_1, τ_2) (which can include \perp 's).

Proof. The unremovability proof is relatively straightforward given the framework of proof in Section 4.2. We will make it brief. Suppose a given C satisfies that $\Pr[\text{Extract}(xk, C) \notin \mathcal{Q}] \geq \epsilon$, then it must be that $\Pr[\text{wPRF.Extract}(xk_i, C_i) \notin \mathcal{Q}]$ for both $i = 1, 2$.

Since the extraction keys of the wPRF scheme can be used to simulate the PRF pseudorandomness game perfectly (see Section 5.4.1), the program C_i built in the Extract algorithm is a stage-2 reduction from weak PRP to weak PRF (with key sk_i). Therefore, for each $i = 1, 2$, we can build a reduction \mathcal{B}_i that samples the key $\text{sk}_{j \neq i, j \in [2]}$ and simulates the $\text{Mark}(\text{mk}, \text{sk}, \cdot)$ queries from \mathcal{A} making queries to the unremovability challenger of wPRF with key sk_i .

After \mathcal{A} outputs program C , \mathcal{B}_i makes a program $C_i(\text{sk}_j)'$ that uses black-box access to C : $C_i(\text{sk}_j)'$ behaves the same as C_i built in Extract except using the real secret key sk_j instead of the extraction key xk_j to simulate the game for C .

By the reduction property, either C_1 or C_2 must be a γ_i -good program for winning the weak pseudorandomness game of weak PRP for some non-negligible γ_i : if C_1 has only negligible advantage, then we can replace $\text{wPRF.Eval}(\text{sk}_1, \cdot)$ in the evaluation algorithm with a real random function (or always give the challenge point evaluation $\text{wPRF.Eval}(\text{sk}_1, x_1^*)$ with a random value) and C should still be $(\gamma - \text{negl}(\lambda))$ -good. Then we use C in a way where $\text{wPRF.Eval}(\text{sk}_1, \cdot)$ is replaced with random to build a program C_2 : if C_2 is also not γ_2 -good for any non-negligible γ_2 , then we must be able to replace $\text{wPRF.Eval}(\text{sk}_2, \cdot)$ with a random function while C should still be $(\gamma - \text{negl}(\lambda))$ -good. But now since all evaluations are random, C should have no advantage. Contradiction. □

C Watermarkable Implementation of Attribute-based Encryption

In this section, we give a modified version of the watermarkable ABE scheme in [GKM⁺19] which will imply a watermarkable implementation of the IBE scheme in Section 6.1.1.

C.1 Watermarkable Attribute-based Encryption

Syntax A watermarkable implementation of ABE consists of the following algorithms:

$\text{WMSetup}(1^\lambda) \rightarrow (\text{mk}, \text{xk}, \text{msk}, \text{mpk})$: on the security parameter, outputs a master public/secret key, a marking key, an extraction key.

$\text{KeyGen}(\text{msk}, f, \text{mode} \in \{\text{Marked}, \text{Unmarked}\}) \rightarrow \text{sk}_f$ on master secret key msk , a policy f , and a mode $\text{mode} \in \{\text{Marked}, \text{Unmarked}\}$ outputs a secret key sk_f .

$\text{Enc}(\text{mpk}, x, m) \rightarrow \text{ct}$: on master public key mpk , attribute $x \in \mathcal{X}$, a message $m \in \mathcal{M}$, outputs a ciphertext ct .

$\text{Dec}(\text{sk}_f, \text{ct}) \rightarrow m'/\perp$: on secret key sk_f and ciphertext ct , output a message $m \in \mathcal{M}$ or \perp .

$\text{Mark}(\text{mk}, \text{sk}, \tau, \text{mode} \in \{\text{MSK}, \text{SKF}\})$: on marking key mk and secret key sk , a message $\tau \in \mathcal{M}_\tau$, and a mode $\text{mode} \in \{\text{MSK}, \text{SKF}\}$, output a marked key sk_τ .

$\text{Extract}(\text{xk}, \text{mpk}, \text{aux}, C)$: on input an extraction key xk , master public key mpk , and program C , output a mark $\tau \in \mathcal{M}_\tau/\perp$.

Remark C.1. The definition in [GKM⁺19] only considers running KeyGen on an unmarked master secret key and marking the policy-embedded key sk_f .

For the use of our watermarkable IBE scheme, we consider marking the master secret key and allowing KeyGen to first "mark" a master secret key.

We additionally allow computing the functional key $\text{sk}_{f,\tau}$ from a marked master secret key msk_τ . When running on Marked mode, the KeyGen algorithm also takes in a symbol τ .

Note that not giving out $\text{KeyGen}(\text{msk}, \cdot)$ oracle to \mathcal{A} , but only $\text{Mark}(\text{msk}, \cdot)$ suffices for our use of watermarkable IBE in Section 6. Here we simply prove a slightly stronger security.

Correctness There exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, all $x \in \mathcal{X}$, $f, g \in \mathcal{F}$, $m \in \mathcal{M}$, when $f(x) = 1$ the following holds:

$$\Pr \left[\text{Dec}(\text{sk}_f, \text{ct}) = m : \begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{WMSetup}(1^\lambda), \text{sk}_f \leftarrow \text{KeyGen}(\text{msk}, f) \\ \text{ct} \leftarrow \text{Enc}(\text{mpk}, x, m) \end{array} \right] \geq 1 - \text{negl}(\lambda)$$

Definition C.2 (ABE security). The standard notion of security for a KP-ABE scheme is that of full or adaptive security. Specifically, a key-policy attribute-based encryption scheme is said to be fully secure if for every stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$, such that for every $\lambda \in \mathbb{N}$ the following holds:

$$\Pr \left[\mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{ct}) = b : \begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{WMSetup}(1^\lambda) \\ ((m_0, m_1), x) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(1^\lambda, \text{mpk}) \\ b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{Enc}(\text{mpk}, x, m_b) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

\mathcal{A} is admissible if all query $f \in \mathcal{F}$ made to oracle $\text{KeyGen}(\text{msk}, \cdot)$ satisfies $f(x) = 0$.

It is selectively secure if \mathcal{A} needs to output x before seeing mpk .

Definition C.3 (γ -Unremovability). The γ -Unremovability for a watermarkable ABE scheme says, for all $\lambda \in \mathbb{N}$, and for all PPT admissible stateful adversary \mathcal{A} , there exists a negligible function $\text{negl}(\lambda)$ such that:

$$\Pr \left[\begin{array}{l} \text{Extract}(\text{xk}, \text{mpk}, \text{aux} = (m_0, m_1, x), C) \notin \mathcal{Q} \\ \wedge C \text{ is } \gamma\text{-good} \end{array} : \begin{array}{l} (\text{msk}, \text{mpk}, \text{xk}, \text{mk}) \leftarrow \text{WMSetup}(1^\lambda) \\ (\{m_0, m_1\}, x, C) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot), \text{Mark}(\text{mk}, \text{msk}, \cdot)}(1^\lambda, \text{mpk}) \end{array} \right] \leq \text{negl}(\lambda).$$

where \mathcal{Q} is the set of marks queried by \mathcal{A} and C is a PPT admissible, stateful γ -good adversary in the security game $G_{\text{ABE}}(\text{msk}, \text{mpk}, \text{aux} = (m_0, m_1, x), \cdot)$, more specifically:

$$\Pr \left[C^{\text{KeyGen}(\text{msk}, \cdot)}(\text{ct}) = b : \text{ct}^* \leftarrow \text{Enc}(\text{mpk}, m_b, x), b \leftarrow \{0, 1\} \right] \geq \frac{1}{2} + \gamma.$$

\mathcal{A}, C is admissible if all query $f \in \mathcal{F}$ made to oracle $\text{KeyGen}(\text{msk}, \cdot)$ satisfies $f(x) = 0$.

C will be given the KeyGen oracle as $\text{KeyGen}(\text{msk}, \cdot, \text{Unmarked})$ where it only queries on unmarked keys.

A selective variant is requiring \mathcal{A} output x before seeing mpk .

Remark C.4. As discussed in Section 5.5.2, the above definition is equivalent to letting C output (m_0, m_1) because we can view any distribution over (m_0, m_1) used by C as a convex combination of different message pairs $\{(m_0, m_1)_i\}_i$ and its corresponding strategy such that the overall winning probability is $1/2 + \gamma$. Thus, we can let \mathcal{A} pick the message-pair and corresponding strategy with the largest winning probability and hardcode them into C instead.

But x needs to be output by \mathcal{A} instead of C , because the Extract algorithm does not get to see queries of \mathcal{A} and thus cannot make sure if $f(x) = 0$ for all f queried.

Remark C.5 (Modification of [GKM⁺19] Watermarkable ABE scheme). In [GKM⁺19], the unremovability game allows \mathcal{A} to be given only one policy-embedded key sk_f and then allowed to query $\text{Mark}(\text{mk}, \text{sk}_f, \cdot)$ for polynomially many times. In our setting, we let \mathcal{A} query a marked master secret key $\text{Mark}(\text{mk}, \text{msk}, \cdot)$ and the KeyGen algorithm can derive a functional key from a "marked" master secret key. We will show how to adapt [GKM⁺19]'s construction for our use.

C.2 Preliminaries: Delegatable ABE and Mixed FE

The [GKM⁺19]'s watermarkable PKE and ABE building blocks are delegatable ABE and mixed FE.

C.2.1 Delegatable ABE

Delegatable ABE has the same syntax with ABE (see Appendix C.1, ignoring the extraction key, marking key and Extract , Mark algorithm) with an additional algorithm Delegate :

$\text{Delegate}(\text{sk}_f, g) \rightarrow \text{sk}_{f,g}$: on input sk_f and a predicate $g \in \mathcal{F}$, output a delegated key $\text{sk}_{f,g}$.

Correctness of Delegation There exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, all $x \in \mathcal{X}$, $f, g \in \mathcal{F}$, $m \in \mathcal{M}$, when $f(x) = g(x) = 1$ the following holds:

$$\Pr \left[\begin{array}{l} \text{Dec}(\text{sk}_{f,g}, \text{ct}) = m : \\ \text{ct} \leftarrow \text{Enc}(\text{mpk}, x, m) \end{array} \quad \begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{WMSetup}(1^\lambda), \text{sk}_f \leftarrow \text{KeyGen}(\text{msk}, f) \\ \text{sk}_{f,g} \leftarrow \text{Delegate}(\text{sk}_f, g) \end{array} \right] \geq 1 - \text{negl}(\lambda)$$

Delegatable ABE Security The DABE security is the same as ABE security except that the adversary \mathcal{A} queries an oracle $\mathcal{O}(\text{msk}, \cdot)$ that has several modes and takes in a query of the form $(f, \text{index}, \text{mode})$:

- If $\text{mode} = \text{Storekey}$, the challenger generates a new key $\text{sk}_f \leftarrow \text{KeyGen}(\text{msk}, f)$ and store the generated (n, f, sk_f) where n is an index; update $n := n + 1$.

- If mode = OutputKey: the challenger first checks if there exists a key tuple of the form (index, g, sk_g) . If no such tuple exists or if $g(x) = 1$, it outputs \perp . Otherwise, it replies with (index, sk_g) .
- If mode = DelegateKey, then the challenger first checks if there exists a key tuple of the form (index, g, sk_g) . If no such tuple exists or if $g(x) = f(x) = 1$, it outputs \perp . Otherwise, it generates $sk_{g,f} \leftarrow \text{Delegate}(sk_g, f)$ and replies with $(\text{index}, sk_{g,f})$.

C.2.2 Mixed Functional Encryption

An MFE scheme consists of the following algorithms:

$\text{Setup}(1^\lambda) \rightarrow (\text{mpk}, \text{msk})$: on input security parameter λ , outputs the public parameters/master public key mpk and master secret key msk .

$\text{Enc}(\text{mpk}) \rightarrow \text{ct}$: on master public key, the normal encryption outputs ciphertext ct .

$\text{SEnc}(\text{msk}, f) \rightarrow \text{ct}$: on master secret key msk and function $f \in \mathcal{F}$, outputs a ciphertext ct .

$\text{KeyGen}(\text{msk}, m) \rightarrow sk_m$: on master secret key msk and message m , outputs a key sk_m .

$\text{Dec}(sk_m, \text{ct}) \rightarrow \{0, 1\}$: on secret key sk_m and ciphertext ct , outputs a bit.

Correctness A mixed functional encryption scheme is correct if there exists a negligible function $\text{negl}(\lambda)$ such that for all $\lambda \in \mathbb{N}, f \in \mathcal{F}, m \in \mathcal{M}$:

$$\Pr \left[\text{Dec}(sk_m, \text{ct}) = 1 : \begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda), sk_m \leftarrow \text{KeyGen}(\text{msk}, m) \\ \text{ct} \leftarrow \text{Enc}(\text{mpk}) \end{array} \right] \geq 1 - \text{negl}(\lambda)$$

$$\Pr \left[\text{Dec}(sk_m, \text{ct}) = f(m) : \begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda), sk_m \leftarrow \text{KeyGen}(\text{msk}, m) \\ \text{ct} \leftarrow \text{SEnc}(\text{msk}, f) \end{array} \right] \geq 1 - \text{negl}(\lambda)$$

q -Bounded function indistinguishability Let $q = q(\lambda)$ be any fixed polynomial. A mixed functional encryption scheme is said to satisfy q -bounded function indistinguishability security if for every stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$, such that for every $\lambda \in \mathbb{N}$ the following holds:

$$\Pr \left[\mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot), \text{SEnc}(\text{msk}, \cdot)}(\text{ct}) = b : \begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda) \\ (f_0, f_1) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot), \text{SEnc}(\text{msk}, \cdot)}(1^\lambda, \text{mpk}) \\ b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{SEnc}(\text{mpk}, f_b) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

where \mathcal{A} can make at most q queries to $\text{SEnc}(\text{msk}, \cdot)$ oracle; every secret key query m made by adversary \mathcal{A} to the $\text{KeyGen}(\text{msk}, \cdot)$ oracle must satisfy $f_0(m) = f_1(m)$.

q -Bounded accept indistinguishability Let $q = q(\lambda)$ be any fixed polynomial. A mixed functional encryption scheme is said to satisfy q -bounded accept indistinguishability security if for every stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$, such that for every $\lambda \in \mathbb{N}$ the following holds:

$$\Pr \left[\mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot), \text{SEnc}(\text{msk}, \cdot)}(\text{ct}_b) = b : \begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda) \\ f^* \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot), \text{SEnc}(\text{msk}, \cdot)}(1^\lambda, \text{mpk}) \\ b \leftarrow \{0, 1\}, \text{ct}_1 \leftarrow \text{SEnc}(\text{msk}, f^*), \text{ct}_1 \leftarrow \text{Enc}(\text{mpk}) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

where \mathcal{A} can make at most q queries to $\text{SEnc}(\text{msk}, \cdot)$ oracle; every secret key query m made by adversary \mathcal{A} to the $\text{KeyGen}(\text{msk}, \cdot)$ oracle must satisfy $f^*(m) = 1$.

QTrace Jump-Finding Algorithm The [GKM⁺19] watermarkable ABE schemes Extract procedure also additionally uses an algorithm QTrace that is widely used in the traitor-tracing literature. QTrace on input parameters $(\lambda, N, q, \delta, \gamma)$ and given black box access to a program Q , runs in time $t = \text{poly}(\lambda, \log N, q, 1/\delta)$. Since detailed discussions are not essential to our presentation on the modified scheme, we refer the readers to [GKM⁺19] Section 4.2.3 for details.

C.3 Construction

Now we give the construction for watermarkable implementation of ABE, which is modified from [GKM⁺19] Section 4.3.

- Given a delegatable ABE scheme $\text{DABE} = (\text{DABE.Setup}, \text{DABE.KeyGen}, \text{DABE.Enc}, \text{DABE.Dec}, \text{DABE.Delegate})$ and a MFE scheme $(\text{MFE.Setup}, \text{MFE.KeyGen}, \text{MFE.SKEnc}, \text{MFE.Enc}, \text{MFE.Dec})$.
- Let \mathcal{M}_τ be the marked space, $\mathcal{X} = \{0, 1\}^{\ell_2}$ the attribute space, and $\mathcal{C} \subseteq \text{Functions}[\mathcal{X}, \{0, 1\}]$ be the predicate class for the target watermarkable ABE scheme, with parameters ℓ_1, ℓ_2 .
- Let $\mathcal{Y} = \{0, 1\}^{\ell_2 + \kappa}$ and $\mathcal{D} = \mathcal{C} \cup \{\text{MFE.Dec}\} \cup \{\text{MFE.Dec} \wedge C\}_{C \in \mathcal{C}}$.
- the DABE scheme has attribute class \mathcal{Y} , message space \mathcal{M} , predicate class cD . Elements in \mathcal{Y} have the following format $(x \in \mathcal{X}, \text{MFE.ct})$.
- Let γ be the unremovability parameter.

$\text{WMSetup}(1^\lambda) \rightarrow (\text{mpk}, \text{msk}, \text{mk}, \text{xk})$:

1. compute $\text{MFE.mpk}, \text{MFE.msk} \leftarrow \text{MFE.Setup}(1^\lambda)$;
2. compute $(\text{DABE.mpk}, \text{DABE.msk}) \leftarrow \text{DABE.Setup}(1^\lambda)$;
3. output $\text{mpk} = (\text{MFE.mpk}, \text{DABE.mpk})$; $\text{msk} = (\text{MFE.msk}, \text{DABE.msk})$;
 $\text{mk} = \text{xk} = (\text{MFE.mpk}, \text{MFE.msk}, \text{DABE.msk})$.

$\text{KeyGen}(\text{msk}, f, \text{mode}, \tau) \rightarrow \text{sk}_f$: Let $\text{msk} = (\text{MFE.msk}, \text{DABE.msk})$

If $\text{mode} = \text{Unmarked}$:

1. Let $\tilde{f} : \{0, 1\}^{\ell_2 + \kappa} \rightarrow \{0, 1\}$ denote the predicate $\tilde{f}(x, c) = f(x), x \in \{0, 1\}^{\ell_2}, c \in \{0, 1\}^\kappa$.
2. output $\text{sk}_f \leftarrow \text{DABE.KeyGen}(\text{msk}, \tilde{f})$.

If $\text{mode} = \text{Marked}$:

1. compute $\text{sk}_\tau \leftarrow \text{MFE.KeyGen}(\text{MFE.msk}, \tau)$
2. Let g_τ denote the mixed FE decryption circuit with sk_τ hardwired, i.e. $g_\tau = \text{MFE.Dec}(\text{sk}_\tau, \cdot)$;
3. compute $\text{sk}_{\tilde{g}_\tau} \leftarrow \text{DABE.KeyGen}(\text{DABE.msk}, \tilde{g}_\tau)$ where $\tilde{g}_\tau(x, c) = g_\tau(c), \forall x \in \{0, 1\}^{\ell_2}, c \in \{0, 1\}^\kappa$;
4. compute and output delegated key $\text{sk}_{f, \tilde{g}_\tau} \leftarrow \text{DABE.Delegate}(\text{sk}_{\tilde{g}_\tau}, f)$

$\text{Enc}(\text{mpk}, x, m)$:

Let $\text{mpk} = (\text{MFE.mpk}, \text{DABE.mpk})$; compute $\text{ct}_{\text{MFE}} \leftarrow \text{MFE.Enc}(\text{MFE.mpk})$.

Next compute ad output $\text{ct} \leftarrow \text{DABE.mpk}, (x, \text{ct}_{\text{MFE}}), m$ where $(x, \text{ct}_{\text{MFE}})$ is the attribute.

$\text{Dec}(\text{sk}, \text{ct}) \rightarrow m/\perp$: output $m \leftarrow \text{DABE.Dec}(\text{sk}, \text{ct})$.

$\text{Mark}(\text{mk}, \text{msk}, \tau)$:

1. $\text{mk} := (\text{MFE.mpk}, \text{MFE.msk})$;
2. compute $\text{sk}_\tau \leftarrow \text{MFE.KeyGen}(\text{MFE.msk})$
3. Let g_τ denote the mixed FE decryption circuit with sk_τ hardwired, i.e. $g_\tau = \text{MFE.Dec}(\text{sk}_\tau, \cdot)$;

4. compute and output $sk_{\tilde{g}_\tau} \leftarrow \text{DABE.KeyGen}(\text{DABE.msk}, \tilde{g}_\tau)$ where $\tilde{g}_\tau(x, c) = g_\tau(c), \forall x \in \{0, 1\}^{\ell_2}, c \in \{0, 1\}^\kappa$

Extract($xk, mpk, aux = (x, m_0, m_1), C, q$):

Let $xk = (\text{MFE.mpk}, \text{MFE.msk}, \text{DABE.msk}); mpk = (\text{DABE.mpk}, \text{MFE.mpk});$

The extraction algorithms runs the QTrace algorithm as $\tau \leftarrow \text{QTrace}^{Q_C}(\lambda, 2^{\ell_1}, q, \delta, \gamma)$ where q is the bounded parameter for the MFE scheme (upper bound on the number of queries to the $\text{SKEnc}(\text{msk}, \cdot)$ oracles), γ is the unremovability paramter and $\delta = \gamma/(5 + 2\ell_1q)$. Q_C is simulated as follows:

On input $\tau \leftarrow [0, 2^{\ell_1}]$:

- compute MFE ciphertext as $ct_{\text{MFE}} \leftarrow \text{MFE.SKEnc}(\text{MFE.msk}, \text{comp}_\tau)$, where comp_τ is the comparison function that on input z , output 1 if and only if $z \geq \tau$.
- sample a bit $b \leftarrow \{0, 1\}$ and compute $ct_b \leftarrow \text{DABE.Enc}(\text{DABE.mpk}, (x, ct_{\text{MFE}}), m_b)$ where (x, ct_{MFE}) is the attribute input.
- compute $b' \leftarrow C^{\text{KeyGen}(\text{msk}, \cdot, \text{Unmarked})}(ct)$ and output 1 if $b' = b$ and 0 otherwise. C has oracle access to $\text{KeyGen}(\text{msk}, \cdot, \text{Unmarked})$, which can be simulated using DABE.msk provided in xk .

Correctness and Other Properties The correctness of decryption, correctness of extraction, functionality-preserving and ABE security can directly follow the same proof as in [GKM⁺19] so we omit them here.

We can observe that the above scheme also satisfies the additional extraction key simulation property and extraction syntax also satisfy our requirement on watermarkable implementation: the extraction key has DABE.msk embedded so enough to simulate the $\text{KeyGen}(\text{msk}, \cdot, \text{Unmarked})$ queries for program C . The entire extraction algorithm follows the format that runs the simulated stage-2 security game for ABE.

Security Proof The security proof will be similar to the proof in [GKM⁺19]. Overall, our changes are minor. To avoid repetitive work of doing a same proof verbatimly, we omit the parts of our proof which are exactly the same as [GKM⁺19] and refer the readers to Section 4.3.1 of [GKM⁺19] "Proof of Theorem 4.15(Unremovability)" part. We only discuss the major changes to make to adjust to our construction:

All discussions and lemmas before Lemma 4.18 will be exactly the same when moved to our case, except some minor changes:

1. We adjust Experiment $\text{GetCircuit}_{\mathcal{A}}(\lambda)$ in Figure 2 of [GKM⁺19] to fit our unremovability game: \mathcal{A} is allowed to query $\text{KeyGen}(\text{msk}, \cdot)$ for arbitrarily polynomial times. We will discuss how this oracle is simulated when doing different reductions.
2. In Lemma 4.16 and 4.17 invokes security of MFE: in our setting, the reduction algorithm will additionally query the $\text{MFE.KeyGen}(\text{msk}, \cdot)$ oracle when dealing with the adversary \mathcal{A} as $\text{KeyGen}(\text{msk}, \cdot, \text{Marked}, \cdot)$ queries. But these queries have exactly the same format as when answering the marking queries $\text{Mark}(\text{mk}, \text{msk}, \tau)$. Moreover, in both the q -bounded function indistinguishability and accept indistinguishability games, the number of queries a reduction can make to oracle $\text{MFE.KeyGen}(\text{msk}, \cdot)$ are unbounded polynomial. Therefore, the proof of Lemma 4.16 and 4.17 are unaffected.

We mainly discuss the changes to Lemma 4.18 and take some steps from [GKM⁺19] verbatimly for completeness: Let \mathcal{B} be the reduction to security of DABE.

1. Adversary \mathcal{A} chooses a challenge attribute x and sends it over to \mathcal{B} , since we consider only selective security.
2. The DABE challenger samples a DABE key pair $(\text{DABE.mpk}, \text{DABE.msk}) \leftarrow \text{DABE.Setup}(1^\lambda)$. Algorithm \mathcal{B} samples a mixed FE key pair $(\text{MFE.mpk}, \text{MFE.msk}) \leftarrow \text{MFE.Setup}(1^\lambda)$, and sends MFE.mpk to \mathcal{B} , and sends the public key $\text{mpk} = (\text{MFE.mpk}, \text{DABE.mpk})$ to \mathcal{A} .
3. \mathcal{B} computes $\text{ct}_{\text{MFE}}^* \leftarrow \text{MFE.SKEnc}(\text{MFE.msk}, \text{comp}_{2^{\ell_1}})$ and commits to the challenge attribute $x^* = (x, \text{ct}_{\text{MFE}}^*)$.
4. When \mathcal{A} makes a query to the $\text{KeyGen}(\text{msk}, \cdot)$ and $\text{Mark}(\text{mk}, \text{msk}, \cdot)$ oracles, \mathcal{B} simulates the answer as follows:
 - (a) If \mathcal{A} makes a marking query on message τ : \mathcal{B} computes the function $g_{\tilde{\text{sk}}_\tau}$ as described in the construction on its own; using MFE.KeyGen ; then it sends query $(g_{\tilde{\text{sk}}_\tau}, \perp, \text{StoreKey})$ to $\mathcal{O}(\text{msk}, \cdot)$.
Note that $\tau \in [0, 2^{\ell_1}]$, so $g_{\tilde{\text{sk}}_\tau}(x^*) = 0$ is satisfied with overwhelming probability.
 \mathcal{B} gets reply in the form (index, \perp) ; \mathcal{B} then makes another query $(g_{\tilde{\text{sk}}_\tau}, \text{index}, \text{OutputKey})$ and gets back $(\text{index}, \text{DABE.sk}_{g_{\tilde{\text{sk}}_\tau}})$.
 \mathcal{B} makes a table that stores the following information $(\text{index}, \tau, \text{sk}_{g_{\tilde{\text{sk}}_\tau}})$ in each entry.
 - (b) If \mathcal{A} makes KeyGen query of the format (Marked, τ, f) : \mathcal{B} first checks if an entry containing τ is stored in its table in the format $(\text{index}, \tau, \text{sk}_{g_{\tilde{\text{sk}}_\tau}})$; if so \mathcal{B} prepares $\tilde{f}(x, c) = f(x)$ and queries the $\mathcal{O}(\text{msk}, \cdot)$ oracle on input $(\text{index}, f, \text{DelegateKey})$ and gets back $(\text{index}, \text{DABE.sk}_{g_{\tilde{\text{sk}}_\tau, \tilde{f}}})$. \mathcal{B} sends $\text{DABE.sk}_{g_{\tilde{\text{sk}}_\tau, \tilde{f}}}$ to \mathcal{A} .
If there is no such entry, \mathcal{B} will query $(g_{\tilde{\text{sk}}_\tau}, \perp, \text{StoreKey})$ first and then query $(\text{index}', f, \text{DelegateKey})$ to get the same output.
 - (c) If \mathcal{A} makes a KeyGen query in the format $(\text{Unmarked}, f)$, then \mathcal{B} prepares $\tilde{f}(x, c) = f(x)$ and first queries $(\tilde{f}, \perp, \text{StoreKey})$ to get back some $(\text{index}_{\tilde{f}}, \perp)$; then queries again $(\text{index}_{\tilde{f}}, \tilde{f}, \text{OutputKey})$ for $\text{sk}_{\tilde{f}}$. Note it is also satisfied that $\tilde{f}(x^*) = f(x) = 0$.
5. At the end of the game, \mathcal{A} outputs a pair of messages (m_0, m_1) and a circuit C^* .
6. \mathcal{B} then sends (m_0, m_1) to DABE challenger and gets a ciphertext ct^* . \mathcal{B} samples a random bit $\beta \leftarrow \{0, 1\}$ and computes a fresh ciphertext $\text{ct} \leftarrow \text{DABE.Enc}(\text{DABE.mpk}, (x, \text{ct}_{\text{MFE}}^*), m_\beta)$, where $\text{ct}_{\text{MFE}}^* \leftarrow \text{MFE.SKEnc}(\text{MFE.msk}, \text{comp}_{2^{\ell_1}})$.
 C^* will continue to make queries but only allowed to query $\text{KeyGen}(\text{msk}, \text{Unmarked}, \cdot)$ oracle. \mathcal{B} can answer these queries by querying $\mathcal{O}(\text{msk}, \cdot)$ with $(\tilde{f}, \perp, \text{StoreKey})$, because \mathcal{B} is allowed to make such queries in DABE game:
 - If C makes a KeyGen query in the format $(\text{Unmarked}, f)$, then \mathcal{B} prepares $\tilde{f}(x, c) = f(x)$ and first queries $(\tilde{f}, \perp, \text{StoreKey})$ to get back some $(\text{index}_{\tilde{f}}, \perp)$; then queries again $(\text{index}_{\tilde{f}}, \tilde{f}, \text{OutputKey})$ for $\text{sk}_{\tilde{f}}$. Note it is also satisfied that $\tilde{f}(x^*) = f(x) = 0$.
7. Finally, \mathcal{B} runs the decryption circuit C^* on ct^* and ct , and if $C^{\text{KeyGen}(\text{msk}, \text{Unmarked}, \cdot)}(\text{ct}^*) = C^{\text{KeyGen}(\text{msk}, \text{Unmarked}, \cdot)}(\text{ct})$, it outputs $b' = \beta$. Otherwise, it outputs $b' = 1 - \beta$.

The rest of the analysis is the same as in the original proof.

D Watermarkable Implementation of Digital Signatures

D.1 Preliminaries: Constrained Signatures

The building block of a bounded collusion resistant watermarkable implementation of digital signatures in [GKM⁺19] relies on the following building block.

Constrained Signatures A constrained signature with message space \mathcal{M} and constraint family $\mathcal{F} \subseteq \text{Function}[\mathcal{M}, \{0, 1\}]$ is a tuple of algorithms:

$\text{Setup}(1^\lambda) \rightarrow (\text{vk}, \text{msk})$. On input the security parameter λ , the setup algorithm outputs the verification key vk and the master secret key msk .

$\text{Sign}(\text{msk}, m) \rightarrow \text{sig}$. On input the master signing key msk and a message $m \in \mathcal{M}$, the signing algorithm outputs a signature sig .

$\text{Verify}(\text{vk}, m, \text{sig}) \rightarrow b$. On input the verification key vk , a message $m \in \mathcal{M}$, and a signature sig , the verification algorithm outputs a bit $b \in \{0, 1\}$.

$\text{Constrain}(\text{msk}, f) \rightarrow \text{sk}_f$: On input the master signing key msk and a function $f \in \mathcal{F}$, the constrain algorithm outputs a constrained key sk_f .

$\text{ConstrainSign}(\text{sk}_f, m) \rightarrow \text{sig}$. On input a constrained key sk_f and a message $m \in \mathcal{M}$, the signing algorithm outputs a signature sig .

Correctness A constrained signature scheme is correct if for all messages $m \in \mathcal{M}$ and key pair $(\text{vk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$:

$$\Pr[\text{Verify}(\text{vk}, m, \text{Sign}(\text{msk}, m)) = 1] = 1.$$

In addition, for all constraints $f \in \mathcal{F}$ where $f(m) = 1$,

$$\Pr[\text{Verify}(\text{vk}, m, \text{ConstrainSign}(\text{sk}_f, m)) = 1 : \text{sk}_f \leftarrow \text{Constrain}(\text{msk}, f)] = 1.$$

Constrained Unforgeability A constrained signature scheme is secure if for every stateful admissible PPT \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$:

$$\Pr \left[\text{Verify}(\text{vk}, m^*, \text{sig}^*) = 1 : \begin{array}{l} (\text{msk}, \text{vk}) \leftarrow \text{Setup}(1^\lambda) \\ (m^*, \text{sig}^*) \leftarrow \mathcal{A}^{\text{Sign}(\text{msk}, \cdot), \text{Constrain}(\text{msk}, \cdot)}(1^\lambda, \text{vk}) \end{array} \right] \leq \text{negl}(\lambda).$$

where \mathcal{A} is admissible if (1) it does not make a signing query on message m^* ; and (2) it does not make a constrained key query for any function $f \in \mathcal{F}$ such that $f(m^*) = 1$.

D.2 Definition: Watermarkable Signatures

A watermarkable signatures scheme consists of the following algorithms:

- $\text{WMSetup}(1^\lambda) \rightarrow (\text{vk}, \text{sk}, \text{xk}, \text{mk})$: on input security parameter outputs verification key vk , signing key sk ; extraction key xk , marking key mk .
- $\text{Mark}(\text{mk}, \text{sk}, \tau) \rightarrow \text{sk}_\tau$. On input the marking key mk , a signing key sk , and a mark $\tau \in \mathcal{M}_\tau$, the marking algorithm outputs a marked key sk_τ .

- $\text{Sign}(\text{sk}, m) \rightarrow \text{sig}$. On input a signing key sk and a message $m \in \mathcal{M}$, the signing algorithm outputs a signature sig . $\text{Verify}(\text{vk}, m, \sigma) \rightarrow 0/1$. On input a verification key vk , a message $m \in \mathcal{M}$, and a signature sig , the verification algorithm outputs a bit to signify whether the signature is valid or not.
- $\text{Extract}(\text{xk}, \text{vk}, C) \rightarrow \tau/\perp$. On input the extraction key xk , a verification key vk , and a circuit C , the extraction algorithm either outputs a mark $\tau \in \mathcal{M}_\tau$ or \perp .

The correctness, meaningfulness and other properties are natural to derive from the general watermarking definition Section 4.1. We refer to [GKM⁺19] for more details and only present functionality-preserving and security definitions.

Functionality Preserving A watermarkable signature scheme satisfies the functionality-preserving property if there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $(\text{vk}, \text{sk}, \text{mk}, \text{xk}) \leftarrow \text{WMSetup}(1^\lambda)$, $m \in \mathcal{M}$, $\tau \in \mathcal{M}_\tau$, the following holds:

$$\Pr \left[\text{Verify}(\text{vk}, m, \text{wSecEval}(\text{sk}_\tau, m)) = 1 : \begin{array}{l} (\text{vk}, \text{sk}, \text{xk}, \text{mk}) \leftarrow \text{WMSetup}(1^\lambda), \\ \text{sk}_\tau \leftarrow \text{Mark}(\text{mk}, \text{sk}, \tau) \end{array} \right] \geq 1 - \text{negl}(\lambda)$$

We provide the unremovability definition, which is different from the one defined in [GKM⁺19], but their scheme can be modified to satisfy this definition.

γ -Unremovability For every stateful γ -unremovable admissible PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, the following holds:

$$\Pr \left[\text{Extract}(\text{xk}, \text{vk}, C) \notin \mathcal{Q} \wedge C \text{ is } \gamma\text{-good} : \begin{array}{l} (\text{sk}, \text{vk}, \text{xk}, \text{mk}) \leftarrow \text{WMSetup}(1^\lambda) \\ C \leftarrow \mathcal{A}^{\text{Mark}(\text{mk}, \text{sk}, \cdot)}(1^\lambda, \text{vk}) \end{array} \right] \leq \text{negl}(\lambda).$$

where \mathcal{Q} is the set of marks queried by \mathcal{A} and C is said to be a (PPT admissible, stateful) γ -good adversary if:

$$\Pr \left[C^{\text{Sign}(\text{sk}, \cdot)} \rightarrow (m^*, \text{sig}^*) : \text{Verify}(\text{vk}, m^*, \text{sig}^*) = 1 \right] \geq \gamma.$$

C is admissible if and only if it does not query $\text{Sign}(\text{sk}, \cdot)$ on m^* .

Discussions on Security Definitions and [GKM⁺19] construction In the [GKM⁺19]’s watermarkable signature scheme, the unmarked Sign function and the circuit used to compute signatures using marked keys have different output distributions, if we let \mathcal{A} query $\text{Sign}(\text{sk}, \cdot)$ oracle in the first stage but allow C to choose the challenge message to sign on, then there can be an attack. Because the [GKM⁺19] scheme ’s $\text{Sign}(\text{sk}, \cdot)$ and marked signing circuit have different output distributions. If C can choose its own m^* but the Extract algorithm does not know which queries \mathcal{A} has made in stage 1, then C can be hardcoded with some m^*, sig' where watermark cannot be extracted. We thus give the above definition, which suffice for our applications 1.

D.3 Construction

We show the following modified construction from [GKM⁺19] based on a constrained signature scheme CSig defined in Appendix D.1:

- Let $\mathcal{T}' = \mathcal{T} \cup \{\perp\}$. For a mark $\tau^* \in \mathcal{T}$, let $f_{\tau^*} : \mathcal{T}' \times \mathcal{M} \rightarrow \{0, 1\}$ be the function $f_{\tau^*}(\tau, m) = 1$ if $\tau = \tau^*$ and 0 otherwise.

$\text{WMSetup}(1^\lambda) \rightarrow (\text{vk}, \text{sk}, \text{mk}, \text{xk})$: outputs a signing/verification key-pair $(\text{vk}, \text{sk}) \leftarrow \text{CSig.Setup}(1^\lambda)$; $\text{mk} = \perp, \text{xk} = \text{sk}$.

$\text{Sign}(\text{sk}, m) \rightarrow \text{sig}$: On input a signing key sk , and a message $m \in \mathcal{M}$, the signing algorithm signs $\text{sig}' \leftarrow \text{CSig.Sign}(\text{sk}, (\perp, m))$, and outputs the signature $\text{sig} = (\perp, \text{sig}')$.

$\text{Verify}(\text{vk}, m, \text{sig}) \rightarrow b$. On input a verification key vk , a message $m \in \mathcal{M}$, and a signature $\text{sig} = (\tau', \text{sig}')$; the verification algorithm outputs $b \leftarrow \text{CSig.Verify}(\text{vk}, (\tau', m), \text{sig}')$.

$\text{Mark}(\text{mk}, \text{sk}, \tau) \rightarrow C$. On input a marking key $\text{mk} = \perp$, a signing key sk , and a mark $\tau \in \mathcal{T}$, the marking algorithm computes $\text{sk}_\tau \leftarrow \text{CSig.Constrain}(\text{sk}, f_\tau)$ and outputs a circuit $C_\tau : \mathcal{M} \rightarrow \text{SIG}$ where $C_\tau(\cdot) := (\tau, \text{CSig.ConstrainSign}(\text{sk}_\tau, (\tau, \cdot)))$.

$\text{Extract}(\text{xk}, \text{vk}, C) \rightarrow \tau/\perp$: on $\text{xk} = \text{CSig.sk}, \text{vk} = \text{CSig.vk}$ and circuit $C : \mathcal{M} \rightarrow \text{SIG}$, perform the following for $T = \lambda/\gamma$ times where γ is the unremovability parameter:

- For $i \in [T]$: compute $(m_i = (\tau'_i, m'_i), \text{sig}_i) \leftarrow C^{\text{Sign}(\text{sk}, \cdot)}$. If $\text{CSig.Verify}(\text{vk}, (\tau_i, m'_i), \text{sig}_i) = 1$ and m'_i has not been queried, abort and output τ'_i .

The correctness, functionality-preserving, meaningfulness proof will all follow exactly from [GKM⁺19]. Even though our extraction algorithm is different, the unremovability proof will follow similarly: any valid signatures provided by C must contain some $\tau_i \in \mathcal{Q}$ of the marking queries, otherwise it helps break the constrained unforgeable security of the constrained signature. The only change is that the reduction does not answer any $\text{Sign}(\text{sk}, \cdot)$ queries in the first stage, and thus do not need to query $\text{CSig.Sign}(\text{sk}, \cdot)$ oracle when interacting with \mathcal{A} , but only when interacting with the circuit C produced by \mathcal{A} .

E Watermarkable CCA-secure Hybrid Encryption (Key Encapsulation Scheme)

Theorem E.1. *Assuming LWE, there exists secure watermarkable implementation of a CCA-secure hybrid encryption scheme.*

The assumption of LWE comes from the need of watermarkable CCA-secure PKE.

Construction A watermarkable implementation of CCA-secure hybrid encryption $\text{wHE} = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Mark}, \text{Extract})$ can be built from a CCA-secure secret key encryption scheme $\text{SKE} = (\text{SKE.KeyGen}, \text{SKE.Enc}, \text{SKE.Dec})$ and a watermarkable implementation of CCA-secure public key encryption $\text{wPKE} = (\text{wPKE.WMSetup}, \text{wPKE.Enc}, \text{wPKE.Dec}, \text{wPKE.Mark}, \text{wPKE.Extract})$ as follows:

- $\text{KeyGen}(\lambda)$: compute $(\text{wPKE.pk}, \text{wPKE.sk}, \text{wPKE.xk}, \text{wPKE.mk}) \leftarrow \text{wPKE.WMSetup}(\lambda)$. Output $\text{sk} = \text{wPKE.sk}; \text{pk} = \text{wPKE.pk}; \text{xk} = \text{wPKE.xk}; \text{mk} = \text{wPKE.mk}$
- $\text{Enc}(\text{pk}, m)$: compute $\text{SKE.sk} \leftarrow \text{SKE.KeyGen}(\lambda)$; then compute $\text{ct}_1 \leftarrow \text{wPKE.Enc}(\text{pk}, \text{SKE.sk})$; compute $\text{ct}_2 \leftarrow \text{SKE.Enc}(\text{SKE.sk}, m)$. Output $\text{ct} = (\text{ct}_1, \text{ct}_2)$.
- $\text{Dec}(\text{sk}, \text{ct})$: parse $\text{ct} = (\text{ct}_1, \text{ct}_2)$; compute $\text{sk}' \leftarrow \text{wPKE.Dec}(\text{sk}, \text{ct}_1)$ and then $m' \leftarrow \text{SKE.Dec}(\text{sk}', \text{ct}_2)$. Output m' .
- $\text{Mark}(\text{mk}, \text{sk}, \tau)$: parse $\text{mk} = \text{wPKE.mk}; \text{sk} = \text{wPKE.sk}$; output $\text{sk}_\tau \leftarrow \text{wPKE.Mark}(\text{sk}, \tau)$;

- $\text{Extract}(xk, pk, \text{aux} = \perp, C)$:
 - parse $xk := \text{wPKE}.xk$; $pk := \text{wPKE}.pk$. Create the following circuit C_{PKE} with black-box access to C :
 - * C_{PKE} is hardcoded with pk and simulates the CCA-PKE game for C as follows:
 - * give pk to C ; for C 's decryption queries $ct = (ct_1, ct_2)$: query an external oracle $sk' \leftarrow \text{wPKE}.Dec(\text{wPKE}.sk, ct_1)$; output $m' \leftarrow \text{SKE}.Dec(sk', ct_2)$.
 - * C submits challenge messages (m_0, m_1) ; C_{PKE} submits $(sk_0 \leftarrow \text{SKE}.KeyGen(1^\lambda), sk_1 \leftarrow \text{SKE}.KeyGen(1^\lambda))$ to the external challenger; C_{PKE} receives challenge ciphertext ct^* from the challenger and C_{PKE} sends $ct^{**} = (ct^*, ct_2^* \leftarrow \text{SKE}.Enc(sk_0, m_b))$ to C , where $m_b \leftarrow \{m_0, m_1\}$.
 - * C continues to simulate the decryption oracle as above to decrypt only valid ciphertexts.
 - * If C outputs $b' = b$, then C_{PKE} outputs guess 0; else it outputs guess 1.
 - output $\tau/\perp \leftarrow \text{wPKE}.Extract(\text{wPKE}.xk, \text{wPKE}.pk, \text{aux} = \perp, C_{\text{PKE}})$.

Now we prove Theorem E.1

Proof. Suppose there exists adversary \mathcal{A} that breaks the γ -unremovability of watermarkable CCA-secure PKE for some non-negligible γ , i.e. $\mathcal{A}^{\text{Dec}(sk, \cdot), \text{Mark}(mk, sk, \cdot)}(pk)$ produces some program C such that $\Pr[\text{Extract}(xk, pk, C) \notin \mathcal{Q}] \geq \epsilon$ for some non-negligible ϵ , whereas C satisfies $\Pr[G_{CCA}(\text{sk}, \text{pk}, C) = 1] \geq \frac{1}{2} + \gamma$ (for G_{CCA} see Section 5.5.2). The probability $\Pr[\text{Extract}(xk, pk, C) \notin \mathcal{Q}]$ is taken over the randomness in WMSetup and the probability $\Pr[G_{CCA}(\text{sk}, \text{pk}, C) = 1]$ is taken over the randomness used in G_{CCA} .

We will show that we can either break γ' -unremovability of the watermarkable wPKE or CCA-security of SKE.

For any (sk, pk, xk, mk) generated by $\text{WMSetup}(1^\lambda)$ and any $\mathcal{A}^{\text{Dec}(sk, \cdot), \text{Mark}(mk, sk, \cdot)}(pk)$ producing a γ -good C such that $\Pr[\text{Extract}(xk, pk, C) \notin \mathcal{Q}] \geq \epsilon$, one of the following cases must hold:

- **Case 1: the program C_{PKE} created during the execution of $\text{Extract}(xk, C)$ is a γ' -good adversary for the stage-2 game $G_{CCA-PKE}(\text{sk}, \text{pk}, \cdot)$ for some non-negligible γ' .** The reduction $\mathcal{B}_{\text{wPKE}}$ works as follows:
 - $\mathcal{B}_{\text{wPKE}}$ receives $\text{wPKE}.pk$ from the challenger and sends to \mathcal{A} . $\mathcal{B}_{\text{wPKE}}$;
 - for \mathcal{A} 's marking queries : query the challenger's marking oracle $\text{Mark}(\text{wPKE}.sk, \cdot)$ and forward the output.

After \mathcal{A} outputs program C , $\mathcal{B}_{\text{wPKE}}$ creates circuit C'_{wPKE} with black-box access to C :

- C'_{wPKE} is hardcoded with pk and makes external queries to simulate the decryption oracle for C :
 - * for C 's decryption queries $ct = (ct_1, ct_2)$: query the challenger's decryption oracle $sk' \leftarrow \text{wPKE}.Dec(\text{wPKE}.sk, ct_1)$; output $m' \leftarrow \text{SKE}.Dec(sk', ct_2)$.
- C submits challenge messages (m_0, m_1) ; C_{PKE} submits $(sk_0 \leftarrow \text{SKE}.KeyGen(1^\lambda), sk_1 \leftarrow \text{SKE}.KeyGen(1^\lambda))$ to the external challenger; C_{PKE} receives challenge ciphertext ct^* from the challenger and C_{PKE} sends $ct^{**} = (ct^*, ct_2^* \leftarrow \text{SKE}.Enc(sk_0, m_b))$ to C , where $m_b \leftarrow \{m_0, m_1\}$.
- C continues to simulate the decryption oracle as above to decrypt only valid ciphertexts.

- If C outputs $b' = b$, then C_{PKE} outputs guess 0; else it outputs guess 1.

By the design of our `Extract` and by our assumption, we must have that $\Pr[\text{wPKE.Extract}(\text{wPKE.xk}, \text{wPKE.pk}, \text{aux}, C_{\text{PKE}}) \notin \mathcal{Q}] \geq \epsilon$ where C_{PKE} is the program created during `Extract(xk, pk, C)`.

For any $(\text{xk}, \text{mk}, \text{sk}, \text{pk})$, it is easy to see that program C_{PKE} created by `Extract(xk, pk, C)`'s input-output behavior is the same as C'_{wPKE} above. Therefore, if $\Pr[\text{Extract}(\text{xk}, \text{pk}, C_{\text{PKE}}) \in \mathcal{Q}]$ is γ' -good and $\Pr[\text{Extract}(\text{xk}, \text{pk}, C_{\text{PKE}})] \geq \epsilon$, so will C'_{wPKE} satisfy these two conditions. Thus $\mathcal{B}_{\text{wPKE}}$ breaks the γ' -unremovability.

- **Case 2: the program C_{PKE} created during the execution of `Extract(xk, C)` is not a γ' -good adversary for the stage-2 game $G_{\text{CCA-PKE}}(\text{sk}, \text{pk}, \cdot)$ for any non-negligible γ' .** In this case, we must have that $|\Pr[C_{\text{PKE}} \rightarrow 0 | \text{ct}^* = \text{PKE.Enc}(\text{pk}, \text{sk}_0)] - \Pr[C_{\text{PKE}} \rightarrow 0 | \text{ct}^* = \text{PKE.Enc}(\text{pk}, \text{sk}_1)]| = |\Pr[C \text{ outputs } b' = b | \text{ct}^* = \text{PKE.Enc}(\text{pk}, \text{sk}_0)] - \Pr[C \text{ outputs } b' = b | \text{ct}^* = \text{PKE.Enc}(\text{pk}, \text{sk}_1)]| \leq \text{negl}(\lambda)$;

Thus we can switch to doing the following in running `Extract(xk, pk, C)`: when C submits challenge messages (m_0, m_1) ; C_{PKE} sends $\text{ct}^{**} = (\text{ct}_1^* = \text{PKE.Enc}(\text{PKE.pk}, \text{sk}_1), \text{ct}_2^* \leftarrow \text{SKE.Enc}(\text{sk}_0, m_b))$ to C , where $m_b \leftarrow \{m_0, m_1\}$ where ct^* is always $\text{PKE.Enc}(\text{PKE.pk}, \text{sk}_1)$.

Since C is still $(\gamma - \text{negl}(\lambda))$ good, we will build a reduction to break IND-CCA security of SKE.

The reduction \mathcal{B}_{SKE} samples `wPKE`'s keys on its own and can query the oracles `SKE.Enc`(sk_0, \cdot), `SKE.Dec`(sk_0, \cdot) provided by the SKE challenger.

After \mathcal{A} outputs C , \mathcal{B}_{SKE} enters stage-2 of reduction and simulates the encryption and decryption oracles by submitting the decryption queries to the encryption, decryption oracles `SKE.Enc`(sk_0, \cdot), `SKE.Dec`(sk_0, \cdot) in the CCA-security game of SKE. In the challenge phase, C submits challenge messages (m_0, m_1) ; $\mathcal{B}_{\text{SKE}}^2$ submits (m_0, m_1) to the external challenger; and samples $\text{sk}_1 \leftarrow \text{SKE.KeyGen}(1^\lambda)$; C_{PKE} receives challenge ciphertext $\text{ct}_2^* \leftarrow \text{SKE.Enc}(\text{sk}_0, m_b)$ from the challenger, where sk_0 is only known to the challenger and C_{PKE} sends $\text{ct}^{**} = (\text{ct}_1^* = \text{wPKE.Enc}(\text{pk}, \text{sk}_1), \text{ct}_2^* \leftarrow \text{SKE.Enc}(\text{sk}_0, m_b))$ to C . C continues to simulate the encryption and decryption oracle as above to decrypt only valid ciphertexts. In the end, if C outputs b' , then $\mathcal{B}_{\text{SKE}}^2$ outputs the same b' .

□

F Watermarkable Functional Encryption from Watermarkable Attribute-Based Encryption

We give a high-level description on how to turn the [GKP⁺13] functional encryption construction from ABE, FHE, garbled circuits into a watermarkable FE based on watermarkable ABE. We will not elaborate details of the construction as we cannot possibly cover all explicit constructions for watermarking-compositions of existing schemes, and our main goal here is to give reference to another example of watermarking-composition when applied to an advanced encryption scheme.

All the main algorithms are the same as the [GKP⁺13] construction. We will mainly remark on how we do extraction and refer interested readers to [GKP⁺13] for the detailed construction.

Watermarkable Functional Encryption: Definitions A watermarkable FE scheme can be given different watermarking definitions. One is to watermark the master secret key and given out an

unmarked functional key to the adversarial program \mathcal{A} ; the other is to let \mathcal{A} submit a function f in the first stage and gets to query on marked versions of the functional key sk_f . Both notions are interesting.

For the first notion, we can consider the following security game: The γ -Unremovability for a watermarkable single-key FE scheme says, for all $\lambda \in \mathbb{N}$, and for all PPT admissible stateful adversary \mathcal{A} , there exists a negligible function $\text{negl}(\lambda)$ such that:

$$\Pr \left[\begin{array}{l} \text{Extract}(\text{xk}, \text{mpk}, \text{aux} = x, C) \notin \mathcal{Q} \\ \wedge C \text{ is } \gamma\text{-good} \end{array} : \begin{array}{l} (\text{msk}, \text{mpk}, \text{xk}, \text{mk}) \leftarrow \text{WMSetup}(1^\lambda) \\ (x, C) \leftarrow \mathcal{A}^{\text{Mark}(\text{mk}, \text{msk}, \cdot)}(1^\lambda, \text{mpk}) \end{array} \right] \leq \text{negl}(\lambda).$$

where \mathcal{Q} is the set of marks queried by \mathcal{A} and $C = (C_1, C_2, D)$ is a PPT admissible, stateful γ -good adversary in the security game $G_{\text{ABE}}(\text{msk}, \text{mpk}, \text{aux} = (x), \cdot)$ if:

$$\Pr \left[\begin{array}{l} D(\text{st}, \text{ct}_b) = b : \\ \begin{array}{l} (f, \text{st}_1) \leftarrow C_1(\text{mpk}) \\ \text{sk}_f \leftarrow \text{FE.KeyGen}(\text{msk}, f) \\ \text{ct}_0 \leftarrow \text{FE.Enc}(\text{mpk}, x), \text{ct}_1 \leftarrow \text{Sim}(\text{mpk}, \text{sk}_f, f, f(x), 1^{|x|}), b \leftarrow \{0, 1\} \\ \text{st} \leftarrow C_2(\text{st}_1, \text{ct}_b) \end{array} \end{array} \right] \geq \frac{1}{2} + \gamma.$$

where Sim is some PPT simulator.

A selective variant is requiring \mathcal{A} output x before seeing mpk . Note that we need \mathcal{A} to output x in the first stage because by relying out construction on ABE, the choice of message x will influence the choice of the ABE reduction's choice of challenge attribute, which must be committed to before entering stage 2. We can also let \mathcal{A} outputs f first so that \mathcal{A} sees sk_f before outputting x (which matches the semantics of a fully secure FE). But then the Extract algorithm will have to take in f as an auxiliary inputs.

The case where sk_f gets watermarked is slightly trickier to define. We need to work with an indistinguishability-based FE security where the function family \mathcal{F} is a "high-entropy" function family. Also, the Extract needs to take in the challenge f as an auxiliary input. We refer to [GKM⁺19] Remark 4.5 for more discussions. [GKM⁺19]'s watermarkable predicate encryption scheme can be extended to a watermarkable FE scheme for watermarking a functional key when working with the above restricted function class.

F.1 Two-Outcome Attribute-Based Encryption

The first building block for [GKP⁺13] FE is a two-outcome ABE, called ABE_2 . It has the similar syntax to a normal ABE except that: the encryption scheme takes in mpk, x, m_0, m_1 ; when decrypting with a functional key sk_f on ciphertext ct_{x, m_0, m_1} , the decryption will output m_b when $f(x) = b$.

The single-key ABE_2 security game is exactly like a single-key ABE security game except that \mathcal{A} provides 3 challenge messages (m, m_0, m_1) and attribute x , then \mathcal{A} is given either $\text{ABE}_2.\text{Enc}(\text{mpk}, x, (m_0, m_1))$ if $f(x) = 1$ for the key generation query f , or $\text{ABE}_2.\text{Enc}(\text{mpk}, x, (m, m_1))$ if $f(x) = 0$, and asked to output b .

A single-key ABE_2 scheme can be built with a simple construction from a regular single-key ABE. We refer the readers to [GKP⁺13] Appendix B for details and give a high level idea: the scheme generates two pairs regular ABE keys; the encryption generates two encryptions of the same message and attribute, each under one of the master public keys respectively; the decryption

algorithm on sk_f and $ct = (ct_0, ct_1)$ will try to decrypt both ciphertexts, and output that one that has a valid decryption result.

In the watermarking construction, we watermark both the master secret keys of the two ABE (regarding single-key watermarkable ABE security: see Definition C.3 except that \mathcal{A} is not give $\text{KeyGen}(\text{msk}, \cdot)$ oracle and C can only query $\text{KeyGen}(\text{msk}, \cdot)$ once). The Extract algorithm in the watermarkable ABE_2 will turn the input program C into a circuit C_i to break the single-key ABE security for each of the keys in the construction. The reduction in [GKP⁺13] Appendix B would go through because each stage of ABE_2 in the unremovability game would map to the stage of the regular ABE unremovability game. That is, making some msk marking queries after receiving mpk , and outputs challenge messages, attribute and program C ; then C gets to query a single sk_f and challenged with the challenge ciphertext (If we choose to work with a watermarkable FE scheme where we let \mathcal{A} choose f , we can also let the adversary in watermarkable ABE scheme query one sk_f and then the program C is not allowed to make any queries).

Watermarkable Functional Encryption Construction, Extraction and Security We refer the readers to [GKP⁺13] Section 3.1 for details on the construction and its preliminaries for definitions of FHE and garbled circuit, since we only plan to give an idea here.

The construction only needs to rely on a watermarkable ABE_2 scheme (which we have shown above can be based on watermarkable regular ABE). The FHE and garbled circuits do not have to be watermarked because they are only generated freshly on each evaluation.

The WMSetup algorithm generates all λ number of $(\text{mpk}_i, \text{msk}_i, \text{mk}_i, \text{xk}_i)$ from the setup of a watermarkable ABE_2 scheme.

The $\text{KeyGen}(\text{msk}, f)$ algorithm generates a $sk_i \leftarrow \text{ABE}_2.\text{KeyGen}(\text{msk}_i, \text{FHE.Eval}_f^i)$ where FHE.Eval_f^i is outputting the i -th bit after the FHE evaluation on some FHE Eval key FHE.pk , the input function f and some FHE ciphertexts c_1, \dots, c_n .

The $\text{Enc}(\text{mpk}, x)$ algorithm takes input attribute $x = x_1 \dots x_n$, generates a fresh FHE key pair, encrypt each x_i to get FHE ciphertext c_i . Then produce garbled circuit for the FHE decryption algorithm $\text{FHE.Dec}(\text{FHE.sk}, \cdot)$ with 2λ labels $\{L_i^0, L_i^1\}_{i \in [\lambda]}$. Then we produce ABE_2 ciphertext $ct_i \leftarrow \text{ABE}_2.\text{Enc}(\text{mpk}_i, (\text{FHE.pk}, \vec{c}), L_i^0, L_i^1)$ for all $i \in [\lambda]$ where $\vec{c} = (c_1, \dots, c_n)$. Output the ciphertext $(ct_1, \dots, ct_\lambda)$ and the garbled circuit.

The decryption algorithm $\text{Dec}(sk_f, ct)$ runs the ABE_2 decryption on each ct_i using sk_i to obtain a label $L_i^{d_i}$, for each $i \in [\lambda]$. Then one can recover $f(x) \leftarrow \text{FHE.Dec}(\text{FHE.sk}, d_1 \dots d_\lambda)$ by using the labels $L_i^{d_i}$ and the garbled circuit.

Now we can talk about how we watermark and extract: the marking scheme simply marks each $\text{ABE}_2.\text{msk}_i$ for $i \in [\lambda]$. To see how we extract, we roughly open up the security proof of the above scheme: On an input program C , the extraction interacts with C as in the stage-2 FE game in our unremovability definition. Recall that in the end of the original security game, we need to prepare either real ciphertext $ct \leftarrow \text{Enc}(\text{mpk}, x)$ or a simulated ciphertext, and ask the final stage adversary D to distinguish. But in our extraction algorithm that tries to create circuit C_i to break the unremovability of the i -the ABE_2 scheme, we instead sample either of the following: (1) generate the ciphertext as in the original scheme; (2) use the label $L_i^{d_i}$ twice in the ciphertext $ct'_i \leftarrow \text{ABE}_2.\text{Enc}(\text{mpk}_i, (\text{FHE.pk}, \vec{c}), L_i^{d_i}, L_i^{d_i})$, where $d_i = \text{FHE.Eval}_f^i(\text{FHE.pk}, \vec{c})$. More specifically, circuit C_i obtains one of these ciphertexts from some external ABE_2 challenger and uses C 's output to distinguish them.

Why would a γ -good C in the original security game still be γ -good when we simulate the

game as in the above extraction? This relies on the hybrid arguments invoking the IND-CPA security of FHE and then the security of garbled circuit. We refer the details to the proofs for Hybrid 0 to Hybrid 2 in Section 3.2 of [GKP⁺13]. By invoking these security properties, we move from giving a completely simulated ciphertext, to give the fake ciphertext described in the extraction algorithm (when flipping the coin to 1), while C should still be γ -good with some noticeable probability.

Then we can argue that at least one of C_i is γ_i -good for some non negligible γ_i , we can break unremovability of ABE_2 with the i -th key pair. In the marking stage, the i -th ABE_2 reduction can simulate the marking queries by making marking query to the ABE_2 marking oracle (and mark the other keys on its own). After \mathcal{A} commits to challenge message x , the reduction can also commit the challenge attribute $(\text{FHE.pk}, \vec{c})$ where \vec{c} is computed from x as in the construction. After \mathcal{A} outputs C , the reduction creates circuit C_i as in the extraction algorithm, and C_i can answer C 's single key query by making a single query to the ABE_2 $\text{KeyGen}(\text{msk}, \cdot)$ oracle.

G Watermarkable Non-malleable PKE from Watermarkable CPA-secure PKE, NIZK and Signatures

G.1 Watermarkable Implementation of Non-malleable CCA2 PKE

We present that an alternative construction to Appendix A, the CCA2 non-malleable PKE in [DDN91] also has a watermarking implementation. We will omit the security proof due to the similarity to Appendix A.

G.2 Construction and Security

Our construction is based on the non-malleable encryption scheme in [DDN91].

Given a watermarkable implementation of a CPA-secure PKE $\text{wPKE} = (\text{WMSetup}, \text{Enc}, \text{Dec}, \text{Mark}, \text{Extract})$, a NIZK scheme $\text{NIZK} = (\text{Setup}, \text{Prove}, \text{Verify})$ and a signature scheme $\text{DS} = (\text{KeyGen}, \text{Sign}, \text{Verify})$. where the verification key size of DS is n .

Construction

$\text{KeyGen}(1^\lambda, 1^n)$: For $i \in [n], b \in \{0, 1\}$: compute $(\text{wPKE.pk}_{i,b}, \text{wPKE.sk}_{i,b}, \text{xk} = \text{wPKE.xk}_{i,b}; \text{mk} = \text{wPKE.mk}_{i,b}) \leftarrow \text{wPKE.WMSetup}(\lambda)$;

Compute $\text{CRS} \leftarrow \text{NIZK.Setup}(1^\lambda)$; Output $\text{sk} = (\text{wPKE.sk}_{i,b})_{i \in [n], b \in \{0,1\}}$; $\text{pk} = (\{\text{wPKE.pk}_{i,b}\}_{i \in [n], b \in \{0,1\}}, \text{CRS})$;

$\{\text{wPKE.pk}_{i,b}\}_{i \in [n], b \in \{0,1\}}$; $\text{mk} = \{\text{wPKE.mk}_{i,b}\}_{i \in [n], b \in \{0,1\}}$

$\text{Enc}(\text{pk}, m)$:

- parse $\text{pk} := \{\text{pk}_{i,b}\}_{i \in [n], b \in \{0,1\}}, \text{CRS}$; compute $(\text{DS.sk}, \text{DS.vk}) \leftarrow \text{DS.KeyGen}(1^\lambda)$; Let us denote i -th bit of DS.vk as vk_i .
- For each $i \in [n]$: compute $\text{ct}_i \leftarrow \text{wPKE.Enc}(\text{pk}_{i,\text{vk}_i}, m)$;
- compute $\pi \leftarrow \text{NIZK.Prove}(\text{CRS}, (\text{ct}_1, \text{ct}_2, \dots, \text{ct}_n, \text{DS.vk}), (r_1, r_2, \dots, r_n, m))$ for the following statement:
 \exists witness $(r_1, r_2, \dots, r_n, m)$ such that $\text{ct}_i = \text{wPKE.Enc}(\text{pk}_{i,\text{vk}_i}, m; r_i)$ for all $i \in [n]$, where r_i is the randomness used in encryption.
- compute a signature $\text{sig} \leftarrow \text{DS.Sign}(\text{DS.sk}, (\{\text{ct}_i\}_{i \in [n]}, \pi))$

- Output $ct = (\{ct_i\}_{i \in n}, DS.vk, \pi, sig)$.

Dec(sk, ct):

- parse $ct = (\{ct_i\}_{i \in n}, DS.vk, \pi)$; $sk = \{sk_{i,b}\}_{i \in [n], b \in \{0,1\}}$;
- if $DS.Verify(DS.vk, (\{ct_i\}_{i \in n}, \pi), sig) = 1$ continue; else abort and output \perp .
- if $NIZK.Verify(CRS, \pi, ((\{ct_i\}_{i \in n}, DS.vk))) = 1$, continue; else abort and output \perp .
- compute $m \leftarrow wPKE.Dec(sk_{1,vk_1}, ct_1)$; output m' .

Mark(mk, sk, τ): parse $mk = \{mk_{i,b}\}_{i \in [n], b \in \{0,1\}}$; $sk = \{sk_{i,b}\}_{i \in [n], b \in \{0,1\}}$; output $\{sk_{i,b,\tau} \leftarrow wPKE_{i,b}.Mark(sk_{i,b}, \tau)\}_{i \in [n]}$

Extract(xk, pk, aux = (m_0, m_1) , C):

- parse $xk := (\{xk_{i,b}\}_{i \in [n], b \in \{0,1\}}, NIZK.td)$; $pk := (\{pk_{i,b}\}_{i \in [n], b \in \{0,1\}}, CRS)$. Initialize an empty tuple $\vec{\tau}$;
- For $i^* = 1, \dots, n$; $b^* = 0, 1$: Create the following circuit C_{i^*,b^*} with black-box access to C :
 - C_{i^*,b^*} is hardcoded with $(\{pk_{i,b}\}_{i \in [n], b \in \{0,1\}}, \{xk_{i,b}\}_{i \neq i^* \vee b \neq b^*}, td, CRS)$
 C_{i^*,b^*} acts as a stage-2 reduction from CCA2-PKE to CPA-PKE with the keys $(sk_{i^*,b^*}, pk_{i^*,b^*})$ simulates the CCA2-PKE stage-2 game for C as follows:
 - For C 's decryption queries $ct = (\{ct_{i,vk_i}\}, vk, \pi, sig)$:
 - * First check if $NIZK.Verify(CRS, \pi) = 1$, if 0 output \perp ; else continue;
 - * By the extraction key simulation property, since $xk_{i \neq i^* \vee b \neq b^*}$ can be used to simulate the oracles used in $G_{CPA}(pk_i, sk_j, \cdot)$ for $i \neq i^*$ or $b \neq b^*$, C_{i^*,b^*} can simulate the oracle $wPKE.Dec(sk_{i^*,1-b^*}, \cdot)$ and thus decrypt $ct_{i^*,1-b^*}$ in the ciphertext ct to output m .
 - C_{i^*,b^*} submits (m_0, m_1) to the external challenger; C_{i^*,b^*} receives challenge ciphertext $ct_{j^*,b^*}^* = wPKE.Enc(pk_{i^*,b^*}, m_\delta)$, $\delta \leftarrow \{0, 1\}$ from the challenger.
 - C_{i^*,b^*} prepares the following ciphertext:
 1. Sample $(DS.vk^*, DS.sk^*) \leftarrow DS.KeyGen(1^\lambda)$ so that the i^* -th bit of vk^* is b^* .
 2. Compute $ct_{i,b_i}^* \leftarrow wPKE.Enc(pk_{i,b_i}, m_{\delta'})$, $\delta' \leftarrow \{0, 1\}$ for all $i \neq i^*$, where $b_i = vk_i^*$.
 3. compute $\hat{\pi} \leftarrow Sim(td, CRS, (\{ct_{i,b_i}^*\}_{i \in [n], b_i = vk_i^*}))$
where Sim is the simulator algorithm for NIZK.
 - Then it sends $ct^* = (\{ct_{i,b_i}^*\}_{i \in [n], b_i = vk_i^*}, vk^*, \hat{\pi})$ to C .
 - C continues to simulate the decryption oracle as above to decrypt only valid ciphertexts $ct \neq ct^*$, except adding the following check:
 - * Check if $vk = vk^*$, if yes, output \perp ; else continue to decrypt.
 - C outputs ℓ ciphertexts ct_1, \dots, ct_ℓ and C_i computes $d_j \leftarrow Dec(sk, ct_j)$ if $ct_j \neq ct^*$, else $d_j = \perp$.
 - In the end, feed (d_1, \dots, d_ℓ) to C and C outputs a bit b' . C_i outputs the same as C outputs.
 - Add $\tau / \perp \leftarrow wPKE.Extract(wPKE.xk_{i^*,b^*}, wPKE.pk_{i^*,b^*}, aux = (m_0, m_1), C_{i^*,b^*})$ to $\vec{\tau}$.
- Output $\vec{\tau}$.

The security proof is highly similar to Appendix A and we omit it here.