A preliminary version of this paper appears in the proceedings of Asiacrypt 2011. This is the full version.

Cryptography Secure Against
Related-Key Attacks and Tampering

MIHIR BELLARE[1]   DAVID CASH[2]   RACHEL MILLER[3]

August 2011

Abstract

We show how to leverage the RKA (Related-Key Attack) security of blockciphers to provide RKA security for a suite of high-level primitives. This motivates a more general theoretical question, namely, when is it possible to transfer RKA security from a primitive P_1 to a primitive P_2? We provide both positive and negative answers. What emerges is a broad and high level picture of the way achievability of RKA security varies across primitives, showing, in particular, that some primitives resist "more" RKAs than others. A technical challenge was to achieve RKA security even for the practical classes of related-key deriving (RKD) functions underlying fault injection attacks that fail to satisfy the "claw-freeness" assumption made in previous works. We surmount this barrier for the first time based on the construction of PRGs that are not only RKA secure but satisfy a new notion of identity-collision-resistance.

[1] Department of Computer Science & Engineering, University of California San Diego, 9500 Gilman Drive, La Jolla, California 92093, USA. Email: `mihir@cs.ucsd.edu`. URL: `http://www.cs.ucsd.edu/users/mihir`. Supported in part by NSF grants CNS-0904380 and CCF-0915675.

[2] IBM T.J. Watson Research Center, Yorktown Heights, NY 10598, USA.. Email: `cdc@gatech.edu`. URL: `http://www.cs.ucsd.edu/users/cdcash`. Supported in part by NSF grant CCF-0915675.

[3] Department of Electrical Engineering and Computer Science, MIT, 32 Vassar Street, Cambridge, Massachusetts 02139, USA. Email: `rmiller@csail.mit.edu`. URL: `http://people.csail.mit.edu/rmiller/`. Supported by a DOD NDSEG Graduate Fellowship.

Contents

1 Introduction

By fault injection [16, 9] or other means, it is possible for an attacker to induce modifications in a hardware-stored key. When the attacker can subsequently observe the outcome of the cryptographic primitive under this modified key, we have a related-key attack (RKA) [4, 20].

The key might be a signing key of a certificate authority or SSL server, a master key for an IBE system, or someone's decryption key. Once viewed merely as a way to study the security of blockciphers [8, 28, 4], RKAs emerge as real threats in practice and of interest for primitives beyond blockciphers.

It becomes of interest, accordingly, to achieve (provable) RKA security for popular high-level primitives. How can we do this?

PRACTICAL CONTRIBUTIONS. One approach to building RKA-secure high-level primitives is to do so directly, based, say, on standard number-theoretic assumptions. This, however, is likely to yield ad hoc results providing security against classes of attacks that are tied to the scheme algebra and may not reflect attacks in practice.

We take a different approach. RKA security is broadly accepted in practice as a requirement for blockciphers; in fact, AES was designed with the explicit goal of resisting RKAs. We currently have blockciphers whose resistance to RKAs is backed by fifteen years of cryptanalytic and design effort. We propose to leverage these efforts.

We will provide a general and systematic way to immunize any given instance of a high-level primitive against RKAs with the aid of an RKA-secure blockcipher, modeling the latter, for the purpose of proofs, as a RKA-secure PRF [4]. We will do this not only for symmetric primitives that are "close" to PRFs like symmetric encryption, but even for public-key encryption, signatures and identity-based encryption. Our methods are cheap, non-intrusive from the software perspective, and able to completely transfer all the RKA security of the blockcipher so that the high-level primitive resists attacks of the sort that arise in practice.

THEORETICAL CONTRIBUTIONS. The ability to transfer RKA security from PRFs to other primitives lead us to ask a broader theoretical question, namely, when is it possible to transfer RKA security from a primitive P_1 to a primitive P_2? We provide positive results across a diverse set of primitives, showing, for example, that RKA-secure IBE implies RKA-secure IND-CCA PKE. We also provide negative results showing, for example, that RKA-secure signatures do not imply RKA-secure PRFs.

All our results are expressed in a compact set-based framework. For any primitive P and class Φ of related-key deriving functions —functions the adversary is allowed to apply to the target key to get a related key— we define what it means for an instance of P to be Φ-RKA secure. We let $\mathbf{RKA}[P]$ be the set of all Φ such that there exists a Φ-RKA secure instance of primitive P. A transfer of RKA security from P_1 to P_2, expressed compactly as a set containment $\mathbf{RKA}[P_1] \subseteq \mathbf{RKA}[P_2]$, is a construction of a Φ-RKA secure instance of P_2 given both a normal-secure instance of P_2 and a Φ-RKA secure instance of P_1. Complementing this are non-containments of the form $\mathbf{RKA}[P_2] \not\subseteq \mathbf{RKA}[P_1]$, which show the existence of Φ such that there exists a Φ-RKA instance of P_2 yet *no* instance of P_1 can be Φ-RKA secure, indicating, in particular, that RKA security cannot be transferred from P_2 to P_1.

As Figure 1 shows, we pick and then focus on a collection of central and representative cryptographic primitives. We then establish these containment and non-containment relations in a comprehensive and systematic way. What emerges is a broad and high level picture of the way achievability of RKA security varies across primitives, showing, in particular, that some primitives resist "more" RKAs than others.

We view these relations between $\mathbf{RKA}[P]$ sets as an analog of complexity theory, where we study relations between complexity classes in order to better understand the computational complexity of particular problems. Let us now look at all this more closely.

BACKGROUND. Related-key attacks were conceived in the context of blockciphers [8, 28]. The first definitions were accordingly for PRFs [4]; for $F \colon \mathcal{K} \times \mathcal{D} \to \mathcal{R}$ they consider the game that picks a

random challenge bit b and random target key $K \in \mathcal{K}$. For each $L \in \mathcal{K}$ the game picks a random function $G(L, \cdot): \mathcal{D} \to \mathcal{R}$, and next allows the adversary multiple queries to an oracle that given a pair (ϕ, x) with $\phi: \mathcal{K} \to \mathcal{K}$ and $x \in \mathcal{D}$ returns $F(\phi(K), x)$ if $b = 1$ and $G(\phi(K), x)$ if $b = 0$. They say that F is Φ-RKA secure, where Φ is a class of functions mapping $\mathcal{K}$ to $\mathcal{K}$, if the adversary has low advantage in predicting b when it is only allowed in its queries to use functions ϕ from Φ.

Let $\mathbf{RKA}[\mathsf{PRF}]$ be the set of all Φ for which there exists a Φ-RKA secure PRF. Which Φ are in this set? All the evidence so far is that this question has no simple answer. Bellare and Kohno [4] gave natural examples of Φ not in $\mathbf{RKA}[\mathsf{PRF}]$, showing the set is not universal. Membership of certain specific Φ in $\mathbf{RKA}[\mathsf{PRF}]$ have been shown by explicit constructions of Φ-RKA PRFs, first under novel assumptions [30] and then under standard assumptions [3]. Beyond this we must rely on cryptanalysis. Modern blockciphers including AES are designed with the stated goal of RKA security. Accordingly we are willing to assume their Φ-RKA security —meaning that $\Phi \in \mathbf{RKA}[\mathsf{PRF}]$— for whatever Φ cryptanalysts have been unable to find an attack.

BEYOND PRFs. Consideration of RKAs is now expanding to primitives beyond PRFs [21, 2, 23]. This is viewed partly as a natural extension of the questions on PRFs, and partly as motivated by the view of RKAs as a class of sidechannel attacks [20]. An RKA results when the attacker alters a hardware-stored key via tampering or fault injection [16, 9] and subsequently observes the result of the evaluation of the primitive on the modified key. The concern that such attacks could be mounted on a signing key of a certificate authority or SSL server, a master key for an IBE system, or decryption keys of users makes achieving RKA security interesting for a wide range of high-level primitives.

DEFINITIONS. We focus on a small but representative set of primitives for which interesting variations in achievability of RKA security emerge. These are PRF (pseudorandom functions), Sig (Signatures), PKE-CCA (CCA-secure public-key encryption), SE-CCA (CCA-secure symmetric encryption), SE-CPA (CPA-secure symmetric encryption), IBE (identity-based encryption) and wPRF (weak PRFs [32]). We define what it means for an instance of P to be Φ-RKA secure for each $\mathsf{P} \in \{\mathsf{wPRF}, \mathsf{IBE}, \mathsf{Sig}, \mathsf{SE\text{-}CCA}, \mathsf{SE\text{-}CPA}, \mathsf{PKE\text{-}CCA}\}$. We follow the definitional paradigm of [4], but there are some delicate primitive-dependent choices that significantly affect the strength of the definitions and the challenge of achieving them (cf. Section 2). We let $\mathbf{RKA}[\mathsf{P}]$ be the set of all Φ for which there exists a Φ-RKA secure instance of P. These sets are all non-trivial.
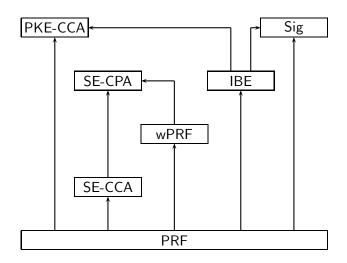
RELATIONS. We establish two kinds of relations between sets $\mathbf{RKA}[\mathsf{P}_1]$ and $\mathbf{RKA}[\mathsf{P}_2]$:

- <u>Containment</u>: A proof that $\mathbf{RKA}[\mathsf{P}_1] \subseteq \mathbf{RKA}[\mathsf{P}_2]$, established by constructing a Φ-RKA secure instance of P_2 from a Φ-RKA secure instance of P_1, usually under the (minimal) additional assumption that one is given a normal-secure instance of P_2. Containments yield constructions of Φ-RKA secure instances of P_2.

- <u>Non-containment</u>: A proof that $\mathbf{RKA}[\mathsf{P}_2] \not\subseteq \mathbf{RKA}[\mathsf{P}_1]$. Here we exhibit a particular Φ for which we (1) construct a Φ-RKA secure instance of P_1 under some reasonable assumption, and (2) show, via attack, that *any* instance of P_2 is Φ-RKA insecure.

We show that RKA-secure PRFs are powerful enablers of RKA-security: Given a Φ-RKA PRF and a normal-secure instance of P, we construct a Φ-RKA secure instance of P for all $\mathsf{P} \in \{\mathsf{wPRF}, \mathsf{IBE}, \mathsf{Sig}, \mathsf{SE\text{-}CCA}, \mathsf{SE\text{-}CPA}, \mathsf{PKE\text{-}CCA}\}$. This is represented by the string of containments in the first row of the table in Figure 1. On the practical side, instantiating the PRF with a blockcipher yields a cheap way to immunize the other primitives against RKAs. On the theoretical side, instantiating the PRF with the construct of [3] yields Φ-RKA secure instances of the other primitives based on standard assumptions.

The separations shown in the first column of the table of Figure 1, however, also show that RKA-PRFs are overkill: *all* the other primitives admit Φ-RKA secure instances for a Φ for which no Φ-RKA PRF exists. This leads one to ask whether there are alternative routes to RKA-secure constructions of beyond-PRF primitives.

We show that IBE is a particularly powerful starting point. We observe that Naor's transform

Figure 1: **Relations between RKA[P] classes.** A containment $\mathbf{RKA}[\mathsf{P}_1] \subseteq \mathbf{RKA}[\mathsf{P}_2]$ is represented in the picture by an arrow $\mathsf{P}_1 \to \mathsf{P}_2$ and in the table by a "$\subseteq$" in the row P_1, column P_2 entry. A non-containment $\mathbf{RKA}[\mathsf{P}_1] \not\subseteq \mathbf{RKA}[\mathsf{P}_2]$ is represented in the table by a "$\not\subseteq$" in the row P_1, column P_2 entry. The picture does not show non-containments. The picture sometimes shows a redundant containment (for example the arrow $\mathsf{PRF} \to \mathsf{Sig}$ when there is already a path $\mathsf{PRF} \to \mathsf{IBE} \to \mathsf{Sig}$) because it corresponds to an interesting direct construction. A blank entry in the table means we do not know.

	PRF	wPRF	IBE	Sig	SE-CCA	SE-CPA	PKE-CCA
PRF	$\subseteq$	$\subseteq$	$\subseteq$	$\subseteq$	$\subseteq$	$\subseteq$	$\subseteq$
wPRF	$\not\subseteq$	$\subseteq$	$\not\subseteq$			$\subseteq$	$\not\subseteq$
IBE	$\not\subseteq$	$\not\subseteq$	$\subseteq$	$\subseteq$	$\not\subseteq$	$\not\subseteq$	$\subseteq$
Sig	$\not\subseteq$	$\not\subseteq$		$\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$
SE-CCA	$\not\subseteq$				$\subseteq$	$\subseteq$	
SE-CPA	$\not\subseteq$		$\not\subseteq$		$\not\subseteq$	$\subseteq$	$\not\subseteq$
PKE-CCA	$\not\subseteq$	$\not\subseteq$			$\not\subseteq$	$\not\subseteq$	$\subseteq$

preserves RKA-security, allowing us to turn a Φ-RKA secure IBE scheme into a Φ-RKA secure Sig scheme. Similarly, we show that the transform of Boneh, Canetti, Halevi and Katz (BCHK) [15] turns a Φ-RKA secure IBE scheme into a Φ-RKA secure PKE-CCA scheme. What lends these transforms well to RKA-security is that they do not change the secret key. We also show that given a Φ-RKA secure wPRF we can build a Φ-RKA secure SE-CPA scheme. (A wPRF is like a PRF except that is only required to be secure on random inputs [32].) These results motivate finding new Φ-RKA secure IBE schemes and wPRFs.

As the table of Figure 1 indicates, we show a number of other non-containments. Sig emerges as a very "RKA-resilient" primitive in the sense that it can be secure against strictly more RKAs than most other primitives. Some of the non-containments, such as $\mathbf{RKA}[\mathsf{PKE\text{-}CCA}] \not\subseteq \mathbf{RKA}[\mathsf{SE\text{-}CPA}]$ might seem odd; doesn't PKE always imply SE? What we are saying is that the trivial transformation of a PKE scheme to an SE one does not preserve RKA-security and, moreover, there are Φ for which *no transform exists* that can do this.

CLAWS OK. All previous constructions of Φ-RKA secure primitives [4, 30, 3, 21, 2, 23, 24] assume Φ is claw-free (distinct functions in ϕ disagree on all inputs) because it is hard to do the proofs otherwise, but the Φ underlying practical fault injection attacks are not claw-free, making it desirable to get constructions avoiding this assumption. For the first time, we are able to do this. In Section 2 we explain the technical difficulties and sketch our solution, which is based on the construction of a Φ-RKA PRG that has a novel property we call identity-collision-resistance (ICR), a variant of the collision-resistance property from [25].

RELATED WORK. The first theoretical treatment of RKAs was by Bellare and Kohno [4]; being inspired by blockciphers, the work addressed PRFs and PRPs. They showed examples of classes not in $\mathbf{RKA}[\mathsf{PRF}]$, gave conditions on Φ for ideal ciphers to be Φ-RKA secure, and provided standard model constructs for some limited classes. Subsequently, constructions of Φ-RKA secure PRFs and PRPs for more interesting Φ were found, first under novel assumptions [30] and then under standard assumptions [3], and the results on ideal ciphers were extended in [1].

We are seeing growing interest in RKA security for primitives other than PRFs. Goldenberg and Liskov [21] study related-secret security of lower-level primitives, namely one-way functions, hardcore bits and pseudorandom generators. Applebaum, Harnik and Ishai [2] define RKA security for (randomized) symmetric encryption, gave several constructions achieving that definition for interesting Φ and then presented numerous applications. Connections with point obfuscation are made by Bitansky and Canetti [10].

Gennaro, Lysyanskaya, Malkin, Micali and Rabin [20] suggest that RKAs may arise by tampering. They show that one can achieve security when related keys are derived via arbitrary key modification, but assume an external trusted authority signs the original secret key and installs the signature on the device together with its own public key, the latter being "off limits" to the attacker. (Meaning, the related-key deriving functions may not modify them.) In our case, no such authority is assumed. The off-limit quantities are confined to pre-installed public parameters. No information that is a function of the parameters and the key is installed on the chip.

Ishai, Prabhakaran, Sahai and Wagner [26] are concerned with tampering of wires in the computation of a circuit while we are concerned with tampering with hardware-stored keys. Dziembowski, Pietrzak and Wichs [19] develop an information theoretic method for preventing tampering and show that a wide class of limited, but non-trivial, Φ can be achieved (unconditionally) for any so-called "interactive stateful system."

INDEPENDENT WORK. Interest in RKA security for higher-level primitives is evidenced by Goyal, O'Neill and Rao [23, 24], who define correlated-input (CI) hash functions, show how to construct them from the q-DHI assumption based on Boneh-Boyen signatures [13, 14] and the Dodis-Yampolskiy PRF [18], and apply this to get Φ-RKA secure signatures from q-DHI for a class Φ consisting of polynomials over a field of prime order. (They indicate their approach would also work for other primitives.) Their construction is similar to ours. Their definitions and results, unlike ours, are restricted to claw-free Φ. Also, we start from Φ-RKA-PRFs and thus get in-practice security for any class Φ for which blockciphers provide them, while they start from a number-theoretic assumption and get security for a specific class Φ, related to the scheme algebra. Their work and ours are concurrent and independent. (Ours was submitted to, and rejected from, Eurocrypt 2011, while theirs was submitted to, and accepted at, TCC 2011.)

Kalai, Kanukurthi and Sahai [27] provide encryption and signature schemes that protect against both tampering and leakage via the idea of key-updates that originated in forward-secure signatures [6]. They allow arbitrary tampering functions but only allow a bounded number of tampering queries within each time period. Their work and ours are again concurrent and independent.

2 Technical approach

Before providing formal definitions, constructions and proofs of our many positive and negative results, we would like to illustrate one technical issue, namely the challenges created by Φ that are not claw-free and how we resolve them. For concreteness, our discussion is restricted to the design of Φ-RKA signatures based on Φ-RKA PRFs.

THE CLAW-FREENESS ASSUMPTION. All known constructions of Φ-RKA-secure primitives [4, 30, 3, 21, 2, 23, 24] are restricted to Φ that are *claw-free*. This means that any two distinct functions in Φ disagree on *all* inputs. This assumption is made for technical reasons; it seems hard to do simulations and proofs without it. Yet the assumption is undesirable, for many natural and practical classes of functions are *not* claw-free. For example, fault injection might be able to set a certain bit of the key to zero, and if Φ contains the corresponding function and the identity function then it is not claw-free. Any Φ that can set the key to a constant value is also not claw-free. Accordingly it is desirable to avoid this assumption. For the first time we are able to do so, via a new technical approach.

DEFINITIONS AND ISSUES. The degree to which claw-freeness is embedded in current approaches is made manifest by the fact that the very *definition* of Φ-RKA secure signatures of [23, 24] assumes it and is unachievable without it. Let us take a closer look to see how.

The signature RKA-security game of [23, 24] picks secret signing key sk and associated public verification key vk. It gives the adversary a signing oracle SIGN that takes m and $\phi \in \Phi$, and returns the signature of message m under key $\phi(sk)$. The adversary eventually outputs m, σ. Besides validity of m, σ under vk, winning requires that m be "new," meaning not "previously signed." The delicate question is, how do we define this? The choice of [23, 24] is to disallow signing query id, m, where id is the identity function. But the adversary can easily define a function ϕ that is the identity on all but a negligible fraction of its inputs. A query ϕ, m is then valid since $\phi \neq \mathsf{id}$, but almost always returns the signature σ of m under sk, so the adversary can output m, σ and win. By assuming Φ is claw-free and contains id, [23, 24] ensure that such a ϕ is not in Φ and the attack is ruled out.

Our altered definition of m being "new" is that there was no signing query ϕ, m with $\phi(sk) = sk$. This seems, indeed, the natural requirement, ruling out nothing more than that m was signed under sk.

We now have a much more general definition that is meaningful even for the non claw-free Φ that arise in practice, but it has a subtle feature that makes achieving it a challenge. Namely, *checking whether the adversary won apparently requires knowing sk* for we have to test whether or not $\phi(sk) = sk$. In the reduction proving security, we will be designing an adversary B attempting to distinguish "real" or "random" instances of some problem given an adversary A breaking the signature scheme; B will see if A won, declaring "real" if so and "random" otherwise. But B will be simulating A and will not know sk, so the difficulty is how it can test that A won.

OVERVIEW OF SOLUTION. We start from a Φ-RKA secure PRF $F \colon \mathcal{K} \times \mathcal{D} \to \mathcal{R}$ that has what we call a key fingerprint for the identity function. This is a relaxation of the notion of a key fingerprint of [3]. It consists of a vector $\mathbf{w}$ over $\mathcal{D}$ such that for all K and all $\phi \in \Phi$ with $\phi(K) \neq K$ there is some i such that $F(K, \mathbf{w}[i]) \neq F(\phi(K), \mathbf{w}[i])$. This allows statistical disambiguation of the original key K from other keys. Such fingerprints exist for the Φ-RKA PRFs of [3] and for blockciphers and are thus a mild assumption.

We now turn F into a PRG (Pseudorandom Generator) $\mathcal{G}$ that has two properties. First, it is Φ-RKA secure; this means the adversary has low advantage in determining the challenge bit b in the game that picks a random target key K and random function R, and then gives the adversary an oracle GEN that on input ϕ returns $\mathcal{G}(\phi(K))$ if $b = 1$ and $R(\phi(K))$ if $b = 0$. This is of course easily obtained from a Φ-RKA PRF. We call the new second property Φ-ICR (Identity-Collision-Resistant); this means that for a hidden key K, it is hard for the adversary to find $\phi \in \Phi$ such that $\phi(K) \neq K$ yet $\mathcal{G}(\phi(K)) = \mathcal{G}(K)$. At first it might seem this follows from Φ-RKA security but Proposition 5.1 shows it does not. However Proposition 5.2 shows how to build a PRG that is both Φ-RKA and Φ-ICR secure from a Φ-RKA PRF with an identity key fingerprint, without assuming Φ is claw-free.

We build our Φ-RKA secure signature scheme from this PRG $\mathcal{G}$ and a base (normal secure) signature scheme, as follows. The secret key of our new signature scheme is a key K for the PRG. The output of the PRG on input K, $\mathcal{G}(K)$, is used as randomness to run the key-generation algorithm $\mathcal{K}$ of the base signature scheme, yielding a public key pk which becomes the public key of our scheme, and the corresponding secret key which is discarded. (Recall the secret key of the new scheme is the PRG key K.) To sign a message m under K, run $\mathcal{G}$ on K to get coins for $\mathcal{K}$, run the latter with these coins to get pk, sk and finally sign m under sk with the base signature scheme. Verification is just as in the base signature scheme.

For the proof we must construct an adversary B breaking the Φ-RKA security of $\mathcal{G}$ given an adversary A breaking the Φ-RKA security of our signature scheme. B thinks of the key K underlying its game as the secret key for our signature scheme and then runs A. When A makes SIGN query ϕ, m, adversary B will call its GEN oracle on ϕ and use the result as coins for $\mathcal{K}$ to get a secret key under which it then signs m for A. Eventually A outputs a forgery attempt m, σ. The assumed security of the base signature scheme will make it unlikely that A's forgery is a winning one when GEN is underlain by a

random function. So B would like to test if A's forgery was a winning one, outputting 1 if so and 0 otherwise, to win its game. The difficulty is that it cannot test this because, not knowing K, it cannot test whether or not A made a SIGN query ϕ, m with $\phi(K) = K$. The Φ-ICR property of $\mathcal{G}$ comes to the rescue, telling us that whether or not $\phi(K) = K$ may be determined by whether or not the outputs of $\mathcal{G}$ on these two inputs, which B *does* have, are the same.

This sketch still pushes under the rug several subtle details which are dealt with in the full proof of Theorem 6.2, to be found in Appendix D.

3 Preliminaries

NOTATION. For sets X, Y, Z let $\mathsf{Fun}(X, Y)$ be the set of all functions mapping X to Y, and let $\mathsf{FF}(X, Y, Z) = \mathsf{Fun}(X \times Y, Z)$. The empty string is denoted ε. If $\mathbf{v}$ is a vector then $|\mathbf{v}|$ denotes the number of its coordinates and $\mathbf{v}[i]$ denotes its i-th coordinate, meaning $\mathbf{v} = (\mathbf{v}[1], \ldots, \mathbf{v}[|\mathbf{v}|])$. A (binary) string x is identified with a vector over $\{0, 1\}$ so that $|x|$ is its length and $x[i]$ is its i-th bit. If $a_1, \ldots, a_n$ are strings then $a_1 \| \cdots \| a_n$ denotes their concatenation. If S is a set then $|S|$ denotes its size and $s \leftarrow\!\!\!{}^{_\$}\, S$ the operation of picking a random element of S and calling it s. We say that a real-valued function on the integers is *negligible* if it vanishes faster than the inverse of any polynomial.

ALGORITHMS. Unless otherwise indicated, an algorithm is PT (Polynomial Time) and may be randomized. An adversary is an algorithm. If A is an algorithm and $\mathbf{x}$ is a vector then $A(\mathbf{x})$ denotes the vector $(A(\mathbf{x}[1]), \ldots, A(\mathbf{x}[|\mathbf{x}|]))$. By $y \leftarrow A(x_1, x_2, \ldots; r)$ we denote the operation of running A on inputs $x_1, x_2, \ldots$ and coins $r \in \{0, 1\}^*$. We denote by $y \leftarrow\!\!\!{}^{_\$}\, A(x_1, x_2, \ldots)$ the operation of picking r at random and letting $y \leftarrow A(x_1, x_2, \ldots; r)$. We denote by $[A(x_1, x_2, \ldots)]$ the set of all possible outputs of A on inputs $x_1, x_2, \ldots$. We denote by $k \in \mathbb{N}$ the security parameter and by 1^k its unary encoding. It is assumed that the length of the output of any algorithm A depends only on the lengths of its inputs. In particular we can associate to single-input algorithm A its *output length* ℓ satisfying $|A(x)| = \ell(|x|)$ for all x. If A, B are algorithms then $A \| B$ denotes the algorithm that on any input x returns $A(x) \| B(x)$.

GAMES. Some of our definitions and proofs are expressed via code-based games [7]. Recall that such a game consists of an INITIALIZE procedure, procedures to respond to adversary oracle queries and a FINALIZE procedure. A game G is executed with an adversary A as follows. First, INITIALIZE executes on input 1^k and its output is the input to A. Then A executes, its oracle queries being answered by the corresponding procedures of G. When A terminates, its output becomes the input to the FINALIZE procedure. The output of the latter, denoted G^A, is called the output of the game. We let "$G^A \Rightarrow d$" denote the event that this game output takes value d. If FINALIZE is absent it is understood to be the identity function, so the game output is the adversary output. Boolean flags are assumed initialized to false.

4 Classes of RKDFs and RKA-PRFs

CLASSES OF RKDFS. In [4], a class Φ of related-key deriving functions (RKDFs) is a finite set of functions, all with the same domain and range. Our more general, asymptotic treatment requires extending this, in particular to allow the functions to depend on public parameters of the scheme. For us a *class* $\Phi = (\mathcal{P}, \mathcal{Q})$ of RKDFs, also called a RKA specification, is a pair of algorithms, the second deterministic. On input 1^k, parameter generation algorithm $\mathcal{P}$ produces parameters π. On input π, a key K and a description ϕ of an RKD function, the evaluation algorithm $\mathcal{Q}$ returns either a modified key or $\perp$. We require that for all ϕ, π, either $\mathcal{Q}(\pi, K, \phi) = \perp$ for all K or for no K. We let $\Phi_{\pi,\phi}(\cdot) = \mathcal{Q}(\pi, \cdot, \phi)$. We require that Φ always includes the identity function. (Formally, there is a special symbol id such that $\Phi_{\pi,\mathsf{id}}(K) = K$ for all K, π. This is to ensure that Φ-RKA security always implies normal security.) We let ID be the class consisting of only the identity function, so that ID-RKA security will be normal security.

8

proc INITIALIZE // PRF	proc INITIALIZE // PRFReal, PRFRand	proc INITIALIZE // IDFP
$\pi \leftarrow_\$ \mathcal{P}(1^k)$; $K \leftarrow_\$ \mathcal{K}(\pi)$	$\pi \leftarrow_\$ \mathcal{P}(1^k)$	$\pi \leftarrow_\$ \mathcal{P}(1^k)$
$b \leftarrow_\$ \{0,1\}$	$K \leftarrow_\$ \mathcal{K}(\pi)$	$K \leftarrow_\$ \mathcal{K}(\pi)$
Return π	Return π	$\mathbf{w} \leftarrow_\$ \mathsf{IKfp}(\pi)$
		Return $\pi, \mathbf{w}$

proc FN(ϕ, x) // PRF	proc FN(ϕ, x) // PRFReal	
$K' \leftarrow \mathcal{Q}(\pi, K, \phi)$	$K' \leftarrow \mathcal{Q}(\pi, K, \phi)$	proc FN(ϕ) // IDFP
If $K' = \bot$ then return $\bot$	If $K' = \bot$ then return $\bot$	$K' \leftarrow \Phi_{\pi,\phi}(K)$
If $b = 1$ then	Return $\mathcal{F}(\pi, K', x)$	If $(K' = \bot)$ then return $\bot$
$\quad T[K', x] \leftarrow \mathcal{F}(\pi, K', x)$		If $(K' \neq K)$ then
If $b = 0$ and $T[K', x] = \bot$ then	proc FN(ϕ, x) // PRFRand	$\quad$ If $(\mathcal{F}(K', \mathbf{w}) = \mathcal{F}(K, \mathbf{w}))$ then
$\quad T[K', x] \leftarrow_\$ \mathsf{Rng}(\pi)$	$K' \leftarrow \mathcal{Q}(\pi, K, \phi)$	$\quad\quad$ WIN $\leftarrow$ true
Return $T[K', x]$	If $K' = \bot$ then return $\bot$	Return $\mathcal{F}(K', \mathbf{w})$
	If $T[K', x] = \bot$ then	
proc FINALIZE(b') // PRF	$\quad T[K', x] \leftarrow_\$ \mathsf{Rng}(\pi)$	proc FINALIZE$()$ // IDFP
Return $(b = b')$	Return $T[K', x]$	Return WIN

Figure 2: Games defining Φ-RKA PRF security and Φ-IDFP security of function family $\mathcal{FF} = (\mathcal{P}, \mathcal{K}, \mathcal{F})$ having range $\mathsf{Rng}(\cdot)$.

A scheme (regardless of the primitive) is a tuple $(\overline{\mathcal{P}}, \cdots)$ of algorithms, the first of which is a parameter generation algorithm that on input 1^k returns a string. If ℓ is the output length of $\mathcal{P}$, we say that $\Phi = (\mathcal{P}, \mathcal{Q})$ is *compatible* with the scheme if the string formed by the first $\ell(k)$ bits of the output of $\overline{\mathcal{P}}(1^k)$ is distributed identically to the output of $\mathcal{P}(1^k)$ for all $k \in \mathbb{N}$. This is done so that, in constructing one Φ-RKA primitive from another, we can extend the parameters of the constructed scheme beyond those of the original one without changing the class of RKDFs.

We say that $\Phi = (\mathcal{P}, \mathcal{Q})$ is *claw-free* if $\phi \neq \phi'$ implies $\mathcal{Q}(\pi, K, \phi) \neq \mathcal{Q}(\pi, K, \phi')$ (or both values are $\bot$) for all π, K. This property has been assumed almost ubiquitously in previous work [4, 30, 21, 3] because of the technical difficulties created by its absence, but its assumption is in fact quite restrictive since many natural classes do not have it. We are able to remove this assumption and provide constructs secure even for non-claw-free classes via new technical approaches. We let **CF** be the set of all Φ that are claw-free.

The class $\Phi^{\mathsf{const}} = (\mathcal{P}, \mathcal{Q}^{\mathsf{const}})$ of constant functions associated to class $\Phi = (\mathcal{P}, \mathcal{Q})$ is defined by $\Phi^{\mathsf{const}}_{\pi,a}(K) = a$ for all $K, a \in \{0,1\}^*$ and all π. The union $\Phi^1 \cup \Phi^2 = (\mathcal{P}, \mathcal{Q})$ of classes $\Phi^1 = (\mathcal{P}, \mathcal{Q}^1)$ and $\Phi^2 = (\mathcal{P}, \mathcal{Q}^2)$ is defined by having $\mathcal{Q}(\pi, K, \phi)$ parse ϕ as $i \, \| \, \phi^*$ for $i \in \{1, 2\}$ and return $\mathcal{Q}^i(\pi, K, \phi^*)$.

DISCUSSION. In a non-asymptotic treatment, there is no formal line between "secure" and "insecure." This makes it unclear how to rigorously define the sets **RKA**[P]. Lead, accordingly, to pursue an asymptotic treatment, we introduce parameter dependence; this allows us to capture constructs in the literature [30, 3] where RKDFs are defined over a group that is now parameter-dependent rather than fixed. (We note that even in the non-asymptotic case, a treatment like ours is needed to capture constructs in [30] relying on a RSA group defined by random primes. This issue is glossed over in [30].) A dividend of our treatment is a separation between an RKDF and its encoding, the latter being what an adversary actually queries, another issue glossed over in previous work.

FUNCTION FAMILIES. A function family $\mathcal{FF} = (\mathcal{P}, \mathcal{K}, \mathcal{F})$ consists of a parameter generator, a key generator, and an evaluator, the last deterministic. For each $k \in \mathbb{N}$ and $\pi \in [\mathcal{P}(1^k)]$, the scheme also defines PT decidable and sampleable sets $\mathsf{Dom}(\pi)$ and $\mathsf{Rng}(\pi)$ such that $\mathcal{F}(\pi, K, \cdot)$ maps elements of $\mathsf{Dom}(\pi)$ to $\mathsf{Rng}(\pi)$. We assume there are polynomials d, l, called the input and output lengths, respectively, such that $\mathsf{Dom}(\pi) \subseteq \{0,1\}^{d(k)}$ and $\mathsf{Rng}(\pi) \subseteq \{0,1\}^{l(k)}$. Unless otherwise indicated we assume $\mathsf{Rng}(\pi) = \{0,1\}^{l(k)}$ and $l(k) = \omega(\log(k))$ and $|\mathsf{Dom}(\pi)| \geq 2^k$ for all $\pi \in [\mathcal{P}(1^k)]$ and all $k \in \mathbb{N}$.

RKA-PRFs. Let $\mathcal{FF} = (\mathcal{P}, \mathcal{K}, \mathcal{F})$ be a function family as above. Game PRF of Figure 2 is associated to $\mathcal{FF}$ and a RKA specification Φ that is compatible with $\mathcal{FF}$. Let $\mathbf{Adv}^{\mathrm{prf\text{-}rka}}_{\mathcal{FF}, A, \Phi}(k)$ equal $2 \Pr[\mathrm{PRF}^A \Rightarrow$

true] $- 1$ when the game has input 1^k. We say $\mathcal{FF}$ is Φ-RKA secure if this advantage function is negligible. For our proofs it is useful to also consider games PRFReal, PRFRand of Figure 2. Standard arguments imply that $\mathbf{Adv}^{\text{prf-rka}}_{\mathcal{FF},A,\Phi}(k)$ equals $\Pr\left[\,\text{PRFReal}^A \Rightarrow 1\,\right] - \Pr\left[\,\text{PRFRand}^A \Rightarrow 1\,\right]$.

IDENTITY KEY FINGERPRINTS. An identity key fingerprint function with vector length $v(\cdot)$ for $\mathcal{FF} = (\mathcal{P}, \mathcal{K}, \mathcal{F})$ is an algorithm IKfp that for every $\pi \in [\mathcal{P}(1^k)]$ and every $k \in \mathbb{N}$ returns, on input π, a $v(k)$-vector over $\mathsf{Dom}(\pi)$ all of whose coordinates are distinct. Game IDFP of Figure 2 is associated to $\mathcal{FF}$ and a RKA specification $\Phi = (\mathcal{P}, \mathcal{Q})$ that is compatible with $\mathcal{FF}$. Let $\mathbf{Adv}^{\text{idfp}}_{\mathcal{FF},A,\Phi}(k)$ equal $\Pr[\text{IDFP}^A \Rightarrow \mathsf{true}]$ when the game has input 1^k. We say $\mathcal{FF}$ is Φ-IDFP secure if this advantage function is negligible.

The key fingerprint notion of [3] can be seen as allowing statistical disambiguation of any pair of keys. They showed that the Naor-Reingold PRF NR had such a fingerprint, but in general, it does not seem common. Interestingly, their own Φ-RKA PRFs, which build on NR, are not known to have such a fingerprint. Our relaxation can be seen as asking for computational disambiguation of the original key from other keys, and ends up being much easier to achieve. In particular, such fingerprints exist for the constructs of [3]. This is a consequence of something more general, namely that any Φ-RKA secure PRF with large enough range is Φ-IDFP secure if Φ is claw-free, using *any* point in the domain functioning as the fingerprint. This is formalized by Proposition 4.1 below, with a proof in Appendix A. Φ-IDFP security for the constructs of [3] follows as the Φ they use is claw-free.

Proposition 4.1 *Suppose Φ is claw-free and $\mathcal{FF}$ is a Φ-RKA secure PRF with associated domain $\mathsf{Dom}(\cdot)$ and super-polynomial size range $\mathsf{Rng}(\cdot)$. Let IKfp be any algorithm that on input π returns a 1-vector over $\mathsf{Dom}(\pi)$. Then $\mathcal{FF}$ is Φ-IDFP secure.*

In practice Φ-IDFP security seems like a mild assumption even when Φ is not claw-free. A vector of a few, distinct domain points ought to be a suitable fingerprint for any practical blockcipher. This does not follow from a standard assumption on it such as PRF, but is consistent with properties assumed by cryptanalysts and can be proved in the ideal cipher model.

Φ-IDFP security of given Φ-RKA PRFs, even for non-claw-free Φ, will be important in the constructions underlying our containment results, and we make it a default assumption on a Φ-RKA PRF. The above shows that this is a mild and reasonable assumption.

RKA SETS. We say that an RKA specification $\Phi = (\mathcal{P}, \mathcal{Q})$ is achievable for the primitive PRF if there exists a Φ-RKA and Φ-IDFP secure PRF that is compatible with Φ. We let $\mathbf{RKA}[\mathsf{PRF}]$ be the set of all Φ that are achievable for PRF.

WHAT CAN ATTACKS MODIFY? We view the system as a whole as having the following components: algorithms (code), parameters, public keys (if any) and secret keys. Of these, our convention is that only secret keys are subject to RKAs. This is not the only possible model, nor is it necessarily the most realistic if considering tampering attacks in practice, but it is a clear and interesting one with some justification. Parameters are systemwide, meaning fixed beforehand and independent of users, and may, in an implementation, be part of the algorithm code. Public keys are accompanied by certificates under a CA public key that is in the parameters, so if parameters are safe, so are public keys. This leaves secret keys as the main target. One consequence of this is that in a public key setting the attack is only on the holder of the secret key, meaning the signer for signatures and the receiver for encryption, while in the symmetric setting, both sender and receiver are under attack, making this setting more complicated.

We could consider attacks on public keys, but these are effectively attacks on parameters. Furthermore the only way for them to succeed is to modify the CA public key in the parameters in a rather special way, replacing it by some other key under which the attack produces signatures for the modified public key. "Natural" attacks caused by fault-injection are unlikely to do this, further supporting our convention of confining attacks to secret keys.

proc INITIALIZE // PRG	proc INITIALIZE // ICR

proc INITIALIZE // PRG
$\pi \leftarrow_\$ \mathcal{P}(1^k)$
$K \leftarrow_\$ \mathcal{K}(\pi)$; $b \leftarrow_\$ \{0,1\}$
Return π

proc GEN(ϕ) // PRG
$K' \leftarrow \Phi_{\pi,\phi}(K)$
If $K' = \bot$ then return $\bot$
If $T[K'] = \bot$ then
  If $b = 1$ then $T[K'] \leftarrow \mathcal{G}(\pi, K')$
  Else $T[K'] \leftarrow_\$ \{0,1\}^{r(k)}$
Return $T[K']$

proc FINALIZE(b') // PRG
Return $(b = b')$

proc INITIALIZE // ICR
$\pi \leftarrow_\$ \mathcal{P}(1^k)$
$K \leftarrow_\$ \mathcal{K}(\pi)$; $T_0 \leftarrow \mathcal{G}(\pi, K)$
Return π

proc GEN(ϕ) // ICR
$K' \leftarrow \Phi_{\pi,\phi}(K)$
If $K' = \bot$ then return $\bot$
$S \leftarrow \mathcal{G}(\pi, K')$
If $((S = T_0) \wedge (K \neq K'))$ then WIN $\leftarrow$ true
Return S

proc FINALIZE() // ICR
Return WIN

Figure 3: Games defining Φ-RKA security and identity-collision-resistance for PRG $\mathcal{PRG} = (\mathcal{P}, \mathcal{K}, \mathcal{G}, r)$.

5 ICR PRGs: A tool in our constructions

We will be using Φ-RKA PRFs to build Φ-RKA instances of many other primitives. An important technical difficulty will be to avoid assuming Φ is claw-free. A tool we introduce and use for this purpose is a Φ-RKA PRG satisfying a weak form of collision-resistance under RKA that we call Φ-ICR. In this section we define this primitive and show how to achieve it based on a Φ-RKA and Φ-IDFP secure PRF.

RKA PRGs. A PRG $\mathcal{PRG} = (\mathcal{P}, \mathcal{K}, \mathcal{G}, r)$ is specified by a parameter generation algorithm, a key generation algorithm, an evaluation algorithm and an output length $r(\cdot)$. Game PRG of Figure 3 is associated to $\mathcal{PRG}$ and an RKA specification Φ that is compatible with $\mathcal{PRG}$. Let $\mathbf{Adv}^{\mathrm{prg}}_{\mathcal{PRG},A,\Phi}(k) = 2\Pr[\mathrm{PRG}^A \Rightarrow \mathsf{true}] - 1$ when the game has input 1^k. We say $\mathcal{PRG}$ is Φ-RKA secure if this advantage function is negligible for all A.

We clarify that unlike a normal PRG [11], we don't require a Φ-RKA PRG to be length extending, meaning that outputs need not be longer than inputs. The utility of the primitive is getting RKA security, for which length extension is irrelevant. If one does want a length extending Φ-RKA PRG (we won't) one can get it by applying a normal-secure PRG to the output of a given Φ-RKA PRG.

ICR. We define and use a weak form of collision-resistance for PRGs which requires that the adversary be unable to find ϕ so that $\Phi_{\pi,\phi}(K) \neq K$ yet $\mathcal{G}(\Phi_{\pi,\phi}(K)) = \mathcal{G}(K)$. Game ICR of Figure 3 is associated to $\mathcal{PRG}$ and a RKA specification Φ that is compatible with $\mathcal{PRG}$. Let $\mathbf{Adv}^{\mathrm{icr}}_{\mathcal{PRG},C,\Phi}(k)$ equal $2\Pr[\mathrm{ICR}^C \Rightarrow \mathsf{true}] - 1$ when the game has input 1^k. We say $\mathcal{PRG}$ is Φ-ICR (Identity-Collision-Resistant) secure if this advantage function is negligible.

DOES RKA SECURITY IMPLY ICR SECURITY? At first glance it would seem that if a PRG $\mathcal{PRG} = (\mathcal{P}, \mathcal{K}, \mathcal{G}, r)$ is Φ-RKA secure then it is also Φ-ICR secure. Indeed, suppose an adversary has ϕ such that $\Phi_{\pi,\phi}(K) \neq K$ yet $\mathcal{G}(\Phi_{\pi,\phi}(K)) = \mathcal{G}(K)$. Let it query $R_0 \leftarrow \mathrm{GEN}(\mathsf{id})$ and $R_1 \leftarrow \mathrm{GEN}(\phi)$ and return 1 if $R_0 = R_1$ and 0 otherwise. In the real ($b = 1$) case R_0, R_1 are equal but in the random ($b = 0$) case they would appear very unlikely to be equal, so that that this strategy would appear to have high advantage in breaking the Φ-RKA security of $\mathcal{PRG}$. The catch is in our starting assumption, which made it appear that $\Phi_{\pi,\phi}(K) \neq K$ yet $\mathcal{G}(\Phi_{\pi,\phi}(K)) = \mathcal{G}(K)$ was an absolute fact, true both for $b = 0$ and $b = 1$. If $\Phi_{\pi,\phi}(K)$ and K are different in the real game but equal in the random game, the adversary sees an output collision in both cases and its advantage disappears. Can this actually happen? It can, and indeed the claim (that Φ-RKA security implies Φ-ICR security) is actually false:

Proposition 5.1 *Suppose there exists a normal-secure PRG $\overline{\mathcal{PRG}} = (\overline{\mathcal{P}}, \overline{\mathcal{K}}, \overline{\mathcal{G}}, r)$ with $r(\cdot) = \omega(\log(\cdot))$. Then there exists a PRG $\mathcal{PRG} = (\overline{\mathcal{P}}, \mathcal{K}, \mathcal{G}, r)$ and a class Φ such that $\mathcal{PRG}$ is Φ-RKA secure but $\mathcal{PRG}$*

is not Φ-ICR secure.

A proof is in Appendix B. Briefly, the constructed PRG $\mathcal{PRG}$ adds a redundant bit to the seed of $\overline{\mathcal{PRG}}$ so that seeds differing only in their first bits yield the same outputs, meaning create non-trivial collisions. But Φ is crafted so that that its members deviate from the identity function only in the real game, so that output collisions appear just as often in both cases but in the real game they are non-trivial while in the random game they are trivial. The difficulty is showing how to do this while retaining Φ-RKA security.

CONSTRUCTION. We saw above that not all Φ-RKA PRGs are Φ-ICR secure. Our containments will rely crucially on ones that are. We obtain them from Φ-RKA PRFs that have key fingerprints for the identity function:

Proposition 5.2 *Let $\mathcal{FF} = (\mathcal{P}, \mathcal{K}, \mathcal{F})$ be a Φ-RKA PRF with output length l. Let IKfp be a Φ-IDFP secure identity key fingerprint function for $\mathcal{FF}$ with vector length v. Let $r = lv$ and let $\overline{\mathcal{K}}$, on input $\pi \,\|\, \mathbf{w}$, return $\mathcal{K}(\pi)$. Define PRG $\mathcal{PRG} = (\mathcal{P} \,\|\, \mathsf{IKfp}, \overline{\mathcal{K}}, \mathcal{G}, r)$ via*

$$\mathcal{G}(\pi \,\|\, \mathbf{w}, K) = \mathcal{F}(\pi, K, \mathbf{w}[1]) \,\|\, \cdots \,\|\, \mathcal{F}(\pi, K, \mathbf{w}[\|\mathbf{w}\|]) \ .$$

Then $\mathcal{PRG}$ is Φ-RKA secure and Φ-ICR secure.

The proof is relatively straightforward and is omitted.

6 Relations

We first present a containment and a non-containment related to Sig. Then, we illustrate a non-containment related to wPRFs that uses a different technique. Then we turn to IBE-related results. We then briefly indicate how other relations may be obtained via the same ideas.

SIGNATURES. A signature scheme $\mathcal{DS} = (\mathcal{P}, \mathcal{K}, \mathcal{S}, \mathcal{V})$ is specified as usual by its parameter generation, key generation, signing and verifying algorithms. Game Sig of Figure 4 is associated to $\mathcal{DS}$ and an RKA specification Φ that is compatible with $\mathcal{DS}$. Let $\mathbf{Adv}^{\text{sig-rka}}_{\mathcal{DS}, A, \Phi}(k) = \Pr[\text{Sig}^A \Rightarrow \mathsf{true}]$ when the game has input 1^k. We say $\mathcal{DS}$ is Φ-RKA secure if this advantage function is negligible. Normal security of a signature scheme is recovered by considering Φ that contains only the identity function. One feature of the definition worth highlighting is the way we decide which messages are not legitimate forgeries. They are the ones signed with the real key sk, which means that oracle SIGN needs to check when a related key equals the real one and record the corresponding message, which is a source of challenges in reduction-based proofs.

ATTACKS. It is instructive to see an attack that shows that there are some (quite simple) Φ such that *no* signature scheme is Φ-RKA secure, meaning $\Phi \notin \mathbf{RKA}[\text{Sig}]$. This indicates that the set $\mathbf{RKA}[\text{Sig}]$ is non-trivial. Our attack is adapted from [5, 20]. Let $\mathcal{DS} = (\mathcal{P}, \mathcal{K}, \mathcal{S}, \mathcal{V})$ be a signature scheme. An RKD function is described by a pair (i, sk^*) where i is an integer and sk^* is a signing key. We let $\Phi_{\pi,(i,sk^*)}(sk)$ return sk if $sk[i] = 1$ and sk^* otherwise. On input π, vk, adversary A will generate a fresh key pair via $(vk^*, sk^*) \leftarrow_{\$} \mathcal{K}(\pi)$. It will pick a message m and then for each i it will query $\text{SIGN}(i, m)$ to get back a signature σ_i. If $\mathcal{V}(\pi, vk, m, \sigma_i) = 1$ then it sets $x[i] = 1$ else it sets $x[i] = 0$. At the end of the loop, it has recovered in x the target secret key sk. (Unless $sk[i] = 0$ yet σ_i is valid under vk for some i, but in this case the adversary would have broken the signature scheme under a normal, meaning not-RKA, attack, and thus we can assume this is unlikely.) Once it has sk it can of course trivially win. Similar attacks can be presented for other primitives. We turn now to positive results.

FROM Φ-RKA PRGS TO Φ-RKA SIGNATURES. We will prove containments of the form $\mathbf{RKA}[\text{PRF}] \subseteq \mathbf{RKA}[\text{P}]$ by proving $\mathbf{RKA}[\text{PRG}] \subseteq \mathbf{RKA}[\text{P}]$ and exploiting the fact that $\mathbf{RKA}[\text{PRF}] \subseteq \mathbf{RKA}[\text{PRG}]$.

   We start with a Φ-RKA PRG $\mathcal{PRG} = (\mathcal{P}, \mathcal{K}, \mathcal{G}, r)$ and a normal-secure signature scheme $\overline{\mathcal{DS}} = (\overline{\mathcal{P}}, \overline{\mathcal{K}}, \overline{\mathcal{S}}, \overline{\mathcal{V}})$ such that $r(\cdot)$ is the number of coins used by $\overline{\mathcal{K}}$. We now build another signature scheme $\mathcal{DS} = (\mathcal{P} \,\|\, \overline{\mathcal{P}}, \mathcal{K}', \mathcal{S}, \mathcal{V})$ as follows:

proc INITIALIZE // Sig
$\pi \leftarrow_\$ \mathcal{P}(1^k) \,;\; M \leftarrow \emptyset$
$(vk, sk) \leftarrow_\$ \mathcal{K}(\pi)$
Return (π, vk)

proc SIGN(ϕ, m) // Sig
$sk' \leftarrow \Phi_{\pi,\phi}(sk)$
If $sk' = \bot$ then return $\bot$
If $sk' = sk$ then $M \leftarrow M \cup \{m\}$
Return $\sigma \leftarrow_\$ \mathcal{S}(\pi, sk', m)$

proc FINALIZE(m, σ) // Sig
Return $((\mathcal{V}(\pi, vk, m, \sigma) = 1) \wedge (m \notin M))$

proc INITIALIZE // IBE
$\pi \leftarrow_\$ \mathcal{P}(1^k) \,;\; (mpk, msk) \leftarrow_\$ \mathcal{M}(\pi)$
$b \leftarrow_\$ \{0,1\} \,;\; id^* \leftarrow \bot \,;\; S \leftarrow \emptyset$
Return (π, mpk)

proc KD(ϕ, id) // IBE
$msk' \leftarrow \Phi_{\pi,\phi}(msk)$
If $msk' = \bot$ then return $\bot$
If $msk' = msk$ then $S \leftarrow S \cup \{id\}$
If $(msk' = msk) \wedge (id = id^*)$ then return $\bot$
Return $dk \leftarrow_\$ \mathcal{K}(\pi, mpk, msk', id)$

proc LR(id, m_0, m_1) // IBE
If $|m_0| \neq |m_1|$ then return $\bot$
$id^* \leftarrow id \,;\;$ If $id^* \in S$ then return $\bot$
Return $C \leftarrow_\$ \mathcal{E}(\pi, mpk, id, m_b)$

proc FINALIZE(b') // IBE
Return $(b = b')$

Figure 4: Games defining Φ-RKA security for primitives Sig, IBE.

1. **Parameters**: Parameters for $\mathcal{DS}$ are the concatenation $\pi \,\|\, \overline{\pi}$ of independently generated parameters for $\mathcal{PRG}$ and $\overline{\mathcal{DS}}$.

2. **Keys**: Pick a random seed $K \leftarrow_\$ \mathcal{K}(\pi)$ and let $(\overline{vk}, \overline{sk}) \leftarrow \overline{\mathcal{K}}(\overline{\pi}; \mathcal{G}(K))$ be the result of generating verifying and signing keys with coins $\mathcal{G}(K)$. The new signing key is K and the verifying key remains $\overline{vk}$. (Key $\overline{sk}$ is discarded.)

3. **Signing**: To sign message m with signing key K, recompute $(\overline{vk}, \overline{sk}) \leftarrow \overline{\mathcal{K}}(\overline{\pi}; \mathcal{G}(K))$ and then sign m under $\overline{\mathcal{S}}$ using $\overline{sk}$.

4. **Verifying**: Just as in the base scheme, namely, verify that σ is a base scheme signature of m under $\overline{vk}$ using $\overline{\mathcal{V}}$.

Signature scheme $\mathcal{DS}$ remains compatible with Φ since the parameters of $\mathcal{PRG}$ prefix those of $\mathcal{DS}$.

We want $\mathcal{DS}$ to inherit the Φ-RKA security of $\mathcal{PRG}$. In fact we will show more, namely that $\mathcal{DS}$ is $(\Phi \cup \Phi_c)$-RKA secure where Φ_c is the class of constant RKDFs associated to Φ. The intuition is deceptively simple. A signing query ϕ, m of an adversary A attacking $\mathcal{DS}$ results in a signature of m under what is effectively a fresh signing key, since it is generated using coins $\mathcal{G}(\phi(K))$ that are computationally independent of $\mathcal{G}(K)$ due to the assumed Φ-RKA security of the PRG. These can accordingly be simulated without access to K. On the other hand, signing queries in which ϕ is a constant function may be directly simulated. The first difficulty is that the adversary attacking the Φ-RKA security of $\mathcal{PRG}$ that we must build needs to know when A succeeds, and for this it needs to know when a derived seed equals the real one, and it is unclear how to do this without knowing the real seed. The second difficulty is that a queried constant might equal the key. We take an incremental approach to showing how these difficulties are resolved, beginning by assuming Φ is claw-free, which makes the first difficulty vanish:

Theorem 6.1 *Let signature scheme $\mathcal{DS} = (\mathcal{P} \,\|\, \overline{\mathcal{P}}, \mathcal{K}', \mathcal{S}, \mathcal{V})$ be constructed as above from Φ-RKA PRG $\mathcal{PRG} = (\mathcal{P}, \mathcal{K}, \mathcal{G}, r)$ and normal-secure signature scheme $\overline{\mathcal{DS}} = (\overline{\mathcal{P}}, \overline{\mathcal{K}}, \overline{\mathcal{S}}, \overline{\mathcal{V}})$ and assume Φ is claw-free. Then $\mathcal{DS}$ is $(\Phi \cup \Phi_c)$-RKA secure.*

A proof of Theorem 6.1 is in Appendix C, and the intuition was discussed in Section 2. This result, however, is weaker than we would like, for, as we have already said, many interesting classes are not claw-free. Also, this result fails to prove $\mathbf{RKA}[\mathsf{PRF}] \subseteq \mathbf{RKA}[\mathsf{Sig}]$ since the first set may contain Φ that are not claw-free. To address this we show that the claw-freeness assumption on Φ can be replaced by the assumption that $\mathcal{PRG}$ is Φ-ICR secure:

Theorem 6.2 *Let signature scheme $\mathcal{DS} = (\mathcal{P} \| \overline{\mathcal{P}}, \mathcal{K}', \mathcal{S}, \mathcal{V})$ be constructed as above from Φ-RKA secure and Φ-ICR secure PRG $\mathcal{PRG} = (\mathcal{P}, \mathcal{K}, \mathcal{G}, r)$ and normal-secure signature scheme $\overline{\mathcal{DS}} = (\overline{\mathcal{P}}, \overline{\mathcal{K}}, \overline{\mathcal{S}}, \overline{\mathcal{V}})$. Then $\mathcal{DS}$ is $(\Phi \cup \Phi_c)$-RKA secure.*

A proof of Theorem 6.2 is in Appendix D. Proposition 5.2 says we can get the PRGs we want from Φ-RKA PRFs so Theorem 6.2 establishes the containment $\mathbf{RKA}[\mathsf{PRF}] \subseteq \mathbf{RKA}[\mathsf{Sig}]$. (Theorem 6.1 only established $\mathbf{RKA}[\mathsf{PRF}] \cap \mathbf{CF} \subseteq \mathbf{RKA}[\mathsf{Sig}] \cap \mathbf{CF}$.)

Our construction has the advantage that the verification process as well as the form of the signatures and public key are unchanged. This means it has minimal impact on software, making it easier to deploy than a totally new scheme. Signing in the scheme now involves evaluation of a Φ-RKA-PRG but this can be made cheap via an AES-based instantiation. However, signing also involves running the key-generation algorithm $\overline{\mathcal{K}}$ of the base scheme which might be expensive.

This construction also meets a stronger notion of Φ-RKA security where the adversary cannot even forge a signature relative to the public keys associated with the derived secret keys. We elaborate on this in Appendix E.

Some base signature schemes lend themselves naturally and directly to immunization against RKAs via Φ-RKA PRFs. This is true for the binary-tree, one-time signature based scheme discussed in [22], where the secret key is already that of a PRF. If the latter is Φ-RKA secure we can show the signature scheme (unmodified) is too, and moreover also meets the strong version of the definition alluded to above. See Appendix E.

SEPARATING Φ-RKA PRFs FROM Φ-RKA SIGNATURES. Having just shown that $\mathbf{RKA}[\mathsf{PRF}] \subseteq \mathbf{RKA}[\mathsf{Sig}]$ it is natural to ask whether the converse is true as well, meaning whether the sets are equal. The answer is no, so $\mathbf{RKA}[\mathsf{Sig}] \not\subseteq \mathbf{RKA}[\mathsf{PRF}]$. The interpretation is that there exist Φ such that there exist Φ-RKA secure signatures, but there are *no* Φ-RKA PRFs. An example is when $\Phi = \Phi_c$ is the set of constant functions. Theorem 6.1 implies that there exists a Φ_c-RKA secure signature scheme by setting $\Phi = \emptyset$ in the theorem, so that $\mathcal{PRG}$ need only be a normal-secure PRG. But attacks from [4] show that no PRF can be Φ_c-RKA secure. Thus, this separation is quite easily obtained. We will later see others which are technically more interesting. This separation motivates finding other avenues to Φ-RKA signatures. Below we will show that IBE is one such avenue.

wPRF. We turn to wPRFs so that we can illustrate an interesting separation, namely with PRFs. Let $\mathcal{wPRF} = (\mathcal{P}, \mathcal{K}, \mathcal{F})$ be a family of functions as defined in Section 3 and $\mathsf{Dom}_{\mathcal{wPRF}}(\pi), \mathsf{Rng}_{\mathcal{wPRF}}(\pi)$ the finite domain and range sets associated to parameters π. Game wPRF of Figure 5 is associated to $\mathcal{wPRF}$ and an RKA specification $\Phi = (\mathcal{P}, \mathcal{Q})$ that is compatible with $\mathcal{wPRF}$. Let $\mathbf{Adv}^{\mathrm{wprf\text{-}rka}}_{\mathcal{wPRF}, A, \Phi}(k)$ equal $2\Pr[\mathrm{wPRF}^A \Rightarrow \mathsf{true}] - 1$ when the game has input 1^k. We say $\mathcal{wPRF}$ is Φ-RKA secure if this advantage function is negligible for all A.

SEPARATING Φ-RKA PRFs FROM Φ-RKA wPRFs. Obviously $\mathbf{RKA}[\mathsf{PRF}] \subseteq \mathbf{RKA}[\mathsf{wPRF}]$. We show that the converse is not true, namely $\mathbf{RKA}[\mathsf{wPRF}] \not\subseteq \mathbf{RKA}[\mathsf{PRF}]$. In this case constant RKDFs do *not* provide the separation since both primitives are insecure under these functions. Instead we create RKDFs which only help the attacker if the function may be invoked on the same input for different RKDFs, which is not possible in the wPRF game. Proceeding to the details, let $\mathcal{wPRF} = (\mathcal{P}, \mathcal{K}, \mathcal{F})$ be a normal-secure wPRF and for simplicity assume that keys are random $(k-1)$-bit strings. (Formally, $\mathcal{K}(\pi)$ induces the uniform distribution on $\{0, 1\}^{k-1}$ for all $\pi \in [\mathcal{P}(1^k)]$ and all $k \in \mathbb{N}$.) If $K \in \{0, 1\}^*$ we let K^- denote the string that is K with the first bit flipped. We construct a class $\Phi = (\mathcal{P}, \mathcal{Q})$ of RKDFs as follows. On input $\pi \in [\mathcal{P}(1^k)]$, key $K \in \{0, 1\}^*$ and description ϕ_i, algorithm $\mathcal{Q}$ returns K if $K[i] = 1$ and K^- otherwise, for $1 \le i \le |K|$. Also given π, K and description flip it returns K^-.

Let $\mathcal{PRF} = (\overline{\mathcal{P}}, \overline{\mathcal{K}}, \overline{\mathcal{F}})$ be *any* PRF that is compatible with Φ. We show that it can be broken under a Φ-RKA. Let $\ell(k)$ denote the length of keys when the security parameter is k. Our adversary, given parameters $\overline{\pi}$, picks x at random from $\mathsf{Dom}_{\mathcal{PRF}}(\overline{\pi})$. It queries $y_{\mathsf{id}} \leftarrow \mathrm{FN}(\mathsf{id}, x)$. Then for each $1 \le i \le \ell(k)$ it queries $y_i \leftarrow \mathrm{FN}(\phi_i, x)$ and sets $K'[i] = 1$ if $y_i = y_{\mathsf{id}}$ and 0 otherwise. (It is crucial that

proc INITIALIZE // PKE
———————————————
$\pi \leftarrow_\$ \mathcal{P}(1^k)$
$b \leftarrow_\$ \{0,1\}$; $C^* \leftarrow \bot$
$(ek, dk) \leftarrow_\$ \mathcal{K}(\pi)$
Return (π, ek)

proc DEC(ϕ, C) // PKE
———————————————
$dk' \leftarrow \Phi_{\pi,\phi}(dk)$
If $dk' = \bot$ then return $\bot$
If $((dk' = dk) \wedge (C = C^*))$
   then return $\bot$
Return $M \leftarrow \mathcal{D}(\pi, dk', C)$

proc LR(m_0, m_1) // PKE
———————————————
If $(|m_0| \neq |m_1|)$ then return $\bot$
Return $C^* \leftarrow_\$ \mathcal{E}(\pi, ek, m_b)$

proc FINALIZE(b') // PKE
———————————————
Return $(b = b')$

---

proc INITIALIZE // wPRF
———————————————
$\pi \leftarrow_\$ \mathcal{P}(1^k)$
$K \leftarrow_\$ \mathcal{K}(\pi)$
$b \leftarrow_\$ \{0,1\}$
Return π

proc FN(ϕ) // wPRF
———————————————
$K' \leftarrow \Phi_{\pi,\phi}(K)$
If $K' = \bot$ then return $\bot$
$x \leftarrow_\$ \mathsf{Dom}_{w\mathcal{PRF}}(\pi)$
If $b = 1$ then $y \leftarrow \mathcal{F}(\pi, K', x)$
Else $y \leftarrow_\$ \mathsf{Rng}_{w\mathcal{PRF}}(\pi)$
Return (x, y)

proc FINALIZE(b') // wPRF
———————————————
Return $(b = b')$

---

proc INITIALIZE // SECPA, SECCA
———————————————
$\pi \leftarrow_\$ \mathcal{P}(1^k)$; $b \leftarrow_\$ \{0,1\}$; $S \leftarrow \emptyset$
$K \leftarrow_\$ \mathcal{K}(\pi)$
Return π

proc DEC(ϕ, C) // SECCA
———————————————
$K' \leftarrow \Phi_{\pi,\phi}(K)$
If $K' = \bot$ then return $\bot$
If $((K', C) \in S)$ then return $\bot$
Return $M \leftarrow \mathcal{D}(\pi, K', C)$

proc LR(ϕ, m_0, m_1) // SECPA, SECCA
———————————————
If $(|m_0| \neq |m_1|)$ then return $\bot$
$K' \leftarrow \Phi_{\pi,\phi}(K)$
If $K' = \bot$ then return $\bot$
$C^* \leftarrow_\$ \mathcal{E}(\pi, K', m_b)$; $S \leftarrow S \cup \{(K', C^*)\}$
Return C^*

proc FINALIZE(b') // SECPA, SECCA
———————————————
Return $(b = b')$

Figure 5: Games defining Φ-RKA security for primitives PKE-CCA, wPRF, SE-CPA, SE-CCA.

---

all queries use the same x.) The result should be that $K' = K$, meaning A has recovered the key. Once it has the key, it can easily win by making a few queries under id at random inputs and returning 1 if the results are consistent with K' and 0 otherwise.

There is one subtlety above, however. Namely if $\overline{\mathcal{F}}(\overline{\pi}, K, x) = \overline{\mathcal{F}}(\overline{\pi}, K^-, x)$, meaning the functions under keys K and K^- agree at x, then the attack will recover the string $K' = 1^{\ell(k)}$ and not recover the key. This is in general problematic because a PRF may very well have the property that the functions defined by two different keys are the same. It is to resolve this problem that we put flip in Φ. Φ-RKA security now ensures that the functions defined by keys K, K^- look like random, independent ones. Thus, our adversary can simply return 1 if $K' = 1^{\ell(k)}$, meaning declare that the challenge bit was 1.

What remains is to show that there exists a wPRF that is Φ-RKA secure. We construct $\overline{w\mathcal{PRF}} = (\mathcal{P}, \overline{\mathcal{K}}, \overline{\mathcal{F}})$ based on the normal-secure $w\mathcal{PRF} = (\mathcal{P}, \mathcal{K}, \mathcal{F})$ whose existence we assumed at the start. Parameter generation is unchanged. Keys are random k-bit strings. Function $\overline{\mathcal{F}}(\pi, K, x)$ lets L be the last $(k-1)$-bits of K and returns $\mathcal{F}(\pi, L, x)$. The result is that functions $\overline{\mathcal{F}}(\pi, K, \cdot)$ and $\overline{\mathcal{F}}(\pi, K^-, \cdot)$ are identical. In the wPRF case this, rather than contradicting Φ-RKA security as in the PRF case, helps us prove it. We claim we can reduce the Φ-RKA security of $\overline{w\mathcal{PRF}}$ to the normal security of $w\mathcal{PRF}$. The reason is that on input K, all the RKDFs we put in Φ return either K or K^- and thus no matter what RKDFs the adversary queries the result can be simulated by access to $\overline{\mathcal{F}}(\pi, K, \cdot)$. As long as no input is repeated, the simulation will be correct, and since the domain has size at least 2^k and inputs are random rather than adversary-selected, the probability of a repeated input in q queries is at most $q^2/2^k$ which is negligible. We omit the details.

IBE. Our specification of an IBE scheme $I\mathcal{BE} = (\mathcal{P}, \mathcal{M}, \mathcal{K}, \mathcal{E}, \mathcal{D})$ adds a parameter generation algorithm $\mathcal{P}$ that given 1^k returns parameters π on which the masterkey generation algorithm $\mathcal{M}$ runs to produce the master public key mpk and master secret key msk. The rest is as usual except that algorithms get π as an additional input. Game IBE of Figure 4 is associated to $I\mathcal{BE}$ and an RKA specification $\Phi = (\mathcal{P}, \mathcal{Q})$ that is compatible with $I\mathcal{BE}$. An adversary is allowed only one query to LR. Let $\mathbf{Adv}^{\text{ibe-rka}}_{I\mathcal{BE}, A, \Phi}(k)$ equal $2\Pr[\text{IBE}^A \Rightarrow \text{true}] - 1$ when the game has input 1^k. We say $I\mathcal{BE}$ is Φ-RKA secure if this advantage function is negligible. Here the feature of the definition worth remarking on is that the adversary loses if it ever issues a query to KD that contains the challenge identity *and* derives the same master secret key.

FROM Φ-RKA IBE TO Φ-RKA SIGNATURES. We show that $\mathbf{RKA}[\mathsf{IBE}] \subseteq \mathbf{RKA}[\mathsf{Sig}]$ by proving that the standard Naor transform preserves RKA security. Given a Φ-RKA IBE $\mathcal{IBE} = (\mathcal{P}, \mathcal{M}, \mathcal{K}, \mathcal{E}, \mathcal{D})$ we obtain signature scheme $\mathcal{DS} = (\mathcal{P}, \mathcal{K}, \mathcal{S}, \mathcal{V})$ as follows: the signature parameters are the IBE parameters; the verifying and signing keys are the IBE master public and secret keys; messages being signed are identities; and signatures are decryption keys for these identities. In the most general case verification is performed by encrypting a random message (this is an IBE message) under the identity (the message whose signature is being verified) and seeing whether the ciphertext decrypts correctly under the decryption key (signature being verified). However with all the specific IBE schemes we know (eg. [17, 12, 34]), there is a direct, deterministic verification algorithm based on the algebra. The signature scheme $\mathcal{DS}$ remains compatible with Φ since its parameters are exactly those of $\mathcal{IBE}$. The following says $\mathcal{DS}$ inherits the Φ-RKA security of $\mathcal{IBE}$.

Theorem 6.3 *Let signature scheme $\mathcal{DS} = (\mathcal{P}, \mathcal{K}, \mathcal{S}, \mathcal{V})$ be constructed as above from Φ-RKA IBE scheme $\mathcal{IBE} = (\mathcal{P}, \mathcal{M}, \mathcal{K}, \mathcal{E}, \mathcal{D})$. Then $\mathcal{DS}$ is Φ-RKA secure.*

When the adversary against the signature scheme makes SIGN queries we can simulate them via KD queries. The delicate point is that the correctness of this simulation relies on the fact that the challenge identity has not yet been defined, which means that the two procedures fail in exactly the same cases. Once the signing adversary outputs its forgery with message id and signature dk, we pick two random k-bit messages m_0, m_1 and query $\mathrm{LR}(id, m_0, m_1)$ to get a challenge ciphertext C^*. We then decrypt C^* using $\mathcal{D}$ with decryption key dk and return 1 if we get back m_1 and 0 otherwise. We omit the details.

This motivates finding Φ-RKA secure IBE schemes, which we leave as an interesting open problem. Another interesting open question is whether or not the converse of the above is true, namely, is $\mathbf{RKA}[\mathsf{Sig}] \subseteq \mathbf{RKA}[\mathsf{IBE}]$?

PKE. A public-key encryption scheme $\mathcal{PKE} = (\mathcal{P}, \mathcal{K}, \mathcal{E}, \mathcal{D})$ is specified as usual by parameter generation, key generation, encryption and decryption algorithms. Game PKE of Figure 5 is associated to $\mathcal{PKE}$ and an RKA specification $\Phi = (\mathcal{P}, \mathcal{Q})$ that is compatible with $\mathcal{PKE}$. An adversary is allowed only one query to LR. Let $\mathbf{Adv}^{\mathrm{pke\text{-}cc\text{-}rka}}_{\mathcal{PKE}, A, \Phi}(k)$ equal $2\Pr[\mathrm{PKE}^A \Rightarrow \mathsf{true}] - 1$ when the game has input 1^k. We say $\mathcal{PKE}$ is Φ-RKA secure if this advantage function is negligible for all A. The attack is on the secret key, which is the decryption key, so that we are considering a chosen-ciphertext related-key attack (CC-RKA). The decryption oracle DEC refuses to act only when the ciphertext it is given matches the challenge one *and* the derived key equals the real one.

FROM Φ-RKA IBE TO Φ-RKA PKE-CCA. We show the containment $\mathbf{RKA}[\mathsf{IBE}] \subseteq \mathbf{RKA}[\mathsf{PKE\text{-}CCA}]$ by showing that the construction of Boneh, Canneti, Halevi, and Katz [15] preserves RKA security. We start with a Φ-RKA IBE scheme $\mathcal{IBE} = (\mathcal{P}, \mathcal{M}, \mathcal{K}, \mathcal{E}, \mathcal{D})$ and a signature scheme $\overline{\mathcal{DS}} = (\overline{\mathcal{P}}, \overline{\mathcal{K}}, \overline{\mathcal{S}}, \overline{\mathcal{V}})$, and construct a PKE scheme $\mathcal{PKE}$ as follows:

1. Parameters: Parameters for $\mathcal{PKE}$ are the concatenation $\pi \| \overline{\pi}$ of parameters for $\mathcal{IBE}$ and $\overline{\mathcal{DS}}$.
2. Keys: Pick $(mpk, msk) \leftarrow \mathcal{M}(\pi)$, and use mpk as the public key and msk as the secret key.
3. Encryption: To encrypt a message m under mpk, generate a signing-verification key pair $(\overline{vk}, \overline{sk}) \leftarrow \overline{\mathcal{K}}(\overline{\pi})$. Then encrypt m for the identity $id = \overline{vk}$ by computing $c \leftarrow_\$ \mathcal{E}(\pi, \overline{vk}, m)$ and sign c under $\overline{\mathcal{S}}$ using $\overline{sk}$ to get a signature $\overline{\sigma}$. The ciphertext is $(c, \overline{vk}, \overline{\sigma})$.
4. Decryption: To decrypt a ciphertext $(c, \overline{vk}, \overline{\sigma})$, first verify that $\overline{\sigma}$ is a valid signature on c under $\overline{vk}$, and output $\perp$ if verification fails. Then compute the user secret key for the identity $id = \overline{vk}$, and decrypt c using that key.

By construction, we have that Φ is compatible with $\mathcal{PKE}$ whenever it is compatible with $\mathcal{IBE}$. In order to prove security, we require that $\overline{\mathcal{DS}}$ satisfy (regular-secure) strongly unforgeability, which we recall in Appendix F.

Theorem 6.4 *Let PKE scheme $\mathcal{PKE} = (\mathcal{P}', \mathcal{K}', \mathcal{E}', \mathcal{D}')$ be constructed as above from Φ-RKA IBE scheme $\mathcal{IBE} = (\mathcal{P}, \mathcal{M}, \mathcal{K}, \mathcal{E}, \mathcal{D})$ and normal-secure strongly-unforgeable signature scheme $\overline{\mathcal{DS}} = (\overline{\mathcal{P}}, \overline{\mathcal{K}}, \overline{\mathcal{S}}, \overline{\mathcal{V}})$. Then $\mathcal{PKE}$ is Φ-RKA secure.*

A proof is given in Appendix F. It is an adaptation of the original proof for the non-RKA version of this theorem. The primary difference is in showing that the RKA games for IBE and PKE-CCA will cooperate, because they each have rules for disallowing certain queries (with IBE one cannot extract a key for id^*, and in PKE-CCA one cannot decrypt C^*, but both rules only hold under the original secret key). Handily, the Boneh et al. construction has a structure that allows these rules to fit together.

SE. Symmetric encryption is interesting because both sender and receiver have the secret key and the RKA can now be mounted on the encryption, not merely on the decryption. A symmetric encryption scheme $\mathcal{SE} = (\mathcal{P}, \mathcal{K}, \mathcal{E}, \mathcal{D})$ is specified as usual by parameter generation, key generation, encryption and decryption algorithms. Games SECPA, SECCA of Figure 5 are associated to $\mathcal{SE}$ and a RKA specification $\Phi = (\pi, \mathcal{Q})$ that is compatible with $\mathcal{SE}$. An adversary is now allowed multiple queries to LR. Let $\mathbf{Adv}^{\text{se-cp-rka}}_{\mathcal{SE},A,\Phi}(k)$ equal $2 \Pr[\text{SECPA}^A \Rightarrow \mathsf{true}] - 1$ when the game has input 1^k. We say $\mathcal{SE}$ is Φ-RKA-CPA secure if this advantage function is negligible. Let $\mathbf{Adv}^{\text{se-cc-rka}}_{\mathcal{SE},A,\Phi}(k)$ equal $2 \Pr[\text{SECCA}^A \Rightarrow \mathsf{true}] - 1$ when the game has input 1^k. We say $\mathcal{SE}$ is Φ-RKA-CCA secure if this advantage function is negligible.

REMAINING RELATIONS. Containments of the form $\mathbf{RKA}^*[\mathsf{PRF}] \subseteq \mathbf{RKA}^*[\mathsf{P}]$ can be established for all the other primitives P using the same idea as illustrated above for $\mathsf{P} = \mathsf{Sig}$. All our non-containments emanate from two basic techniques. One is based on constants, as we illustrated with $\mathbf{RKA}[\mathsf{Sig}] \not\subseteq \mathbf{RKA}[\mathsf{PRF}]$. Another example use of this technique is to show $\mathbf{RKA}[\mathsf{PKE\text{-}CCA}] \not\subseteq \mathbf{RKA}[\mathsf{SE\text{-}CCA}]$. This is true because constant functions are fine for the first but, due to the RKA on the encryption oracle, not for the second. However sometimes we want to show $\mathbf{RKA}[\mathsf{P}_1] \not\subseteq \mathbf{RKA}[\mathsf{P}_2]$ when constant functions are either fine for both or not fine for either. We do this based on variants of the technique used above to establish $\mathbf{RKA}[\mathsf{wPRF}] \not\subseteq \mathbf{RKA}[\mathsf{PRF}]$ where the idea can be viewed as using RKDFs that flip-flop between two keys depending on a bit of the original key in such a way that there is a Φ-RKA on P_2 yet the two keys are functionally equivalent under a Φ-RKA on P_1. This can be used for example to show that $\mathbf{RKA}[\mathsf{SE\text{-}CPA}] \not\subseteq \mathbf{RKA}[\mathsf{SE\text{-}CCA}]$. In this case the attack exploits the fact that the decryption oracle rejects if a query (ϕ, C) results in (K', C) being in S, meaning that C was a challenge ciphertext created under K', and this can be made to happen depending on a certain bit of K so that the rejection leaks information about K that eventually allows the attacker to recover K. We omit details on these and the remaining relations shown in Figure 1.

Acknowledgments

References

[1] M. Albrecht, P. Farshim, K. Paterson, and G. Watson. On cipher-dependent related-key attacks in the ideal-cipher model. Cryptology ePrint Archive, Report 2011/213, 2011. `http://eprint.iacr.org/`. (Cited on page 5.)

[2] B. Applebaum, D. Harnik, and Y. Ishai. Semantic security under related-key attacks and applications. In A. C.-C. Yao, editor, *ICS 2011*. Tsinghua University Press, 2011. (Cited on page 4, 5, 6.)

[3] M. Bellare and D. Cash. Pseudorandom functions and permutations provably secure against related-key attacks. In T. Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 666–684. Springer, Aug. 2010. (Cited on page 4, 5, 6, 7, 9, 10.)

[4] M. Bellare and T. Kohno. A theoretical treatment of related-key attacks: RKA-PRPs, RKA-PRFs, and applications. In E. Biham, editor, *EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 491–506. Springer, May 2003. (Cited on page 3, 4, 5, 6, 8, 9, 14.)

[5] M. Bellare and T. Kohno. Hash function balance and its impact on birthday attacks. In C. Cachin and J. Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 401–418. Springer, May 2004. (Cited on page 12.)

[6] M. Bellare and S. K. Miner. A forward-secure digital signature scheme. In M. J. Wiener, editor, *CRYPTO'99*, volume 1666 of *LNCS*, pages 431–448. Springer, Aug. 1999. (Cited on page 6.)

[7] M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In S. Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, May / June 2006. (Cited on page 8, 20, 23.)

[8] E. Biham. New types of cryptanalytic attacks using related keys (extended abstract). In T. Helleseth, editor, *EUROCRYPT'93*, volume 765 of *LNCS*, pages 398–409. Springer, May 1993. (Cited on page 3.)

[9] E. Biham and A. Shamir. Differential fault analysis of secret key cryptosystems. In B. S. Kaliski Jr., editor, *CRYPTO'97*, volume 1294 of *LNCS*, pages 513–525. Springer, Aug. 1997. (Cited on page 3, 4.)

[10] N. Bitansky and R. Canetti. On strong simulation and composable point obfuscation. In T. Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 520–537. Springer, Aug. 2010. (Cited on page 6.)

[11] M. Blum and S. Micali. How to generate cryptographically strong sequences of pseudorandom bits. *SIAM Journal on Computing*, 13(4):850–864, 1984. (Cited on page 11.)

[12] D. Boneh and X. Boyen. Efficient selective-ID secure identity based encryption without random oracles. In C. Cachin and J. Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 223–238. Springer, May 2004. (Cited on page 16.)

[13] D. Boneh and X. Boyen. Short signatures without random oracles. In C. Cachin and J. Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 56–73. Springer, May 2004. (Cited on page 6.)

[14] D. Boneh and X. Boyen. Short signatures without random oracles and the SDH assumption in bilinear groups. *Journal of Cryptology*, 21(2):149–177, Apr. 2008. (Cited on page 6.)

[15] D. Boneh, R. Canetti, S. Halevi, and J. Katz. Chosen-ciphertext security from identity-based encryption. *SIAM Journal on Computing*, 36(5):915–942, 2006. (Cited on page 5, 16.)

[16] D. Boneh, R. A. DeMillo, and R. J. Lipton. On the importance of checking cryptographic protocols for faults (extended abstract). In W. Fumy, editor, *EUROCRYPT'97*, volume 1233 of *LNCS*, pages 37–51. Springer, May 1997. (Cited on page 3, 4.)

[17] D. Boneh, B. Lynn, and H. Shacham. Short signatures from the Weil pairing. In C. Boyd, editor, *ASIACRYPT 2001*, volume 2248 of *LNCS*, pages 514–532. Springer, Dec. 2001. (Cited on page 16.)

[18] Y. Dodis and A. Yampolskiy. A verifiable random function with short proofs and keys. In S. Vaudenay, editor, *PKC 2005*, volume 3386 of *LNCS*, pages 416–431. Springer, Jan. 2005. (Cited on page 6.)

[19] S. Dziembowski, K. Pietrzak, and D. Wichs. Non-malleable codes. In A. C.-C. Yao, editor, *ICS 2010*. Tsinghua University Press, 2010. (Cited on page 6.)

[20] R. Gennaro, A. Lysyanskaya, T. Malkin, S. Micali, and T. Rabin. Algorithmic tamper-proof (ATP) security: Theoretical foundations for security against hardware tampering. In M. Naor, editor, *TCC 2004*, volume 2951 of *LNCS*, pages 258–277. Springer, Feb. 2004. (Cited on page 3, 4, 6, 12.)

[21] D. Goldenberg and M. Liskov. On related-secret pseudorandomness. In D. Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 255–272. Springer, Feb. 2010. (Cited on page 4, 5, 6, 9.)

[22] O. Goldreich. *Foundations of Cryptography: Basic Applications*, volume 2. Cambridge University Press, Cambridge, UK, 2004. (Cited on page 14, 27.)

[23] V. Goyal, A. O'Neill, and V. Rao. Correlated-input secure hash functions. In Y. Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 182–200. Springer, Mar. 2011. (Cited on page 4, 5, 6, 7, 18.)

[24] V. Goyal, A. O'Neill, and V. Rao. Correlated-input secure hash functions. Cryptology ePrint Archive, Report 2011/233, 2011. Full version of [23], http://eprint.iacr.org/. (Cited on page 5, 6, 7.)

[25] S. Halevi and H. Krawczyk. Security under key-dependent inputs. In P. Ning, S. D. C. di Vimercati, and P. F. Syverson, editors, *ACM CCS 07*, pages 466–475. ACM Press, Oct. 2007. (Cited on page 5.)

[26] Y. Ishai, M. Prabhakaran, A. Sahai, and D. Wagner. Private circuits II: Keeping secrets in tamperable circuits. In S. Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 308–327. Springer, May / June 2006. (Cited on page 6.)

[27] Y. T. Kalai, B. Kanukurthi, and A. Sahai. Cryptography with tamperable and leaky memory. In P. Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 373–390. Springer, Aug. 2011. (Cited on page 6.)

[28] L. R. Knudsen. Cryptanalysis of LOKI91. In J. Seberry and Y. Zheng, editors, *AUSCRYPT'92*, volume 718 of *LNCS*, pages 196–208. Springer, Dec. 1992. (Cited on page 3.)

[29] L. Lamport. Constructing digital signatures from a one-way function. Technical Report SRI-CSL-98, SRI International Computer Science Laboratory, Oct. 1979. (Cited on page 27.)

[30] S. Lucks. Ciphers secure against related-key attacks. In B. K. Roy and W. Meier, editors, *FSE 2004*, volume 3017 of *LNCS*, pages 359–370. Springer, Feb. 2004. (Cited on page 4, 5, 6, 9.)

[31] R. C. Merkle. A digital signature based on a conventional encryption function. In C. Pomerance, editor, *CRYPTO'87*, volume 293 of *LNCS*, pages 369–378. Springer, Aug. 1988. (Cited on page 27.)

[32] M. Naor and O. Reingold. Synthesizers and their application to the parallel construction of pseudo-random functions. *J. Comput. Syst. Sci.*, 58(2):336–375, 1999. (Cited on page 4, 5.)

[33] J. Rompel. One-way functions are necessary and sufficient for secure signatures. In *22nd ACM STOC*, pages 387–394. ACM Press, May 1990. (Cited on page 27.)

[34] B. R. Waters. Efficient identity-based encryption without random oracles. In R. Cramer, editor, *EURO-CRYPT 2005*, volume 3494 of *LNCS*, pages 114–127. Springer, May 2005. (Cited on page 16.)

A Proof of Proposition 4.1

Given A attacking the Φ-IDFP security of $\mathcal{FF}$ we construct B such that for all k we have

$$\mathbf{Adv}^{\mathrm{idfp}}_{\mathcal{FF},A,\Phi}(k) \;\leq\; \mathbf{Adv}^{\mathrm{prf\text{-}rka}}_{\mathcal{FF},B,\Phi}(k) + \frac{q(k)}{|\mathsf{Rng}(\pi)|} \;.$$

On input π, adversary B lets $\mathbf{w} \leftarrow_{\$} \mathsf{IKfp}(\pi)$ and $\mathbf{y}[1] \leftarrow \mathrm{FN}(\mathsf{id}, \mathbf{w}[1])$. It initializes set T to empty and then runs A on inputs $\pi, \mathbf{w}$. When A makes oracle query ϕ, adversary B replies via the following procedure:

> proc $\mathrm{FNSIM}(\phi)$
> _____
> If $(\phi = \mathsf{id})$ then return $\mathbf{y}$
> $\mathbf{z}[1] \leftarrow \mathrm{FN}(\phi, \mathbf{w}[1])$; $T \leftarrow T \cup \{\mathbf{z}\}$
> Return $\mathbf{z}$

When A halts, adversary B returns 1 if $\mathbf{y} \in T$ and 0 otherwise. For the analysis, assume (wlog) that A never repeats an oracle query. Then

$$\Pr[\mathrm{PRFReal}^B \Rightarrow 1] \;=\; \mathbf{Adv}^{\mathrm{idfp}}_{\mathcal{FF},A,\Phi}(k)$$

$$\Pr[\mathrm{PRFRand}^B \Rightarrow 1] \;\leq\; \frac{q(k)}{|\mathsf{Rng}(\pi)|}$$

where $q(\cdot)$ is the number of oracle queries made by A. The last equation is justified by the claw-freeness of Φ which ensures that $\Phi_{\pi,\phi}(K) = K$ only if $\phi = \mathsf{id}$. Thus in the random case all queries of A yield keys different from K.

proc INITIALIZE $/\!/$ G_0	proc INITIALIZE $/\!/$ $\boxed{\mathrm{G}_1}$, G_2
$\overline{\pi} \leftarrow_{\$} \overline{\mathcal{P}}(1^k)$	$\overline{\pi} \leftarrow_{\$} \overline{\mathcal{P}}(1^k)$
$\overline{K} \leftarrow_{\$} \overline{\mathcal{K}}(\overline{\pi})$	$\overline{K} \leftarrow_{\$} \overline{\mathcal{K}}(\overline{\pi})$
$S_0 \leftarrow \overline{\mathcal{G}}(\pi, \overline{K})$	$S_0, S_1 \leftarrow_{\$} \{0,1\}^{r(k)}$
Return $\overline{\pi}$	Return $\overline{\pi}$

proc GEN(R) $/\!/$ G_0	proc GEN(R) $/\!/$ $\mathrm{G}_1, \mathrm{G}_2$
Return S_0	$S \leftarrow S_0$
	If $\overline{\mathcal{G}}(\overline{K}) = R$ then bad $\leftarrow$ true ; $\boxed{S \leftarrow S_1}$
proc FINALIZE(b') $/\!/$ G_0	Return S
Return b'	proc FINALIZE(b') $/\!/$ $\mathrm{G}_1, \mathrm{G}_2$
	Return b'

Figure 6: Games for proof of Proposition 5.1. Game G_1 includes the boxed code and game G_2 does not.

B Proof of Proposition 5.1

Let $\ell(k)$ be the length of the seed returned by $\overline{\mathcal{K}}(\overline{\pi})$ when the security parameter underlying $\overline{\pi}$ is k. On input $\overline{\pi}$ let the key-generation algorithm $\mathcal{K}$ of $\mathcal{PRG} = (\overline{\mathcal{P}}, \mathcal{K}, \mathcal{G}, r)$ pick a random bit c and return $c \,\|\, \overline{\mathcal{K}}(\overline{\pi})$. Algorithm $\mathcal{G}(\overline{\pi}, K)$ lets $\overline{K}$ be the last $\ell(k)$ bits of K and returns $\overline{\mathcal{G}}(\overline{\pi}, K)$. Since $\mathcal{G}$ ignores the first bit of its input seed, any pair of seeds that differ only in their first bits result in the same output, meaning are non-trivial collisions.

Now we define Φ so that $\mathcal{PRG}$ is not Φ-ICR secure yet is Φ-RKA secure. For any $R \in \{0,1\}^{r(k)}$ we let $\Phi_{\overline{\pi}, R}$ be the function that on input $K \in \{0,1\}^{\ell(k)+1}$ returns K if $\mathcal{G}(\overline{\pi}, K) \neq R$ and otherwise returns the string formed by flipping the first bit of K.

The adversary A that makes queries $X \leftarrow$ GEN(id) and $S \leftarrow$ GEN(ϕ_X) and returns 1 iff $S = T$ has icr-advantage 1, meaning $\mathcal{PRG}$ is not Φ-ICR secure. Now we show that $\mathcal{PRG}$ is Φ-RKA secure under the assumption that $\overline{\mathcal{PRG}}$ was normal (i.e. ID-RKA) secure.

The functions $\Phi_{\overline{\pi}, R}$, on input $c \,\|\, \overline{K}$, where $c \in \{0,1\}$, have only two possible outputs, namely $c \,\|\, \overline{K}$ itself, or this string with its first bit c flipped. But by design $\mathcal{G}(\overline{\pi}, \cdot)$ has the same output on these two inputs, so in the real case, the response to a query ϕ_R is always $\overline{\mathcal{G}}(\overline{\pi}, \overline{K})$, regardless of R. This would seem to make it easy for an adversary B against the normal-security of $\overline{\mathcal{PRG}}$ to simulate A. The difficulty is the random case. Here the answer to the same query is one of two possible random strings that we call S_0, S_1, but B does not know which. We resolve this by considering two cases in the random case. If A does not query $\mathcal{G}(\overline{\pi}, K)$ then the simulation is again easy since the response to all queries is the same. If A succeeds in querying $\mathcal{G}(\overline{\pi}, K)$, a different strategy is used. Let us now provide the actual proof.

Given A attacking the Φ-RKA security of $\mathcal{PRG}$ we construct B_1, B_2 such that for all k we have

$$\mathbf{Adv}^{\mathrm{prg}}_{\mathcal{PRG}, A, \Phi}(k) \leq \mathbf{Adv}^{\mathrm{prg}}_{\mathcal{PRG}, B_1, \mathsf{ID}}(k) + \mathbf{Adv}^{\mathrm{prg}}_{\mathcal{PRG}, B_2, \mathsf{ID}}(k) + \frac{q(k)}{2^{r(k)}} \,.$$

where $q(\cdot)$ is the number of GEN queries made by A. The proof considers games $\mathrm{G}_0, \mathrm{G}_1, \mathrm{G}_2$ of Figure 6. We begin by claiming that

$$\mathbf{Adv}^{\mathrm{prg}}_{\mathcal{PRG}, A, \Phi}(k) = \Pr[\mathrm{G}_0^A \Rightarrow 1] - \Pr[\mathrm{G}_1^A \Rightarrow 1] \,.$$

This is justified by the discussion above. Now games $\mathbf{G}_1, \mathbf{G}_2$ are identical-until-bad so by the Fundamental Lemma of Game Playing [7] we have

$$\Pr[\mathrm{G}_0^A \Rightarrow 1] - \Pr[\mathrm{G}_1^A \Rightarrow 1] = \Pr[\mathrm{G}_0^A \Rightarrow 1] - \Pr[\mathrm{G}_2^A \Rightarrow 1] + \Pr[\mathrm{G}_2^A \Rightarrow 1] - \Pr[\mathrm{G}_1^A \Rightarrow 1]$$

$$\leq \Pr[\mathrm{G}_0^A \Rightarrow 1] - \Pr[\mathrm{G}_2^A \Rightarrow 1] + \Pr[\mathrm{BAD}(\mathrm{G}_2^A)] \,.$$

We construct B_1, B_2 such that for all k we have

$$\Pr[G_0^A \Rightarrow 1] - \Pr[G_2^A \Rightarrow 1] \leq \mathbf{Adv}_{\mathcal{PRG}, B_1, \mathsf{ID}}^{\mathrm{prg}}(k)$$

$$\Pr[\mathrm{BAD}(G_2^A)] - \frac{q(k)}{2^{r(k)}} \leq \mathbf{Adv}_{\mathcal{PRG}, B_2, \mathsf{ID}}^{\mathrm{prg}}(k) .$$

On input $\overline{\pi}$, adversary B_1 lets $S_0 \leftarrow\!\!\$\ \mathrm{GEN}(\mathsf{id})$ and runs A on input $\overline{\pi}$. When A makes a query R, adversary B_1 responds with S_0, regardless of the value of R. When A halts, B_1 halts with the same output as A.

On input $\overline{\pi}$, adversary B_2 lets $X \leftarrow\!\!\$\ \mathrm{GEN}(\mathsf{id})$ and $S_0 \leftarrow\!\!\$\ \{0,1\}^{r(k)}$. It initializes a set $T \leftarrow \emptyset$ and runs A on input $\overline{\pi}$. When A makes a query R, adversary B_1 responds with S_0 regardless of the value of R, and puts R in T. When A halts, B_2 returns 1 if $X \in T$ and 0 otherwise.

C   Proof of Theorem 6.1

Let ID be the RKA specification consisting of just the identity function, so that an ID-RKA secure signature scheme is just a normal-secure signature scheme, the condition we assume is met by $\overline{\mathcal{DS}}$. Given an adversary A mounting a $(\Phi \cup \Phi_c)$-RKA on $\mathcal{DS}$ we construct adversaries P_0, P_1, S_0, S_1 such that for every $k \in \mathbb{N}$

$$\mathbf{Adv}_{\mathcal{DS}, A, \Phi \cup \Phi_c}^{\mathrm{sig\text{-}rka}}(k) \leq \mathbf{Adv}_{\mathcal{PRG}, P_0, \Phi}^{\mathrm{prg}}(k) + \mathbf{Adv}_{\overline{\mathcal{DS}}, S_0, \mathsf{ID}}^{\mathrm{sig\text{-}rka}}(k) + \mathbf{Adv}_{\mathcal{PRG}, P_1, \Phi}^{\mathrm{prg}}(k) + \mathbf{Adv}_{\overline{\mathcal{DS}}, S_1, \mathsf{ID}}^{\mathrm{sig\text{-}rka}}(k) . \quad (1)$$

This proves the theorem.

In the real game, where the secret signing key is a seed K for $\mathcal{PRG}$, the response to a query $\mathrm{SIGN}(\phi, m)$ is computed by letting $K' \leftarrow \Phi_{\pi, \phi}(K)$ and $T[K'] \leftarrow \mathcal{G}(\pi, K')$ and then using the latter as coins for the key generation algorithm $\overline{\mathcal{K}}$ to get the secret key $\overline{sk}'$ under which the returned signature is generated. A natural approach to proving security is to consider a game R which replaces $T[K']$ with a random $r(k)$ bit string associated to K'. We expect to be able to prove that the assumed Φ-RKA security of $\mathcal{PRG}$ implies that A's success probability in R is about the same as in the real game and then prove that A will not succeed in R based on the assumed security of $\overline{\mathcal{DS}}$. However, we run into a subtle difficulty in the first step. In order to be able to test whether A succeeded, the adversary P attacking the Φ-RKA security of $\mathcal{PRG}$ needs to know the set M, and to do this it must be able to test when a derived seed K' resulting from a signing query ϕ, m equals the real seed K so that it can put the corresponding message m in M. However, it does not know either seed so it is unclear how it can perform this test. To get around this problem we resort to assuming claw-freeness of Φ. This ensures that $K' = K$ only when $\phi = \mathsf{id}$ since no function different from id is allowed to map K to the same thing that id maps it to, namely K. Now the test is independent of K', K.

Let ϕ_a denote the (description of the) constant function that returns a. Simulation of signatures under constant functions will be done directly. For this purpose, both the games and the constructed adversaries will test whether or not ϕ in a signing query is ϕ_a for some a and branch accordingly. A delicate point is that for this to be correct it must be unambiguous, meaning no constant function is in Φ. The latter turns out to be a consequence of the claw-freeness and the fact that the identity function is in Φ. However, constant functions create another difficulty, namely that the adversary might query ϕ_a with $a = K$ being the secret signing key. The difficulty is not for a simulation to return the signature, but to know to put the message m from the query into M, just as above. This turns out to be more of a bother than we would have liked. One's first thought would be that a query of K hands us the seed and should lead at once to breaking the Φ-RKA security of $\mathcal{PRG}$. Letting $R = \mathrm{GEN}(\mathsf{id})$, an adversary P, on seeing query ϕ_a, could test whether or not $\mathcal{G}(\pi, a) = R$, declaring itself to be in the real game if so and the random game otherwise. But for this to be correct we need to show that the probability

21

proc INITIALIZE $/\!/$ G_0, G_1, G_2
$\pi \leftarrow_\$ \mathcal{P}(1^k) \,;\, \overline{\pi} \leftarrow_\$ \overline{\mathcal{P}}(1^k) \,;\, M \leftarrow \emptyset$
$K \leftarrow_\$ \mathcal{K}(\pi) \,;\, T[K] \leftarrow \mathcal{G}(\pi, K)$
$(\overline{vk}, \overline{sk}) \leftarrow \overline{\mathcal{K}}(\overline{\pi}; T[K])$
Return $(\pi \,\|\, \overline{\pi}, \overline{vk})$

proc SIGN(ϕ, m) $/\!/$ $\boxed{G_0}$, G_1
If $(\phi = \phi_a)$ for some a then
$\quad (\overline{vk}', \overline{sk}') \leftarrow \overline{\mathcal{K}}(\overline{\pi}; \mathcal{G}(\pi, a))$
$\quad \sigma \leftarrow_\$ \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}', m)$
$\quad$If $(a = K)$ then
$\quad\quad$bad $\leftarrow$ true $;$ $\boxed{M \leftarrow M \cup \{m\}}$
Else
$\quad K' \leftarrow \Phi_{\pi, \phi}(K)$
$\quad$If $K' = \bot$ then return $\bot$
$\quad$If $T[K'] = \bot$ then
$\quad\quad T[K'] \leftarrow \mathcal{G}(\pi, K')$
$\quad\quad (\overline{vk}', \overline{sk}') \leftarrow \overline{\mathcal{K}}(\overline{\pi}; T[K'])$
$\quad \sigma \leftarrow_\$ \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}', m)$
$\quad$If $(\phi = \mathsf{id})$ then
$\quad\quad M \leftarrow M \cup \{m\}$
Return σ

proc FINALIZE(m, σ) $/\!/$ G_0, G_1, G_2, G_3
Return $((\overline{\mathcal{V}}(\overline{\pi}, \overline{vk}, m, \sigma) = 1) \wedge (m \notin M))$

proc SIGN(ϕ, m) $/\!/$ G_2
If $(\phi = \phi_a)$ for some a then
$\quad (\overline{vk}', \overline{sk}') \leftarrow \overline{\mathcal{K}}(\overline{\pi}; \mathcal{G}(\pi, a))$
$\quad \sigma \leftarrow_\$ \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}', m)$
$\quad$If $(\overline{vk}' = \overline{vk})$ then
$\quad\quad$bad $\leftarrow$ true
Else
$\quad K' \leftarrow \Phi_{\pi, \phi}(K)$
$\quad$If $K' = \bot$ then return $\bot$
$\quad$If $T[K'] = \bot$ then
$\quad\quad T[K'] \leftarrow \mathcal{G}(\pi, K')$
$\quad\quad (\overline{vk}', \overline{sk}') \leftarrow \overline{\mathcal{K}}(\overline{\pi}; T[K'])$
$\quad \sigma \leftarrow_\$ \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}', m)$
$\quad$If $(\phi = \mathsf{id})$ then
$\quad\quad M \leftarrow M \cup \{m\}$
Return σ

proc INITIALIZE $/\!/$ G_3
$\pi \leftarrow_\$ \mathcal{P}(1^k) \,;\, \overline{\pi} \leftarrow_\$ \overline{\mathcal{P}}(1^k) \,;\, M \leftarrow \emptyset$
$K \leftarrow_\$ \mathcal{K}(\pi) \,;\, T[K] \leftarrow_\$ \{0,1\}^{r(k)}$
$(\overline{vk}, \overline{sk}) \leftarrow \overline{\mathcal{K}}(\overline{\pi}; T[K])$
Return $(\pi \,\|\, \overline{\pi}, \overline{vk})$

proc SIGN(ϕ, m) $/\!/$ G_3
If $(\phi = \phi_a)$ for some a then
$\quad (\overline{vk}', \overline{sk}') \leftarrow \overline{\mathcal{K}}(\overline{\pi}; \mathcal{G}(\pi, a))$
$\quad \sigma \leftarrow_\$ \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}', m)$
$\quad$If $(\overline{vk}' = \overline{vk})$ then
$\quad\quad$bad $\leftarrow$ true
Else
$\quad K' \leftarrow \Phi_{\pi, \phi}(K)$
$\quad$If $K' = \bot$ then return $\bot$
$\quad$If $T[K'] = \bot$ then
$\quad\quad T[K'] \leftarrow_\$ \{0,1\}^{r(k)}$
$\quad\quad (\overline{vk}', \overline{sk}') \leftarrow \overline{\mathcal{K}}(\overline{\pi}; T[K'])$
$\quad \sigma \leftarrow_\$ \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}', m)$
$\quad$If $(\phi = \mathsf{id})$ then
$\quad\quad M \leftarrow M \cup \{m\}$
Return σ

Figure 7: Games for proof of Theorem 6.1. Game G_0 includes the boxed code and game G_1 does not.

of the test succeeding when R is random is small. We might hope to argue this unconditionally. But the only way one would appear to be able to do so is to show that for most R there does not *exist* a seed a such that $\mathcal{G}(\pi, a) = R$. But this is only true if the generator is expanding, meaning the number of possible seeds is much less than $2^{r(k)}$. But we are using PRGs to get RKA security and there is no reason that they should be, or should be required to be, expanding, so this argument does not work. What we do instead is prove the claim (that the success probability of the test is small when R is random) computationally, based on the assumed security of the base signature scheme. This is why we end up with two adversaries of each kind. We note that given any Φ-RKA PRG one can easily make one that is expanding by applying a normal PRG to the output of the given one, but it seemed to us a poor choice to change and add burden to the scheme for the sake of the proof when we could in fact prove it without this.

Proceeding now to the actual proof, consider games G_0, G_1, G_2, G_3 of Figure 7. We claim that

$$\begin{aligned}
\mathbf{Adv}_{\mathcal{DS}, A, \Phi \cup \Phi_c}^{\text{sig-rka}}(k) &= \Pr[G_0^A \Rightarrow \mathsf{true}] \qquad\qquad (2)\\
&= \Pr[G_1^A \Rightarrow \mathsf{true}] + \Pr[G_0^A \Rightarrow \mathsf{true}] - \Pr[G_1^A \Rightarrow \mathsf{true}]\\
&\leq \Pr[G_1^A \Rightarrow \mathsf{true}] + \Pr[\textsc{Bad}(G_1^A)] \qquad\qquad (3)\\
&\leq \Pr[G_1^A \Rightarrow \mathsf{true}] + \Pr[\textsc{Bad}(G_2^A)] . \qquad\qquad (4)
\end{aligned}$$

In game G_0 we focus on what is difficult to simulate, namely testing when signing queries result in use of the real key and thus of addition of the message m to the set M. When this results from queries of constant functions we set a flag but, since the boxed code is included, take the appropriate action anyway. When the query involves a function $\phi \in \Phi$, rather than test whether $\phi(K) = K$, we test whether ϕ is the identity, which is equivalent because Φ is claw-free. This justifies Equation (2). Games G_0, G_1 are identical-until-bad so Equation (3) follows from the Fundamental Lemma of Game

22

Playing [7]. When bad is set in G_1 it is also set in game G_2 (although not necessarily vice-versa). Thus $\Pr[\text{BAD}(G_1^A)] \leq \Pr[\text{BAD}(G_2^A)]$ which justifies Equation (4). In game G_2 the test no longer depends on K and hence can be simulated. This justifies the above. Now we will design P_0, P_1, S_0, S_1 so that for every $k \in \mathbb{N}$

$$\Pr[G_1^A \Rightarrow \text{true}] - \Pr[G_3^A \Rightarrow \text{true}] \quad \leq \quad \mathbf{Adv}_{\mathcal{PRG}, P_0, \Phi}^{\text{prg}}(k) \tag{5}$$

$$\Pr[G_3^A \Rightarrow \text{true}] \quad \leq \quad \mathbf{Adv}_{\overline{\mathcal{DS}}, S_0, \text{ID}}^{\text{sig-rka}}(k) \tag{6}$$

$$\Pr[\text{BAD}(G_2^A)] - \Pr[\text{BAD}(G_3^A)] \quad \leq \quad \mathbf{Adv}_{\mathcal{PRG}, P_1, \Phi}^{\text{prg}}(k) \tag{7}$$

$$\Pr[\text{BAD}(G_3^A)] \quad \leq \quad \mathbf{Adv}_{\overline{\mathcal{DS}}, S_1, \text{ID}}^{\text{sig-rka}}(k) \; . \tag{8}$$

Now, Equations (5) and (6) bound the first term of Equation (4) while Equations (7) and (8) bound the second term. Let us now describe the claimed adversaries.

For $b \in \{0, 1\}$, adversary P_b gets input π. It begins with the initializations

$$\overline{\pi} \leftarrow_{\$} \overline{\mathcal{P}}(1^k) \; ; \; M \leftarrow \emptyset \; ; \; R \leftarrow_{\$} \text{GEN}(\text{id}) \; ; \; (\overline{vk}, \overline{sk}) \leftarrow \overline{\mathcal{K}}(\overline{\pi}; R)$$

It now runs A on inputs $\pi \| \overline{\pi}, \overline{vk}$, responding to its oracle queries via the following procedure:

proc $\text{SIGNSIM}(\phi, m)$
If $(\phi = \phi_a)$ for some a then
$\quad (\overline{vk}', \overline{sk}') \leftarrow \overline{\mathcal{K}}(\overline{\pi}; \mathcal{G}(\pi, a))$
$\quad \sigma \leftarrow_{\$} \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}', m)$
$\quad$ If $(vk' = \overline{vk})$ then $\text{bad} \leftarrow \text{true}$
Else
$\quad R' \leftarrow_{\$} \text{GEN}(\phi)$
$\quad$ If $R' = \bot$ then return $\bot$
$\quad (\overline{vk}', \overline{sk}') \leftarrow \overline{\mathcal{K}}(\overline{\pi}; R')$
$\quad \sigma \leftarrow_{\$} \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}', m)$
$\quad$ If $(\phi = \text{id})$ then $M \leftarrow M \cup \{m\}$
Return σ

When A makes query $\text{FINALIZE}(m, \sigma)$, adversary P_0 returns 1 if $(\overline{\mathcal{V}}(\overline{\pi}, \overline{vk}, m, \sigma) = 1) \wedge (m \notin M)$ and 0 otherwise, while adversary P_1 returns 1 if it set bad and 0 otherwise.

For $b \in \{0, 1\}$, adversary S_b gets input $\overline{\pi}, \overline{vk}$. It begins with the initializations

$$\pi \leftarrow_{\$} \mathcal{P}(1^k) \; ; \; K \leftarrow_{\$} \mathcal{K}(\pi)$$

It now runs A on inputs $\pi \| \overline{\pi}, \overline{vk}$, responding to its oracle queries via the following procedure:

proc $\text{SIGNSIM}(\phi, m)$
If $(\phi = \phi_a)$ for some a then
$\quad (\overline{vk}', \overline{sk}') \leftarrow \overline{\mathcal{K}}(\overline{\pi}; \mathcal{G}(\pi, a))$
$\quad \sigma \leftarrow_{\$} \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}', m)$
$\quad$ If $(\overline{vk}' = \overline{vk})$ then $\text{bad} \leftarrow \text{true} \; ; \; \overline{sk}^* \leftarrow \overline{sk}'$
If $(\phi = \text{id})$ then $M \leftarrow M \cup \{m\} \; ; \; \sigma \leftarrow_{\$} \text{SIGN}(\text{id}, m)$
$\quad$ Else
$\quad\quad K' \leftarrow \Phi_{\pi, \phi}(K)$
$\quad\quad$ If $K' = \bot$ then return $\bot$

23

proc INITIALIZE $/\!/$ G_0, G_1, G_2, G_3
$\pi \leftarrow_\$ \mathcal{P}(1^k)$; $\overline{\pi} \leftarrow_\$ \overline{\mathcal{P}}(1^k)$; $M \leftarrow \emptyset$
$K \leftarrow_\$ \mathcal{K}(\pi)$; $T[K] \leftarrow \mathcal{G}(\pi, K)$
$(\overline{vk}, \overline{sk}) \leftarrow \overline{\mathcal{K}}(\overline{\pi}; T[K])$
Return $(\pi \parallel \overline{\pi}, \overline{vk})$

proc SIGN(ϕ, m) $/\!/$ $\boxed{G_0}$, G_1
If $(\phi = \phi_a)$ for some a then
$\quad (\overline{vk}', \overline{sk}') \leftarrow \overline{\mathcal{K}}(\overline{\pi}; \mathcal{G}(\pi, a))$
$\quad \sigma \leftarrow_\$ \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}', m)$
$\quad$ If $(a = K)$ then $\mathsf{bad} \leftarrow \mathsf{true}$; $\boxed{M \leftarrow M \cup \{m\}}$
Else
$\quad K' \leftarrow \Phi_{\pi, \phi}(K)$
$\quad$ If $K' = \perp$ then return $\perp$
$\quad$ If $T[K'] = \perp$ then $T[K'] \leftarrow \mathcal{G}(\pi, K')$
$\quad (\overline{vk}', \overline{sk}') \leftarrow \overline{\mathcal{K}}(\overline{\pi}; T[K'])$; $\sigma \leftarrow_\$ \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}', m)$
$\quad$ If $((T[K'] = T[K]) \wedge (m \notin M))$ then
$\quad\quad M \leftarrow M \cup \{m\}$
$\quad\quad$ If $(K' \neq K)$ then $\mathsf{bad} \leftarrow \mathsf{true}$; $\boxed{M \leftarrow M \setminus \{m\}}$
Return σ

proc FINALIZE(m, σ) $/\!/$ G_0, G_1, G_2, G_3
Return $((\overline{\mathcal{V}}(\overline{\pi}, \overline{vk}, m, \sigma) = 1) \wedge (m \notin M))$

proc SIGN(ϕ, m) $/\!/$ G_2
If $(\phi = \phi_a)$ for some a then
$\quad (\overline{vk}', \overline{sk}') \leftarrow \overline{\mathcal{K}}(\overline{\pi}; \mathcal{G}(\pi, a))$
$\quad \sigma \leftarrow_\$ \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}', m)$
$\quad$ If $(\overline{vk}' = \overline{vk})$ then $\mathsf{bad} \leftarrow \mathsf{true}$
Else
$\quad K' \leftarrow \Phi_{\pi, \phi}(K)$
$\quad$ If $K' = \perp$ then return $\perp$
$\quad$ If $T[K'] = \perp$ then $T[K'] \leftarrow \mathcal{G}(\pi, K')$
$\quad (\overline{vk}', \overline{sk}') \leftarrow \overline{\mathcal{K}}(\overline{\pi}; T[K'])$; $\sigma \leftarrow_\$ \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}', m)$
$\quad$ If $((T[K'] = T[K]) \wedge (m \notin M))$ then
$\quad\quad M \leftarrow M \cup \{m\}$
Return σ

proc SIGN(ϕ, m) $/\!/$ G_3
If $(\phi = \phi_a)$ for some a then
$\quad (\overline{vk}', \overline{sk}') \leftarrow \overline{\mathcal{K}}(\overline{\pi}; \mathcal{G}(\pi, a))$
$\quad \sigma \leftarrow_\$ \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}', m)$
Else
$\quad K' \leftarrow \Phi_{\pi, \phi}(K)$
$\quad$ If $K' = \perp$ then return $\perp$
$\quad$ If $T[K'] = \perp$ then $T[K'] \leftarrow \mathcal{G}(\pi, K')$
$\quad (\overline{vk}', \overline{sk}') \leftarrow \overline{\mathcal{K}}(\overline{\pi}; T[K'])$; $\sigma \leftarrow_\$ \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}', m)$
$\quad$ If $((T[K'] = T[K]) \wedge (m \notin M))$ then
$\quad\quad M \leftarrow M \cup \{m\}$
$\quad\quad$ If $(K' \neq K)$ then $\mathsf{bad} \leftarrow \mathsf{true}$
Return σ

Figure 8: Games for proof of Theorem 6.2. Game G_0 includes the boxed code and game G_1 does not.

$\quad$ If $T[K'] = \perp$ then $T[K'] \leftarrow_\$ \{0, 1\}^{r(k)}$
$\quad (\overline{vk}', \overline{sk}') \leftarrow \overline{\mathcal{K}}(\overline{\pi}; T[K'])$
$\quad \sigma \leftarrow_\$ \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}', m)$
Return σ

When A makes query FINALIZE(m, σ), adversary S_0 does too. Adversary S_1, however, sees whether it set bad. If so it picks some message $m \notin M$, lets $\sigma \leftarrow_\$ \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}^*, m)$ and makes query FINALIZE(m, σ). Otherwise, it halts without output.

D   Proof of Theorem 6.2

Again let ID be the RKA specification consisting of just the identity function, so that an ID-RKA secure signature scheme is just a normal-secure signature scheme, the condition we assume is met by $\overline{\mathcal{DS}}$. Given an adversary A mounting a $(\Phi \cup \Phi_c)$-RKA on $\mathcal{DS}$ we construct adversaries P_0, P_1, S_0, S_1, C such that for every $k \in \mathbb{N}$

$$\mathbf{Adv}^{\text{sig-rka}}_{\mathcal{DS}, A, \Phi \cup \Phi_c}(k)$$
$$\leq \mathbf{Adv}^{\text{prg}}_{\mathcal{PRG}, P_0, \Phi}(k) + \mathbf{Adv}^{\text{sig-rka}}_{\overline{\mathcal{DS}}, S_0, \text{ID}}(k) + \mathbf{Adv}^{\text{prg}}_{\mathcal{PRG}, P_1, \Phi}(k) + \mathbf{Adv}^{\text{sig-rka}}_{\overline{\mathcal{DS}}, S_1, \text{ID}}(k) + \mathbf{Adv}^{\text{icr}}_{\mathcal{PRG}, C, \Phi}(k) \ . \ (9)$$

This proves the theorem.

As we discussed in Appendix C, an important difficulty is that an adversary P against the PRG cannot test whether $K' = K$. There we solved this by assuming claw-freeness of Φ, which allowed us to instead test whether $\phi = \mathsf{id}$. Now that Φ may not be claw-free, we attempt to exploit the collision-resistance of $\mathcal{PRG}$. We test instead whether $T[K'] = T[K]$, hoping this happens only when $K' = K$. To ensure that the game is faithful to the original one, we put in a correction, setting bad and removing m from M if $K' \neq K$. We have

$$
\begin{aligned}
\mathbf{Adv}^{\text{sig-rka}}_{\mathcal{DS},A,\Phi\cup\Phi_c}(k) &= \Pr[G_0^A \Rightarrow \mathsf{true}] \\
&= \Pr[G_1^A \Rightarrow \mathsf{true}] + \Pr[G_0^A \Rightarrow \mathsf{true}] - \Pr[G_1^A \Rightarrow \mathsf{true}] \\
&\leq \Pr[G_1^A \Rightarrow \mathsf{true}] + \Pr[\text{BAD}(G_1^A)] \,.
\end{aligned}
$$

The boxed code being dropped in G_1, a simulator can ignore the corresponding tests that set bad, since they do not affect the outcome of the game. Thus, as in Appendix C, we can construct P_0, S_0 so that

$$
\Pr[G_1^A \Rightarrow \mathsf{true}] \leq \mathbf{Adv}^{\text{prg}}_{\mathcal{PRG},P_0,\Phi}(k) + \mathbf{Adv}^{\text{sig-rka}}_{\mathcal{DS},S_0,\mathsf{ID}}(k) \,.
$$

It remains to bound $\Pr[\text{BAD}(G_1^A)]$. The probability that the first occurence of bad is set in G_1 is at most the probability that it is set in G_2. The probability that second occurence of bad is set in G_1 equals the probability that it is set in G_3. Thus

$$
\Pr[\text{BAD}(G_1^A)] \leq \Pr[\text{BAD}(G_2^A)] + \Pr[\text{BAD}(G_3^A)] \,.
$$

As in Appendix C, we can construct P_1, S_1 so that

$$
\Pr[\text{BAD}(G_2^A)] \leq \mathbf{Adv}^{\text{prg}}_{\mathcal{PRG},P_1,\Phi}(k) + \mathbf{Adv}^{\text{sig-rka}}_{\mathcal{DS},S_1,\mathsf{ID}}(k) \,.
$$

Finally we construct C so that

$$
\Pr[\text{BAD}(G_3^A)] \leq \mathbf{Adv}^{\text{icr}}_{\mathcal{PRG},C,\Phi}(k) \,.
$$

Adversary C gets input π. It begins with the initializations

$$
\overline{\pi} \leftarrow_\$ \overline{\mathcal{P}}(1^k) \,;\ R \leftarrow_\$ \text{GEN}(\mathsf{id}) \,;\ (\overline{vk}, \overline{sk}) \leftarrow \overline{\mathcal{K}}(\overline{\pi}; R)
$$

It now runs A on inputs $\pi \,\|\, \overline{\pi}, \overline{vk}$, responding to its oracle queries via the following procedure:

proc $\text{SIGNSIM}(\phi, m)$
If $(\phi = \phi_a)$ for some a then
$\quad (\overline{vk}', \overline{sk}') \leftarrow \overline{\mathcal{K}}(\overline{\pi}; \mathcal{G}(\pi, a)) \,;\ \sigma \leftarrow_\$ \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}', m)$
Else
$\quad R' \leftarrow_\$ \text{GEN}(\phi)$
$\quad$ If $R' = \bot$ then return $\bot$
$\quad (\overline{vk}', \overline{sk}') \leftarrow \overline{\mathcal{K}}(\overline{\pi}; R') \,;\ \sigma \leftarrow_\$ \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}', m)$
Return σ

When A makes query $\text{FINALIZE}(m, \sigma)$, adversary P halts without output.

E Strong Φ-RKA Security

We suggest a stronger notion of Φ-RKA security for signature schemes and other primitives. Let us start with the former. In strong Φ-RKA security for signatures, forgery is not only hard relative to the

```
proc INITIALIZE  ∥ SSig
π ←$ P(1^k) ; M ← ∅ ; (vk, sk) ←$ K(π) ; i ← 0
Return (π, vk)

proc SIGN(φ, m)  ∥ SSig
sk' ← Φ_{π,φ}(sk)
If sk' = ⊥ then return ⊥
i ← i + 1 ; vk_i ← T(π, sk')
M ← M ∪ {(vk_i, m)}
σ ←$ S(π, sk', m)
Return (vk_i, σ)

proc FINALIZE(j, m, σ)  ∥ SSig
If (j > i) then return false
Return ((V(π, vk_j, m, σ) = 1) ∧ ((vk_j, m) ∉ M))
```

Figure 9: Game defining strong Φ-RKA security for $(\mathcal{DS}, \mathcal{T})$ where $\Phi = (\mathcal{P}, \mathcal{Q})$.

original public key, but even for ones associated with the derived keys produced in the attack. We are not aware of this having any application-relevance but wish to highlight it because our constructions possess it and it may be of some utility in the future.

In the usual syntax of digital signatures, the public and secret keys are produced, together, by the (randomized) key generation algorithm. In most "real" schemes, however, the secret key is produced first and the public key is a deterministic function of the (parameters and) secret key. We call any signature scheme with this property separable and we call the algorithm that deterministically produces the public key from the parameters and secret key the public-key generator. Note that even if a scheme is not separable to start with, it can be made so by using the coins of the key-generation algorithm as the secret key in an obvious way.

DEFINITIONS. A signature scheme $\mathcal{DS} = (\mathcal{P}, \mathcal{K}, \mathcal{S}, \mathcal{V})$ is separable if there is a deterministic algorithm $\mathcal{T}$, called the public-key generator, such that for all $\pi \in [\mathcal{P}(1^k)]$ and all $k \in \mathbb{N}$ the output of the process

$$(vk, sk) \leftarrow_\$ \mathcal{K}(\pi) ; \ vk \leftarrow \mathcal{T}(\pi, sk) ; \ \text{Return } (vk, sk)$$

is distributed identically to the output of $\mathcal{K}(\pi)$. Let $\mathcal{DS} = (\mathcal{P}, \mathcal{K}, \mathcal{S}, \mathcal{V})$ be a separable signature scheme with public-key generator $\mathcal{T}$. Game SSig of Figure 9 is associated to $(\mathcal{DS}, \mathcal{T})$ and a RKA specification $\Phi = (\mathcal{P}, \mathcal{Q})$ that is compatible with $\mathcal{DS}$. Let $\mathbf{Adv}^{\text{ssig-rka}}_{\mathcal{DS}, \mathcal{T}, A, \Phi}(k)$ be the probability that SSig^A returns true when the game has input 1^k. We say $(\mathcal{DS}, \mathcal{T})$ is strongly Φ-RKA secure if this advantage function is negligible. A novel aspect of this definition is that the SIGN oracle returns not just the signature but the public key associated to the secret key under which the signature was created. This is necessary because the adversary may not be able to compute this public key on its own. To win, it need only forge a signature under one of the public keys derived via its SIGN queries.

RESULT. We first observe that the signature scheme $\mathcal{DS} = (\mathcal{P} \,\|\, \overline{\mathcal{P}}, \mathcal{K}', \mathcal{S}, \mathcal{V})$ constructed from Φ-RKA PRG $\mathcal{PRG} = (\mathcal{P}, \mathcal{K}, \mathcal{G}, r)$ and normal-secure signature scheme $\overline{\mathcal{DS}} = (\overline{\mathcal{P}}, \overline{\mathcal{K}}, \overline{\mathcal{S}}, \overline{\mathcal{V}})$ as described in Section 6 is separable. This is true regardless of whether or not $\overline{\mathcal{DS}}$ is separable. Now we claim that $\mathcal{DS}$ is strongly Φ-RKA secure.

Theorem E.1 *Let separable signature scheme $\mathcal{DS} = (\mathcal{P} \,\|\, \overline{\mathcal{P}}, \mathcal{K}', \mathcal{S}, \mathcal{V})$ be constructed as in Section 6 from Φ-RKA PRG $\mathcal{PRG} = (\mathcal{P}, \mathcal{K}, \mathcal{G}, r)$ and normal-secure signature scheme $\overline{\mathcal{DS}} = (\overline{\mathcal{P}}, \overline{\mathcal{K}}, \overline{\mathcal{S}}, \overline{\mathcal{V}})$ Let $\mathcal{T}$ denote its public-key generation algorithm and assume Φ is claw-free. Then $(\mathcal{DS}, \mathcal{T})$ is strongly Φ-RKA secure.*

Notice that we do *not* claim $(\mathcal{DS}, \mathcal{T})$ is strongly $(\Phi \cup \Phi_c)$-RKA secure. Indeed, strongly Φ_c-RKA signature schemes do not exist. The strong security of any signature scheme breaks, since an adversary

can itself compute signatures under a constant secret key and can then forge under them. The proof is similar to that of Theorem 6.1 and is omitted. An analogue of Theorem 6.2 may also be stated and proved.

OTHER PRIMITIVES. The analogous security definitions for strong Φ-RKA CCA asymmetric encryption also exists. For Φ-RKA CCA asymmetric encryption, the adversary selects not only two messages m_0 and m_1, but also a $\phi \in \Phi$, and then receives the encryption of one of the two messages under the related key derived using ϕ; the adversary is said to succeed if it can guess with non-negligible advantage which message was encrypted.

The generic construction that transforms a secure CCA asymmetric encryption scheme using a Φ-RKA-PRG to generate the secret key before each use gives this strong security definition. We omit the proof here, but it is very similar to the proof of strong security for the signature scheme. Intuitively, each distinct $\phi \in \Phi$ cause the Φ-RKA-PRG to output values indistinguishable from random. When this pseudorandomness is used to generate the secret key, it appears as if it was a new randomly generated instance of the primitive. Since the adversary can only interact with polynomially many of these, we can guess which one they will choose, and embed the standard Φ-RKA security challenge here.

A DEDICATED SCHEME. In addition to our general transformation from any signature scheme to a strongly Φ-RKA secure signature scheme, we present a specific and direct construction of a strongly Φ-RKA secure signature scheme based on a Φ-RKA PRF. We use the Merkle Tree [31] transformation from a secure one-time signature scheme to a secure many time signature scheme. By replacing the PRF in this construction with a Φ-RKA-PRF, we note that the scheme becomes a $(\phi \cup \Phi_c)$-RKA secure signature scheme. It is interesting that creating a Φ-RKA secure scheme here only requires replacing the PRF with a Φ-RKA-PRF. This means the scheme, unmodified, will be RKA secure if the underlying PRF is RKA secure. No modification to the scheme is needed.

To sign n-bit messages with security parameter k, we require a one-time signature scheme that can sign messages twice the length of its public key, and that supports a public-key generator π. We additionally require $\mathcal{PRF} = (\mathcal{P}, \mathcal{K}, \mathcal{F})$ with $\mathsf{Rng}(\pi)$ equal to the private key space of the one-time signature scheme for each value of k. Such a one-time signature scheme exists under the minimal assumption that one-way functions exist [33]. This scheme uses a hash and sign paradigm based on a universal one-way hash function (UOWHF), and the Lamport one-time signature scheme [29], which signs a bit by revealing one of two random OWF preimages.

A Merkle tree [31] is used to create a many-time signature scheme from a depth k binary tree of one-time signature instances. To sign a message, a path through the tree is randomly selected. Each internal signature scheme instance of the tree is used to sign a hash of the public keys of its children; the children's public keys are also included in the signature. The final signature instance, in a leaf of the tree, is used to sign the desired message.

Goldreich noted [22] that instead of storing the secret keys for each one-time signature scheme instance, the secret keys can generated as needed by a PRF. Given secret key K for the overall signature scheme, each instance's secret key is the PRF evaluated at the binary path to the instance in the tree. For example, the root node's secret key is $\mathcal{F}(\pi, K, \epsilon)$, and its two children's secret keys are $\mathcal{F}(\pi, K, 0)$ and $\mathcal{F}(\pi, K, 1)$. This allows the secret key for the entire construction to be comprised only of the key for the PRF. See [22] for a formal description of this scheme.

Replacing the PRF with a Φ-RKA-PRF, this scheme is a $C \cup \Phi$-RKA secure and strongly Φ-RKA secure signature scheme. The intuition for these proofs of security follows the proof for Theorem E.1.

F   Proof of Theorem 6.4

We first state the definition of strong unforgeability. Let $\overline{\mathcal{DS}} = (\overline{\mathcal{P}}, \overline{\mathcal{K}}, \overline{\mathcal{S}}, \overline{\mathcal{V}})$ be a signature scheme and consider the game SUSig in Figure 10. For an adversary A, let $\mathbf{Adv}^{\mathrm{susig}}_{\overline{\mathcal{DS}}, A}(k)$ be the probability that SUSig^A outputs true. We say that $\overline{\mathcal{DS}}$ is strongly unforgeable if this advantage function is negligible

proc INITIALIZE // SUSig
$\overline{\pi} \leftarrow_\$ \overline{\mathcal{P}}(1^k)$; $M \leftarrow \emptyset$; $(\overline{vk}, \overline{sk}) \leftarrow_\$ \overline{\mathcal{K}}(\overline{\pi})$
Return $(\overline{\pi}, \overline{vk})$

proc SIGN(m) // SUSig
$\overline{\sigma} \leftarrow_\$ \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}, m)$
$M \leftarrow M \cup \{(m, \overline{\sigma})\}$
Return $\overline{\sigma}$

proc FINALIZE(m, σ) // SUSig
Return $((\overline{\mathcal{V}}(\overline{\pi}, \overline{vk}, m, \overline{\sigma}) = 1) \wedge ((m, \overline{\sigma}) \notin M))$

Figure 10: Game SUSig defining strong unforgeability for $\overline{\mathcal{DS}}$.

proc INITIALIZE // G_0, G_1, G_2, G_3
$\pi \leftarrow_\$ \mathcal{P}(1^k)$; $\overline{\pi} \leftarrow_\$ \overline{\mathcal{P}}(1^k)$
$(mpk, msk) \leftarrow_\$ \mathcal{M}(\pi)$
$b \leftarrow_\$ \{0, 1\}$; $C^* \leftarrow \perp$; $S \leftarrow \emptyset$
Return $(\pi \| \overline{\pi}, mpk)$

proc LR(m_0, m_1) // G_0, G_1, G_2
If $|m_0| \neq |m_1|$ then return $\perp$
$(\overline{vk}^*, \overline{sk}^*) \leftarrow_\$ \overline{\mathcal{K}}(\overline{\pi})$
$c^* \leftarrow_\$ \mathcal{E}(\pi, mpk, \overline{vk}^*, m_b)$
$\overline{\sigma}^* \leftarrow \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}^*, c^*)$
$C^* \leftarrow (c^*, \overline{vk}^*, \overline{\sigma}^*)$
Return C^*

proc FINALIZE(b') // G_0, G_1, G_2, G_3
Return $(b = b')$

proc DEC$(\phi, (c, \overline{vk}, \overline{\sigma}))$ // $G_0, \boxed{G_1}$
$msk' \leftarrow \Phi_{\pi,\phi}(msk)$
If $msk' = \perp$ then return $\perp$
If $((msk' = msk) \wedge ((c, \overline{vk}, \overline{\sigma}) = C^*))$
  then return $\perp$
If $\overline{\mathcal{V}}(\overline{\pi}, \overline{vk}, c, \overline{\sigma}) = 0$
  then return $\perp$
If $(\overline{vk} = \overline{vk}^*) \wedge ((c, \overline{\sigma}) \neq (c^*, \overline{\sigma}^*))$
  then bad $\leftarrow$ true ; $\boxed{\text{return } \perp}$
$dk \leftarrow \mathcal{K}(\pi, msk', \overline{vk})$
Return $m \leftarrow \mathcal{D}(\pi, \overline{vk}, dk, c)$

proc DEC$(\phi, (c, \overline{vk}, \overline{\sigma}))$ // G_2, G_3
$msk' \leftarrow \Phi_{\pi,\phi}(msk)$
If $msk' = \perp$ then return $\perp$
If $(msk' = msk) \wedge (\overline{vk} = \overline{vk}^*)$
  then return $\perp$
If $\overline{\mathcal{V}}(\overline{\pi}, \overline{vk}, c, \overline{\sigma}) = 0$
  then return $\perp$
If $msk' = msk$ then $S \leftarrow S \cup \{\overline{vk}\}$
$dk \leftarrow \mathcal{K}(\pi, msk', \overline{vk})$
Return $m \leftarrow \mathcal{D}(\pi, \overline{vk}, dk, c)$

proc LR(m_0, m_1) // $G_2, \boxed{G_3}$
If $|m_0| \neq |m_1|$ then return $\perp$
$(\overline{vk}^*, \overline{sk}^*) \leftarrow_\$ \overline{\mathcal{K}}(\overline{\pi})$
If $\overline{vk}^* \in S$ then
  bad $\leftarrow$ true ; $\boxed{\text{return } \perp}$
$c^* \leftarrow_\$ \mathcal{E}(\pi, mpk, \overline{vk}^*, m_b)$
$\overline{\sigma}^* \leftarrow \overline{\mathcal{S}}(\overline{\pi}, \overline{sk}^*, c^*)$
$C^* \leftarrow (c^*, \overline{vk}^*, \overline{\sigma}^*)$
Return C^*

Figure 11: Games for the proof of Theorem 6.4. G_1 includes the boxed code and G_0 does not.

for all adversaries. Note that we have stated this definition without its RKA version because we only require the usual definition. Also note that this notion is orthogonal to the notions in Appendix C (the usage of the word "strong" is coincidental).

The proof uses the sequence of games in Figure 11. Game G_0 implements PKE^A with $\mathcal{PKE}$, and game G_1 is the same, except it includes the boxed code. Thus we have

$$
\begin{aligned}
\mathbf{Adv}^{\text{pke-cca}}_{\mathcal{PKE}, A, \Phi}(k) &= 2 \Pr[G_0^A \Rightarrow \text{true}] - 1 \\
&= 2(\Pr[G_1^A \Rightarrow \text{true}] + \Pr[G_0^A \Rightarrow \text{true}] - \Pr[G_1^A \Rightarrow \text{true}]) - 1 \\
&\leq 2(\Pr[G_1^A \Rightarrow \text{true}] + \Pr[\text{BAD}(G_0^A)]) - 1
\end{aligned}
\tag{10}
$$

The inequality follows from the fundamental lemma of game playing since G_1 and G_0 are identical until bad. We construct an adversary B such that

$$
\Pr[\text{BAD}(G_0^A)] \leq \mathbf{Adv}^{\text{susig}}_{\overline{\mathcal{DS}}, B}(k)
\tag{11}
$$

B has oracle SIGN and takes as input $(\overline{\pi}, \overline{vk}^*)$. It starts by running

28

$$\pi \leftarrow_\$ \mathcal{P}(1^k) \,;\, (mpk, msk) \leftarrow_\$ \mathcal{M}(\pi) \,;\, C^* \leftarrow \bot \,;\, b \leftarrow_\$ \{0, 1\}.$$

It then runs A with input $(\pi \,\|\, \overline{\pi}, mpk)$. Recall that A has oracles LR and DEC for which B must produce answers. When A queries $\text{LR}(m_0, m_1)$, B does

> If $|m_0| \neq |m_1|$ then return $\bot$
> $c^* \leftarrow_\$ \mathcal{E}(\pi, mpk, \overline{vk}^*, m_b) \,;\, \overline{\sigma}^* \leftarrow \text{SIGN}(c^*) \,;\, C^* \leftarrow (c^*, \overline{vk}^*, \sigma^*)$
> Return C^*,

where the call to SIGN is a call to the oracle provided to B. When A queries $\text{LR}(\phi, (c, \overline{vk}, \overline{\sigma}))$, B responds via

> $msk' \leftarrow \Phi_{\pi, \phi}(sk)$; If $msk' = \bot$ then return $\bot$
> If $((msk' = msk) \wedge ((c, \overline{vk}, \overline{\sigma}) = C^*))$ then return $\bot$
> If $\overline{\mathcal{V}}(\overline{\pi}, \overline{vk}, c, \overline{\sigma}) = 0$ then return $\bot$
> If $(\overline{vk} = \overline{vk}^*) \wedge ((c, \overline{\sigma}) \neq (c^*, \overline{\sigma}^*))$ then halt with output $(c, \overline{\sigma})$
> $dk \leftarrow \mathcal{K}(\pi, msk', \overline{vk}) \,;\, M \leftarrow \mathcal{D}(\pi, dk, c)$
> Return M

Adversary B runs until A halts. At this point, if B has not already halted, it halts with no output. We now argue that B satisfies the claimed advantage bound. This amounts to arguing that, when B halts early during a DEC query, then it produces a valid forgery. In the SUSig game, B will win as long as $\overline{\mathcal{V}}(\overline{\pi}, \overline{vk}, c, \overline{\sigma}) = 1$ and $(c, \overline{\sigma}) \notin M$. By the construction of B, both of these will be satisfied. The first will be true because the second "If" test must have failed in DEC processing, and the second considition is satisfied because the set M will consist only of (c^*, σ^*), a point that is distinct from the output of B when the third "If" test passes. This concludes the description of B and establishes (11).

Game G_2 collapses the second and fourth "If" lines in the DEC oracle processing into a single equivalent test. It also maintains a set S of values $\overline{vk}$ queried by the adversary, and in the LR the bad flag is set if $\overline{vk}^* \in S$. None of this affects the output of the oracles, and hence does not affect the output of the game, so we have

$$\Pr[G_2^A \Rightarrow \text{true}] = \Pr[G_1^A \Rightarrow \text{true}]. \tag{12}$$

Now, since G_2 and G_3 are identical until bad, we have

$$\Pr[G_2^A \Rightarrow \text{true}] - \Pr[G_3^A \Rightarrow \text{true}] \leq \Pr[\text{BAD}(G_2)]. \tag{13}$$

We claim there is an adversary C such that

$$\Pr[\text{BAD}(G_2)] \leq \mathbf{Adv}_{\overline{\mathcal{DS}}, C}^{\text{susig}}(k). \tag{14}$$

We only give the intuition for how C works. Intuitively, the bad flag is set only when A guesses $\overline{vk}^*$ before it is generated. This means, in particular, that the public keys for $\overline{\mathcal{DS}}$ must be in a relatively small set with good probability, and hence there is a good chance that $\overline{\mathcal{K}}$ outputs the same $\overline{vk}^*$ twice on independent runs. Thus, our adversary C can simply run $\overline{\mathcal{K}}$ itself and hope to re-generate the challenge key, and when this does happen, it can forge a message using the secret key. This C succeeds by the correctness of the signature scheme.

We are now in a position to show that an adversary winning G_3 with good probability can be used to break the Φ-RKA security of $\mathcal{IBE}$. We construct D such that

$$\mathbf{Adv}_{\mathcal{IBE}, D, \Phi}^{\text{ibe-rka}}(k) = 2 \Pr[G_3^A \Rightarrow \text{true}] - 1. \tag{15}$$

Adversary D takes input (π, mpk) and has access to oracles LR, KD in the game PKE. It first executes

> $\overline{\pi} \leftarrow_\$ \overline{\mathcal{P}}(1^k) \,;\, \overline{vk}^* \leftarrow \bot$

and then runs adversary A on input $(\pi \,\|\, \overline{\pi}, mpk)$. A will expect oracles LR, DEC in the game IBE, and D provides answers to the queries of A as follows. When A queries $\text{LR}(m_0, m_1)$, D responds via

If $|m_0| \neq |m_1|$ then return $\perp$

$(\overline{vk}^*, \overline{sk}^*) \leftarrow_\$ \overline{\mathcal{K}}(\pi)$; $c^* \leftarrow \mathrm{LR}(\overline{vk}^*, m_0, m_1)$

If $c^* \perp$ then return $\perp$

$\overline{\sigma}^* \leftarrow \overline{\mathcal{S}}(\pi, \overline{sk}^*, c^*)$

Return $(c^*, \overline{vk}^*, \overline{\sigma}^*)$.

Here the call to LR is a call to the oracle of D in game IBE. When A queries $\mathrm{DEC}(\phi, (c, \overline{vk}, \overline{\sigma}))$, D responds via

If $\overline{\mathcal{V}}(\overline{\pi}, \overline{vk}, c, \overline{\sigma}) = 0$ then return $\perp$

$dk' \leftarrow \mathrm{KD}(\phi, \overline{vk})$; If $dk' = \perp$ then return $\perp$

$m \leftarrow \mathcal{D}(\pi, \overline{vk}, dk', c)$

Return m

We have that D properly simulates G_3 for A. In particular, the rules for returning $\perp$ in the LR and DEC oracles are correct, by construction.

Returning to (10), and using (11), (12), (13), (14), (15), we have

$$
\begin{aligned}
\mathbf{Adv}^{\mathrm{pke\text{-}cca}}_{\mathcal{PKE}, A, \Phi}(k) \;\; &= \;\; 2\Pr[\mathrm{G}_0^A \Rightarrow \mathsf{true}] - 1 \\[2mm]
&\leq \;\; 2(\Pr[\mathrm{G}_1^A \Rightarrow \mathsf{true}] + \Pr[\mathrm{BAD}(\mathrm{G}_0^A)]) - 1 \\[2mm]
&\leq \;\; 2(\Pr[\mathrm{G}_1^A \Rightarrow \mathsf{true}] + \mathbf{Adv}^{\mathrm{susig}}_{\overline{\mathcal{DS}}, B}(k)) - 1 & (16) \\[2mm]
&= \;\; 2(\Pr[\mathrm{G}_2^A \Rightarrow \mathsf{true}] + \mathbf{Adv}^{\mathrm{susig}}_{\overline{\mathcal{DS}}, B}(k)) - 1 & (17) \\[2mm]
&= \;\; 2(\Pr[\mathrm{G}_3^A \Rightarrow \mathsf{true}] + \Pr[\mathrm{G}_2^A \Rightarrow \mathsf{true}] - \Pr[\mathrm{G}_3^A \Rightarrow \mathsf{true}] + \mathbf{Adv}^{\mathrm{susig}}_{\overline{\mathcal{DS}}, B}(k)) - 1 & (18) \\[2mm]
&= \;\; 2(\Pr[\mathrm{G}_3^A \Rightarrow \mathsf{true}] + \Pr[\mathrm{G}_2^A \Rightarrow \mathsf{true}] - \Pr[\mathrm{G}_3^A \Rightarrow \mathsf{true}] + \mathbf{Adv}^{\mathrm{susig}}_{\overline{\mathcal{DS}}, B}(k)) - 1 & (19) \\[2mm]
&\leq \;\; 2(\Pr[\mathrm{G}_3^A \Rightarrow \mathsf{true}] + \Pr[\mathrm{BAD}(\mathrm{G}_2^A)] + \mathbf{Adv}^{\mathrm{susig}}_{\overline{\mathcal{DS}}, B}(k)) - 1 & (20) \\[2mm]
&\leq \;\; 2(\Pr[\mathrm{G}_3^A \Rightarrow \mathsf{true}] + \mathbf{Adv}^{\mathrm{susig}}_{\overline{\mathcal{DS}}, C}(k) + \mathbf{Adv}^{\mathrm{susig}}_{\overline{\mathcal{DS}}, B}(k)) - 1 & (21) \\[2mm]
&= \;\; 2\Pr[\mathrm{G}_3^A \Rightarrow \mathsf{true}] - 1 + 2\mathbf{Adv}^{\mathrm{susig}}_{\overline{\mathcal{DS}}, C}(k) + 2\mathbf{Adv}^{\mathrm{susig}}_{\overline{\mathcal{DS}}, B}(k) & (22) \\[2mm]
&= \;\; \mathbf{Adv}^{\mathrm{ibe\text{-}rka}}_{\mathcal{IBE}, D, \Phi}(k) + 2\mathbf{Adv}^{\mathrm{susig}}_{\overline{\mathcal{DS}}, C}(k) + 2\mathbf{Adv}^{\mathrm{susig}}_{\overline{\mathcal{DS}}, B}(k). & (23)
\end{aligned}
$$

This establishes the theorem.