Improved Constructions of PRFs Secure Against Related-Key Attacks

Kevin Lewi, Hart Montgomery, and Ananth Raghunathan

Stanford University

Abstract. Building cryptographic primitives that are secure against related-key attacks (RKAs) is a well-studied problem by practitioners and theoreticians alike. Practical implementations of block ciphers take into account RKA security to mitigate fault injection attacks. The theoretical study of RKA security was initiated by Bellare and Kohno (Eurocrypt '03). In Crypto 2010, Bellare and Cash introduce a framework for building RKA-secure pseudorandom functions (PRFs) and use this framework to construct RKA-secure PRFs based on the decision linear and DDH assumptions.

We build RKA-secure PRFs by working with the Bellare-Cash framework and the LWE- and DLIN-based PRFs recently constructed by Boneh, Lewi, Montgomery, and Raghunathan (Crypto '13). As a result, we achieve the first RKA-secure PRFs from lattices. In addition, we note that our DLIN-based PRF (based on multilinear maps) is the first RKA-secure PRF for affine classes under the DLIN assumption, and the first RKA-secure PRF against a large class of polynomial functions under a natural generalization of the DLIN assumption. Previously, RKA security for higher-level primitives (such as signatures and IBEs) were studied in Bellare, Paterson, and Thomson (Asiacrypt '12) for affine and polynomial classes, but the question of RKA-secure PRFs for such classes remained open.

Although our RKA-secure LWE-based PRF only applies to a restricted linear class, we show that by weakening the notion of RKA security, we can handle a significantly larger class of affine functions. Finally, the results of Bellare, Cash, and Miller (Asiacrypt '11) show that all of our RKA-secure PRFs can be used as building blocks for a wide variety of public-key primitives.

Keywords: related-key attacks, pseudorandom functions, learning with errors.

1 Introduction

The usual notions of security for cryptographic primitives do not address the possibility that an attacker could adversarially modify the internal state of hardware devices that

implement the primitive. Indeed, fault injection attacks (and other types of side-channel attacks including cold-boot attacks [22], timing attacks [24, 16], and power analysis attacks [27]) have shown that our traditional security definitions are not sufficient for most practical implementations of provably secure cryptographic primitives [12, 13, 33, 6].

To deal with fault injection attacks, cryptographers have developed the notion of related-key attack (RKA) security. RKA security definitions [9] capture the following notion: in addition to allowing the adversary to make input queries on the primitive for a randomly chosen secret key, the adversary is allowed to make input queries on the primitive for adversarially chosen "related-key deriving" functions $\phi \in \Phi$ of a randomly chosen secret key (where Φ is a function family specified in advance). This notion can be used to show that certain classes of tampering attacks are ineffective against primitives proven secure in the presence of RKAs.

In the past few years, there has been much work in constructing RKA-secure primitives [7, 8, 3, 11, 35, 10]. In addition, RKA security is also of interest to practitioners, particularly in the design of block ciphers [19, 23, 36]. In this work, we will focus our attention on building one of the most basic of the RKA primitives—pseudorandom functions (PRFs). Not only do PRFs find applications in many real-world implementations where side-channel attacks are possible (and hence RKA security becomes relevant) [6], but RKA-secure PRFs are also known to imply RKA security for a wide range of more advanced primitives, including signatures, identity-based encryption, and both public-key and private-key chosen ciphertext secure encryption [8].

1.1   Background and Related Work

Bellare and Cash [7] developed the first RKA-secure PRF for a non-trivial class of functions. Instantiations prior to [7] on RKA-secure PRFs required ideal ciphers, random oracles, or non-standard assumptions [26, 9]. In addition, Bellare and Cash develop a novel framework (which we call the BC framework) for building RKA-secure PRFs, and show how the DDH assumption implies an RKA-secure PRF for the class $\Phi_{\text{lin}} = \{\phi_{\mathbf{a}} : \mathbb{Z}_q^m \to \mathbb{Z}_q^m \mid \phi_{\mathbf{a}}(\mathbf{k}) = \mathbf{k} + \mathbf{a}\}_{\mathbf{a} \in \mathbb{Z}_q^m}$, the class of all linear transformations to the key. Additionally, they construct an RKA-secure PRF under the DLIN assumption [34, 30] for an interesting multiplicative class Φ (where related keys are derived from scalar multiples of components of the key).

Bellare *et al.* [8] explore the possibilities of transferring RKA security from one primitive to another (while preserving the class Φ of related-key deriving functions). In particular, they show that RKA-secure PRFs can be used to construct a wide variety of higher-level RKA-secure primitives. Thus, improvements in building RKA-secure PRFs have wide applicability to RKA-secure public-key cryptographic primitives.

Applebaum *et al.* [3] show how to build RKA-secure symmetric encryption from a variety of hardness assumptions for linear related-key attacks. Wee [35] presents chosen ciphertext RKA-secure public-key encryption scheme constructions from the DBDH and LWE assumptions, also for linear related-key attacks. Finally, Bellare *et al.* [11] show how to build RKA-secure variants from a variety of primitives discussed in [8] for more expressive classes Φ including affine and polynomial function families. However, constructing RKA-secure PRFs for affine or polynomial Φ is notably left open. Concurrently, Bellare *et al.* [10] build RKA-secure signature schemes against related-key

deriving functions drawn from such classes of polynomials. Their construction relies on RKA-secure one-way functions which appear to be easier to build under standard assumptions (as opposed to RKA-secure PRFs).

PRFs are extremely well-studied primitives and have been built from a wide variety of assumptions [29, 18, 25, 15, 5, 14]. Currently known RKA-secure PRFs only consider the Naor-Reingold [29] and Lewko-Waters [25] PRFs. We note that PRFs constructed by Boneh *et al.* [14] satisfy an additional "key homomorphism" property which we find useful in constructing RKA-secure PRFs. Our constructions are based on the PRFs considered in this work.

1.2   Our Contributions

Lattice-based RKA-secure PRFs. We present the first lattice-based PRFs secure against related-key attacks. Our construction achieves RKA security under the standard LWE assumption against the class of related-key functions $\Phi_{\mathsf{lin}^*} = \{\phi_{\mathbf{a}} : \mathbb{Z}_q^m \to \mathbb{Z}_q^m \mid \phi_{\mathbf{a}}(\mathbf{k}) = \mathbf{k} + \mathbf{a}\}_{\mathbf{a} \in (\frac{q}{p})\mathbb{Z}_q^m}$ over the key space $\mathcal{K} = \mathbb{Z}_q^m$. The class $(\frac{q}{p})\mathbb{Z}_q^m$ here denotes the vectors in $\mathbb{Z}_q^m$ whose entries are all multiples of q/p (where p divides q). This linear RKA class Φ_{lin^*} is a restricted case of the linear class in [7, Section 6], but our construction offers two advantages: it is the first LWE-based RKA-secure PRF (as opposed to the DDH-based construction in [7]) and its proof does not require a simulator that runs in time exponential in the input length.[1] Ideally we would like to address RKA security for the entire class of linear key shifts, but we only achieve a weaker notion of security. However, these restrictions are quite plausible as they translate to an adversary that can inject faults into the higher order bits of the key.[2]

RKA security against an affine class of related keys. Next, we show how the powerful multilinear map abstraction by Garg *et al.* [20] along with the DLIN assumption in this abstraction can be used to construct PRFs with RKA security against a very large and natural class of affine key transformations $\Phi_{\mathsf{aff}} = \{\phi_{\mathbf{C},\mathbf{B}} : \mathbb{Z}_p^{m \times \ell} \to \mathbb{Z}_p^{m \times \ell} \mid \phi_{\mathbf{C},\mathbf{B}}(\mathbf{K}) = \mathbf{CK} + \mathbf{B}\}$ over the key space $\mathcal{K} = \mathbb{Z}_p^{m \times \ell}$. For Φ_{aff}, we require that $\mathbf{C}$ comes from a family of *invertible* matrices and that Φ_{aff} be claw-free—for all $\phi_1, \phi_2 \in \Phi_{\mathsf{aff}}$ and $\mathbf{K} \in \mathcal{K}$, $\phi_1(\mathbf{K}) \neq \phi_2(\mathbf{K})$.

Both restrictions arise from a technical requirement under the BC framework. As noted in [7, 11], some restrictions must be placed on Φ_{aff} in order for PRFs to achieve RKA security against them (for example, Φ_{aff} cannot include constant functions $\phi(\mathbf{K}) = \mathbf{B}$). Hence, our class Φ_{aff} is essentially the most expressive affine class of transformations for which RKA PRF security is still attainable under the Bellare-Cash framework. In fact, there are no known PRFs which are RKA-secure against a class which does not have the claw-free restriction. Bellare *et al.* [11] constructed higher-level primitives RKA-secure

---

[1] We note that we require the LWE assumption to hold over superpolynomially-sized modulus q, but this is a well-studied and widely-used assumption [31, 5, 1, 14].

[2] We note that when q and p are powers of 2, Φ_{lin^*} captures all functions that perform linear shifts on the entries of the key that do not modify the $\log(q/p)$-least significant bits of each entry.

against affine classes, but left open the problem of constructing such a PRF (for which we provide an answer).

Unique-input RKA security against an affine class. We note, however, that the assumption that there exists an instantiation of the Garg *et al.* multilinear map abstraction [20] for which DLIN holds is a fairly strong assumption. This raises the following question: Can we achieve a similar result for RKA PRF security against affine transformations from a more standard assumption? We answer this question in the affirmative by considering a slightly weaker notion of RKA security, denoted *unique-input* RKA security, where adversary queries are restricted to unique inputs. We build RKA-secure PRFs from the LWE assumption that can handle the class of transformations $\Phi_{\mathsf{ln\text{-}aff}} = \{\phi_{\mathbf{C},\mathbf{B}} : \phi_{\mathbf{C},\mathbf{B}}(\mathbf{K}) = \mathbf{C}\mathbf{K} + \mathbf{B}\}$, where $\mathbf{C}$ is a full-rank "low-norm" matrix and $\mathbf{B}$ is an arbitrary matrix in $\mathbb{Z}_q^{m \times m}$ from the LWE assumption. We observe that under this weaker notion of security, our class is significantly more expressive than our first result from lattices because it allows for the addition of arbitrary vectors. However, this requires us to work outside the Bellare-Cash framework. We leave it as an open problem to construct "truly" RKA-secure PRFs from LWE (or other standard assumptions, such as DDH) for an affine class of key transformations.

Unique-input RKA security against a class of polynomials. We further explore the connection between key homomorphism and unique-input RKA security by using the multilinear map abstraction to tackle a polynomial class of related-key functions. More specifically, we consider the class of polynomials $\Phi_{\mathsf{poly}(d)}$ of bounded degree d over matrices $\mathbb{Z}_q^{m \times m}$ and consider a natural exponent assumption over multilinear maps called the Multilinear Diffie-Hellman Exponent (MDHE) assumption. For technical reasons, we require that at least one of the polynomial's non-constant coefficient matrices is full-rank. This natural restriction simply ensures that the output of the polynomial is sufficiently random given a uniformly drawn input of a special form. We note that the MDHE assumption is a natural and fairly plausible generalization of the DLIN assumption.

Finally, we can apply the results of [8] to get Φ-RKA security for signatures, identity-based encryption, and public and private key CCA encryption from our Φ-RKA-secure PRFs.

1.3   Our Techniques

At a high level, we use the Bellare-Cash framework with the (LWE- and DLIN-based) key homomorphic PRFs from Boneh *et al.* [14] to construct RKA-secure PRFs against the classes Φ_{lin^*} and Φ_{aff}. Below, we give an outline of the framework and note that key homomorphic PRFs are a natural starting point due to the malleability requirement of the framework.

Bellare-Cash framework. The only known construction of RKA-secure PRFs to date is that of Bellare and Cash [7]. In their framework, Bellare and Cash identify sufficient properties for constructing an RKA-secure PRF. They first consider PRFs $F : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$ that are *key malleable*—PRFs which have an efficient algorithm (denoted a transformer T) that when given an input $(\phi, x) \in \Phi \times \mathcal{X}$ and oracle access to $F(k, \cdot)$ computes $F(\phi(k), x)$. In addition, T must satisfy a *uniformity* property, namely, when $F(k, \cdot)$ is

replaced with a random function $f(\cdot)$, the outputs of T on inputs $(\phi_1, x_1), \ldots, (\phi_Q, x_Q)$ for distinct $x_1, \ldots, x_Q$ are uniform and independently distributed. The framework also requires the existence of a *key fingerprint*—an input $w \in \mathcal{X}$ such that for all $k \in \mathcal{K}$ and distinct $\phi_1, \phi_2 \in \Phi$, $F(\phi_1(k), w) \neq F(\phi_2(k), w)$.

For a class Φ with a suitable key malleable PRF, a fingerprint w, and a collision-resistant hash function that satisfies a simple *compatiblity* property H_{com} (see Definition 2.8), under the Bellare-Cash framework, the authors show that the PRF $F_{\mathsf{rka}}(k, x) = F(k, H_{\mathsf{com}}(x, F(k, w)))$ is Φ-RKA-secure.

Applying the BC framework to the DLIN-based PRF. Our starting point is the construction of a DLIN-based key homomorphic PRF by Boneh *et al.* [14], who note that key homomorphic PRFs are key malleable. In this work, we generalize this PRF to operate with the key space $\mathcal{K} = \mathbb{Z}_p^{m \times \ell}$ instead of $\mathbb{Z}_p^\ell$. The PRF has public parameters $\mathbf{A}_0, \mathbf{A}_1 \in \mathbb{Z}_q^{\ell \times \ell}$. On input x, the PRF is of the form $(g_\ell)^{\mathbf{W}}$ for $\mathbf{W} = \mathbf{KP}$ where $\mathbf{P} \in \mathbb{Z}_p^{\ell \times \ell}$ is the publicly computable matrix $\mathbf{A}_{x_\ell} \mathbf{A}_{x_{\ell-1}} \cdots \mathbf{A}_{x_1}$ (that only depends on the bits of x) and g_ℓ is the generator of a group with a multilinear map. This additional algebraic structure allows us to consider the class of affine related-key deriving functions of the form $\mathbf{CK} + \mathbf{B}$ for matrices $\mathbf{C} \in \mathbb{Z}_q^{m \times m}$ and $\mathbf{B} \in \mathbb{Z}_q^{m \times \ell}$. The pseudorandomness of the PRF holds by a straightforward hybrid argument, noting that the rows of $\mathbf{K}$ are now identical to independent keys of the original PRF.

Working in the exponent, given access to an oracle that computes $\mathbf{W}$ and an input $\phi_{\mathbf{C}, \mathbf{B}}$, it is easy to construct a transformer that computes $\mathbf{W}' = \mathbf{CW} + \mathbf{BP}$. From some simple algebra, one can verify that this indeed computes the exponent $\mathbf{W}'$ corresponding to $F_{\mathrm{DLIN}}(\phi(\mathbf{K}), x)$. In addition, as long as $\mathbf{C}$ is restricted to the set of full-rank matrices, it follows that the transformer described above outputs uniform matrices if $\mathbf{W}$ corresponds to the outputs of a random function. From this, the rest of the BC framework can be applied and is shown in Section 3.2. We note here that the restriction that Φ is claw-free seems to be inherently required in applying the BC framework (here, we require it in constructing a suitable fingerprint), and we do not overcome this limitation in our construction either.[3]

Applying the BC framework to the LWE-based PRF. Recollect that Boneh *et al.* construct an "almost" key homomorphic LWE-based PRF F which on input x is of the form $\lfloor \mathbf{Pk} \rfloor_p$, where $\mathbf{P} = \mathbf{A}_{x_\ell} \mathbf{A}_{x_{\ell-1}} \cdots \mathbf{A}_{x_1}$. (Here, $\lfloor x \rfloor_p$ for $x \in \mathbb{Z}_q$ denotes multiplying x by p/q and rounding the result to $\mathbb{Z}_p$.) Unfortunately, the "almost"-ness of the key homomorphism disallows a direct argument of key malleability. Furthermore, a transformer which is "almost" key malleable (in the same sense) is still insufficient for instantiating the BC framework.

This limitation can be overcome by observing that $F(\mathbf{k}_1, x) + F(\mathbf{k}_2, x) = F(\mathbf{k}_1 + \mathbf{k}_2, x)$ if the entries of either $\mathbf{k}_1$ or $\mathbf{k}_2$ are all multiples of q/p. This property is sufficient to show that F is key malleable with respect to the class Φ_{lin^*}, where $\mathbf{k}_2$ is required to be an element of $(\frac{q}{p})\mathbb{Z}_q^m$. Additionally, this restriction is needed show that any fixed input $w \in \{0, 1\}^\ell$ acts as a key fingerprint for F under the class Φ_{lin^*}. It seems likely

---

[3] However, in [8], the authors overcome this barrier and achieve RKA security for PRGs, not PRFs, against a class Φ which is not claw-free.

that this restriction is in fact necessary for applying the BC framework, leaving this the most expressive class achievable for the LWE-based PRF F.

One natural question to ask is whether the Banerjee *et al.* [5] LWE-based PRF can be used instead of F. We note that their PRF is not key homomorphic and hence the above approach does not apply. However, we leave open the question of achieving unique-input RKA security for their PRF (see Section 6).

Unique-input adversaries. As was observed by Bellare and Cash, key malleability is intuitively useful in constructing RKA security because it allows us to simulate $F(\phi(k), \cdot)$ without access to the key k but also leads to a simple related-key attack against any class that contains the functions ϕ_{id} (the identity function) and any $\phi' \neq \phi_{\mathrm{id}}$. The difficulty in achieving security lies in the adversary's ability to request multiple related-key deriving functions on the same input x. Given ϕ_{id}, to attack the pseudorandomness, the adversary can run the transformer for ϕ' himself and compare the output of the transformer to the output of the oracle on (ϕ', x). Thus, Bellare and Cash require additional tools.

However, the notion of key malleability suffices to show security against unique-input adversaries, where the adversary's queries are restricted to distinct x's. In extending the RKA-secure LWE-based PRF to a class of affine functions, as discussed earlier in this section, the presence of the rounding does not directly imply key malleability. However, in Section 4, we work through the proof of security of the pseudorandomness of F, along the lines of the proof in [14], to consider its RKA security against the larger class $\Phi_{\mathsf{ln\text{-}aff}}$. We show that the structure of the PRF allows us to simulate, in addition to PRF queries on input x, RKA queries for functions $\phi \in \Phi_{\mathsf{ln\text{-}aff}}$. As in [14], the proof works through several hybrid arguments that modify a challenger from a truly random function to a pseudorandom function that also provides answers to RKA queries $(\phi, x) \in \Phi_{\mathsf{ln\text{-}aff}} \times \{0,1\}^{\ell}$.

The low-norm restriction on the matrix $\mathbf{C}$ in $\phi_{\mathbf{C},\mathbf{B}} \in \Phi_{\mathsf{ln\text{-}aff}}$ is required to ensure that when using LWE challenges in the hybrids, the noise does not grow larger than what the rounding allows. In the final hybrid, the adversary interacts with uniform and independently chosen outputs corresponding to inputs x_i. As long as the adversary is restricted to unique inputs, this interaction is identical to the game where the adversary receives uniform and independent (consistent) values on queries (ϕ, x). This is sufficient to show RKA security. Whether we can take advantage of the algebraic structure of other pseudorandom functions to directly prove unique-input RKA security is an interesting question.

Unique-input security against a class of polynomials. We have shown how under the DLIN and LWE assumptions we can build RKA-secure PRFs for classes of affine functions, but unfortunately we do know how to extend these results to handle classes of polynomials. However, in Section 5, we show that the PRF F_{DLIN} (defined in Section 3.2) is RKA-secure against unique-input adversaries under the (new) d-MDHE assumption (see Definition 2.5) for a class of degree-d polynomials.

For integers ℓ, d, and a prime p, we consider the class $\Phi_{\mathsf{poly}(d)}$ consisting of all degree-d polynomials over $\mathbb{Z}_p^{\ell \times \ell}$ of the form $P(\mathbf{K}) = \sum_{i=0}^{d} \mathbf{C}_i \cdot \mathbf{K}^i$, where $\mathbf{C}_0, \dots, \mathbf{C}_d, \mathbf{K} \in \mathbb{Z}_p^{\ell \times \ell}$ and at least one of $\mathbf{C}_1, \dots, \mathbf{C}_d$ is full rank. To prove the RKA security of F_{DLIN} against unique-input adversaries, we consider a series of hybrid experiments which

respond to queries $(\phi_{P(\cdot)}, x) \in \Phi_{\mathsf{poly}(d)} \times \{0,1\}^\ell$, where $P(\mathbf{S}) = \sum_{i=0}^d \mathbf{C}_i \cdot \mathbf{S}^i$, by choosing d uniformly random, *independent* secrets $\mathbf{K}_1, \ldots, \mathbf{K}_d$ and computing the weighted sum $\mathbf{C}_0 + \sum_{i=1}^d \mathbf{C}_i \cdot \mathbf{K}_i$, as opposed to choosing a single uniformly random secret $\mathbf{S}$ and computing $P(\mathbf{S})$. We show how an adversary which distinguishes between these two cases can be used to break the d-MDHE assumption, and then we use the techniques used to prove the pseudorandomness of F_{DLIN} to complete the argument.

The additional requirement of at least one of $\mathbf{C}_1, \ldots, \mathbf{C}_d$ being full rank is only needed to ensure that a sufficient amount of entropy from the secret key will remain in the output of the PRF. Note that this restriction on $\Phi_{\mathsf{poly}(d)}$ rules out polynomials P for which the output of P on randomly chosen key can be predicted (as an example consider *constant* polynomials $P(\mathbf{K}) = \mathbf{C}$ for some fixed $\mathbf{C} \in \mathbb{Z}_p^{\ell \times \ell}$), for which achieving RKA security is impossible. We believe $\Phi_{\mathsf{poly}(d)}$ captures what is essentially the most expressive class of bounded-degree polynomials for RKA-secure PRFs.

Organization. In Section 2 we introduce preliminary notation and definitions. In Section 3 we construct RKA-secure LWE- and DLIN-based PRFs using the BC framework. Then, in Section 4, we give an LWE-based RKA-secure PRF against unique-input adversaries for an affine class of transformations. In Section 5, we show how the DLIN-based PRF is secure against unique-input adversaries where the related-key attacks come from a class of bounded-degree polynomials. We conclude in Section 6. In the full version, we give additional preliminaries, missing proofs, and more details.

2 Preliminaries

2.1 Notation

Rounding. We define $\lfloor \cdot \rfloor$ to round a real number to the largest integer which does not exceed it. For integers q and p where $q \geq p \geq 2$, we define the function $\lfloor \cdot \rfloor_p : \mathbb{Z}_q \to \mathbb{Z}_p$ as $\lfloor x \rfloor_p = i$ where $i \cdot \lfloor q/p \rfloor$ is the largest multiple of $\lfloor q/p \rfloor$ which does not exceed x. For a vector $\mathbf{v} \in \mathbb{Z}_q^m$, we define $\lfloor \mathbf{v} \rfloor_p$ as the vector in $\mathbb{Z}_p^m$ obtained by rounding each coordinate of the vector individually.

When $p \mid q$, we let $(\frac{q}{p})\mathbb{Z}_q$ denote the subgroup of $\mathbb{Z}_q$ comprising the set $\{(q/p) \cdot x \mid x \in \mathbb{Z}_q\}$. The following lemma follows from some elementary arithmetic.

Lemma 2.1. *For any $u \in (\frac{q}{p})\mathbb{Z}_q$ and $x \in \mathbb{Z}_q$ such that $u \equiv x(q/p) \mod q$ and any $y \in \mathbb{Z}_q$,*

$$\lfloor y + u \rfloor_p = \lfloor y \rfloor_p + \lfloor u \rfloor_p = \lfloor y \rfloor_p + x \pmod{p}.$$

Groups. For a matrix $\mathbf{M}$, we let the component-wise exponentiation $g^{\mathbf{M}}$ denote a matrix with entries $g^{\mathbf{M}_{i,j}}$. We let $\left(g^{\mathbf{A}}\right)^{\mathbf{B}}$ denote the matrix with entries $g^{(\mathbf{AB})_{i,j}}$. We let $\mathbf{Rk}_i(\mathbb{Z}_p^{a \times b})$ denote the set of all $a \times b$ matrices over $\mathbb{Z}_p$ of rank i.

Pseudorandom Functions. Informally, a PRF [21] is an efficiently computable function $F : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$ such that no efficient adversary can distinguish the function from a truly random function given only black-box access. In this paper, we allow the PRF to additionally take public parameters pp. The advantage $\mathbf{Adv}_F^{\mathsf{prf}}(\cdot)$ against the PRF is defined in a standard manner and deferred to the full version due to space constraints.

2.2   RKA-secure PRFs

For a class of related-key deriving functions $\Phi = \{\phi : \mathcal{K} \to \mathcal{K}\}$, the notion of Φ-RKA security for a PRF $F : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$ is defined using an experiments between a challenger and an adversary $\mathcal{A}$. For $b \in \{0, 1\}$ define the following experiment $\mathsf{Expt}_b^{\mathsf{prf\text{-}rka}}$:

1. Given security parameter λ, the challenger samples and publishes public parameters pp to the adversary. Next, the challenger chooses a random key $k \in \mathcal{K}$ and if $b = 0$, sets $f(\cdot) \stackrel{\mathsf{def}}{=} F(k, \cdot)$. Otherwise, if $b = 1$, the challenger chooses a random keyed function $f : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$.
2. The adversary (adaptively) sends input queries $(\phi_1, x_1), \ldots, (\phi_Q, x_Q)$ in $\Phi \times \mathcal{X}$ and receives back $f(\phi_1(k), x_1), \ldots, f(\phi_Q(k), x_Q)$.
3. The adversary outputs a bit $b' \in \{0, 1\}$, and the experiment also outputs b'.

Definition 2.2 (RKA-secure PRF for Φ). *A PRF $F : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$ is RKA-secure with respect to class Φ if for all efficient adversaries $\mathcal{A}$ the quantity*

$$\mathbf{Adv}_{\Phi, F}^{\mathsf{prf\text{-}rka}} (\mathcal{A}) \stackrel{\mathsf{def}}{=} \left| \Pr\left[\mathsf{Expt}_0^{\mathsf{prf\text{-}rka}} = 1 \right] - \Pr\left[\mathsf{Expt}_1^{\mathsf{prf\text{-}rka}} = 1 \right] \right|$$

is negligible.

Unique-input RKA security (cf. [7]). We say that an adversary is *unique-input* in the above security game if the input queries $(\phi_1, x_1), \ldots, (\phi_Q, x_Q) \in \Phi \times \mathcal{X}$ are such that $x_1, \ldots, x_Q$ are distinct. A PRF is *unique-input* RKA-secure if it is RKA secure against unique-input adversaries.

2.3   Security Assumptions

Learning with errors (LWE) assumption. The LWE problem was introduced by Regev [32] who showed that solving the LWE problem on average is as hard as (quantumly) solving several standard lattice problems in the worst case.

Definition 2.3 (Learning With Errors). *For integers $q > 2$ and a noise distribution χ over $\mathbb{Z}_q$, the learning with errors problem (LWE) over n-dimensional vectors is to distinguish between the distributions $\{\mathbf{A}, \mathbf{A}^\intercal \mathbf{s} + \chi\}$ and $\{\mathbf{A}, \mathbf{u}\}$, where $m = \mathrm{poly}(n)$, $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\chi \leftarrow \chi^m$, and $\mathbf{u} \leftarrow \mathbb{Z}_q^m$.*

Regev [32] shows that for a certain noise distribution $\chi = \overline{\Psi}_\alpha$,[4] for n polynomial in λ and $q > 2\sqrt{n}/\alpha$, the LWE problem is as hard as the worst-case SIVP and GapSVP under a quantum reduction (see also [31, 17] for classical reductions). These results have been extended to show that $\mathbf{s}$ can be sampled from a low-norm distribution (in particular, from the noise distribution χ) and the resulting problem is as hard as the

---

[4] For an $\alpha \in (0, 1)$ and a prime q, let $\overline{\Psi}_\alpha$ denote the distribution over $\mathbb{Z}_q$ of the random variable $\lceil qX \rfloor \pmod{q}$ where X is a normal random variable with mean 0 and standard deviation $\alpha/\sqrt{2\pi}$.

basic LWE problem [2]. Similarly, the noise distribution χ can be a simple low-norm distribution [28]. Boneh *et al.* [14] show that the variant of LWE where the entries of $\mathbf{A}$ are binary and $m > n \log q$ is equivalent (modulo a $\log q$-factor loss in dimension) to LWE over n-dimensional vectors. In this work, we let $B \in \mathbb{R}$ be an error bound such that for $\chi \leftarrow \overline{\Psi}_\alpha$, $|\chi| \leq B$ with overwhelming probability.

Low-norm matrix LWE. We work with the right-multiplied matrix form of (low-norm) LWE, namely, that for a uniformly drawn $\mathbf{A} \leftarrow \{0,1\}^{m \times 2m}$, $\mathbf{U} \leftarrow \mathbb{Z}_q^{m \times 2m}$, $\mathbf{S} \leftarrow \mathbb{Z}_q^{m \times m}$, and $\mathbf{X} \leftarrow \chi^{m \times 2m}$, the problem is to distinguish between the distributions $\{\mathbf{A}, \mathbf{SA} + \mathbf{X}\}$ and $\{\mathbf{A}, \mathbf{U}\}$.

To compare it to the low-norm LWE variant in [14], we note that $\{\mathbf{A}, \mathbf{SA} + \mathbf{X}\}$ and $\{\mathbf{A}, \mathbf{A}^\mathsf{T}\mathbf{S} + \mathbf{X}^\mathsf{T}\}$ are distributed identically, and a standard hybrid argument shows that any adversary which can distinguish $\{\mathbf{A}, \mathbf{A}^\mathsf{T}\mathbf{S} + \mathbf{X}^\mathsf{T}\}$ from $\{\mathbf{A}, \mathbf{U}\}$ can be used to distinguish $\{\mathbf{A}, \mathbf{A}^\mathsf{T}\mathbf{s} + \chi\}$ from $\{\mathbf{A}, \mathbf{u}\}$ with only a $(1/m)$-factor loss in advantage.

The DLIN assumption in multilinear groups. In Section 3.2, we rely on the decisional linear (DLIN) assumption (as stated in Boneh *et al.* [14]) for the Garg *et al.* abstraction of graded multilinear maps [20]. Consider a sequence of groups $\vec{\mathbb{G}} = (\mathbb{G}_1, \ldots, \mathbb{G}_\ell)$ with a set of bilinear maps $\hat{e}_i$ for $i \in [1, \ell-1]$, and a generator g of $\mathbb{G}_1$.

Definition 2.4 (Decisional Linear). *The κ-decisional linear (κ-DLIN) assumption in the presence of a graded ℓ-linear map states that for any integers $a, b \geq \kappa$, and for any $\ell \leq j < \kappa$ the distributions*

$$\left\{g, g^\mathbf{X}\right\}_{\mathbf{X} \leftarrow \mathbf{Rk}_j\left(\mathbb{Z}_p^{a \times b}\right)} \quad and \quad \left\{g, g^\mathbf{Y}\right\}_{\mathbf{Y} \leftarrow \mathbf{Rk}_\kappa\left(\mathbb{Z}_p^{a \times b}\right)}$$

are computationally indistinguishable, in the presence of $\vec{\mathbb{G}}$ and $\{\hat{e}_i\}_{i \in [1, \ell-1]}$.

The Multilinear Diffie-Hellman Exponent assumption. In Section 5, we will use the Multilinear Diffie-Hellman Exponent (MDHE) assumption, defined as follows. Consider a sequence of groups $\vec{\mathbb{G}} = (\mathbb{G}_1, \ldots, \mathbb{G}_\ell)$ with a set of bilinear maps $\hat{e}_i$ for $i \in [1, \ell-1]$, and a generator g of $\mathbb{G}_1$.

Definition 2.5 (Multilinear Diffie-Hellman Exponent). *The d-Multilinear Diffie-Hellman Exponent (d-MDHE) assumption in the presence of a graded ℓ-linear map (as abstracted by [20]) states that, in the presence of $\vec{\mathbb{G}}$ and $\{\hat{e}_i\}_{i \in [1, \ell-1]}$, for any integer $j \geq \ell$, the distribution*

$$\left\{g^\mathbf{A}, \left\langle g^{\mathbf{S}^i \cdot \mathbf{A}}\right\rangle_{i \in [d]}, g^\mathbf{B}, \left\langle g^{\mathbf{S}^i \cdot \mathbf{B}}\right\rangle_{i \in [d]}\right\}_{\mathbf{A}, \mathbf{B} \leftarrow \mathbf{Rk}_j\left(\mathbb{Z}_p^{j \times j}\right), \mathbf{S} \leftarrow \mathbb{Z}_p^{j \times j}}$$

is computationally indistinguishable from the distribution

$$\left\{g^\mathbf{A}, \left\langle g^{\mathbf{U}_i}\right\rangle_{i \in [d]}, g^\mathbf{B}, \left\langle g^{\mathbf{V}_i}\right\rangle_{i \in [d]}\right\}_{\mathbf{A}, \mathbf{B} \leftarrow \mathbf{Rk}_j\left(\mathbb{Z}_p^{j \times j}\right), \forall i \in [d], \mathbf{U}_i, \mathbf{V}_i \leftarrow \mathbb{Z}_p^{j \times j}}.$$

We note that the 1-MDHE assumption is essentially equivalent to the 2ℓ-DLIN assumption (where $j = \ell$ and $\kappa = 2\ell$ as in [14]), and hence the d-MDHE assumption can be seen as a generalization of DLIN assumption to the d^{th} exponent of the secret.

2.4   The Bellare-Cash Framework

Bellare and Cash [7] give a general framework (denoted the BC framework) for constructing RKA-secure PRFs for a class Φ using a key malleable PRF, a key fingerprint, and a collision-resistant hash function. We review their definitions and main theorem here.

Definition 2.6 (Key Malleable PRF). *A PRF $F : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$ is* key malleable *if there exists an efficient algorithm T, which on input $\phi \in \Phi$ and $x \in \mathcal{X}$ and with oracle access to $F(k, \cdot)$, which satisfies $\mathsf{T}^{F(k,\cdot)}(\phi, x) = F(\phi(k), x)$, for all $k \in \mathcal{K}$. Also, we require that for any distinct $x_1, \ldots, x_Q \in \mathcal{X}$, if $f : \mathcal{X} \to \mathcal{Y}$ is a truly random function, then $\mathsf{T}^{f(\cdot)}(\phi, x_1), \ldots, \mathsf{T}^{f(\cdot)}(\phi, x_Q)$ are distributed independently and uniformly in $\mathcal{Y}$.*

Definition 2.7 (Key Fingerprint). *An element $w \in \mathcal{X}$ is a* key fingerprint *if for all $k \in \mathcal{K}$ and distinct $\phi_1, \phi_2 \in \Phi$, $F(\phi_1(k), w) \neq F(\phi_2(k), w)$.*

Definition 2.8 (Compatible Hash Function). *For a fingerprint w, a hash function $H_{\mathsf{com}} : \mathcal{X} \times \mathcal{Y} \to \mathcal{R}$ is* compatible *if the set of oracle queries made by $\mathsf{T}^{F(k,\cdot)}(\phi, w)$ over all $\phi \in \Phi$ is disjoint from the set of oracle queries made by $\mathsf{T}^{F(k,\cdot)}(\phi, z)$ over all $z \in \mathcal{R}$ and $\phi \in \Phi$.*

Theorem 2.9 (c.f. [7, Theorem 3.1], paraphrased). *For a fixed class Φ of related-key deriving functions, let $F : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$ be a key malleable PRF for Φ, $w \in \mathcal{X}$ a key fingerprint for F and Φ, and $H_{\mathsf{com}} : \mathcal{X} \times \mathcal{Y} \to \mathcal{X}$ a compatible hash function. Define $F_{\mathsf{rka}} : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$ as*

$$F_{\mathsf{rka}}(k, x) = F(k, H_{\mathsf{com}}(x, F(k, w))).$$

For any probabilistic polynomial-time (PPT) adversary $\mathcal{A}$ against the RKA PRF F_{rka} for the class Φ, there exist PPT adversaries $\mathcal{B}$ against the PRF security of F_{LWE} and $\mathcal{C}$ against the collision-resistance of the hash function H_{com} such that

$$\mathbf{Adv}^{\mathsf{prf\text{-}rka}}_{\Phi, F_{\mathsf{rka}}}(\mathcal{A}) \leq \mathbf{Adv}^{\mathsf{prf}}_{F}(\mathcal{B}) + \mathbf{Adv}^{\mathsf{cr}}_{H_{\mathsf{com}}}(\mathcal{C}).$$

3   New RKA-secure PRFs Using the BC Framework

In this section, we use the BC framework [7] to construct new RKA-secure PRFs. We introduce two classes of related-key functions, a linear (Φ_{lin^*}) and an affine (Φ_{aff}) class, and show that the key homomorphic PRFs from Boneh *et al.* [14] can be used to instantiate the BC framework. The main technical challenge requires using the key homomorphism property to construct appropriate *transformers* required in the BC framework.

3.1   RKA-secure PRFs for a Restricted Linear Class Φ_{lin^*}

Boneh, Lewi, Montgomery, and Raghunathan [14] constructed the following PRF that is *almost* key homomorphic and showed its pseudorandomness under the LWE assumption.

The PRF F_{LWE}. For parameters m, p, and $q \in \mathbb{N}$ such that $p \mid q$, the public parameters of the PRF are binary matrices $\mathbf{A}_0, \mathbf{A}_1 \in \mathbb{Z}_p^{m \times m}$. The PRF key is a vector $\mathbf{k} \in \mathbb{Z}_q^m$. The PRF $F_{\text{LWE}} \colon \mathbb{Z}_q^m \to \mathbb{Z}_p^m$ is defined as follows:

$$F_{\text{LWE}}(\mathbf{k}, x) = \left\lfloor \prod_{i=1}^{\ell} \mathbf{A}_{x_i} \cdot \mathbf{k} \right\rfloor_p. \tag{3.1}$$

Theorem 3.1 (cf. [14], paraphrased). *The function F_{LWE} is pseudorandom under the LWE assumption for suitable choices of the parameters.*

The class $\Phi_{\text{lin}*}$. Recall the definition of $(\frac{q}{p})\mathbb{Z}_q$. We consider a class of linear RKA functions defined as follows:

$$\Phi_{\text{lin}*} = \{\phi_{\mathbf{a}} : \mathbb{Z}_q^m \to \mathbb{Z}_q^m \mid \phi_{\mathbf{a}}(\mathbf{k}) = \mathbf{k} + \mathbf{a}\}_{\mathbf{a} \in (\frac{q}{p})\mathbb{Z}_q^m}. \tag{3.2}$$

In other words, $\Phi_{\text{lin}*}$ is identical to the class $\Phi_{\text{lin}} = \{\phi_{\mathbf{a}} : \mathbb{Z}_q^m \to \mathbb{Z}_q^m \mid \phi_{\mathbf{a}}(\mathbf{k}) = \mathbf{k} + \mathbf{a}\}_{\mathbf{a} \in \mathbb{Z}_q^m}$ of all possible linear transformations of the key (the class for which an RKA-secure PRF is given in [7] under the DDH assumption), except that in $\Phi_{\text{lin}*}$ we have the added restriction that the transformation must be an element of $(\frac{q}{p})\mathbb{Z}_q^m$.

We use the homomorphic property of the PRF to construct a transformer, that we denote $\mathsf{T}_{\text{lin}}^{f(\cdot)}$, in a straightforward manner: $\mathsf{T}_{\text{lin}}^{f(\cdot)}(\phi_{\mathbf{a}}, x) := f(x) + F_{\text{LWE}}(\mathbf{a}, x)$. To use the BC framework, it is necessary to show that for the class of RKA functions $\Phi_{\text{lin}*}$, the PRF and the transformer satisfy the malleability and uniformity properties.

Lemma 3.2 (Malleability). *For all $\mathbf{k} \in \mathbb{Z}_q^m$, $\phi \in \Phi_{\text{lin}*}$, and $x \in \{0,1\}^{\ell}$, it holds that*

$$\mathsf{T}_{\text{lin}}^{F_{\text{LWE}}(\mathbf{k},\cdot)}(\phi, x) = F_{\text{LWE}}(\phi(\mathbf{k}), x). \tag{3.3}$$

Proof. Fix a key $\mathbf{k} \in \mathbb{Z}_q^m$ and $x \in \{0,1\}^{\ell}$. Let $\phi_{\mathbf{a}}$ denote a function in $\Phi_{\text{lin}*}$ corresponding to $\mathbf{a} \in (\frac{q}{p})\mathbb{Z}_q^m$. Define the product of matrices $\mathbf{P} = \prod_{i=1}^{\ell} \mathbf{A}_{x_i}$. From the definition of the transformer $\mathsf{T}_{\text{lin}}^{F_{\text{LWE}}(\mathbf{k},\cdot)}$ the left side of equation (3.3) equals $\lfloor \mathbf{Pk} \rfloor_p + \lfloor \mathbf{Pa} \rfloor_p$. The right side of the equation is $\lfloor \mathbf{P}(\mathbf{k} + \mathbf{a}) \rfloor_p = \lfloor \mathbf{Pk} + \mathbf{Pa} \rfloor_p$. As $\mathbf{a} \in (\frac{q}{p})\mathbb{Z}_q^m$, it holds that $\mathbf{Pa} \in (\frac{q}{p})\mathbb{Z}_q^m$. Applying Lemma 2.1 on each coordinate, it holds that $\lfloor \mathbf{Pk} + \mathbf{Pa} \rfloor_p = \lfloor \mathbf{Pk} \rfloor_p + \lfloor \mathbf{Pa} \rfloor_p$, as required. ∎

The following lemma follows straightforwardly from the definition of $\mathsf{T}_{\text{lin}}^{f(\cdot)}$.

Lemma 3.3 (Uniformity). *If $f : \{0,1\}^{\ell} \to \mathbb{Z}_p^m$ is a random function and $x_1, \ldots, x_Q \in \{0,1\}^{\ell}$ are distinct, for any functions $\phi_1, \ldots, \phi_Q \in \Phi_{\text{lin}*}$, the values $\mathsf{T}_{\text{lin}}^{f(\cdot)}(\phi_i, x_i)$ are independently and uniformly distributed in $\mathbb{Z}_p^m$.*

Next, we show that any $w \in \{0,1\}^{\ell}$ is a key fingerprint for $\Phi_{\text{lin}*}$.

Lemma 3.4 (Fingerprint). *For any $w \in \{0,1\}^{\ell}$, $\mathbf{k} \in \mathbb{Z}_q^m$, for any distinct $\phi_1, \phi_2 \in \Phi_{\text{lin}*}$, it holds that $F_{\text{LWE}}(\phi_1(\mathbf{k}), w) \neq F_{\text{LWE}}(\phi_2(\mathbf{k}), w)$.*

Proof. For $i \in \{1, 2\}$, let $\phi_i = \phi_{\mathbf{a}_i}$ for vectors $\mathbf{a}_i \in (\frac{q}{p})\mathbb{Z}_q^m$. Let $\mathbf{P} = \prod_{i=1}^{\ell} \mathbf{A}_{w_i}$, the product of full-rank matrices. As ϕ_1 and ϕ_2 are *distinct* and $\mathbf{P}$ is full-rank over $\mathbb{Z}_q$, it holds that $\mathbf{P}(\mathbf{a}_1 - \mathbf{a}_2) = \mathbf{u}$ for some *non-zero* $\mathbf{u}$. Moreover, as $\mathbf{a}_1$ and $\mathbf{a}_2$ are in $(\frac{q}{p})\mathbb{Z}_q^m$, the difference $(\mathbf{a}_1 - \mathbf{a}_2)$ and therefore $\mathbf{u}$ are in $(\frac{q}{p})\mathbb{Z}_q^m$. Now, note that $F_{\mathrm{LWE}}(\phi_1(\mathbf{k}), w) = \lfloor \mathbf{P} \cdot \mathbf{k} + \mathbf{P} \cdot \mathbf{a}_1 \rfloor_p = \lfloor \mathbf{P} \cdot \mathbf{k} + \mathbf{P} \cdot \mathbf{a}_2 + \mathbf{u} \rfloor_p$. Applying Lemma 2.1, this in turn equals $\lfloor \mathbf{P} \cdot \mathbf{k} + \mathbf{P} \cdot \mathbf{a}_2 \rfloor_p + \lfloor \mathbf{u} \rfloor_p = F_{\mathrm{LWE}}(\phi_2(\mathbf{k}), w) + \lfloor \mathbf{u} \rfloor_p$. As $\mathbf{u} \in (\frac{q}{p})\mathbb{Z}_q^m$ and is non-zero, $\lfloor \mathbf{u} \rfloor_p$ is also non-zero in $\mathbb{Z}_p^m$ concluding the proof of the lemma. ∎

Consider a collision-resistant hash function $H\colon \{0,1\}^\ell \times \mathbb{Z}_q^m \to \{0,1\}^{\ell-1}$ and the fingerprint $w = 0^\ell$. We define $H_{\mathrm{com}}^{(\Phi_{\mathrm{lin}^*})}\colon \{0,1\}^\ell \times \mathbb{Z}_q^m \to \{0,1\}^\ell$ as $H_{\mathrm{com}}^{(\Phi_{\mathrm{lin}^*})}(x, y) = 1\|H(x, y)$ and note that it is a compatible hash function. Applying Lemmas 3.2–3.4 and Theorem 3.1 to the BC framework, Theorem 2.9 implies the following result.

Theorem 3.5. *Under the LWE assumption and the collision-resistance of the hash function H, the function $F_{\mathrm{rka\text{-}lin}}\colon \mathbb{Z}_q^m \times \{0,1\}^\ell \to \mathbb{Z}_p^m$ defined as:*

$$F_{\mathrm{rka\text{-}lin}}(\mathbf{k}, x) = F_{\mathrm{LWE}}\left(\mathbf{k}, H_{\mathrm{com}}^{(\Phi_{\mathrm{lin}^*})}\left(x, F_{\mathrm{LWE}}\left(\mathbf{k}, 0^\ell\right)\right)\right)$$

is an RKA-secure PRF with respect to Φ_{lin^*}.

3.2   RKA-secure PRFs for an Affine Class Φ_{aff}

In addition to the LWE-based almost key homomorphic PRF, Boneh *et al.* [14] also constructed a "fully" homomorphic PRF under the DLIN assumption over groups equipped with a multilinear map.

The PRF F_{DLIN}. For parameters m and $\ell \in \mathbb{N}$, let $\vec{\mathbb{G}} = (\mathbb{G}_1, \ldots, \mathbb{G}_\ell)$ be a sequence of groups equipped with a graded ℓ-multilinear map $\{\hat{e}_i\}_{i\in[\ell-1]}$. The public parameters comprise $pp = (g^{\mathbf{A}_0}, g^{\mathbf{A}_1})$, where $\mathbf{A}_0, \mathbf{A}_1 \leftarrow \mathbf{Rk}_\ell\left(\mathbb{Z}_p^{\ell\times\ell}\right)$. The PRF key $\mathbf{K}$ is a matrix in $\mathbb{Z}_p^{m\times\ell}$. Define $F_{\mathrm{DLIN}}\colon \mathbb{Z}_p^{m\times\ell} \times \{0,1\}^\ell \to (\mathbb{G}_\ell)^{m\times\ell}$ as follows:

$$F_{\mathrm{DLIN}}(\mathbf{K}, x) = (g_\ell)^{\mathbf{W}}, \text{ where } \mathbf{W} = \mathbf{K} \cdot \left(\prod_{i=1}^{\ell} \mathbf{A}_{x_i}\right). \tag{3.4}$$

Theorem 3.6 (cf. [14], paraphrased). *The function F_{DLIN} is pseudorandom under the DLIN assumption for suitable choices of parameters.*

As noted by Boneh *et al.*, the PRF can be evaluated at a point $x = x_1 \ldots x_\ell \in \{0,1\}^\ell$ given the the public parameters pp and secret key $\mathbf{k} \in \mathbb{Z}_p^\ell$ using the graded bilinear maps $\hat{e}_i : \mathbb{G}_1 \times \mathbb{G}_i \to \mathbb{G}_{i+1}$. The matrix multiplication is carried out one step at a time by nesting these bilinear maps as follows:

$$F_{\mathrm{DLIN}}(\mathbf{K}, x) = \hat{e}_{\ell-1}\left(g^{\mathbf{K}\mathbf{A}_{x_1}}, \hat{e}_{\ell-2}\left(g^{\mathbf{A}_{x_2}}, \ldots \hat{e}_2\left(g^{\mathbf{A}_{x_{\ell-2}}}, \hat{e}_1\left(g^{\mathbf{A}_{x_{\ell-1}}}, g^{\mathbf{A}_{x_\ell}}\right)\right)\right)\right),$$

where $g^{\mathbf{K}\mathbf{A}_{x_1}}$ is computed "in the exponent" given $\mathbf{K}$ and $g^{\mathbf{A}_{x_1}}$. A pairing $\hat{e}\left(g^{\mathbf{A}_0}, g^{\mathbf{A}_1}\right)$ of matrices given in the exponent is done by computing the component-wise dot products of rows of $\mathbf{A}_0$ with columns of $\mathbf{A}_1$ using the bilinear map $\hat{e}$.

Observe that this PRF is identical to the DLIN-based PRF in [14] except that the key $\mathbf{K}$ is now a matrix. This is required to define a meaningful affine class over the key space. The pseudorandomness extends to the case where $\mathbf{K}$ is a matrix by considering the rows of $\mathbf{K}$, $\mathbf{k}_1{}^\mathsf{T}, \ldots, \mathbf{k}_m{}^\mathsf{T}$ to be m independent keys of the original DLIN-based PRF. The key homomorphism also extends in a straightforward manner.

The affine class Φ_{aff}. With the above DLIN-based PRF, we can consider the following affine class of related-key deriving functions. We define

$$\Phi_{\mathsf{aff}} = \{\phi_{\mathbf{C},\mathbf{B}} : \mathbb{Z}_p^{m \times \ell} \to \mathbb{Z}_p^{m \times \ell} \mid \phi_{\mathbf{C},\mathbf{B}}(\mathbf{K}) = \mathbf{CK} + \mathbf{B}\}, \tag{3.5}$$

for matrices $\mathbf{C} \in \mathbb{Z}_p^{m \times m}$ and $\mathbf{B} \in \mathbb{Z}_p^{m \times \ell}$ constrained as follows: (a) the class Φ_{aff} is *claw-free*, and (b) $\mathbf{C}$ is a *full-rank* matrix.

As in Section 3.1, the key homomorphism of F_{DLIN} allows us to construct a transformer, denoted $\mathsf{T}_{\mathsf{aff}}^{f(\cdot)}$, in the following manner: $\mathsf{T}_{\mathsf{aff}}^{f(\cdot)}(\phi_{\mathbf{C},\mathbf{B}}, x)$ sets $f(x) = (g_\ell)^{\mathbf{F}}$ and computes $(g_\ell)^{\mathbf{CF}} \cdot F_{\mathrm{DLIN}}(\mathbf{B}, x)$. In other words, we left-multiply (in the exponent) the output of $f(\cdot)$ with entries from $\mathbf{C}$ and then use the homomorphism of F_{DLIN} to incorporate $\mathbf{B}$. We use the BC framework and show that for the class of related-key functions Φ_{aff}, the PRF and the transformer satisfy the malleability and uniformity properties.

Lemma 3.7 (Malleability). *For all $\mathbf{K} \in \mathbb{Z}_p^{m \times \ell}$, $\phi \in \Phi_{\mathsf{aff}}$, and $x \in \{0,1\}^\ell$, it holds that*

$$\mathsf{T}_{\mathsf{aff}}^{f(\cdot)}(\phi, x) = F_{\mathrm{DLIN}}(\phi(\mathbf{k}), x). \tag{3.6}$$

Proof. The proof follows from elementary algebra in the exponent. Let $\phi = \phi_{\mathbf{C},\mathbf{B}}$ for arbitrary $\mathbf{C}$ and $\mathbf{B}$. For a key $\mathbf{K}$ and input x, let $\mathbf{W}$ be the matrix in equation (3.4). By definition, $\mathsf{T}_{\mathsf{aff}}^{f(\cdot)}(\phi, x) = (g_\ell)^{\mathbf{C} \cdot \mathbf{W}} \cdot F_{\mathrm{DLIN}}(\mathbf{B}, x) = F_{\mathrm{DLIN}}(\mathbf{CK} + \mathbf{B}, x)$ as required. The last equality follows from the key homomorphism of F_{DLIN}. ∎

The following lemma follows straightforwardly from the definition of $\mathsf{T}_{\mathsf{aff}}^{f(\cdot)}$.

Lemma 3.8 (Uniformity). *If $f : \{0,1\}^\ell \to (\mathbb{G}_\ell)^{m \times \ell}$ is a random function and $x_1, \ldots, x_Q \in \{0,1\}^\ell$ are distinct, for any functions $\phi_1, \ldots, \phi_Q \in \Phi_{\mathsf{aff}}$, the values $\mathsf{T}_{\mathsf{aff}}^{f(\cdot)}(\phi_i, x_i)$ are independently and uniformly distributed in $(\mathbb{G}_\ell)^{m \times \ell}$.*

Next, we show that any $w \in \{0,1\}^\ell$ is a key fingerprint for Φ_{lin^*}.

Lemma 3.9 (Fingerprint). *For any $w \in \{0,1\}^\ell$, for any $\mathbf{K} \in \mathbb{Z}_q^{m \times \ell}$, and for any two distinct $\phi_1, \phi_2 \in \Phi_{\mathsf{aff}}$, it holds that $F_{\mathrm{DLIN}}(\phi_1(\mathbf{K}), w) \neq F_{\mathrm{DLIN}}(\phi_2(\mathbf{K}), w)$.*

Proof. We use the fact that the family Φ_{aff} is claw-free. For any key $\mathbf{K}$, this implies that $\phi_1(\mathbf{K}) \neq \phi_2(\mathbf{K})$. For $i \in \{1, 2\}$, let $\mathbf{W}_i$ denote the matrix $\phi_i(\mathbf{K}) \cdot \left(\prod_{i=1}^\ell \mathbf{A}_{w_i}\right)$. The product of full-rank matrices $\mathbf{A}_{w_i}$ is full-rank and as $\phi_1(\mathbf{K}) \neq \phi_2(\mathbf{K})$, it follows that $\mathbf{W}_1 \neq \mathbf{W}_2$. As F_{DLIN} is defined as $(g_\ell)^{\mathbf{W}}$ for generator g_ℓ, it holds that if $\mathbf{W}_1 \neq \mathbf{W}_2$, then $(g_\ell)^{\mathbf{W}_1} \neq (g_\ell)^{\mathbf{W}_2}$ concluding the proof of the lemma. ∎

Consider a collision-resistant hash function $H: \{0,1\}^\ell \times (\mathbb{G}_\ell)^{m \times \ell} \to \{0,1\}^{\ell-1}$ and the fingerprint $w = 0^\ell$. We define $H_{\mathsf{com}}^{(\Phi_{\mathsf{aff}})}: \{0,1\}^\ell \times (\mathbb{G}_\ell)^{m \times \ell} \to \{0,1\}^\ell$ as $H_{\mathsf{com}}^{(\Phi_{\mathsf{aff}})}(x, y) = 1 \| H(x, y)$ and note that it is a compatible hash function. Applying Lemmas 3.7–3.9 and Theorem 3.6 to the BC framework, Theorem 2.9 implies the following result.

Theorem 3.10. *Under the DLIN assumption and the collision-resistance of the hash function H, the function $F_{\mathsf{rka\text{-}aff}}: \mathbb{Z}_p^{m \times \ell} \times \{0,1\}^\ell \to (\mathbb{G}_\ell)^{m \times \ell}$ defined as:*

$$F_{\mathsf{rka\text{-}aff}}(\mathbf{K}, x) = F_{\mathsf{DLIN}}\left(\mathbf{K}, H_{\mathsf{com}}^{(\Phi_{\mathsf{aff}})}\left(x, F_{\mathsf{DLIN}}\left(\mathbf{K}, 0^\ell\right)\right)\right)$$

is an RKA-secure PRF with respect to Φ_{aff}.

4   Unique-Input RKA-secure PRFs for an Affine Class

In this section, we construct RKA-secure PRFs from the LWE assumption for a slightly more restricted notion of RKA security, denoted unique-input RKA security. As explained in Section 1.3, we work directly with the pseudorandomness proof of F_{LWE} to show unique-input RKA security against a larger class of affine related-key functions rather than the restricted linear class Φ_{lin^*} from Section 3.1. To do this, we use the algebraic structure that suits the key homomorphism of F_{LWE} to overcome the restrictions of Φ_{lin^*} required in order to apply the Bellare-Cash framework. We prove unique-input RKA security for the affine class $\Phi_{\mathsf{ln\text{-}aff}} = \{\phi_{\mathbf{C},\mathbf{B}} : \phi_{\mathbf{C},\mathbf{B}}(\mathbf{K}) = \mathbf{CK} + \mathbf{B}\}$, where $\mathbf{C}$ is a full rank matrix in $[-c, c]^{m \times m}$ for a small constant c, and $\mathbf{B}$ is an arbitrary matrix in $\mathbb{Z}_q^{m \times m}$.

We consider the PRF F_{LWE} where the key $\mathbf{k}$, originally a vector, is replaced by a matrix $\mathbf{K}$ in order to obtain the algebraic structure required for $\Phi_{\mathsf{ln\text{-}aff}}$. Recollect the definition of F_{LWE} from Equation (3.1). For parameters $m, p, q \in \mathbb{N}$ such that $p \mid q$, the public parameters of the PRF are binary matrices $\mathbf{A}_0, \mathbf{A}_1 \in \mathbb{Z}_p^{m \times m}$. The key is now a *matrix* $\mathbf{K} \in \mathbb{Z}_q^{m \times m}$, and the PRF $F_{\mathsf{LWE}}: \mathbb{Z}_q^{m \times m} \times \{0,1\}^\ell \to \mathbb{Z}_p^{m \times m}$ is defined as follows:

$$F_{\mathsf{LWE}}(\mathbf{K}, x) = \left\lfloor \mathbf{K} \cdot \prod_{i=1}^{\ell} \mathbf{A}_{x_i} \right\rceil_p . \tag{4.1}$$

Recollect the bound B for samples drawn from the LWE error distribution $\overline{\Psi}_\alpha$. In the rest of the section, we set the parameters of the system $q, p, m, c, B, \lambda, \ell > 0$ such that the quantity $(2m)^\ell cBp/q$ is negligible in the security parameter λ. This is along the lines of the parameters chosen in [14]. We state the following theorem for this choice of parameters:

Theorem 4.1. *Under the LWE assumption, the PRF F_{LWE} defined in Equation (4.1) is RKA-secure against unique-input adversaries for the class $\Phi_{\mathsf{ln\text{-}aff}}$.*

Proof of Theorem 4.1. In what follows, for a bit string x on ℓ bits, we use $x|_j$ to denote the bit string comprising bits j through ℓ of x. Let $x|_{\ell+1}$ denote the empty string ε^*.

Let $\mathcal{A}$ be a probabilistic polynomial time unique-input RKA adversary. We consider the following experiments interacting with $\mathcal{A}$.

Experiment G_j for $j \in [1, \ell + 1]$.

1. The challenger samples as public parameters full-rank matrices $\mathbf{A}_0, \mathbf{A}_1 \in \{0, 1\}^{m \times m} \subset \mathbb{Z}_q^{m \times m}$ which are sent to the adversary.
2. The challenger creates a lookup table L of pairs $(w, \mathbf{Z}) \in \{0, 1\}^{\ell - j + 1} \times \mathbb{Z}_q^{m \times m}$, and initializes L to contain only the pair $(\varepsilon^*, \mathbf{R})$ for some randomly chosen $\mathbf{R} \in \mathbb{Z}_q^{m \times m}$.
3. For $k \in [Q]$, the adversary (adaptively) sends input queries $\left(\phi_{\mathbf{C}, \mathbf{B}}^{(k)}, x^{(k)}\right) \in \Phi_{\mathsf{ln\text{-}aff}} \times \{0, 1\}^{\ell}$ to the challenger. For each input query, the challenger checks to see if there is a pair $\left(x^{(k)}|_j, \mathbf{Z}\right)$ in L for some $\mathbf{Z} \in \mathbb{Z}_q^{m \times m}$. If there is no such pair, then the challenger chooses a random $\mathbf{Y} \in \mathbb{Z}_q^{m \times m}$, adds the pair $\left(x^{(k)}|_j, \mathbf{Y}\right)$ to L, and sets $\mathbf{Z} = \mathbf{Y}$. The challenger returns $\mathbf{N} = \left\lfloor \mathbf{CZ} \prod_{i=1}^{j-1} \mathbf{A}_{x_i^{(k)}} + \mathbf{B} \prod_{i=1}^{\ell} \mathbf{A}_{x_i^{(k)}} \right\rceil_p$ to the adversary.
4. The adversary outputs a bit $b' \in \{0, 1\}$, which the experiment also outputs.

Experiment H_j for $j \in [1, \ell + 1]$.

1. The challenger samples as public parameters full-rank matrices $\mathbf{A}_0, \mathbf{A}_1 \in \{0, 1\}^{m \times m} \subset \mathbb{Z}_q^{m \times m}$ which are sent to the adversary.
2. The challenger creates a lookup table L of triples $(w, \mathbf{Y}, \mathbf{Z}) \in \{0, 1\}^{\ell - j + 1} \times \mathbb{Z}_q^{m \times m} \times \mathbb{Z}_q^{m \times m}$, and initializes L to contain only the triple $(\varepsilon^*, \mathbf{R}, \mathbf{\Delta})$ for some randomly chosen $\mathbf{R} \in \mathbb{Z}_q^{m \times m}$ and $\mathbf{\Delta} \leftarrow \overline{\Psi}_{\alpha}^{m \times m}$.
3. For $k \in [Q]$, the adversary (adaptively) sends input queries $\left(\phi_{\mathbf{C}, \mathbf{B}}^{(k)}, x^{(k)}\right) \in \Phi_{\mathsf{ln\text{-}aff}} \times \{0, 1\}^{\ell}$ to the challenger. For each input query, the challenger checks to see if there is a triple $\left(x^{(k)}|_{j-1}, \mathbf{Z}, \mathbf{\Delta}\right)$ in L for some $\mathbf{Z} \in \mathbb{Z}_q^m$ and $\mathbf{\Delta} \leftarrow \overline{\Psi}_{\alpha}^{m \times m}$. If there is no such triple, then the challenger chooses a random $\mathbf{Y} \in \mathbb{Z}_q^{m \times m}$ and random $\mathbf{V}_0, \mathbf{V}_1 \leftarrow \overline{\Psi}_{\alpha}^{m \times m}$, adds the triples $\left(0 \parallel \left(x^{(k)}|_j\right), \mathbf{Y}, \mathbf{V}_0\right)$ and $\left(1 \parallel \left(x^{(k)}|_j\right), \mathbf{Y}, \mathbf{V}_1\right)$ to L, and sets $\mathbf{Z} = \mathbf{Y}$ and $\mathbf{\Delta} = \mathbf{V}_{x_{j-1}^{(k)}}$ (i.e., $\mathbf{V}_0$ or $\mathbf{V}_1$ depending on the $j - 1^{\text{th}}$ bit of $x^{(k)}$). The challenger returns to the adversary the value:

$$\mathbf{N} = \left\lfloor \mathbf{C} \left(\mathbf{Z} \mathbf{A}_{x_{j-1}^{(k)}} + \mathbf{\Delta} \right) \cdot \prod_{i=1}^{j-2} \mathbf{A}_{x_i^{(k)}} + \mathbf{B} \cdot \prod_{i=1}^{\ell} \mathbf{A}_{x_i^{(k)}} \right\rceil_p .$$

4. The adversary outputs a bit $b' \in \{0, 1\}$, which the experiment also outputs.

Observe that $\mathsf{G}_{\ell+1}$ responds to the adversary's queries identically as in $\mathsf{Expt}_0^{\mathsf{prf\text{-}rka}}$. Hence, $\Pr\left[\mathsf{Expt}_0^{\mathsf{prf\text{-}rka}} = 1\right] = \Pr[\mathsf{G}_{\ell+1} = 1]$.

Lemma 4.2. *For all $j \in [2, \ell+1]$, it holds that $\left|\Pr[\mathsf{G}_j = 1] - \Pr[\mathsf{H}_j = 1]\right|$ is negligible.*

Proof. In Experiment H_j, let $\mathbf{M}_k = \mathbf{CZA}_{x_{j-1}^{(k)}} \cdot \prod_{i=1}^{j-2} \mathbf{A}_{x_i^{(k)}}$ and $\mathbf{W}_k = \mathbf{C\Delta} \cdot \prod_{i=1}^{j-2} \mathbf{A}_{x_i^{(k)}}$. Since the entries of $\mathbf{C}$ lie within $[-c, c]$, the entries of $\mathbf{\Delta}$ lie within

$[-B, B]$, and the entries of each of the $j - 2$ matrices $A_{x_i^{(k)}}$ lie within $\{0, 1\}$, the entries of $\mathbf{W}_k$ must lie within $[-cBm^{j-2}, cBm^{j-2}]$.[5] Since $\mathbf{A}_0$ and $\mathbf{A}_1$ are full rank, the product of these matrices is also full rank. Since $\mathbf{Z}$ is drawn uniformly at random from $\mathbb{Z}_q^{m \times m}$, the matrix $\mathbf{M}_k$ is distributed uniformly in $\mathbb{Z}_q^{m \times m}$. Thus, the probability that $\lfloor \mathbf{M}_k + \mathbf{W}_k \rfloor_p \neq \lfloor \mathbf{M}_k \rfloor_p$ is at most $m^2(cBm^{j-2})p/q$. By taking a union bound over all $x \in \{0, 1\}^\ell$, we have that the probability that there exists some input $x \in \{0, 1\}^\ell$ for which $\lfloor \mathbf{M}_k + \mathbf{W}_k \rfloor_p \neq \lfloor \mathbf{M}_k \rfloor_p$ is at most $(2m)^\ell cBp/q$. Conditioned on the above event not occurring, it holds that for all x, $\lfloor \mathbf{M}_k + \mathbf{W}_k \rfloor_p = \lfloor \mathbf{M}_k \rfloor_p$ which implies that G_j and H_j respond identically to adversary queries. Therefore $|\Pr[\mathsf{G}_j = 1] - \Pr[\mathsf{H}_j = 1]|$ is bounded by the probability of the above "bad" event, which is negligible for a suitable choice of parameters. ∎

We now state Lemmas 4.3 and 4.4, the proofs of which are deferred to the full version. Applying Lemmas 4.2–4.4 with suitable parameters yields Theorem 4.1.

Lemma 4.3. *Under the LWE assumption, for all $j \in [2, \ell + 1]$, it holds that the quantity $|\Pr[\mathsf{G}_{j-1} = 1] - \Pr[\mathsf{H}_j = 1]|$ is negligible.*

Lemma 4.4. $\Pr[\mathsf{G}_1 = 1] = \Pr\left[\mathsf{Expt}_1^{\mathsf{prf\text{-}rka}} = 1\right].$

5   Unique-Input RKA-secure PRFs for a Class of Polynomials

In this section, under the d-MDHE assumption, we show that F_{DLIN} is RKA-secure against unique-input adversaries with respect to the following class of bounded-degree polynomials. For positive integers ℓ, d and prime p we define

$$\Phi_{\mathsf{poly}(d)} = \left\{ \phi_{P(\cdot)} : \mathbb{Z}_p^{\ell \times \ell} \to \mathbb{Z}_p^{\ell \times \ell} \mid \phi_{P(\cdot)}(\mathbf{K}) = P(\mathbf{K}) \right\},$$

for polynomials P over $\mathbb{Z}_p^{\ell \times \ell}$ of degree at most d which have at least one coefficient matrix (excluding the constant coefficient matrix) which is full rank. In other words, if $P(\mathbf{K}) = \sum_{i=0}^d \mathbf{C}_i \cdot \mathbf{K}^i$ for matrices $\mathbf{C}_i \in \mathbb{Z}_p^{\ell \times \ell}$, then there exists a $j > 0$ such that $\mathbf{C}_j \in \mathbf{Rk}_\ell\left(\mathbb{Z}_p^{\ell \times \ell}\right)$. The proof of the following theorem is given in the full version.

Theorem 5.1. *Under the d-MDHE assumption, the PRF F_{DLIN} is RKA-secure against unique-input adversaries for the class $\Phi_{\mathsf{poly}(d)}$.*

6   Conclusions

We construct the first lattice-based PRFs secure against related-key attacks. We achieve RKA security under the standard (super-polynomial) LWE assumption for a restricted linear class of related-key functions and this result is comparable to the DDH-based RKA-secure PRF construction by Bellare and Cash [7]. Under the powerful multilinear map

---

[5] The fact that entries of $\boldsymbol{\Delta}$ lie within $[-B, B]$ holds only with overwhelming probability, but we will ignore this detail for ease of presentation, as it does not affect the final theorem.

abstraction [20], we construct RKA-secure PRFs against a large and natural class of affine related-key deriving functions with minimal restrictions. We believe this to be the most expressive affine class of transformations attainable under the Bellare-Cash framework. We also achieve the weaker notion of unique-input RKA security for an affine class of related-key deriving functions by considering the LWE-based key homomorphic PRF by Boneh *et al.* [14]. We show that by working with the proof of pseudorandomness and utilizing the algebraic structure of the PRF, we can overcome restrictions on the related-key class that are necessary to apply the Bellare-Cash framework. Finally, we show how, under the d-MDHE assumption in the presence of multilinear maps, we can achieve RKA security against unique-input adversaries for the class of degree-d polynomials. Our work on constructing new RKA-secure PRFs leads to several interesting open problems:

⋄ Can we construct LWE-based PRFs under the Bellare-Cash framework for a class less restrictive than Φ_{lin^*}? The only known LWE-based PRFs [5, 14] both require rounding and have "error terms" in proofs that have to be carefully dealt with. This will require a more careful application of the Bellare-Cash framework.
⋄ Can we construct unique-input RKA-secure PRFs from other LWE-based PRFs by Banerjee *et al.* [5] and (more recently) Banerjee and Peikert [4]?
⋄ Can we construct RKA-secure PRFs against unique-input adversaries for classes of polynomials from more standard assumptions such as LWE or DLIN?

References

[1] Shweta Agrawal, Xavier Boyen, Vinod Vaikuntanathan, Panagiotis Voulgaris, and Hoeteck Wee. Functional encryption for threshold functions (or fuzzy IBE) from lattices. In *PKC*, 2012.
[2] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *CRYPTO*, 2009.
[3] Benny Applebaum, Danny Harnik, and Yuval Ishai. Semantic security under related-key attacks and applications. In *ICS*, 2011.
[4] Abhishek Banerjee and Chris Peikert. New and improved key-homomorphic pseudorandom functions. Cryptology ePrint Archive, Report 2014/074, 2014. `http://eprint.iacr.org/`.
[5] Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. In *EUROCRYPT*, 2012.
[6] Alessandro Barenghi, Luca Breveglieri, Israel Koren, and David Naccache. Fault injection attacks on cryptographic devices: Theory, practice, and countermeasures. *Proceedings of the IEEE*, 100(11), 2012.
[7] Mihir Bellare and David Cash. Pseudorandom functions and permutations provably secure against related-key attacks. In *CRYPTO*, 2010.
[8] Mihir Bellare, David Cash, and Rachel Miller. Cryptography secure against related-key attacks and tampering. In *ASIACRYPT*, 2011.
[9] Mihir Bellare and Tadayoshi Kohno. A theoretical treatment of related-key attacks: RKA-PRPs, RKA-PRFs, and applications. In *EUROCRYPT*, 2003.
[10] Mihir Bellare, Sarah Meiklejohn, and Susan Thomson. Key-versatile signatures and applications: RKA, KDM and joint enc/sig. *To appear in EUROCRYPT*, 2014.
[11] Mihir Bellare, Kenneth G. Paterson, and Susan Thomson. RKA security beyond the linear barrier: IBE, encryption and signatures. In *ASIACRYPT*, 2012.

[12] Eli Biham and Adi Shamir. Differential fault analysis of secret key cryptosystems. In *CRYPTO*, 1997.

[13] Dan Boneh, Richard A. DeMillo, and Richard J. Lipton. On the importance of eliminating errors in cryptographic computations. *J. Cryptology*, 14(2), 2001.

[14] Dan Boneh, Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Key homomorphic PRFs and their applications. In *CRYPTO (1)*, 2013.

[15] Dan Boneh, Hart William Montgomery, and Ananth Raghunathan. Algebraic pseudorandom functions with improved efficiency from the augmented cascade. In *ACM CCS*, 2010.

[16] Joseph Bonneau and Ilya Mironov. Cache-collision timing attacks against aes. In *CHES*, 2006.

[17] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In *STOC*, 2013.

[18] Yevgeniy Dodis and Aleksandr Yampolskiy. A verifiable random function with short proofs and keys. In *PKC*, 2005.

[19] Niels Ferguson, John Kelsey, Stefan Lucks, Bruce Schneier, Michael Stay, David Wagner, and Doug Whiting. Improved cryptanalysis of rijndael. In *FSE*, 2000.

[20] Sanjam Garg, Craig Gentry, Shai Halevi, Amit Sahai, and Brent Waters. Attribute-based encryption for circuits from multilinear maps. In *CRYPTO (2)*, 2013.

[21] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 34(4), 1986.

[22] J. Alex Halderman, Seth D. Schoen, Nadia Heninger, William Clarkson, William Paul, Joseph A. Calandrino, Ariel J. Feldman, Jacob Appelbaum, and Edward W. Felten. Lest we remember: cold-boot attacks on encryption keys. *Commun. ACM*, 52(5), 2009.

[23] Goce Jakimoski and Yvo Desmedt. Related-key differential cryptanalysis of 192-bit key aes variants. In *Selected Areas in Cryptography*, 2003.

[24] Paul C. Kocher. Timing attacks on implementations of diffie-hellman, rsa, dss, and other systems. In *CRYPTO*, 1996.

[25] Allison B. Lewko and Brent Waters. Efficient pseudorandom functions from the decisional linear assumption and weaker variants. In *CCS*, 2009.

[26] Stefan Lucks. Ciphers secure against related-key attacks. In *FSE*, 2004.

[27] Thomas S. Messerges, Ezzy A. Dabbish, and Robert H. Sloan. Examining smart-card security under the threat of power analysis attacks. *IEEE Trans. Computers*, 51(5), 2002.

[28] Daniele Micciancio and Chris Peikert. Hardness of sis and lwe with small parameters. In *CRYPTO (1)*, 2013.

[29] Moni Naor and Omer Reingold. Number-theoretic constructions of efficient pseudo-random functions. In *FOCS*, 1997.

[30] Moni Naor and Gil Segev. Public-key cryptosystems resilient to key leakage. *SIAM J. Comput.*, 41(4), 2012.

[31] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In *STOC*. ACM, 2009.

[32] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, 2005.

[33] Thomas Ristenpart, Eran Tromer, Hovav Shacham, and Stefan Savage. Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. In *CCS*, 2009.

[34] Hovav Shacham. A cramer-shoup encryption scheme from the linear assumption and from progressively weaker linear variants. *IACR Cryptology ePrint Archive*, 2007.

[35] Hoeteck Wee. Public key encryption against related key attacks. In *PKC*, 2012.

[36] Wentao Zhang, Lei Zhang, Wenling Wu, and Dengguo Feng. Related-key differential-linear attacks on reduced aes-192. In *INDOCRYPT*, 2007.