

Related-Key Secure Pseudorandom Functions: The Case of Additive Attacks

Benny Applebaum* Eyal Widder*

August 30, 2015

Abstract

In a related-key attack (RKA) an adversary attempts to break a cryptographic primitive by invoking the primitive with several secret keys which satisfy some known relation. The task of constructing provably RKA secure PRFs (for non-trivial relations) under a standard assumption has turned to be challenging. Currently, the only known provably-secure construction is due to Bellare and Cash [7]. This important feasibility result is restricted, however, to linear relations over relatively complicated groups (e.g., \mathbb{Z}_q^* where q is a large prime) that arise from the algebraic structure of the underlying cryptographic assumption (DDH/DLIN). In contrast, applications typically require RKA-security with respect to simple additive relations such as XOR or addition modulo a power-of-two.

In this paper, we partially fill this gap by showing that it is possible to deal with simple additive relations at the expense of restricting the model of the attack. We introduce several natural relaxations of RKA-security, study the relations between these notions, and describe efficient constructions either under lattice assumptions or under general assumptions. Our results enrich the landscape of RKA security and suggest useful trade-offs between the attack model and the family of possible relations.

*School of Electrical Engineering, Tel-Aviv University, {benny.applebaum,wieyal}@gmail.com. Supported by ISF grant 1155/11, Israel Ministry of Science and Technology (grant 3-9094), and GIF grant 1152/2011.

1 Introduction

In a related-key attack (RKA) an adversary attempts to break a cryptographic primitive by invoking the primitive with several secret keys which satisfy some known relation. For example, the adversary may query a pseudorandom function under a tuple of keys (k_1, \dots, k_t) whose differences $\Delta_i = k_i - k_1 \pmod{2^n}$ are known, or even chosen by the adversary during the attack. Apart from being theoretically interesting, the study of RKAs is motivated by several scenarios.

- (Tampering attacks) In some cases, an adversary can mount a related-key attack by modifying the bits of the secret key (e.g., by tampering the hardware) and observing the resulting behavior. (See [21] and references therein.)
- (RKA as a working hypothesis) In practice, security against RKA has become a de-facto standard for block-ciphers. (The design of Rijndael, for example, explicitly stated RKA-security as a goal [16].) Correspondingly, systems are often designed while implicitly relying on the RKA-security of the underlying block-cipher. (See [7] and references therein.)
- (RKA-security as a resource) Security against RKA allows *cheap and structured* re-keying mechanism — a feature which turns to be extremely useful for higher-level applications (cf. [34, 4, 3]).

RKAs were originally considered by Biham [11] and Knudsen [28] in the early 1990's, and since then have become commonly used in the cryptanalysis of symmetric cryptography [12, 13, 20, 27, 35]. Motivated by this state of affairs, Bellare and Kohno [9] initiated a theoretical study of RKA security for block ciphers, theoretically modeled by pseudorandom functions (PRFs) and pseudorandom permutations (PRPs) [23]. They defined RKA security with respect to a class of related-key-deriving (RKD) functions Φ which specify the key-relations available to the adversary, and considered an active (and adaptive) adversary who can choose the relation from Φ during the attack. The notion of RKA security was further extended to many other cryptographic primitives, and several constructions were introduced (cf. [22, 4, 24, 8, 10, 3]).

The task of constructing provably RKA secure PRFs (for non-trivial relations) under a standard assumption has turned to be challenging. For a while, the only positive results were based on ideal models or non-standard assumptions [9, 29, 22]. The situation has changed with the beautiful work of Bellare and Cash [7] who showed that RKA secure block ciphers can be based on standard cryptographic assumptions (e.g., the hardness of the DDH/DLIN problem). This important feasibility result is restricted, however, to linear relations over relatively complicated groups (e.g., \mathbb{Z}_q^* where q is a large prime) that arise from the algebraic structure of the underlying cryptographic assumption (DDH/DLIN).

From an application point of view, such relations are somewhat unnatural. It is hard to imagine, for example, a hardware-tampering attack which manipulates the key in a way that corresponds to multiplication modulo a prime. Similarly, in practice, a cipher's key is typically "tweaked" by XOR-ing it with some public value or by treating it as a counter and increasing it by 1. Indeed, Bellare and Kohno [9] suggested XOR and addition modulo 2^n as the canonical examples of useful relations. The existence of RKA secure PRFs under such simple relations has been left open by [7] and its follow-ups [2, 1]. Interestingly, the situation is completely different for randomized encryption, for which RKA-security under additive relations can be achieved in a relatively simple way [4].

1.1 Our Contribution

In this paper, we partially fill this gap by showing that it is possible to deal with simple additive relations at the expense of relaxing the model of the attack. We introduce various relaxations of RKA-security by putting different natural restrictions on the set of legal RKA queries (Δ, x) . This include “passive” adversaries which are restricted to query $F_{k+\Delta}(x)$ with respect to random points x ’s or random shifts Δ ’s, non-adaptive adversaries which prepare their queries ahead of time and “bounded” adversaries which can use only a bounded number of different keys. We study the relations between these notions and present generic transformations between them. We also describe efficient constructions either for modular addition (over any large modulus) or, in the case of bounded adversaries, for any linear relation (including XOR). The security of these constructions is based on lattice assumptions or on general one-wayness assumptions. (A detailed account of our results appears below; see also Figure 1 for a summary.)

In a sense, our study takes the one-dimensional RKA-security game (which is solely parameterized by the RKD family Φ) and turns it into a multidimensional game. This enriches the landscape of RKA security, and provides a useful trade-off between the attack model and the family of possible relations.

1.1.1 RKA Secure Weak PRF

A weak PRF (wPRF) [30] is a relaxation of standard PRF which remains indistinguishable from a truly random function as long as it is being queried on random points. We construct a wPRF which offers RKA security with respect to mod- p addition over any (sufficiently large) integer modulus (including the case of $p = 2^n$). Our construction relies on the Learning with Rounding (LWR) assumption, introduced by Banerjee et al. [5], whose security can be based on the worst-case hardness of lattice problems. The construction and its proof are quite simple and efficient.

In a high-level, the LWR-based wPRF $F_s(x)$ computes the inner-product of the key $s \in \mathbb{Z}_q^n$ and the point $x \in \mathbb{Z}_q^n$ and outputs the result rounded to some integer grid (i.e., q is divided to r equal intervals and the output is rounded to the starting point of the corresponding interval). This function has an almost linear form which provides an “approximate” key-homomorphism, namely, $F_s(x) + F_{s'}(x)$ is close to $F_{s+s'}(x)$. This property, which was also used in [15], is extremely useful for proving RKA-security. Indeed, if we had an exact homomorphism we could emulate the value of $F_{s+\Delta}(x)$ given $F_s(x)$. A natural way to turn the approximate key-homomorphism into an exact homomorphism is to further round the result according to a grid with larger intervals. This gives an efficient mapping h from $F_s(x)$ and Δ to $F'_{s+\Delta}(x)$ where F' is the LWR-based wPRF parameterized with different “rounding resolution”. The mapping h is almost always correct with respect to a random key and a random point. However, in our context the adversary can shift the key *arbitrarily*, and when the modulus p is composite (e.g., $p = 2^k$ and rounding is applied with respect to multiples of $2^{k'}$) there are shift vectors Δ for which h fails miserably. Fortunately, we note that such a rounding error happens only when the output of h is close to the end of an interval, hence we can detect a potential failure. Using sufficiently large intervals we can make sure that when h is applied to a truly random function the result is unlikely to fall next to an end of an interval. Hence, the failure of h (which prevents us from emulating an RKA oracle) allows us to directly distinguish the wPRF from a truly random function.

We further demonstrate the usefulness of our LWR-based RKA-wPRF by constructing a simple and efficient message-authentication code (MAC). Our construction follows the outline suggested

in [17] which relies on a key-homomorphic wPRF. While our wPRF does not support a full key-homomorphism in the strict sense of [17], we show how to adopt the RKA-security proof to prove the security of the resulting MAC. (This result appears in Section A.)

1.1.2 Passive-RKA secure PRF

Passive-RKA (pRKA) is a weak form of RKA which was previously considered in the cryptanalytic literature (see discussion in [14]) and in the context of randomized encryption [4]. In such attacks the key-relations are chosen randomly and are not controlled by the adversary. For example, in the case of linear relations, an adversary who asks for an RKA query on a point x will get as a result the value of $(\Delta, F_{k+\Delta}(x))$ where Δ is a random group element and k is the target key. (This is essentially dual to the case of RKA secure wPRF where the point x is chosen at random and the shift Δ is chosen by the adversary.) We formalize the notion of pRKA security for PRFs, and show how to transform an RKA wPRF into a pRKA secure pseudorandom function or pseudorandom permutations (PRP). Combining this with our LWR-based construction we derive a PRF which is pRKA-secure under modular addition.

1.1.3 Non-Adaptive RKA secure PRF

A natural strengthening of a passive adversary is a non-adaptive adversary which prepares its (arbitrary) queries ahead of time. The resulting notion of *non-adaptive* RKA security (na-RKA) is therefore an intermediate notion between full-RKA security and pRKA security. It is not hard to imagine scenarios where na-RKA security may suffice (e.g., when a PRF is employed as a building block in a larger system whose design defines a set of RKA queries for which the underlying PRF should be secure). One may further consider a weaker notion of security in which the non-adaptive adversaries never queries the same point x under different keys. Observe that this notion of *distinct* non-adaptive RKA security (dna-RKA) is still strictly stronger from pRKA security.

We show that, in some circumstances, one can upgrade an RKA-wPRF into distinct non-adaptive RKA PRF (dna-RKA), which in turn can be upgraded to non-adaptive RKA security. In more detail, our first transformation maps an RKA-wPRF and a standard PRF into a new PRF. We prove that if the resulting PRF fails to achieve dna-RKA-security, then a public-key encryption can be established (whose correctness holds for infinitely-many input lengths).¹ Thus, in “Minicrypt” — the hypothetical world in which private-key cryptography exists but public-key cryptography does not exist — the existence of RKA-wPRF implies the existence of dna-RKA-PRFs. (This is similar in spirit to the results of [31, 32, 6].) Our transformation is generic and works for any class of RKD functions Φ .

Next, we show that if the key is a vector of field elements (as in the case of the XOR relation), it is possible to transform a dna-RKA secure PRF into a na-RKA PRF. The transformation is highly efficient and it incurs only a minor loss in the security. The question of upgrading non-adaptive RKA security to standard PRF security (for linear relations) is left open for further research.

1.1.4 Bounded-RKA secure PRF

Finally, we consider the notion of *bounded-RKA* (bRKA) where there exists an a-priori bound t on the number of different related keys that the adversary can generate. (Each related key k_i can be

¹In fact, we can even get an (infinitely often) Oblivious Transfer protocol. See Remark 5.3.

queried with respect to arbitrarily many different x 's.) We show that any PRF F can be immunized against t -bRKA attack, where $t(n)$ can be an arbitrary (a-priory fixed) polynomial. Furthermore, this holds for an arbitrary linear relation over any group (including the case of XOR).

The idea is to use a long key s of length $\Omega(tn)$ and derive an n -bit key k for F via an appropriate (public) key-derivation mechanism H . An adversary that uses the shifts $\vec{\Delta} = (\Delta_1, \dots, \Delta_t)$ gets an access to F keyed by the keys $k_1 = H(s + \Delta_1), \dots, k_t = H(s + \Delta_t)$, and, in addition, sees F keyed under the master key $k_0 = H(s)$. We show that if H is chosen from a family of $\Omega(nt^2)$ -wise independent hash functions, then, for any choice of $\vec{\Delta}$, the joint distribution of the keys (k_0, \dots, k_t) is statistically indistinguishable from uniform. Hence, the bRKA security of the scheme reduces to the security of F under t independent keys, which follows from standard PRF security.

Our usage of ℓ -wise independent hash function is inspired by the work of Faust et al. [19], who showed that such hash functions give rise to *non-malleable key derivation*. In this setting the adversary is restricted to *one-time* tampering with respect to a *large* family of tempering functions G . Namely, it is allowed to view the underlying primitive under the key $h(g(s))$ where $g \in G$. In contrast, we allow the adversary to access many keys but restrict it to linear relations. As a result, we have to overcome some non-trivial technicalities which do not appear in [19].

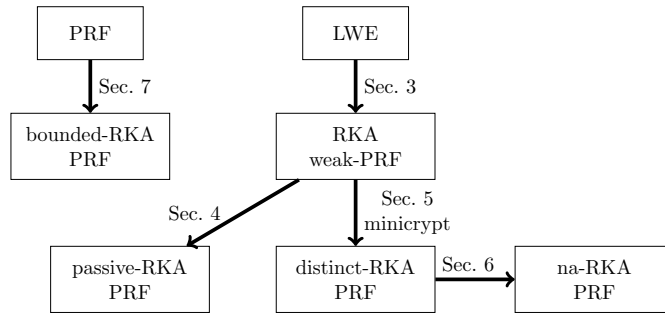


Figure 1: Results and Organization.

2 Preliminaries

Standard definitions. We say that a function $\mu : \mathbb{N} \rightarrow [0, 1]$ is *negligible* if for every positive polynomial p there exists $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$, $\mu(n) \leq \frac{1}{p(n)}$. We let $\text{neg}(n)$ denote some unspecified negligible function. A function $\mu : \mathbb{N} \rightarrow [0, 1]$ is *noticeable* if for some positive polynomial p there exists $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$, $\mu(n) \geq \frac{1}{p(n)}$.

The statistical distance between two random variables X and Y over a domain Ω , denoted $\Delta(X; Y)$, is defined by

$$\Delta(X; Y) = \frac{1}{2} \cdot \sum_{u \in \Omega} |\Pr[X = u] - \Pr[Y = u]|.$$

We will say that X and Y are ε -close if $\Delta(X; Y) \leq \varepsilon$. Let $X = \{X_n\}_{n \in \mathbb{N}}$ and $Y = \{Y_n\}_{n \in \mathbb{N}}$ be sequences of probability distributions. Then X and Y are said to be *statistically indistinguishable* if $\Delta(X_n; Y_n) = \text{neg}(n)$, and *computationally indistinguishable* if for every PPT distinguisher algorithm D , the distinguishing advantage $|\Pr[D(1^n, X_n) = 1] - \Pr[D(1^n, Y_n) = 1]|$ is negligible in n .

Definition 2.1 (Public-Key Encryption). *A public-key encryption scheme consists of three PPT algorithms: A key-generation algorithm KG which, given a security parameter 1^n outputs a public-key/secret-key pair (pk, sk) , an encryption algorithm E which maps a message m and a public-key pk to a ciphertext c , and a decryption algorithm D which maps a ciphertext c and secret-key sk to a plaintext m . We require standard semantic security and (potentially weak form of) correctness as follows:*

- ε -Correctness: For a random (single-bit) message $b \xleftarrow{R} \{0, 1\}$ and random public-key/secret-key pair $(\text{pk}, \text{sk}) \xleftarrow{R} \text{KG}(1^n)$

$$\Pr[\text{D}_{\text{sk}}(\text{E}_{\text{sk}}(b)) = b] \geq (1 + \varepsilon(n))/2,$$

for all n 's.

- Security: For every efficient adversary (eavesdropper) E and for all n

$$\left| \Pr_{(\text{pk}, \text{sk}) \xleftarrow{R} \text{KG}(1^n)} [E(\text{pk}, \text{E}_{\text{pk}}(0)) = 1] - \Pr_{(\text{pk}, \text{sk}) \xleftarrow{R} \text{KG}(1^n)} [E(\text{pk}, \text{E}_{\text{pk}}(1)) = 1] \right| < \text{neg}(n)$$

By default, we require that for every polynomial $p(n)$, $\varepsilon > 1 - 1/p(n)$ for all sufficiently large n 's, (this corresponds to the standard notion of correctness). If the above holds for infinitely many n 's, then the protocol is called an an infinitely-often correct Public-Key Encryption (io-PKE, in short).

2.1 Pseudorandom Functions

Syntax. Let $X = \{X_n\}_{n \in \mathbb{N}}$, $Y = \{Y_n\}_{n \in \mathbb{N}}$, and $K = \{K_n\}_{n \in \mathbb{N}}$ be a sequence of finite sets. A function ensemble \mathcal{F} from X to Y , indexed by keys from K , is a sequence of collections of functions $\{\mathcal{F}_n\}$ where for every n (security parameter) $\mathcal{F}_n = \{F_s : X_n \rightarrow Y_n\}_{s \in K_n}$. We always assume that the ensemble is efficiently computable and efficiently samplable. Namely, there exists an efficient sampler that given 1^n outputs a uniformly chosen element from K_n , and there exists an evaluation algorithm that given $s \in K_n$ and $x \in X_n$ outputs $y = F_s(x)$ in time polynomial in n . In the special case of *permutation* ensemble we assume that F_s forms a permutation over X_n and that there exists an efficient inversion algorithm which computes F_s^{-1} . Finally, we assume that the set of keys K_n is an additive Abelian group, and denote its group operation by $+$. By default, addition over K_n is assumed to be efficiently computable (in time $\text{poly}(n)$).

In the following, we define different security notions for PRFs. Since we mainly deal with linear relations, we will define RKA security explicitly for such relations. (All the following definitions can be extended to general related-key-deriving function families Φ as in [9].)

Weak PRF. An ensemble \mathcal{F} is a *weak PRF* (wPRF) if for every PPT adversary A

$$\left| \Pr_{s \xleftarrow{R} K_n} [A^{\text{Sam}F_s}(1^n) = 1] - \Pr[A^{\text{R}(X_n, Y_n)}(1^n) = 1] \right| < \text{neg}(n),$$

where the sampling oracle $\text{Sam}F_s$ ignores its input and outputs a pair (x, y) where $x \xleftarrow{R} X_n$ and $y = F_s(x)$, and the random oracle $\text{R}(X_n, Y_n)$ ignores its input and outputs a random pair $(x, y) \xleftarrow{R} (X_n \times Y_n)$.

RKA secure Weak PRF. An ensemble \mathcal{F} is an *RKA weak PRF* (RKA-wPRF) if for every PPT adversary A

$$\left| \Pr_{s \xleftarrow{R} K_n} [A^{\text{Sam}^+ F_s}(1^n) = 1] - \Pr[A^{\text{R}(X_n, Y_n)}(1^n) = 1] \right| < \text{neg}(n),$$

where the related-key oracle $\text{Sam}^+ F_s$ takes $\Delta \in K_n$ as an input, and outputs a pair (x, y) where $x \xleftarrow{R} X_n$ and $y = F_{s+\Delta}(x)$.

Strong (standard) PRFs. An ensemble \mathcal{F} is a *strong PRF* (or simply PRF) if for every PPT adversary A

$$\left| \Pr_{s \xleftarrow{R} K_n} [A^{F_s}(1^n) = 1] - \Pr[A^{\text{R}(X_n \rightarrow Y_n)}(1^n) = 1] \right| < \text{neg}(n),$$

where the (stateful) oracle $\text{R}(X_n \rightarrow Y_n)$ initializes a random function $\rho : X_n \rightarrow Y_n$, and given a query $x \in X_n$ outputs the value $\rho(x)$. Without loss of generality, we may restrict our attention to adversaries A which do not repeat the same query twice, and in this case we may replace $\text{R}(X_n \rightarrow Y_n)$ with the oracle $\text{R}(Y_n)$ which ignores its query and outputs a random value in Y_n .

RKA secure PRFs. An ensemble \mathcal{F} is a *RKA secure PRF* (RKA-PRF) if for every PPT adversary A

$$\left| \Pr_{s \xleftarrow{R} K_n} [A^{F_s^+}(1^n) = 1] - \Pr[A^{\text{R}^+(X_n \rightarrow Y_n)}(1^n) = 1] \right| < \text{neg}(n), \quad (1)$$

where the oracles are defined as follows. The oracle F_s^+ takes $x \in X_n$ and $\Delta \in K_n$ as inputs, and outputs the value $y = F_{s+\Delta}(x)$. The stateful oracle $\text{R}^+(X_n \rightarrow Y_n)$ initializes for each $s \in K_n$ a random function $\rho_s : X_n \rightarrow Y_n$, and, given a query $(x, \Delta) \in X_n \times K_n$ responds with $y = \rho_{s+\Delta}(x)$. Again, we may assume, without loss of generality, that a query (x, Δ) is never repeated twice, and in this case the oracle $\text{R}^+(X_n \rightarrow Y_n)$ can be replaced with the oracle $\text{R}(Y_n)$ which ignores its query and outputs a random value in Y_n .

Relaxations (bounded, passive, non-adaptive and distinct non-adaptive RKA). We say that a PRF is *t(n)-bounded-RKA secure* (bRKA) if Eq. (1) holds for adversaries that use at most t different shift vectors (but may apply the same shift vector more than once with different inputs.) We say that a PRF is *passive-RKA secure* (pRKA) if (1) holds for adversaries that each of their shift queries Δ_i is either set to zero or is chosen uniformly at random from K_n as part of the adversary's *public-coins*. Equivalently, we may assume that the random shift values are given to the adversary by the oracle. We say that a PRF is *non-adaptively RKA secure* (naRKA) if Eq. (1) holds for adversaries that prepare their oracle queries in advance and without inspecting the outcome of previous queries. The notion of *distinct non-adaptively RKA security* (dna-RKA) corresponds to non-adaptive adversaries which never query the same point twice; Namely, there are no two queries of the form (x, Δ) and (x, Δ') .

Public parameters. In some cases, it is useful to consider collections whose index consists of a *public parameter* $p \in P_n$ in addition to a secret key $s \in K_n$. Namely, $\mathcal{F}_n = \{F_{s,p} : X_n \rightarrow Y_n\}_{s \in K_n, p \in P_n}$ where $P = \{P_n\}_{n \in \mathbb{N}}$ is the domain of the public parameter. In this case, we assume the existence of an efficient procedure which samples $p \stackrel{R}{\leftarrow} P_n$ in time $\text{poly}(n)$. When defining security, we always assume that the public parameter p is chosen once and for all, and is given to the adversary as an additional input which is not subject to related-key attacks (i.e., the adversary cannot shift p). In the case of non-adaptive RKAs, we assume that the queries are chosen before p is chosen, and so they are independent of p . Public parameters can be used to set-up a number-theoretic construction (e.g., for selecting some group parameters), and can be also used as “one-time short salt” for system initialization. We mention that the RKA-PRFs of [7] also employed public parameters both to set-up a DDH group and to choose a collision resistance hash function.

2.2 A weak-PRF based on Learning with Rounding

For any integer modulus $q \geq 2$ we let \mathbb{Z}_q denote the quotient ring of integers modulo q . We let U_q denote the uniform distribution of \mathbb{Z}_q and U_q^n denote the uniform distribution over \mathbb{Z}_q^n .

Learning with Errors and Learning with Rounding. For dimension n , modulus $q = q(n)$ and noise distribution $\chi = \chi(n)$ over \mathbb{Z}_q , the (*decisional*) *Learning With Errors* $\text{LWE}_{n,q,\chi}$ [33] assumption asserts that the random oracle which outputs uniform samples from $\mathbb{Z}_q^n \times \mathbb{Z}_q$ is computationally indistinguishable from the (stateful) LWE oracle T_s which is initialized with a random $s \stackrel{R}{\leftarrow} \mathbb{Z}_q$ and outputs samples $x \stackrel{R}{\leftarrow} \mathbb{Z}_q^n, y \stackrel{R}{\leftarrow} \langle s, x \rangle + \chi$. Namely, for every PPT A ,

$$\left| \Pr[A^{T_s} = 1] - \Pr[A^{U_q^n \times U_q} = 1] \right| \leq \text{neg}(n).$$

Following [5], we modify the LWE problem such that the noise distribution over \mathbb{Z}_q is replaced with a deterministic rounding operation. The idea is to partition the elements $0, \dots, q-1$ to p consecutive intervals, and map an element to the starting point of its interval. Formally, for parameters $q > p$, we define the following mapping from $x \in \mathbb{Z}_q$ to an element in \mathbb{Z}_p by

$$\lceil x \rceil_{q/p} := \lfloor \bar{x} \cdot p/q \rfloor \bmod p,$$

where $\bar{x} \in \mathbb{Z}$ is any integer congruent to $x \bmod q$. Our notation slightly deviates from [5] as we would like to keep q explicit as part of the operation.

Definition 2.2 (The LWR Function). *For integers $q > p$ and n , and a key $s \in \mathbb{Z}_q^n$ we let $F_s^{q,p} : \mathbb{Z}_q^n \rightarrow \mathbb{Z}_p$ be the mapping*

$$F_s^{q,p} : x \mapsto \lceil \langle s, x \rangle \rceil_{q/p},$$

where $\langle x, s \rangle$ denotes inner-product, i.e., $\sum_i x_i \cdot s_i \bmod q$. For integer functions $q(n) > p(n)$, we define the function ensemble $\mathcal{F}^{q,p}$ over the domain/key-space $\mathbb{Z}_{q(n)}^n$ and range set $\mathbb{Z}_{p(n)}$, which for every n , consists of the functions $\left\{ F_s^{q(n),p(n)} : s \in \mathbb{Z}_{q(n)}^n \right\}$.

For ease of notation, we typically omit the dependency in the security parameter and write $F_s^{q,p}$.

In [5] it is proven that under the LWE assumption, $\mathcal{F}^{q,p}$ is a weak PRF. In the following, we say that a noise distribution $\chi = \chi_n$ over \mathbb{Z} is $B(n)$ -bounded if $\Pr[|\chi_n| > B(n)] < \text{neg}(n)$.

Proposition 2.3 ([5]). *Assume that $\text{LWE}_{n,q,\chi}$ holds for some B -bounded noise distribution χ and $q(n) > p(n) \cdot B(n) \cdot n^{\omega(1)}$. Then the collection $\mathcal{F}^{q,p}$ is a weak-PRF.*

The proposition does not impose any restriction regarding the structure of p and q (except for their size) and, in particular, holds for the case where p and q are powers of two.

The following proposition shows that when a uniformly chosen element from \mathbb{Z}_q is projected down to \mathbb{Z}_p , either by taking modulo or by the rounding operation, the resulting element is almost uniformly distributed over \mathbb{Z}_p .

Proposition 2.4. *Let $p < q$ and let $\delta(p, q)$ be zero if p divides q , and $\frac{p}{2q}$ otherwise. Then, the random variable $U_q \bmod p$ (resp., $\lceil U_q \rceil_{q/p}$) is $\delta(p, q)$ -close to the uniform distribution over \mathbb{Z}_p .*

Proof. Observe that the modulo operation and the rounding operation are both “almost balanced” in the sense that each element in \mathbb{Z}_p has either $\lceil q/p \rceil$ preimages or $\lfloor q/p \rfloor$ preimages. Assume that we have r values of the first type and $p - r$ values of the second type. Then, the statistical distance between the two distributions is

$$\frac{1}{2} \left| r \cdot \frac{\lceil q/p \rceil - q/p}{q} + (p - r) \cdot \frac{q/p - \lfloor q/p \rfloor}{q} \right|,$$

which is zero when p divides q , and is upper-bounded by $\frac{p}{2q}$, otherwise. \square

3 RKA secure Weak-PRF

In this section we show that if the collection $\mathcal{F}^{q,r}$ is a weak PRF then, for a proper choice of parameters $q > r > p$, the collection $\mathcal{F}^{q,p}$ is a weak PRF which is secure under related-key attacks. The key idea is to map a random input/output pair (x, y) of $F_s^{q,r}$ and an RKA query Δ into a random input/output pair (x, y') of $F_{s+\Delta}^{q,p}$. The transformation is based on the following mapping.

The mapping h . For integers $q > r > p$ and n , we define the mapping $h : \mathbb{Z}_q^n \times \mathbb{Z}_r \times \mathbb{Z}_q^n \rightarrow \mathbb{Z}_p$ as follows:

$$h : (x, y, \Delta) \mapsto \lceil y + F_{\Delta}^{q,r}(x) \pmod{r} \rceil_{r/p}. \quad (2)$$

Before we prove our main theorem, we collect several useful facts about the mapping h . We begin by showing that when h is applied to a random pair $(x, y) \stackrel{R}{\leftarrow} \mathbb{Z}_q^n \times \mathbb{Z}_r$ then, regardless of Δ , the result is almost uniform over \mathbb{Z}_p and independent of x .

Claim 3.1. *For $(x, y) \stackrel{R}{\leftarrow} \mathbb{Z}_q^n \times \mathbb{Z}_r$ and every $\Delta \in \mathbb{Z}_q^n$ the pair $(x, h(x, y, \Delta))$ is $\delta(p, r)$ -close to $U_p^n \times U_p$, where $\delta(p, r)$ is zero if p divides r , and $\frac{p}{2r}$ otherwise.*

Proof. Since y is uniform and independent of x , the random variable $y + F_{\Delta}^{q,r}(x)$ is uniform (over \mathbb{Z}_r) and independent of x and so, by Proposition 2.4, for every fixed x the random variable $h(x, y, \Delta) = \lceil y + F_{\Delta}^{q,r}(x) \pmod{r} \rceil_{r/p}$ is $\delta(p, r)$ -close to uniform over \mathbb{Z}_p . The claim follows. \square

Next we would like to show that when (x, y) is an input/output pair of $F_s^{q,r}$ then, whp, the output $h(x, y, \Delta)$ equals to $F_{s+\Delta}^{q,p}(x)$. Unfortunately, this claim does not hold for a non-prime modulus q . (See the discussion in the introduction.) Instead, we define an (efficiently recognizable) event that guarantees that h succeeds. Specifically, we will need the following claim.

Claim 3.2. *Let $p < r < q$. For every $s, \Delta, x \in \mathbb{Z}_q^n$, if the difference between $(F_s^{q,r}(x) + F_\Delta^{q,r}(x))$ to any multiple of r/p is at least 2, then $h(x, F_s^{q,r}(x), \Delta) = F_{s+\Delta}^{q,r}(x)$.*

Proof. First, we claim that the difference between $F_s^{q,r}(x) + F_\Delta^{q,r}(x)$ and $F_{s+\Delta}^{q,r}(x)$ is at most 2. Indeed,

$$\begin{aligned} F_s^{q,r}(x) + F_\Delta^{q,r}(x) &= \left\lfloor \langle x, s \rangle \cdot \frac{r}{q} \right\rfloor + \left\lfloor \langle x, \Delta \rangle \cdot \frac{r}{q} \right\rfloor \\ &= \langle x, s \rangle \cdot \frac{r}{q} + \langle x, \Delta \rangle \cdot \frac{r}{q} + e_1 + e_2 \\ &= \left\lfloor \langle x, s + \Delta \rangle \cdot \frac{r}{q} \right\rfloor + e_1 + e_2 + e_3 \\ &= F_{s+\Delta}^{q,r}(x) + e \end{aligned}$$

where the “error” terms $e_1, e_2 \in [-1, 0]$ and $e_3 \in [0, 1]$ compensate the omission/addition of the floor operation, and $e = e_1 + e_2 + e_3 \in [\pm 2]$ denotes the accumulated error.

Since $F_s^{q,r}(x) + F_\Delta^{q,r}(x)$ is 2-close to $F_{s+\Delta}^{q,r}(x)$, it follows, by our hypothesis, that both elements reside in the same r/p -interval. We can therefore write:

$$h(x, F_s^{q,r}(x), \Delta) = \lceil F_s^{q,r}(x) + F_\Delta^{q,r}(x) \rceil_{r/p} = \lceil F_{s+\Delta}^{q,r}(x) \rceil_{r/p} = F_{s+\Delta}^{q,p}(x)$$

and the claim follows. \square

We can now prove the main theorem of this section.

Theorem 3.3. *Let $p < r < q$ such that $r/q < \text{neg}(n)$, $p/r < \text{neg}(n)$ and $1/p < \text{neg}(n)$. Then if the collection $\mathcal{F}^{q,r}$ is a weak PRF then the collection $\mathcal{F}^{q,p}$ is RKA secure weak PRF.*

Proof. Fix some p, r , and q that satisfy the hypothesis. Assume, towards a contradiction, that A is an adversary that breaks the security of function $F_s^{q,p}$ as an RKA wPRF. Namely,

$$\Pr[A^{\text{Sam}^+ F_s^{q,p}}(1^n) = 1] - \Pr[A^{\text{R}(\mathbb{Z}_q^n, \mathbb{Z}_p)}(1^n) = 1] > \varepsilon(n),$$

for some non-negligible ε . We derive a contradiction by constructing an adversary B which distinguishes the oracle $\text{Sam}^+ F_s^{q,r}$ from the oracle $\text{R}(\mathbb{Z}_q^n, \mathbb{Z}_r)$ with advantage $\varepsilon - \text{neg}(n)$. The algorithm $B^{\mathcal{O}}$ emulates A , if A makes a query Δ to its oracle, the algorithm B asks for a sample $(x, y) \stackrel{R}{\leftarrow} \mathcal{O}$, and continues as follows: If the difference between $y + F_\Delta^{q,r}(x)$ to some multiple of r/p is smaller or equal to 2, then B quits with the output 1; Otherwise, B answers A with $(x, h(x, y, \Delta))$, where h is the mapping defined in Eq. (2). At the end of the emulation, B halts with the same output as A .

Let us analyze the distinguishing advantage of B assuming that A makes $t = \text{poly}(n)$ queries. First, observe that when \mathcal{O} is the uniform oracle $\text{R}(\mathbb{Z}_q^n, \mathbb{Z}_r)$, the probability that B quits is $t \cdot (4r/p)/r = \text{neg}(n)$. Also, by Claim 3.1, conditioned on not quitting, the view of A in the emulation is within statistical distance of $t \cdot \delta(p, r) = \text{neg}(n)$ from the view of A when the oracle is $\text{R}(\mathbb{Z}_q^n, \mathbb{Z}_p)$. It, therefore, follows that

$$\Pr[B^{\text{R}(\mathbb{Z}_q^n, \mathbb{Z}_r)}(1^n) = 1] \leq \Pr[A^{\text{R}(\mathbb{Z}_q^n, \mathbb{Z}_p)}(1^n) = 1] + \text{neg}(n).$$

We move on to the case where the oracle \mathcal{O} is $\text{Sam}F_s^{q,r}$. By Claim 3.2, if $B^{\text{Sam}F_s^{q,r}}$ does not quit, the emulation is perfect. Letting $\delta = \Pr_s[B^{\text{Sam}F_s^{q,r}}(1^n) \text{ quits}]$, and recalling that when B quits it always outputs 1, we can write

$$\Pr_s[B^{\text{Sam}F_s^{q,r}}(1^n) = 1] = \delta + \Pr_s[B^{\text{Sam}F_s^{q,r}}(1^n) = 1 \mid \text{not quitting}] \cdot (1 - \delta) \geq \Pr_s[A^{\text{Sam}^+ F_s^{q,p}}(1^n) = 1].$$

Overall, it follows that

$$\Pr_s[B^{\text{Sam}F_s^{q,r}}(1^n) = 1] - \Pr_s[B^{\text{R}(\mathbb{Z}_q^n, \mathbb{Z}_r)}(1^n) = 1] > \varepsilon(n) - \text{neg}(n),$$

as required. \square

4 Passive-RKA secure PRF and PRP

In this section we show how to convert any RKA-wPRF \mathcal{F} (such as the one constructed in Section 3) into a strong PRF \mathcal{G} with Passive-RKA security. The idea is similar to the one proposed in [22] (See also [7, 8]). Essentially, \mathcal{F} is used as a key derivation mechanism for a standard PRF PRF. That is, our new PRF \mathcal{G} will be keyed by pairs (k_1, k_2) and $\mathcal{G}_{k_1, k_2}(x) = \text{PRF}_k(x)$ where the pseudorandom key k is taken to be $\mathcal{F}_{k_1}(k_2)$. A passive-RKA adversary that applies random shifts (Δ_1, Δ_2) to (k_1, k_2) will get an access to the function $\text{PRF}_{k'}$ keyed under $k' = \mathcal{F}_{k_1 + \Delta_1}(k_2 + \Delta_2)$. Due to the RKA-security of \mathcal{F} , the key k' will still be pseudorandom, and so, intuitively, $\text{PRF}_{k'}$ remains secure. Furthermore, if PRF is a *pseudorandom permutation* (PRP), the resulting construction yields a passive-RKA secure PRP. We move on to a formal definition of our construction.

Construction 4.1. Let $\text{PRF} = \{\text{PRF}_k\}$ be a function ensemble with key space K , domain X and range Y , and let $\mathcal{F} = \{F_{k_1}\}$ be a function ensemble with key space K_1 , domain K_2 , and range K . We define a new function ensemble \mathcal{G} with key space $K_1 \times K_2$ domain X and range Y via the mapping:

$$G_{(k_1, k_2)} : x \mapsto \text{PRF}_k(x), \quad \text{where } k = F_{k_1}(k_2).$$

Theorem 4.2. If PRF is a PRF and \mathcal{F} is an RKA secure wPRF, then \mathcal{G} is a passive-RKA secure PRF. Moreover, if PRF is a PRP then \mathcal{G} is a passive-RKA secure PRP.

Proof. The proof relies on a hybrid argument. We define four sampling oracles which support standard PRF queries and passive-RKA (pRKA) queries. Syntactically, a PRF query $x \in X$ should be answered with $y \in Y$, and pRKA query $x \in X$ should be answered with $(\Delta_1, \Delta_2, y) \in K_1 \times K_2 \times Y$. The oracles are defined as follows. (For simplicity we omit the dependencies on n .)

- $\mathcal{O}_{(k_1, k_2)}^0$ is the passive-RKA oracle for the original G construction:
 - pRKA query $x \in X$ is answered with $(\Delta_1, \Delta_2) \stackrel{R}{\leftarrow} K_1 \times K_2$ and $y = \text{PRF}_{k'}(x)$ where $k' = F_{k_1 + \Delta_1}(k_2 + \Delta_2)$.
 - PRF query $x \in X$ is answered with $y = \text{PRF}_k(x)$ where $k = F_{k_1}(k_2)$.
- \mathcal{O}_k^1 is a hybrid oracle:
 - pRKA query $x \in X$ is answered with $(\Delta_1, \Delta_2) \stackrel{R}{\leftarrow} K_1 \times K_2$ and $y = \text{PRF}_{k'}(x)$ where $k' \stackrel{R}{\leftarrow} K$ is a fresh random key.

- PRF query $x \in X$ is answered with $y = \text{PRF}_k(x)$.
- \mathcal{O}_k^2 is a hybrid oracle:
 - pRKA query $x \in X$ is answered with $(\Delta_1, \Delta_2) \xleftarrow{R} K_1 \times K_2$ and $y \xleftarrow{R} Y$.
 - PRF query $x \in X$ is answered with $y = \text{PRF}_k(x)$.
- \mathcal{O}^3 is the random sampling oracle:
 - pRKA query $x \in X$ is answered with $(\Delta_1, \Delta_2) \xleftarrow{R} K_1 \times K_2$ and $y \xleftarrow{R} Y$.
 - New PRF query $x \in X$ is answered with $y \xleftarrow{R} Y$, old PRF query is answered consistently with the previous answers.

Our goal is to show that the oracles \mathcal{O}^0 and \mathcal{O}^3 are indistinguishable. We begin by showing that the oracle \mathcal{O}^0 is indistinguishable from the oracle \mathcal{O}^1 .

Claim 4.3. *For every PPT adversary A ,*

$$\left| \Pr_{k_1, k_2} [A^{\mathcal{O}_{k_1, k_2}^0}(1^n) = 1] - \Pr_k [A^{\mathcal{O}_k^1}(1^n) = 1] \right| \leq \varepsilon(n),$$

for some negligible function $\varepsilon(n)$.

Proof. Assume, towards contradiction, that A has a non-negligible distinguishing advantage $\varepsilon(n)$. We construct an efficient PPT $B^{\mathcal{O}}$ that breaks the RKA wPRF security of \mathcal{F} . Namely, B distinguishes between the random oracle $\text{R}(K_2, K)$ to the RKA wPRF oracle $\text{Sam}^+ F_{k_1}$ where $k_1 \xleftarrow{R} K_1$. The idea is to emulate A as follow:

1. B initially queries \mathcal{O} with $\Delta = \mathbf{0}$ and obtains a random input/output pair $(k_2, k) \in K_2 \times K$.
2. If A makes a PRF query $x \in X$, then B returns the value $y = \text{PRF}_k(x)$.
3. If A makes pRKA query $x \in X$, then B computes the answer (Δ_1, Δ_2, y) as follows. B chooses a random shift $\Delta_1 \xleftarrow{R} K_1$, uses the oracle to compute $(k'_2, k') = \mathcal{O}(\Delta_1)$ and sets $\Delta_2 = k'_2 - k_2$ and $y = \text{PRF}_{k'}(x)$.
4. At the end of the emulation B terminates with the same output as A .

It is not hard to verify that if $\mathcal{O} = \text{R}(K_2, K)$ then B perfectly emulates the oracle \mathcal{O}_k^1 . On the other hand, when \mathcal{O} is $\text{Sam}^+ F_{k_1}$ we obtain a perfect emulation of $\mathcal{O}_{(k_1, k_2)}^0$ for a randomly chosen $k_2 \xleftarrow{R} K_2$. Indeed, since $k = F_{k_1}(k_2)$, PRF queries are answered properly by $\text{PRF}_k(x)$. Similarly, the answer (Δ_1, Δ_2, y) to a pRKA query x is computed properly, since the shifts (Δ_1, Δ_2) are uniform over $K_1 \times K_2$ and $y = \text{PRF}_{k'}(x)$ for $k' = F_{k_1 + \Delta_1}(k_2 + \Delta_2)$, as required. It follows that B breaks the RKA wPRF security of \mathcal{F} in contraction to our assumption. \square

We proceed by showing that \mathcal{O}^1 is indistinguishable from the oracle \mathcal{O}^2 .

Claim 4.4. For every PPT adversary A ,

$$\left| \Pr_k[A^{\mathcal{O}_k^1}(1^n) = 1] - \Pr_k[A^{\mathcal{O}_k^2}(1^n) = 1] \right| \leq \varepsilon(n),$$

for some negligible function $\varepsilon(n)$.

Proof. The security is reduced to the security of the PRF via a standard hybrid argument. Assume, towards contradiction, that A has a non-negligible distinguishing advantage $\varepsilon(n)$. Let $t(n) = \text{poly}(n)$ denote an upper bound on the number of queries that A performs, and assume, wlog, that A never repeats the same PRF query twice. Fix some n

We construct an efficient PPT $B^{\mathcal{O}}(1^n)$ that breaks the security of PRF, where its oracle \mathcal{O} is either PRF_{k_1} or $R(Y_n)$.

1. B samples a random key $k \xleftarrow{R} K_n$.
2. B chooses a random hybrid index $i \in \{1, \dots, t(n)\}$, and emulates A as follows.
 - (a) For the first $i - 1$ queries B ignores its oracle and answers as in \mathcal{O}_k^2 :
 - If x is a pRKA query then B outputs $(\Delta_1, \Delta_2) \xleftarrow{R} K_1 \times K_2$ and $y \xleftarrow{R} Y$.
 - If x is a PRF query then B answers with $\text{PRF}_k(x)$.
 - (b) On the i -th query B answers as follows:
 - If x is a pRKA query then B outputs $(\Delta_1, \Delta_2) \xleftarrow{R} K_1 \times K_2$ and $y = \mathcal{O}(x)$.
 - If x is a PRF query then B answers with $\text{PRF}_k(x)$.
 - (c) For the remaining queries B ignores its oracle and answers as in \mathcal{O}_k^1 :
 - If x is a pRKA query then B outputs $(\Delta_1, \Delta_2) \xleftarrow{R} K_1 \times K_2$ and $y = \text{PRF}_{k'}(x)$ for a randomly sampled $k' \xleftarrow{R} K$.
 - If x is a PRF query then B answers with $\text{PRF}_k(x)$.
3. When A terminates B outputs the same output of A .

For $j \in \{1, \dots, t\}$, let

$$\alpha_j(n) = \Pr_{k_1 \xleftarrow{R} K_n} [B^{\text{PRF}_{k_1}}(1^n) = 1 | i = j] \quad \text{and} \quad \beta_j(n) = \Pr[B^{R(Y_n)}(1^n) = 1 | i = j].$$

It is not hard to verify that $\alpha_1 = \Pr_k[A^{\mathcal{O}_k^1}(1^n) = 1]$ and that $\beta_t = \Pr_k[A^{\mathcal{O}_k^2}(1^n) = 1]$. Also, by definition, for every j , $\beta_j(n) = \alpha_{j+1}(n)$. Therefore, the distinguishing advantage of B (for a randomly chosen i) is $1/t \sum_j \alpha_j(n) - \beta_j(n) = (\alpha_1(n) - \beta_t(n))/t$ and so B breaks the security of PRF with noticeable advantage of $\varepsilon(n)/t(n)$, and we derive a contradiction. \square

We proceed by showing that \mathcal{O}^2 is indistinguishable from the oracle \mathcal{O}^3 .

Claim 4.5. For every PPT adversary A ,

$$\left| \Pr_{k_1}[A^{\mathcal{O}_{k_1}^2}(1^n) = 1] - \Pr[A^{\mathcal{O}^3}(1^n) = 1] \right| \leq \varepsilon(n),$$

for some negligible function $\varepsilon(n)$.

Proof. Assume, towards contradiction, that A has a non-negligible distinguishing advantage $\varepsilon(n)$, and assume without loss of generality that A does not repeat the same query twice. We construct an efficient PPT $B^\mathcal{O}$ which distinguishes with advantage ε , between the oracle $\mathsf{R}(Y)$ to the oracle PRF_k for a randomly chosen $k \xleftarrow{R} K$.

1. B emulates A and answers its queries as follows:

- If x is a pRKA query then B outputs $(\Delta_1, \Delta_2) \xleftarrow{R} K_1 \times K_2$ and $y \xleftarrow{R} Y$.
- If x is a PRF query then B answers with a sample from the oracle $\mathcal{O}(x)$.

2. When A terminates B outputs the same output of A .

It is not hard to verify that when \mathcal{O} is PRF_k we perfectly emulate \mathcal{O}_k^2 and when $\mathcal{O} = \mathsf{R}(Y)$ we perfectly emulate \mathcal{O}^3 and so the claim follows. \square

By combining the three claims we complete the proof of the theorem. The proof naturally extends to the special case where PRF is a permutation since both claims still hold, and so, in this case, a pRKA secure PRP is obtained. \square

5 Constructing dna-RKA secure PRFs in Minicrypt

In this section, we show that in Minicrypt one can transform an RKA secure wPRF \mathcal{F} into a distinct non-adaptive RKA secure PRF. That is, we transform RKA wPRF \mathcal{F} (as provided in Section 3) and a standard PRF PRF into a function ensemble \mathcal{H} , and show that if \mathcal{H} is not a dna-RKA-PRF, then an (infinitely-often) public-key encryption (io-PKE) exists.

Construction 5.1. For function ensembles $\mathsf{PRF} : K \times X \rightarrow Y$ and $F : S \times Y \rightarrow Z$, define the ensemble H with key-space S , domain X , range Y and public parameter space K via the mapping

$$H_{s,k} : x \mapsto F_s(\mathsf{PRF}_k(x)), \quad \text{where } s \in S_n, x \in X_n, \text{ and } k \in K_n \text{ is a public parameter.}$$

Theorem 5.2. Assume that F is an RKA wPRF and PRF is a (standard) PRF. Then, either H is a dna-RKA-PRF or there exists an io-PKE.

Proof. Assume, towards a contradiction, that the construction H is not dna-RKA secure then there exists a distinct non-adaptive adversary that breaks H with non-negligible advantage $\varepsilon(n)$. Since the adversary is non-adaptive we can describe it by a pair of adversaries: a query generator Q and a distinguisher D . These adversaries share the same randomness r and operate as follows: At the first stage of the attack $Q(1^n, r)$ generates x -distinct RKA queries $(\Delta_1, x_1), \dots, (\Delta_t, x_t) \in (S_n \times X_n)$ for some polynomial $t = t(n)$; At the second phase, D uses r , the public parameter $k \xleftarrow{R} K_n$ to ε -distinguish between a random tuple $(u_i)_{i=1}^t \xleftarrow{R} \mathsf{R}(Z_n)^t$ to an RKA tuple $(H_{s+\Delta_i}^k(x_i))_{i=1}^t$ chosen with respect to a random secret-key $s \xleftarrow{R} S_n$. That is, by our assumption, we have

$$\Pr_{r,k \xleftarrow{R} K_n, s \xleftarrow{R} S_n} [D(1^n, r, k, (z_1, \dots, z_t)) = 1] - \Pr_{k \xleftarrow{R} K_n} [D(1^n, r, k, (u_1, \dots, u_t)) = 1] \geq \varepsilon(n),$$

where $(\Delta_i, x_i)_{i=1}^t = Q(1^n, r)$ and for every i , $z_i = H_{s+\Delta_i}^k(x_i) = F_{s+\Delta_i}(\mathsf{PRF}_k(x_i))$ and $u_i \xleftarrow{R} \mathsf{R}(Z_n)$.

We define a public-key encryption scheme as follows.

- (Key Generation): Set the secret key to be $\text{sk} = (r, k)$ where r is a fresh randomness for the adversary and $k \stackrel{R}{\leftarrow} K_n$. Generate the queries $(\Delta_i, x_i)_{i=1}^t = Q(1^n, r)$ and set the public-key pk to be $(\Delta_i, y_i = \text{PRF}_k(x_i))_{i=1}^t$.
- (Encryption) To encrypt the bit zero output t random values (u_1, \dots, u_t) sampled from Z_n ; To encrypt one output the vector $(H_{s+\Delta_i}^k(x_i))_{i=1}^t = (F_{s+\Delta_i}(y_i))_{i=1}^t$.
- (Decryption) The ciphertext $z = (z_1, \dots, z_t)$ is decrypted to the distinguisher's bit $a = D(1^n, r, k, z)$.

Correctness. It is not hard to see that on a random message $b \stackrel{R}{\leftarrow} \{0, 1\}$, decryption succeeds with probability $\frac{1}{2} + \varepsilon/2$. Indeed, letting $a = D_{\text{sk}}(\text{E}_{\text{pk}}(b))$ denote the outcome of the decryption algorithm (which is a random variable jointly distributed with b), we have

$$\begin{aligned}
\Pr[a = b] &= \Pr[b = 1] \cdot \Pr[a = 1|b = 1] + \Pr[b = 0] \cdot \Pr[a = 0|b = 0] \\
&= \frac{1}{2} \cdot (\Pr[a = 1|b = 1] + 1 - \Pr[a = 1|b = 0]) \\
&= \frac{1}{2} + \frac{1}{2} \cdot \left(\Pr_{s,k,r} [D(1^n, r, k, (H_{s+\Delta_i}^k(x_i))_{i=1}^t) = 1] - \Pr_{k,r,(u_i)_{i=1}^t} [D(1^n, r, k, (u_i)_{i=1}^t) = 1] \right) \\
&\geq \frac{1 + \varepsilon}{2}.
\end{aligned}$$

where the last equality follows by noting that when $b = 1$ the queries are answered according to $H_{s,k}^+$, and when $b = 0$ they are being answered according to $R(Z)$.

Security. Next we show that our PKE is semantically secure. Intuitively, an external adversary E who sees a public-key pk and a ciphertext $c = \text{E}_{\text{pk}}(b)$ just sees F_s evaluated on a pseudorandom points y_i (without knowing k), and therefore it should not be able to guess b due to the security of F_s as an RKA wPRF. We formalize this intuition below. Assume that E is an efficient adversary which distinguishes with advantage $\delta(n)$ between the distributions $(\text{pk}, \text{E}_{\text{pk}}(0))$ and $(\text{pk}, \text{E}_{\text{pk}}(1))$ where $\text{pk} \stackrel{R}{\leftarrow} \text{KG}(1^n)$.

Consider a modified version of the key-generation $\text{KG}'(1^n)$ in which the public-key is $\text{pk}' = (\Delta_i, y'_i)_{i=1}^t$ where y'_i are uniformly chosen from $R(Y)$ and the Δ_i 's are chosen as in the original scheme $(\Delta_i, x_i)_{i=1}^t = Q(1^n, r)$. We claim that the modified public-key is indistinguishable from the original public key. Indeed, one can implement the original key-generation algorithm given only an oracle to PRF_k , by switching this oracle to a random oracle $R(X \rightarrow Y)$ we get an indistinguishable key-generation algorithm. Since all the queries of Q are x -distinct we can further switch this oracle to $R(Y)$. The resulting public-key distribution corresponds exactly to the one generated by $\text{KG}'(1^n)$.

Since, $\text{pk} \stackrel{R}{\leftarrow} \text{KG}'(1^n)$ is indistinguishable from the modified key $\text{pk}' \stackrel{R}{\leftarrow} \text{KG}'(1^n)$, it follows that E distinguishes $(\text{pk}', \text{E}_{\text{pk}'}(0))$ from $(\text{pk}', \text{E}_{\text{pk}'}(1))$, i.e.,

$$\Pr_{\text{pk}' \stackrel{R}{\leftarrow} \text{KG}'(1^n)} [E(\text{pk}', \text{E}_{\text{pk}'}(1)) = 1] - \Pr_{\text{pk}' \stackrel{R}{\leftarrow} \text{KG}'(1^n)} [E(\text{pk}', \text{E}_{\text{pk}'}(0)) = 1] \geq \delta(n) - \text{neg}(n). \quad (3)$$

We can now combine E and Q into an adversary \mathcal{A} which violates the RKA security of F as a wPRF. That is, we define an oracle-aided adversary \mathcal{A}^O which distinguishes the RKA wPRF oracle

$\text{Sam}^+ F_s$ from the random oracle $R(Y, Z)$. The adversary \mathcal{A} uses $Q(1^n, r)$ to generate the queries $(\Delta_1, x_1), \dots, (\Delta_t, x_t)$, then, for every i , the adversary \mathcal{A} ignores the x_i 's and calls its oracle with Δ_i and obtains the value $(y_i, z_i) = \mathcal{O}(\Delta_i)$. At the end of the emulation, apply E to the transcript $(\Delta_i, y_i, z_i)_{i=1}^t$ and output the result. It is not hard to verify that if $\mathcal{O} = \text{Sam}^+ F_s$ then E receives a transcript which is distributed exactly as $(\text{pk}', E_{\text{pk}'}(1))$, and if $\mathcal{O} = R(Y, Z)$ then the transcript is distributed exactly as $(\text{pk}', E_{\text{pk}'}(0))$. It follows, by Eq. (3), that E breaks the RKA security of F as a wPRF, and we derive a contradiction.

Note that the above argument fails if Q is allowed to query the same point x twice, since in this case we cannot replace the value $y = \text{PRF}_k(x)$ with a random point, and so we cannot switch the PRF oracle PRF_k with the oracle $R(Y)$.

We conclude that the PKE is secure and that the decryption succeeds with probability $\frac{1}{2} + 1/\text{poly}(n)$ for infinitely many n 's. Such a PKE can be amplified into to standard (infinitely-often) PKE with negligible error probability via the use of error-correcting codes, e.g., by repetition. (See also [26] for a more general transformation). \square

Remark 5.3 (Oblivious Transfer). *Our proof of security shows that one can generate a “bad public key” that looks indistinguishable from the valid public key, but does not allow the generating party to distinguish between the encryption of 0 and the encryption of 1. It is shown in [18] that such an encryption scheme can be used to construct an Oblivious-Transfer protocol. Hence, if H fails to achieve dna-RKA-PRF security one can construct an (infinitely-often) Oblivious-Transfer protocol.*

6 From dna-RKA to na-RKA

In this section, we show how to transform a dna-RKA secure PRF \mathcal{F} into a non-adaptive RKA secure PRF. Let F be a dna-PRF where the keys, messages, and outputs are all n -long vectors over a field \mathbb{F} , e.g., the binary field \mathbb{F}_2 .² Let $m = \lceil \delta n \rceil$ for some small constant $\delta \in (0, 1)$. We define an ensemble H over the key space \mathbb{F}^{n+m} , message-space \mathbb{F}^{n-m} , and range \mathbb{F}^n . The ensemble will be parameterized by a public $(n+m) \times (n+m)$ invertible matrix $M \xleftarrow{R} \text{GL}(n+m, \mathbb{F})$ which will be parsed into a top $n \times (n+m)$ sub-matrix $M_1 \in \mathbb{F}^{n \times (n+m)}$ and a bottom $m \times (n+m)$ sub-matrix $M_2 \in \mathbb{F}^{m \times (n+m)}$. The ensemble H is defined via the mapping

$$H_{k,M} : x \mapsto F_{M_1 \cdot k}(x, M_2 \cdot k),$$

where $k \in \mathbb{F}^{n+m}$, $x \in \mathbb{F}^{n-m}$, and $M \in \text{GL}(n+m, \mathbb{F})$ and $M \cdot k$ denotes matrix-vector multiplication over \mathbb{F} , and “comma” denotes concatenation.

Theorem 6.1. *Assume that F is dna-RKA PRF then H is na-RKA PRF.*

Intuition. Note that the key k is mapped (via M) into two components: an “actual key” $k_1 = M_1 \cdot k \in \mathbb{F}^n$ which will be used as the key for F , and a padding part $k_2 = M_2 \cdot k \in \mathbb{F}^m$ which will be used to pad the input to F . Intuitively, the padding part should allow us to eliminate x -identical RKA queries. To see this first note that an RKA query (Δ, x) for $H_{k,M}$ correspond to the value $H_{k+\Delta}(x) = F_{k_1+M_1 \cdot \Delta}(x, k_2 + M_2 \cdot \Delta)$. Hence, a pair of H -RKA queries (x, Δ) and (x, Δ') will be

²The construction generalizes to the case where the lengths of the messages, keys and outputs are different.

mapped to a pair of F -RKA queries on the points $X = (x, k_2 + M_2 \cdot \Delta)$ and $X' = (x, k_2 + M_2 \cdot \Delta')$ since the choice of Δ, Δ' is independent of M_1 which is (almost) uniform matrix in $\mathbb{F}^{(n+m) \times n}$, the points X and X' are likely to differ on their suffix. Furthermore, since the mapping $k \mapsto (k_1, k_2)$ is linear, we can translate linear related key queries against H into linear related-key queries for F . Finally, the fact that the mapping is invertible allows us to map the oracle F_{k_1} into the oracle $H_{k,M}$ while preserving the uniform distribution of the keys. Indeed, this can be done by sampling k_2 and M and thinking of k as the (unique) preimage of (k_1, k_2) under the linear mapping induce by M . We proceed with a formal proof.

Proof. Let \mathcal{A} be a an adversary which breaks the non-adaptive RKA security of H with non-negligible advantage ε . Since \mathcal{A} is non-adaptive we can partition it to two adversaries (with joint randomness r): a query generator $Q(1^n; r)$ which samples $t = t(n)$ pairs $(\Delta_i, x_i)_{i \in [t]}$ and a distinguisher $D(1^n, M, (\Delta_i, x_i, y_i)_{i \in [t]}; r)$ which outputs a verdict b . We are guaranteed that D distinguishes with advantage $\varepsilon(n)$ between the following distributions

$$(M, (\Delta_i, x_i, y_i)_{i \in [t]}; r) \quad \text{and} \quad M, (\Delta_i, x_i, z_i)_{i \in [t]}; r) \quad (4)$$

where $M \stackrel{R}{\leftarrow} \text{GL}(n+m, \mathbb{F})$, $(\Delta_i, x_i)_{i \in [t]} \stackrel{R}{\leftarrow} Q(1^n; r)$, $k \stackrel{R}{\leftarrow} \mathbb{F}^{n+m}$, $y_i = H_{k+\Delta_i, M}(x_i)$, $z_i \stackrel{R}{\leftarrow} \mathbb{F}^{n+m}$ and the coins r of the distinguisher are chosen uniformly at random. (Without loss of generality, we assume that there are no fully identical queries $(\Delta_i, x_i) = (\Delta_j, x_j)$ for some $i \neq j$.) We will construct a non-adaptive *distinct* adversary \mathcal{B} for $F_{k'}$ ($k' \in \mathbb{F}^n$) as follows:

- Given 1^n sample $(\Delta_i, x_i)_{i \in [t]} \stackrel{R}{\leftarrow} Q(1^n; r)$.
- Sample a random vector $k_2 \stackrel{R}{\leftarrow} \mathbb{F}^m$ and an invertible square matrix $M \in \text{GL}(n+m, \mathbb{F})$ and parse M as $M_1 \in \mathbb{F}^{n \times 2n}$ and $M_2 \in \mathbb{F}^{m \times 2n}$ as in the construction.
- Let $\Delta'_i = M_1 \Delta_i$ and $x'_i = (x_i, k_2 + M_2 \cdot \Delta_i)$.
- If there is a pair of identical inputs $x'_i = x'_j$, abort with failure.
- Call the oracle with (Δ'_i, x'_i) and let w_i be its answers.
- Output the value of the distinguisher $D(M, (\Delta_i, x_i, y_i)_{i \in [t]}; r)$.

Note that \mathcal{B} is a distinct non-adaptive RKA adversary for F . We will show that \mathcal{B} breaks F with advantage $\varepsilon - \text{neg}(n)$ and derive a contradiction. The analysis relies on the following claims.

Claim 6.2. *The algorithm aborts with probability of at most $t^2 \cdot (|\mathbb{F}|^{-m} + m|\mathbb{F}|^{m-n}) < \text{poly}(n) \cdot 2^{-\delta n}$.*

Proof. Fix some $\Delta \neq \Delta'$ we show that, with all but negligible probability (over the choice of M) it holds that $M_2 \cdot \Delta \neq M_2 \Delta'$ the claim then follows by a union-bound over all the $t^2 = \text{poly}(n)$ pairs of queries. We begin by proving the claim for the case that M_2 is chosen uniformly from $\mathbb{F}^{m \times (n+m)}$. (We will later show that the actual distribution of M_2 is statistically-close to uniform and so the claim will follow.) Indeed, the bad event $M_2 \cdot \Delta = M_2 \Delta'$ happens only if $M_2(\Delta - \Delta') = \mathbf{0}$ which happens with probability $|\mathbb{F}|^{-m}$.

It is left to show that M_2 is statistically-close to a uniform $m \times (n+m)$ matrix in \mathbb{F} . Indeed, consider the following procedure for sampling M : sample each row uniformly and independently

from \mathbb{F}^{n+m} and keep it if it is not in the row-span of all previous rows. Observe that in each of the first $m = \delta n$ steps the probability of rejecting a row is at most $|\mathbb{F}|^{m-n}$. Hence, the probability of rejection in the first m steps is at most $m \cdot |\mathbb{F}|^{m-n}$. It follows that M_2 is statistically-close to uniform. \square

Claim 6.3. *Assume that \mathcal{B} does not abort and that its RKA queries are answered by F'_k for a random $k' \xleftarrow{R} \mathbb{F}^n$. Then, the resulting joint distribution $(M, (\Delta_i, x_i, w_i)_{i \in [t]}, r)$ is identical to the input distribution $(M, (\Delta_i, x_i, y_i)_{i \in [t]}, r)$ of D when attacking H as defined in the left part of Eq. (4).*

Proof. Conditioned on not aborting the joint distribution of $M, (\Delta_i, x_i)$ is identical in both cases. For $k' = k_1 \xleftarrow{R} \mathbb{F}^n$ and $k_2 \xleftarrow{R} \mathbb{F}^m$, let $k = M^{-1}(k_1, k_2)$. Then, k is uniformly distributed over \mathbb{F}^{n+m} and, by linear algebra, for every $i \in [t]$

$$w_i = F_{k'+\Delta'_i}(x'_i) = F_{(k+\Delta_i) \cdot M_1}(x_i, M_2 \cdot (k + \Delta_i)) = F_{k+\Delta_i}(x_i) = y_i,$$

as required. \square

The following claim follows immediately from the definition of \mathcal{B} .

Claim 6.4. *Assume that \mathcal{B} does not abort and that its RKA queries are answered randomly by $R(\mathbb{F}^n)$. Then, the resulting joint distribution $(M, (\Delta_i, x_i, w_i)_{i \in [t]}, r)$ is identical to the input distribution $(M, (\Delta_i, x_i, z_i)_{i \in [t]}, r)$ of D when attacking H as defined in the right part of Eq. (4).*

It follows the distinguishing advantage of \mathcal{B} is the same (up to a negligible loss due to the probability of aborting) as the distinguishing advantage of \mathcal{A} , and the theorem follows. \square

7 Bounded-RKA secure PRF

In this section we show how to immunize any standard PRF (or PRP) $\text{PRF} : K \times X \rightarrow Y$ against a bounded related-key attack which makes use of at most t related keys, where $t(n)$ is an arbitrary (a-priori fixed) polynomial in the security parameter. (This will hold for any arbitrary additive relation, including XOR, and addition modulo 2^n .) The idea is to use a long key s taken from a large key-space S (larger than K^t) and use some public hash function h to derive a shorter key $h(s) \in K$ for PRF.

Let us say that h is t -good if for any t -tuple of distinct shifts $(\Delta_1, \dots, \Delta_t)$, the joint distribution of all the keys

$$(h(s), h(s + \Delta_1), \dots, h(s + \Delta_t)), \quad (5)$$

induced by a random choice of $s \xleftarrow{R} S_n$, is ε -close in statistical distance to the uniform distribution over K_n^{t+1} for some negligible function $\varepsilon(n)$. It is not hard to verify that if h is t -good, then the resulting PRF is t -bounded RKA secure. In fact, it will be useful to consider a collection of hash function $H = \{H_z\}$ such that, with all but negligible probability, a randomly selected hash function $H_z \xleftarrow{R} H$ is t -good. We refer to such an ensemble as t -good.

Lemma 7.1. *For a pseudorandom function $\text{PRF} : K \times X \rightarrow Y$ and t -good ensemble $H : Z \times S \rightarrow K$, define the ensemble G with key-space S , domain X , range Y and public parameter space Z via the mapping*

$$G_{s,z} : x \mapsto \text{PRF}_k(x), \quad \text{where } k = H_z(s), s \in S_n, x \in X_n, \text{ and } z \in Z_n \text{ is a public parameter.}$$

Then, the PRF G is secure against t -bounded RKA.

Proof. Let A be a PPT RKA-adversary that uses at most t -distinct shifts and ε -breaks G , namely,

$$\Pr_{s \stackrel{R}{\leftarrow} K_n, z \stackrel{R}{\leftarrow} Z_n} [A^{G_{s,z}^+}(1^n, z) = 1] - \Pr_{z \stackrel{R}{\leftarrow} Z_n} [A^{\mathbf{R}^+(X_n \rightarrow Y_n)}(1^n, z) = 1] > \varepsilon(n),$$

for some non-negligible function $\varepsilon(n)$. For $\vec{k} = (k_0, \dots, k_t) \in K_n^{t(n)+1}$, define the following (stateful) oracle $\mathcal{O}_{\vec{k}}$. Given a standard (non-RKA) query x , the oracle $\mathcal{O}_{\vec{k}}$ answers with $\text{PRF}_{k_0}(x)$. To deal with RKA queries the oracle keeps track of the t different shifts $(\Delta_1, \dots, \Delta_t)$ used by the adversary, and answers an RKA query (Δ_i, x) with $\text{PRF}_{k_i}(x)$. (The oracle dynamically builds the list of seen Δ 's.) Since H is t -good, we have that

$$\Pr_{\vec{k} \stackrel{R}{\leftarrow} K_n^{t(n)+1}} [A^{\mathcal{O}_{\vec{k}}}(1^n, z) = 1] > \Pr_{s \stackrel{R}{\leftarrow} K_n, z \stackrel{R}{\leftarrow} Z_n} [A^{G_s^+}(1^n, z) = 1] - \text{neg}(n).$$

Assume (without loss of generality) that A does not repeat the same query (Δ, x) twice. Then, A distinguishes with non-negligible advantages between $\mathcal{O}_{\vec{k}}$, for a random key-vector \vec{k} , to a random oracle $\mathbf{R}(Y_n)$. It is not hard to show (via a standard hybrid argument) that this violates the security of the PRF PRF. \square

The following lemma shows that an $\Omega(t^2 \log(|K|))$ -wise independent hash function with a sufficiently large domain is t -good.

Lemma 7.2. *Let $t = \text{poly}(n)$, and let H be an ensemble of ℓ -wise independent hash function with domain $S = \{S_n\}$ and range $K = \{K_n\}$ where $\ell \geq n(2t+2)(t+1)$, $|K_n| = 2^n$ and $|S_n| = 2^{(2t+6)n}$. Then H is t -good. In particular, for all but a 2^{-n} fraction of the functions in H , the distribution (5) is $2^{-0.99n}$ -close to uniform.*

Proof. Fix a sequence of distinct non-zero shifts $(\Delta_1, \dots, \Delta_t)$. It will be convenient to let $\Delta_0 = 0$ and set $\vec{\Delta} = (\Delta_0, \Delta_1, \dots, \Delta_t)$. To simplify notation, we use lower-case h to denote a function H_z from the collection H . We say that $h \in H$ is ε -good for $\vec{\Delta}$ if for a random s , the distribution $(h(s + \Delta_i))_{0 \leq i \leq t}$ is ε -close to the uniform distribution over K^{t+1} . In order to bound the statistical distance, we prove the following claim.

Claim 7.3. *For all but $2^{-n(t+1)}$ -fraction of the h 's the following holds. For every vector of (not necessarily distinct) keys $\vec{k} = (k_0, \dots, k^t) \in K^{t+1}$,*

$$\Pr_s \left[\bigwedge_{i=0}^t h(s + \Delta_i) = k_i \right] \in \left(\frac{1}{|K|^{t+1}} \cdot (1 \pm 2^{-0.99n}) \right).$$

Proof. Fix some vector of keys $\vec{k} \in K^{t+1}$. For every $s \in S$ define a random variable χ_s which takes the value 1 if $h(s + \Delta_i) = k_i$ for every $0 \leq i \leq t$, where $h \stackrel{R}{\leftarrow} H$. Observe that the random variable (induced by a choice of h) $\Pr_s[\bigwedge_{i=0}^t h(s + \Delta_i) = k_i]$ can be written as $\bar{\chi} = \sum_{s \in S} \chi_s / |S|$. We will show that

$$\Pr_h \left[\bar{\chi} \notin \left(\frac{1}{|K|^{t+1}} (1 \pm 2^{-0.99n}) \right) \right] \leq 2^{-n(2t+2)}. \quad (6)$$

Note that Claim 7.3 follows from (6) by applying a union-bound over all $2^{(t+1)n}$ possible $\vec{k} \in K^{t+1}$.

To prove (6) we make the following observations. First, since H is ℓ -wise independent and $\ell > t+1$, we have that $E[\chi_s] = 1/|K|^{t+1}$ for every s , and, so, by the linearity of expectation, $E(\bar{\chi}) = 1/|K|^{t+1}$. We would like to show that the average of χ_s is concentrated around its expectation. Note that the random variables χ_s 's are not even pair-wise independent. Indeed, consider s and s' for which $s + \Delta_i = s' + \Delta_j$ for some $i \neq j$. Then, the i -th coordinate of the random variable $(h(s + \Delta_0), \dots, h(s + \Delta_t))$ equals to the j -th coordinate of the random variable $(h(s' + \Delta_0), \dots, h(s' + \Delta_t))$ and as a result χ_s is not statistically independent from $\chi_{s'}$. Fortunately, we can show that, apart from these local dependencies, the χ 's are r -wise independent, for $r = (\ell/(t+1)) \geq 2t+2$, which still yield a strong concentration (cf. [25]). We proceed with a formal proof.

Define a simple graph G over $s \in S$ by putting an edge between s and s' if $s + \Delta_i = s' + \Delta_j$ for some $i \neq j$. Observe that the degree of each node in the graph is at most $d = (t+1)^2$. We claim that for every independent set I in the graph, the random variables $\{\chi_s : s \in I\}$ are r -wise independent. In the terminology of [25], the random variables r -agree with G . Indeed, for any independent set I , and for any r -subset $(s_1, \dots, s_r) \subseteq I$, the value of each of the random variables $\chi_{s_1}, \dots, \chi_{s_r}$ solely depends on the value of h on a set of $t+1$ distinct points $\{s_j + \Delta_0, \dots, s_j + \Delta_{t+1}\}$. These sets of points are distinct (since I is an independent set) and so their images under h are statistically independent since h is ℓ -wise independent for $\ell = r \cdot (t+1)$. It therefore, follows that $\chi_{s_1}, \dots, \chi_{s_r}$ are statistically independent.

Having shown that the χ_s 's agree with G , we can apply Corollary 3.2 of [25] and conclude that

$$\Pr_h[\bar{\chi} \notin \frac{1}{|K|^{t+1}}(1 \pm \delta)] \leq 4\sqrt{\pi r} \left(\frac{|K|^{t+1} \sqrt{(d+1)r}}{\delta \sqrt{|S|}} \right)^r.$$

For $\delta = 2^{-0.99n}$, $|K| = 2^n$, $|S| = 2^{(2t+6)n}$, and $r, t \in \text{poly}(n)$, the RHS is upper-bounded by $2^{-nr} \leq 2^{-n(2t+2)}$, for all sufficiently large n 's, and (6) follows. \square

We can now complete the proof of the lemma. First, observe that any h that satisfies the lemma is $2^{-0.99n}$ -good for $\vec{\Delta}$. Indeed, for such an h , the probability distribution $(h(s + \Delta_i))_{0 \leq i \leq t}$ assigns to each outcome \vec{k} a weight which corresponds to its weight under the uniform distribution up to a multiplicative factor of $1 \pm 2^{-0.99n}$. It is not hard to show that in this case the statistical distance is at most $2^{-0.99n}$. Therefore, all but $2^{-n(t+1)}$ -fraction of the h 's are $2^{-0.99n}$ -good for $\vec{\Delta}$. By applying a union bound over all possible 2^{nt} shift vectors, we conclude that all but 2^{-n} -fraction of the h 's are $2^{-0.99n}$ -good, and the lemma follows. \square

By combining Lemmas 7.1, 7.2 with any efficient construction of ℓ -wise independent hash function (e.g., based on Reed-Solomon codes), we obtain the following theorem.

Theorem 7.4. *Let $K = \{(K_n, +_n)\}$ be a sequence of efficiently computable additive groups, and let $t(n)$ be an arbitrary polynomial. Then, assuming the existence of a PRF $\text{PRF} : K \times X \rightarrow Y$, there exists a t -bounded-RKA secure PRF with respect to addition over K .*

Specifically, assuming the existence of a one-way function, we can achieve bounded-RKA security with respect to XOR or addition modulu 2^n .

Acknowledgement. We thank Daniel Wichs for a detailed discussion about [19].

References

- [1] Michel Abdalla, Fabrice Benhamouda, and Alain Passelègue. An algebraic framework for pseudorandom functions and applications to related-key security. Cryptology ePrint Archive, Report 2015/554, 2015. <http://eprint.iacr.org/2015/554>.
- [2] Michel Abdalla, Fabrice Benhamouda, Alain Passelègue, and Kenneth G. Paterson. Related-key security for pseudorandom functions beyond the linear barrier. In *CRYPTO 2014, Part I*, LNCS, pages 77–94, Santa Barbara, CA, USA, August 2014. Springer, Berlin, Germany.
- [3] Benny Applebaum. Garbling XOR gates “for free” in the standard model. In Amit Sahai, editor, *TCC 2013*, volume 7785 of LNCS, pages 162–181, Tokyo, Japan, March 3–6, 2013. Springer, Berlin, Germany.
- [4] Benny Applebaum, Danny Harnik, and Yuval Ishai. Semantic security under related-key attacks and applications. In Bernard Chazelle, editor, *ICS 2011*, pages 45–60, Tsinghua University, Beijing, China, January 7–9, 2011. Tsinghua University Press.
- [5] Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of LNCS, pages 719–737, Cambridge, UK, April 15–19, 2012. Springer, Berlin, Germany.
- [6] Boaz Barak, Yevgeniy Dodis, Hugo Krawczyk, Olivier Pereira, Krzysztof Pietrzak, François-Xavier Standaert, and Yu Yu. Leftover hash lemma, revisited. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of LNCS, pages 1–20, Santa Barbara, CA, USA, August 14–18, 2011. Springer, Berlin, Germany.
- [7] Mihir Bellare and David Cash. Pseudorandom functions and permutations provably secure against related-key attacks. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of LNCS, pages 666–684, Santa Barbara, CA, USA, August 15–19, 2010. Springer, Berlin, Germany.
- [8] Mihir Bellare, David Cash, and Rachel Miller. Cryptography secure against related-key attacks and tampering. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of LNCS, pages 486–503, Seoul, South Korea, December 4–8, 2011. Springer, Berlin, Germany.
- [9] Mihir Bellare and Tadayoshi Kohno. A theoretical treatment of related-key attacks: RKA-PRPs, RKA-PRFs, and applications. In Eli Biham, editor, *EUROCRYPT 2003*, volume 2656 of LNCS, pages 491–506, Warsaw, Poland, May 4–8, 2003. Springer, Berlin, Germany.
- [10] Mihir Bellare, Kenneth G. Paterson, and Susan Thomson. RKA security beyond the linear barrier: IBE, encryption and signatures. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of LNCS, pages 331–348, Beijing, China, December 2–6, 2012. Springer, Berlin, Germany.
- [11] Eli Biham. New types of cryptanalytic attacks using related keys (extended abstract). In Tor Helleseeth, editor, *EUROCRYPT’93*, volume 765 of LNCS, pages 398–409, Lofthus, Norway, May 23–27, 1993. Springer, Berlin, Germany.

- [12] Eli Biham, Orr Dunkelman, and Nathan Keller. New cryptanalytic results on IDEA. In Xuejia Lai and Kefei Chen, editors, *ASIACRYPT 2006*, volume 4284 of *LNCS*, pages 412–427, Shanghai, China, December 3–7, 2006. Springer, Berlin, Germany.
- [13] Eli Biham, Orr Dunkelman, and Nathan Keller. A unified approach to related-key attacks. In Kaisa Nyberg, editor, *FSE 2008*, volume 5086 of *LNCS*, pages 73–96, Lausanne, Switzerland, February 10–13, 2008. Springer, Berlin, Germany.
- [14] Alex Biryukov, Orr Dunkelman, Nathan Keller, Dmitry Khovratovich, and Adi Shamir. Key recovery attacks of practical complexity on AES-256 variants with up to 10 rounds. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 299–319, French Riviera, May 30 – June 3, 2010. Springer, Berlin, Germany.
- [15] Dan Boneh, Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Key homomorphic PRFs and their applications. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 410–428, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Berlin, Germany.
- [16] Joan Daemen and Vincent Rijmen. *The Design of Rijndael: AES - The Advanced Encryption Standard*. Information Security and Cryptography. Springer, 2002.
- [17] Yevgeniy Dodis, Eike Kiltz, Krzysztof Pietrzak, and Daniel Wichs. Message authentication, revisited. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 355–374, Cambridge, UK, April 15–19, 2012. Springer, Berlin, Germany.
- [18] Shimon Even, Oded Goldreich, and Abraham Lempel. A randomized protocol for signing contracts. In David Chaum, Ronald L. Rivest, and Alan T. Sherman, editors, *CRYPTO’82*, pages 205–210, Santa Barbara, CA, USA, 1982. Plenum Press, New York, USA.
- [19] Sebastian Faust, Pratyay Mukherjee, Daniele Venturi, and Daniel Wichs. Efficient non-malleable codes and key-derivation for poly-size tampering circuits. In *EUROCRYPT 2014*, *LNCS*, pages 111–128. Springer, Berlin, Germany, 2014.
- [20] Niels Ferguson, John Kelsey, Stefan Lucks, Bruce Schneier, Michael Stay, David Wagner, and Doug Whiting. Improved cryptanalysis of Rijndael. In Bruce Schneier, editor, *FSE 2000*, volume 1978 of *LNCS*, pages 213–230, New York, NY, USA, April 10–12, 2000. Springer, Berlin, Germany.
- [21] Rosario Gennaro, Anna Lysyanskaya, Tal Malkin, Silvio Micali, and Tal Rabin. Algorithmic tamper-proof (ATP) security: Theoretical foundations for security against hardware tampering. In Moni Naor, editor, *TCC 2004*, volume 2951 of *LNCS*, pages 258–277, Cambridge, MA, USA, February 19–21, 2004. Springer, Berlin, Germany.
- [22] David Goldenberg and Moses Liskov. On related-secret pseudorandomness. In Daniele Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 255–272, Zurich, Switzerland, February 9–11, 2010. Springer, Berlin, Germany.
- [23] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM*, 33:792–807, 1986.

- [24] Vipul Goyal, Adam O’Neill, and Vanishree Rao. Correlated-input secure hash functions. In Yuval Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 182–200, Providence, RI, USA, March 28–30, 2011. Springer, Berlin, Germany.
- [25] Ronen Gradwohl and Amir Yehudayoff. t -wise independence with local dependencies. *Information Processing Letters*, 106(5):208–212, 2008.
- [26] Thomas Holenstein. Key agreement from weak bit agreement. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 664–673, Baltimore, Maryland, USA, May 22–24, 2005. ACM Press.
- [27] Goce Jakimoski and Yvo Desmedt. Related-key differential cryptanalysis of 192-bit key AES variants. In Mitsuru Matsui and Robert J. Zuccherato, editors, *SAC 2003*, volume 3006 of *LNCS*, pages 208–221, Ottawa, Ontario, Canada, August 14–15, 2003. Springer, Berlin, Germany.
- [28] Lars R. Knudsen. Cryptanalysis of LOKI. In Hideki Imai, Ronald L. Rivest, and Tsutomu Matsumoto, editors, *ASIACRYPT’91*, volume 739 of *LNCS*, pages 22–35, Fujiyoshida, Japan, November 11–14, 1991. Springer, Berlin, Germany.
- [29] Stefan Lucks. Ciphers secure against related-key attacks. In Bimal K. Roy and Willi Meier, editors, *FSE 2004*, volume 3017 of *LNCS*, pages 359–370, New Delhi, India, February 5–7, 2004. Springer, Berlin, Germany.
- [30] Moni Naor and Omer Reingold. Synthesizers and their application to the parallel construction of pseudo-random functions. In *36th FOCS*, pages 170–181, Milwaukee, Wisconsin, October 23–25, 1995. IEEE Computer Society Press.
- [31] Krzysztof Pietrzak. Composition implies adaptive security in minicrypt. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 328–338, St. Petersburg, Russia, May 28 – June 1, 2006. Springer, Berlin, Germany.
- [32] Krzysztof Pietrzak and Johan Sjödin. Weak pseudorandom functions in minicrypt. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfssdóttir, and Igor Walukiewicz, editors, *ICALP 2008, Part II*, volume 5126 of *LNCS*, pages 423–436, Reykjavik, Iceland, July 7–11, 2008. Springer, Berlin, Germany.
- [33] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93, Baltimore, Maryland, USA, May 22–24, 2005. ACM Press.
- [34] Liting Zhang, Wenling Wu, Peng Wang, Lei Zhang, Shuang Wu, and Bo Liang. Constructing rate-1 MACs from related-key unpredictable block ciphers: PGV model revisited. In Seokhie Hong and Tetsu Iwata, editors, *FSE 2010*, volume 6147 of *LNCS*, pages 250–269, Seoul, Korea, February 7–10, 2010. Springer, Berlin, Germany.
- [35] Wentao Zhang, Lei Zhang, Wenling Wu, and Dengguo Feng. Related-key differential-linear attacks on reduced AES-192. In K. Srinathan, C. Pandu Rangan, and Moti Yung, editors, *INDOCRYPT 2007*, volume 4859 of *LNCS*, pages 73–85, Chennai, India, December 9–13, 2007. Springer, Berlin, Germany.

A Simple Message Authentication Code

In [17] it is shown how to construct an efficient MAC based a key-homomorphic wPRF. While our wPRF does not support a full key-homomorphism in the strict sense of [17], we show how to adopt the RKA-security proof from Section 3 to prove the security of the resulting MAC. The result is a very simple MAC based on the hardness of lattices.

A.1 Definitions

Syntax. A message authentication code $\text{MAC} = (\text{KG}, \text{TAG}, \text{VRFY})$ is a triple of algorithms associated with key space \mathcal{K} , message space \mathcal{M} and tag space \mathcal{T} :

- KG is a probabilistic algorithm which takes as input the security parameter 1^n and outputs a secret key $k \in \mathcal{K}$.
- $\text{TAG} : \mathcal{K} \times \mathcal{M} \rightarrow \mathcal{T}$ is a probabilistic authentication algorithm which takes as input a secret key and a message and outputs a corresponding tagging. (Since the algorithm is probabilistic the tag may not be unique.)
- $\text{VRFY} : \mathcal{K} \times \mathcal{M} \times \mathcal{T} \rightarrow \{\text{accept}, \text{reject}\}$ is a deterministic verification algorithm which receives a key, a message and an authentication tag triplet and returns `accept` or `reject`. We require that for every $k \in \mathcal{K}, m \in \mathcal{M}$

$$\Pr[\text{VRFY}_k(m, \text{TAG}_k(m)) = \text{accept}] = 1.$$

Security. We consider two models of forgery attacks: *universal forgery under chosen message and verification attacks* (`ufcmva`) and *selective universal-forgability under chosen message attacks* (`sufcma`). For an adversary A and message authentication code $\text{MAC} = (\text{KG}, \text{TAG}, \text{VRFY})$, we define the following games.

- Game $G_{\text{MAC}}^{\text{ufcmva}}(1^n, A)$:
 1. Set $k \leftarrow \text{KG}(1^n)$.
 2. Invoke $A^{\text{TAG}_k(\cdot), \text{VRFY}_k(\cdot, \cdot)}$ and answer its `TAG` and `VRFY` queries.
 3. Output 1 iff A queries (m^*, σ^*) such that $\text{VRFY}_k(m^*, \sigma^*) = \text{accept}$ and A did not receive $\text{TAG}_k(m^*)$ as a query answer.
- Game $G_{\text{MAC}}^{\text{sufcma}}(1^n, A)$:
 1. Set $k \leftarrow \text{KG}(1^n)$.
 2. A announces a target message $m^* \in \mathcal{M}$.
 3. Invoke $A^{\text{TAG}_k(\cdot)}$ and answer only queries $m \neq m^*$.
 4. At the end of the game A outputs σ^* , the value of the game is 1 iff $\text{VRFY}_k(m^*, \sigma^*) = \text{accept}$.

Definition A.1 (`ufcmva` and `sufcma` security). *We say that a $\text{MAC} = (\text{KG}, \text{TAG}, \text{VRFY})$ is `ufcmva` secure (resp., `sufcma` secure) if for every PPT adversary A the winning probability $\Pr[G_{\text{MAC}}^{\text{ufcmva}}(1^n, A)$ (resp., $\Pr[G_{\text{MAC}}^{\text{sufcma}}(1^n, A) = 1]$) is negligible in n .*

It is not hard to see that ufcma security implies sufcma security. In [17] it is shown that if a sufcma secure MAC is also “message-hiding” (as our construction will be) then it can be efficiently converted to a ufcma secure MAC.

A.2 Construction

Let $\ell < p < q$ be integer-valued functions such that $\ell(n)/p(n) < \text{neg}(n)$, and let $F^{q,p}$ be the LWR-based weak-PRF function (See Definition 2.2).

Construction A.2 (MAC). Define $\text{MAC} = (\text{KG}, \text{TAG}, \text{VRFY})$ with key-space $\mathbb{Z}_q^n \times \mathbb{Z}_q^n$ and message space $\{-\frac{\ell-1}{2}, \dots, \frac{\ell+1}{2}\} \subseteq \mathbb{Z}_q$ as follows.

- *Key Generation:* $\text{KG}(1^n)$ chooses $k_1, k_2 \stackrel{R}{\leftarrow} \mathbb{Z}_q^n$ uniformly at random and outputs $k = (k_1, k_2)$.
- *Tagging:* $\text{TAG}_{(k_1, k_2)}(m)$ chooses $x \stackrel{R}{\leftarrow} \mathbb{Z}_q^n$ uniformly at random and sets $y = F_{m \cdot k_1 + k_2}(x) = \lceil \langle m \cdot k_1 + k_2, x \rangle \rceil_{q/p}$. Output $\sigma = (x, y)$.
- *Verification:* $\text{VRFY}_{(k_1, k_2)}(m, \sigma)$ parses $\sigma = (x, y)$ and accepts iff $F_{m \cdot k_1 + k_2}(x) = \lceil \langle m \cdot k_1 + k_2, x \rangle \rceil_{q/p} = y$.

To prove the security of our construction we rely on the following mapping. For $\ell < p < r < q$, let

$$h : (\Delta, x, \alpha, y) \mapsto \alpha \cdot y + F_{\Delta}^{q,r}(x) \pmod{r}, \quad (7)$$

where $\alpha \in \{-\frac{\ell-1}{2}, \dots, \frac{\ell+1}{2}\}$, $\Delta, x \in \mathbb{Z}_q^n$ and $y \in \mathbb{Z}_r$.

Claim A.3. For every $s, \Delta, x \in \mathbb{Z}_q^n$ and every $\alpha \in \{-\frac{\ell-1}{2}, \dots, \frac{\ell+1}{2}\}$, if the integer $h(\Delta, x, \alpha, F_s^{q,r}(x))$ is at least $\frac{\ell+3}{2}$ -far from a multiple of r/p then

$$\lceil h(\Delta, x, \alpha, F_s^{q,r}(x)) \rceil_{r/p} = F_{\alpha \cdot s + \Delta}^{q,p}(x).$$

Proof.

$$\begin{aligned} \alpha \cdot F_s^{q,r}(x) + F_{\Delta}^{q,r}(x) &= \alpha \cdot \left\lfloor \langle x, s \rangle \cdot \frac{r}{q} \right\rfloor + \left\lfloor \langle x, \Delta \rangle \cdot \frac{r}{q} \right\rfloor \\ &= \alpha \cdot \langle x, s \rangle \cdot \frac{r}{q} + \langle x, \Delta \rangle \cdot \frac{r}{q} + \alpha \cdot e_1 + e_2 \\ &= \left\lfloor \langle x, \alpha \cdot s + \Delta \rangle \cdot \frac{r}{q} \right\rfloor + \alpha \cdot e_1 + e_2 + e_3 \\ &= F_{\alpha \cdot s + \Delta}^{q,r}(x) + e \end{aligned}$$

where the “error” terms $e_1, e_2 \in [-1, 0]$ and $e_3 \in [0, 1]$ compensate the omission/addition of the floor operation, and $e = \alpha \cdot e_1 + e_2 + e_3 \in [\pm(\alpha - 1)]$ denotes the accumulated error.

Since $|\alpha - 1| \leq \frac{\ell+3}{2}$, $\alpha \cdot F_s^{q,r}(x) + F_{\Delta}^{q,r}(x)$ is $\frac{\ell+3}{2}$ -close to $F_{\alpha \cdot s + \Delta}^{q,r}(x)$, it follows, by our hypothesis, that both elements reside in the same r/p -interval. We can therefore write:

$$\lceil h(\Delta, x, \alpha, F_s^{q,r}(x)) \rceil_{r/p} = \lceil \alpha \cdot F_s^{q,r}(x) + F_{\Delta}^{q,r}(x) \rceil_{r/p} = \lceil F_{\alpha \cdot s + \Delta}^{q,r}(x) \rceil_{r/p} = F_{\alpha \cdot s + \Delta}^{q,p}(x)$$

and the claim follows. \square

In the following we let $\varphi(p)$ denote Euler's phi function that counts the totatives of p , i.e., $\varphi(p) = |\mathbb{Z}_p^*|$. Also, recall that $\delta(p, q)$ is zero if p divides q , and $\frac{p}{2q}$ otherwise.

Theorem A.4. *Let $\ell < p < r < q$ be integer valued functions such that $r(n)$ is a prime and all the following quantities are negligible in n :*

$$\ell(n)p(n)/r(n), \quad 1/p(n), \quad \text{and} \quad \left(1 - \frac{\varphi(q(n))}{q(n)}\right)^n.$$

Then, if $F^{q,r}$ is a weak PRF, then Construction A.2 is a sufcma-MAC.

For example, one can let $\ell = p = 2^n$, r be a prime of magnitude 2^{3n} , and $q = 2^{4n}$.

Proof. Fix some ℓ, p, r, q that satisfy the hypothesis. Assume, towards a contradiction, that A is an adversary that breaks the security of the construction above as sufcma-MAC. Namely, the adversary A forges a successful tag on a predeclared message m^* with a non negligible probability ε . We construct an adversary B that distinguishes the oracle $\text{Sam}F_s^{q,r}$ from the oracle $\text{R}(\mathbb{Z}_q^n, \mathbb{Z}_r)$ with advantage $\varepsilon - \text{neg}(n)$. We begin with a high-level description.

Let us think of B 's oracle as $\text{Sam}F_s^{q,r}$. The key idea is to emulate A 's attack against $\text{MAC}_{k_1=s, k_2=\Delta-m^* \cdot s}$ where $\Delta \xleftarrow{R} \mathbb{U}_q^n$ is chosen at random by B and m^* is the forgery target message of A . To answer A 's queries we use the (approximate) homomorphism h defined in (7). We will show that the emulation is close to the real game, and so the adversary A forges a tag y^* on m^* with non-negligible probability. Note that B can compute the forged tag by herself as it is equal to

$$\sigma = (x^*, y^*) = (x^*, F_{m^* \cdot s + \Delta - m^* \cdot s}^{q,r}(x^*)) = (x^*, F_{\Delta}^{q,r}(x^*)).$$

Therefore, B can identify a successful forgery. On the other hand, if the oracle \mathcal{O} is a random oracle, then A is supplied with random tags which are almost independent of the correct tag, and so it cannot win the game with more than negligible probability. Therefore, B can distinguish the wPRF from a random oracle by checking if A 's attack succeeds. In fact, the above description is not fully accurate as the emulation itself may fail (since the homomorphism is not perfect), however, we show that failure is unlikely to happen when \mathcal{O} is a random oracle, and so failure allows to break the wPRF as well.

We move on to a formal description of $B^{\mathcal{O}}(1^n)$.

1. B chooses at random $\Delta \xleftarrow{R} \mathbb{U}_q^n$.
2. B emulates A , which announces $m^* \in \{-\frac{\ell-1}{2}, \dots, \frac{\ell+1}{2}\}$ as its forgery target message.
3. If A makes a query $m \in \{-\frac{\ell-1}{2}, \dots, \frac{\ell+1}{2}\}$ to its tagging oracle, B asks for a sample $(x, y) \xleftarrow{R} \mathcal{O}$, and continues as follows:
 - If $h(\Delta, x, \alpha = (m - m^*), y)$ is $\frac{\ell+3}{2}$ -close to a multiple of r/p then B quits with the output 1.
 - Otherwise, B computes $y' = \lceil h(\Delta, x, \alpha = (m - m^*), y) \rceil_{r/p}$ and returns the tag (x, y') .
 - When reached the end of the emulation B intercepts the forged pair (x^*, y^*) , B checks and outputs 1 iff $y^* = F_{\Delta}^{q,p}(x^*) = \lceil \langle x^*, \Delta \rangle \rceil_{q/p}$.

Let us analyze the distinguishing advantage of B assuming that A makes $t = \text{poly}(n)$ queries.

Claim A.5. $\Pr[B^{\mathbf{R}(\mathbb{Z}_q^n, \mathbb{Z}_r)}(1^n) = 1] \leq t(\ell + 3)p/r + t\delta(p, r) + \left(1 - \frac{\varphi(q)}{q}\right)^n + 1/p + \delta(p, q) = \text{neg}(n)$.

Proof. For each query, y is uniform over \mathbb{Z}_r and since $(m - m^*)$ is non-zero, the product $(m - m^*) \cdot y$ is also uniform over \mathbb{Z}_r (recall that r is a prime). It follows that the value $h(\Delta, x, (m - m^*), y) = (m - m^*) \cdot y + F_{\Delta}^{q,r}(x)$ is uniform over \mathbb{Z}_r . We therefore conclude that, for each query, the probability that B halts is at most $\frac{(\ell+3)p}{r}$, and the overall halting probability is

$$\Pr[B^{\mathbf{R}(\mathbb{Z}_q^n, \mathbb{Z}_r)}(1^n) \text{ halts}] \leq t \frac{(\ell + 3)p}{r}.$$

Conditioned on not halting, A is simulated with the tags (x, y') where $y' = \lceil h(\Delta, x, (m - m^*), y) \rceil_{q/p}$. Since $h(\Delta, x, (m - m^*), y)$ is uniform over \mathbb{Z}_r , by Proposition 2.4, y' is $\delta(p, r)$ -close to uniform over \mathbb{Z}_p . Therefore, by Claim 3.1, we have that

$$\Pr[B^{\mathbf{R}(\mathbb{Z}_q^n, \mathbb{Z}_r)}(1^n) = 1 \mid \text{not halting}] \leq \Pr[B'(1^n) = 1] + t\delta(p, r),$$

where B' is similar to B except that A 's queries are answered with random values $\mathbf{R}(\mathbb{Z}_q^n, \mathbb{Z}_p)$. Finally, we claim that

$$\Pr[B'(1^n) = 1] \leq \left(1 - \frac{\varphi(q)}{q}\right)^n + 1/p + \delta(q, p).$$

Indeed, let us further condition on the event that $\Delta \in \mathbb{Z}_q^n$ has at least one invertible element, which happens with probability $1 - (1 - \frac{\varphi(q)}{q})^n$. In this case, $F_{\Delta}^{q,p}(x^*) = \lceil \langle x^*, \Delta \rangle \rceil_{q/p}$ is $\delta(p, q)$ -close to uniform (and independent of A 's view) and so the probability that A guesses $F_{\Delta}^{q,p}(x^*)$ is at most $(1/p + \delta(p, q))$. By applying a union bound over all the above terms, we conclude that $\Pr[B^{\mathbf{R}(\mathbb{Z}_q^n, \mathbb{Z}_r)}(1^n) = 1]$ is upper-bounded by

$$t(\ell + 3)p/r + t\delta(p, r) + \left(1 - \frac{\varphi(q)}{q}\right)^n + 1/p + \delta(p, q) \leq t(\ell + 3)p/r + 2tp/r + \left(1 - \frac{\varphi(q)}{q}\right)^n + 1/p + 2p/q,$$

which, by our choice of the parameters, is negligible in n . \square

We move on to the case where the oracle \mathcal{O} is $\text{Sam}F_s^{q,r}$.

Claim A.6. $\Pr[B^{\text{Sam}F_s^{q,r}}(1^n) = 1] \geq \varepsilon$.

Proof. Let $\delta = \Pr_s[B^{\text{Sam}F_s^{q,r}}(1^n) \text{ quits}]$. Since B outputs 1 when it quits we have that

$$\begin{aligned} \Pr_s[B^{\text{Sam}F_s^{q,r}}(1^n) = 1] &= \delta + \Pr_s[B^{\text{Sam}F_s^{q,r}}(1^n) = 1 \mid \text{not quitting}] \cdot (1 - \delta) \\ &\geq \Pr_s[B^{\text{Sam}F_s^{q,r}}(1^n) = 1 \mid \text{not quitting}]. \end{aligned}$$

Let us condition on the event that all the queries that A made were answered by B without quitting. Then, A receives tags of the form (x, y') where $y' = \lceil (m - m^*) \cdot y + F_{\Delta}^{q,r}(x) \rceil_{q/p}$ and where $(m - m^*) \cdot y + F_{\Delta}^{q,r}(x)$ is at least $\frac{\ell+3}{2}$ far from a multiple of r/p . By Claim A.3 it follows that B

perfectly emulates the attack of A against MAC_{k_1, k_2} where $(k_1 = s, k_2 = \Delta - m^* \cdot s)$ is a random pair. It therefore follows that

$$\Pr_s[B^{\text{Sam}_{F_s^{q,r}}}(1^n) = 1 \mid \text{not quitting}] = \varepsilon,$$

and the claim follows. □

The theorem follows by combining Claim A.5 and A.6. □