An extended abstract of this paper appears in *Advances in Cryptology – EUROCRYPT '03*, Lecture Notes in Computer Science Vol. 2656, E. Biham ed., Springer-Verlag, 2003. This is the full version.

A Theoretical Treatment of Related-Key Attacks: RKA-PRPs, RKA-PRFs, and Applications

Mihir Bellare* Tadayoshi Kohno†

August 2, 2021

Abstract

We initiate a theoretical investigation of the popular block-cipher design-goal of security against "related-key attacks" (RKAs). We begin by introducing definitions for the concepts of PRPs and PRFs secure against classes of RKAs, each such class being specified by an associated set of "related-key deriving (RKD) functions." Then for some such classes of attacks, we prove impossibility results, showing that no block-cipher can resist these attacks while, for other, related classes of attacks that include popular targets in the block cipher community, we prove possibility results that provide theoretical support for the view that security against them is achievable. Finally we prove security of various block-cipher based constructs that use related keys, including a tweakable block cipher given in [22]. We believe this work helps block-cipher designers and cryptanalysts by clarifying what classes of attacks can and cannot be targets of design. It helps block-cipher users by providing guidelines about the kinds of related keys that are safe to use in constructs, and by enabling them to prove the security of such constructs. Finally, it puts forth a new primitive for consideration by theoreticians with regard to open questions about constructs based on minimal assumptions.

Keywords: Block ciphers, related-key attacks, pseudorandom permutations, tweakable block ciphers, concrete security, ideal-ciphers, Shannon-ciphers, Shannon-security.

*Dept. of Computer Science & Engineering, University of California at San Diego, 9500 Gilman Drive, La Jolla, California 92093, USA. E-Mail: mihir@cs.ucsd.edu. URL: http://www-cse.ucsd.edu/users/mihir. Supported in part by NSF Grant CCR-0098123, NSF Grant ANR-0129617 and an IBM Faculty Partnership Development Award.

†Dept. of Computer Science & Engineering, University of California at San Diego, 9500 Gilman Drive, La Jolla, California 92093, USA. E-mail: tkohno@cs.ucsd.edu. URL: http://www-cse.ucsd.edu/users/tkohno. Supported by a National Defense Science and Engineering Graduate Fellowship.

Contents

1 Introduction

Most modern block ciphers, including AES [9], are designed with the explicitly stated goal of resisting what are called "related-key attacks (RKAs)" [5]. However, it is not clear exactly what types of attacks this encompasses, and against which of these security is even achievable.

Towards answering such questions, this paper provides a theoretical treatment of related-key attacks. Via notions of RKA secure PRPs and PRFs parameterized by a class of "related-key deriving functions," we provide a formal definition of what it means for a block cipher to be secure against a given class of related-key attacks. Then for some classes of attacks, we prove impossibility results, showing that no block-cipher can resist these attacks while, for other, related classes of attacks that include popular targets in the block cipher community, we prove possibility results that provide theoretical support for the view that security against them is achievable. We also prove security of some specific related-key-using block-cipher-based constructs based on assumptions about the security of the block cipher under an appropriate class of RKAs.

This work can help block-cipher designers and cryptanalysts by clarifying what classes of attacks can and cannot be targets of design. It can help block-cipher users by providing guidelines about the kinds of related keys that are safe to use in constructs, and by enabling them to prove security of the resulting constructs. Finally, it puts forth a new primitive for consideration by theoreticians with regard to constructions based on minimal complexity assumptions.

Overall our results indicate that there is a thin dividing line between unachievable and achievable goals in this area, and thus a need for care on the part of both designers and users.

Let us now discuss the background and our results in more detail.

RKAs. Under a related-key attack, an adversary can obtain input-output examples of the block cipher E, not just under the target key K, but under keys $K_1, K_2, \ldots$ related to K. However the understanding of what "related" means seems currently to be based only on specific examples, such as K_i being $K + i \bmod 2^k$ where k is the key-length, or $K \oplus \Delta_i$ where $\Delta_1, \Delta_2, \ldots$ are known values. We ask what a related-key attack might mean in general, and how one might model it and define a corresponding notion of security.

Motivation for definitions. There is significant value in capturing block-cipher security goals via formal definitions of security. It provides cryptanalysts with clear attack models, and it enables theorists to prove the security of block-cipher based constructs. The best example to date is the pseudorandom permutation (PRP) model for a block cipher [23, 2] which has been instrumental in both these ways. We seek something similar with regard to RKAs.

Definition. We propose an extension of the notion of a PRP. Let $E\colon\ \mathcal{K} \times \mathcal{D} \to \mathcal{D}$ be the block cipher whose security we are trying to capture. We allow the adversary A to make *related-key oracle* queries consisting of a *related-key-deriving (RKD) function* $\phi\colon\ \mathcal{K} \to \mathcal{K}$ and a point $x \in \mathcal{D}$. It is placed in one of two "worlds." In world 1, a key K is chosen at random from $\mathcal{K}$, and query (ϕ, x) is answered by $E(\phi(K), x)$. In world 0, a key K is again chosen at random from $\mathcal{K}$ but we also choose at random, for each key $L \in \mathcal{K}$, a permutation $G(L, \cdot)\colon\ \mathcal{D} \to \mathcal{D}$, and the query is answered by $G(\phi(K), x)$. The advantage of A is the difference between the probabilities that it returns 1 in the two worlds. For any set Φ of functions mapping $\mathcal{K}$ to $\mathcal{K}$, we say that E is secure against Φ-restricted RKAs if the advantage of an adversary of limited resources, restricted to drawing the RKD functions in its oracle queries from Φ, is small. See Section 3 for formal definitions.

Φ as a parameter. An important definitional choice above was to have made the set Φ of allowed RKD functions a parameter of the definition rather than, say, letting Φ be the set of all functions from $\mathcal{K}$ to $\mathcal{K}$. The reason is that the power of attacks depends significantly on the types of related-key-deriving functions the adversary uses. (In particular we will see that security when Φ is the set

of all functions, or even just all permutations, is impossible to achieve.) The main question is thus for what classes Φ of RKD functions security against Φ-restricted RKAs is achievable.

CANONICAL CLASSES OF RKD FUNCTIONS. Examples of Φ, corresponding to the example attacks discussed above, include Φ_k^+, the set of functions $K \mapsto K + i \bmod 2^k$ for $0 \leq i < 2^k$, and $\Phi_k^{\oplus}$, the set of functions $K \mapsto K \oplus \Delta$ for $\Delta \in \{0,1\}^k$, where $\mathcal{K} = \{0,1\}^k$. These classes are important because security against Φ-restricted RKAs appears not only to be a design target but is useful in applications, and hence we will pay extra attention to these classes.

DES: A TEST CASE. The goal of resisting related-key attack seems to have been a design target only relatively recently. It is well-known that DES, due to its complementation property ($\overline{\mathrm{DES}_K(P)} = \mathrm{DES}_{\overline{K}}(\overline{P})$ for all keys K and plaintexts P) is insecure against related-key attacks. It is worth noting that our model and definition capture this. One can design an adversary that, in just two oracle queries, wins the game outlined above with advantage almost 1, as long as Φ contains the identity function and the map $K \mapsto \overline{K}$. In other words, DES is insecure against Φ-restricted RKAs for any such Φ.

IMPOSSIBILITY RESULTS. We show that there are inherent limitations to the security one can achieve against related-key attacks. Namely, we identify some relatively simple classes Φ of RKD functions such that for any block cipher E, there exist successful Φ-restricted RKAs against E. This means it is impossible to design a block cipher to resist these attacks.

This is relatively easy to see when Φ includes a non-injective function such as a constant function (cf. Proposition 4.1). One would expect better, however, if Φ consists only of permutations on the key space, because the result of applying a permutation to a random key is itself a random key. However, Proposition 4.2 identifies small, simple classes of permutations Φ for which we can present successful Φ-restricted RKAs on any block cipher, and Proposition 4.3 shows that there are successful $(\Phi_k^+ \cup \Phi_k^{\oplus})$-restricted RKAs on almost any block cipher of key-length k. (That is, it is impossible to design a block-cipher that is simultaneously resistant to the two basic classes of RKAs that we noted above.) Furthermore, in the last two cases, our attacks not only break the pseudorandomness of the block cipher, but are stronger in that they recover the target key.

THE NEED FOR POSSIBILITY RESULTS. Block-cipher designers seem to believe that security against Φ_k^+ and $\Phi_k^{\oplus}$-restricted RKAs is achievable. Nothing above contradicts this, but the unachievability of security against the closely related class of $(\Phi_k^+ \cup \Phi_k^{\oplus})$-restricted RKAs leads us to desire better evidence of the achievability of security against Φ-restricted RKAs on these classes, as well as other classes, than the mere inability to find attacks as above.

However, while unachievability of a security goal can be conclusively established via attacks as above, it is harder to find ways of gauging achievability that are better than merely saying that we have not found attacks. Our approach is based on the thesis that the minimal requirement for a block-cipher security goal to be considered feasible is that it should be *provably* achievable for an *ideal* (ie. Shannon) cipher. (We may not, in practice, be able to realize all properties of an ideal cipher in a real cipher, but certainly we should be wary of targeting goals that are *not* achieved by ideal ciphers, and thus it is a good idea to ensure that goals we target are at least achieved by ideal ciphers.) Accordingly, we seek to determine classes Φ of RKD functions for which we can prove that ideal ciphers resist Φ-restricted RKAs.

A GENERAL POSSIBILITY RESULT. We define two properties of a class Φ of RKD functions that we call *collision-resistance* and *output-unpredictability*. Theorem 6.3 then shows that an ideal cipher is secure against Φ-restricted RKAs for any Φ having these two properties. We consider this the main result of the paper.

The properties themselves are fairly simple. Roughly, collision-resistance (cf. Definition 5.2)

asks that for any small subset P of Φ, the probability, over a random choice of key K, that there exist distinct $\phi_1, \phi_2 \in P$ with $\phi_1(K) = \phi_2(K)$, is small. Output-unpredictability (cf. Definition 5.1) asks that for any small subset P of Φ and any small subset X of the key-space, the probability, over a random choice of key K, that there exists $\phi \in P$ with $\phi(K) \in X$, is small. The actual definitions and results in Section 6 are quantitative, upper bounding the advantage of a related-key attack in terms of advantages with respect to the two underlying properties of Φ.

Lemma 5.3 says that any Φ consisting only of permutations has the output-unpredictability property, so that in this (common) case, collision-resistance of Φ alone suffices for an ideal cipher to be secure against Φ-restricted RKAs.

Implications. A corollary of these results is that an ideal cipher is secure against Φ-restricted related-key attacks both when $\Phi = \Phi_k^+$ and when $\Phi = \Phi_k^\oplus$. Corollary 6.5 establishes this by showing that these two sets of related-key-deriving permutations have the collision-resistance property and then applying our main result. (We clarify that this does not contradict the impossibility result of Proposition 4.3 since in the latter the adversary could use RKD functions from both classes in its attack, and in the current possibility result it can use RKD functions from one or the other, but not both simultaneously.)

Applications. One consequence of having a notion of security for block ciphers with respect to RKAs is that we can now prove the security of protocols that use a block cipher with multiple, but related, keys. The proofs are standard reductions that assume that the underlying block cipher resists Φ-restricted RKAs for some suitable set of RKD functions Φ. An important point is that because Φ is a parameter of our definitions, and because different applications use keys with different relationships, these proofs precisely identify what assumptions we are making about the underlying block cipher. When Φ is some small set (eg. with two or three elements) or when Φ is some set whose RKA-resistance is commonly targeted as a design goal (eg. $\Phi_k^\oplus$), then we may have reasonable confidence that the protocol is secure. We now discuss some specific results in this vein.

Tweakable block ciphers. Liskov, Rivest and Wagner [22] introduce the notion of a tweakable block cipher and argue that use of this primitive enables conceptually simpler designs and proofs of security for modes of operation. They suggest a simple way to construct a tweakable block cipher out of a block cipher resistant to related-key attacks: simply xor the tweak into the key. Having no definitions for security against related-key attack, however, they are not able to prove the security of their construction. As an application of our notions, we prove that their construct yields a secure tweakable PRP under the assumption that the original block cipher is a PRP resistant to $\Phi_k^\oplus$-restricted related-key attacks.

Simplifying constructs. Some block-cipher based schemes such as Black and Rogaway's three-key CBC MAC constructions [7] use several independent block-cipher keys. In such schemes it is possible to use related keys instead and thereby both reduce the key-length of the scheme and conceptually simplify it. We present related-key using modifications of these schemes and prove that they retain their security if the block cipher is assumed to be a PRP secure against Φ-restricted related key attacks, where Φ is some fixed three-element subset of Φ_k^+ or $\Phi_k^\oplus$ (eg. $\{K \mapsto K, K \mapsto K + 1 \bmod 2^k, K \mapsto K + 2 \bmod 2^k\}$).

Analysis of legacy protocols. Constructions using related keys also show up in existing cryptographic applications. (For example, [17] mentions a proprietary application that uses different, related keys to encrypt different messages.) Our notions can be used to retroactively analyze such protocols, thus providing formal justification for those protocols in the case they are secure, or insights into their insecurity if they are not secure.

EXTENSIONS. The first part of the paper focuses on PRPs secure against chosen-plaintext RKAs, since this is the simplest goal related to the question of the security of block ciphers under RKAs. Later in the paper we provide definitions for PRPs secure against chosen-ciphertext RKAs, and also for PRFs secure against RKAs, and discuss how our results extend to them. It is straightforward to extend our definitions to encryption schemes and MACs secure against RKAs.

TOWARDS CONSTRUCTS. The central theoretical question raised by this work is whether it is possible, for some non-trivial classes Φ, to construct PRPs or PRFs that are provably secure against Φ-restricted related-key attacks under some standard assumption, such as the existence of one-way functions or the hardness of an algebraic problem like factoring or Decision-Diffie-Hellman (DDH). Related-key attacks are so different from standard ones that this appears to be a challenging problem.

Towards this, we note in Proposition 9.1 that it is possible to solve this problem for some very simple classes Φ, such as if Φ consists of functions that modify only the second half of their input key. In that case, we show how a standard PRP can be modified to be provably resistant to Φ-restricted related-key attack. This is already of some interest for applications, since an example of a class Φ meeting the desired condition is the subset of $\Phi_k^{\oplus}$ given by the set of all maps $K \mapsto K \oplus \Delta$ where $\Delta = 0^{k/2} \| \Delta'$ and Δ' is any $k/2$-bit string. However, we would like such results for broader classes Φ like $\Phi_k^{\oplus}$ or Φ_k^{+}.

A natural approach is to examine existing proven-secure constructions of PRFs and PRPs and see whether they resist related-key attacks. In this regard, we note that although Luby and Rackoff proved that a three-round Feistel network with independent round keys and a PRF-secure round function is a secure pseudorandom permutation in the standard model [23], any Feistel network (regardless of the number of rounds) with independent round keys is not resistant to $\Phi_k^{\oplus}$-restricted related-key attacks. We then look at DDH-based PRF constructions such as those of Naor-Reingold [25] and Nielsen [27] and show that they succumb to related-key attacks restricted to trivial classes Φ. (We stress that these constructs were never designed with the goal or claim of resisting any kind of related-key attack, so the attacks we present do not contradict their provable-security. However, in the search for constructs secure against related-key attacks it makes sense to analyze existing constructs and learn from how they fail in the new model.)

DISCUSSION. Whether to accept these new notions of pseudorandomness may be controversial since they are certainly stronger than the standard notions. But we hope this work will stimulate more interest in the continued analysis of the security of block ciphers against related-key attacks, in the design and analysis of protocols using related keys, and, from a foundational perspective, in proving the existence, based on standard assumptions, of PRFs secure against Φ-restricted RKAs for non-trivial classes Φ.

VERSIONS. An extended abstract of this paper appeared in [3]. This is the full version.

RELATED WORK. Prior to the work of Courtois and Pieprzyk [8], the best (in terms of the number of rounds) known attack against AES was a $\Phi_k^{\oplus}$-restricted related key attack that uses 256 different related keys and that extends through nine (out of 14) rounds of AES with 128-bit blocks and 256-bit keys [11]; Biham, Dunkelman, and Keller recently discovered a related-key attack against nine (out of 12) rounds of AES with 128-bit blocks and 192-bit keys [6]. Daemen and Rijmen discuss related-key attacks in their book [10] and in their AES submission documents [9] and comment that the diffusion and non-linearity of the AES key schedule makes it difficult for related-key attacks to pass through the entire cipher. In [17] Kelsey, Schneier, and Wagner give a related-key key-recovery attack against 3DES (or 3AES) using resources roughly that of an exhaustive search for a single DES (or AES) key.

Following the extended abstract of this paper [3], Iwata and Kohno [13] used our definitions to

prove that the 3GPP $f8$ encryption scheme and the 3GPP $f9$ MAC [1] are secure if the underlying block cipher is a PRP secure against Φ-restricted related key attacks, where Φ is a two-element subset of $\Phi_k^{\oplus}$; the results in [13] clarify the assumptions necessary for $f8$ and $f9$ to be secure, which is significant since Iwata and Kurosawa [15] previously showed that it is impossible to prove $f8$ and $f9$ secure under the standard PRP assumption. Jaulmes and Lercier [16] introduce a new randomized MAC, FRMAC, and prove FRMAC secure assuming that the underlying block cipher is a PRP secure against $\Phi_k^{\oplus}$-restricted related key attacks. To complement our construction in Proposition 9.1, Lucks [24] presents another method for creating a PRP provably secure against Φ-restricted related key attacks from a standard PRP where, as with our construction, Φ consists of functions that modify only part of their input key. Lucks also proposes two new PRF constructions and proves, under new hardness assumptions, that his constructions are secure against Φ-restricted related key attacks for certain sets Φ that contain RKD functions that modify the entire underlying key.

2 Notation and standard definitions

We denote by $s \xleftarrow{\$} S$ the operation of selecting s at random from set S and by $x \leftarrow y$ the assignment of value y to x. If S is a set then $|S|$ denotes its size, while if s is a string then $|s|$ denotes its length.

PRFs were introduced by [12] and PRPs by [23]. We recall the latter, but since our goal is to model block ciphers, we adopt the concrete approach of [2] rather than the asymptotic approach of the original papers. Let $\mathrm{Perm}(\mathcal{D})$ denote the set of all permutations on $\mathcal{D}$ and let $\mathrm{Perm}(l)$ be shorthand for $\mathrm{Perm}(\{0,1\}^l)$. Let $F\colon\ \mathcal{K} \times \mathcal{D} \to \mathcal{R}$ be a family of functions from $\mathcal{D}$ to $\mathcal{R}$ indexed by keys $\mathcal{K}$. We use $F_K(D)$ as shorthand for $F(K, D)$. F is a family of permutations (ie. a block-cipher), if $\mathcal{D} = \mathcal{R}$ and $F_K(\cdot)$ is a permutation on $\mathcal{D}$ for each $K \in \mathcal{K}$. If F is a family of permutations, we use $F_K^{-1}(\cdot)$ to denote the inverse of $F_K(\cdot)$ and we use $F^{-1}(\cdot,\cdot)$ to denote the function that takes as input (K, D) and computes $F_K^{-1}(D)$.

Suppose $E\colon\ \mathcal{K} \times \mathcal{D} \to \mathcal{D}$ is a family of functions. If A is an adversary with access to an oracle, we let

$$\mathbf{Adv}^{\mathrm{prp}}_E(A) \ = \ \Pr\left[\, K \xleftarrow{\$} \mathcal{K} \ :\ A^{E_K(\cdot)} = 1 \,\right] - \Pr\left[\, g \xleftarrow{\$} \mathrm{Perm}(\mathcal{D}) \ :\ A^{g(\cdot)} = 1 \,\right]$$

denote the *prp-advantage* of A in attacking E. Under this concrete security approach [2], there is no formal definition of what it means for E to be a "secure PRP," but in discussions this phrase should be taken to mean that, for any A attacking E with resources (running time, size of code, number of oracle queries) limited to "practical" amounts, the prp-advantage of A is "small." Formal results are stated with concrete bounds.

The above is a definition for pseudorandom permutations under chosen-plaintext attack. In Section 7 we shall recall the definition of pseudorandom functions and in Section 8 we shall recall the definition of pseudorandom permutations under chosen-ciphertext attacks.

3 New notions

In this section we introduce our formalizations for capturing the security of block ciphers under related-key attacks.

We let $\mathrm{Perm}(\mathcal{K}, \mathcal{D})$ denote the set of all block-ciphers with domain $\mathcal{D}$ and key-space $\mathcal{K}$. Thus the notation $G \xleftarrow{\$} \mathrm{Perm}(\mathcal{K}, \mathcal{D})$ corresponds to selecting a random block-cipher. In more detail, it comes down to defining G via

$$\text{For each } K \in \mathcal{K} \text{ do:}\ \ G_K \xleftarrow{\$} \mathrm{Perm}(\mathcal{D})\ .$$

Let $\mathrm{Perm}(k,l)$ be shorthand for $\mathrm{Perm}(\{0,1\}^k, \{0,1\}^l)$. Given a family of functions $E\colon\ \mathcal{K}\times\mathcal{D}\to\mathcal{D}$ and a key $K\in\mathcal{K}$, we define the *related-key oracle* $E_{\mathrm{RK}(\cdot,K)}(\cdot)$ as an oracle that takes two arguments, a function $\phi\colon\ \mathcal{K}\to\mathcal{K}$ and an element $M\in\mathcal{D}$, and that returns $E_{\phi(K)}(M)$. In pseudocode,

Oracle $E_{\mathrm{RK}(\phi,K)}(M)$ // where $\phi\colon\ \mathcal{K}\to\mathcal{K}$ is a function and $M\in\mathcal{D}$
$K' \leftarrow \phi(K)$; $C \leftarrow E_{K'}(M)$
Return C

We shall refer to ϕ as a *related-key-deriving (RKD)* function or a *key transformation*. We let Φ be a set of functions mapping $\mathcal{K}$ to $\mathcal{K}$. We call Φ the set of *allowed RKD functions*, or *allowed key-transformations*, and it will be a parameter of our definition.

Definition 3.1 [Pseudorandomness with respect to related-key attacks.] Let $E\colon\ \mathcal{K}\times\mathcal{D}\to\mathcal{D}$ be a family of functions and let Φ be a set of RKD functions over $\mathcal{K}$. Let A be an adversary with access to a related-key oracle, and restricted to queries of the form (ϕ,x) in which $\phi\in\Phi$ and $x\in\mathcal{D}$. Then

$$\begin{aligned}\mathbf{Adv}^{\text{prp-rka}}_{\Phi,E}(A) &= \Pr\left[K \xleftarrow{\$} \mathcal{K} \ :\ A^{E_{\mathrm{RK}(\cdot,K)}(\cdot)}=1\right]\\ &\quad -\Pr\left[K \xleftarrow{\$} \mathcal{K}\,;\, G \xleftarrow{\$} \mathrm{Perm}(\mathcal{K},\mathcal{D}) \ :\ A^{G_{\mathrm{RK}(\cdot,K)}(\cdot)}=1\right].\end{aligned}$$

is defined as the *prp-rka-advantage* of A in a Φ-restricted related-key attack (RKA) on E. ▌

The attack model allows the adversary A to choose a function ϕ which transforms the target key K into the key $\phi(K)$, and then to obtain the value of the block cipher, on an input of A's choice, under this transformed key. We measure its success at determining whether its oracle queries are being answered via the block cipher E or via a random block cipher.

One might think that the appropriate definition would be to allow the adversary to choose any related-key-deriving functions for its queries rather than restrict them to a set Φ which parameterizes the definition, but as we will see later, without a restriction, security is simply impossible, and thus the most interesting questions pertain to the manner in which security behaves as a function of Φ. Furthermore, making Φ a parameter enables us, when proving the security of a construct that uses related keys (see Section 7), to make assumptions only about the security of the given block cipher under Φ-restricted related-key attacks for some specific Φ.

Remark 3.2 [Concrete security versus asymptotics] Since our goal is to model block ciphers, our definition uses the concrete security approach rather than the asymptotic approach. Under the concrete security approach there is no formal definition of what it means for E to be a "secure PRP under Φ-restricted related-key attack," but in discussions, this phrase should be taken to mean that for any A attacking E with resources (running time, size of code, number of oracle queries) limited to "practical" amounts, and obeying the restriction that the related-key deriving functions in all its oracle queries are from the set Φ, the prp-rka-advantage of A is "small." We remark that for other considerations, such as the design of RKA-secure PRPs based on complexity-assumptions, an asymptotic definition is likely to be more appropriate, but it is trivial to extend our definitions to asymptotic ones. One would consider families of functions indexed by a security parameter, and families of RKD functions, also indexed by the same security parameter. Then one would view the advantage above as function of this security parameter, and ask that it be negligible for all polynomial-time adversaries. ▌

The following proposition shows that the notion of pseudorandomness under related-key attacks is stronger than the standard notion of pseudorandomness, assuming that the set of RKD functions Φ includes any permutation on the key space. As a special case, this proposition shows that if Φ contains the identity permutation and if a block cipher is secure against Φ-restricted RKAs, then it is also secure under the standard notion of pseudorandomness. (Furthermore, the RKA notion and the standard notion are equivalent when $|\Phi| = 1$ and the function in Φ is a permutation.)

Proposition 3.3 *Let $E\colon\ \mathcal{K} \times \mathcal{D} \to \mathcal{D}$ be any block cipher, and let Φ be any set of RKD functions over $\mathcal{K}$ that contains at least one permutation. Then given any PRP adversary A against E, we can construct a Φ-restricted RKA adversary B_A against E such that*

$$\mathbf{Adv}^{prp}_{E}(A) \le \mathbf{Adv}^{prp\text{-}rka}_{\Phi,E}(B_A)$$

and adversary B_A uses the same resources as adversary A. ∎

Proof of Proposition 3.3: Let $\phi \in \Phi$ be a permutation on $\mathcal{K}$. Let B_A be an adversary that runs A and, when A makes an oracle query M, B_A makes oracle query (ϕ, M) and returns the response to A. The equality

$$\Pr\left[\, K \xleftarrow{\$} \mathcal{K} \ :\ A^{E_K(\cdot)} = 1 \,\right] = \Pr\left[\, K \xleftarrow{\$} \mathcal{K} \ :\ B_A^{E_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \,\right]$$

holds because ϕ is a permutation on $\mathcal{K}$ and, therefore, in both experiments A is given oracle access to E with a randomly selected key. The equality

$$\Pr\left[\, g \xleftarrow{\$} \mathrm{Perm}(\mathcal{D}) \ :\ A^{g(\cdot)} = 1 \,\right] = \Pr\left[\, K \xleftarrow{\$} \mathcal{K}\,;\ G \xleftarrow{\$} \mathrm{Perm}(\mathcal{K},\mathcal{D}) \ :\ B_A^{G_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \,\right]$$

holds because, in both experiments A is given access to a random permutation on $\mathcal{D}$. The proposition follows. ∎

Definition 3.1 is for pseudorandom permutations under chosen-plaintext related-key attack. It is straight forward to extend this definition to pseudorandom functions under related-key attack, and also to pseudorandom permutations under chosen-ciphertext related-key attack. For simplicity, we stick for the bulk of the paper to the basic notion of Definition 3.1, but shall discuss these other notions in Section 8.

Since we shall often consider XOR and additive differences on k-bit keys, we give the corresponding classes of RKD functions special names. Let $\mathcal{K} = \{0,1\}^k$ where $k \ge 1$ is an integer. For any integer i with $0 \le i < 2^k$ we let $\mathrm{ADD}_i\colon\ \mathcal{K} \to \mathcal{K}$ denote the function which on input K returns $K + i \bmod 2^k$. (Here K is first interpreted as an integer and then the final result is interpreted as a k-bit string.) For any $\Delta \in \{0,1\}^k$ we let $\mathrm{XOR}_\Delta\colon\ \mathcal{K} \to \mathcal{K}$ denote the function which on input K returns $K \oplus \Delta$. Then we let

$$\Phi^{+}_k \ =\ \{\, \mathrm{ADD}_i \ :\ 0 \le i < 2^k \,\} \qquad \text{and} \qquad \Phi^{\oplus}_k \ =\ \{\, \mathrm{XOR}_\Delta \ :\ \Delta \in \{0,1\}^k \,\}\ .$$

4 Impossibility results

There are inherent limitations to security against related-key attacks. We show here that there exist relatively simple sets of RKD functions Φ over $\mathcal{K}$ such that no block cipher $E\colon\ \mathcal{K} \times \mathcal{D} \to \mathcal{D}$ (with $|\mathcal{D}|$ sufficiently large) can resist Φ-restricted related-key attacks. (One consequence of this is that it is impossible to design a block cipher that resists Φ-restricted related-key attacks for all Φ.) The first and obvious example is when Φ contains a constant function.

Proposition 4.1 *Let Φ be any class of RKD functions that contains a constant function. (Meaning there exists a $C \in \mathcal{K}$ such that Φ contains the function ϕ defined by $\phi(K) = C$ for all $K \in \mathcal{K}$.) Let*

$E\colon\ \mathcal{K} \times \mathcal{D} \to \mathcal{D}$ be any block cipher. Then there exists an adversary A such that

$$\mathbf{Adv}^{\textit{prp-rka}}_{\Phi,E}(A) \ \geq\ 1 - \frac{1}{|\mathcal{D}|}\ ,$$

and A makes only one oracle query and has running time that of one computation of E. ∎

Proof of Proposition 4.1: Let A be an adversary that first queries its related-key oracle with (ϕ, M) for some $M \in \mathcal{D}$ (where ϕ is the function described in the proposition statement: ie. ϕ maps all keys to some constant $C \in \mathcal{K}$). Let R be the result of that oracle query. The adversary A then computes $R' = E_C(M)$. If $R = R'$ the adversary A returns 1; otherwise A returns 0. Clearly

$$\Pr\left[\ K \xleftarrow{\$} \mathcal{K}\ :\ A^{E_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1\ \right] = 1$$

since the related-key oracle will return $E_{\phi(K)}(M) = E_C(M)$. Furthermore

$$\Pr\left[\ K \xleftarrow{\$} \mathcal{K}\,;\ G \xleftarrow{\$} \mathrm{Perm}(\mathcal{K},\mathcal{D})\ :\ A^{G_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1\ \right] = 1/|\mathcal{D}|$$

since G is a randomly selected family of permutations and $G_C(M)$ takes on each element in $\mathcal{D}$ with equal probability. The proposition follows. To reduce the error further, an adversary could repeat the above process but with additional, different $M' \in \mathcal{D}$. ∎

One might expect better if Φ consists only of permutations (since the result of applying a permutation to a random key is again a random key). The following indicates, however, that there are simple sets Φ of permutations on $\mathcal{K}$ such that there exist Φ-restricted related-key attacks against any block cipher.

Proposition 4.2 *Let $E\colon\ \{0,1\}^k \times \mathcal{D} \to \mathcal{D}$ be any block cipher. Then there exists an adversary A and a set of RKD functions Φ such that Φ consists only of permutations on $\{0,1\}^k$ and*

$$\mathbf{Adv}^{\textit{prp-rka}}_{\Phi,E}(A) \geq 1 - \frac{k+1}{|\mathcal{D}|}\ ,$$

and A makes $2k+1$ oracle queries (using $2k+1$ different key transformations) and has running time $O(k)$ plus the time for one computation of E. ∎

Proof of Proposition 4.2: Let $\Phi = \{\ \phi_i^c\ :\ c \in \{0,1\}, i \in \{1,\ldots,k\}\ \} \cup \{\mathrm{id}\}$ where id is the identity function, $\phi_i^0(K)$ maps K to K if the i-th bit of K is 0 and complements all but the i-th bit of K if the i-th bit of K is 1, and $\phi_i^1(K)$ maps K to K if the i-th bit of K is 1 and complements all but the i-th bit of K if the i-th bit of K is 0. Note that all the functions in Φ are permutations on $\{0,1\}^k$. Let D be any element in $\mathcal{D}$. Let A be defined as follows:

Adversary $A^{f_{\mathrm{RK}(\cdot,K)}(\cdot)}$
  $R_{\mathrm{id}} \leftarrow f_{\mathrm{RK}(\mathrm{id},K)}(D)$
  For $i = 1$ to k do
    $R_0 \leftarrow f_{\mathrm{RK}(\phi_i^0,K)}(D)\ ;\ R_1 \leftarrow f_{\mathrm{RK}(\phi_i^1,K)}(D)$
    If $R_0 = R_1$ then return 1 and halt
    If $R_0 = R_{\mathrm{id}}$ then $b_i \leftarrow 0$ else $b_i \leftarrow 1$
  $K' \leftarrow b_k \| \cdots \| b_1$ // $K' = K$
  If $E_{K'}(D) = f_{\mathrm{RK}(\mathrm{id},K)}(D)$ then return 1 else return 0

To see that

$$\Pr\left[\ K \xleftarrow{\$} \mathcal{K}\ :\ A^{E_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1\ \right] = 1$$

note that if $E_{\phi_i^0(K)}(D) = E_{\phi_i^1(K)}(D)$ for any index i, then A always returns 1. If this event does not occur, then the new key K' will equal K and, therefore, $E_{K'}(D) = f_{\mathrm{RK}(\mathrm{id},K)}(D)$ and A will return 1. Additionally,

$$\Pr\left[\, K \xleftarrow{\$} \mathcal{K}\,;\, G \xleftarrow{\$} \mathrm{Perm}(\mathcal{K},\mathcal{D}) \;:\; A^{G_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \,\right] \leq (k+1)/|\mathcal{D}|$$

since G is a random family of permutations and therefore, for each $i \in \{1,\ldots,k\}$, the probability that $G_{\phi_i^0(K)}(D) = G_{\phi_i^1(K)}(D)$ is $1/|D|$ and the probability that $G_K(D) = E_K(D)$ is also $1/|D|$. The proposition follows. As with Proposition 4.1, the error can be reduced by performing the above attack with additional points in $\mathcal{D}$. ▌

While one might consider the above attack somewhat artificial, we remark that there exist other, more natural sets Φ of permutations on $\mathcal{K}$ such that an adversary can mount a Φ-restricted related-key attack against most block ciphers $E\colon\ \mathcal{K}\times\mathcal{D}\to\mathcal{D}$. Namely we will show this for $\mathcal{K} = \{0,1\}^k$ and $\Phi = \Phi_k^+ \cup \Phi_k^{\oplus}$. To state the result we first need a definition. If $E\colon\ \{0,1\}^k\times\mathcal{D}\to\mathcal{D}$ is a block cipher, we let

$$\mathsf{KC}_E \;=\; \max_{L\neq M}\left\{\, \Pr\left[\, D \xleftarrow{\$} \mathcal{D} \;:\; E_L(D) = E_M(D) \,\right] \right\} .$$

The maximum is over all pairs of distinct keys in $\mathcal{K}$. Above, when we said our result applied to "most" block ciphers, we meant ones for which KC_E is assumed small. In practice this does not seem to be a restriction. (We would expect the above probability to be about $1/|D|$.) The formal result below applies to any block cipher and is stated quantitatively. From the result one sees that the advantage of the adversary is high as long as KC_E is small.

Proposition 4.3 *Let $E\colon\ \{0,1\}^k\times\mathcal{D}\to\mathcal{D}$ be any block cipher. Let $\Phi = \Phi_k^+ \cup \Phi_k^{\oplus}$. Then there exists an adversary A such that*

$$\mathbf{Adv}^{\textit{prp-rka}}_{\Phi,E}(A) \geq 1 - \frac{(k-1)\cdot \mathsf{KC}_E}{2} - \frac{2}{|\mathcal{D}|}\ ,$$

and A makes $2k-1$ oracle queries, each with a different key transformation, and has running time $O(k)$ plus the time for two computations of E. ▌

Proof of Proposition 4.3: The critical observation is that $\mathrm{XOR}_{0^{k-i}10^{i-1}}(K)$ and $\mathrm{ADD}_{2^{i-1}}(K)$ are equal if and only if the i-th bit of K is 0 or $i = k$ (for addition we assume that the most significant bit is on the left). In more detail, consider the following adversary:

Adversary $A^{f_{\mathrm{RK}(\cdot,K)}(\cdot)}$
  $D \xleftarrow{\$} \mathcal{D}$
  For $i = 1$ to $k-1$ do
    $R_i^+ \leftarrow f_{\mathrm{RK}(\mathrm{ADD}_{2^{i-1}},K)}(D)\,;\ R_i^{\oplus} \leftarrow f_{\mathrm{RK}(\mathrm{XOR}_{0^{k-i}10^{i-1}},K)}(D)$
    If $R_i^+ = R_i^{\oplus}$ then $b_i \leftarrow 0$ else $b_i \leftarrow 1$
  $K_0 \leftarrow 0\|b_{k-1}\|\cdots\|b_1\ ;\ K_1 \leftarrow 1\|b_{k-1}\|\cdots\|b_1$
  If $E_{K_0}(D) = f_{\mathrm{RK}(\mathrm{ADD}_0,K)}(D)$ or $E_{K_1}(D) = f_{\mathrm{RK}(\mathrm{ADD}_0,K)}(D)$ then return 1 else return 0

We first claim that

$$\Pr\left[\, K \xleftarrow{\$} \mathcal{K} \;:\; A^{E_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \,\right] \geq 1 - 2^{-1}\cdot(k-1)\cdot \mathsf{KC}_E \tag{1}$$

and

$$\Pr\left[\, K \xleftarrow{\$} \mathcal{K}\,;\, G \xleftarrow{\$} \mathrm{Perm}(\mathcal{K},\mathcal{D}) \;:\; A^{G_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \,\right] \leq 2/|\mathcal{D}| \tag{2}$$

from which the proposition follows. To justify Equation (1), let H_i, $i \in \{1, \ldots, k-1\}$, be the event that the i-th bit of K is 1 and $R_i^+ = R_i^{\oplus}$. When the event $H_1 \vee \cdots \vee H_{k-1}$ does not occur, the adversary learns the last $k-1$ bits of K exactly, implying that $K_0 = K$ or $K_1 = K$ and that the adversary will return 1. If the i-th bit of K is 1 and $i < k$, then $\mathrm{ADD}_{2^{i-1}}(K) \neq \mathrm{XOR}_{0^{k-i}10^{i-1}}(K)$. This latter observation, the definition of KC, and the fact that the i-th bit of K is 1 with probability 1/2, implies that $\Pr[\, H_i \,] \leq \mathsf{KC}_E/2$. To justify Equation (2) we note that f_K is a random permutation independent of the block cipher E, the keys K_0, K_1, and the element D, and that there are $|\mathcal{D}|$ equally-likely possibilities for $f_K(D)$. ▮

5 Properties of RKD transformations

The attack in Proposition 4.1 works because the adversary is able to predict the output of the function $\phi(K)$ for a random key K. And the attacks in Proposition 4.2 and Proposition 4.3 work because the adversary is able to find two different functions in Φ that sometimes produced the same output key (eg. if the i-th bit of K is 0 then $\phi_i^0(K) = \mathrm{id}(K)$ in the attack for Proposition 4.2). In this section we introduce two security notions capturing these properties. We will use these definitions in Section 6 when we present possibility results in the Shannon model.

In both security notions, we associate to a given set of RKD transformations Φ a measure of the extent to which Φ fails to have the property in question. The measure function takes resource bounds and returns a number between 0 and 1. The higher this number, the more "insecure" is Φ with regard to the property in question. We name the first property we measure *output-unpredictability*. Intuitively, a set Φ is output-unpredictable if, for all reasonably-sized sets $P \subseteq \Phi$ and $X \subseteq \mathcal{K}$, the probability, over a random choice of key K, that there exists a $\phi \in P$ and $K' \in X$ such that $\phi(K) = K'$, is small. The set Φ used in Proposition 4.1 was not output-unpredictable.

Definition 5.1 [Output-unpredictability for Φ.] Let $\mathcal{K}$ be a set of keys and let Φ be a set of RKD functions over $\mathcal{K}$. Let r, r' be positive integers. Then

$$\mathbf{InSec}^{\mathrm{up}}_{\Phi}(r, r') = \max_{P \subseteq \Phi, X \subseteq \mathcal{K}, |P| \leq r, |X| \leq r'} \left\{ \Pr\left[K \xleftarrow{\$} \mathcal{K} \,:\, \{\, \phi(K) \,:\, \phi \in P \,\} \cap X \neq \emptyset \right] \right\}$$

is defined as the (r, r')-*output-unpredictability* of Φ. ▮

We name the second property we measure *collision-resistance*. Intuitively, a set Φ is collision-resistant if, for all reasonably-sized sets $P \subseteq \Phi$, the probability, over a random choice of key K, that there exist distinct $\phi_1, \phi_2 \in P$ such that $\phi_1(K) = \phi_2(K)$, is small. The attacks in Proposition 4.2 and Proposition 4.3 both exploit collisions of this form.

Definition 5.2 [Collision resistance for Φ.] Let $\mathcal{K}$ be a set of keys and let Φ be a set of RKD functions over $\mathcal{K}$. Let r be a positive integer. Then

$$\mathbf{InSec}^{\mathrm{cr}}_{\Phi}(r) = \max_{P \subseteq \Phi, |P| \leq r} \left\{ \Pr\left[K \xleftarrow{\$} \mathcal{K} \,:\, |\{\, \phi(K) \,:\, \phi \in P \,\}| < |P| \right] \right\}$$

is defined as the r-*collision resistance* of Φ. ▮

Upper-bounding output-unpredictability and collision-resistance. The following lemma shows that if Φ contains only permutations and if the key-space is large, then output-unpredictability is assured for reasonable r, r'.

Lemma 5.3 Let Φ be a set of permutations on some keys space $\mathcal{K}$. Let r, r' be positive integers. Then

$$\mathbf{InSec}^{\mathrm{up}}_{\Phi}(r, r') \leq rr' |\mathcal{K}|^{-1} \,. \quad ▮$$

Proof of Lemma 5.3: Consider any set $P \subseteq \Phi$ and $X \subseteq \mathcal{K}$ such that $|P| \leq r$ and $|X| \leq r'$. Label the permutations in P as $\phi_1, \phi_2, \ldots, \phi_{|P|}$. Then

$$\begin{aligned}\Pr\left[K \xleftarrow{\$} \mathcal{K} \,:\, \{\, \phi(K) \,:\, \phi \in P \,\} \cap X \neq \emptyset \right] &\leq \Pr\left[K \xleftarrow{\$} \mathcal{K} \,:\, \phi_1(K) \in X \vee \cdots \vee \phi_{|P|}(K) \in X \right] \\ &\leq \sum_{i=1}^{|P|} \Pr\left[K \xleftarrow{\$} \mathcal{K} \,:\, \phi_i(K) \in X \right] \\ &= \frac{|P| \cdot |X|}{|\mathcal{K}|} \leq \frac{rr'}{|\mathcal{K}|}\end{aligned}$$

as desired. ∎

For the canonical sets of RKD functions in which we are interested, namely Φ_k^+ and $\Phi_k^\oplus$, the following lemma shows that collision-resistance is guaranteed.

Lemma 5.4 Let $\mathcal{K} = \{0,1\}^k$ and let Φ be either Φ_k^+ or $\Phi_k^\oplus$. Then for any positive integer r,

$$\mathbf{InSec}^{\mathrm{cr}}_{\Phi}(r) \;=\; 0\,. \qquad \blacksquare$$

Proof of Lemma 5.4: $\mathbf{InSec}^{\mathrm{cr}}_{\Phi_k^+}(r) = 0$ follows from the fact that if $K+i \bmod 2^k = K+j \bmod 2^k$, $0 \leq i, j < 2^k$, then $i = j$. Similarly for $\mathbf{InSec}^{\mathrm{cr}}_{\Phi_k^\oplus}(r)$. ∎

LOWER-BOUNDS. It is possible to lower-bound the insecurity of a block cipher against Φ-restricted RKAs as a function of the output-unpredictability of Φ.

We first state the following lemma, which lower-bounds the output-unpredictability of Φ_k^+ and $\Phi_k^\oplus$.

Lemma 5.5 Let $\mathcal{K} = \{0,1\}^k$ and Φ be Φ_k^+ or $\Phi_k^\oplus$. Then for any positive integers r, r', where $rr' \leq 2^k$ when Φ is Φ_k^+, and where $\lceil \lg r \rceil + \lceil \lg r' \rceil \leq k$ when Φ is $\Phi_k^\oplus$,

$$\mathbf{InSec}^{\mathrm{up}}_{\Phi}(r, r') \geq rr' 2^{-k}\,. \qquad \blacksquare$$

Proof of Lemma 5.5: Consider first $\mathbf{InSec}^{\mathrm{up}}_{\Phi_k^+}(r, r')$. We interchange between $\{0,1\}^k$ and $\mathbb{Z}_{2^k}$ in some standard way (eg. big-endian or little-endian). Let

$$X = \{\, a \cdot \lfloor 2^k / r' \rfloor \,:\, a \in \{0, 1, \ldots, r'-1\} \,\}$$

and let

$$P = \{\, \mathrm{ADD}_i \,:\, i \in \{0, 1, \ldots, r-1\} \,\}\,.$$

For each point $K' \in X$ there are r different values for K such that $K' \in \{\, \phi(K) \,:\, \phi \in P \,\}$. For any two distinct points $K', K'' \in X$, $\{K : \exists \phi \in P \text{ s.t. } \phi(K) = K'\} \cap \{K : \exists \phi \in P \text{ s.t. } \phi(K) = K''\} = \emptyset$ since $rr' \leq 2^k$ and difference between any two points in X is at least r. Together, these imply that

$$\Pr\left[K \xleftarrow{\$} \mathcal{K} \,:\, \{\, \phi(K) \,:\, \phi \in P \,\} \cap X \neq \emptyset \right] \geq rr' 2^{-k}\,.$$

The theorem statement for $\Phi = \Phi_k^+$ follows.

The case for $\Phi_k^\oplus$ is similar. Recall that for $\Phi_k^\oplus$ we assume that $\lceil \lg r \rceil + \lceil \lg r' \rceil \leq k$. If r or r' is 1, then the appropriate set P or X has a single element and the other set has distinct elements. Otherwise we let X be a set of r' keys such that the high order $\lceil \lg r' \rceil$ bits of each key is unique.

We let P be a subset of $\Phi_k^{\oplus}$ of cardinality r such that each permutation in P XORs a different value into the low order $\lceil \lg r \rceil$ bits of the key and does not modify any of the high order bits. ▮

The following proposition provides a lower bound on the insecurity of any block cipher under Φ-restricted RKAs. Let $E\colon\ \mathcal{K} \times \mathcal{D} \to \mathcal{D}$ be a block cipher and Φ a set of RKD functions over $\mathcal{K}$. These results show that if Φ is not output-unpredictable (ie. $\mathbf{InSec}^{\mathrm{up}}_{\Phi}(r, r')$ is high for reasonable r, r'), then an adversary can exploit this lack of output-unpredictability to distinguish E from a random family of permutations.

Proposition 5.6 *Let $E\colon\ \mathcal{K} \times \mathcal{D} \to \mathcal{D}$ be a block cipher, let Φ be a set of RKD functions over $\mathcal{K}$, and let $q \leq |\mathcal{D}|$ be a positive integer. Then there exists an adversary A that queries its related-key oracle with r different key transformations and q times per transformation, that performs $r'q$ offline applications of E, that runs in time $O(qr + qr')$, and that has advantage*

$$\mathbf{Adv}^{\mathit{prp\text{-}rka}}_{\Phi,E}(A) \geq \mathbf{InSec}^{\mathit{up}}_{\Phi}(r, r') - \frac{rr'(|\mathcal{D}| - q)!}{|\mathcal{D}|!}\ . \qquad ▮$$

Proof of Proposition 5.6: Let $\Phi' \subseteq \Phi$ and $X \subseteq \mathcal{K}$, $|\Phi'| \leq r$ and $|X| \leq r'$, such that

$$\Pr\left[\ K \xleftarrow{\$} \mathcal{K}\ :\ \{\ \phi(K)\ :\ \phi \in \Phi'\ \} \cap X \neq \emptyset\ \right] = \mathbf{InSec}^{\mathrm{up}}_{\Phi}(r, r')$$

(such a set exists by the definition of $\mathbf{InSec}^{\mathrm{up}}_{\Phi}(r, r')$). Let A be a related-key adversary that works as follows

Adversary $A^{f_{\mathrm{RK}(\cdot,K)}(\cdot)}$
  $d_1, \ldots, d_q \leftarrow$ distinct elements in $\mathcal{D}$
  If $\exists K' \in X, \phi \in \Phi'$ s.t. $E_{K'}(d_1) \| \cdots \| E_{K'}(d_q) = f_{\mathrm{RK}(\phi,K)}(d_1) \| \cdots \| f_{\mathrm{RK}(\phi,K)}(d_q)$ then return 1
  Else return 0

The collision can be found either by sorting or by using a hash table. Now

$$\begin{aligned}
\mathbf{InSec}^{\mathrm{up}}_{\Phi}(r, r') &\leq \Pr\left[\ K \xleftarrow{\$} \mathcal{K}\ :\ A^{E_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1\ \right] \\
&= \mathbf{Adv}^{\mathrm{prp\text{-}rka}}_{\Phi,E}(A) + \Pr\left[\ K \xleftarrow{\$} \mathcal{K}, G \xleftarrow{\$} \mathrm{Perm}(\mathcal{K}, \mathcal{D})\ :\ A^{G_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1\ \right] \\
&\leq \mathbf{Adv}^{\mathrm{prp\text{-}rka}}_{\Phi,E}(A) + \frac{|\Phi'| \cdot |X| \cdot (|\mathcal{D}| - q)!}{|\mathcal{D}|!} \leq \mathbf{Adv}^{\mathrm{prp\text{-}rka}}_{\Phi,E}(A) + \frac{rr'(|\mathcal{D}| - q)!}{|\mathcal{D}|!}
\end{aligned}$$

The first equation comes from the fact that if the randomly selected key is such that $\{\ \phi(K)\ :\ \phi \in \Phi'\ \} \cap X \neq \emptyset$, then the adversary A will return 1. The third equation comes from the fact that the oracle G will return to A random strings from a set of cardinality $|\mathcal{D}|!/(|\mathcal{D}| - q)!$ for each unique key $\phi(K)$ for $\phi \in \Phi'$. ▮

The following corollary presents a lower-bound on the insecurity of a block cipher with respect to Φ_k^{+}- and $\Phi_k^{\oplus}$-restricted related-key attacks. Note that the lower-bound increases as a function of the number of related-keys accessed by the adversary.

Corollary 5.7 *Let $E\colon\ \mathcal{K} \times \mathcal{D} \to \mathcal{D}$ be a block cipher with key-space $\mathcal{K} = \{0,1\}^k$. Let Φ be either Φ_k^{+} or $\Phi_k^{\oplus}$. Let $q \leq |\mathcal{D}|$, r, and r' be positive integers such that $rr' \leq 2^k$ when Φ is Φ_k^{+} and $\lceil \lg r \rceil + \lceil \lg r' \rceil \leq k$ when Φ is $\Phi_k^{\oplus}$, Then there exists an adversary A that runs in $O(qr + qr')$ time and that queries its oracle with r different related-key permutations and at most q times per related-key permutation, that performs $r'q$ offline applications of E, and that has advantage*

$$\mathbf{Adv}^{\mathit{prp\text{-}rka}}_{\Phi,E}(A) \geq rr'2^{-k} - \frac{rr'(|\mathcal{D}| - q)!}{|\mathcal{D}|!}\ . \ ▮$$

Proof of Corollary 5.7: Combine Lemma 5.5 and Proposition 5.6. ▌

Remark 5.8 Proposition 5.6 and Corollary 5.7 both extend to PRPs with respect to Φ-restricted chosen-ciphertext RKAs and to Φ-restricted RKAs against the pseudorandomness of function families. See Section 8. ▌

6 Possibility results: The Shannon model

In this section we show that if a set of RKD transformations Φ over $\mathcal{K}$ is both output-unpredictable and collision-resistant, then security against Φ-restricted RKAs is achievable in the Shannon model. This suggests that security against Φ-restricted RKAs for such Φ is a reasonable block cipher design goal.

THE SHANNON MODEL. We begin by extending Definition 3.1 to the Shannon model. This is easily done: we simply provide the adversary with oracles for E and E^{-1}, in both worlds, where E, the target block cipher, is chosen at random from the class of all block ciphers. Note the choice of G remains as before.

Definition 6.1 [RKA pseudorandomness in the Shannon model.] Fix sets $\mathcal{K}$ and $\mathcal{D}$ and let Φ be a set of RKD functions over $\mathcal{K}$. Let A be an adversary with access to three oracles, and restricted to queries of the form (K', x) for the first two oracles and (ϕ, x) for the last, where $K' \in \mathcal{K}$, $\phi \in \Phi$, and $x \in \mathcal{D}$. Then

$$\mathbf{Adv}^{\text{prp-rka}}_{\Phi,\mathcal{K},\mathcal{D}}(A) = \Pr\left[K \xleftarrow{\$} \mathcal{K} \,;\, E \xleftarrow{\$} \text{Perm}(\mathcal{K},\mathcal{D}) \,:\, A^{E(\cdot,\cdot),E^{-1}(\cdot,\cdot),E_{\text{RK}(\cdot,K)}(\cdot)} = 1 \right]$$
$$- \Pr\left[K \xleftarrow{\$} \mathcal{K} \,;\, E \xleftarrow{\$} \text{Perm}(\mathcal{K},\mathcal{D}) \,;\, G \xleftarrow{\$} \text{Perm}(\mathcal{K},\mathcal{D}) \,:\, A^{E(\cdot,\cdot),E^{-1}(\cdot,\cdot),G_{\text{RK}(\cdot,K)}(\cdot)} = 1 \right]$$

is defined as the *prp-rka-advantage* of A in a Φ-restricted related-key attack on a Shannon cipher with keys $\mathcal{K}$ and domain $\mathcal{D}$. ▌

Remark 6.2 The attacks in Section 4 apply in the Shannon model as well. (This is as one would expect since the attacks exploit properties of Φ and not properties of the block cipher in question.) For example, the equations in Proposition 4.1, Proposition 4.2 and Proposition 4.3 become, respectively

$$\mathbf{Adv}^{\text{prp-rka}}_{\Phi,\mathcal{K},\mathcal{D}}(A) \geq 1 - \frac{1}{|D|} \;,\; \mathbf{Adv}^{\text{prp-rka}}_{\Phi,\mathcal{K},\mathcal{D}}(A) \geq 1 - \frac{k+1}{|D|} \text{ and } \mathbf{Adv}^{\text{prp-rka}}_{\Phi,\mathcal{K},\mathcal{D}}(A) \geq 1 - \frac{k+3}{2\cdot|D|} \,. \;\blacksquare$$

POSSIBILITY RESULTS. We are now able to present our main result: if Φ is both output-unpredictable and collision-resistant, then security against Φ-restricted RKAs is a reasonable design goal for a real block cipher.

More formally, we show that the Φ-restricted prp-rka-advantage of an adversary A in the Shannon model is upper-bounded by $\mathbf{InSec}^{\text{up}}_{\Phi}(r, r')$ plus $\mathbf{InSec}^{\text{cr}}_{\Phi}(r)$ where r' is the number of different keys A queries its Shannon cipher with and r is the number of different RKD functions with which the adversary queries its related-key oracle. This implies that if $\mathbf{InSec}^{\text{up}}_{\Phi}(r, r')$ and $\mathbf{InSec}^{\text{cr}}_{\Phi}(r)$ are small, then any attack on a real block cipher that succeeds with high probability must exploit a property of the block cipher itself and not just a property of the related-key transformations Φ.

Theorem 6.3 *Fix a key space $\mathcal{K}$ and domain $\mathcal{D}$. Let Φ be a set of RKD functions over $\mathcal{K}$. Let A be a Shannon adversary that queries its first two oracles with a total of at most r' different keys and that queries its last oracle with a total of at most r different RKD functions from Φ. Then*

$$\mathbf{Adv}^{\textit{prp-rka}}_{\Phi,\mathcal{K},\mathcal{D}}(A) \;\leq\; \mathbf{InSec}^{\textit{up}}_{\Phi}(r, r') + \mathbf{InSec}^{\textit{cr}}_{\Phi}(r) \,. \quad \blacksquare$$

The proof of Theorem 6.3 is in Section 6.1. Note that this result is independent of the number of queries A performs with respect to each key (for its first two oracles) or key transformation (for the last oracle). That is, the parameters of interest are only the number of different keys with which A queries its Shannon cipher and the number of different RKD functions with which A queries its related key oracle.

Remark 6.4 Theorem 6.3 extends to PRPs with respect to Φ-restricted chosen-ciphertext RKAs and to Φ-restricted RKAs against the pseudorandomness of function families. See Section 8. ▌

The value of this general result is that one can now, given a class Φ of RKD functions, determine whether security against Φ-restricted RKAs is achievable by testing whether Φ has the collision-resistance and output-unpredictability properties. This is typically easy to do, as we saw in Section 5.

Results about the security against Φ-restricted RKAs in the Shannon model for $\Phi = \Phi_k^+$ or $\Phi = \Phi_k^\oplus$ follow. These results are important because they provide evidence that security against RKAs restricted to the classes of RKD functions that are popular targets in the block cipher community, is achievable. They also provide a quantitative indication of how well such attacks might be expected to fare.

Corollary 6.5 *Fix key-space $\mathcal{K} = \{0,1\}^k$ and domain $\mathcal{D}$. Let Φ be either Φ_k^+ or $\Phi_k^\oplus$. Then, for all Shannon prp-rka adversaries A that query their last oracle with a total of at most r different key transformations and that query their first two oracles with a total of at most r' different keys,*

$$\mathbf{Adv}^{\textit{prp-rka}}_{\Phi,\mathcal{K},\mathcal{D}}(A) \leq rr'2^{-k} \ . \ ▌$$

Proof of Corollary 6.5: Combine Lemma 5.3, Lemma 5.4, and Theorem 6.3. ▌

6.1 Proof of Theorem 6.3

Before proving Theorem 6.3, we first introduce two alternative definitions of security for sets of RKD functions Φ.

Definition 6.6 [Output Unpredictability-2] Let Φ be a set of RKD functions on the key-space $\mathcal{K}$. Let $\mathcal{UP}^p_K(\cdot)$ and $\mathcal{UP}^x(\cdot)$ be a pair of oracles. The oracle $\mathcal{UP}^p_K(\cdot)$ takes as input an element $\phi \in \Phi$ and the oracle $\mathcal{UP}^x(\cdot)$ takes as input an element $K' \in \mathcal{K}$. Neither oracle returns a value. An adversary "wins" if it queries its $\mathcal{UP}^x(\cdot)$ oracle with a key K' and if it queries its $\mathcal{UP}^p_K(\cdot)$ oracle with a function ϕ such that $\phi(K) = K'$. We define the *output-unpredictability-2-advantage* of an adversary A as

$$\mathbf{Adv}^{\mathrm{up2}}_{\Phi}(A) = \Pr\left[\, K \xleftarrow{\$} \mathcal{K} \ : \ A^{\mathcal{UP}^p_K(\cdot),\mathcal{UP}^x(\cdot)} \text{ "wins"} \,\right] \ . \ ▌$$

Definition 6.7 [Collision Resistance-2] Let Φ be a set of functions on the key-space $\mathcal{K}$. Let $\mathcal{CR}_K(\cdot)$ be an oracle that takes as input a function $\phi \in \Phi$ and that returns no value. An adversary "wins" if it queries its oracle with two distinct functions $\phi_1, \phi_2 \in \Phi$ such that $\phi_1(K) = \phi_2(K)$. We define the *collision resistance-2-advantage* of an adversary A as

$$\mathbf{Adv}^{\mathrm{cr2}}_{\Phi}(A) = \Pr\left[\, K \xleftarrow{\$} \mathcal{K} \ : \ A^{\mathcal{CR}_K(\cdot)} \text{ "wins"} \,\right] \ . \ ▌$$

The following lemmas map between our previous notions of security for Φ and these new interactive definitions (the omitted directions are obvious).

Lemma 6.8 Let Φ be a set of functions on the key space $\mathcal{K}$. For all output-unpredictability-2 adversaries A that make at most r queries to $\mathcal{UP}^p_K(\cdot)$ and at most r' queries to $\mathcal{UP}^x(\cdot)$, it is the case that

$$\mathbf{Adv}^{\mathrm{up2}}_{\Phi}(A) \leq \mathbf{InSec}^{\mathrm{up}}_{\Phi}(r, r') \ . \ \blacksquare$$

Lemma 6.9 Let Φ be a set of functions on the key space $\mathcal{K}$. For all collision resistance-2 adversaries A that make at most r oracle queries, it is the case that

$$\mathbf{Adv}^{\mathrm{cr2}}_{\Phi}(A) \leq \mathbf{InSec}^{\mathrm{cr}}_{\Phi}(r) \ . \ \blacksquare \tag{3}$$

Proof of Lemma 6.8 and Lemma 6.9: We show the proof for Lemma 6.9; the proof of Lemma 6.8 is analogous. Note that A may be randomized and may, upon each invocation, query its oracle with different functions from Φ. Let T be the set of all possible random inputs for A and, for each $t \in T$, let Φ_t be the subset of Φ corresponding to A's oracle queries when run with randomness t (recall that A makes at most r oracle queries). Let p_t be the probability that A is run with randomness t. Then by conditioning we have

$$\begin{aligned}\mathbf{Adv}^{\mathrm{cr2}}_{\Phi}(A) &= \sum_{t \in T} \Pr\left[\, K \xleftarrow{\$} \mathcal{K} \ : \ |\{\, \phi(K) \ : \ \phi \in \Phi_t \,\}| < |\Phi_t| \,\right] \cdot p_t \\ &\leq \sum_{t \in T} \mathbf{InSec}^{\mathrm{cr}}_{\Phi}(r) \cdot p_t = \mathbf{InSec}^{\mathrm{cr}}_{\Phi}(r) \cdot \sum_{t \in T} p_t = \mathbf{InSec}^{\mathrm{cr}}_{\Phi}(r)\end{aligned}$$

as desired. $\blacksquare$

Proof of Theorem 6.3: We are now in a position to prove Theorem 6.3. Let A be the adversary as specified in the theorem statement. Let $\Pr_e[\,\cdot\,]$ denote the probability in the experiment

$$K \xleftarrow{\$} \mathcal{K} \,;\ E \xleftarrow{\$} \mathrm{Perm}(\mathcal{K}, \mathcal{D}) \,;\ A^{E(\cdot,\cdot), E^{-1}(\cdot,\cdot), E_{\mathrm{RK}(\cdot,K)}(\cdot)}$$

and let $\Pr_g[\,\cdot\,]$ denote the probability in the experiment

$$K \xleftarrow{\$} \mathcal{K} \,;\ E \xleftarrow{\$} \mathrm{Perm}(\mathcal{K}, \mathcal{D}) \,;\ G \xleftarrow{\$} \mathrm{Perm}(\mathcal{K}, \mathcal{D}) \,;\ A^{E(\cdot,\cdot), E^{-1}(\cdot,\cdot), G_{\mathrm{RK}(\cdot,K)}(\cdot)} \ .$$

Let D denote the event that A queries its related-key oracle with a function ϕ and queries its Shannon cipher (in either the forward or backward directions) with a key K' such that $\phi(K) = K'$. By definition we have that

$$\mathbf{Adv}^{\text{prp-rka}}_{\Phi,\mathcal{K},\mathcal{D}}(A) = \Pr_e\left[\, A^{E(\cdot,\cdot), E^{-1}(\cdot,\cdot), E_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \,\right] - \Pr_g\left[\, A^{E(\cdot,\cdot), E^{-1}(\cdot,\cdot), G_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \,\right] \ .$$

Note that

$$\Pr_e\left[\, A^{E(\cdot,\cdot), E^{-1}(\cdot,\cdot), E_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \wedge \overline{D} \,\right] = \Pr_g\left[\, A^{E(\cdot,\cdot), E^{-1}(\cdot,\cdot), G_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \wedge \overline{D} \,\right]$$

since A's view is the same in both experiments as long as the event D does not occur. Therefore, by conditioning we have that

$$\mathbf{Adv}^{\text{prp-rka}}_{\Phi,\mathcal{K},\mathcal{D}}(A) \leq \Pr_e\left[\, A^{E(\cdot,\cdot), E^{-1}(\cdot,\cdot), E_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \wedge D \,\right] \leq \Pr_e[\, D \,] \ .$$

We upper bound $\Pr_e[\, D \,]$ as follows. Let C_A be a output-unpredictability-2 adversary that submits the keys of A's Shannon queries to its (C_A's) own $\mathcal{UP}^x(\cdot)$ oracle and that submits the functions ϕ in A's related-key queries to its (C_A's) own $\mathcal{UP}^p_K(\cdot)$ oracle. Let H_A be a collision-resistance-2 adversary that submits the functions ϕ in A's related-key queries to its (H_A's) own $\mathcal{CR}_K(\cdot)$ oracle.

For each of the keys K' that A queries its Shannon oracle with (in either the forward or backward directions), C_A and H_A reply to A's queries using an independently selected random permutation

(this permutation is randomly selected for each key K' but, for a given key K', is the same for both the forward and backward directions). They do this by picking and returning random points in $\mathcal{D}$ subject to the constraint that they always return the same point for two identical queries and that they do not deviate from the properties of a permutation (injective and surjective). Similarly, for each of the RKD functions ϕ' that A queries to its related-key oracle, C_A and H_A reply to A's queries with an independently selected random permutation.

Let E_1 denote the event that A queries its oracle with two distinct functions $\phi_1, \phi_2 \in \Phi$ such that $\phi_1(K) = \phi_2(K)$ and that A does this *before* it constructs a query that would cause D to occur (ie. before it queries its Shannon cipher with a key K' and queries its related-key oracle with an RKD function ϕ such that $\phi(K) = K'$). Let E_2 be the event $D \wedge \overline{E_1}$. Note that before the events E_1 or E_2 occur, C_A and H_A run A exactly as A should be run in the experiment

$$K \xleftarrow{\$} \mathcal{K}\ ;\ E \xleftarrow{\$} \mathrm{Perm}(\mathcal{K}, \mathcal{D})\ ;\ A^{E(\cdot,\cdot),E^{-1}(\cdot,\cdot),E_{\mathrm{RK}(\cdot,K)}(\cdot)}\ .$$

Now

$$\Pr_e[\,D\,] \le \Pr_e[\,E_1\,] + \Pr_e[\,E_2\,] \le \mathbf{Adv}^{\mathrm{cr2}}_{\Phi}(H_A) + \mathbf{Adv}^{\mathrm{up2}}_{\Phi}(C_A)$$

Applying Lemma 6.8 and Lemma 6.9 we get

$$\Pr_e[\,D\,] \le \mathbf{InSec}^{\mathrm{cr}}_{\Phi}(r) + \mathbf{InSec}^{\mathrm{up}}_{\Phi}(r, r')\ .$$

as desired. ∎

7 Applications of RKA-secure PRPs

Above we have been able to formally define a notion of security of block ciphers against Φ-restricted RKAs, and to determine for which classes Φ it is reasonable to assume security against Φ-restricted RKAs. Based on this we can approach the analysis of block cipher based constructions that use related keys with the goal of proving their security based on assumptions about the security against Φ-restricted RKAs of the underlying block cipher. As per the above we will certainly want to confine the choices of Φ to classes with low output-unpredictability and collision-resistance. But typically we do more than that. We confine our assumptions on the security of the block cipher against Φ-restricted RKAs to $\Phi = \Phi^{+}_{k}$ or $\Phi^{\oplus}_{k}$, or, even better, to small subsets of these classes.

We begin by showing how to use our new notions of security to prove the security of a tweakable block-cipher constructions suggested in [22].

Proof of security for a tweakable block cipher. In [22] Liskov, Rivest, and Wagner suggest that if a block cipher resists related key attacks, then one could construct a tweakable block cipher by XORing the tweak into the key. Here we provide formal justification for their belief.

Let us recall some definitions from [22]. A tweakable block cipher $\widetilde{E}$ is a function mapping $\{0,1\}^k \times \{0,1\}^t \times \{0,1\}^l$ to $\{0,1\}^l$. For each $K \in \{0,1\}^k$ and $T \in \{0,1\}^t$, we require that $\widetilde{E}(K, T, \cdot)$ is a permutation on $\{0,1\}^l$. We shall use $\widetilde{E}_K(\cdot,\cdot)$ as shorthand for $\widetilde{E}(K,\cdot,\cdot)$. If A is an adversary with access to one oracle, we let

$$\mathbf{Adv}^{\text{tweak-prp}}_{\widetilde{E}}(A) \;=\; \Pr\left[\, K \xleftarrow{\$} \mathcal{K} \;:\; A^{\widetilde{E}_K(\cdot,\cdot)} = 1 \,\right] - \Pr\left[\, G \xleftarrow{\$} \mathrm{Perm}(t,l) \;:\; A^{G(\cdot,\cdot)} = 1 \,\right]$$

denote the *tweak-prp-advantage* of A in attacking $\widetilde{E}$. We can now state the following theorem, namely that if $E\colon\ \{0,1\}^k \times \{0,1\}^l \to \{0,1\}^l$ is a secure block cipher under $\Phi^{\oplus}_{k}$-restricted related-key attacks, then $\widetilde{E}\colon\ \{0,1\}^k \times \{0,1\}^k \times \{0,1\}^l \to \{0,1\}^l$ defined as $\widetilde{E}_K(T, M) = E_{K \oplus T}(M)$ will be a secure tweakable block cipher.

Theorem 7.1 *Let $E : \{0,1\}^k \times \{0,1\}^l \to \{0,1\}^l$ be a block cipher and let $\widetilde{E} : \{0,1\}^k \times \{0,1\}^k \times \{0,1\}^l \to \{0,1\}^l$ be a tweakable block cipher defined as $\widetilde{E}(K,T,M) = E(K \oplus T, M)$. Then given a tweak-prp adversary A against $\widetilde{E}$ we can construct an $\Phi_k^{\oplus}$-restricted prp-rka adversary B against E such that*

$$\mathbf{Adv}^{\textit{tweak-prp}}_{\widetilde{E}}(A) \leq \mathbf{Adv}^{\textit{prp-rka}}_{\Phi_k^{\oplus},E}(B) \; .$$

If A queries its oracle with at most r tweaks and at most q times per tweak, then B runs in the same time as A and queries its oracle with at most r key transformations and at most q times per transformation. ▌

Proof of Theorem 7.1: Let adversary B be defined as:

Adversary $B^{f_{\mathrm{RK}(\cdot,K)}(\cdot)}$
  Run A, responding to A's request (T, M) as follows:
    Return $f_{\mathrm{RK}(\mathrm{XOR}_T,K)}(M)$ to A
  Until A halts returning a bit b
  Return b

The equality

$$\Pr\left[\, K \xleftarrow{\$} \mathcal{K} \;:\; A^{\widetilde{E}_K(\cdot,\cdot)} = 1 \,\right] = \Pr\left[\, K \xleftarrow{\$} \mathcal{K} \;:\; B^{E_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \,\right]$$

holds since B computes $\widetilde{E}_K$ exactly. Furthermore

$$\Pr\left[\, G \xleftarrow{\$} \mathrm{Perm}(k,l) \;:\; A^{G(\cdot,\cdot)} = 1 \,\right] = \Pr\left[\, K \xleftarrow{\$} \mathcal{K}\,;\, G' \xleftarrow{\$} \mathrm{Perm}(k,l) \;:\; B^{G'_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \,\right]$$

because for each tweak T in A's queries of G, B replies to A using an independently selected random permutation on $\{0,1\}^l$ (because $\mathbf{InSec}^{\mathrm{cr}}_{\Phi_k^{\oplus}}(|\Phi_k^{\oplus}|) = 0$ by Lemma 5.4). The theorem follows. ▌

Single-key CBC MACs for Arbitrary-Length Messages. In addition to proving the security of existing constructions (eg. the examples in [17] and the tweakable block cipher above), related-keys can also be used to reduce the number of keys in constructs that are defined to use several independent keys, thereby conceptually simplifying the designs. We present an example here.

Black and Rogaway's [7] "three-key constructions" are efficient CBC-MAC variants for messages of arbitrary bit-lengths, but, as their name indicates, use three independent block-cipher keys. We show here how to modify two of those variants (ECBC and FCBC) so as to use a single key. This is done by having a single "master" key and then using keys related to it to key the constructs of [7]. We call our new constructions ECBC′ and FCBC′. (From a pragmatic perspective, one may now wish to use OMAC [14] or CMAC [28], the recently-proposed one-key variants of one of Black and Rogaway's three-key CBC-MAC constructions. We present ECBC′ and FCBC′ primarily because they illustrate the use of RKA-secure PRPs in constructs.)

Let $E\colon\ \{0,1\}^k \times \{0,1\}^l \to \{0,1\}^l$ be the underlying block cipher and let Φ be a set of permutations on $\{0,1\}^k$ such that $|\Phi| \geq 3$ and $\mathbf{InSec}^{\mathrm{cr}}_{\Phi}(3) = 0$. (For concreteness, one can think of Φ as consisting of three functions, namely $K \mapsto K$, $K \mapsto K + 1 \bmod 2^k$, and $K \mapsto K + 2 \bmod 2^k$.) See Figure 1 for a description of the tagging algorithm; the key generation algorithm selects a random key K from $\{0,1\}^k$ and the verification algorithm is defined in the natural way.

We briefly recall the security goal for a message authentication code $\mathcal{MA} = (\mathcal{K}, \mathcal{T}, \mathcal{V})$, following [2]; $\mathcal{K}$ is the randomized key generation algorithm, $\mathcal{T}$ is the stateless and deterministic tagging algorithm, and $\mathcal{V}$ is the stateless and deterministic verification algorithm. Let $\mathcal{D}$ be the message

Algorithm $\mathsf{ECBC}'_K(M)$	Algorithm $\mathsf{FCBC}'_K(M)$
If $M \in (\{0,1\}^l)^+$	If $M \in (\{0,1\}^l)^+$
Then $K' \leftarrow \phi_2(K)$; $P \leftarrow M$	Then $K' \leftarrow \phi_2(K)$; $P \leftarrow M$
Else $K' \leftarrow \phi_3(K)$; $P \leftarrow M\|10^i$	Else $K' \leftarrow \phi_3(K)$; $P \leftarrow M\|10^i$
where $i = l - 1 - \|M\| \bmod l$	where $i = l - 1 - \|M\| \bmod n$
Parse P into l-bit blocks $P_1\| \cdots \|P_m$	Parse P into l-bit blocks $P_1\| \cdots \|P_m$
$C_0 \leftarrow 0^l$	$C_0 \leftarrow 0^l$
For $i \leftarrow 1$ to m do	For $i \leftarrow 1$ to $m-1$ do
$C_i \leftarrow E_{\phi_1(K)}(P_i \oplus C_{i-1})$	$C_i \leftarrow E_{\phi_1(K)}(P_i \oplus C_{i-1})$
Return $E_{K'}(C_m)$	Return $E_{K'}(P_m \oplus C_{m-1})$

Figure 1: The ECBC' and FCBC' MACs for arbitrary-length messages. We require that $\Phi = \{\phi_1, \phi_2, \phi_3\}$ is a set of permutations on the underlying block cipher E's key space and that $\mathbf{InSec}^{\mathrm{cr}}_{\Phi}(3) = 0$. For example $\Phi = \{\mathrm{ADD}_0, \mathrm{ADD}_1, \mathrm{ADD}_2\} \subset \Phi^+_k$.

---

space of $\mathcal{T}$. Let A be an adversary with access to a tagging oracle with a randomly selected key K. We say A *forges* if it outputs a pair $(x, \mathcal{T}_K(x))$ where $x \in \mathcal{D}$ and A never queried its oracle $\mathcal{T}_K(\cdot)$ at x. Then

$$\mathbf{Adv}^{\mathrm{mac}}_{\mathcal{MA}}(A) = \Pr\left[\, K \xleftarrow{\$} \mathcal{K} \,:\, A^{\mathcal{T}_K(\cdot)} \text{ forges} \,\right]$$

is defined as the *forging-advantage* of A. In the concrete setting, we consider a MAC secure if the advantage of all forging adversaries with reasonable resources is small.

The following theorem shows that the ECBC' and FCBC' constructions as described are secure assuming that the underlying block cipher is secure against Φ-restricted related-key attacks, where Φ is as above.

Theorem 7.2 [ECBC' **and** FCBC' **are secure MACs.**] *Fix $l \geq 1$ and $k \geq 2$ and let $N = 2^l$. Let $E\colon\ \{0,1\}^k \times \{0,1\}^l \to \{0,1\}^l$ be a block cipher. Let $\Phi = \{\phi_1, \phi_2, \phi_3\}$ be a set of permutations on $\{0,1\}^k$ such that $\mathbf{InSec}^{cr}_{\Phi}(3) = 0$. Let ECBC' and FCBC' be the constructions from Figure 1 that use E and Φ. For any adversaries A and C that make at most q oracle queries each, where each query is at most ml bits long, m an integer at most $N/4$, we can construct adversaries B and D such that*

$$\mathbf{Adv}^{mac}_{\mathsf{ECBC}'}(A) \leq \frac{2m^2q^2 + q^2 + 1}{N} + \mathbf{Adv}^{prp\text{-}rka}_{\Phi,E}(B)$$

and

$$\mathbf{Adv}^{mac}_{\mathsf{FCBC}'}(C) \leq \frac{2m^2q^2 + q^2 + 1}{N} + \mathbf{Adv}^{prp\text{-}rka}_{\Phi,E}(D)\ .$$

and B and D query their oracles at most mq times for relationship ϕ_1 and at most q times for the relationships ϕ_2 and ϕ_3. Furthermore, B and D run in approximately the same time as A and C. ∎

To interpret the above results, let us consider the case of ECBC' and let A be the adversary as specified in the theorem statement; the case for FCBC' is analogous. This theorem says is that if A is able to forge a message with probability ϵ, then we can construct a Φ-restricted RKA adversary B against E that succeeds in distinguishing E from a random family of permutations with probability $\epsilon' \geq \epsilon - (2m^2q^2 + q^2 + 1)/2^l$. If $\Phi = \{K \mapsto K, K \mapsto K + 1 \bmod 2^k, K \mapsto K + 2 \bmod 2^k\}$, it seems reasonable to assume that E is secure against Φ-restricted related-key attacks and, therefore, that

ϵ' is small for all adversaries using reasonable resources. This would imply that ϵ is also small for all adversaries A against ECBC' using reasonable resources and, therefore, that the ECBC' instantiation with this set Φ is a viable alternative to [7]'s original ECBC construction, at least under the standard model of unforgeability (see Remark 7.6 for some caveats).

PSEUDORANDOM FUNCTIONS. Before proving Theorem 7.2, we first recall the standard notion of a PRF, modified appropriately for the concrete setting [12, 2]. Let $\mathrm{Rand}(\mathcal{D},\mathcal{R})$ be the set of all functions from $\mathcal{D}$ to $\mathcal{R}$. Let $\mathrm{Rand}(l,l')$ be shorthand for $\mathrm{Rand}(\{0,1\}^l,\{0,1\}^{l'})$. Let $F\colon\ \mathcal{K}\times\mathcal{D}\to\mathcal{R}$ be a family of functions. If A is an adversary with access to an oracle, we let

$$\mathbf{Adv}^{\mathrm{prf}}_{F}(A) = \Pr\left[\, K \xleftarrow{\$} \mathcal{K} \ :\ A^{F_K(\cdot)}=1 \,\right] - \Pr\left[\, g \xleftarrow{\$} \mathrm{Rand}(\mathcal{D},\mathcal{R}) \ :\ A^{g(\cdot)}=1 \,\right]$$

denote the *prf-advantage* of A in attacking F. We say that F is a "secure PRF" if the prf-advantage of all adversaries attacking F and using reasonable resources is "small."

PROOF OF THEOREM 7.2. Let $\mathsf{ECBC}[\mathrm{Perm}(l)]$ be a variant of ECBC' that, instead of using $E_{\phi_1(K)}$, $E_{\phi_2(K)}$, and $E_{\phi_3(K)}$ as in Figure 1, uses three random and independent permutations (π_1, π_2, and π_3) on l bits. Similarly for $\mathsf{FCBC}[\mathrm{Perm}(l)]$. Let CONS' and $\mathsf{CONS}[\mathrm{Perm}(l)]$ be either ECBC' and $\mathsf{ECBC}[\mathrm{Perm}(l)]$ or FCBC' and $\mathsf{FCBC}[\mathrm{Perm}(l)]$, respectively.

Theorems 1 and 3 of [7] (Lemma 7.3 below) upper bound the probability of an adversary distinguishing $\mathsf{ECBC}[\mathrm{Perm}(l)]$ or $\mathsf{FCBC}[\mathrm{Perm}(l)]$ from a random function from $\{0,1\}^*$ to $\{0,1\}^l$.

Lemma 7.3 [Theorems 1 and 3 of [7].] Let $N=2^l$ and $l\geq 1$. Let $\mathsf{CONS}[\mathrm{Perm}(l)]$ be as described above. Let A be an adversary which asks at most q queries each of which is at most ml-bits, m an integer. Assume $m\leq N/4$. Then

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{CONS}[\mathrm{Perm}(l)]}(A) \leq \frac{(2m^2+1)q^2}{N}\ . \quad \blacksquare$$

We now present a modification of Lemma 2 of [7], adapted for use with our notion of related key pseudorandomness (Lemma 7.4). This lemma states that if $\mathsf{ECBC}[\mathrm{Perm}(l)]$ (resp., $\mathsf{FCBC}[\mathrm{Perm}(l)]$) is a secure pseudorandom function and if the underlying block cipher resists Φ-restricted related key attacks, then ECBC' (resp., FCBC') is a secure pseudorandom function.

Lemma 7.4 [Inf. Th. PRF $\to$ Comp. Th. PRF.] Let E, Φ, CONS', and $\mathsf{CONS}[\mathrm{Perm}(l)]$ be as before. Then given a PRF adversary A against CONS' we can construct a PRF adversary B against $\mathsf{CONS}[\mathrm{Perm}(l)]$ and a related-key adversary C against E such that

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{CONS}'}(A) \leq \mathbf{Adv}^{\mathrm{prf}}_{\mathsf{CONS}[\mathrm{Perm}(l)]}(B) + \mathbf{Adv}^{\text{prp-rka}}_{\Phi,E}(C)\ .$$

If A makes at most q oracle queries, each of which is at most ml bits long, m an integer, then B makes the same type of oracle queries as A and C makes at most mq oracle queries for relationship ϕ_1 and at most q oracle queries for the relationships ϕ_2 and ϕ_3. Furthermore, B and C run in approximately the same time as A. $\blacksquare$

Proof of Lemma 7.4: Let B be an adversary against the pseudorandomness of $\mathsf{CONS}[\mathrm{Perm}(l)]$ that runs A and that replies to A's oracle queries using its own oracle and that returns the same bit that A returns. Let C be a related-key adversary against E that runs A and that replies to A's oracle queries by computing CONS but using its related-key oracle for the underlying permutations.

Then

$$\begin{aligned}
\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{CONS}'}(A) &= \Pr\left[K \xleftarrow{\$} \{0,1\}^k \ : \ A^{\mathsf{CONS}'_K(\cdot)} = 1 \right] - \Pr\left[G \xleftarrow{\$} \mathrm{Rand}(\{0,1\}^*, \{0,1\}^l) \ : \ A^{G(\cdot)} = 1 \right] \\
&= \Pr\left[K \xleftarrow{\$} \{0,1\}^k \ : \ A^{\mathsf{CONS}'_K(\cdot)} = 1 \right] \\
&\quad - \Pr\left[\pi_1, \pi_2, \pi_3 \xleftarrow{\$} \mathrm{Perm}(l) \ : \ A^{\mathsf{CONS}[\mathrm{Perm}(l)]_{\pi_1,\pi_2,\pi_3}(\cdot)} = 1 \right] \\
&\quad + \Pr\left[\pi_1, \pi_2, \pi_3 \xleftarrow{\$} \mathrm{Perm}(l) \ : \ A^{\mathsf{CONS}[\mathrm{Perm}(l)]_{\pi_1,\pi_2,\pi_3}(\cdot)} = 1 \right] \\
&\quad - \Pr\left[G \xleftarrow{\$} \mathrm{Rand}(\{0,1\}^*, \{0,1\}^l) \ : \ A^{G(\cdot)} = 1 \right] .
\end{aligned}$$

Since $\mathbf{InSec}^{\mathrm{cr}}_{\Phi}(|\Phi|) = 0$, we have that

$$\begin{aligned}
&\Pr\left[\pi_1, \pi_2, \pi_3 \xleftarrow{\$} \mathrm{Perm}(l) \ : \ A^{\mathsf{CONS}[\mathrm{Perm}(l)]_{\pi_1,\pi_2,\pi_3}(\cdot)} = 1 \right] = \\
&\qquad\qquad \Pr\left[K \xleftarrow{\$} \{0,1\}^k, G \xleftarrow{\$} \mathrm{Perm}(k,l) \ : \ C^{G_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \right] .
\end{aligned}$$

Therefore

$$\begin{aligned}
\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{CONS}'}(A) &= \Pr\left[K \xleftarrow{\$} \{0,1\}^k \ : \ C^{E_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \right] \\
&\quad - \Pr\left[K \xleftarrow{\$} \{0,1\}^k, G \xleftarrow{\$} \mathrm{Perm}(k,l) \ : \ C^{G_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \right] \\
&\quad + \Pr\left[\pi_1, \pi_2, \pi_3 \xleftarrow{\$} \mathrm{Perm}(l) \ : \ B^{\mathsf{CONS}[\mathrm{Perm}(l)]_{\pi_1,\pi_2,\pi_3}(\cdot)} = 1 \right] \\
&\quad - \Pr\left[G \xleftarrow{\$} \mathrm{Rand}(\{0,1\}^*, \{0,1\}^l) \ : \ B^{G(\cdot)} = 1 \right] \\
&\leq \mathbf{Adv}^{\text{prp-rka}}_{\Phi,E}(C) + \mathbf{Adv}^{\mathrm{prf}}_{\mathsf{CONS}[\mathrm{Perm}(l)]}(B)
\end{aligned}$$

as desired. ∎

Proof of Theorem 7.2: Combine Lemma 7.3, Lemma 7.4 and the known relationship between MACs and PRFs (see [2]). ∎

Discussion. We end this section with some observations about constructs that use related keys.

Remark 7.5 If Φ' is a subset of Φ, then the insecurity of E with respect to Φ'-restricted related-key attacks can be no greater than the insecurity of E with respect to Φ-restricted related-key attacks (and may, in fact, be much smaller). Take Φ' to be $\{\mathrm{ADD}_0, \mathrm{ADD}_1, \mathrm{ADD}_2\} \subset \Phi^+_k$. While one may not wish to base the security of a protocol on the security of a block cipher against Φ^+_k-restricted related-key attacks, one may feel more comfortable basing the security of a protocol on the security of a block cipher against Φ'-restricted related-key attacks, as we did with our CBC-MAC variants. See also Corollary 5.7, which shows that the insecurity of block ciphers under Φ^+_k- or $\Phi^{\oplus}_k$-restricted related-key attacks is (essentially) lower-bounded by a birthday-like term of the form $rr'2^{-k}$ (r is the number of different related-key transformations with which an adversary queries its related-key oracle, and r' is the number of different keys K with which the attacker computes $E_K(\cdot)$ directly). ∎

Remark 7.6 Consider a construct that uses a block cipher with related keys and that is provably secure under some standard notion of security (eg, unforgeability for MACs or indistinguishability for encryption schemes) assuming that the block cipher resists Φ-restricted RKAs for some appropriate set Φ. It is important to note that even though that construct is provably secure under some

standard notion, that construct may be vulnerable to a construction-level related-key attack; this is not a contradiction since construction-level related-key attacks are outside the standard models of security for MACs and encryption schemes.

Consider, for example, the construction-level related-key attack against RMAC in [21]. Consider also a CTR mode encryption scheme that generates the keystream as follows (since we are summarizing an attack, we only describe the relevant aspects of the keystream-generation process):

$$E_{\mathrm{XOR}_{0^k}(K)}(\langle \mathsf{ctr} \rangle) \| E_{\mathrm{XOR}_{0^{k-1}1}(K)}(\langle \mathsf{ctr} \rangle) \| E_{\mathrm{XOR}_{0^k}(K)}(\langle \mathsf{ctr}+1 \rangle) \| E_{\mathrm{XOR}_{0^{k-1}1}(K)}(\langle \mathsf{ctr}+1 \rangle) \| \cdots .$$

Here we assume that $E\colon\ \{0,1\}^k \times \{0,1\}^l \to \{0,1\}^l$ is a block cipher, K is a k-bit key, ctr is some counter, initially zero, and $\langle x \rangle$ denotes an l-bit encoding of the integer x modulo 2^l. Let $\Phi = \{\mathrm{XOR}_{0^k}, \mathrm{XOR}_{0^{k-1}1}\}$. Although it is clear that one can use the above keystream generation method in a provably-secure CTR mode encryption scheme assuming that E is secure against Φ-restricted RKAs, there is a construction-level related-key attack against the above keystream generator. In particular, two keys that differ only in their last bit will produce similar keystreams, the only difference being the order of the keystream blocks. We can use this property in a construction-level related-key attack against the privacy of the encryption scheme, under a suitable definition of privacy under construction-level related-key attacks (such a notion of privacy follows naturally from the standard notions of privacy and our new formalisms for pseudorandomness under related-key attacks; it is also straightforward to extend our notions and define security for MACs against related-key attacks). As another example, the tweakable block cipher in Theorem 7.1 is vulnerable to construction-level related-key attacks. Namely, $\widetilde{E}(K,T,M) = \widetilde{E}(K \oplus X, T \oplus X, M)$ for any k-bit string X.

Whether or not construction-level related-key attacks are of a concern depends on the application in question. ▌

Remark 7.7 While most modern block ciphers, including the AES, are designed with the explicitly stated goal of resisting related-key attacks, some block cipher constructions do not resist related-key attacks (or are more vulnerable to related-key attacks than one would expect). Consider, for example, the complementation property with DES, or [17]'s attack against three-key triple DES. Developers of protocols that use related-keys should be aware of this problem and realize that some block ciphers may not be good candidates for use with their constructions. See, for example, the problems with using 3DES in RMAC [21]. ▌

8 PRFs and PRPs under chosen-ciphertext RKAs

In addition to considering pseudorandom permutations under chosen-plaintext RKAs, it is also possible to consider pseudorandom permutations under chosen-ciphertext RKAs, and pseudorandom functions under RKAs.

PRPs UNDER CHOSEN-CIPHERTEXT ATTACKS. We first recall the standard definition of a pseudorandom permutation under chosen-ciphertext attacks (aka super-pseudorandom permutations [23] or strong-PRPs [26]): Suppose $E\colon\ \mathcal{K} \times \mathcal{D} \to \mathcal{D}$ is a block cipher with domain $\mathcal{D}$ and keys $\mathcal{K}$. If A is an adversary with access to two oracles, we let

$$\mathbf{Adv}^{\text{prp-cca}}_E(A) \ =\ \Pr\left[\, K \xleftarrow{\$} \mathcal{K} \ :\ A^{E_K(\cdot),E_K^{-1}(\cdot)} = 1 \,\right] - \Pr\left[\, g \xleftarrow{\$} \mathrm{Perm}(\mathcal{D}) \ :\ A^{g(\cdot),g^{-1}(\cdot)} = 1 \,\right]$$

denote the *prp-cca-advantage* of A in attacking E. We say that E is a "secure PRP under chosen-ciphertext attacks" if the prp-cca-advantage of all adversaries attacking E and using reasonable resources is "small."

We recall the standard definition of a pseudorandom function in Section 7 before proving Theorem 7.2.

PRP-CCRKA AND PRF-RKA. Our notion of pseudorandomness for permutations under chosen-ciphertext related-key attacks is as follows:

Definition 8.1 Let $E\colon\ \mathcal{K}\times\mathcal{D}\to\mathcal{D}$ be a block cipher and let Φ be a set of RKD functions over $\mathcal{K}$. Let A be an adversary with access to two related-key oracles, and restricted to queries of the form (ϕ, x) in which $\phi\in\Phi$ and $x\in\mathcal{D}$. Then

$$\begin{aligned}\mathbf{Adv}^{\text{prp-ccrka}}_{\Phi,E}(A) = &\Pr\left[\, K\xleftarrow{\$}\mathcal{K}\ :\ A^{E_{\text{RK}(\cdot,K)}(\cdot),E^{-1}_{\text{RK}(\cdot,K)}(\cdot)}=1\,\right]\\ &-\Pr\left[\, K\xleftarrow{\$}\mathcal{K}\,;\ G\xleftarrow{\$}\text{Perm}(\mathcal{K},\mathcal{D})\ :\ A^{G_{\text{RK}(\cdot,K)}(\cdot),G^{-1}_{\text{RK}(\cdot,K)}(\cdot)}=1\,\right]\end{aligned}$$

is defined as the *prp-ccrka-advantage* of A in a Φ-restricted related-key attack on E. ▮

As usual with the concrete security approach, we say that E is a "secure PRP under Φ-restricted chosen-ciphertext related-key attack" if the prp-ccrka-advantage of all adversaries attacking E and using reasonable resources is "small."

Let $\text{Rand}(\mathcal{K},\mathcal{D},\mathcal{R})$ be the set of all families of functions with domain $\mathcal{D}$, range $\mathcal{R}$ and keys $\mathcal{K}$. Let $\text{Rand}(k,l,l')$ be shorthand for $\text{Rand}(\{0,1\}^k,\{0,1\}^l,\{0,1\}^{l'})$. Our notion of pseudorandomness for function families under related-key attacks is as follows:

Definition 8.2 Let $F\colon\ \mathcal{K}\times\mathcal{D}\to\mathcal{R}$ be a family of functions and let Φ be a set of RKD functions over $\mathcal{K}$. Let A be an adversary with access to a related-key oracle and restricted to queries of the form (ϕ, x) in which $\phi\in\Phi$ and $x\in\mathcal{D}$. Then

$$\begin{aligned}\mathbf{Adv}^{\text{prf-rka}}_{\Phi,F}(A) = &\Pr\left[\, K\xleftarrow{\$}\mathcal{K}\ :\ A^{F_{\text{RK}(\cdot,K)}(\cdot)}=1\,\right]\\ &-\Pr\left[\, K\xleftarrow{\$}\mathcal{K}, G\xleftarrow{\$}\text{Rand}(\mathcal{K},\mathcal{D},\mathcal{R})\ :\ A^{G_{\text{RK}(\cdot,K)}(\cdot)}=1\,\right]\end{aligned}$$

is defined as the *prf-rka-advantage* of A in a Φ-restricted related-key attack on F. ▮

We say that F is a "secure PRF under Φ-restricted related-key attack" if the prf-rka-advantage of all adversaries attacking E and using reasonable resources is "small."

PRP-CCRKA IMPLIES PRP-RKA. The following proposition shows that if a block cipher is secure against Φ-restricted chosen-ciphertext RKAs then it is also secure against Φ-restricted chosen-plaintext attacks.

Proposition 8.3 *Let $E\colon\ \mathcal{K}\times\mathcal{D}\to\mathcal{D}$ be any block cipher, and let Φ be any set of RKD functions over $\mathcal{K}$. Then given any Φ-restricted chosen-plaintext RKA adversary A against E, we can construct a Φ-restricted chosen-ciphertext RKA adversary B_A against E such that*

$$\mathbf{Adv}^{\textit{prp-rka}}_{\Phi,E}(A)\le\mathbf{Adv}^{\textit{prp-ccrka}}_{\Phi,E}(B_A)$$

and adversary B_A uses the same resources as adversary A. ▮

Proof of Proposition 3.3: Let B_A be an adversary that runs A and, when A makes an oracle query (ϕ, M), B_A makes a query (ϕ, M) to its first oracle and returns the response to A. The equalities

$$\Pr\left[\, K\xleftarrow{\$}\mathcal{K}\ :\ A^{E_{\text{RK}(\cdot,K)}(\cdot)}=1\,\right]=\Pr\left[\, K\xleftarrow{\$}\mathcal{K}\ :\ B_A^{E_{\text{RK}(\cdot,K)}(\cdot),E^{-1}_{\text{RK}(\cdot,K)}(\cdot)}=1\,\right]$$

and

$$\Pr\left[\, K \xleftarrow{\$} \mathcal{K}\,;\, G \xleftarrow{\$} \mathrm{Perm}(\mathcal{K},\mathcal{D}) \;:\; A^{G_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \,\right]$$
$$= \Pr\left[\, K \xleftarrow{\$} \mathcal{K}\,;\, G \xleftarrow{\$} \mathrm{Perm}(\mathcal{K},\mathcal{D}) \;:\; B_A^{G_{\mathrm{RK}(\cdot,K)}(\cdot),G^{-1}_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \,\right]$$

holds because B_A's first oracle is the same as A's oracle. The proposition follows. ▌

Impossibility results. Since the prp-ccrka notion is stronger than the prp-rka notion, the impossibility results in Section 4, as well as the lower-bounds in Proposition 5.6 and Corollary 5.7, apply in the prp-ccrka setting.

Furthermore, the attacks against block ciphers in Section 4 can be applied to function families. As with block ciphers, this means that it is impossible to design function families that resist Φ-restricted RKAs for all sets of RKD functions Φ. We can also state the following analog of Proposition 5.6 for function families. The proof is a straight-forward adaptation of the proof of Proposition 5.6 and is omitted.

Proposition 8.4 *Let $F\colon\ \mathcal{K}\times\mathcal{D}\to\mathcal{R}$ be a family of functions, let Φ be a set of RKD functions over $\mathcal{K}$, and let $q \le |\mathcal{D}|$ be a positive integer. Then there exists an adversary A that queries its related-key oracle with r different key transformations and q times per transformation, that performs $r'q$ offline applications of E, that runs in time $O(qr + qr')$, and that has advantage*

$$\mathbf{Adv}^{\mathit{prf\text{-}rka}}_{\Phi,F}(A) \ge \mathbf{InSec}^{\mathit{up}}_{\Phi}(r,r') - \frac{rr'}{|\mathcal{R}|^q}\,. \qquad ▌$$

Possibility results: The Shannon model. Theorem 6.3 can be extended to the case where the adversary is allowed a chosen-ciphertext attack without needing to increase the assumptions made on Φ. Theorem 6.3 can also be extended to pseudorandom functions. We first present the following definitions in the Shannon model.

Definition 8.5 [RKA pseudorandomness under CCA in the Shannon model.] Fix sets $\mathcal{K}$ and $\mathcal{D}$ and let Φ be a set of RKD functions over $\mathcal{K}$. Let A be an adversary with access to four oracles and restricted to queries of the form (K', x) for the first two oracles and (ϕ, x) for the last two oracles, $K' \in \mathcal{K}$, $\phi \in \Phi$, and $x \in \mathcal{D}$. Then

$$\mathbf{Adv}^{\text{prp-ccrka}}_{\Phi,\mathcal{K},\mathcal{D}}(A) \;=\; \Pr\left[\, K \xleftarrow{\$} \mathcal{K}\,;\, E \xleftarrow{\$} \mathrm{Perm}(\mathcal{K},\mathcal{D}) \;:\; A^{E(\cdot,\cdot),E^{-1}(\cdot,\cdot),E_{\mathrm{RK}(\cdot,K)}(\cdot),E^{-1}_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \,\right]$$
$$- \Pr\left[\, K \xleftarrow{\$} \mathcal{K}\,;\, E \xleftarrow{\$} \mathrm{Perm}(\mathcal{K},\mathcal{D})\,;\, G \xleftarrow{\$} \mathrm{Perm}(\mathcal{K},\mathcal{D}) \;:\; A^{E(\cdot,\cdot),E^{-1}(\cdot,\cdot),G_{\mathrm{RK}(\cdot,K)}(\cdot),G^{-1}_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \,\right]$$

is defined as the *prp-ccrka-advantage* of A in a Φ-restricted related-key attack on a Shannon cipher with keys $\mathcal{K}$ and domain $\mathcal{D}$. ▌

Definition 8.6 [RKA pseudorandom functions in the Shannon model.] Fix sets $\mathcal{K}$, $\mathcal{D}$, and $\mathcal{R}$ and let Φ be a set of RKD functions over $\mathcal{K}$. Let A be an adversary with access to two oracles and restricted to queries of the form (K', x) for the first oracle and (ϕ, x) for the second oracle, $K' \in \mathcal{K}$, $\phi \in \Phi$, and $x \in \mathcal{D}$, then

$$\mathbf{Adv}^{\text{prf-rka}}_{\Phi,\mathcal{K},\mathcal{D},\mathcal{R}}(A) = \Pr\left[\, K \xleftarrow{\$} \mathcal{K}, F \xleftarrow{\$} \mathrm{Rand}(\mathcal{K},\mathcal{D},\mathcal{R}) \;:\; A^{F(\cdot,\cdot),F_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \,\right]$$
$$- \Pr\left[\, K \xleftarrow{\$} \mathcal{K}, F \xleftarrow{\$} \mathrm{Rand}(\mathcal{K},\mathcal{D},\mathcal{R}), G \xleftarrow{\$} \mathrm{Rand}(\mathcal{K},\mathcal{D},\mathcal{R}) \;:\; A^{F(\cdot,\cdot),G_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \,\right]$$

is the *prf-rka-advantage* of A in a Φ-restricted related-key attack on a Shannon function family with keys $\mathcal{K}$, domain $\mathcal{D}$, and range $\mathcal{R}$. ▌

We now state the following upper-bound:

Theorem 8.7 *Fix a key space $\mathcal{K}$ and domain $\mathcal{D}$. Let Φ be a set of RKD functions over $\mathcal{K}$. Let A be a Shannon adversary that queries its first two oracles with a total of at most r' different keys and that queries its last two oracles with a total of at most r different RKD functions from Φ. Then*

$$\mathbf{Adv}^{\textit{prp-ccrka}}_{\Phi,\mathcal{K},\mathcal{D}}(A) \;\leq\; \mathbf{InSec}^{\textit{up}}_{\Phi}(r,r') + \mathbf{InSec}^{\textit{cr}}_{\Phi}(r)\;. \qquad \blacksquare$$

To modify the proof of Theorem 6.3, as given in Section 6.1, to prove the above, one must simply ensure that the new adversaries C_A and H_A keep track of the RKD functions used in both the forward and backward directions of the related-key oracle queries.

It is straightforward to modify the proof of Theorem 6.3 for the prf-rka Shannon setting. The resulting theorem statement reads:

Theorem 8.8 *Fix a key space $\mathcal{K}$, domain $\mathcal{D}$, and range $\mathcal{R}$. Let Φ be a set of RKD functions over $\mathcal{K}$. Let A be a Shannon adversary that queries its first oracle with a total of at most r' different keys and that queries its last oracle with a total of at most r different RKD functions from Φ. Then*

$$\mathbf{Adv}^{\textit{prf-rka}}_{\Phi,\mathcal{K},\mathcal{D},\mathcal{R}}(A) \;\leq\; \mathbf{InSec}^{\textit{up}}_{\Phi}(r,r') + \mathbf{InSec}^{\textit{cr}}_{\Phi}(r)\;. \qquad \blacksquare$$

PRF/PRP Switching. When proving the security of a block cipher-based protocol, it is often simpler to first assume that the block cipher is a family of random functions. The following proposition can be used to replace a family of random functions with a family of random permutations, or visa-versa. This result is similar to the result in Proposition 2.5 in [2] for the standard notions of pseudorandomness. We use this result in the proof of Proposition 9.1.

We first need a definition. Let Φ be a set of RKD functions over $\mathcal{K}$. Let

$$\mathsf{NM}_{\Phi} \;=\; \max_{K,K' \in \mathcal{K}} \left\{\; |\{\; \phi \in \Phi \;:\; \phi(K) = K' \;\}| \;\right\}$$

be the maximum, over all $K, K' \in \mathcal{K}$, of the number of key transformations in Φ mapping K to K'. Note that, for the special case of $\Phi = \Phi^{+}_{k}$ or $\Phi = \Phi^{\oplus}_{k}$, $\mathsf{NM}_{\Phi} = 1$. We now state the following proposition:

Proposition 8.9 ***[PRF/PRP Switching.]*** *Fix sets $\mathcal{K}$ and $\mathcal{D}$ and RKD set Φ. Let A be an adversary that queries its related-key oracle with at most r different transformations from Φ and at most q times per transformation. Then*

$$\Big|\, \Pr\Big[\; K \xleftarrow{\$} \mathcal{K}, F \xleftarrow{\$} \mathit{Rand}(\mathcal{K},\mathcal{D},\mathcal{D}) \;:\; A^{F_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \;\Big] \\ - \Pr\Big[\; K \xleftarrow{\$} \mathcal{K}, G \xleftarrow{\$} \mathit{Perm}(\mathcal{K},\mathcal{D}) \;:\; A^{G_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \;\Big] \,\Big| \;\leq\; \frac{rq'(q'-1)}{2|\mathcal{D}|}$$

where $q' = q \cdot \mathsf{NM}_{\Phi}$. $\blacksquare$

Proof of Proposition 8.9: The structure of our proof follows the structure of Bellare and Rogaway's [4] and Shoup's [29] game-based proofs of Proposition 2.5 from [2]. We begin by defining the game $\mathbf{GameS1}_{\mathcal{K},\mathcal{D}}(A)$:

Game $\mathbf{GameS1}_{\mathcal{K},\mathcal{D}}(A)$
  $\mathsf{bad} \leftarrow \mathsf{false}$
  $K \xleftarrow{\$} \mathcal{K}$
  For $L \in \mathcal{K}$ and $X \in \mathcal{D}$ do $\pi(L,X) \leftarrow \mathsf{undefined}$

For $L \in \mathcal{K}$ do $\mathsf{Range}_L \leftarrow \emptyset$
Run $A^{f(\cdot,\cdot)}$, replying to $f(\phi, X)$ oracle queries as follows:
  $K' \leftarrow \phi(K)$
  If $\pi(K', X) = \mathsf{undefined}$ then
    $Y \xleftarrow{\$} \mathcal{D}$
    If $Y \in \mathsf{Range}_{K'}$ then $\mathsf{bad} \leftarrow \mathsf{true}$ $\boxed{\,; Y \xleftarrow{\$} \mathcal{D} - \mathsf{Range}_{K'}}$
    $\mathsf{Range}_{K'} \leftarrow \mathsf{Range}_{K'} \cup \{Y\}$
    $\pi(K', X) \leftarrow Y$
  Return $\pi(K', X)$ to A
Until A halts returning a bit b
Return b

Let $\mathbf{GameS0}_{\mathcal{K},\mathcal{D}}(A)$ be exactly $\mathbf{GameS1}_{\mathcal{K},\mathcal{D}}(A)$ minus the statement inside the box.

Consider the experiment

$$K \xleftarrow{\$} \mathcal{K} \,;\; F \xleftarrow{\$} \mathrm{Rand}(\mathcal{K}, \mathcal{D}, \mathcal{D}) \,;\; A^{F_{\mathrm{RK}(\cdot,K)}(\cdot)}$$

and the game $\mathbf{GameS0}_{\mathcal{K},\mathcal{D}}(A)$. In both the experiment and the game, let $(\phi_1, X_1), (\phi_2, X_2), \ldots$ denote A's sequence of queries to its oracle, and let $Y_1, Y_2, \ldots$ denote the oracle's responses. Consider the oracle's response to the adversary's j-th oracle query. In both the experiment and the game, if $\phi_j(K) = \phi_i(K)$ and $X_j = X_i$ for any index $i < j$, then $Y_j = Y_i$; otherwise both the experiment and the game select Y_j uniformly at random from $\mathcal{D}$. This equivalence in behavior means that

$$\Pr\left[\, K \xleftarrow{\$} \mathcal{K}, F \xleftarrow{\$} \mathrm{Rand}(\mathcal{K}, \mathcal{D}, \mathcal{D}) \;:\; A^{F_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \,\right] \;=\; \Pr\left[\, \mathbf{GameS0}_{\mathcal{K},\mathcal{D}}(A) = 1 \,\right] . \tag{4}$$

Consider now the experiment

$$K \xleftarrow{\$} \mathcal{K} \,;\; G \xleftarrow{\$} \mathrm{Perm}(\mathcal{K}, \mathcal{D}) \,;\; A^{G_{\mathrm{RK}(\cdot,K)}(\cdot)}$$

and the game $\mathbf{GameS1}_{\mathcal{K},\mathcal{D}}(A)$. As before, let $(\phi_1, X_1), (\phi_2, X_2), \ldots$ denote A's sequence of queries to its oracle, and let $Y_1, Y_2, \ldots$ denote the oracle's responses. Consider the oracle's response to the adversary's j-th oracle query. In both the experiment and the game, if $\phi_j(K) = \phi_i(K)$ and $X_j = X_i$ for any index $i < j$, then $Y_j = Y_i$; otherwise both the experiment and the game select Y_j at random from $\mathcal{D}$ subject to the constraint that $Y_j \neq Y_k$ for all indices $k < j$ where $\phi_j(K) = \phi_k(K)$. Consequently,

$$\Pr\left[\, K \xleftarrow{\$} \mathcal{K}, G \xleftarrow{\$} \mathrm{Perm}(\mathcal{K}, \mathcal{D}) \;:\; A^{G_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \,\right] \;=\; \Pr\left[\, \mathbf{GameS1}_{\mathcal{K},\mathcal{D}}(A) = 1 \,\right] . \tag{5}$$

Let $\Pr_0[\,\cdot\,]$ denote the probability over the game $\mathbf{GameS0}_{\mathcal{K},\mathcal{D}}(A)$, and let B denote the event that $\mathbf{GameS0}_{\mathcal{K},\mathcal{D}}(A)$ sets bad to true. From the fundamental lemma of game-playing (Lemma 4 of [4] and Lemma 1 of [29]), we have that

$$\left|\, \Pr\left[\, \mathbf{GameS0}_{\mathcal{K},\mathcal{D}}(A) = 1 \,\right] - \Pr\left[\, \mathbf{GameS1}_{\mathcal{K},\mathcal{D}}(A) = 1 \,\right] \right| \;\leq\; \Pr_0\left[\, B \,\right] . \tag{6}$$

Combining Equations (4), (5), and (6), it follows that

$$\left|\, \Pr\left[\, K \xleftarrow{\$} \mathcal{K}, F \xleftarrow{\$} \mathrm{Rand}(\mathcal{K}, \mathcal{D}, \mathcal{D}) \;:\; A^{F_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \,\right] - \Pr\left[\, K \xleftarrow{\$} \mathcal{K}, G \xleftarrow{\$} \mathrm{Perm}(\mathcal{K}, \mathcal{D}) \;:\; A^{G_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \,\right] \right| \;\leq\; \Pr_0\left[\, B \,\right] .$$

To upper bound $\Pr_0[\,B\,]$ we first recall that given any two keys $K, K' \in \mathcal{K}$, there are at most NM_Φ RKD transformations in Φ that map K to K'. This means that $\mathbf{GameS0}_{\mathcal{K},\mathcal{D}}(A)$ may randomly

pick up to $q' = q \cdot \mathsf{NM}_\Phi$ range points for each key K'. By [2] we know that the probability of a collision in the range points for each key K' is therefore at most $q'(q'-1)/(2|\mathcal{D}|)$. Since A queries its oracle with at most r key-transformations, $\Pr_0[B] \leq rq'(q'-1)/(2|\mathcal{D}|)$, as desired. ∎

9 Existence of RKA-secure function families

In practice, it seems reasonable to assume that most modern block ciphers resist Φ-restricted related-key attacks for reasonable sets of RKD functions Φ (eg. Φ a subset of Φ_k^+ or $\Phi_k^\oplus$). This assumption is supported by the fact that the cryptographic primitives community frequently evaluates the security of block ciphers against related-key attacks (eg. [5, 19, 20, 17, 18] and numerous recent works) and modern block cipher are designed with the explicitly-stated goal of resisting related-key attacks (eg. [9]). (Please see the caveats discussed in Remark 7.7: eg, the complementation property of DES and the fact that 3DES provides less security than one would hope from a cipher with 168-bit keys [17].)

From a theoretical perspective, however, an interesting question is whether RKA-secure block ciphers (or function families) exist in theory (ie., assuming something else exists) for reasonable sets Φ of allowable RKD transformations.

Block ciphers. The following theorem shows that if the functions in Φ are restricted to only modifying some fixed portion of a block cipher's key, and if the transformations do not depend on the other portion of the key, then there exist block ciphers secure against Φ-restricted related-key attacks assuming that there exists block ciphers secure under the standard notion of pseudorandomness. This is already a useful result since, as shown in Section 7 in the context of ECBC' and FCBC', applications that use related-keys may only use key transformations that modify a small portion of the key.

Proposition 9.1 *Let $E\colon \{0,1\}^k \times \{0,1\}^l \to \{0,1\}^l$ be a block cipher. Let $E'\colon \{0,1\}^{k+l} \times \{0,1\}^l \to \{0,1\}^l$ be the block cipher defined as $E'_{K_1\|K_2}(M) = E_{K_1}(E_{K_1}(M) \oplus K_2)$ where K_1 is k-bits long and K_2 is l-bits long. Let Φ be any set of RKD functions over $\{0,1\}^{k+l}$ that modify only the last l-bits of the key and that are independent of the first k bits. Then assuming E is a secure block cipher, E' is a secure block cipher with respect to Φ-restricted related-key attacks. Formally, for any adversary A against E' that queries its related-key oracle with at most r different RKD transformations and at most q times per transformation, we can construct an adversary B against E such that*

$$\mathbf{Adv}^{\textit{prp-rka}}_{\Phi,E'}(A) \leq \mathbf{Adv}^{\textit{prp}}_{E}(B) + \frac{16r^2q^2 + rq'(q'-1)}{2^{l+1}}$$

and B makes $2rq$ oracle queries and runs in the same time as A and q' is q times the maximum, over all $K, K' \in \{0,1\}^{k+l}$, of the number of $\phi \in \Phi$ mapping K to K'. ∎

Proof of Proposition 9.1: The proof is a corollary of [2]'s pseudorandomness theorem for CBC MAC. Let $\mathsf{CBC}[E]\colon \{0,1\}^k \times \{0,1\}^{2l} \to \{0,1\}^l$ denote the CBC MAC construction defined as $M_0\|M_1 \mapsto E_K(E_K(M_0) \oplus M_1)$ where $|M_0| = |M_1| = l$. In [2] it was shown that given any PRF adversary A attacking $\mathsf{CBC}[E]$ making q oracle queries, one could construct a PRP adversary B attacking E such that

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{CBC}[E]}(A) \leq \mathbf{Adv}^{\mathrm{prp}}_{E}(B) + \frac{4q^2}{2^{l-1}}$$

and B makes $2q$ oracle queries.

Let us now consider a Φ-restricted RKA adversary C against E'. Let r denote the maximum number of different RKD transformations C queries its related-key oracle with, and let q denote the maximum number of queries per transformation. Let A be the $\mathsf{CBC}[E]$ adversary that works as follows:

Adversary $A^{f(\cdot)}$
  $T \xleftarrow{\$} \{0,1\}^l$
  Run C, responding to C's related-key request (ϕ, M) as follows
    // Recall: ϕ only modifies the last l bits of its input, and does not depend on the first k bits
    $T' \xleftarrow{\$}$ last l bits of $\phi(0^k \| T)$
    Return $f(M \| T')$ to A
  Until A halts returning a bit b
  Return b

Adversary A uses makes at most rq oracle queries. We now have that

$$\begin{aligned}
\mathbf{Adv}^{\text{prp-rka}}_{\Phi,E'}(C) &= \Pr\left[K \xleftarrow{\$} \{0,1\}^{k+l} \ : \ C^{E'_{\text{RK}(\cdot,K)}(\cdot)} = 1 \right] \\
&\quad - \Pr\left[K \xleftarrow{\$} \{0,1\}^{k+l} ; \ G \xleftarrow{\$} \text{Perm}(k+l,l) \ : \ C^{G_{\text{RK}(\cdot,K)}(\cdot)} = 1 \right] \\
&= \Pr\left[K \xleftarrow{\$} \{0,1\}^{k+l} \ : \ C^{E'_{\text{RK}(\cdot,K)}(\cdot)} = 1 \right] \\
&\quad - \Pr\left[K \xleftarrow{\$} \{0,1\}^{k+l} ; \ F \xleftarrow{\$} \text{Rand}(k+l,l,l) \ : \ C^{F_{\text{RK}(\cdot,K)}(\cdot)} = 1 \right] \\
&\quad + \Pr\left[K \xleftarrow{\$} \{0,1\}^{k+l} ; \ F \xleftarrow{\$} \text{Rand}(k+l,l,l) \ : \ C^{F_{\text{RK}(\cdot,K)}(\cdot)} = 1 \right] \\
&\quad - \Pr\left[K \xleftarrow{\$} \{0,1\}^{k+l} ; \ G \xleftarrow{\$} \text{Perm}(k+l,l) \ : \ C^{G_{\text{RK}(\cdot,K)}(\cdot)} = 1 \right]
\end{aligned}$$

and applying Proposition 8.9,

$$\begin{aligned}
&\leq \Pr\left[K \xleftarrow{\$} \{0,1\}^{k+l} \ : \ C^{E'_{\text{RK}(\cdot,K)}(\cdot)} = 1 \right] \\
&\quad - \Pr\left[K \xleftarrow{\$} \{0,1\}^{k+l} ; \ F \xleftarrow{\$} \text{Rand}(k+l,l,l) \ : \ C^{F_{\text{RK}(\cdot,K)}(\cdot)} = 1 \right] + \frac{rq'(q'-1)}{2^{l+1}} \ .
\end{aligned}$$

Note that

$$\Pr\left[K \xleftarrow{\$} \{0,1\}^{k+l} \ : \ C^{E'_{\text{RK}(\cdot,K)}(\cdot)} = 1 \right] = \Pr\left[K \xleftarrow{\$} \{0,1\}^{k} \ : \ A^{\mathsf{CBC}[E]_K(\cdot)} = 1 \right] \ .$$

Furthermore

$$\Pr\left[K \xleftarrow{\$} \{0,1\}^{k+l} ; \ F \xleftarrow{\$} \text{Rand}(k+l,l,l) \ : \ C^{F_{\text{RK}(\cdot,K)}(\cdot)} = 1 \right] = \Pr\left[f \leftarrow \text{Rand}(2l,l) \ : \ A^{f(\cdot)} = 1 \right] \ .$$

To see the last, note that when A runs C it replies to each query exactly as if C were run in the experiment on the left. In more detail, for each unique value T' for A in the second experiment, $f(\cdot \| T')$ is a random function from $\{0,1\}^l$ to $\{0,1\}^l$ and, for each unique key $K' = K_0 \| T'$ in the second experiment, $|T'| = l$, $F_{K'}(\cdot)$ is a random function from $\{0,1\}^l \to \{0,1\}^l$. Also recall that Φ contains only RKD functions that modify the last l-bits of E''s $k+l$-bit keys and that the transformations are independent of the first k bits.

This implies that there exists a PRP adversary B against E, making $2rq$ oracle queries, such that

$$\mathbf{Adv}^{\text{prp-rka}}_{\Phi,E'}(C) \leq \mathbf{Adv}^{\text{prp}}_{E}(B) + \frac{4r^2q^2}{2^{l-1}} + \frac{rq'(q'-1)}{2^{l+1}}$$

as desired. ∎

Another route for proving the existence of block ciphers secure against Φ-restricted related-key attacks (for similarly constrained Φ) is to note that (1) tweakable block ciphers exist assuming block ciphers exist [22] and (2) if the tweak of a tweakable block cipher E is made part of the key of a new block cipher E', then E' is secure against any Φ-restricted related-key attack provided that E is a secure tweakable block cipher, the functions in Φ only modify the tweak portion of E''s key, and the functions in Φ are not affected by the other portion of E''s key. We chose to prove Proposition 9.1 directly in order to obtain concrete bounds.

FUNCTION FAMILIES. The following proposition shows how to construct a Φ-restricted RKA-secure PRF from a standard PRF, where the functions in Φ only modify a portion of the new PRF's key and the functions in Φ are independent of the other portion of the new PRF's key.

Proposition 9.2 *Let $F : \{0,1\}^k \times \{0,1\}^l \to \{0,1\}^L$ be a family of functions, let $m \in \{1,\ldots,l-1\}$, and let $F' : \{0,1\}^{k+m} \times \{0,1\}^{l-m} \to \{0,1\}^L$ be another family of functions defined as $F'_K(M) = F_{K_1}(K_2 \| M)$ where K_1 is the first k bits of K and K_2 are the remaining m bits of K. If F is a secure PRF, then F' is a secure PRF with respect to Φ-restricted RKAs when the functions in Φ only modify the last m bits of F''s $k+m$-bit key and the functions in Φ do not depend on the first k-bits of the key. In particular, given a related-key adversary A attacking F', we can construct a standard PRF adversary B attacking F such that*

$$\mathbf{Adv}^{\textit{prf-rka}}_{\Phi,F'}(A) \leq \mathbf{Adv}^{\textit{prf}}_{F}(B)$$

and B takes the same amount of time and makes the same number of oracle queries as A. ∎

Proof of Proposition 9.2: The adversary B works as follows:

Adversary $B^f(\cdot)$
  $K_2 \xleftarrow{\$} \{0,1\}^m$
  Run A, responding to A's related-key request (ϕ, M) as follows
    $K'_2 \leftarrow$ last m bits of $\phi(0^k \| K_2)$
    Return $f(K'_2 \| M)$ to A
  Until A halts returning a bit b
  Return b

The equality

$$\Pr\left[K \xleftarrow{\$} \{0,1\}^k \ : \ B^{F_K(\cdot)} = 1 \right] = \Pr\left[K \xleftarrow{\$} \{0,1\}^{k+m} \ : \ A^{F'_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \right]$$

follows from the construction of F' and the fact that the permutations in Φ only modify the last m bits of F''s key and are independent of the first m bits (which allows B, when given access to an instance of F, to present A with an F' oracle). The equality

$$\begin{aligned}&\Pr\left[G \xleftarrow{\$} \mathrm{Rand}(l, L) \ : \ B^{G(\cdot)} \right] \\ &\quad = \Pr\left[K \xleftarrow{\$} \{0,1\}^{k+m}, G \xleftarrow{\$} \mathrm{Rand}(k+m, l-m, L) \ : \ A^{G_{\mathrm{RK}(\cdot,K)}(\cdot)} = 1 \right]\end{aligned}$$

comes from the fact that each distinct K'_2 generated by B will induce a random function from $\{0,1\}^{l-m}$ to $\{0,1\}^L$ (since the function $G : \{0,1\}^l \to \{0,1\}^L$ given to B is random). The theorem statement follows. ∎

We note that this same approach can be used to make a tweakable PRF.

9.1 RKA-attacks against existing PRFs

Although the above propositions show that, for limited RKD functions Φ, RKA-secure block ciphers and function families exist (assuming, respectively, that secure block ciphers and function families exist), an interesting question is whether one can do better. By "better," we mean a block cipher or function family that is provably secure even when an attacker is able to modify any portion of the underlying key (as we assume is the case for block ciphers).

We motivate our interest in such families by presenting related-key attacks against several provably secure (under the standard models) block ciphers and pseudorandom function families. We stress that we mount these attacks *outside of the model in which these function families were proven secure*; ie. these attacks do not invalidate the proofs of security for these constructions.

Following the extended abstract of this paper [3], Lucks [24] presents two function families that are provably secure against Φ-restricted related key attacks, for specific sets Φ that contain RKD functions that modify the entire underlying key. The security proofs in Lucks [24] are under new and non-standard complexity assumptions; it remains an open problem to find block ciphers and function families that are provably RKA-secure under standard assumptions when Φ consists of RKD functions that can modify the entire underlying key.

Luby-Rackoff [23]. In [23] Luby and Rackoff showed that a three round Feistel network with independent round keys is a secure pseudorandom permutation under chosen-plaintext attacks assuming that the round function is a secure PRF. And in [23] they showed that a four round Feistel network under the same assumptions is a secure PRP under chosen-ciphertext attacks.

Unfortunately, if the set Φ includes functions that modify only the last round key of any Feistel network E (regardless of the number of rounds), then one can easily mount a distinguishing Φ-restricted related-key attack against E: simply modify the last round key and see if the portion of the output corresponding to the input strand of the last round function is modified. If so, then the adversary is interacting with a family of random permutations. If not, then with high probability the adversary is interacting with an instance of E.

This observation is not new; it was used in [17] in the context of a key-recovery attack against DES with independent round subkeys. The novelty here is lifting the block cipher cryptanalytic ideas from [17] to the formal, distinguishing setting and, in the process, motivating the goal of a Φ-restricted RKA-secure block ciphers where the functions in Φ modify the block cipher's entire key (eg. $\Phi_k^{\oplus}$).

Naor-Reingold [25]. In [25] Naor and Reingold introduce a now well-known method of constructing a secure PRF based on the DDH assumption. We quote their construction exactly:

Construction 9.3 The Diffie-Hellman instance generator, IG, is a probabilistic polynomial-time algorithm such that on input 1^n the output of IG is distributed over triplets $\langle P, Q, g\rangle$, where P is an n-bit prime, Q a (large) prime divisor of $P-1$ and g an element of order Q in $\mathbb{Z}_P^*$.

We define the function ensemble $F = \{F_n\}_{n\in\mathbb{N}}$. For every n, a key of a function in F_n is a tuple, $\langle P, Q, g, \overrightarrow{a}\rangle$, where P is an n-bit prime, Q is a prime divisor of $P-1$, g is an element of order Q in $\mathbb{Z}_P^*$, and $\overrightarrow{a} = \langle a_0, a_1, \ldots, a_n\rangle$ a sequence of $n+1$ elements of $\mathbb{Z}_Q$. For any n-bit input $x = x_1x_2\cdots x_n$, the function $f_{P,Q,g,\overrightarrow{a}}$ is defined by:

$$f_{P,Q,g,\overrightarrow{a}}(x) = (g^{a_0})^{\Pi_{x_i=1}a_i} .$$

The distribution of functions in F_n is induced by the following distribution on their keys: $\overrightarrow{a}$ is uniform in its range and the distribution of $\langle P, Q, g\rangle$ is $IG(1^n)$. ∎

Let A be an adversary against this construction and let $h_{\mathrm{RK}(\cdot,P,Q,g,\overrightarrow{a})}(\cdot)$ be its related-key oracle. Let ϕ be a function that, on input a key $P, Q, g, \overrightarrow{a}$, computes a new key $P, Q, g, \overrightarrow{a'}$ such that

$\overrightarrow{a'} = \langle a_0, a_1 + 1 \bmod Q, a_2, \ldots, a_n \rangle$. Let id be the identity function.

Our $\{\phi, \mathrm{id}\}$-restricted related-key attack against Construction 9.3 works as follows. The adversary first obtains the values $y_1 = h_{\mathrm{id}(P,Q,g,\overrightarrow{a})}(0^n)$ and $y_2 = h_{\phi(P,Q,g,\overrightarrow{a})}(0^n)$ via its related-key oracle. If the adversary is given related-key access to Construction 9.3, then $y_1 = y_2$ with probability 1. If the adversary is given related-key access to a random family of functions, then with very high probability $y_1 \neq y_2$. This allows the adversary to distinguish between an instance of Construction 9.3 and a random family of functions.

NIELSEN [27]. We can also attack the function family described by Nielsen in [27], which is based on Construction 9.3. Informally, for every n the key of a function F_n is $\langle P, Q, g, \overrightarrow{a} \rangle$ where P, Q, g are as in Construction 9.3 and $\overrightarrow{a} = \langle (a_{1,0}, a_{1,1}), (a_{2,0}, a_{2,1}), \ldots, (a_{n,0}, a_{n,1}) \rangle$ is a sequence of n pairs of elements from $\mathbb{Z}_Q$. For any n-bit input $x = x_1 x_2 \cdots x_n$, the function $f_{P,Q,g,\overrightarrow{a}}$ is defined by

$$f_{P,Q,g,\overrightarrow{a}}(x) = g^{\prod_{i=1}^{n} a_{i,x_i}} \ .$$

Let ϕ be a function that, on input a key $P, Q, g, \overrightarrow{a}$, computes a new key $P, Q, g, \overrightarrow{a'}$ such that $\overrightarrow{a'} = \langle (a_{1,0} + 1 \bmod Q, a_{1,1}), (a_{2,0}, a_{2,1}), \ldots, (a_{n,0}, a_{n,1}) \rangle$. Let id be the identity function. Note that $f_{\phi(P,Q,g,\overrightarrow{a})}(1^n) = g^{\prod_{i=1}^{n} a_{i,1}} = f_{\mathrm{id}(P,Q,g,\overrightarrow{a})}(1^n)$, which allows an adversary to distinguish between an instance of Nielsen's function family and a random family of functions with very high probability.

References

[1] 3GPP. 3GPP TS 35.201 v 3.1.1, specification of the 3GPP confidentiality and integrity algorithms, document 1: $f8$ and $f9$ specification. Available at `http://www.3gpp.org/tb/other/algorithms.htm`.

[2] M. Bellare, J. Kilian, and P. Rogaway. The security of the cipher block chaining message authentication code. *Journal of Computer and System Sciences*, 61(3):362–399, 2000.

[3] M. Bellare and T. Kohno. A theoretical treatment of related-key attacks: RKA-PRPs, RKA-PRFs, and applications. In E. Biham, editor, *Advances in Cryptology – EUROCRYPT 2003*, volume 2656 of *Lecture Notes in Computer Science*, pages 491–506. Springer-Verlag, Berlin Germany, May 2003.

[4] M. Bellare and P. Rogaway. The game-playing technique. Cryptology ePrint Archive `http://eprint.iacr.org/`: Report 2004/331, 2004.

[5] E. Biham. New types of cryptanalytic attacks using related keys. In T. Helleseth, editor, *Advances in Cryptology – EUROCRYPT '93*, volume 765 of *Lecture Notes in Computer Science*, pages 398–409. Springer-Verlag, Berlin Germany, 1993.

[6] E. Biham, O. Dunkelman, and N. Keller. Related-key boomerang and rectangle attacks. In R. Cramer, editor, *Advances in Cryptology – EUROCRYPT 2005*, volume 3494 of *Lecture Notes in Computer Science*, pages 507–525. Springer-Verlag, Berlin Germany, May 2005.

[7] J. Black and P. Rogaway. CBC MACs for arbitrary-length messages: The three-key construction. In M. Bellare, editor, *Advances in Cryptology – CRYPTO 2000*, volume 1880 of *Lecture Notes in Computer Science*, pages 197–215. Springer-Verlag, Berlin Germany, 2000.

[8] N. Courtois and J. Pieprzyk. Cryptanalysis of block ciphers with overdefined systems of equations. Cryptology ePrint Archive `http://eprint.iacr.org/`: Report 2002/044, 2002.

[9] J. Daemen and V. Rijmen. AES proposal: Rijndael. `http://csrc.nist.gov/encryption/aes/rijndael/Rijndael.pdf`, 1999.

[10] J. Daemen and V. Rijmen. *The Design of Rijndael*. Springer-Verlag, Berlin Germany, 2002.

[11] N. Ferguson, J. Kelsey, S. Lucks, B. Schneier, M. Stay, D. Wagner, and D. Whiting. Improved cryptanalysis of Rijndael. In B. Schneier, editor, *Fast Software Encryption 2000*, volume 1978 of *Lecture Notes in Computer Science*, pages 213–230. Springer-Verlag, Berlin Germany, 2000.

[12] O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. *Journal of the ACM*, 33(4):210–217, 1986.

[13] T. Iwata and T. Kohno. New security proofs for the 3GPP confidentiality and integrity algorithms. In B. Roy and W. Meier, editors, *Fast Software Encryption 2004*, volume 3017 of *Lecture Notes in Computer Science*, pages 427–445. Springer-Verlag, Berlin Germany, 2004.

[14] T. Iwata and K. Kurosawa. OMAC: One-key CBC MAC. In T. Johansson, editor, *Fast Software Encryption 2003*, volume 2887 of *Lecture Notes in Computer Science*, pages 129–153. Springer-Verlag, Berlin Germany, 2003.

[15] T. Iwata and K. Kurosawa. On the correctness of security proofs for the 3GPP confidentiality and integrity algorithms. In K. G. Paterson, editor, *Cryptography and Coding, Ninth IMA International Conference*, volume 2898 of *Lecture Notes in Computer Science*, pages 306–318. Springer-Verlag, Berlin Germany, 2004.

[16] É. Jaulmes and R. Lercier. FRMAC, a fast randomized message authentication code. Cryptology ePrint Archive `http://eprint.iacr.org/`: Report 2004/166, 2004.

[17] J. Kelsey, B. Schneier, and D. Wagner. Key-schedule cryptanalysis of IDEA, G-DES, GOST, SAFER, and Triple-DES. In N. Koblitz, editor, *Advances in Cryptology – CRYPTO '96*, volume 1109 of *Lecture Notes in Computer Science*, pages 237–251. Springer-Verlag, Berlin Germany, 1996.

[18] J. Kelsey, B. Schneier, and D. Wagner. Related-key cryptanalysis of 3-WAY, Biham-DES, CAST, DES-X, NewDES, RC2, and TEA. In Y. Han, T. Okamoto, and S. Qing, editors, *Information and Communications Security '97*, volume 1334 of *Lecture Notes in Computer Science*, pages 233–246. Springer-Verlag, Berlin Germany, 1997.

[19] L. Knudsen. Cryptanalysis of LOKI91. In J. Seberry and Y. Zheng, editors, *Advances in Cryptology – AUSCRYPT '92*, volume 718 of *Lecture Notes in Computer Science*, pages 196–208. Springer-Verlag, Berlin Germany, 1992.

[20] L. Knudsen. A key-schedule weakness in SAFER K-64. In D. Coppersmith, editor, *Advances in Cryptology – CRYPTO '95*, volume 963 of *Lecture Notes in Computer Science*, pages 274–286. Springer-Verlag, Berlin Germany, 1995.

[21] L. Knudsen and T. Kohno. Analysis of RMAC. In T. Johansson, editor, *Fast Software Encryption 2003*, volume 2887 of *Lecture Notes in Computer Science*, pages 182–191. Springer-Verlag, Berlin Germany, 2003.

[22] M. Liskov, R. Rivest, and D. Wagner. Tweakable block ciphers. In M. Yung, editor, *Advances in Cryptology – CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 31–46. Springer-Verlag, Berlin Germany, 2002.

[23] M. Luby and C. Rackoff. How to construct pseudorandom permutations from pseudorandom functions. *SIAM J. Computation*, 17(2), Apr. 1988.

[24] S. Lucks. Ciphers secure against related-key attacks. In B. Roy and W. Meier, editors, *Fast Software Encryption 2004*, volume 3017 of *Lecture Notes in Computer Science*, pages 359–370. Springer-Verlag, Berlin Germany, 2004.

[25] M. Naor and O. Reingold. Number-theoretic constructions of efficient pseudo-random functions. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science*, pages 458–467. IEEE Computer Society Press, 1997.

[26] M. Naor and O. Reingold. On the construction of pseudorandom permutations: Luby-rackoff revisited. *J. Cryptology*, 12(1):29–66, 1999.

[27] J. B. Nielsen. A threshold pseudorandom function construction and its applications. In M. Yung, editor, *Advances in Cryptology – CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 401–416. Springer-Verlag, Berlin Germany, 2002.

[28] NIST. Recommendation for block cipher modes of operation: The CMAC mode for authentication. NIST Special Publication 800-38B, May 2005.

[29] V. Shoup. Sequences of games: A tool for taming complexity in security proofs. Cryptology ePrint Archive `http://eprint.iacr.org/`: Report 2004/332, 2004.