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

Mihir Bellare and Tadayoshi Kohno

Dept. of Computer Science & Engineering,
University of California at San Diego,
9500 Gilman Drive, La Jolla, California 92093, USA.
http://www-cse.ucsd.edu/users/{mihir,tkohno}

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 [14].

1 Introduction

Most modern block ciphers, including the AES [6], are designed with the explicitly stated goal of resisting what are called "related-key attacks (RKAs)" [3]. 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

E. Biham (Ed.): EUROCRYPT 2003, LNCS 2656, pp. 491–506, 2003.

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

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.

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 [15,1] 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 $0 \le i < 2^k$, 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 2). 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 3 identifies small, simple classes of permutations Φ for which we can present successful Φ-restricted RKAs on any block cipher, and Proposition 4 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 1 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 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 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.

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 1 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 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 [14] 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 [4] 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. In the full version of this paper 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, [11] 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. This extended abstract 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. In the full version of this 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.

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. Our results along these lines appear in the full version of this paper.

In the full version of this paper we first note 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 [15], any Feistel networks (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 [16] and Nielsen [17] 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.

Related work. Prior to the recent work of Courtois and Pieprzyk [5], the best (in terms of the number of rounds) known attack against Rijndael 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 Rijndael with 128-bit blocks and 256-bit keys [8]. Daemen and Rijmen discuss related-key attacks in their book [7] and in their AES submission documents [6] and comment that the diffusion and non-linearity of Rijndael's key schedule makes it difficult for related-key attacks to pass through the entire cipher. In [11] 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.

Full version. This abstract omits much of the content of the associated full paper [2].

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 [9] and PRPs by [15]. We recall the latter, but since our goal is to model block ciphers, we adopt the concrete approach of [1] rather than the asymptotic approach of the original papers. Let $\mathrm{Perm}(\mathcal{D})$ denote the set of all permutations on $\mathcal{D}$. 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 (i.e. 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 [1], 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.

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

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

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 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

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

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.

Remark 1. **[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.) The proof of Proposition 1 appears in [2].

Proposition 1. *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. ▮

Definition 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, in this extended abstract we stick to the basic notion of Definition 1, but shall discuss the other notions in [2].

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

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

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. The formal proof of the following appears in [2].

Proposition 2. *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*

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

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

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 3. *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*

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

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. ∎

An example set Φ for which the above proposition holds is $\Phi=\{\ \phi_i^c\ :\ c\in\{0,1\}, i\in\{1,\ldots,k\}\ \}\cup\{\mathrm{id}\}$ 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. See the full version of this paper for the complete proof.

While one might consider the above set of RKD permutations somewhat contrived, 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. The proof of the following proposition appears in [2].

Proposition 4. *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*

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

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*.* ∎

5 Properties of RKD Transformations

The attack in Proposition 2 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 3 and Proposition 4 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 3). 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 2 was not output-unpredictable.

Definition 2. [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 3 and Proposition 4 both exploit collisions of this form.

Definition 3. [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 1. *Let Φ be a set of permutations on some keys space $\mathcal{K}$. Let r, r' be positive integers. Then*

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

For the canonical sets of RKD functions in which we are interested, namely Φ^{+}_{k} and $\Phi^{\oplus}_{k}$, the following lemma shows that collision-resistance is guaranteed.

Lemma 2. *Let $\mathcal{K} = \{0,1\}^k$ and let Φ be either Φ^{+}_{k} or $\Phi^{\oplus}_{k}$. Then for any positive integer r,*

$$\mathbf{InSec}^{cr}_{\Phi}(r) \ = \ 0 \ . \quad ∎$$

See [2] for the proof of Lemma 1 and Lemma 2.

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. See [2] for details.

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 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 4. [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 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 lower bounds on $\mathbf{Adv}^{\text{prp-rka}}_{\Phi,\mathcal{K},\mathcal{D}}(A)$ in Proposition 2, Proposition 3 and Proposition 4 become, respectively

$$1 - \frac{1}{|D|} \ ,\ 1 - \frac{k+1}{|D|} \ \text{ and } \ 1 - \frac{k+3}{2 \cdot |D|} \ . \ ∎$$

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 1. *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*

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

The proof of Theorem 1 is in [2]. 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 3. Theorem 1 extends to PRPs with respect to Φ-restricted chosen-ciphertext RKAs and to Φ-restricted RKAs against the pseudorandomness of function families. See [2]. ∎

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 1. *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,*

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

Proof (Corollary 1). Combine Lemma 1, Lemma 2, and Theorem 1.

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 $\Phi_k^\oplus$, 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 [14].

Proof of security for a tweakable block cipher. In [14] 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 [14]. 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

$$\mathbf{Adv}^{\text{tweak-prp}}_{\widetilde{E}}(A) = \Pr\left[K \xleftarrow{\$} \mathcal{K} \ : \ A^{\widetilde{E}_K(\cdot,\cdot)} = 1 \right] - \Pr\left[G \xleftarrow{\$} \text{Perm}(\{0,1\}^t, \{0,1\}^l) \ : \ A^{G(\cdot,\cdot)} = 1 \right]$$

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 $\Phi_k^{\oplus}$-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. The proof of Theorem 2 appears in [2].

Theorem 2. *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. ∎

Single-key CBC MACs for Arbitrary-Length Messages. In addition to proving the security of existing constructions (eg. the examples mentioned in [11] 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 exemplify this by providing, in the full version of this paper, some CBC-MAC variants that are provably secure for variable-length messages yet use just one key. (While these MAC constructions have only one key, they still require running the block cipher's key schedule algorithm for three related keys.)

We base our constructions on two of Black and Rogaway's [4] "three-key constructions," and our constructions, keyed with a single key K, invoke the underlying block cipher with keys $\phi_1(K), \phi_2(K), \phi_3(K)$, where $\Phi = \{\phi_1, \phi_2, \phi_3\}$ is some set of RKD permutations and $\mathbf{InSec}^{\mathrm{cr}}_{\Phi}(3) = 0$. For example, Φ might consist of the three functions $\mathrm{ADD}_0, \mathrm{ADD}_1, \mathrm{ADD}_2$.

From a pragmatic perspective, one may now wish to use TMAC [13] or OMAC [10], two recently-proposed two-key and one-key CBC-MAC variants, rather than the constructions we present in the full version of this paper. We present our constructions primarily because they illustrate the use of RKA-secure PRPs in constructs. See the full version of this paper for details.

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

Remark 4. 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 the full version of this paper, which shows that the insecurity of block ciphers under Φ_k^+- or $\Phi_k^\oplus$-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 5. 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 [12]. As another example, note that the tweakable block cipher in Theorem 2 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 6. While most modern block ciphers, including the AES, are designed with the explicitly stated goal of resisting related-key attacks, it is important to note that 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 [11]'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 [12]. ∎

Acknowledgments. Mihir Bellare is supported in part by NSF Grant CCR-0098123, NSF Grant ANR-0129617 and an IBM Faculty Partnership Development Award. Tadayoshi Kohno is supported by a National Defense Science and Engineering Graduate Fellowship.

References

1. 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.
2. M. Bellare and T. Kohno. A theoretical treatment of related-key attacks: RKA-PRPs, RKA-PRFs, and applications. Full version of this paper, available at `http://www-cse.ucsd.edu/users/tkohno/papers/RKA/`, 2003.

3. 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.
4. 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.
5. 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.
6. J. Daemen and V. Rijmen. AES proposal: Rijndael. `http://csrc.nist.gov/encryption/aes/rijndael/Rijndael.pdf`, 1999.
7. J. Daemen and V. Rijmen. *The Design of Rijndael*. Springer-Verlag, Berlin Germany, 2002.
8. 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.
9. O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. *Journal of the ACM*, 33(4):210–217, 1986.
10. T. Iwata and K. Kurosawa. OMAC: One-key CBC MAC. In T. Johansson, editor, *Fast Software Encryption 2003*, Lecture Notes in Computer Science. Springer-Verlag, Berlin Germany, 2003.
11. 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.
12. L. Knudsen and T. Kohno. Analysis of RMAC. In T. Johansson, editor, *Fast Software Encryption 2003*, Lecture Notes in Computer Science. Springer-Verlag, Berlin Germany, 2003.
13. K. Kurosawa and T. Iwata. TMAC: Two-key CBC MAC. NIST submission, available at `http://csrc.nist.gov/CryptoToolkit/modes/`, June 2002.
14. 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.
15. M. Luby and C. Rackoff. How to construct pseudorandom permutations from pseudorandom functions. *SIAM J. Computation*, 17(2), Apr. 1988.
16. 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.
17. 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.