A Formal Treatment of Remotely Keyed Encryption
(Extended Abstract)*

Matt Blaze,[1] Joan Feigenbaum,[1] Moni Naor[2]

[1] AT&T Labs – Research
180 Park Avenue
Florham Park, NJ 07932 USA
{mab,jf}@research.att.com

[2] Dept. Applied Math. and Computer Science
Weizmann Institute of Science
Rehovot 76100, Israel
naor@wisdom.weizmann.ac.il

Abstract. Remotely keyed encryption schemes (RKESs), introduced by Blaze [6], support high-bandwidth cryptographic applications (such as encrypted video conferences) in which long-lived secrets (such as users' private keys) never leave lower-bandwidth environments such as secure smart-cards. We provide a formal framework in which to study the security of RKESs and give an RKES that satisfies our formal security requirements. Our RKES is efficient in that the amount of communication and computation required of the smart-card is independent of the input size. Our proof of security uses the pseudorandom permutation framework of Naor and Reingold [14] in an essential way.

Keywords: Block Ciphers, Pseudorandomness, Remotely Keyed Encryption, Session Keys, Smart-cards

1 Introduction

No cryptographic protocol is stronger than the mechanism protecting its secret keys. However, in many computing and communication systems, there is no "safe place" in which secret keys can be stored and cryptographic computations can be performed. This is especially true of modern networked computers; in some sense, every computer that communicates extensively with the world is bound at some point to be partly controlled by an unfriendly entity. Therefore, it is natural to consider adding an external, special-purpose device, such as a smart-card or a PCMCIA card, for storing cryptographic keys and computing

* A preliminary version of the full paper is available from the authors. The third author was supported by a RAND2 grant from the EC.

cryptographic functions. Because they have only one purpose and communicate only via a limited set of functions, such devices can be made much more secure than their general-purpose host machines. However, it is not always practical to rely on such devices to perform all sensitive cryptographic operations. Inexpensive smart-cards, for example, are also characterized by their limited bandwidth, memory, and processor speed. If the host computer used such devices simple-mindedly, by just encrypting all external communication and all disk traffic, then the bandwidth of the link between the host and the cryptographic module would have to be at least as high as that between the host and the outside world. Even if the engineering problems of developing inexpensive high-bandwidth, high-performance cryptographic modules were completely solved, it would still be the case that, whenever the host's link to the outside world was upgraded, the modules and the secret keys they store would have to be changed, because cryptographic modules are typically designed never to reveal their keys.

This paper provides a formal treatment of the *remotely keyed encryption problem*: how to do bulk encryption and decryption for high-bandwidth applications in a way that takes advantage of both the superior power of the host and the superior security of the smart-card? If adversary $\mathcal{A}$ takes control of the host for a certain period, then clearly $\mathcal{A}$ will obtain whatever plaintext or ciphertext is resident in the host during that period. We would like to say formally that this is all it obtains: Once $\mathcal{A}$ loses control of the host, it cannot compute anything that it couldn't compute before it took control, except for the values it obtained explicitly while it was in control.

Note that we are concerned with attacks on the host but *not* with direct attacks on the card; we assume that the card owner wants to safeguard the "remote keys" and that an attacker can only communicate with the card via its official communication channels. See, *e.g.*, Boneh *et al.* [7] and Biham-Shamir [5], for a discussion of direct attacks on cards. Note as well that the remotely keyed encryption problem is different from the one of having a smart-card take advantage of a host's superior processing power in order to do a *public-key* computation without leaking the *input* to the host. For a discussion of the (well-studied) problem of host-assisted public-key cryptography, see *e.g.*, Feigenbaum [8], Matsumoto *et al.* [13], and the references therein. Finally, note that the goal of a remotely keyed encryption scheme (RKES) is not "session-key exchange" between two *different* hosts each connected to a card, where the two cards share a key. In an RKES application, such as the encryption of disk traffic, there is only one host; it encrypts at some point in time and then decrypts the stored ciphertext later. In settings in which the reason for exchanging keys is encryption, an RKES may replace a session-key exchange protocol and has the advantage of no interaction; however, there are other reasons for key exchange. See, for example, Shoup and Rubin [15] for a rigorous treatment of session-key exchange (following Bellare and Rogaway [3]), in which the adversary is similar in power to the one we consider here. The Shoup-Rubin protocol requires several rounds of communication between the hosts.

We give a formal definition that captures the notion of "security" needed in our scenario and an RKES that satisfies the definition. Our RKES is efficient as well as secure, because the amount of communication and computation required of the smart-card is independent of the input size.

History

Blaze [6] was the first to use the term "remotely keyed encryption" and to focus attention on the fact that many high-bandwidth applications need symmetric-key encryption schemes that store long-lived secret keys in low-bandwidth smart-cards.[1] He proposed a specific scheme but did not give a formal statement of the properties that an RKES should satisfy. Although the scheme in [6] does not satisfy the formal security requirements that we give in this paper, the basic idea of the scheme is sound, and we use it as a starting point in the design of an RKES that does satisfy our formal requirements. One weakness of the original scheme in [6] is that it may enable an adversary that has controlled the host during m interactions with the card subsequently to "forge" a plaintext/ciphertext pair that is not one of the m pairs he has obtained during the interaction.

Lucks [12] first noted that the RKES in [6] was not completely satisfactory; in particular, he noted the forgery weakness just described. Lucks [12] attempted to formalize the security properties that an RKES should have and to construct schemes that have them. Although the properties proposed in [12] are indeed desirable, we believe that the overall formalism and construction are flawed. Roughly speaking, [12] proposed that an RKES should have three properties: (i) Forgery security: If the adversary has controlled the host for m interactions, then it cannot produce $m+1$ plaintext/ciphertext pairs; (ii) Inversion security: Access to encryption should not imply the ability to decrypt and *vice versa*; (iii) Pseudorandomness: The encryption function should be a pseudorandom permutation. We suggest that an RKES might have these three properties but still be "insecure" in an intuitive sense.

In fact, the scheme in [12] is a good example of one that has properties (i), (ii), and (iii) but is intuitively insecure. That scheme uses the first two plaintext blocks in order to define an encryption key for the rest of the message; the encryption of these two blocks depends on the entire message, thus allowing property (i) to be satisfied. However, because the encryption key depends only on the first two plaintext blocks, an arbitrarily large set of messages all of which start with the same two blocks will always be encrypted with the same key. This is not a hypothetical situation: A set of files in a computer file system, for example, might always start with the same few bytes of structural information. An adversary that controls the host during the encryption or decryption of *one* file in such a set could subsequently decrypt the encryption of *any* file in the set. More fundamentally, the framework in [12] fails to recognize that it is nontrivial in this scenario to give a precise meaning to the statement that "the encryption

[1] Blaze [6] used the phrase "remotely keyed encryption *protocol*," and we use "remotely keyed encryption *scheme*." The terms are interchangeable.

function is a pseudorandom permutation." Once the adversary has witnessed one host/card interaction, it can subsequently distinguish between the encryption function and a random permutation by asking for the value on a single point.

We thus conclude that the formalism in [12] is inadequate. In this paper, we develop a formal framework that is both more precise and more stringent than those in the previous literature. In particular, we define pseudorandomness in a way that is meaningful for remotely keyed encryption.

In this Extended Abstract, we restrict attention to *length-preserving* RKESs, in which ciphertexts are of exactly the same length as the corresponding plaintexts. The full paper also treats *length-increasing* RKESs, in which ciphertexts are slightly longer than the corresponding plaintexts. The advantage of allowing increased length is that one can also achieve a stronger notion of security.

2 Notation, Terminology, and Building Blocks

Definitions of standard cryptographic and complexity theoretic terms can be found in, for example, Goldreich [9], Luby [11], and Naor and Reingold [14]. The following is a description of the building blocks and terminology used throughout.

- The plaintext and the ciphertext are, respectively, X and Y. Usually, both are given in blocks and hence are denoted $X = (X_1, \ldots, X_n)$ and $Y = (Y_1, \ldots, Y_n)$, where each of X_i and Y_i is in $\{0,1\}^b$.
- The encryption and decryption functions of a block cipher are E and D. $E_S(X_j)$ denotes the encryption of plaintext block X_j with encryption key S, *i.e.*, $E_S, D_S : \{0,1\}^b \mapsto \{0,1\}^b$ and $D_S(E_S(X_j)) = X_j$.
  The security property required of $E_S(\cdot)$ is that it should be a strong pseudorandom permutation, *i.e.*, that any probabilistic, polynomial-time adversary given access to $E_S(\cdot)$ and $D_S(\cdot)$ cannot distinguish them from a truly random permutation. A thorough treatment of strong pseudorandom permutations is given by Luby [11], who calls them "super" pseudorandom permutations.
- A pseudorandom function $F_S : \{0,1\}^b \mapsto \{0,1\}^b$; it may or may not be identical to the encryption function E of the block cipher. (Note that every pseudorandom permutation is also a pseudorandom function, where the added advantage of a distinguisher is bounded by $m^2/2^b$.) We use F_S, rather than E_S, in situations that never require the function to be inverted.
- A length-preserving method G_S for encrypting an n-block plaintext $(X_1, \ldots, X_n)$ using encryption key S. $G_S^j(X_1, \ldots, X_n)$ denotes the jth block of the resulting ciphertext. The corresponding decryption function is denoted $\hat{G}_S$, and the jth block of the plaintext that results from decrypting $(Y_1, \ldots, Y_n)$ is denoted $\hat{G}_S^j(Y_1, \ldots, Y_n)$.
  The security requirement for G_S is that, for any $X_1, X_2, \ldots, X_n$, if S is chosen uniformly at random, then $G_S(X_1, \ldots, X_n)$ is pseudorandom (*i.e.*, indistinguishable from a random string of similar length). We impose a similar requirement on $\hat{G}_S$. Possible realizations of G_S are:
  - Apply a pseudorandom generator to S and Xor the resulting sequence with $X_1, \ldots, X_n$.

- Use E_S with some sort of chaining, *e.g.*, CBC. The security of such an operation follows from [2].

– A collision-intractable hash function $H : \{0,1\}^* \mapsto \{0,1\}^b$. "Collision-intractability" means that it is computationally infeasible to find distinct X and X' such that $H(X) = H(X')$.

– The adversary is in general an oracle machine M where $M^{(f,f^{-1})}$ that has access to the *function pair* (f, f^{-1}). As in [14], M may submit two forms of queries to the function-pair oracle: A query of the form $(+, x)$ results in the answer $f(x)$, and one of the form $(-, y)$ results in the answer $f^{-1}(y)$.

As usual, various parameters are needed in order to express things in full detail. In particular, there is an underlying size (security) parameter u, and there are three polynomially bounded functions that measure the key length $\kappa(u)$, the block length $b(u)$, and the number of blocks $n(u)$. The total length of the input to any of our protocols is a polynomial function of $\kappa(u)$, $b(u)$, and $n(u)$. For clarity of presentation, we suppress these parameters whenever possible, but they are an implicit part of everything that follows. For example, the statement "F is a pseudorandom function" means that $F : \{0,1\}^{\kappa(u)} \times \{0,1\}^{b(u)} \to \{0,1\}^{b(u)}$ is a pseudorandom function generator, in the sense of [10] or [11, Lecture 12]. Similarly, additional (but standard) detail is also required to say precisely what is meant by a "random" function, permutation, or function pair. These details can be found in, for example, [9,11,14].

3 Definition of Secure Remotely Keyed Encryption

Intuitively, we would like an RKES to resist the following form of attack. Adversary $\mathcal{A}$ may gain control of the host temporarily. During this *host phase* of his attack, $\mathcal{A}$ may have a total of m interactions with the card, where m is polynomially bounded. He may send any message to the card during one of these interactions and may deviate from the protocol. However, since m is an upper bound on the total number of interactions, $\mathcal{A}$ obtains at most m plaintexts and m ciphertexts during the host phase. After these m interactions with the card, $\mathcal{A}$ loses control of the host. He should subsequently have no advantage in his attempts to find the encryption (resp., decryption) of plaintexts (resp., ciphertexts) other than those m that he found explicitly during the host phase.

One step in formalizing this intuition is to make precise what we mean by "no advantage." We will do this in terms of pseudorandomness. That is, the encryption and decryption protocols of the RKES compute some function pair (f, f^{-1}), and this function pair should appear truly random to $\mathcal{A}$. During the host phase, $\mathcal{A}$ learns the value of f and f^{-1} each at m points. This should give him "no advantage" in the sense that if, in a *distinguishing phase* that takes place after he loses control of the host, $\mathcal{A}$ is asked to distinguish between (f, f^{-1}) and a truly random function pair, he should be able to do so only with negligible probability.

Because the amount of communication between the host and the card in an efficient RKES should be much shorter than the input length, there is a

complication that is missing in the standard definition of pseudorandomness. If, during the host phase, $\mathcal{A}$ learns the value of f on m points $X^1, X^2, \ldots, X^m$, then f almost certainly does *not* look random to $\mathcal{A}$ on these points, because the transcript of the host phase and the description of the protocol constitute a short description of $(X^1, f(X^1)), (X^2, f(X^2)), \ldots, (X^m, f(X^m))$. Our formalism addresses this by requiring that, between the host phase and the distinguishing phase, a nondeterministic choice occurs: Either (f, f^{-1}) is replaced with a truly random function pair or it is kept the same. $\mathcal{A}$'s challenge is thus to decide whether or not a switch occurred or not.

Now a new complication arises: We cannot allow the adversary, during the distinguishing phase, to query the oracle about any of the m plaintexts and m ciphertexts that he obtained during the host phase. If the adversary $\mathcal{A}$ does so and receives the same answers as he did the first time around, then $\mathcal{A}$ will know, with high probability, that there was no switch and (f, f^{-1}) remained unchanged. If $\mathcal{A}$ receives a different answer on one of these queries during the distinguishing phase than he did during the host phase, then he can conclude with certainty that the oracle is not (f, f^{-1}).

We would like therefore to "filter" those values that appeared in the host phase. The problem with making this discussion rigorous is that the adversary's actions during the host phase do not necessarily correspond to specific inputs, and certainly there are many inputs that yield the same (host,card) transcript. To overcome this problem, we introduce an *arbiter* into our definition of secure remotely keyed encryption. The purpose of the arbiter, which we denote by $\mathcal{B}$, is to make sure that $\mathcal{A}$ does not ask during the distinguishing phase any of the queries that it asked during the host phase. $\mathcal{B}$ should be a simple function of the transcript of the communication that occurred during the host phase and should have limited filtering ability. Instead of saying that the inputs queried during the host phase are excluded (which is *not* well-defined), we say that if $\mathcal{A}$ has had m interactions with the card during the host phase, then $\mathcal{B}$ is allowed to filter no more than m queries during the distinguishing phase. Note that there is no need actually to implement this arbiter; rather, a (host,card) protocol is secure if there exists such an arbiter.

This discussion can be summarized as follows.

Definition 1. A **length-preserving RKES** is a pair of protocols, one for encryption and one for decryption, to be executed by a *host* and a *card*. The length of a ciphertext must be the same as that of the corresponding plaintext. The RKES is **secure** if there is a polynomial-time *arbiter* $\mathcal{B}$ that can enforce the following restriction on any probabilistic, polynomial-time *adversary* $\mathcal{A}$ and any polynomial bound m: During the *host phase*, $\mathcal{A}$ may play the role of the host in a total of m interactions with the card. During this phase, $\mathcal{A}$ may send any message to the card and does not necessarily follow the encryption or decryption protocol. Between the host phase and the *distinguishing phase*, a nondeterministic choice is made between continuing to use the RKES or switching to a random function pair. The arbiter $\mathcal{B}$ receives as input the transcript of the host-phase communication between the host and the card. During the distinguishing phase,

$\mathcal{A}$ may run any probabilistic, polynomial-time test T that submits plaintexts or ciphertexts to $\mathcal{B}$; on at most m of the plaintexts and m of the ciphertexts, $\mathcal{B}$ may choose to run the RKES, even if a switch to a random function pair was made between phases. Otherwise the plaintext (resp. ciphertext) is given to the encryption (resp. decryption) protocol if no switch was made between phases and to the random function f (resp. f^{-1}) if a switch was made. The difference between the probability that T accepts on a continuation of the RKES and the probability that T accepts on a switch to a random function pair must be negligible.

A natural way to relax the above requirement is to allow $\mathcal{B}$ to reject polynomially in m many input/output, instead of exactly m. However, the constructions given in Section 4 achieve the stricter notion.

Three remarks are in order about this definition. First, it generalizes the corresponding definitions in standard (*i.e.*, not remotely keyed) encryption. If one fixes $m = 0$, *i.e.*, makes the host phase trivial, then Definition 1 reduces to the definition of a strong pseudorandom permutation. Second, Definition 1 is concerned with security rather than efficiency. Note, for instance, that a strong pseudorandom permutation evaluated solely by the card satisfies the Definition. Clearly, an RKES is most useful if the computational, memory, and bandwidth demands on the card are small. In particular, it is desirable for all to be slowly growing functions of the block length b and key length κ and to be independent of n, the number of blocks in the plaintext. Finally, we assume that the length of the message is known. It may be implicitly known, *e.g.*, the size of a disk sector or a data packet, or it may be conveyed by some other protocol. Note that it is possible to set our protocols so that they yield a different permutation for each message length.

4 A Secure, Length-Preserving RKES

Figure 1 below contains Scheme P, a secure, length-preserving RKES.[2] The card's secret key has four components k_1, k_2, k_3, and k_4. We have designed Scheme P for maximum clarity and have not sought to optimize by, for example, minimizing the number of distinct key components or cryptographic building blocks.

The scheme is best understood as part of the Naor-Reingold [14] framework for constructing and proving the security of pseudorandom permutations. In this framework, the pseudorandom permutation Π is the composition of three permutations: $\Pi \equiv p_2^{-1} \circ A \circ p_1$. In general, p_1 and p_2^{-1} are "lightweight," and A is where most of the work is done. In our setting, A will be the part performed mostly by the host, and p_1 and p_2^{-1} will be done mostly by the card.

The "heavyweight" building block A should behave as a random permutation on most inputs. An important step in applying the Naor-Reingold framework is

[2] The "P" is for "preserving." In the full paper, we also present Schemes I1 and I2, which are both length-increasing.

the identification of a collection of input-output sequences that are called "A-good." For an input-output sequence $\langle (X^1, Y^1), \ldots, (X^m, Y^m) \rangle$ to be A-good, $Pr_A[Y^i = A(X^i), 1 \leq i \leq m]$ should be close to $2^{-\ell \cdot m}$, where $\ell = n \cdot b$, *i.e.*, the probability should be close to what it would be if A were a truly random function. The role of the permutations p_1 and p_2 is to ensure that, with overwhelming probability, the inputs and outputs to A form an A-good sequence, even if the inputs to Π are chosen by an adaptive adversary, under a chosen plaintext and ciphertext attack. Thus a sequence is A-good or not based on p_1 and p_2.

Encryption protocol: input $X_1, \ldots, X_n$**; output** $Y_1, \ldots, Y_n$

P1 Host: $h_x \leftarrow H(X_2, \ldots, X_n)$
P2 Host $\rightarrow$ Card: h_x, X_1
P3 Card: $W \leftarrow E_{F_{k_1}(h_x)}(X_1)$
P4 Card: $Z \leftarrow E_{k_2}(W)$
P5 Card: $S \leftarrow F_{k_3}(W)$
P6 Card $\rightarrow$ Host: S
P7 Host: For $j \leftarrow 2$ to n, $Y_j \leftarrow G^j_S(X_2, \ldots, X_n)$.
P8 Host: $h_y \leftarrow H(Y_2, \ldots, Y_n)$
P9 Host $\rightarrow$ Card: h_y
P10 Card: $Y_1 \leftarrow E_{F_{k_4}(h_y)}(Z)$
P11 Card $\rightarrow$ Host: Y_1

Decryption protocol: input $Y_1, \ldots, Y_n$**; output** $X_1, \ldots, X_n$

P12 Host: $h_y \leftarrow H(Y_2, \ldots, Y_n)$
P13 Host $\rightarrow$ Card: h_y, Y_1
P14 Card: $Z \leftarrow D_{F_{k_4}(h_y)}(Y_1)$
P15 Card: $W \leftarrow D_{k_2}(Z)$
P16 Card: $S \leftarrow F_{k_3}(W)$
P17 Card $\rightarrow$ Host: S
P18 Host: For $j \leftarrow 2$ to n, $X_j \leftarrow \hat{G}^j_S(Y_2, \ldots, Y_n)$.
P19 Host: $h_x \leftarrow H(X_2, \ldots, X_n)$
P20 Host $\rightarrow$ Card: h_x
P21 Card: $X_1 \leftarrow D_{F_{k_1}(h_x)}(W)$
P22 Card $\rightarrow$ Host: X_1

Fig. 1. Scheme P: Length-Preserving RKES

In our construction, p_1 and p_2 produce output that depends on all the input blocks, but they change only the first block. That is

$$p_1 : (X_1, X_2, \ldots, X_n) \mapsto (W, X_2, \ldots, X_n),$$

where W is a function of X_1 and $h_x = H(X_2, \ldots, X_n)$, and

$$p_2 : (Y_1, Y_2, \ldots, Y_n) \mapsto (Z, Y_2, \ldots, Y_n),$$

where Z is a function of Y_1 and $h_y = H(Y_2, \ldots, Y_n)$.

Good sequences will be those in which different $X_1, \ldots, X_n$ and $X'_1, \ldots, X'_n$ are mapped by p_1 to different W and W', and similarly different Y's are mapped by p_2 to different Z's. To obtain permutations p_1 and p_2 with the right properties, we define a new primitive called *non-colliding encryption*.

Definition 2. A *non-colliding encryption scheme* is a pair of keyed functions $C_k : \{0,1\}^b \times \{0,1\}^b \mapsto \{0,1\}^b$ and $\hat{C}_k : \{0,1\}^b \times \{0,1\}^b \mapsto \{0,1\}^b$ with the following two properties.

1. For all $V \in \{0,1\}^b$ and $h \in \{0,1\}^b$, the functions satisfy $\hat{C}_k(C_k(V,h),h) = V$ and $C_k(\hat{C}_k(V,h),h) = V$. (Note that this property allows us to use C to "store" V, provided h is retrievable.)
2. Let $\mathcal{A}$ be a probabilistic, polynomial-time adversary that is allowed to query C_k and $\hat{C}_k$ adaptively. We say that "(V,h) appears in a (polynomial-length) sequence of queries" if $\mathcal{A}$ asks for $C_k(V,h)$ directly or if V is the reply to some direct query $\hat{C}_k(U,h)$. If the key k is chosen at random, then $\mathcal{A}$ has only a negligible probability of finding two pairs $(V,h) \neq (V',h')$ such that (a) $C_k(V,h) = C_k(V',h')$, and (b) at least one of (V,h) and (V',h') did not appear in the sequence of queries.

In Scheme P, the permutations p_1 and p_2 are determined by $(C_k, \hat{C}_k)$. For example, $p_1 : (X_1, X_2, \ldots, X_n) \mapsto (W = C_{k_1}(X_1, H(X_2, \ldots, X_n)), X_2, \ldots, X_n)$, and p_2 is defined similarly. The "storage" capability of non-colliding encryption ensures that p_1 is indeed a permutation, because $p_1^{-1} : (W, X_2, \ldots, X_n) \mapsto (X_1 = \hat{C}_{k_1}(W, H(X_2, \ldots, X_n)), X_2, \ldots, X_n)$. The permutation A depends on the two key components k_3 and k_4. $A : (W, X_2, \ldots, X_n) \mapsto (Z, Y_2, \ldots, Y_n)$, where $Z = E_{k_3}(W)$, $(Y_2, \ldots, Y_n) = G_{F_{k_4}(W)}(X_2, \ldots, X_n)$, and $G_S : \{0,1\}^{(n-1)b} \mapsto \{0,1\}^{(n-1)b}$. The overall permutation computed by the encryption protocol is $\Pi = p_2^{-1} \circ A \circ p_1$.

Our main result is as follows.

Theorem 3. Scheme P is a secure, length-preserving RKES.

Proof. To prove this result, we must define an arbiter $\mathcal{B}$, construct a non-colliding encryption scheme, and apply the Naor-Reingold framework [14]. Applying the framework entails identifying A-good sequences and proving that the overall construction gives a strong pseudorandom permutation. The identification of A-good sequences has to take into account the "two-phase" aspect of the definition of security of RKESs; this is a complication that is not present in the original Naor-Reingold paper. We address each of these issues in turn.

Arbiter:

$\mathcal{B}$ records all the pairs (h_x, X_1) and (h_y, Y_1) that appear in the host phase. The list of pairs is easy to deduce from the transcript. During the distinguishing phase, $\mathcal{B}$ does the following for each encryption query $(+, (X_1, \ldots, X_n))$. First, it computes $h_x = H(X_2, \ldots, X_n)$. If the pair (h_x, X_1) appeared in the host

phase, then $\mathcal{B}$ answers the query using the encryption protocol; otherwise, it uses either the encryption protocol or the random permutation, depending on whether the decision between phases was to continue or to switch. Similarly, when it receives a decryption query $(-,(Y_1,\ldots,Y_n))$ during the distinguishing phase, $\mathcal{B}$ first computes $h_y = H(Y_2,\ldots,Y_n)$; then, if the pair (h_y, Y_1) appeared in the host phase, $\mathcal{B}$ uses the decryption protocol to answer the query, and otherwise it uses either the decryption protocol or the random function inverse, depending on whether the decision between phases was to continue or to switch.

Non-colliding encryption:

We provide one construction here and an additional one in the full paper. Assume without loss of generality that the key k required by the non-colliding encryption scheme is the same length (b bits) as the key for the block cipher E. If the block-cipher keys are too short, they can be stretched using a pseudorandom generator, and if they are too long, they can be truncated. Recall that F is a pseudorandom function.

Lemma 4. Let $C_k(V,h) = E_{F_k(h)}(V)$ and $\hat{C}_k(V,h) = D_{F_k(h)}(V)$. Then $(C_k, \hat{C}_k)$ is a non-colliding encryption scheme.

Proof. This construction obviously satisfies Property 1 of Definition 2.

To prove that is also satisfies Property 2, consider an adaptive adversary $\mathcal{A}$ that makes a sequence of m queries to C_k and $\hat{C}_k$. Let ε_m^1 (resp. ε_m^2) be an upper bound on the probability that, with m queries, $\mathcal{A}$ can distinguish F from a truly random function (resp. an upper bound on the probability that, with m queries, $\mathcal{A}$ can distinguish a collection of $m+2$ pseudorandom permutations from a collection of $m+2$ truly random permutations). Then $\mathcal{A}$'s chance of finding two pairs $(V,h) \neq (V',h')$ such that (a) $C_k(V,h) = C_k(V',h')$, and (b) at least one of (V,h) and (V',h') did not appear in the sequence of queries is at most

$$\frac{m^2}{2^b} + \frac{1}{2^b - m} + \varepsilon_m^1 + \varepsilon_m^2. \tag{1}$$

In see this, compare $(C_k, \hat{C}_k)$ to the following process $(C', \hat{C}')$, which is defined in terms of *random* functions and permutations rather than pseudorandom functions and permutations. Let $E_1, E_2, \ldots, E_{m+2}$ be random permutations, and let $D_1, D_2, \ldots, D_{m+2}$ be the corresponding inverse permutations. The process $(C', \hat{C}')$ acts as follows on $\mathcal{A}$'s i^{th} query (V_i, h_i), given that (V_1,h_1), (V_2,h_2), ..., (V_{i-1},h_{i-1}) were $\mathcal{A}$'s $i-1$ previous queries. Suppose that h_i is the jth distinct element in the set $\{h_1, h_2, \ldots, h_i\}$. If (V_i,h_i) is a query is to C', then respond with $E_j(V_i)$, and if it is a query to $\hat{C}'$, then respond with $D_j(V_i)$.

We would like to bound the probability that, after m queries, $\mathcal{A}$ can find $(V,h) \neq (V',h')$ such that $C'(V,h) = C'(V',h')$ but at least one of (V,h) or (V',h') did not appear in the sequence (V_1,h_1), (V_2,h_2), ..., (V_m,h_m). Suppose that h is the j_1^{th} distinct element and h' is the j_2^{th} distinct element among $h_1, h_2, \ldots, h_m, h, h'$. If $j_1 = j_2$, we are done, because $E_{j_1}(V) \neq E_{j_1}(V')$. Otherwise, assume without loss of generality that it is the query (V,h) that did not

appear in the sequence. Then the probability that $E_{j_1}(V) = E_{j_2}(V')$ is at most $1/(2^b - m)$, because the value of E_{j_1} has been specified on at most m points, and $E_{j_1}(V)$ is uniformly distributed among the remaining $2^b - m$ points in the range.

We now bound the probability that a polynomial-time adversary can distinguish between (truly random) $(C', \hat{C}')$ and (pseudorandom) $(C_k, \hat{C}_k)$. Essentially, we use a hybrid argument. If instead of the pseudorandom F_k, a truly random function f were used, the probability that the adversary could find two different h and h' such that $f(h) = f(h')$ would be at most $m^2/2^b$. If this does not happen, then the randomness of f implies that the keys of the pseudorandom permutations are random; if the adversary could distinguish between such a process and $(C', \hat{C}')$, then it could distinguish between a collection of m pseudorandom permutations and a collection of m truly random permutations – this happens with probably at most ε_m^2. Distinguishing between the case in which a random f is used and the one in which a pseudorandom F_k is used adds probability at most ε_m^1, yielding (1).

A-good sequences:

Recall that we would like these to be the sequences in which different X^i's correspond to different W^i's and different Y^i's correspond to different Z^i's. Furthermore, the W's and Z's of the distinguishing phase should be different from those obtained during the host phase, except in those inputs filtered by the arbiter. Intuitively, these sequences are "good" for the pseudorandom permutation construction, because distinct W's produce distinct S's with overwhelming probability. The properties of the building blocks E, F, and G then ensure that there is a process $\tilde{A}$, indistinguishable from A, such that, for all A-good input-output sequences $\langle (X^1, Y^1), \ldots, (X^m, Y^m) \rangle$

$$\Pr_{\tilde{A}}[p_2(Y^i) = \tilde{A}(p_1(X^i)), 1 \leq i \leq m] \approx 2^{-\ell \cdot m}.$$

More precisely, let $\mathcal{A}$ be a probabilistic, polynomial-time adversary that has m_1 interactions with the card during the host phase and makes m_2 oracle queries during the distinguishing phase. The sequences we are interested in consist of

$$(X_1^1, h_x^1, Y_1^1, h_y^1), (X_1^2, h_x^2, Y_1^2, h_y^2), \ldots, (X_1^{m_1}, h_x^{m_1}, Y_1^{m_1}, h_y^{m_1})$$

from the host phase and

$$(X^{m_1+1}, Y^{m_1+1}), (X^{m_1+2}, Y^{m_1+2}), \ldots, (X^{m_1+m_2}, Y^{m_1+m_2})$$

from the distinguishing phase. For any such sequence, the permutations p_1 and p_2 determine W^1, ..., W^{m_1}, W^{m_1+1}, ..., $W^{m_1+m_2}$ and Z^1, ..., Z^{m_1}, Z^{m_1+1}, ..., $Z^{m_1+m_2}$. An encryption query $(+, X^{m_1+i})$ or decryption query $(-, Y^{m_1+i})$ is filtered during the distinguishing phase by $\mathcal{B}$ if there is a $1 \leq j \leq m_1$ for which

$$(X_1^{m_1+i}, H(X_2^{m_1+i}, \ldots, X_n^{m_1+i})) = (X_1^j, h_x^j)$$

(or analogously $(Y_1^{m_1+i}, H(Y_2^{m_1+i}, \ldots, Y_n^{m_1+i})) = (Y_1^j, h_y^j)$). Note that the adversary should not be able to find more than m_1 inputs and m_1 outputs that are filtered – otherwise, the pigeonhole principle implies that the adversary would have found in the distinguishing phase two encryption queries $(+, X^{m_1+j_1}) \neq (+, X^{m_1+j_2})$ such that $(X_1^{m_1+j_1}, h_x^{m_1+j_1}) = (X_1^{m_1+j_2}, h_x^{m_1+j_2}) = (X_i, h_x^i)$ for some $1 \leq i \leq m_1$ and $1 \leq j_1, j_2 \leq m_2$ (or two analogous decryption queries). However, that would mean that it had broken the collision-intractable hash function H.

We say that a sequence is A-good for p_1 and p_2 if

1. For all $1 \leq i < j \leq m_2$, if $X^{m_1+i} \neq X^{m_1+j}$ and X^{m_1+i} and X^{m_1+j} are not filtered by $\mathcal{B}$, then $W^{m_1+i} \neq W^{m_1+j}$, and, if $Y^{m_1+i} \neq Y^{m_1+j}$ and Y^{m_1+i} and Y^{m_1+j} are not filtered by $\mathcal{B}$, then $Z^{m_1+i} \neq Z^{m_1+j}$.
2. For all $1 \leq i \leq m_2$, if X^{m_1+i} is not filtered by $\mathcal{B}$, then $W^{m_1+i} \neq W^j$ for all $1 \leq j \leq m_1$, and, if Y^{m_1+i} is not filtered by $\mathcal{B}$, then $Z^{m_1+i} \neq Z^j$ for all $1 \leq j \leq m_1$.

In other words, the W's and Z's of the distinguishing phase are different from one another and from those of the host phase. We must show that the adversary is not able to find bad sequences, except with negligible probability.

Lemma 5. For any permutation A, for any probabilistic, polynomial-time adversary $\mathcal{A}$ that has m_1 interactions with the card during the host phase and makes m_2 oracle queries during the distinguishing phase, the probability that $\mathcal{A}$ finds a sequence that is not A-good for p_1 and p_2 is negligible. The probability is computed over the choice of p_1 and p_2 and the random coin-flips of $\mathcal{A}$. Note that A is not necessarily secret.

Proof. Let ε_1 be an upper bound on the probability that an adversary with $\mathcal{A}$'s resources breaks the collision-intractable hash function H, and let ε_2 be an upper bound on the probability that an adversary with $\mathcal{A}$'s resources breaks the non-colliding encryption scheme. Then $\mathcal{A}$'s probability of finding a sequence that is not A-good for p_1 and p_2 is upper-bounded by $\varepsilon_1 + \varepsilon_2$.

Suppose that the first query that witnesses the fact that this sequence is not A-good occurs at the jth step of the distinguishing phase, and assume without loss of generality that it is an encryption query. We divide this event into two cases. In case 1, Property 1 is violated, *i.e.*, there are i and j, $1 \leq i < j \leq m_2$, such that $X^{m_1+i} \neq X^{m_1+j}$ but $X_1^{m_1+i} = X_1^{m_1+j}$ and $H(X_2^{m_1+i}, \ldots, X_n^{m_1+i}) = H(X_2^{m_1+j}, \ldots, X_n^{m_1+j})$. This means that $\mathcal{A}$ has broken the collision-intractable function H, which happens with probability at most ε_1. The other possibility is that $Z^{m_1+j} = Z^i$ but $(X_1^{m_1+j}, h_x^{m_1+j}) \neq (X_1^i, h_x^i)$ where i $(1 \leq i \leq m_1 + j - 1)$ is either from the host phase or from the distinguishing phase. However, note that $(X_1^{m_1+j}, h_x^{m_1+j})$ did not appear explicitly before in the sequence (if it had, it would have been filtered, or j would not be the first "bad location" in the sequence), but $C_{k_1}(X_1^{m_1+j}, h_x^{m_1+j}) = C_{k_1}(X_1^i, h_x^i)$. Thus $\mathcal{A}$ could break the non-colliding encryption scheme $(C_k, \hat{C}_k)$, which happens with probability at most ε_2.

Indistinguishability:

As in the original Naor-Reingold paper, we consider what happens when A is replaced with a "more random" process. Let $\tilde{A}_1$ be the obtained from A by replacing E_{k_3} with a random permutation and replacing G_S with a process $\tilde{G}$ that, on input S, produces a random string of length $(n-1)\cdot b$ (*i.e.*, a random function $\{0,1\}^b \mapsto \{0,1\}^{(n-1)\cdot b}$) and Xors the string with $X_2, \ldots, X_n$. Let $\tilde{A}_2$ be a random permutation. Note that, if A is replaced with $\tilde{A}_2$, the composition $p_2^{-1} \circ \tilde{A}_2 \circ p_1$ is a random permutation for any p_1 and p_2.

We complete the proof of Theorem 3 by showing (i) when A is replaced with $\tilde{A}_1$, the result is indistinguishable by probabilistic, polynomial-time adversaries, and (ii) when $\tilde{A}_1$ is replaced by $\tilde{A}_2$ the result is indistinguishable to adversaries restricted to good sequences.

Lemma 6. Suppose that, following the host phase of an attack on Scheme P, a nondeterministic choice is made between replacing the function A by $\tilde{A}_1$ or continuing to use A. Then any probabilistic, polynomial-time adversary $\mathcal{A}$ has only a negligible probability of determining whether a switch was made, where the probability is over the choice of $A, p_1, p_2, \tilde{A}_1$ and $\mathcal{A}$'s coin-flips.

Proof. Observe first that A can be partly transformed without detection: Suppose that E_{k_3} is replaced with a random permutation and F_{k_4} is replaced with a random function prior to the beginning of the host phase. This should be indistinguishable to $\mathcal{A}$, because the only information $\mathcal{A}$ has about E_{k_3}, D_{k_3}, and F_{k_4} is their values (or some function of them) at some specific points. Therefore the important part of a potential switch is the replacement of G by $\tilde{G}$.

By Lemma 5, except with negligible probability, all the W's of the distinguishing phase are different from those of the host phase and different from each other, except those that were filtered. These W^i's are assigned a random value S^i. Recall that G has the property that, if S is chosen at random, then $G_S(X_2, \ldots, X_n)$ is indistinguishable from a truly random string of the same length for any $X_2, \ldots, X_n$, and similarly for $\hat{G}$. In case a switch is not made between phases, then G (or $\hat{G}$) produces an input that is indistinguishable from a truly random one. If a switch is made, then $\tilde{G}$ is used, and the result is a random and independent string (except when S_i is a collision, which happens with probability at most $m^2/2^b$). Therefore, the overall probability with which a switch is detected is negligible.

Lemma 7. Suppose that, following the host phase of an attack on Scheme P, a nondeterministic choice is made between replacing A by $\tilde{A}_1$ or replacing A by $\tilde{A}_2$. Then any probabilistic, polynomial-time adversary $\mathcal{A}$ has a negligible probability of distinguishing between the two cases, where the probability is computed over the choice of $A, p_1, p_2, \tilde{A}_1, \tilde{A}_2$, and $\mathcal{A}$'s coin-flips.

Proof. Fix the adversary $\mathcal{A}$ to be the best deterministic machine, and fix p_1, p_2, and A. This determines the queries made during the host phase. Now consider any sequence

$$SEQ = \langle (X^{m_1+1}, Y^{m_1+1}), \ldots, (X^{m_1+m_2}, Y^{m_1+m_2}) \rangle$$

such that SEQ (together with the host phase queries) is A-good for p_1 and p_2. If SEQ is a *possible* outcome for $\mathcal{A}$, given the fixed A, p_1, and p_2, then

$$\Pr_{\tilde{A}_1}[SEQ \text{ is produced}] = \frac{1}{2^{bn}} \cdot \frac{1}{(2^b - 1) \cdot 2^{b(n-1)}} \cdots \frac{1}{(2^b - m_2 + 1) \cdot 2^{b(n-1)}}$$

and

$$\Pr_{\tilde{A}_2}[SEQ \text{ is produced}] = \frac{1}{2^{bn}} \cdot \frac{1}{2^{bn} - 1} \cdots \frac{1}{2^{bn} - m_2 + 1}.$$

Therefore, $\Pr_{\tilde{A}_1}[SEQ \text{ is produced}]$ is at least

$$\Pr_{\tilde{A}_2}[SEQ \text{ is produced}] \geq (1 - \frac{m_2}{2^b})^{m_2} \cdot \Pr_{\tilde{A}_1}[SEQ \text{ is produced}].$$

Now consider any collection $\mathcal{C}$ of sequences that are A-good for p_1 and p_2. The probabilities that a member of this collection is produced by $\tilde{A}_1$ and $\tilde{A}_2$ are close, because the ratio of $\Pr_{\tilde{A}_1}[\mathcal{C}]$ and $\Pr_{\tilde{A}_2}[\mathcal{C}]$ is between 1 and $1 - (m_2^2/2^b)$.

Suppose without loss of generality that, just before $\mathcal{A}$ has to guess whether it is querying process $\tilde{A}_1$ or process $\tilde{A}_2$, the permutations p_1 and p_2 are revealed. This can only help $\mathcal{A}$. If we consider all possible executions of $\mathcal{A}$, then with all but negligible probability (over p_1, p_2), the sequence generated is A-good for p_1 and p_2. Partition the executions that end with an A-good sequence into $\mathcal{C}_1$ (those for which $\mathcal{A}$ announces 1), and $\mathcal{C}_2$ (those where it announces 2). Summing over p_1, p_2, and A, the probabilities $\Pr_{\tilde{A}_1}[\mathcal{C}_1]$ and $\Pr_{\tilde{A}_2}[\mathcal{C}_1]$ are close (and similarly for $\Pr_{\tilde{A}_1}[\mathcal{C}_2]$ and $\Pr_{\tilde{A}_2}[\mathcal{C}_2]$). Therefore, we can conclude that $\tilde{A}_1$ and $\tilde{A}_2$ are indistinguishable for $\mathcal{A}$.

To complete the proof of Theorem 3, note that $p_2^{-1} \circ \tilde{A}_2 \circ p_1$ is a random permutation. Thus, $\mathcal{A}$ has at most a negligible probability of determining whether A was switched with a random permutation between phases.

5 Open Questions

Two remaining questions include: (1) The protocols in Section 4 require two rounds of interaction between the host and the card. Is there a secure RKES that requires only one round of interaction? (2) Is the existence of a one-way function sufficient for the construction of a provably secure RKES? Note that a collision-intractable hash function is used in our construction and that it is not known how to build such a hash function based only on the assumption that a one-way function exists. See [4] for a recent discussion of the desirability of using UOWHF instead of collision-intractable hash functions.

Acknowledgments

We thank Omer Reingold for useful discussions and the Eurocrypt '98 PC members for their comments.

References

1. M. Bellare, A. Desai, E. Jokipii, and P. Rogaway, "A Concrete Security Treatment of Symmetric Encryption," in *Proceedings of the 38th Symposium on Foundation of Computer Science*, IEEE Computer Society Press, Los Alamitos, pp. 394–403, 1997.
2. M. Bellare, J. Kilian, and P. Rogaway, "The Security of Cipher Block Chaining," in *Advances in Cryptology – Crypto '94*, Lecture Notes in Computer Science, vol. 839, Springer, Berlin, pp. 341–358, 1994.
3. M. Bellare and P. Rogaway, "Provably Secure Session Key Distribution – The Three Party Case," in *Proceedings of the 27th Symposium on Theory of Computing*, ACM, New York, pp. 57–66, 1995.
4. M. Bellare and P. Rogaway, "Collision Resistant Hashing, Towards Making UOWHFs practical," in *Advances in Cryptology – Crypto '97*, Lecture Notes in Computer Science, vol. 1294, Springer, Berlin, pp. 470–484, 1997.
5. E. Biham and A. Shamir, "Differential Fault Analysis of Secret Key Cryptosystems," in *Advances in Cryptology – Crypto '97*, Lecture Notes in Computer Science, vol. 1294, Springer, Berlin, pp. 513–525, 1997.
6. M. Blaze, "High-Bandwidth Encryption with Low-Bandwidth Smartcards," in *Proceedings of the Fast Software Encryption Workshop*, Lecture Notes in Computer Science, vol. 1039, Springer, Berlin, pp. 33–40, 1996.
7. D. Boneh, R. A. Demillo, and R. J. Lipton, "On the Importance of Checking Protocols for Faults," in *Advances in Cryptology – Eurocrypt '97*, Lecture Notes in Computer Science vol. 1233, Springer, Berlin, pp. 37–51, 1997.
8. J. Feigenbaum, "Locally Random Reductions in Interactive Complexity Theory," in *Advances in Computational Complexity Theory*, DIMACS Series on Discrete Mathematics and Theoretical Computer Science, vol. 13, American Mathematical Society, Providence, 1993, pp. 73–98.
9. O. Goldreich, **Foundations of Cryptography (Fragments of a Book)**, 1995. `http://www.eccc.uni-trier.de/eccc/info/ECCC-Books/eccc-books.html`
10. O. Goldreich S. Goldwasser, and S. Micali, "How to Construct Random Functions," *J. of the ACM*, 33 (1986), pp. 792-807.
11. M. Luby, **Pseudorandomness and Cryptographic Applications**, Princeton University Press, Princeton, 1996.
12. S. Lucks, "On the Security of Remotely Keyed Encryption," in *Proceedings of the Fast Software Encryption Workshop*, Lecture Notes in Computer Science, vol. 1267, Springer, Berlin, pp. 219–229, 1997.
13. T. Matsumoto, K. Kato, and H. Imai, "Speeding Up Secret Computations with Insecure Auxiliary Devices," in *Advances in Cryptology – Crypto '88*, Lecture Notes in Computer Science, vol. 403, Springer, Berlin, pp. 497–506, 1990.
14. M. Naor and O. Reingold, "On the Construction of Pseudo-Random Permutations: Luby-Rackoff Revisited," to appear in *J. Cryptology*. Extended abstract appears in *Proceedings of the 29th Symposium on Theory of Computing*, ACM, New York, pp. 189–199, 1997.
15. V. Shoup and A. Rubin, "Session Key Distribution Using Smart Cards," in *Advances in Cryptology – Eurocrypt '96*, Lecture Notes in Computer Science vol. 1070, Springer, Berlin, pp. 321–331, 1996.