Eliminating Random Permutation Oracles in the Even-Mansour Cipher

Craig Gentry and Zulfikar Ramzan

DoCoMo Communications Laboratories USA, Inc.
{cgentry, ramzan}@docomolabs-usa.com

Abstract. Even and Mansour [EM97] proposed a block cipher construction that takes a publicly computable random permutation oracle P and XORs different keys prior to and after applying P: $C = k_2 \oplus P(M \oplus k_1)$. They did not, however, describe how one could instantiate such a permutation securely. It is a fundamental open problem whether their construction could be proved secure outside the random permutation oracle model. We resolve this question in the affirmative by showing that the construction can be proved secure in the random *function* oracle model. In particular, we show that the random permutation oracle in their scheme can be replaced by a construction that utilizes a four-round Feistel network (where each round function is a random function oracle publicly computable by all parties including the adversary). Further, we prove that the resulting cipher is super pseudorandom – the adversary's distinguishing advantage is at most $2q^2/2^n$ if he makes q total queries to the cipher, its inverse, as well as any random oracles. Even and Mansour, on the other hand, only showed security against inversion and forgery. *One noteworthy aspect of this result is that the cipher remains secure even though the adversary is permitted separate oracle access to all of the round functions.* One can achieve a two-fold and four-fold reduction respectively in the amount of key material by a closer inspection of the proof and by instantiating the scheme using group operations other than exclusive-OR. On the negative side, a straightforward adaption of an advanced slide attack recovers the $4n$-bit key with approximately $\sqrt{2} \cdot 2^n$ work using roughly $\sqrt{2} \cdot 2^n$ known plaintexts. Finally, if only three Feistel rounds are used, the resulting cipher is pseudorandom, but not super pseudorandom.

1 Introduction

THE EVEN-MANSOUR CONSTRUCTION. Even and Mansour [EM97] proposed a block cipher construction based on XORing secret key material just prior to and just after applying a random permutation oracle P: $C = k_2 \oplus P(M \oplus k_1)$, where M is the plaintext, C is the ciphertext, and k_1, k_2 is the key material. The permutation P (as well as its inverse P^{-1}) is computable by all parties, including the adversary (see fig. 1). Even-Mansour proved that a polynomial-time adversary with black-box query access to the cipher and its inverse, as well as black-box query access to the internal permutation and its inverse cannot invert an un-queried ciphertext of his choice, except with negligible probability. They also proved an analogous result about computing the cipher's forward direction.

While there are practical limitations to their construction [Dae91, BW00], the Even-Mansour work is well known and theoretically interesting. In particular, it is an example of a cipher for which an adversary has black-box access to the only real "cryptographic" component; i.e., the random permutation oracle. The only secrets are simply XORed at the beginning and the end, and everything else is publicly accessible.

Fundamental Open Problems. The Even-Mansour work may be described within the framework of the random-oracle model [BR93] in which their cipher makes use of a random *permutation* oracle. Naturally, the need for such a permutation oracle is unpleasant, especially since Even and Mansour did not indicate how one might instantiate such a random permutation oracle while maintaining security. This motivates the following problem:

Open Problem 1: How can one go about instantiating the random permutation oracle in the Even-Mansour scheme?

Furthermore, Even and Mansour only proved security against inversions and forgeries. However, for block ciphers, the current bar is to prove super pseudorandomess [LR88]. That is, the cipher should be indistinguishable from a randomly chosen permutation on the same message space even if the adversary is granted black-box access to the forward and inverse directions of the cipher.[1] This motivates a second problem:

Open Problem 2: Can one prove that an Even-Mansour type construction yields a super pseudorandom permutation?

Our Contributions. We address the first question by demonstrating that the random permutation oracle can be replaced by a construction involving random *function* oracles; i.e., the underlying oracle (which must be accessible to all parties) does not have to be bijective, but we construct a permutation using it that is bijective. By supplanting the use of random permutation oracles by random function oracles, we have a result based on a less restrictive model. Our construction uses a Feistel ladder in which the random function oracle is used as a round function (see fig. 1). *However, what is different in this setting is that the adversary not only has access to the forward and reverse directions of the cipher, but also to each of the individual round functions.*

We address the second problem by proving that the construction is *super pseudorandom*. We remark the one can construe the Kilian-Rogaway analysis of DESX [KR96] as a proof that Even-Mansour is pseudorandom. Recall that in DESX, the Even-Mansour random permutation is supplanted with a *keyed*

[1] Their is also a notion of *pseudorandomness* for block ciphers wherein the adversary must distinguish it from a random permutation given black-box access to only the forward direction of the cipher.

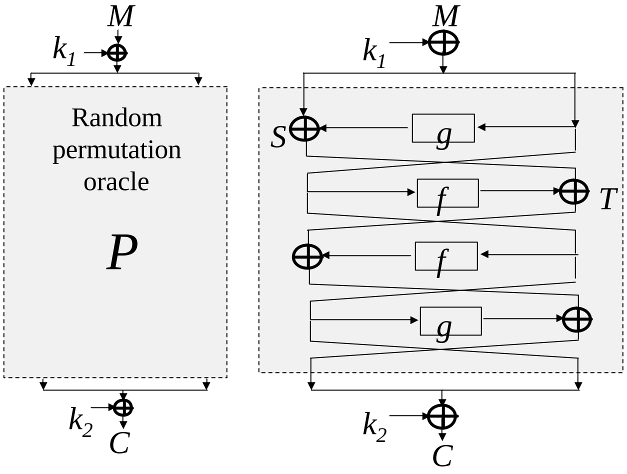

Fig. 1. The diagram on the left depicts the Even-Mansour scheme where P is a random *permutation* oracle; i.e., the adversary has black-box access to P and P^{-1}. The diagram on the right depicts our scheme in which the permutation oracle is instantiated by a Feistel network consisting of *publicly-accessible* random *function* oracles f, g.

block cipher, such as DES. The Kilian-Rogaway proof allowed the adversary oracle access to the internal permutation P (modeled as an ideal block cipher) as well as P^{-1}, to simulate that an adversary had correctly guessed the key – this maneuver isolates the benefits of the pre- and post-whitening keys. However, in their published proof the adversary was not given access to the inverse of the block cipher – so *super* pseudorandomness *was not* proved.[2]

In addition, Ramzan-Reyzin [RR00] noted that one could use their round security framework to prove that Even-Mansour is super pseudorandom, but their focus was different, so no proof was supplied. Also their comment was limited to the original Even-Mansour construction (which used the random permutation oracle). Therefore, we consider addressing the first fundamental open problem as our main technical contribution; a side benefit of our work is a proof of super-pseudorandomness for Even-Mansour style block ciphers.

Our results help us better understand block cipher design. First, they point to the benefit of pre- and post- whitening. In particular, our construction shows that, in the random function oracle model, one can construct a super pseudorandom block cipher in which the all key material is only incorporated during the pre- and post-whitening phases and in a very simple way. This is despite the fact that the adversary has access to the internals of the cipher. Second, our constructions show that it may be possible to obtain a middle ground between pure black-box analysis and one in which an adversary has some meaningful

[2] Kilian and Rogaway mentioned that one could extend their proof to address chosen ciphertext queries, however, they did not provide the proof, nor did they state a formal security theorem where such access is given.

knowledge about the internal design of the black box. This can be thought of as a "gray-box" analysis. We also remark that the random permutation oracle model seems less appealing than the random function oracle model. Instantiating a random function oracle while maintaining security seems more plausible since such functions could be made sufficiently complex that their behavior is ill understood. On the other hand, when instantiating a random permutation oracle with an actual permutation, one is limited in the complexity of the design since the function must remain bijective and efficient to invert. Our results give hope that one may be able to base future cryptosystems on random permutation oracles and replace them with constructions based on random function oracles in a provably secure way. Finally, our work helps bridge the gap between the theory and practice of Feistel ciphers. In particular, the theoretical work on Feistel ciphers (e.g., [LR88]) considers round functions that are strong (e.g., pseudorandom) and potentially complex keying mechanisms (e.g., the functions themselves are keyed). This departs from practice in two ways. First, round functions in practice are weak. Second, block cipher round keys are introduced in some simple way, for example by XORing them prior to applying an un-keyed function (c.f., DES [FIPS46]). Our work sits somewhere in between since it considers complex round functions (random oracles), but simple keying procedures (XORing). Therefore, we can view our work as providing better mathematical insight into the security of DES-like ciphers.

Other Results. Our proof of security holds even if the amount of key material is reduced twofold. Also, if we permit group operations other than XOR, we can recycle keying material, yielding a fourfold reduction; interestingly, if XOR is used with recycled keying material, the cipher behaves like an involution and is trivially distinguishable from a random permutation. This idea of considering different group operations has previously been applied to Luby-Rackoff ciphers [PRS02]. On the negative side, a "sliding with a twist" attack [BW00] allows an adversary to recover the key using $\sqrt{2} \cdot 2^n$ *known* plaintexts and $\sqrt{2} \cdot 2^n$ work. Finally, if we instantiate the permutation with three Feistel rounds, the construction is pseudorandom, but is not super pseudorandom. The attack adapts the standard distinguisher for three-round Luby-Rackoff ciphers [LR88]. Due to space constraints, as well as the fact that these results follow easily from existing techniques, we omit a further discussion. For details, see the full version of the paper [GR04].

Caveat(s) Emptor. While the random-oracle model is an extremely useful cryptographic tool, there are instances of schemes that are secure in the random oracle model, but are insecure for any instantiation of the random oracle by a polynomial-time computable function [CGH98, GK03, BBP04]. We further note that the lower bounds we present indicate that n should be chosen so that $2^{n/2}$ is sufficiently large to thwart distinguishing attacks. We also remark that Even and Mansour gave a $O(2^{-n})$ upper bound on the adversary's success probability, whereas our bound resembles $O(2^{-n/2})$. However, Even and Mansour only proved security against inversions and forgeries whereas we show super pseudorandomness. Moreover, we eliminate the random permutation oracle requirement and also give the adversary access to the innards of the cipher. Therefore, we expect there to be a sizeable gap in the respective security guarantees. In light of these caveats, *we stress that our main contribution is in resolving fundamental issues from the Even-Mansour work and gaining more theoretical insight into block cipher design; we do not recommend this as a practical approach to building a block cipher.* In fact, efficient random oracle model based block ciphers are desired, then Ramzan and Reyzin have a four-round Feistel block cipher construction in which the middle two rounds use a random oracle, and the outer two rounds involve universal hash functions [RR00].

Organization. Section 2 reviews prior definitions and constructions. Section 3 discusses our main construction and security proof. Finally, we make concluding remarks in Section 4.

2 Prior Definitions and Constructions

We describe definitions and prior constructions that are germane to our work. We avoid asymptotic analysis in favor of the "concrete" (or "exact") security model as laid out by Bellare, Kilian, and Rogaway [BKR94], and Bellare, Canetti, Krawczyk [BCK96]. However, our results can be adapted to either model.

Notation. For a natural number n, we let I_n denote the set of bit strings of length n: $\{0,1\}^n$. For a bit string x, we let $|x|$ denote its length. If $|x|$ is even, then x^L and x^R denote the left and right halves of the bits respectively; we sometimes write $x = (x^L, x^R)$. If x and y are two bit strings with $|x| = |y|$, we denote by $x \oplus y$ their bitwise exclusive OR. If S is a probability space, then $x \overset{R}{\leftarrow} S$ denotes the process of picking an element from S according to the underlying probability distribution. Unless otherwise specified, the underlying distribution is assumed to be uniform. By a finite function (or permutation) family $\mathcal{F}$, we denote a set of functions with common domain and common range. Let $\mathsf{Rand}^{k \to \ell}$ be the set of all functions going from I_k to I_ℓ, and let Perm^m be the set of all permutations on I_m. We call a finite function (or permutation) family *keyed* if every function in it can be specified (not necessarily uniquely) by a key a. We denote the function given by a as f_a. We assume that given a, it is possible to efficiently evaluate f_a at any point (as well as f_a^{-1} in case of a keyed permutation family). For a given keyed function family, a key can be any string from I_s, where s is known as "key length." (Sometimes it is convenient to have keys from a set other than I_s; we do not consider such function families simply for clarity of exposition, but our results continue to apply in such cases.) For functions f and g, $g \circ f$ denotes the function $x \mapsto g(f(x))$.

Model of Computation. We model the adversary $\mathcal{A}$ as a program for a Random Access Machine (RAM) with black-box access to some number k of oracles, each computing some specified function. If $(f_1, \ldots, f_k)$ is a k-tuple of functions, then $\mathcal{A}^{f_1,\ldots,f_k}$ denotes a k-oracle adversary who is given black-box oracle access to each of the functions $f_1, \ldots, f_k$. We define $\mathcal{A}$'s "running time" to be the number of time steps it takes plus the length of its description (to prevent one from embedding arbitrarily large lookup tables in $\mathcal{A}$'s description).

Pseudorandom Functions and Block Ciphers. The pseudorandomness of a keyed function family $\mathcal{F}$ with domain I_k and range I_ℓ captures its computational indistinguishability from $\mathsf{Rand}^{k\to\ell}$. The following definition is adapted from [GGM84]:

Definition 1. *A pseudorandom function family $\mathcal{F}$ is a keyed function family with domain I_k, range I_ℓ, and key length s. Let $\mathcal{A}$ be a 1-oracle adversary. Then we define $\mathcal{A}$'s advantage as*

$$\mathsf{Adv}^{\mathsf{prf}}_{\mathcal{F}}(\mathcal{A}) \triangleq \left|\Pr[a \overset{R}{\leftarrow} I_s : \mathcal{A}^{f_a} = 1] - \Pr[f \overset{R}{\leftarrow} \mathsf{Rand}^{k\to\ell} : \mathcal{A}^{f} = 1]\right|.$$

For any integers $q, t \geq 0$, we define $\mathsf{Adv}^{\mathsf{prf}}_{\mathcal{F}}(q,t) \triangleq \max_{\mathcal{A}}\{\mathsf{Adv}^{\mathsf{prf}}_{\mathcal{F}}(\mathcal{A})\}$, as an insecurity function, where the maximum is taken over choices of adversary $\mathcal{A}$ such that:

- *$\mathcal{A}$ makes at most q oracle queries, and*
- *the running time of $\mathcal{A}$, plus the time necessary to select $a \overset{R}{\leftarrow} I_s$ and answer $\mathcal{A}$'s queries, is at most t.*

Recall that the Even-Mansour cipher [EM97] operates on a $2n$-bit string x as follows $E(x) = k_2 \oplus P(x \oplus k_1)$ where $k_1, k_2 \in I_{2n}$ constitutes the keying material and P is a random *permutation* oracle. Here P and P^{-1} are publicly computable (in a black-box fashion) by all parties. Even and Mansour proved that E is hard to invert on a point C_0 of the adversary's choice even if the adversary has oracle access to E, E^{-1}, P, P^{-1} subject to the restriction that the adversary cannot query the E^{-1} oracle on the point C_0; i.e., it is hard to find M_0 such that $M_0 = E^{-1}_{k_1,k_2}(C_0)$. Similarly, they showed that the adversary cannot compute the ciphertext corresponding to a message point M_0 of its choice with access to these same oracles, but this time subject to the restriction that the adversary cannot query the E oracle on point M_0; i.e., it is hard to find C_0 such that $C_0 = E_{k_1,k_2}(M_0)$. While these results capture some of the security requirements needed for a block cipher, there are stronger notions of security for a block cipher. One such notion, proposed by Luby and Rackoff [LR88], is called *super pseudorandomness.* The notion captures the pseudorandomness of a permutation family on I_ℓ in terms of its indistinguishability from Perm^ℓ, where the adversary is given access to both directions of the permutation. In other words, it measures security of a block cipher against chosen plaintext and ciphertext attacks. We now describe such notions and how to achieve them.

Definition 2. *A block cipher $\mathcal{P}$ is a keyed permutation family with domain and range I_ℓ and key length s. Let $\mathcal{A}$ be a 2-oracle adversary. Then we define $\mathcal{A}$'s advantage as*

$$\mathsf{Adv}^{\mathsf{sprp}}_{\mathcal{P}}(\mathcal{A}) \triangleq \left|\Pr[a \overset{R}{\leftarrow} I_s : \mathcal{A}^{p_a, p_a^{-1}} = 1] - \Pr[p \overset{R}{\leftarrow} \mathsf{Perm}^\ell : \mathcal{A}^{p,p^{-1}} = 1]\right|.$$

For any integers $q, t \geq 0$, $\mathsf{Adv}_F^{\mathsf{sprp}}(q,t)$ specifies the insecurity function (analogous to Definition 1).

Luby and Rackoff showed how to construct a secure block cipher using Feistel permutations.

Definition 3 (Basic Feistel Permutation). *Let $\mathcal{F}$ be a function family with domain and range I_n. Let $f \in \mathcal{F}$. Let $x = (x^L, x^R)$ with $x^L, x^R \in I_n$. We denote by $\overline{f}$ the permutation on I_{2n} defined as $\overline{f}(x) = (x^R, x^L \oplus f(x^R))$. Note that it is a permutation because $\overline{f^{-1}}(y) = (y^R \oplus f(y^L), y^L)$. Similarly, let $\overline{\mathcal{F}} = \{\overline{f} \mid f \in \mathcal{F}\}$.*

Definition 4 (Feistel Network). *If $f_1, \ldots, f_s$ are mappings with domain and range I_n, then we denote by $\Phi(f_1, \ldots, f_s)$ the permutation on I_{2n} defined as $\Phi(f_1, \ldots, f_s) = \overline{f_s} \circ \ldots \circ \overline{f_1}$.*

Theorem 1 (Luby-Rackoff). *Let h_1, f_1, f_2, h_2 be independently-keyed functions from a keyed function family $\mathcal{F}$ with domain and range I_n and key space I_s. Let $\mathcal{P}$ be the family of permutations on I_{2n} with key space I_{4s} defined by $\mathcal{P} = \Phi(h_1, f_1, f_2, h_2)$ (the key for an element of $\mathcal{P}$ is simply the concatenation of keys for h_1, f_1, f_2, h_2). Then, $\mathsf{Adv}_{\mathcal{P}}^{\mathsf{sprp}}(q,t) \leq \mathsf{Adv}_{\mathcal{F}}^{\mathsf{prf}}(q,t) + \binom{q}{2}\left(2^{-n+1} + 2^{-2n+1}\right)$.*

The Luby-Rackoff result proved security when the adversary has access to the permutation and its inverse. In our case, we will show security of the Even-Mansour cipher when the adversary has black-box access to the cipher, its inverse, and to each of the internal round functions.

Having presented the classical definitions of block ciphers and Feistel networks, we now describe notions of the Ramzan-Reyzin round security framework [RR00] which we make use of in the present work.

Definition 5 (Round Decomposition [RR00]). *Let $\mathcal{P}, \mathcal{F}^1, \mathcal{F}^2, \ldots, \mathcal{F}^r$ be keyed permutation families, each with domain and range I_ℓ and key length s, such that for any key $a \in I_s$, $p_a = f_a^r \circ \cdots \circ f_a^1$. Then $\mathcal{F}^1, \ldots, \mathcal{F}^r$ is called an* r-round decomposition *for $\mathcal{P}$. For $i \leq j$, denote by $(i \rightarrow j)_a$ the permutation $f_a^j \circ \ldots \circ f_a^i$, and by $(i \leftarrow j)_a$ the permutation $\left(f_a^j \circ \ldots \circ f_a^i\right)^{-1}$. Denote by $i \rightarrow j$ and $i \leftarrow j$ the corresponding keyed function families.*

Note that having oracle access to a member of $i \rightarrow j$ means being able to give inputs to round i of the forward direction of a block cipher and view outputs after round j. Likewise, having oracle access to $i \leftarrow j$ corresponds to being able to give inputs to round j of the *reverse* direction of the block cipher and view outputs after round i. Thus, the oracle for $1 \rightarrow r = \mathcal{P}$ corresponds to the oracle for a chosen plaintext attack, and the oracle for $1 \leftarrow r = \mathcal{P}^{-1}$ corresponds to the oracle for a chosen ciphertext attack.

We now give a formal security definition of a block cipher when an adversary has access to internal rounds. Note that the adversary is allowed oracle access to some subset K of the set $\{i \rightarrow j, i \leftarrow j : 1 \leq i \leq j \leq r\}$, and the insecurity function additionally depends on K.

Definition 6 (Round Security [RR00]). *Let $\mathcal{P}$ be a block cipher with domain and range I_ℓ, key length s, and some r-round decomposition $\mathcal{F}^1, \ldots, \mathcal{F}^r$. Fix some subset $K = \{\pi^1, \ldots, \pi^k\}$ of the set $\{i \to j, i \leftarrow j : 1 \leq i \leq j \leq r\}$, and let $\mathcal{A}$ be a $k+2$-oracle adversary. Then we define $\mathcal{A}$'s advantage as*

$$\mathsf{Adv}^{\mathsf{sprp}}_{\mathcal{P},\mathcal{F}^1,\ldots,\mathcal{F}^r,K}(\mathcal{A}) =$$
$$\left|\Pr[a \overset{R}{\leftarrow} I_s : \mathcal{A}^{p_a, p_a^{-1}, \pi_a^1, \ldots, \pi_a^k} = 1] - \Pr[p \overset{R}{\leftarrow} \mathsf{Perm}^\ell, a \overset{R}{\leftarrow} I_s : \mathcal{A}^{p, p^{-1}, \pi_a^1, \ldots, \pi_a^k} = 1]\right|$$

For any integers $q, t \geq 0$ and set K, $\mathsf{Adv}^{\mathsf{sprp}}_{\mathcal{P},\mathcal{F}^1,\ldots,\mathcal{F}^r,K}(q,t)$ specifies our insecurity function (analogous to Definition 2).

Ramzan and Reyzin [RR00] were the first to consider what happens when internal round functions of a Feistel network are available to an external adversary.

Theorem 2 (Ramzan-Reyzin). *Let f_1, f_2, f_3, f_4 be independently-keyed functions from a keyed function family $\mathcal{F}$ with domain and range I_n and key space I_s. Let $\mathcal{P}$ be the family of permutations on I_{2n} with key space I_{4s} defined by $\mathcal{P} = \Phi(f_1, f_2, f_3, f_4)$ with the natural 4-round decomposition $\overline{\mathcal{F}}, \overline{\mathcal{F}}, \overline{\mathcal{F}}, \overline{\mathcal{F}}$. Let $K = \{i \to j, i \leftarrow j : 2 \leq i \leq j \leq 3\}$. Then*

$$\mathsf{Adv}^{\mathsf{sprp}}_{\mathcal{P},\overline{\mathcal{F}},\overline{\mathcal{F}},\overline{\mathcal{F}},\overline{\mathcal{F}},K}(q,t) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathcal{F}}(q,t) + \binom{q}{2}\left(2^{-n+1} + 2^{-2n+1}\right) + q^2\left(2^{-n-1}\right).$$

Ramzan-Reyzin consider the case where all parties have black-box access to the internal permutations $\overline{f_2}, \overline{f_3}$. They noted that if the underlying round functions f_1, and f_2 are chosen from $\mathsf{Rand}^{n \to n}$, then one could translate their results to the random oracle model wherein f_2, f_3 are modeled as random function oracles that are accessible to all parties, including the adversary.

3 Our Main Result

We now prove our main result. We use the Ramzan-Reyzin round-security framework [RR00] to analyze our construction and leverage their techniques to obtain the desired result. However, for technical reasons, the proof must also incorporate an additional hybrid distribution into the argument. Now, let $\Psi^{f,g}_{k_1,k_2}$ denote the Even-Mansour construction when the internal permutation is replaced by a four-round Feistel network with outer round function g and inner round function f: $\Psi^{f,g}_{k_1,k_2}(x) = k_2 \oplus \Phi(g, f, f, g)(x \oplus k_1)$. Here $k_1, k_2 \in I_{2n}$ are the keys and f, g are modeled as random function oracles; i.e., they are publicly accessible to all parties (including the adversary) and behave like random functions. Observe then that the adversary can compute not only the Even-Mansour permutation, but also knows its internal structure and has black-box access to the functions f and g around which it is designed. We view this construction as consisting of the composition of six round permutations:

- $\pi_1^{k_1}(x) = x \oplus k_1$

- $\pi_3, \pi_4 = \overline{f}$. Recall that $\overline{f}$ denotes a permutation on I_{2n} defined as $\overline{f}(x) = (x^R, x^L \oplus f(x^R))$.
- $\pi_2, \pi_5 = \overline{g}$.
- $\pi_6^{k_2}(x) = x \oplus k_2$.

Observe that $\Psi^{f,g}_{k_1,k_2}(M) = \pi_6^{k_2} \circ \pi_5 \circ \cdots \circ \pi_2 \circ \pi_1^{k_1}$. We now state our main result in the following theorem:

Theorem 3 (Main Result). *Suppose $K \subseteq \{i \rightarrow j, i \leftarrow j \mid 2 \le i \le j \le 5\}$. Let f be modeled as a random oracle, and let k_1 and k_2 be picked randomly and independently from I_{2n}. Let $\Psi^{f,g}_{k_1,k_2}(x) = k_2 \oplus \Phi(g, f, f, g)(x \oplus k_1)$, and R be a random permutation on I_{2n}. Then*

$$\mathsf{Adv}^{\mathsf{sprp}}_{\mathcal{P},\pi_1,\pi_2,\ldots,\pi_6}(q, t, K) \le \left(2q^2 - q\right) \cdot 2^{-n} + \binom{q}{2} \cdot 2^{-2n+1}.$$

Observe that we do not consider any terms of the form $\mathsf{Adv}^{\mathsf{prf}}_{\mathcal{F}}(q, t)$ since we assume that the underlying round functions are modeled as random oracles in which case such terms will evaluate to 0.

Recasting the problem in the round-security framework allows us to apply the techniques of Ramzan and Reyzin [RR00] (who generalized the techniques of Naor and Reingold [NR99] to deal with the extra queries from an oracle with internal access). We note that access to the oracles of K is equivalent to access to the oracles for f and g.[3] Now, consider the following theorem.

Theorem 4. *Let f and g be modeled as random oracles, and let k_1 and k_2 be picked randomly and independently from I_{2n}. Let $\Psi^{f,g}_{k_1,k_2}(x) = k_2 \oplus \Phi(g, f, f, g)(x \oplus k_1)$, and let R be a random element of Perm^{2n}. Then, for any 4-oracle adversary $\mathcal{A}$ (we do not restrict the running time of $\mathcal{A}$) that makes at most q_c queries to its first two oracles (either Ψ, Ψ^{-1} or R, R^{-1}) and at most q_{of} and q_{og} queries to its second two oracles (f and g) respectively, it follows that:*

$$\left|\Pr[\mathcal{A}^{\Psi,\Psi^{-1},f,g} = 1] - \Pr[A^{R,R^{-1},f,g} = 1]\right|$$
$$\le (q_c^2 + 2q_{of}q_c + 2q_{og}q_c + q_c^2 - q_c)2^{-n} + \binom{q_c}{2}\left(2 \cdot 2^{-n} + 2^{-2n+1}\right).$$

Observing that the total number of queries $q = q_c + q_{of} + q_{og}$, it is straightforward to see that

$$(q_c^2 + 2q_{of}q_c + 2q_{og}q_c + q_c^2 - q_c) \le 2q^2 - q.$$

Therefore, we see that theorem 4 implies theorem 3. In the sequel, we describe the proof of theorem 4. The first part of the proof focuses on the adversary's

[3] We remark, however, that one query to an oracle in K may need to be simulated by multiple queries to f, g. Therefore, the total number of queries made to f and g is an upper bound on the number of queries that would need to be made to an oracle in K.

transcript (i.e., his “view”) and shows that each possible transcript is about as likely to occur when $\mathcal{A}$ is given Ψ, f, g as when $\mathcal{A}$ is given R, f, g. This part of the proof also relies on a hybrid distribution $\tilde{\Psi}$ to facilitate the proof. The second part uses a standard probability argument to show that if the distributions on transcripts are similar, then $\mathcal{A}$ will have a small advantage in distinguishing Ψ from R.

Proof of Theorem 4. To start with, let P denote the permutation oracle (either Ψ or R) that $\mathcal{A}$ accesses. From now on, for notational convenience we ignore the superscripts f, g and the subscripts k_1, k_2 associated with Ψ. Let $\mathcal{O}^f$ and $\mathcal{O}^g$ denote the oracles that compute the functions f and g (note that when $\mathcal{A}$ gets Ψ as its permutation oracle, f and g are actually used as the round function in the computation of the oracle $P = \Psi$; when $\mathcal{A}$ gets R as its permutation oracle, f and g are independent of $P = R$). The machine $\mathcal{A}$ makes two types of queries to the oracle P: $(+, x)$ which asks to obtain the value of $P(x)$, or $(-, y)$ which asks to obtain the value of $P^{-1}(y)$ – where both x and y are in I_{2n}. We call these *cipher queries.* We define the query-answer pair for the i^{th} cipher query as $\langle x_i, y_i\rangle \in I_{2n} \times I_{2n}$ if $\mathcal{A}$'s query was $(+, x_i)$ and y_i is the answer it received from P or its query was $(-, y_i)$ and x_i is the answer it received. We assume that $\mathcal{A}$ makes exactly q_c cipher queries and we call the sequence $\{\langle x_1, y_1\rangle, \ldots, \langle x_{q_c}, y_{q_c}\rangle\}_P$ the cipher-transcript of $\mathcal{A}$. In addition, $\mathcal{A}$ can make queries to $\mathcal{O}^f$ and $\mathcal{O}^g$. We call these *oracle queries.* We denote these queries as: $(\mathcal{O}^f, x')$ (*resp.* $(\mathcal{O}^g, x')$) which asks to obtain $f(x')$ (*resp.* $g(x')$). We define the query-answer pair for the i^{th} oracle query as $\langle x'_i, y'_i\rangle \in I_n \times I_n$ if $\mathcal{A}$'s query was $(\mathcal{O}^f, x')$ and the answer it received was y' and as $\langle x''_i, y''_i\rangle \in I_n \times I_n$ if $\mathcal{A}$'s query was $(\mathcal{O}^g, x'')$ and the answer it received was y''. We assume that $\mathcal{A}$ makes q_{of} and q_{og} queries to $\mathcal{O}^f$ and $\mathcal{O}^g$ respectively. We call the sequence $\{\langle x'_1, y'_1\rangle, \ldots, \langle x'_{q_{of}}, y'_{q_{of}}\rangle\}_{\mathcal{O}^f}$ the f-oracle-transcript of $\mathcal{A}$ and $\{\langle x''_1, y''_1\rangle, \ldots, \langle x''_{q_{og}}, y''_{q_{og}}\rangle\}_{\mathcal{O}^g}$ the g-oracle-transcript of $\mathcal{A}$. Note that since $\mathcal{A}$ is computationally unbounded, we can make the standard assumption that $\mathcal{A}$ is a deterministic machine. Under this assumption, the exact next query made by $\mathcal{A}$ can be determined by the previous queries and the answers received. We formalize this as follows:

Definition 7. *We denote the $i + j + k + 1^{st}$ query $\mathcal{A}$ makes as a function of the first $i + j + k$ query-answer pairs in $\mathcal{A}$'s cipher and oracle transcripts (where either $i < q_c$ or $j < q_{of}$ or $k < q_{og}$) by:*

$$C_{\mathcal{A}}[\{\langle x_1, y_1\rangle, \ldots, \langle x_i, y_i\rangle\}_P, \{\langle x'_1, y'_1\rangle, \ldots, \langle x'_j, y'_j\rangle\}_{\mathcal{O}^f}, \{\langle x''_1, y''_1\rangle, \ldots, \langle x''_k, y''_k\rangle\}_{\mathcal{O}^g}].$$

For the case that all queries have been made (i.e., $i = q_c, j = q_{of}, k = q_{og}$), we define the above expression to denote $\mathcal{A}$'s output as a function of its cipher and oracle transcripts.

Definition 8. *Let $\sigma = (T_P, T_f, T_g)$ be a three tuple comprising the sequences $T_P = \{\langle x_1, y_1\rangle, \ldots, \langle x_{q_c}, y_{q_c}\rangle\}_P$, $T_f = \{\langle x'_1, y'_1\rangle, \ldots, \langle x'_{q_{of}}, y'_{q_{of}}\rangle\}_{\mathcal{O}^f}$, and $T_g = \{\langle x''_1, y''_1\rangle, \ldots, \langle x''_{q_{og}}, y''_{q_{og}}\rangle\}_{\mathcal{O}^g}$, and where for $1 \le i \le q_c$ we have that $\langle x_i, y_i\rangle \in I_{2n} \times I_{2n}$, for $1 \le j \le q_{of}$, we have that $\langle x'_j, y'_j\rangle \in I_n \times I_n$, and for $1 \le k \le q_{og}$,* *we have that* $\langle x''_k, y''_k \rangle \in I_n \times I_n$. *Then,* σ *is a possible* $\mathcal{A}$*-transcript if for every* $1 \le i \le q_c$*, for every* $1 \le j \le q_{of}$ *and for every* $1 \le k \le q_{og}$*,*

$$C_{\mathcal{A}}[\{\langle x_1, y_1 \rangle, \ldots, \langle x_i, y_i \rangle\}_P, \{\langle x'_1, y'_1 \rangle, \ldots, \langle x'_j, y'_j \rangle\}_{\mathcal{O}^f} \{\langle x''_1, y''_1 \rangle, \ldots, \langle x''_k, y''_k \rangle\}_{\mathcal{O}^g}] \in \{(+, x_{i+1}), (-, y_{i+1}), (\mathcal{O}^f, x'_{j+1}), (\mathcal{O}^g, x''_{k+1})\}.$$

We now consider two useful processes for answering $\mathcal{A}$'s *cipher* queries.

Definition 9. *Let* $\tilde{\Psi}$ *denote the process in which the cipher queries and* f*-oracle queries are answered as they would be for* Ψ*, however the* g*-oracle queries are answered by another independent random function oracle* h*.*

Definition 10. *Let* $\tilde{R}$ *denote the process that answers all oracle queries as* Ψ *would, but answers the* i^{th} *cipher query of* $\mathcal{A}$ *as follows:*

1. *If* $\mathcal{A}$*'s query is* $(+, x_i)$ *and for some* $1 \le j < i$ *the* j^{th} *query-answer pair is* $\langle x_i, y_i \rangle$*, then* $\tilde{R}$ *answers* y_i*.*
2. *If* $\mathcal{A}$*'s query is* $(-, y_i)$ *and for some* $1 \le j < i$ *the* j^{th} *query-answer pair is* $\langle x_i, y_i \rangle$*, then* $\tilde{R}$ *answers* x_i*.*
3. *If neither of the above happens, then* $\tilde{R}$ *answers with a uniformly chosen element in* I_{2n}*.*

We formalize the fact that $\tilde{R}$'s answers may not be consistent with any function, let alone any permutation.

Definition 11. *Let* $\sigma' = \{\langle x_1, y_1 \rangle, \ldots, \langle x_{q_c}, y_{q_c} \rangle\}_P$ *be any possible* $\mathcal{A}$*-cipher transcript. We say that* σ' *is inconsistent if for some* $1 \le j < i \le q_c$ *the corresponding query-answer pairs satisfy* $x_i = x_j$ *but* $y_i \ne y_j$*, or* $x_i \ne x_j$ *but* $y_i = y_j$*. Likewise, we call any* $\mathcal{A}$*-transcript* σ *that contains* σ' *inconsistent.*

Note 1. If $\sigma = (T_P, T_f, T_g)$, with sub-transcripts $T_P = \{\langle x_1, y_1 \rangle, \ldots, \langle x_{q_c}, y_{q_c} \rangle\}_P$, $T_f = \{\langle x'_1, y'_1 \rangle, \ldots, \langle x'_{q_{of}}, y'_{q_{of}} \rangle\}_{\mathcal{O}^f}$, and $T_g = \{\langle x''_1, y''_1 \rangle, \ldots, \langle x''_{q_{og}}, y''_{q_{og}} \rangle\}_{\mathcal{O}^g}$, is a possible $\mathcal{A}$-transcript, we assume from now on that if σ is consistent and if $i \ne j$ then $x_i \ne x_j$, $y_i \ne y_j$, $x'_i \ne x'_j$, and $x''_i \ne x''_j$. This formalizes the concept that $\mathcal{A}$ never repeats a query if it can determine the answer from a previous query-answer pair.

Fortunately, the process $\tilde{R}$ often behaves like a permutation. It turns out that if $\mathcal{A}$ is given oracle access to either $\tilde{R}$ or R to answer its cipher queries, it will have a negligible advantage in distinguishing between the two. Proposition 1 states this more formally. Before doing so, we first consider the distributions on the various transcripts seen by $\mathcal{A}$ as a function of the different distributions on answers it can get.

Definition 12. *The discrete random variables* $T_\Psi, T_{\tilde{\Psi}}, T_R, T_{\tilde{R}}$ *denote the cipher and oracle transcripts seen by* $\mathcal{A}$ *when its cipher queries are answered by* Ψ*,* $\tilde{\Psi}$*,* R*,* $\tilde{R}$ *respectively, and its oracle queries are answered by* $\mathcal{O}^f$ *or* $\mathcal{O}^g$*.*

Remark 1. Observe that according to our definitions and assumptions, $\mathcal{A}^{\Psi,\Psi^{-1},f,g}$ and $C_{\mathcal{A}}(T_{\Psi})$ denote the same random variable. The same is true for the other discrete random variables.

Proposition 1. $|\Pr_{\tilde{R}}[C_{\mathcal{A}}(T_{\tilde{R}})=1]-\Pr_{R}[C_{\mathcal{A}}(T_{R})=1]| \leq \binom{q_c}{2}\cdot 2^{-2n}$.

The proof of this proposition has appeared in numerous places [NR99, RR00]. The idea is to observer that T_R, $T_{\tilde{R}}$ have the same distribution conditioned on $T_{\tilde{R}}$ being consistent. One can then bound the probability that $T_{\tilde{R}}$ is inconsistent by $\binom{q_c}{2}\cdot 2^{-2n}$. The proof can be completed by a standard probability argument. We omit the details, though they are available in the full version [GR04]. We now proceed to obtain a bound on the advantage that $\mathcal{A}$ will have in distinguishing between T_{Ψ} and $T_{\tilde{R}}$. We first show that T_{Ψ} and $T_{\tilde{\Psi}}$ are identically distributed, unless the input to g in a cipher query related to Ψ matches the input to g in an oracle query related to Ψ. We can compute the likelihood of such an event as a function of only k_1 and k_2 – we term this event BadG and define it next; we then compute the probability that it occurs.

Definition 13. *For every specific pair of keys $k_1, k_2 \in I_{2n}$, we define $\mathsf{BadG}(k_1,k_2)$ to be the set of all possible and consistent transcripts $\sigma=(T_P,T_f,T_g)$, with sub-transcripts $T_P=\{\langle x_1,y_1\rangle,\ldots,\langle x_{q_c},y_{q_c}\rangle\}_P$, $T_f=\{\langle x'_1,y'_1\rangle,\ldots,\langle x'_{q_{of}},y'_{q_{of}}\rangle\}_{\mathcal{O}^f}$, and $T_g=\{\langle x''_1,y''_1\rangle,\ldots,\langle x''_{q_{og}},y''_{q_{og}}\rangle\}_{\mathcal{O}^g}$ satisfying at least one of the following events:*

- **BG1:** *there exists $1\leq i\leq q_c$, $1\leq j\leq q_{og}$ such that $x_i^R \oplus k_1^R = x''_j$, or*
- **BG2:** *there exists $1\leq i\leq q_c$, $1\leq j\leq q_{og}$ such that $y_i^L \oplus k_2^L = x''_j$.*

Proposition 2. *Let k_1,k_2 be randomly and independently chosen from I_{2n}. For any possible and consistent $\mathcal{A}-$transcript $\sigma=(T_P,T_f,T_g)$, with sub-transcripts $T_P=\{\langle x_1,y_1\rangle,\ldots,\langle x_{q_c},y_{q_c}\rangle\}_P$, $T_f=\{\langle x'_1,y'_1\rangle,\ldots,\langle x'_{q_{of}},y'_{q_{of}}\rangle\}_{\mathcal{O}^f}$, and $T_g=\{\langle x''_1,y''_1\rangle,\ldots,\langle x''_{q_{og}},y''_{q_{og}}\rangle\}_{\mathcal{O}^g}$, we have that*

$$\Pr_{k_1,k_2}[\sigma\in\mathsf{BadG}(k_1,k_2)] \leq 2q_{og}q_c\cdot 2^{-n}.$$

Proof. (Sketch) A transcript σ is in $\mathsf{BadG}(k_1,k_2)$ if one of **BG1** or **BG2** occur. We obtain an upper bound on the probabilities of each of these events separately by using the fact that k_1,k_2 are chosen uniformly at random from I_{2n}. Applying the union bound to sum the individual probabilities yields the desired result.

We now show that T_{Ψ} and $T_{\tilde{\Psi}}$ are identically distributed if neither **BG1** nor **BG2** occur.

Lemma 1. *Let $\sigma=(T_P,T_f,T_g)$, where $T_P=\{\langle x_1,y_1\rangle,\ldots,\langle x_{q_c},y_{q_c}\rangle\}_P$, $T_f=\{\langle x'_1,y'_1\rangle,\ldots,\langle x'_{q_{of}},y'_{q_{of}}\rangle\}_{\mathcal{O}^f}$, and $T_g=\{\langle x''_1,y''_1\rangle,\ldots,\langle x''_{q_{og}},y''_{q_{og}}\rangle\}_{\mathcal{O}^g}$, be any possible and consistent $\mathcal{A}-$transcript, then*

$$\Pr_{\Psi}[T_{\Psi}=\sigma|\sigma\notin\mathsf{BadG}(k_1,k_2)]=\Pr_{\tilde{\Psi}}[T_{\tilde{\Psi}}=\sigma].$$

Proof. (Sketch) Observe that $x_i^R \oplus k_1^R \neq x_j''$ and $y_i^L \oplus k_2^L \neq x_j''$ for all i, j whenever $\sigma \notin \mathsf{BadG}(k_1, k_2)$. In such a case, the inputs to g during the cipher queries are distinct from the inputs to g during the g-oracle queries. Since there is no overlap in the two sets of queries, since g is modeled as a random oracle, and since the events depend only on the choice of k_1 and k_2 (which are chosen independently of g), the distribution is identical to one in which g is replaced by another independently chosen random oracle h.

We now focus on $T_{\tilde{\Psi}}$. It turns out that $T_{\tilde{\Psi}}$ and $T_{\tilde{R}}$ are identically distributed unless the same value is input to the inner random oracle f on different occasions (we show this in Lemma 2). We can compute the likelihood of this event as a function of *only* k_1, k_2, and g. We call this event "Bad" (in the next definition) and obtain a bound on the probability that it actually occurs (in Proposition 3).

Definition 14. *For every specific pair of keys $k_1, k_2 \in I_{2n}$ and oracle $g \in \mathsf{Rand}^{n \to n}$, define $\mathsf{Bad}(k_1, k_2, g)$ to be the set of all possible and consistent transcripts $\sigma = (T_P, T_f, T_g)$, with sub-transcripts $T_P = \{\langle x_1, y_1\rangle, \ldots, \langle x_{q_c}, y_{q_c}\rangle\}_P$, $T_f = \{\langle x_1', y_1'\rangle, \ldots, \langle x_{q_{of}}', y_{q_{of}}'\rangle\}_{\mathcal{O}^f}$, and $T_g = \{\langle x_1'', y_1''\rangle, \ldots, \langle x_{q_{og}}'', y_{q_{og}}''\rangle\}_{\mathcal{O}^g}$, satisfying at least one of the following events:*

- **B1:** $\exists\ 1 \leq i < j \leq q_c$ *such that* $g(x_i^R \oplus k_1^R) \oplus x_i^L = g(x_j^R \oplus k_1^R) \oplus x_j^L$
- **B2:** $\exists\ 1 \leq i < j \leq q_c$ *such that* $y_i^R \oplus g(y_i^L \oplus k_2^L) = y_j^R \oplus g(y_j^L \oplus k_2^L)$
- **B3:** $\exists\ 1 \leq i, j \leq q_c$ *such that* $g(x_i^R \oplus k_1^R) \oplus x_i^L \oplus k_1^L = k_2^R \oplus y_j^R \oplus g(y_j^L \oplus k_2^L)$
- **B4:** $\exists\ 1 \leq i \leq q_c$, $1 \leq j \leq q_{of}$ *such that* $g(x_i^R \oplus k_1^R) \oplus x_i^L \oplus k_1^L = x_j'$
- **B5:** $\exists\ 1 \leq i \leq q_c$, $1 \leq j \leq q_{of}$ *such that* $k_2^R \oplus y_i^R \oplus g(y_i^L \oplus k_2^L) = x_j'$.

Proposition 3. *Let k_1, k_2 be randomly and independently chosen from I_{2n}. Then, for any possible and consistent $\mathcal{A}-$transcript $\sigma = (T_P, T_f, T_g)$, with sub-transcripts $T_P = \{\langle x_1, y_1\rangle, \ldots, \langle x_{q_c}, y_{q_c}\rangle\}_P$, $T_f = \{\langle x_1', y_1'\rangle, \ldots, \langle x_{q_{of}}', y_{q_{of}}'\rangle\}_{\mathcal{O}^f}$, and $T_g = \{\langle x_1'', y_1''\rangle, \ldots, \langle x_{q_{og}}'', y_{q_{og}}''\rangle\}_{\mathcal{O}^g}$, we have that*

$$\Pr_{k_1, k_2, g}[\sigma \in \mathsf{Bad}(k_1, k_2, g)] \leq \left(q_c^2 + 2q_{of}q_c + 2 \cdot \binom{q_c}{2}\right) \cdot 2^{-n}.$$

Proof. (Sketch) Recall that a transcript $\sigma \in \mathsf{Bad}(k_1, k_2, g)$ if at least one of the above events occurs. We obtain an upper bound on the probabilities of each of these events separately using the fact that k_1, k_2 are chosen uniformly at random from I_{2n} and that g is chosen uniformly at random from $\mathsf{Rand}^{n \to n}$. Applying the union bound to sum the probabilities for each event yields the desired result.

Lemma 2. *Let σ be defined as in Lemma 1. Then,*

$$\Pr_{\tilde{\Psi}}[T_{\tilde{\Psi}} = \sigma | \sigma \notin \mathsf{Bad}(k_1, k_2, g)] = \Pr_{\tilde{R}}[T_{\tilde{R}} = \sigma].$$

Proof. It is easy to see that $\Pr_{\tilde{R}}[T_{\tilde{R}} = \sigma] = 2^{-(2q_c + q_{of} + q_{og})n}$ (following an argument in [NR99], [RR00]). Now, fix k_1, k_2, g to be such that $\sigma \notin \mathsf{Bad}(k_1, k_2, g)$. We will now compute $\Pr_{f,h}[T_{\tilde{\Psi}} = \sigma]$ (recall that in the definition of $\tilde{\Psi}$, h is a random oracle independent of f and g, and note that the probability is now only over the choice of f and h). Since σ is a possible $\mathcal{A}$-transcript, it follows that $T_{\tilde{\Psi}} = \sigma$ if and only if $y_i = k_1 \oplus \tilde{\Psi}(g, f, f, g)(x_i \oplus k_2)$ for all $1 \leq i \leq q_c$, $y'_j = f(x'_j)$, for all $1 \leq j \leq q_{of}$, and $y''_j = g(x''_j)$ for all $1 \leq j \leq q_{og}$. If we define $S_i = k_1^L \oplus x_i^L \oplus g(x_i^R \oplus k_1^R)$ and $T_i = k_2^R \oplus y_i^R \oplus g(y_i^L \oplus k_2^L)$, then

$$(y_i^L, y_i^R) = \tilde{\Psi}(x_i^L, x_i^R) \Leftrightarrow f(S_i) \oplus k_1^R = T_i \oplus x_i^R \text{ and } f(T_i) \oplus k_2^L = y_i^L \oplus S_i.$$

Now observe that for all $1 \leq i < j \leq q_c$, $S_i \neq S_j$ and $T_i \neq T_j$ (otherwise $\sigma \in \mathsf{Bad}(k_1, k_2, g)$). Similarly, for all $1 < i, j < q_c$, $S_i \neq T_j$. In addition, it follows again from the fact that $\sigma \notin \mathsf{Bad}(k_1, k_2, g)$ that for all $1 \leq i \leq q_c$ and $1 \leq j \leq q_{og}$, $x'_i \neq S_j$ and $x'_i \neq T_j$. So, if $\sigma \notin \mathsf{Bad}(k_1, k_2, g)$ all the inputs to f are distinct. Since f is modeled as a random oracle, $\Pr_{f,h}[T_{\tilde{\Psi}} = \sigma] = 2^{-(2q_c+q_{of}+q_{og})n}$ (the cipher transcript contributes 2^{-2nq_c} and the oracle transcripts contribute $2^{-q_{of}n-q_{og}n}$ to the probability). Thus, for every choice of k_1, k_2, g such that $\sigma \notin \mathsf{Bad}(k_1, k_2, g)$, the probability that $T_{\tilde{\Psi}} = \sigma$ is exactly the same: $2^{-(2q_c+q_{of}+q_{og})n}$. Therefore: $\Pr_{\tilde{\Psi}}[T_{\tilde{\Psi}} = \sigma | \sigma \notin \mathsf{Bad}(k_1, k_2, g)] = 2^{-(2q_c+q_{of}+q_{og})n}$.

The rest of the proof consists of using the above lemma and Propositions 1, 2 and 3, as well as Lemmas 1 and 2, in a probability argument. The idea is to first express the adversary's advantage as a function of how its distinguishing machine behaves on specific transcripts. Then, these probabilities are re-expressed to incorporate the conditions Bad and BadG. By basic manipulation of probabilities, we can show that the adversary's advantage is bounded above by the probability of the conditions Bad or BadG occurring, plus the probability that the transcript is inconsistent. An additional term of the form $\binom{q_c}{2} \cdot 2^{-2n}$ also appears because of an application of the triangle inequality. The complete details are omitted due to space constraints, though are available in the full version [GR04].

4 Conclusions

We resolved a fundamental open problem of the Even-Mansour work by demonstrating that the underlying random permutation oracle could be instantiated with a construction based on random function oracles. There are many avenues for future work. For example, we may be able to apply our techniques to other situations where random permutation oracles are useful. Also, there is a sizeable gap between the best known key-recovery attack and the bound achieved in our security proof. Perhaps that gap can be decreased by developing a variant on the slide-with-twist that exploits the structure of our construction.

References

[BBP04] M. Bellare, A. Boldyreva, and A. Palacio. An un-instantiable random-oracle-model scheme for a hybrid-encryption problem. In C. Cachin and J. Camenisch, editors, *Proc. EUROCRYPT 2004*, Lecture Notes in Computer Science. Springer-Verlag, 2004.

[BCK96] Mihir Bellare, Ran Canetti, and Hugo Krawczyk. Pseudorandom functions revisited: The cascade construction and its concrete security. In *37th Annual Symposium on Foundations of Computer Science*, pages 514–523. IEEE, 1996.

[BKR94] Mihir Bellare, Joe Kilian, and Phillip Rogaway. The security of cipher block chaining. In Yvo G. Desmedt, editor, *Advances in Cryptology—CRYPTO '94*, volume 839 of *Lecture Notes in Computer Science*, pages 341–358. Springer-Verlag, 21–25 August 1994.

[BR93] M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *First ACM Conference on Computer and Communications Security*, pages 62–73, Fairfax, 1993.

[BW00] A. Biryukov and D. Wagner. Advanced slide attacks. In Advances in Cryptology – Proc. of Eurocrypt 2000. Springer-Verlag.

[CGH98] R. Canetti, O. Goldreich, and S. Halevi. The random oracle methodology, revisited. In *Proc.* 30*th ACM Symp. on Theory of Computing*, 1998.

[Dae91] J. Daemen. Limitations of the Even-Mansour construction. In *Advances in Cryptology – ASIACRYPT '91, vol. 739, 495–498, 1992.* Springer-Verlag. Initially Presented at the Rump Session.

[EM97] S. Even and Y. Mansour. A construction of a cipher from a single pseudorandom permutation. *Journal of Cryptology*, 10(3):151–162, Summer 1997. Earlier version in *Proc. ASIACRYPT 1991.* Lecture Notes in Computer Science, vol. 739, 210–224, Springer Verlag (1992).

[GGM84] O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, October 1984.

[GK03] S. Goldwasser and Y. Tauman Kalai. On the (in)security of the Fiat-Shamir Paradigm. In *Proceedings of FOCS 2003*, 2003.

[GR04] C. Gentry and Z. Ramzan. Eliminating random permutation oracles in the Even-Mansour cipher. *Cryptology ePrint archive*, 2004.

[KR96] J. Kilian and P. Rogaway. How to protect against exhaustive search. In *Proc. CRYPTO 96*, Lecture Notes in Computer Science. Springer-Verlag, 1996.

[LR88] M. Luby and C. Rackoff. How to construct pseudorandom permutations and pseudorandom functions. *SIAM J. Computing*, 17(2):373–386, April 1988.

[FIPS46] National Bureau of Standards. FIPS publication 46: Data encryption standard, 1977. Federal Information Processing Standards Publication 46.

[NR99] M. Naor and O. Reingold. On the construction of pseudo-random permutations: Luby-Rackoff revisited. *J. of Cryptology*, 12:29–66, 1999. Preliminary version in: *Proc. STOC 97.*

[PRS02] S. Patel, Z. Ramzan, and G. Sundaram. Luby-Rackoff ciphers: Why XOR is not exclusive. In Proc. of Selected Areas of Cryptography, Lecture Notes in Computer Science, Vol. 2595, pages 271–290.

[RR00] Z. Ramzan and L. Reyzin. On the Round Security of Symmetric-Key Cryptographic Primitives. In *Proc. CRYPTO 00*, Lecture Notes in Computer Science. Springer-Verlag, 2000.