

Luby-Rackoff Backwards: Increasing Security by Making Block Ciphers Non-invertible

Mihir Bellare¹, Ted Krovetz², and Phillip Rogaway²

¹ Dept. of Computer Science & Engineering, University of California at San Diego, 9500 Gilman Drive, La Jolla, California 92093, USA. E-Mail: mihir@cs.ucsd.edu
URL: <http://www-cse.ucsd.edu/users/mihir>

² Dept. of Computer Science, University of California at Davis, Engineering II Bldg., One Shields Ave., Davis, CA 95616, USA
E-mail: {krovetz,rogaway}@cs.ucdavis.edu
URL: <http://www.cs.ucdavis.edu/~krovetz,~rogaway>

Abstract. We argue that the invertibility of a block cipher can reduce the security of schemes that use it, and a better starting point for scheme design is the non-invertible analog of a block cipher, that is, a pseudorandom function (PRF). Since a block cipher may be viewed as a pseudorandom permutation, we are led to investigate the *reverse* of the problem studied by Luby and Rackoff, and ask: “how can one transform a PRP into a PRF in as security-preserving a way as possible?” The solution we propose is *data-dependent re-keying*. As an illustrative special case, let $E : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be the block cipher. Then we can construct the PRF F from the PRP E by setting $F(k, x) = E(E(k, x), x)$. We generalize this to allow for arbitrary block and key lengths, and to improve efficiency. We prove strong quantitative bounds on the value of data-dependent re-keying in the Shannon model of an ideal cipher, and take some initial steps towards an analysis in the standard model.

1 Introduction

This paper describes a transformation — turning a “pseudorandom permutation” (PRP) into a “pseudorandom function” (PRF) using “data-dependent re-keying.” It can be applied to a block cipher to increase the block cipher’s security in certain ways, and, in particular, the method leads to block cipher based message encryption and authentication techniques which are approximately as efficient as ones in current use, but have better security.

In Section 2 we explain our (at first paradoxical sounding) thesis: that invertibility of a block cipher can be a liability, not an asset, when it comes to the security of schemes that use the cipher. We will then explain what are PRFs and PRPs, how the former are a better starting point for constructions but the latter a better model for block ciphers, and how all this leads us to consider the problem of transforming PRPs into PRFs in a security-preserving way.

In Section 3 we describe our way to do the PRP to PRF transformation. We call our transform F_n^d , where d is a parameter on which the construction

depends. (The impatient reader can jump to Section 3 to see how F_n^d works. It is very simple.)

Our main result is an analysis¹ in the Shannon model which shows that if the block cipher is ideal then its transform under F_n^d is close to an ideal random function. The provided bounds are strong, showing the transform is close to security preserving.

The interpretation of the above is that the F_n^d transform gives good security against “generic” attacks. To gauge its strength against cryptanalytic attacks we also analyze it in the standard complexity theoretic or “reductionist” framework. We do succeed in providing a reduction, but the quality of the bounds is not as good as in the Shannon model, and thus we view these results as preliminary, hopefully to be improved.

The results are presented, discussed, and displayed graphically in Section 5. Just before that, in Section 4, we provide the precise definitions of the security notions, but these can be skipped at first reading, or skipped entirely by an expert. In this truncated version of the paper we do not have space to include proofs of theorems. The reader interested in these is referred to the full version of this paper [7] which is available on the web.

2 The Problem

We begin with a simple example, then relate these issues to PRFs and PRPs, then describe the problem that results, and conclude with a discussion of related work.

2.1 Invertibility can hurt when using block ciphers: An example

A block cipher is a function $E: \{0, 1\}^\kappa \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ which transforms an n -bit message block x into an n -bit string y under the control of a κ -bit key k : $y = E(k, x)$. The function is invertible in the sense that for each key the map $E_k \stackrel{\text{def}}{=} E(k, \cdot)$ is a permutation of $\{0, 1\}^n$, and knowledge of k permits computation of E_k^{-1} . Concrete examples are DES, triple-DES and RC5.

Message encryption is done by using the block cipher in some mode of operation, such as “CBC.” Using even a very “good” block cipher (say triple-DES, or even an ideal cipher), CBC encryption becomes insecure once $2^{n/2}$ blocks have been encrypted, in the sense that at this point partial information about the message begins to leak², due to birthday attacks.³ Furthermore, this is true for many other common modes of operation, too. Thus direct use of a 64-bit block

¹ All analyses in this paper are concrete and quantitative, meaning providing explicit, non-asymptotic bounds on the success probability of an adversary as a function of its resources.

² A good encryption scheme is much more than one that prevents key recovery from a ciphertext: it should have the property that even partial information about the plaintext is not revealed [10, 4].

³ The attacks are well known. See [4] for an analysis of their effectiveness relative to formal notions of security for encryption.

size block cipher usually enables one to safely encrypt no more than 2^{32} blocks, which is quite small.

We stress that these attacks arise because the cipher is a permutation, and their cost depends only on the block length, not the key length or the security of the block cipher. So the attacks are just as effective for triple-DES, or even an ideal block cipher, as they are for DES. In summary, block cipher based schemes are often subject to birthday attacks arising from the very nature of block ciphers as permutations.

So how can we safely encrypt more than $2^{n/2}$ blocks? One answer is to use a slightly different type of primitive in an appropriate mode of operation: specifically, a “pseudorandom function” (PRF) in CTR (counter) mode, as discussed in [4, 12] and explained further below. This way to encrypt is easy and has no extra overhead if a PRF of cost comparable to the block cipher is available.

The above is only one example of an issue that arises in many places: that the permutivity of a block cipher can hinder the security of schemes which use it. To effectively address this we need to explain what are PRFs and PRPs and how they relate to block ciphers.

2.2 PRPs, PRFs, and their relation to block ciphers

Let us first back up and look at how the security of a block cipher is best captured.

SECURITY OF A BLOCK CIPHER: PRPs. It is natural to view a real block cipher as constructed to “approximate”, as closely as possible, an ideal block cipher (that is, a random permutation) in the sense that if you don’t know the key k and only see input/output examples of E_k then these should appear like input/output examples of a random permutation. The quality of a given block cipher E as a PRP (pseudorandom permutation) is thus captured by a function $\text{SEC}_E^{\text{PRP}}(q, t)$ which returns the maximum “advantage” that one can obtain in distinguishing E_k from a random permutation if you see q input/output examples and are allowed further computational resources bounded by t . (In the complexity-theoretic model, t will bound computing time; in the information-theoretic model, t will bound the number of known $(k, x, E_k(x))$ values. The advantage is a number between 0 and 1 given as the difference of two probabilities: the probability that the adversary outputs 1 given a random function E_k from E , and the probability that the adversary outputs 1 given a random permutation π . See Section 4 for more details.)

Each specific cipher (eg. DES) will have such an associated security function, which depends on (and to a large extent comprises) its cryptanalytic strength. Of course we won’t know for sure what is this function, but we can work with what we know from cryptanalytic results. For example, if the linear cryptanalysis of [14] is the best attack on DES, we might assume $\text{SEC}_{\text{DES}}^{\text{PRP}}(q, t)$ stays small (close to 0) until q, t reaches around 2^{43} . From now on, “block cipher” and “PRP” are synonymous, from the security point of view.

CIPHERS WITHOUT INVERTIBILITY: PRFs. Like a block cipher, a pseudorandom function (PRF) is a map $F: \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$, but now $F_k \stackrel{\text{def}}{=} F(k, \cdot)$ is

not required to be invertible. The required security property is to approximate, as closely as possible, a random function. The quality of a given function F is captured by $\text{SEC}_F^{\text{prf}}(q, t)$ which returns the maximum “advantage” that one can obtain in distinguishing F_k from a random function if you see q input-output examples and are allowed computational resources t . (This advantage is the difference between probability that the adversary outputs 1 given a random function F_k from F and the probability that the adversary outputs 1 given a random function ρ . See Section 4 for more details.)

THE EXAMPLE REVISITED. Counter mode encryption with a PRF F means that to encrypt an m -block plaintext $M = x_1 \cdots x_m$, send

$$(\text{ctr}, F_k(\langle \text{ctr} + 1 \rangle) \oplus x_1 \parallel \cdots \parallel F_k(\langle \text{ctr} + m \rangle) \oplus x_m)$$

where $\langle i \rangle$ is the binary encoding of i into n bits, “ \parallel ” denotes concatenation, and where you increment ctr by m after doing each encryption. (Notice that to decrypt you need only apply F_k , so that you don’t need this function to be invertible.) Counter-mode encryption with a good PRF is pretty much “ideal encryption”: it is shown in [4] that an adversary’s chance of obtaining partial information about some plaintext, after q blocks have been encrypted, is at most $\text{SEC}_F^{\text{prf}}(q, t)$, the strength of F as a PRF. In particular if we had a PRF F with the same numerical security as DES but as a PRF not a PRP, namely $\text{SEC}_F^{\text{prf}}(q, t) \approx \text{SEC}_{\text{DES}}^{\text{prp}}(q, t)$, then we could encrypt nearly 2^{43} blocks, well above the birthday bound.

In contrast, when we use a block cipher (PRP) directly in CBC (or CTR) mode, we are not able to recoup all of the cryptographic strength captured by its $\text{SEC}_F^{\text{prp}}(\cdot, \cdot)$ value, because at $q = 2^{n/2}$ (which is $q = 2^{32}$ for DES) birthday attacks kill the encryption scheme.

The conclusion can be put like this: to get quantitatively good security, what is most useful and convenient about F is that $\text{SEC}_F^{\text{prf}}(q, t)$ be small, not $\text{SEC}_F^{\text{prp}}(q, t)$. To make the former as low as possible the family F must *not* be a family of permutations, since no family of permutation will have a good value of $\text{SEC}_F^{\text{prf}}(q, t)$ if $q \geq 2^{n/2}$. This is because of birthday attacks: if F is a family of permutations then the adversary $A(q)$ who guesses “random function” if and only if she sees a collision in the answers returned from q distinct but otherwise arbitrary queries already accrues advantage of about $1/e$ if $q = 2^{n/2}$. The adversary’s advantage then goes quickly to 1 with $q \gg 2^{n/2}$.

2.3 Luby-Rackoff backwards

The above is part of an emerging view or understanding, emanating from works like [4–6, 21], that when it comes to designing higher-level primitives (like encryption schemes or MACs) a PRF is a better tool than a PRP, from two points of view: it permits easier and more effective *analysis* of the designed scheme, and the resulting schemes have a greater proven *quantitative security*. This leads us to suggest that for the purpose of protocol design, what we really want are PRFs, not block ciphers (PRPs).

So the question is how to get PRF families of high security and low cost. One possibility is to make these directly, in the same way we make block ciphers now.

We suggest that this indeed be kept in mind for the future, but at the moment is not a very pragmatic view, for two reasons. First, we have lots of (good) block ciphers available, and we want to use them well. Second, permutivity may be important to the design process of block ciphers; for example, using the round structure of a Feistel-network gives rise to a permutation.⁴

We propose instead to transform PRPs into PRFs. That is, starting with a good PRP E (realized by a block cipher), convert it into a good PRF F . This is effectively the reverse of the problem considered by Luby and Rackoff [13], who wanted to turn PRFs into PRPs.

A crucial issue is to make transformations that are as “security preserving” as possible. We want $\text{SEC}_F^{\text{prf}}(q, t)$ to remain low even for $q \gg 2^{n/2}$. Ideally, $\text{SEC}_F^{\text{prf}}(q, t)$ would be close to $\text{SEC}_E^{\text{prp}}(q, t)$.

Let us now discuss some related work. Following that we present our construction.

2.4 History and related work

Our construction is related to the cascade construction of [3].

The notion of a PRF was first defined in the polynomial-time framework by Goldreich, Goldwasser and Micali [9]. A concrete security treatment of PRFs, together with the idea that concretely defined PRFs/PRPs can be used to model block ciphers, originates with [6]. Luby and Rackoff use the term PRP to refer to a family of permutations that is a PRF family in the sense of [9]. Our notion is different in that we measure the advantage relative to random permutations, not functions. This makes no difference in the polynomial-time framework, but in the concrete-security framework the difference is crucial; indeed, if concrete security is ignored, the problem we are considering does not exist.

The ideal block cipher model we use for some of our results is that of [20], used also in [8, 11].

There are many natural ways to try to do the PRP-to-PRF conversion. One of the first to come to mind is to define $F_k(x) = x \oplus E_k(x)$. This construction is of value in some contexts, but not in ours. For if you are given an oracle for this $F_k(\cdot)$ you effectively have an oracle for $E_k(\cdot)$: for any query x you can compute $E_k(x)$ as $x \oplus F_k(x)$. So F_k will resemble a random function no more than E_k does.

There are many natural alternatives to the F_n^d transformation. For example, truncate $E_k(x)$, defining $F_k(x)$ to be some appropriate-length prefix of $E_k(x)$. This scheme was partially analyzed by [2]. Another natural method is $F_{k_1 k_2}(x) = E_{k_1}(x) \oplus E_{k_2}(x)$. This has not been analyzed.

Aiello and Venkatesan [1] give a general construction for turning a PRF $E : \{0, 1\}^\kappa \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ into a PRF $F : \{0, 1\}^{6\kappa} \times \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$. But this is a different problem. Although they too want to circumvent some

⁴ Another possibility is to make sure that the block size n is large enough ($n \geq 128$) that attacks of complexity $2^{n/2}$ are irrelevant. This too is a good idea, but the construction we give has merit which goes beyond the birthday attacks which we have been using to motivate this problem.

birthday attacks, their starting point is a random function (not permutation) and the problem is to double the number of bits the function can take as input. They are bound by the original security of the starting function as a PRF: birthday attacks are only prevented in the sense that the construction does not induce such attacks itself. So if a block cipher is the starting point, it is viewed as a PRF, meaning the security is only $2^{n/2}$. There is no notion of modeling a cipher as a random permutation. In contrast, we go above the original birthday threshold, to a security close to 2^n . Our construction is also more efficient, and it yields a map of the same key size and block length as the original one.

In constructing a Wegman-Carter message authentication code (MAC) [22] one needs to symmetrically encrypt the universal-hash of each message M . If a PRP is in hand for doing the encryption, one could define $\text{MAC}_{k_1, k_2}(M) = (\text{ctr}, E_{k_2}(\text{ctr}) \oplus h_{k_1}(M))$, but the security would degrade by $\Theta(q^2 2^{-n})$ compared to using a PRF. (Here q is the number of MACed messages.) Shoup [21] describes an alternative with better exact security. Our methods allow the simpler and more general $(\text{ctr}, F_{k_2}(\text{ctr}) \oplus h_{k_1}(M))$, where F is the result of PRP-to-PRF conversion starting from E .

As we explained, Luby and Rackoff consider the complementary problem of turning a PRF into a block cipher [13]. Luby and Rackoff spawned much further work, including [15–18, 23], and our work shares their emphasis on concrete bounds, efficiency, and tight reductions.

3 The Fn Construction

We have described in Section 2.4 some simple suggestions that don't work and some related constructions. Now we present our solution. We let $E: \{0, 1\}^\kappa \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be the given block cipher (PRP).

The values n and κ vary across real block ciphers; for example, for DES we have $\kappa = 56$ and $n = 64$; for (two-key) triple DES we have $\kappa = 112$ and $n = 64$. We want to handle all these cases. Accordingly, our construction depends on the relative values of κ and n . It also depends on a parameter d , where $0 \leq d < n$.

SIMPLE CASE. The simplest case of our construction is when the given PRP has the property that $\kappa = n$, and we choose $d = 0$. One then defines $F = \text{Fn}^0 E$ by $F(k, x) = E(E(k, x), x)$. That is, $F_k(x) = E_{k'}(x)$, where $k' = E_k(x)$. We call this “data-dependent re-keying” since we are applying E to x , but using the data-dependent “derived key” $k' = E_k(x)$. The cost of computing F is twice the cost of computing E , in the sense that there are two applications of E for each application of F . The general construction includes a provision aimed at reducing this overhead.

THE GENERAL CASE. Let $0 \leq d < n$ be given. If x' is an n -bit string, let $x' \gg d$ denote x' shifted to the right by d positions, with 0-bits filling the vacated positions. If k' is any string of length at least i , let $[k']_{1..i}$ be the string consisting of the first i bits of k' . Set $j = \lceil \kappa/n \rceil$. The function $F = \text{Fn}^d E$ takes a κj bit key $k_1 \cdots k_j$ and an n -bit input x to return an n -bit output y as follows:

```

function  $F(k_1 \cdots k_j, x)$ 
begin
   $x' \leftarrow x \gg d$  // Shift away low  $d$  bits
   $k' \leftarrow E(k_1, x') \parallel \cdots \parallel E(k_j, x')$  // Construct the "extended" derived key
   $k \leftarrow [k']_{1 \dots \kappa}$  // We only need  $\kappa$  bits of derived key
   $y \leftarrow E(k, x)$  // Use derived key on the input
  return  $y$ 
end

```

We call x' the *group selector* and k the *derived key*. The j applications of E_{k_i} are to deal with the possibility that $\kappa > n$, and the truncating of k' to κ bits is to handle the possibility that the key length κ might not be a multiple of the block length n . (More strange is the discarding of bits from the x , namely the $x \gg d$. This is for efficiency, as we will explain below.) As an example, if $E = \text{DES}$, so that $\kappa = 56$ and $n = 64$, we would have $j = 1$, so the key of F is just a 56-bit DES key k_1 , the derived key k' is the first 56 bits of $\text{DES}_{k_1}(x')$, and the output is $\text{DES}_{k'}(x)$. If E is TDES (two-key triple-DES), so that $\kappa = 112$ and $n = 64$, we would have $j = 2$, so the key for F is a pair $k_1 k_2$ of TDES keys, the derived key k' is the first 112 bits of $\text{TDES}_{k_1}(x') \text{TDES}_{k_2}(x')$, and the output is $\text{TDES}_{k'}(x)$.

Notice that for fixed $k_1 \cdots k_j$, if two n -bit strings determine the same group selector then they generate the same derived key, and this happens if the two strings agree in the first $n - d$ bits. Accordingly, we cluster together all points that have the same group selector into what we call a *common key group*. Thus there are a total of 2^{n-d} common key groups. For any $\alpha \in \{0, 1\}^{n-d}$ we define $\text{ckg}_\alpha = \{x : [x]_{1 \dots n-d} = \alpha\}$ as the α -th common key group. Identifying strings with integers in the natural way, the i -th common key group consists of the integers $(i-1)2^d, \dots, i2^d - 1$.

EFFICIENCY. Recall that the nominal way to encrypt using $F = \text{Fn}^d E$ involves applying F to a single key k and successive ctr-values. By dropping the least significant d bits of this counter, one needs to recompute k' only once every 2^d invocations of F . Of course an implementation would need to record the last derived key and refrain from re-computing it. Doing this makes the amortized cost to compute F just $(1 + j2^{-d})$ times the cost of computing E . For many ciphers this is an underestimate because of additional cost associated to changing the key. In fact, the cost of changing the key some block ciphers is high, which is why we don't want to do it very often.

VARIATIONS. How exactly one drops bits of x is not so important. For example, instead of shifting to the right one could zero-out the least significant d bits. This makes no difference in the analysis.

We have constructed $F = \text{Fn}^d E$ to be a map $F : \{0, 1\}^{j\kappa} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$. If one prefers, define k_i by $E_k((i))$ and then let $\overline{F}_k(x) = E_{k'}(x)$, where k' is the first κ bits of $E_{k_1}(x \gg d) \parallel \cdots \parallel E_{k_j}(x \gg d)$. Now \overline{F} uses a κ -bit key, just like F . The analysis of F lifts to \overline{F} with just a tiny loss in quantitative security.

4 Definitions

Here we give the more precise definitions of security in the two models in which we will be analyzing our construction, namely the (standard) “complexity theoretic” model and the Shannon model.

Recall that in Section 2 we discussed security of F and E via functions $\text{SEC}_F^{\text{prf}}(q, t)$ and $\text{SEC}_E^{\text{prp}}(q, t)$. Their meaning changes according to the model in a simple way:

- In the complexity theoretic model they are $\text{CSEC}_F^{\text{prf}}(q, t)$ and $\text{CSEC}_E^{\text{prp}}(q, t)$, respectively, these quantities being defined in Section 4.1 below, and
- In the ideal cipher model, they are $\text{ISEC}_F^{\text{prf}}(q, t)$ and $\text{ISEC}_{\kappa, n}^{\text{prp}}(q, t)$, respectively, these quantities being defined in Section 4.2 below, where F refers to the transformation that takes E into F . (In our case, $F = \text{Fn}^d$).

PRELIMINARIES. If S is a probability space then $g \leftarrow S$ denotes the operation of selecting g at random according to the distribution specified by S . If S is a set it is viewed as endowed with the uniform distribution, so that $g \leftarrow S$ means that g is selected uniformly at random from set S . If y is not a set then $g \leftarrow y$ is a simple assignment statement, assigning g the value y . (It is thus equivalent to $g \leftarrow \{y\}$.) Let Perm_n denote the set of all permutations $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$. Let Rand_n denote the set of all functions $\rho : \{0, 1\}^n \rightarrow \{0, 1\}^n$. Let $\text{BC}_{\kappa, n}$ be the set of all maps $E : \{0, 1\}^\kappa \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ such that $E(k, \cdot) \in \text{Perm}_n$ for all $k \in \{0, 1\}^\kappa$. Let $\text{RF}_{\kappa, n}$ be the set of all maps $R : \{0, 1\}^\kappa \times \{0, 1\}^n \rightarrow \{0, 1\}^n$.

A family of functions with key length κ and block length n is a map $G : \{0, 1\}^\kappa \times \{0, 1\}^n \rightarrow \{0, 1\}^n$, that is, $G \in \text{RF}_{\kappa, n}$. Each κ -bit key k specifies the map $G_k \stackrel{\text{def}}{=} G(k, \cdot) \in \text{Rand}_n$. This map is not necessarily a permutation. If G_k is a permutation for each $k \in \{0, 1\}^\kappa$ (ie., $G \in \text{BC}_{\kappa, n}$) then we call G a family of permutations, or a block cipher. We view G as a probability space over Rand_n given by choosing functions via a uniform choice of the underlying key; that is, $g \leftarrow G$ is the same as $k \leftarrow \{0, 1\}^\kappa$; $g \leftarrow G_k$.

Given a block cipher E , the block cipher $E^{-1} : \{0, 1\}^\kappa \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ is defined by $E^{-1}(k, y)$ being the unique point x such that $E(k, x) = y$. We interchangeably write $E_k^{-1}(y)$ and $E^{-1}(k, y)$.

An adversary is an algorithm A with access to some number of oracles. Oracles are denoted as superscripts to A , as in $A^{E, E^{-1}, F}$. An oracle responds to its query in unit time.

4.1 Complexity theoretic model

We will have two measures of security: the strength of G as a PRF and the strength of G as a PRP. We follow [6] in the manner in which the basic notion of [9] is “concretized.”

First, we need the concept of advantage, which for emphasis we call the “computational advantage” and write CADV . Let D be an algorithm (a “distinguisher”) taking an oracle for a function g , and let G_1, G_2 be two families of functions with the same block length. We define

$$\text{CADV}_{G_1, G_2}(D) = \Pr[g \leftarrow G_1 : D^g = 1] - \Pr[g \leftarrow G_2 : D^g = 1].$$

Now, suppose F is a family of functions, and E is a family of permutations. We let

$$\left. \begin{aligned} \text{CADV}_F^{\text{prf}}(D) &= \text{CADV}_{F, \text{Rand}_n}(D) \\ \text{SEC}_F^{\text{prf}}(q, t) &= \max_D \{ \text{CADV}_F^{\text{prf}}(D) \} \end{aligned} \right| \left. \begin{aligned} \text{CADV}_E^{\text{prp}}(D) &= \text{CADV}_{E, \text{Perm}_n}(D) \\ \text{SEC}_E^{\text{prp}}(q, t) &= \max_D \{ \text{CADV}_E^{\text{prp}}(D) \} \end{aligned} \right\}$$

Here the first quantity measures the advantage D has in distinguishing random members of F (resp. E) from truly random functions (resp. permutations) of the same block length. The second quantity is the maximum advantage attainable using some amount of resources, in this case the number q of oracle queries and the running time t . For simplicity, when we speak of an adversary's time we mean the adversary's actual running time plus the size of the encoding of the adversary (relative to some fixed encoding scheme), so we have a single parameter t to capture time plus description size. The maximum here is over all distinguishers D that make up to q oracle queries and have running time bounded by t .

4.2 Ideal block cipher model

The Shannon model [20] treats E as a random block cipher. This means that each E_k is taken to be random permutation on n -bit strings. Let FE be some operator on E which returns a new family of functions, and say the new family has key length κ^* but the block length is still n . (For us, $F = \text{Fn}^d$ and $\kappa^* = j\kappa$ where $j = \lceil \kappa/n \rceil$.) As modeled by [8], the adversary that attacks F is given oracles for $E(\cdot, \cdot)$ and $E^{-1}(\cdot, \cdot)$ — as well as an oracle f where either $f(\cdot) = F(k^*, \cdot)$ for $F = FE$ and k^* a randomly chosen key in $\{0, 1\}^{\kappa^*}$, or else $f(\cdot) = \rho(\cdot)$, for a random function $\rho : \{0, 1\}^n \rightarrow \{0, 1\}^n$. We investigate the adversary's advantage in determining what type of oracle f is. This is defined as:

$$\begin{aligned} \text{IADV}_F^{\text{prf}}(A) &= \Pr \left[E \leftarrow \text{BC}_{\kappa, n}; k^* \leftarrow \{0, 1\}^{\kappa^*}; f \leftarrow (FE)_{k^*} : A^{E, E^{-1}, f} = 1 \right] \\ &\quad - \Pr \left[E \leftarrow \text{BC}_{\kappa, n}; f \leftarrow \text{Rand}_n : A^{E, E^{-1}, f} = 1 \right]. \end{aligned}$$

The advantage A gains depends, in part, on the number of queries q she asks of f and the total number of queries t she asks of E and E^{-1} . We are interested in

$$\text{ISEC}_F^{\text{prf}}(q, t) = \max_A \{ \text{IADV}_F^{\text{prf}}(A) \},$$

the maximum being over all adversaries that make up to q queries to the f oracle and up to t queries to the E and E^{-1} oracles.

This is an information-theoretic setting: the adversary has unlimited computational power. If we think of E as a concrete block cipher, and not an idealized one, then attacks in this model correspond to attacks in which the adversary exploits no characteristics specific to the block cipher, only “generic” features of the construction F we are analyzing. Thus, security guarantees from results in this model are weaker than those from results in the model above, yet they do have some meaning. We use the Shannon model when technical difficulties prevent us from getting bounds as good as we would like in the complexity theoretic model.

NOTE. The goal will be to upper bound $\text{ISEC}_F^{\text{prf}}(q, t)$ as a function of t, q, κ, n . As such we don't really need any notion of $\text{ISEC}_{\kappa, n}^{\text{prp}}(q, t)$, the security of the block cipher, because the latter is assumed ideal, but there are two reasons to define it anyway. First, to maintain a uniform security treatment across the models, and in particular be consistent with Section 2; second, because it is indeed the quantity with which we wish to compare $\text{ISEC}_F^{\text{prf}}(q, t)$.

We define $\text{ISEC}_{\kappa, n}^{\text{prp}}(q, t)$ as the maximum, over all adversaries A of the specified resources, of the quantity

$$\Pr \left[E \leftarrow \text{BC}_{\kappa, n}; k \leftarrow \{0, 1\}^{\kappa}; f \leftarrow E_k : A^{E, E^{-1}, f} = 1 \right] \\ - \Pr \left[E \leftarrow \text{BC}_{\kappa, n}; f \leftarrow \text{Perm}_n : A^{E, E^{-1}, f} = 1 \right].$$

Notice that this quantity is not zero. For $q > 1$ and large n we would expect it to be about $t \cdot 2^{-\kappa}$, corresponding to an exhaustive key search attack.

5 Security of the Fn Construction

We summarize both proven security guarantees and attacks that indicate the tightness of the bounds in them.

5.1 Security in the complexity theoretic model

Here we refer to the notions of security of Section 4.1. We assume E is a PRP family and show our construction is a PRF family, via a reduction. We do this only for the case where the key length, κ , is identical to the block length, n , and we drop no bits, namely $d = 0$.

Theorem 1. *Let $\kappa = n$ be a positive integer and let $E: \{0, 1\}^{\kappa} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a family of permutations whose security as a PRP family is described by security function $\text{CSEC}_E^{\text{prp}}(\cdot, \cdot)$. Let $F: \{0, 1\}^{\kappa} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be our construction for the case of no bit dropping, namely $F = \text{Fn}^0 E$. Its security as a PRF is described by function $\text{CSEC}_F^{\text{prf}}(\cdot, \cdot)$ which for any number of queries $q \leq 2^n/2$ and time t can be bounded as follows:*

$$\text{CSEC}_F^{\text{prf}}(q, t) \leq \text{CSEC}_E^{\text{prp}}(q, t') + q \cdot \text{CSEC}_E^{\text{prp}}(3, t') + \frac{q^2}{2^{2n}}$$

where $t' = t + O(q) \cdot (\kappa + n + \text{Time}_E)$.

Proof. See [7]. □

The bound here looks good at first glance. The first term, namely $\text{CSEC}_E^{\text{prp}}(q, t')$, is saying the security of F as a PRF is related to that of E as a PRP for essentially the same resources: we can't ask better. The last term, namely $q^2/2^{2n}$, is negligible. What about the middle term, namely $q \cdot \text{CSEC}_E^{\text{prp}}(3, t')$? Intuitively, $\text{CSEC}_E^{\text{prp}}(3, t')$ is small: what can you do in three queries? This view is deceptive because one should not forget the time t' . One can spend it in exhaustive key search, and thus $\text{CSEC}_E^{\text{prp}}(3, t')$ can be $\Omega(t'2^{-\kappa})$. But (dropping constants) this is at least $q2^{-\kappa}$ so the second term in our bound looks like $q^22^{-\kappa}$. Since $\kappa = n$ this is q^22^{-n} .

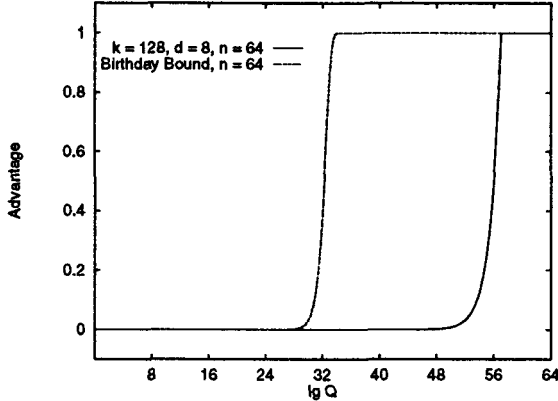


Fig. 1. *Right curve:* Illustrating Theorem 2, our upper bound on the adversary's advantage in distinguishing $F = Fn^d E$ from a random function, assuming $n = 64$, $\kappa = 128$, and $d = 8$. Here E is a random permutation and the horizontal axis $Q = \max(q, t)$ is the maximum of the number of consecutive f -queries and the total number of E, E^{-1} queries. *Left curve:* The birthday bound for the same choice of parameters.

So these bounds are not proof that the security of F goes beyond the birthday bound. It would be nice to improve the above result. However, even the proof of the above is not exactly trivial, and this is one reason we include the result in this paper: we hope its ideas are food for thought towards an extension.

As far as we can tell, the difficulties in extending the above result are technical rather than arising from any weakness in the construction. (We could be wrong.) Is there any other way we can give some meaningful evidence of the strength of the construction? We do this by analyzing it in the Shannon model.

5.2 Security in the ideal block cipher model

The theorem below looks at the most general version of the $F = Fn^d E$ construction, when the number d of bits dropped is arbitrary and no restrictions are made on κ, n , in the model of Section 4.2, where E is an ideal cipher. We obtain very strong results, showing security not only beyond the birthday bound, but nearly as good as one could hope for.

As we noted in Section 2, an important mode of operation for our construction will be when the values to which $F_{k_1 \dots k_j}$ are being applied are successive counter values. Indeed, the bit dropping is done precisely to have maximum efficiency in this mode: as explained in Section 3, in this case, the amortized cost of computing F is just $(1 + j/2^d)$ times that of computing E , a negligible overhead. Accordingly, this is the case to which the following security analysis pertains. (Though some later analyses are more general.)

Theorem 2. *Let n, κ be positive integers and d, q, t, \hat{t} be non-negative integers with $0 \leq d < n$ and let $F = Fn^d$. Let A be an adversary with three oracles, $E(\cdot, \cdot)$, $E^{-1}(\cdot, \cdot)$, and $f(\cdot)$, who asks the numbers $0, \dots, q - 1$ of its f -oracle (so*




Fig. 2. Varying the parameters of Theorem 2 — our upper bound on the adversary’s advantage in distinguishing $F = F_n^d E$ from a random function, with the horizontal axis $Q = \max(q, t)$ as in the previous figure. **Left:** Varying key length κ . **Right:** Varying bits dropped d . For both pictures $n = 64$.

that these refer to $\hat{q} = \lceil q2^{-d} \rceil$ common key groups), and asks at most t total queries of its E - and E^{-1} -oracles, these referring to no more than \hat{t} common key groups. Let $j = \lceil \kappa/n \rceil$. Then $\text{IADV}_F^{\text{prf}}(A) \leq$

$$\frac{\hat{q}^5 + \hat{t}^5}{120} \cdot 2^{-4\kappa} + (j^2 + 2j\hat{q} + tj + t) \cdot 2^{-\kappa} + \hat{q}^2 2^{2d-n+3} + \hat{t}^2 2^{d-n-\kappa+2} .$$

Proof. See [7]. □

Thus the advantage remains low until $q \approx 2^{4\kappa/5}$, $t \approx 2^{4\kappa/5}$, $q \approx 2^n$, or $t \approx 2^{(n+\kappa)/2}$. We speculate that the first two of these conditions can be further improved to $2^{(1-\epsilon)\kappa}$ (and they are already very small in their current form), so a reasonable summary is to say that the construction is good until $q \approx \min\{2^\kappa, 2^{n-2d}\}$ or $t \approx \min\{2^\kappa, 2^{(d-n-\kappa)/2}\}$.

In Fig. 1 we illustrate our bound for the case of a block cipher with parameters $n = 64$, $\kappa = 128$, and dropping $d = 8$ bits. The bound indicates that one must ask about 2^{55} queries before one can hope to distinguish F_k from a random function with advantage $1/e$. (This $1/e$ -convention is a convenient way to summarize security.) For comparison, if you let $F = E$ you get the usual birthday-attack curve, which indicates that it takes but 2^{32} queries before an adversary can get like advantage at distinguishing E_k from a random function.

In Fig. 2 we illustrate our bound by showing the effect on advantage of changing either the key length (left-hand plot) or the value of d (right-hand plot). We assume a block size of $n = 64$ bits. The adversary’s maximum advantage decreases with increasing key length, but this effect soon saturates. The construction has worse demonstrated security for larger values of d , but the effect is not that dramatic, and there is little reason to select a very large value of d , anyway.

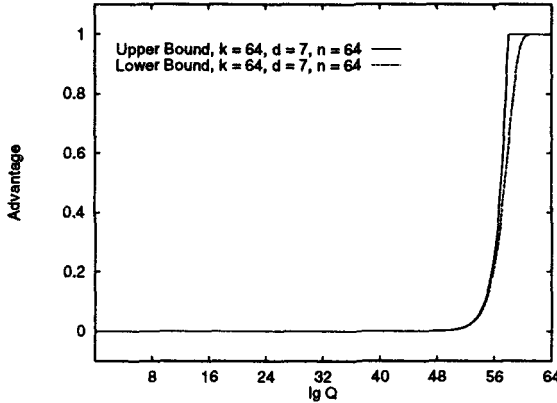


Fig. 3. With typical parameters our bounds are tight. Illustrating Propositions 1 and 2 and Theorem 2 for $n = 64$, $\kappa = 64$, $d = 7$. The horizontal axis Q is the same as in the previous figures.

It is important to understand the difference between the results here and the result of Section 5.1. The “type” of security guarantee is better in the latter, since we are saying that security in the sense of a PRP (using the standard notion of a PRP) translates into security in the sense of a PRF (using the standard notion of a PRF). The results here are only about ideal ciphers, which only guarantees security against generic attacks. Yet, generic attacks are an important and easy to mount class of attacks, and a proof of security against them, especially with such strong bounds, is certainly meaningful. Eventually we hope strong results will emerge in the other model (as well as for other PRP-to-PRF constructions).

5.3 Attacks / Lower bounds

In Propositions 1 and 2 we present the best attacks that we know on our construction. These translate into lower bounds on the security of $\text{Fn}^d E$. We present two adversaries: one which becomes successful when $q \approx 2^{n-d}$, and one which becomes successful when $t \approx 2^\kappa$. This is done in the Shannon model, but in this case (of attacks) this is not a restriction; if we can attack ideal ciphers we can certainly attack real ones. Thus, the results here should be viewed as complementing Theorem 2, telling us how close to tight is the analysis in the latter.

Proposition 1. *Let n, κ be positive integers and d, q non-negative integers with $0 \leq d < n$, and let $F = \text{Fn}^d$. Then there is an adversary CS which asks at most q of an f oracle, no queries of the E or E^{-1} oracles, and achieves advantage*

$$\text{IADV}_F^{\text{prf}}(\text{CS}) \geq 1 - e^{-\lfloor q2^{-d} \rfloor \cdot (2^d - 1) \cdot 2^{d-n-1}}.$$

Proof. See [7].

□

Proposition 2. *Let n, κ be positive integers and t, d, c be non-negative integers with $0 \leq d < n$, let $F = \text{Fn}^d$, and let $j = \lceil \kappa/n \rceil$. Then there is an adversary KS which asks c queries of her f oracle, t queries of her E oracle, and achieves advantage*

$$\text{IADV}_F^{\text{prf}}(KS) = \min\{1, \lfloor t/(cj + c) \rfloor \cdot 2^{-j\kappa}\} - t2^{-cn}.$$

Proof. See [7]. □

The first lower bound is around $1 - e^{-q2^{d-n-1}}$, while the second one is around $t2^{-j\kappa}$. These become significant when $q \approx 2^{n-d}$ or $t \approx 2^{j\kappa}$. The point of giving these lower bounds is to see how tight is Theorem 2. As Fig. 3 illustrates, the bounds are quite close for realistic parameters. On the same plot we graph our upper and lower bound for $\kappa = 56$, $n = 64$, and $d = 7$. The curves almost coincide.

Acknowledgments

The first author was supported by a 1996 Packard Foundation Fellowship in Science and Engineering, and by NSF CAREER Award CCR-9624439. The second two authors were supported by NSF CAREER Award CCR-9624560 and a MICRO grant from RSA Data Security, Inc.

References

1. W. AIELLO AND R. VANKETESAN, "Foiling birthday attacks in output-doubling transformations." *Advances in Cryptology – Eurocrypt 96 Proceedings*, Lecture Notes in Computer Science Vol. 1070, U. Maurer ed., Springer-Verlag, 1996.
2. M. BELLARE, O. GOLDREICH AND H. KRAWCZYK, personal communications, 1995.
3. M. BELLARE, R. CANETTI AND H. KRAWCZYK, "Pseudorandom functions revisited: The cascade construction and its concrete security." *Proceedings of the 37th Symposium on Foundations of Computer Science*, IEEE, 1996.
4. M. BELLARE, A. DESAI, E. JOKIPII AND P. ROGAWAY, "A concrete security treatment of symmetric encryption." *Proceedings of the 38th Symposium on Foundations of Computer Science*, IEEE, 1997.
5. M. BELLARE, R. GUÉRIN AND P. ROGAWAY, "XOR MACs: New methods for message authentication using a finite pseudorandom function." *Advances in Cryptology – Crypto 95 Proceedings*, Lecture Notes in Computer Science Vol. 963, D. Coppersmith ed., Springer-Verlag, 1995.
6. M. BELLARE, J. KILIAN AND P. ROGAWAY, "The security of cipher block chaining." *Advances in Cryptology – Crypto 94 Proceedings*, Lecture Notes in Computer Science Vol. 839, Y. Desmedt ed., Springer-Verlag, 1994.
7. M. BELLARE, T. KROVETZ AND P. ROGAWAY, "Luby-Rackoff backwards: Increasing security by making block ciphers non-invertible." Full version of this paper, available at <http://www-cse.ucsd.edu/users/mihir>
8. S. EVEN AND Y. MANSOUR, "A construction of a cipher from a single pseudorandom permutation." *Advances in Cryptology – ASIACRYPT 91 Proceedings*, Lecture Notes in Computer Science Vol. 739, H. Imai, R. Rivest and T. Matsumoto ed., Springer-Verlag, 1991.

9. O. GOLDBREICH, S. GOLDWASSER AND S. MICALI, "How to construct random functions," *Journal of the ACM*, Vol. 33, No. 4, 1986, pp. 210–217.
10. S. GOLDWASSER AND S. MICALI, "Probabilistic encryption." *J. of Computer and System Sciences*, Vol. 28, April 1984, pp. 270–299.
11. J. KILIAN AND P. ROGAWAY, "How to protect DES against exhaustive key search." *Advances in Cryptology – Crypto 96 Proceedings*, Lecture Notes in Computer Science Vol. 1109, N. Kobnitz ed., Springer-Verlag, 1996.
12. M. LUBY, *Pseudorandomness and Cryptographic Applications*. Princeton University Press, 1996.
13. M. LUBY AND C. RACKOFF, "How to construct pseudorandom permutations from pseudorandom functions." *SIAM J. Comput*, Vol. 17, No. 2, April 1988.
14. M. Matsui, "The first experimental cryptanalysis of the Data Encryption Standard." *Advances in Cryptology – Crypto 94 Proceedings*, Lecture Notes in Computer Science Vol. 839, Y. Desmedt ed., Springer-Verlag, 1994, pp. 1–11.
15. U. MAURER, "A simplified and generalized treatment of Luby-Rackoff pseudorandom permutation generators." *Advances in Cryptology – Eurocrypt 92 Proceedings*, Lecture Notes in Computer Science Vol. 658, R. Rueppel ed., Springer-Verlag, 1992, pp. 239–255.
16. M. NAOR AND O. REINGOLD, "On the construction of pseudo-random permutations: Luby-Rackoff revisited." *Proceedings of the 29th Annual Symposium on Theory of Computing*, ACM, 1997.
17. J. PATARIN, "Improved security bounds for pseudorandom permutations." Fourth ACM Conference on Computer and Communications Security, 1997.
18. J. PATARIN, "About Feistel schemes with six (or more) rounds." To appear in *Fast Software Encryption (FSE5)*, March 1998.
19. J. PIEPRZYK, "How to construct pseudorandom permutations from single pseudorandom functions." *Advances in Cryptology – Eurocrypt 90 Proceedings*, Lecture Notes in Computer Science Vol. 473, I. Damgård ed., Springer-Verlag, 1990 pp. 140–150.
20. C. SHANNON, "Communication theory of secrecy systems." *Bell Systems Technical Journal*, 28(4), 656–715 (1949).
21. V. SHOUP, "On fast and provably secure message authentication based on universal hashing." *Advances in Cryptology – Crypto 96 Proceedings*, Lecture Notes in Computer Science Vol. 1109, N. Kobnitz ed., Springer-Verlag, 1996.
22. M. WEGMAN AND L. CARTER, "New hash functions and their use in authentication and set equality." *J. of Computer and System Sciences* 22, 265–279 (1981).
23. Y. ZHENG, T. MATSUMOTO AND H. IMAI, "Impossibility and optimality results on constructing pseudorandom permutations." *Advances in Cryptology – Crypto 90 Proceedings*, Lecture Notes in Computer Science Vol. 537, A. J. Menezes and S. Vanstone ed., Springer-Verlag, 1990, pp. 412–422.