Concrete Security Characterizations of PRFs and PRPs: Reductions and Applications

Anand Desai[1] and Sara Miner[2]

[1] Bell Labs Research Silicon Valley, 3180 Porter Drive, Palo Alto, CA 94304, USA.
adesai@research.bell-labs.com
[2] Dept. of Computer Science & Engineering, University of California at San Diego, 9500 Gilman Drive, La Jolla, CA 92093, USA.
sminer@cs.ucsd.edu

Abstract. We investigate several alternate characterizations of pseudorandom functions (PRFs) and pseudorandom permutations (PRPs) in a concrete security setting. By analyzing the concrete complexity of the reductions between the standard notions and the alternate ones, we show that the latter, while equivalent under polynomial-time reductions, are weaker in the concrete security sense. With these alternate notions, we argue that it is possible to get better concrete security bounds for certain PRF/PRP-based schemes. As an example, we show how using an alternate characterization of a PRF could result in tighter security bounds for some types of message authentication codes. We also use this method to give a simple concrete security analysis of the counter mode of encryption. In addition, our results provide some insight into how injectivity impacts pseudorandomness.

1 Introduction

Pseudorandom functions (PRFs) and pseudorandom permutations (PRPs) are extremely useful and widely used tools in cryptographic protocol design, particularly in the setting of private-key cryptography. In this paper, we study several different notions of security for these objects. Specifically, we study these notions in a concrete security framework, and we show how different characterizations may be used to derive better security bounds for some commonly used private-key cryptographic protocols.

1.1 Descriptions of Notions

The notion of a PRF family was proposed by Goldreich, Goldwasser and Micali [8]. In such a family, each function is specified by a short key, and can be easily computed given the key. Yet it has the property that telling apart a function sampled from the PRF family and one from a random function family, given adaptive access to the function as a black-box, is computationally infeasible (for someone who does not know the key). This is the standard notion of a PRF, and (to distinguish it from alternate notions) we refer to it in this paper as the

T. Okamoto (Ed.): ASIACRYPT 2000, LNCS 1976, pp. 503–516, 2000.

PRF notion. Luby and Rackoff extended the above to permutation families by introducing the notion of a PRP family [11]. The reference family for defining the security of a PRP family can be that of random functions, as in [11], or that of random *permutations*, a practice started by Bellare, Kilian and Rogaway [5]. We adopt the definition of [5] and refer to it here as the PRP notion.

Alternate characterizations. PRFs may be characterized in several ways other than the standard notion. We are particularly interested in one way suggested in the very paper that introduced the standard notion [8]. This alternate notion can be described informally through the following interactive protocol: a distinguisher who is given adaptive oracle access to the function obtains the output of the function on some points of its choice through oracle queries. It then outputs a point that has not yet been queried and gets back, based on a hidden coin flip, either the output of the function on that point or a uniformly distributed point in the range of the function. It should be computationally infeasible for the distinguisher to guess which of the two possibilities it was presented. We call this notion indistinguishable-uniform functions or IUF, to distinguish it from the standard notion PRF. A similar notion may be defined for permutation families, and we call this IUP for indistinguishable-uniform permutations.

We also consider another notion that is normally associated with the security of encryption schemes. In this notion too, the distinguisher is given adaptive oracle access to the function. It then outputs *two* new points and, based on a hidden coin flip, is presented with the output of the function on one of them. We require that a computationally-restricted distinguisher have negligible success in telling apart the two cases. In this paper, we refer to this notion as IPF, for indistinguishable-point functions. We show that this notion does not imply pseudorandomness for functions. However, when we consider the analogous notion for permutations, which we call IPP (indistinguishable-point permutations), we find that pseudorandomness is captured.

1.2 Concrete Security and Reductions Among the Notions

Making a break from the traditional approach of presenting PRF families in an asymptotic way, Bellare, Kilian and Rogaway began the practice of explicitly specifying the resources determining security and paying particular attention to the quality of security reductions [5]. This approach forms the basis of concrete security analysis and has been used in many subsequent works [4,2,3]. One benefit of this approach is that it enables the comparison (and classification as weaker or stronger) of polynomially-equivalent notions in cryptography. Paying attention to the concrete complexity of reductions between notions is important in practice, as inefficient reductions translate to a penalty either in security assurance or in running time.

Reductions Among the Notions. Under polynomial-time reductions, the equivalence between the notions of PRF and IUF has been established by Goldreich et al [8]. (In fact, the concrete security bounds we derive in our reductions between these notions are implicit in theirs.) We establish that our reductions

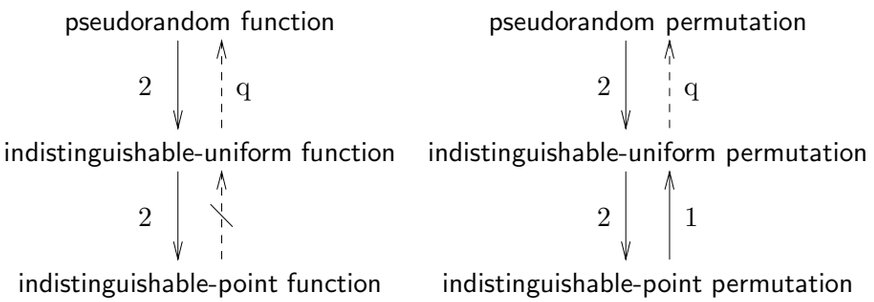

Fig. 1. *Relating the notions. A solid arrow from notion A to notion B means that there is a security-preserving reduction from A to B. A broken arrow indicates a reduction that is not security-preserving. The arrows are labeled by the loss-factor of the reduction. A hatched arrow means that there is no polynomial-time reduction.*

are tight. Additionally, we relate the notions of PRP and IUP. The reductions between these two permutation notions are the same as those between the corresponding notions for functions.

Furthermore, we show that IUP and IPP are equivalent, up to a small constant factor in the reduction. However, as mentioned above, a different picture emerges when we look at the corresponding notions for functions. It turns out that IPF and IUF (or PRF) are not equivalent, even in just an asymptotic sense. We show that IPF is a strictly weaker notion, in that there are function families which are secure in the IPF sense, but completely insecure in the IUF sense. A summary of the reductions is given in Figure 1.

1.3 Motivation: Tighter Security Analyses

Our demonstration that the alternate notions we consider here are weaker in the concrete security sense than the standard notions might be seen as an argument *against* using any of them. Yet we will recommend their use in certain circumstances (to complement, rather than replace the standard notions).

In a concrete security analysis of a protocol which is based on a particular primitive, the security of the protocol is related to that of the underlying primitive in a precise way. If we know the concrete security of a protocol in terms of the security of the underlying primitive under one notion, it is easy to translate this to the security of the protocol in terms of the security of the primitive under a weaker notion. We simply use the appropriate security reduction between the notions. We then see a drop in the translated security, reflecting the gap in the reduction between the notions. However, we show that it is sometimes possible to directly reduce the security of the protocol to that of the underlying primitive under a weaker notion *without* the expected drop in security. Such a situation exists when the weaker notion somehow "meshes" better with the notion of security for the protocol.

We make the above discussion more concrete with two examples: message authentication codes and symmetric encryption schemes. The security of (deterministic) message authentication codes (MACs) is captured by the notion of unpredictable functions [10,1,12]. In the context of MACs, this means that an adversary who is given valid MACs on some messages of its choice will be unlikely to succeed in outputting a "new" message (that is, one different from those whose MACs it was been given) along with a valid MAC on that message. It is well-known that any PRF is unpredictable (i.e. a secure MAC) [8]. Moreover, the reduction from unpredictable functions to PRFs is almost tight [5]. We show that using a direct reduction from unpredictable functions to IUF, one can obtain *exactly* the same bounds. This represents a tightening of the analysis, as we expect security of a PRF in the IUF sense to be smaller than the security in the standard PRF sense. (Our reductions show that the security in the IUF sense will never be more than a constant factor of 2 greater than the security in the PRF sense and will typically be a quantitative factor less.)

Now let us examine in what sense IUF "meshes" better with the notion of unpredictable functions. The quantitative drop in security in the reduction from PRF to IUF can be traced to the fact that under IUF the distinguisher must decide given *one* challenge whereas, under PRF, every response to a query potentially constitutes a "challenge". Like IUF, the notion of unpredictable functions also has a single distinguished challenge. In the reduction to PRF, however, we cannot really take any advantage of the source of the strength of this notion, and hence the bounds derived are not as tight as what could be achieved otherwise.

Another example of a notion with a distinguished challenge phase is the standard indistinguishability of encryptions notion of security for encryption schemes [9,3]. Here again, using the notion of IUF instead of the standard PRF, we can hope to tighten analysis of PRF-based encryption schemes. We do this for the *counter mode* of encryption.

1.4 Related Work

We have already mentioned the foundational work on PRFs and PRPs [8,11] and the concrete security analysis of these objects [5,4]. Our approach in this work follows that of Bellare et al [3], who compared and classified notions of security for symmetric encryption schemes according to the concrete complexity of reductions. A concrete security analysis of various symmetric encryption schemes, including the counter mode, is given in that paper. Naor and Reingold have explored the relationship between unpredictable functions and PRFs under different attack models [12].

2 Definitions and Notation

We describe different notions of security for (finite) function families in this section. A function family is a keyed multi-set F of functions where all functions have the same domain and range. To pick a function f from family F means to

pick a key a, uniformly from key space $\mathsf{Keys}(F)$ of F, and let $f = F_a$. A family F has input length l and output length L if each $f \in F$ maps $\{0,1\}^l$ to $\{0,1\}^L$.

We let $R_{l,L}$ denote the function family consisting of all functions with input length l and output length L. Similarly, we let P_l denote the set of all permutations on l-bit strings.

A function family F is pseudorandom if the input-output behavior of F_a is indistinguishable from the behavior of a random function of the same domain and range. This is formalized via the notion of statistical tests of Goldreich et al [8]. Our concrete security formalizations follow those of Bellare et al [5].

We first informally describe the two additional notions (IUF and IPF) for function families considered in this paper. The corresponding notions for permutation families (IUP and IPP) are analogous to these, and so we skip their description. At the end of this section, we formally define all these notions (for both function and permutation families).

Indistinguishable-Uniform Functions. This is an adaptation of a notion given by Goldreich et al [8]. The idea is that a distinguisher should not be able to distinguish the output of the PRF from a uniformly distributed value in the range of the function. The formalization considers two different experiments. In both experiments we start by choosing a random key $a \leftarrow \mathsf{Keys}(F)$, specifying a function F_a. In the first phase, the distinguisher is given an oracle for F_a and allowed to query this oracle on points of its choice. It then outputs a point x that has not been queried yet and some state information s that it may want to preserve for use during the second phase. In one experiment, it receives in response the value $F_a(x)$. In the other experiment, it receives a uniformly distributed value in the range of F. The PRF family is "good" if no "reasonable" distinguisher can obtain significant advantage in distinguishing the two experiments.

Indistinguishable-Point Functions. This is an adaptation of the indistinguishability of encryptions notion of security for encryption schemes. Here again we imagine a distinguisher A that runs in two phases. In the find phase, given adaptive access to an oracle for the function, it comes up with a pair of points x_0, x_1 that it has not queried yet and some state information s. In the guess phase, given the output of the function y on one of these points and s, it must identify which of the two points goes with y.

It is interesting that the notion IPP does capture pseudorandomness for permutation families. For most other primitives, we find that an indistinguishable-point-based characterization is weaker than an indistinguishable-uniform-based characterization. This is true for encryption schemes and turns out to be true for function families, as well. Observe that, for encryption schemes, we are usually concerned with this weaker characterization, because it captures the desired security requirements.

Formal Definitions. For each of the six notions we consider in this paper, we give definitions using the experiments defined in Figure 2. First, we consider the function family notions: PRF, IUF, and IPF.

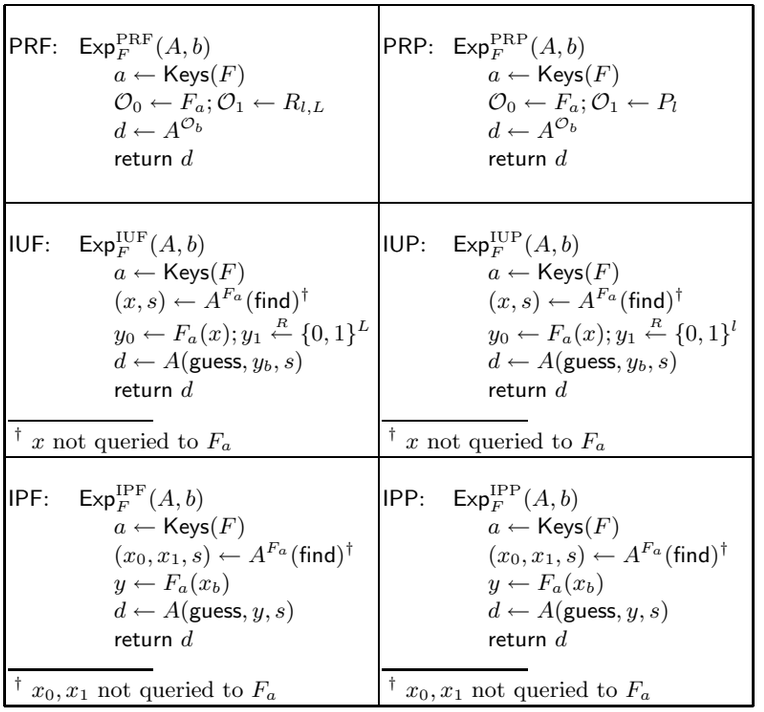

Fig. 2. Experiments defining each of the notions considered in this paper.

Definition 1. *For each notion* $\mathsf{N} \in \{\mathsf{PRF}, \mathsf{IUF}, \mathsf{IPF}\}$*, let* $F\colon\ \mathsf{Keys}(F) \times \{0,1\}^l \to \{0,1\}^L$ *be a finite function family. For an adversary* A *and* $b = 0, 1$ *define the experiment* $\mathsf{Exp}_F^{\mathsf{N}}(A,b)$*, as given in Figure 2. Define the* advantage *of* A *and the advantage function of* F*, respectfully, as follows. For any integers* $t, q \geq 0$*,*

$$\mathsf{Adv}_F^{\mathsf{N}}(A) = \Pr[\,\mathsf{Exp}_F^{\mathsf{N}}(A,0) = 0\,] - \Pr[\,\mathsf{Exp}_F^{\mathsf{N}}(A,1) = 0\,]$$
$$\mathsf{Adv}_F^{\mathsf{N}}(t,q) = \max_A \{\mathsf{Adv}_F^{\mathsf{N}}(A)\}$$

where the maximum is over all A *with time complexity* t*, making* $\leq q$ *queries.* ∎

Here the "time-complexity" is the worst-case total execution time of the experiment, plus the size of the code of the adversary, in some fixed RAM model of computation. This convention is used for all definitions in this paper.

Next, we turn our attention to the definitions for the corresponding permutation family notions: PRP, IUP, and IPP.

Definition 2. *For each notion* $\mathsf{N} \in \{\mathsf{PRP}, \mathsf{IUP}, \mathsf{IPP}\}$*, let* $F\colon\ \mathsf{Keys}(F) \times \{0,1\}^l \to \{0,1\}^l$ *be a finite permutation family. For an adversary* A *and* $b = 0, 1$ *define*

the experiment $\mathsf{Exp}^{\mathsf{N}}_F(A, b)$*, as given in Figure 2. Define the advantage of* A *and the advantage function of* F*, respectfully, as follows. For any integers* $t, q \geq 0$*,*

$$\mathsf{Adv}^{\mathsf{N}}_F(A) = \Pr[\,\mathsf{Exp}^{\mathsf{N}}_F(A, 0) = 0\,] - \Pr[\,\mathsf{Exp}^{\mathsf{N}}_F(A, 1) = 0\,]$$
$$\mathsf{Adv}^{\mathsf{N}}_F(t, q) = \max_A \{\mathsf{Adv}^{\mathsf{N}}_F(A)\}$$

where the maximum is over all A *with time complexity* t*, making* $\leq q$ *queries.* ▮

3 Reductions Among the Notions

In this section, we formally state the relations shown in Figure 1. The proofs for these results are given in the full version of this paper [7].

We use the notation $A \Rightarrow B$ to indicate a security-preserving reduction from notion A to notion B. $A \rightarrow B$ indicates a reduction (not necessarily security-preserving) from A to B. $A \not\Rightarrow B$ and $A \not\rightarrow B$ are the natural interpretations given the above. This convention is followed for all reductions given in this paper.

3.1 Function Family Notions

The first theorem says that if a function family has certain security in the standard PRF sense, then it has essentially the same security in the IUF sense.

Theorem 1. [PRF $\Rightarrow$ IUF] *For any function family* F *and integers* $t, q \geq 1$*,*

$$\mathsf{Adv}^{\mathrm{IUF}}_F(t, q) \leq 2 \cdot \mathsf{Adv}^{\mathrm{PRF}}_F(t', q)$$

where $t' = t + O(l + L)$*.* ▮

Our next theorem says that if a function family is secure in the IUF sense, then it is also secure in the PRF sense, but the security is quantitatively lower.

Theorem 2. [IUF $\rightarrow$ PRF] *For any function family* F *and integers* $t, q \geq 1$*,*

$$\mathsf{Adv}^{\mathrm{PRF}}_F(t, q) \leq q \cdot \mathsf{Adv}^{\mathrm{IUF}}_F(t', q)$$

where $t' = t + O(l + L)$*.* ▮

The following proposition establishes that the drop in security in the previous theorem was not due to any weakness of our reduction but is, in fact, intrinsic to the notions. We give a concrete example of a function family that has higher security in the PRF sense, with a gap of the same order as in Theorem 2.

Proposition 1. [IUF $\not\Rightarrow$ PRF] *There exists a function family* F *such that*

$$\mathsf{Adv}^{\mathrm{PRF}}_F(t, q) \geq \frac{1}{2} \text{ and } \mathsf{Adv}^{\mathrm{IUF}}_F(t, q) \leq \frac{1}{q}$$

for any integers $t \geq 1$ *and* $1 \leq q \leq 2^{L-1}$*.* ▮

Our next two results demonstrate that the IPF notion is weaker than the other two notions we have considered, and hence does not capture pseudorandomness.

Theorem 3. [IUF ⇒ IPF] *For any function family F and integers $t, q \geq 1$,*

$$\mathsf{Adv}_F^{\mathrm{IPF}}(t, q) \leq 2 \cdot \mathsf{Adv}_F^{\mathrm{IUF}}(t', q)$$

where $t' = t + O(l + L)$. ∎

Proposition 2. [IPF ⇏ IUF] *There exists a function family F such that,*

$$\mathsf{Adv}_F^{\mathrm{IUF}}(t, q) \geq 1 - 2^{-qL} \text{ and } \mathsf{Adv}_F^{\mathrm{IPF}}(t, q) = 0$$

for any integers $t, q \geq 1$. ∎

3.2 Permutation Family Notions

We first give the reductions between PRP and IUP. Our next three claims show that the security bounds we had derived between the notions for function families also hold between the corresponding notions for permutation families.

Theorem 4. [PRP ⇒ IUP] *For any permutation family F and integers $t, q \geq 1$,*

$$\mathsf{Adv}_F^{\mathrm{IUP}}(t, q) \leq 2 \cdot \mathsf{Adv}_F^{\mathrm{PRP}}(t', q)$$

where $t' = t + O(l)$. ∎

Theorem 5. [IUP → PRP] *For any permutation family F and integers $t, q \geq 1$,*

$$\mathsf{Adv}_F^{\mathrm{PRP}}(t, q) \leq q \cdot \mathsf{Adv}_F^{\mathrm{IUP}}(t', q)$$

where $t' = t + O(l)$. ∎

Proposition 3. [IUP ⇏ PRP] *There exists a permutation family F such that*

$$\mathsf{Adv}_F^{\mathrm{PRP}}(t, q) \geq \frac{1}{2} \text{ and } \mathsf{Adv}_F^{\mathrm{IUP}}(t, q) \leq \frac{1}{q}$$

for any integers $t \geq 1$ and $1 \leq q \leq 2^{L-1}$. ∎

Next, we establish that IUP and IPP are of essentially equivalent strength. Note that this is a departure from the relationship that exists between the corresponding function family notions.

Theorem 6. [IUP ⇒ IPP] *For any permutation family F and integers $t, q \geq 1$,*

$$\mathsf{Adv}_F^{\mathrm{IPP}}(t, q) \leq 2 \cdot \mathsf{Adv}_F^{\mathrm{IUP}}(t', q)$$

where $t' = t + O(l)$. ∎

Theorem 7. [IPP ⇒ IUP] *For any permutation family F and integers $t, q \geq 1$,*

$$\mathsf{Adv}_F^{\mathrm{IUP}}(t, q) \leq \mathsf{Adv}_F^{\mathrm{IPP}}(t', q)$$

where $t' = t + O(l)$. ∎

4 Applications

Here, we give some motivation for the use of the IUF characterization of PRF families. As discussed in Section 1, use of this notion gives tighter security bounds for certain cryptographic protocols. We give two such examples in this section.

4.1 The Case of Message Authentication Codes

A message authentication code (MAC) enables two parties who share a secret key to authenticate their transmissions. To be secure, MACs must resist existential forgery under chosen-message attacks [10,5]. For deterministic MACs, this notion matches that of unpredictable functions (UPF) [1,12].

Formally, the notion is captured by allowing a distinguisher A to query a MAC oracle, F_a, where F is a function family and a is a random MAC key. A must then output a point x that has not been queried yet, along with its prediction y for the value of $F_a(x)$.

Definition 3. [Message authentication security: UPF] *Let* F: $\mathsf{Keys}(F)\times\{0,1\}^l \to \{0,1\}^L$ *be a MAC. For an adversary* A *define the following experiment:*

Experiment $\mathsf{Exp}^{\mathrm{UPF}}_F(A)$
$a \leftarrow \mathsf{Keys}(F)$; $(x,y) \leftarrow A^{F_a}$ *//where* x *is a point that* A *has not queried*
If $y = F_a(x)$ *then* $d \leftarrow 0$ *else* $d \leftarrow 1$; Return d.

Define the advantage of A *and the advantage function of* F*, respectfully, as follows. For any integers* $t, q \geq 0$*,*

$$\mathsf{Adv}^{\mathrm{UPF}}_F(A) = \Pr[\,\mathsf{Exp}^{\mathrm{UPF}}_F(A) = 0\,]$$
$$\mathsf{Adv}^{\mathrm{UPF}}_F(t,q) = \max_A \{\mathsf{Adv}^{\mathrm{UPF}}_F(A)\}$$

where the maximum is over all A *with time complexity* t*, making* $\leq q$ *queries.* ▮

PRF families are more well-studied than unpredictable function families and, moreover, are widely available. Hence, the observation that a PRF family constitutes a secure MAC [8] has proven very useful in practice. The following exact security reduction is already known [5].

Proposition 4. [PRF ⇒ UPF] *For any function family* F *and integers* $t, q \geq 1$*,*

$$\mathsf{Adv}^{\mathrm{UPF}}_F(t,q) \leq \mathsf{Adv}^{\mathrm{PRF}}_F(t',q) + 2^{-L}$$

where $t' = t + O(l + L)$. ▮

The reduction is almost tight. Consider now translating the above, to get security as a MAC in terms of the security as a PRF family in the IUF sense. Using Theorem 2 will lead to a drop in security by a factor q. However, by applying a direct reduction, we avoid this expected loss.

Proposition 5. [IUF $\Rightarrow$ UPF] *For any function family F and integers $t, q \geq 1$,*

$$\mathsf{Adv}_F^{\mathrm{UPF}}(t, q) \leq \mathsf{Adv}_F^{\mathrm{IUF}}(t', q) + 2^{-L}$$

where $t' = t + O(l + L)$.

Proof. The reduction is standard. Let A be a forger attacking the MAC F, making at most q oracle queries and running in time at most t, in the experiment $\mathsf{Exp}_F^{\mathrm{UPF}}(A)$. We construct a distinguisher A', making at most q queries and running in time at most t', using the forger A as a subroutine.

Let $\mathcal{O}_f$ be A''s oracle. $A'^{\mathcal{O}_f}$ will run A using $\mathcal{O}_f$ to provide an appropriate simulation of A's oracle, as indicated below.

Algorithm $A'^{\mathcal{O}_f}$

(1) Run A, answering any query u with $\mathcal{O}_f(u)$.
(2) Let $(x, y) \leftarrow A$.
(3) Output (x, y) and receive y' as the challenge.
(4) If $y' = y$ then output 0, else output 1.

For simplicity, we assume that A makes exactly q queries in $\mathsf{Exp}_F^{\mathrm{UPF}}(A)$. It is easy to check that the time and query complexity are as claimed. Next, we compute the advantage of A'.

$$\begin{aligned}\mathsf{Adv}_F^{\mathrm{IUF}}(A') &= \Pr[\,\mathsf{Exp}_F^{\mathrm{IUF}}(A', 0) = 0\,] - \Pr[\,\mathsf{Exp}_F^{\mathrm{IUF}}(A', 1) = 0\,] \\ &= \Pr[\,\mathsf{Exp}_F^{\mathrm{UPF}}(A) = 0\,] - 2^{-L} \;=\; \mathsf{Adv}_F^{\mathrm{UPF}}(A) - 2^{-L}\end{aligned}$$

Given that A was any arbitrary forger, the claimed relation follows. ∎

We say that Proposition 5 represents a tightening of the security bounds given in Proposition 4 since, from Theorems 1 and 2, we know that $\mathsf{Adv}_F^{\mathrm{IUF}}(t', q)$ is at most $2 \cdot \mathsf{Adv}_F^{\mathrm{PRF}}(t', q)$ and can be as small as $\frac{1}{q} \cdot \mathsf{Adv}_F^{\mathrm{PRF}}(t', q)$.

4.2 The Case of Symmetric Encryption Schemes

In the following discussion, we use the standard syntax and notion of security for encryption schemes given by Bellare et al [3], which is an adaptation of one given by Goldwasser and Micali [9]. In the indistinguishability of encryptions under chosen-plaintext attack (IND) notion, the adversary A is imagined to run in two phases. In the find phase, given adaptive access to an encryption oracle, A produces a pair of equal-length messages x_0, x_1, along with some state information s. In the guess phase, given the encryption y of one of the messages and s, it must identify which of the two messages goes with y.

Definition 4. [Symmetric encryption security: IND] *Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an encryption scheme. For an adversary A and $b = 0, 1$ define the experiment:*

Experiment $\mathsf{Exp}_\Pi^{\mathrm{IND}}(A, b)$
$a \leftarrow \mathcal{K}$; $(x_0, x_1, s) \leftarrow A^{\mathcal{E}_a}(\mathsf{find})$; $y \leftarrow \mathcal{E}_a(x_b)$; $d \leftarrow A^{\mathcal{E}_a}(\mathsf{guess}, y, s)$; Return d.

It is mandated that $|x_0| = |x_1|$ above. Define the advantage of A and the advantage function of Π, respectfully, as follows. For any integers $t, q, \mu \geq 0$,

$$\mathsf{Adv}^{\mathrm{IND}}_{\Pi}(A) = \Pr[\,\mathsf{Exp}^{\mathrm{IND}}_{\Pi}(A,0) = 0\,] - \Pr[\,\mathsf{Exp}^{\mathrm{IND}}_{\Pi}(A,1) = 0\,]$$
$$\mathsf{Adv}^{\mathrm{IND}}_{\Pi}(t,q,\mu) = \max_{A} \{\mathsf{Adv}^{\mathrm{IND}}_{\Pi}(A)\,\}$$

where the maximum is over all A with time complexity t, making $\leq q$ oracle queries which total $\leq \mu$ bits. ∎

We analyze the counter mode of encryption based on a finite PRF. In practice, the finite PRF may be instantiated by a block cipher. For a finite PRF F, the counter mode $\mathrm{CTR}(F) = (\mathcal{E}\text{-CTR}, \mathcal{D}\text{-CTR}, \mathcal{K}\text{-CTR})$ can be described as follows. The key generation algorithm $\mathcal{K}$-CTR outputs a random key a for the underlying PRF family F, thereby specifying a function $f = F_a$ of l-bits to L-bits. The sender maintains a l bit counter ctr that is initially -1 and is incremented after each encryption by the number of blocks encrypted. The message x to be encrypted is regarded as a sequence of L-bit blocks (padding is done first, if necessary), $x = x_1 \cdots x_n$. We define $\mathcal{E}\text{-CTR}_a(x, ctr) = \mathcal{E}\text{-CTR}^{F_a}(x, ctr)$ and $\mathcal{D}\text{-CTR}_a(z) = \mathcal{D}\text{-CTR}^{F_a}(z)$, where:

Algorithm $\mathcal{E}\text{-CTR}^f(x, ctr)$	Algorithm $\mathcal{D}\text{-CTR}^f(z)$
for $i = 1, \ldots, n$ do	Parse z as $ctr\,,\ y_1 \cdots y_n$
$\quad y_i = f(ctr + i) \oplus x_i$	for $i = 1, \ldots, n$ do
$ctr \leftarrow ctr + n$	$\quad x_i = f(ctr + i) \oplus y_i$
return $(ctr\,,\ y_1 y_2 \cdots y_n)$	return $x = x_1 \cdots x_n$

We show that $\mathrm{CTR}(F)$ is secure in the IND sense if F is secure in the IUF sense. As with our previous example, the reduction achieves the same concrete security bounds as those possible using the standard notion of PRF families.

Theorem 8. [Security of CTR using an IUF function family] *For any function family F and integers $t, q \geq 1$ and $L \leq \mu \leq L2^l$,*

$$\mathsf{Adv}^{\mathrm{IND}}_{\mathrm{CTR}(F)}(t,q,\mu) \leq 2 \cdot \mathsf{Adv}^{\mathrm{IUF}}_{F}(t',q')$$

where $t' = t + O(\frac{\mu}{L}(l+L))$ and $q' = \frac{\mu}{L}$.

Proof. We want to show that if $\mathrm{CTR}(F)$ is not secure in the IND sense, then it must be the case that F is not secure in the IUF sense. Let A be an adversary attacking the $\mathrm{CTR}(F)$, running in time at most t and making at most q oracle queries, these totalling at most μ bits, in the experiment $\mathsf{Exp}^{\mathrm{IND}}_{\mathrm{CTR}(F)}(A)$. We construct a distinguisher A', making at most q' queries and running in time at most t', using the adversary A as a subroutine.

Let $\mathcal{O}_f$ be A''s oracle. $A'^{\mathcal{O}_f}$ will run A using $\mathcal{O}_f$ to provide an appropriate simulation of A's encryption oracle. We assume, for the sake of simplicity of the exposition, that the two messages A outputs at the end of its first phase are

exactly L bits in length (i.e. of the size of one block). In the following, $\mu_G < \mu$, is the amount of ciphertext A needs to see in its guess phase.

Algorithm $A'^{\mathcal{O}_f}$

(1) Initialize counter: $ctr \leftarrow -1$.
(2) Run $A(\mathsf{find})$, answering any query u with $\mathcal{E}\text{-CTR}^{\mathcal{O}_f}(u)$.
(3) Let $(x_0, x_1, s) \leftarrow A(\mathsf{find})$.
(4) Let the current value of the counter be ctr_0.
(5) Compute $\mathcal{F} = \{\mathcal{O}_f(ctr_0 + i) : 1 \leq i \leq \frac{\mu_G}{L}\}$.
(6) Let $s' = (s, x_0, x_1, ctr_0, \mathcal{F})$.
(7) Output (ctr_0, s') and receive y as the challenge.
(8) Let $d \leftarrow \{0, 1\}$.
(9) Run $A(\mathsf{guess}, y \oplus x_d, s)$, answering any query u, using $\mathcal{F}$, with $\mathcal{E}\text{-CTR}(u)$.
(10) Let $d' \leftarrow A(\mathsf{guess}, y \oplus x_d, s)$.
(11) If $d = d'$ then output 0, else output 1.

In the reduction above, A' maintains the counter ctr, incrementing it appropriately. It is important here that A' can implement $\mathcal{E}\text{-CTR}^f(\cdot, ctr)$ given an oracle for f. At the end of the find phase queries of A, it picks the current value of counter ctr_0 to be the output of its own find phase, along with the state information. A slight problem that comes up here is that A' does not have access to $\mathcal{O}_f$ in its guess phase but it will still need to provide a simulation of the encryption oracle during A's guess phase queries. We get around this by having A' pre-compute the value of $\mathcal{O}_f$ on as many points as necessary, starting from $ctr_0 + 1$, to answer all of A's guess phase encryption oracle queries. These pre-computed values are in the set $\mathcal{F}$ which is passed to A''s guess phase via state information s. Note that it is important that A' did not query $\mathcal{O}_f$ with ctr_0, since otherwise it could not output ctr_0 as the point on which it gets its challenge. The counter mode guarantees that, as long as fewer than $\frac{\mu}{L}$ queries are made (i.e the counter does not loop around), the function will always be invoked on a new point.
The total number of oracle queries made by A' is at most $\frac{\mu}{L}$, which by assumption is q'. Given this, one can check that the running time of A' is as claimed. The advantage of A' is given by,

$$\begin{aligned}
\mathsf{Adv}^{\mathrm{IUF}}_{F}(A') &= \Pr[\,\mathsf{Exp}^{\mathrm{IUF}}_{F}(A', 0) = 0\,] - \Pr[\,\mathsf{Exp}^{\mathrm{IUF}}_{F}(A', 1) = 0\,] \\
&= \Pr[\,\mathsf{Exp}^{\mathrm{IND}}_{\mathrm{CTR}(F)}(A, 0) = 0\,] + \Pr[\,\mathsf{Exp}^{\mathrm{IND}}_{\mathrm{CTR}(F)}(A, 1) = 1\,] - \frac{1}{2} \\
&= \frac{1}{2}(1 + \mathsf{Adv}^{\mathrm{IND}}_{\mathrm{CTR}(F)}(A)) - \frac{1}{2} \;=\; \frac{1}{2} \cdot \mathsf{Adv}^{\mathrm{IND}}_{\mathrm{CTR}(F)}(A)
\end{aligned}$$

Given that A is an arbitrary adversary, the claimed relation follows. ∎

5 Discussion

We stress that the benefits of tighter security analyses, such as those we have presented here, are real. For example, using the standard notion of a PRF, the security of a protocol may appear to be marginal, prompting the use of a larger security parameter. However, using a tighter characterization, such as IUF, the security might have been determined to be adequate.

In criticism to our approach to getting tighter bounds for MACs and symmetric encryption schemes, one may suggest that we are looking at the wrong notions of security for these protocols. Indeed, there are alternate notions for which our gains would disappear. However, the notions of security we consider for both MACs and symmetric encryption are, in practice, the notions which are most widely used.

Future Directions. Unlike the case with the counter mode of encryption, in our first example we view the entire MAC as being the primitive, when in fact it too may be built on a PRF (for example, the CBC-MAC based on a block cipher). While it seems unlikely that we can achieve a tighter security analysis for the CBC-MAC scheme using the same approach, it may be possible for other message authentication schemes. Then there are other schemes, besides those for message authentication and symmetric encryption, to which our techniques could be applied. For example, it may be possible to improve the security bounds of variable-length input pseudorandom functions (VI-PRFs) [2] and variable-input-length ciphers [6].

Using similar techniques as above, we can also get tighter bounds for PRP-based protocols. In a sense, this is more interesting, given that PRP families provide a more natural model for block ciphers [5]. Viewing a block cipher as a PRP family rather than a PRF family itself can lead to tighter security bounds. However, our examples were motivated by the fact that analysis of a block-cipher-based scheme is, as far as possible, done modeling the block cipher as a PRF. This is because the analysis using PRFs is usually significantly simpler.

We remark that it seems somewhat significant that, in the indistinguishability of points characterization, there is a difference between function and permutation families. This seems to be the first such distinction, as far as we know, when asymptotic measures are used. It may be interesting to investigate further the impact of injectivity upon pseudorandomness.

Acknowledgements

We are grateful to Mihir Bellare for his advice and assistance with this work. We also thank the Asiacrypt 2000 program committee for their helpful comments.

This work was completed while the first author was a student at the University of California at San Diego, USA. Both authors were supported in part by Mihir Bellare's 1996 Packard Foundation Fellowship in Science and Engineering and NSF CAREER Award CCR-9624439.

References

1. M. Bellare, R. Canetti and H. Krawczyk, "Keying hash functions for message authentication," *Advances in Cryptology - Crypto '96*, LNCS Vol. 1109, N. Koblitz ed., Springer-Verlag, 1996.
2. 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.
3. M. Bellare, A. Desai, E. Jokipii and P. Rogaway, "A concrete security treatment of symmetric encryption: Analysis of the DES modes of operation," *Proceedings of the 38th Symposium on Foundations of Computer Science*, IEEE, 1997.
4. M. Bellare, R. Guérin and P. Rogaway, "XOR MACs: New methods for message authentication using finite pseudorandom functions," *Advances in Cryptology - Crypto '95*, LNCS Vol. 963, D. Coppersmith ed., Springer-Verlag, 1995.
5. M. Bellare, J. Kilian and P. Rogaway, "The security of the cipher block chaining message authentication code," *Advances in Cryptology - Crypto '94*, LNCS Vol. 839, Y. Desmedt ed., Springer-Verlag, 1994.
6. M. Bellare and P. Rogaway, "On the construction of variable-input-length ciphers," *Proceedings of the Sixth Workshop on Fast Software Encryption*, L. Knudsen ed., 1999.
7. A. Desai and S. Miner, "Concrete security characterizations of PRFs and PRPs: Reductions and applications," Full version of this paper, available via: `http://www-cse.ucsd.edu/users/sminer/`.
8. O. Goldreich, S. Goldwasser and S. Micali, How to construct random functions. *Journal of the ACM*, 33(4): 792-807, 1986.
9. S. Goldwasser and S. Micali, "Probabilistic encryption," *Journal of Computer and System Science*, Vol. 28, pp. 270-299, 1984.
10. S. Goldwasser, S. Micali and R. Rivest, "A digital signature signature scheme secure against adaptive chosen-message attacks," *SIAM J. of Computing*, 17(2): 281-308, April 1988.
11. M. Luby and C. Rackoff, "How to construct pseudorandom permutations from pseudorandom functions," *SIAM J. Computing*, 17(2), April 1988.
12. M. Naor and O. Reingold, "From unpredictability to indistinguishability: A simple construction of PRFs from MACs," *Advances in Cryptology - Crypto '98*, LNCS Vol. 1462, H. Krawczyk ed., Springer-Verlag, 1998.