

Format-Preserving Encryption

Mihir Bellare¹, Thomas Ristenpart¹, Phillip Rogaway², and Till Stegers²

¹ Dept. of Computer Science & Engineering, UC San Diego, La Jolla, CA 92093, USA

² Dept. of Computer Science, UC Davis, Davis, CA 95616, USA

Abstract. Format-preserving encryption (FPE) encrypts a plaintext of some specified format into a ciphertext of identical format—for example, encrypting a valid credit-card number into a valid credit-card number. The problem has been known for some time, but it has lacked a fully general and rigorous treatment. We provide one, starting off by formally defining FPE and security goals for it. We investigate the natural approach for achieving FPE on complex domains, the “rank-then-encipher” approach, and explore what it can and cannot do. We describe two flavors of unbalanced Feistel networks that can be used for achieving FPE, and we prove new security results for each. We revisit the cycle-walking approach for enciphering on a non-sparse subset of an encipherable domain, showing that the timing information that may be divulged by cycle walking is not a damaging thing to leak.

1 Introduction

BACKGROUND. During the last few years, *format-preserving encryption* (FPE) has emerged as a useful tool in applied cryptography. The goal is this: under the control of a symmetric key K , deterministically encrypt a plaintext X into a ciphertext Y that has the same *format* as X . Examples include encryption of US social security numbers (SSNs), credit card numbers (CCNs) of a given length, 512-byte disk sectors, postal addresses of some particular country, and jpeg files of some given length. In our formalization of FPE, the format of a plaintext X will be a name N describing a finite set \mathcal{X}_N over which the encryption function induces a permutation. For example, with SSNs this is the set of all nine-decimal-digit numbers.

The FPE goal is actually quite old. For one thing, a blockcipher itself can be seen as one kind of FPE: each N -bit string, where N is the block size, is mapped to some N -bit string. But what makes FPE an interesting and powerful idea is that the notion reaches far beyond blockciphers, which normally encipher strings of some one, convenient length.

SOME PRIOR WORK. In FIPS 74 (1981) [29], a DES-based approach is described to encipher strings over some fixed alphabet, say the decimal digits $D = \{0, 1, \dots, 9\}$. Each plaintext $X \in D^N$ would be mapped to a ciphertext $Y \in D^N$. Here each plaintext $X \in D^*$ has a unique format $N = |X|$ and we must encipher X relative to the set $\mathcal{X}_N = D^N$.

Brightwell and Smith (1997) [6] considered a more general scenario, identifying what they termed *datatype-preserving encryption*. They wanted to encrypt database entries of some particular datatype without disrupting that datatype. A field containing an SSN (a nine-digit decimal string) should get mapped to another SSN. The authors colorfully explain the difficulty of doing this, saying that, with conventional encryption schemes, a “Ciphertext . . . bears roughly the same resemblance to plaintext . . . as a hamburger does to a T-bone steak. A social security number, encrypted using the DES encryption algorithm, not only does not resemble a social security number but will likely not contain any numbers at all” [6, p. 142]. The authors provide a proposed solution, though, as with FIPS 74, definitions or proofs for it are not likely or claimed.

Black and Rogaway [4] provided a provable-security investigation of a special case of FPE, asking how to make a cipher $E: \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{X}$ with an arbitrary domain \mathcal{X} . Their solutions focused on $\mathcal{X} = \mathbb{Z}_N$, the integers $\{0, 1, \dots, N - 1\}$. The authors offer no general definition for FPE but they

clearly intend that ciphers with domains of \mathbb{Z}_N be used to construct schemes with other domains, like the set of valid CCNs of a given length.

The term *format-preserving encryption* is due to Terence Spies, Voltage Security’s CTO [42]. Voltage, Semtek and other companies have been active in productizing FPE and explaining its utility [41]. FPE can enable a simpler migration path when encryption is added to legacy systems and databases, as required, for example, by the payment-card industry’s data security standard (PCI DSS) [36]. Use of FPE enables upgrading database security in a way transparent to many applications, and minimally invasive to others. Spies has gone on to submit to NIST a proposed mechanism, FFSEM, that combines cycle walking and an AES-based balanced Feistel network [43].

SYNTAX. The current paper aims to help cryptographic theory “catch up” with cryptographic practice in this FPE domain. We initiate a general treatment of the problem, doing this within the provable-security tradition of modern cryptography.

We begin with a very general definition for FPE. Unlike a conventional cipher, an FPE scheme has associated to it a *collection* of domains, $\{\mathcal{X}_N\}_{N \in \mathcal{N}}$. We call each \mathcal{X}_N a *slice* (the overall domain is their union, $\mathcal{X} = \bigcup_N \mathcal{X}_N$). The set \mathcal{N} is the *format space*. For every key K , format N , and tweak T the FPE scheme E names a permutation $E_K^{N,T}$ on \mathcal{X}_N . We are careful to make FPEs tweakable [22] because, in this context, use of a tweak can significantly enhance security.

Returning to the CCN example, suppose we want to do FPE of CCNs with a zero Luhn-checksum [20]. Let’s assume that the map should be length-preserving and that the possible lengths range from 12 to 19 decimal digits. Then we could let $\mathcal{N} = \{12, \dots, 19\}$ and let \mathcal{X}_N be the set of all N -digit numbers X such that $\text{LuhnOK}(X)$ is true. Now an FPE scheme E with slices $\{\mathcal{X}_N\}_{N \in \mathcal{N}}$ does the job. You encrypt CCN X with key K and tweak T by letting $Y = E_K^{N,T}(X)$, where $N = \text{len}(X)$.

SECURITY NOTIONS. We define multiple notions of security for FPE schemes. Our strongest adapts the traditional PRP notion to capture the idea that FPE is a good approximation for a family of uniform permutations on the slices. Our weaker notions are denoted SPI, MP, and MR. SPI (single-point indistinguishability) is a variant of the PRP notion in which there is only a single challenge point. MP (message privacy) lifts semantic security to the FPE setting by adapting earlier notions of deterministic encryption [2, 5]. MR (message recovery) formalizes an adversary’s inability to recover a challenge message, in its entirety, from the message’s ciphertext. All of these notions can be made with respect to an adaptive or nonadaptive adversary, and can also be strengthened to allow chosen-ciphertext attacks (for PRP, this would result in what is called a *strong* PRP).

Why bother with SPI, MP, and MR when they are implied by PRP? SPI is useful because it is easy to work with and implies MP and MR with a tight bound. MP and MR are interesting because they, even in their nonadaptive form, are what an application will most typically need. An attack against the PRP notion may be no threat in practice, and achieving good PRP security may be overkill. Good concrete security bounds become particularly a focus when slices are small: a bound permitting $q \approx 2^{n/4}$ queries provides limited assurance when $n = 20$ bits.

CONSTRUCTIONS. We next investigate the construction of FPE schemes. Suppose we wish to build an FPE scheme \mathcal{E} with a complex *specification*—the slices $\{\mathcal{X}_N\}$ on which it should encipher. A natural approach is to arbitrarily number the points in each \mathcal{X}_N , say $\mathcal{X}_N = \{X_0, X_1, \dots, X_{n-1}\}$ where $n = |\mathcal{X}_N|$. Then, to encipher $X \in \mathcal{X}_N$, find its index i in the enumeration, encipher i to j in \mathbb{Z}_n , and then return X_j as the encryption of X . We call this strategy the *rank-then-encipher* approach. It’s the obvious, one could say folklore, approach. To implement it, we need an *integer* FPE that can encipher on \mathbb{Z}_n for any needed n , as well a *ranking function*, rank , that maps each (N, X) with $X \in \mathcal{X}_N$ to a point in \mathbb{Z}_n with $\text{rank}(N, \cdot): \mathcal{X}_N \rightarrow \mathbb{Z}_n$ a bijection for all $N \in \mathcal{N}$.

We will show how to build ranking functions for any FPE problem whose domain is a regular language (the slices being strings of each possible length). This includes many practical problems. This can be extended to domains that are context-free languages having unambiguous grammars.

Our starting point for building integer FPEs is the construction of Black and Rogaway [4], which combines a generalization of an unbalanced Feistel network (the left and right hand side are numbers in \mathbb{Z}_a and \mathbb{Z}_b rather than strings) and a technique the authors call *cycle walking*, a method apparently going back to the rotor machines of the early 1900's [39]. We extend their work to handle multiple slices with the same key, and to incorporate tweaks.

The type of unbalanced Feistel network that was extended in [4] is the type due to Lucks [24]. It is not the only kind of unbalanced Feistel network. An equally natural possibility is the unbalanced Feistel design of Schneier and Kelsey [38]. Extended to \mathbb{Z}_N where $N = ab$, we call this a *type-1* Feistel, as opposed to the *type-2* unbalanced Feistel network of [4, 24]. Our FPE schemes FE1 and FE2, based on type-1 and type-2 unbalanced Feistel networks, comprise a flexible, efficient, and customizable means for enciphering domains \mathbb{Z}_N where $N = ab$ is the product of integers greater than one. Its round function can be based, for example, on AES. Combining FE1 or FE2 with the rank-then-encipher approach lets one achieve FPE in a wide variety of contexts.

SECURITY. Ideally, we would like to prove good bounds on the strong-PRP security for FE1 and FE2, assuming the round function to be a good PRF. But we run into a limitation, namely that the proven strength of Feistel ciphers [4, 23, 26, 28, 30–35], in terms of quality of bounds, falls short of what is wanted, and what appears to be the actual strength of the techniques. We address this in a couple of ways.

First, proofs have always targeted PRP. Instead, we target MP and MR, thereby getting better bounds more easily. We prove that FE2 with only *three* rounds hides all partial information with respect to a nonadaptive chosen-plaintext attack: one achieves nonadaptive SPI, MP, and MR security with reasonable bounds. Even then, we feel that being guided purely by what can be proved would lead to an overly quite pessimistic security estimate. The most realistic picture may be obtained by also assessing resistance to attacks. We consider known attacks and discuss their implications for our parameter choices (principally the number of rounds). We also provide a novel attack against (heavily) unbalanced type-2 Feistel networks, one that achieves message recovery with success probability exponentially small in the number of rounds. The attack is damaging if the number of rounds is too small.

Finally, reaching beyond PRP/SPI/MP/MR security, we consider a particular kind of side-channel attack. The use of cycle-walking in the rank-then-encipher approach raises the fear of timing attacks: might the number of times one has to apply the underlying cipher leak adversarially valuable information? We prove that cycle-walking will *not*, on its own, give rise to timing attacks. This is because the correct distribution on the number of iterations of the cipher on any input can be computed by a simulator that does not attend to the inputs.

THE FUTURE. We expect FPE to be increasingly deployed. The complex systems that process financial transactions impose a powerful legacy constraint. Using classical blockcipher-based modes would require far larger changes to these systems, which is costly and error-prone. FPE can be realized by simple, AES-based modes of operation, avoiding the need to design and review any fundamentally new primitive. Besides the enciphering of database fields, FPE may prove useful in networking applications, allowing datagrams to have their fields protected without changing their format. What one might lose in security when employing a *deterministic* encryption scheme can often be erased by sensibly tweaking the FPE scheme [22]. Moreover, such loss of security may be entirely overshadowed by the reduced need for random bits and disruption in infrastructure, protocols, and code.

2 FPE Syntax

SYNTAX. A scheme for *format-preserving encryption* (FPE) is a function $E: \mathcal{K} \times \mathcal{N} \times \mathcal{T} \times \mathcal{X} \rightarrow \mathcal{X} \cup \{\perp\}$ where the sets \mathcal{K} , \mathcal{N} , \mathcal{T} , and \mathcal{X} are called the *key space*, *format space*, *tweak space*, and *domain*, respectively. All of these sets are nonempty and $\perp \notin \mathcal{X}$. We write $E_K^{N,T}(X) = E(K, N, T, X)$ for the encryption of X with respect to key K , format N , and tweak T . We require that whether or not $E_K^{N,T}(X) = \perp$ depends only on N, X and not on K, T , and let

$$\mathcal{X}_N = \{X \in \mathcal{X} : E_K^{N,T}(X) \in \mathcal{X} \text{ for all } (K, T) \in \mathcal{K} \times \mathcal{T}\}$$

be the N -indexed *slice* of the domain. We demand that a point $X \in \mathcal{X}$ live in at least one slice, $X \in \mathcal{X}_N$ for some N (if X is in no slice it should not be included in E 's domain). We demand that there be finitely many points in each slice, meaning \mathcal{X}_N is finite for all $N \in \mathcal{N}$. We require that $E_K^{N,T}(\cdot)$ be a permutation on \mathcal{X}_N for any $(K, T) \in \mathcal{K} \times \mathcal{T}$. Its inverse $D: \mathcal{K} \times \mathcal{N} \times \mathcal{T} \times \mathcal{X} \rightarrow \mathcal{X} \cup \{\perp\}$ is defined by $D_K^{N,T}(Y) = D(K, N, T, Y) = X$ if $E_K^{N,T}(X) = Y$. In summary, an FPE enciphers the points within each of the (finite) slices that collectively comprise its domain.

A practical FPE scheme $E: \mathcal{K} \times \mathcal{N} \times \mathcal{T} \times \mathcal{X} \rightarrow \mathcal{X} \cup \{\perp\}$ must be realizable by efficient algorithms: an algorithm E to encrypt, an algorithm D to decrypt, and an algorithm to sample uniformly from the key space \mathcal{K} . Thus \mathcal{K} , \mathcal{N} , \mathcal{T} , and \mathcal{X} should consist of strings or points easily encoded as strings, and E and D should return \perp when presented a point outside of $\mathcal{K} \times \mathcal{N} \times \mathcal{T} \times \mathcal{X}$. We will not draw any distinction between an integer element of \mathcal{X} , say, and a string that encodes such a point.

THE FORMAT OF A POINT. Let $E: \mathcal{K} \times \mathcal{N} \times \mathcal{T} \times \mathcal{X} \rightarrow \mathcal{X} \cup \{\perp\}$ be an FPE scheme. Then we can speak of $X \in \mathcal{X}$ as having *format* N if $X \in \mathcal{X}_N$. One could associate to E a *format function* $\varphi: \mathcal{X} \rightarrow \mathcal{P}(\mathcal{N}) \setminus \{\emptyset\}$ that maps each $X \in \mathcal{X}$ to its possible formats; formally, $\varphi(X) = \{N \in \mathcal{N} : X \in \mathcal{X}_N\}$.

Note that, under our definitions, a point may have multiple formats. But often this will not be the case: each $X \in \mathcal{X}$ will belong to exactly one \mathcal{X}_N . In that case we can regard the format function as mapping $\varphi: \mathcal{X} \rightarrow \mathcal{N}$ and interpret $\varphi(X)$ as *the* format of X . FPE is somewhat simpler to understand for such *unique-format* FPEs: you can examine an X and know from it the slice $\mathcal{X}_{\varphi(X)}$ on which you mean to encipher it. For a unique format FPE one can write $E_K^T(X)$ rather than $E_K^{N,T}(X)$ since N is determined by X .

SPECIFICATIONS. An FPE problem, as needed by some application, will specify the desired collection of slices, $\{\mathcal{X}_N\}_{N \in \mathcal{N}}$. It will also specify the desired tweak space \mathcal{T} . Typically it is easy to support whatever tweak space one wants, but it may be quite hard to support a given collection of slices $\{\mathcal{X}_N\}_{N \in \mathcal{N}}$ (indeed it may be hard to accommodate a single slice, depending on what it is). We therefore call the collection of slices $\{\mathcal{X}_N\}_{N \in \mathcal{N}}$ the *specification* for an FPE scheme. We will write $\mathcal{X} = \{\mathcal{X}_N\}_{N \in \mathcal{N}}$ for a specification, only slightly abusing notation because the domain \mathcal{X} is the union of slices in $\{\mathcal{X}_N\}_{N \in \mathcal{N}}$. The question confronting the cryptographer is *how to design an FPE scheme with a given specification*. We now provide some example possibilities.

EXAMPLES. **(1)** AES-128 can be regarded as an FPE with a single slice, $\{0, 1\}^{128}$. The key space is $\mathcal{K} = \{0, 1\}^{128}$ and the format space and tweak space are trivial (have size one). **(2)** To encipher 16-digit decimal numbers, take $\mathcal{X} = \{0, 1, \dots, 9\}^{16}$ and just the one slice. **(3)** To encipher 512-byte disk sectors using an 8-byte sector index as the tweak, let $\mathcal{X} = \{0, 1\}^{4096}$, $\mathcal{T} = \{0, 1\}^{64}$, and just the one slice. **(4)** Suppose you want to encipher CCNs of 12–19 digits with a proper Luhn checksum, the ciphertext having the same length as the plaintext. Then the specification could be $\mathcal{X} = \{\mathcal{X}_N\}_{N \in \mathcal{N}}$ where $\mathcal{N} = \{12, 13, \dots, 18, 19\}$ and \mathcal{X}_N is the set of all strings $X \in \{0, 1, \dots, 9\}^N$ satisfying the predicate $\text{LuhnOK}(X)$. Here $|\mathcal{X}_N| = 10^{N-1}$. **(5)** One nice FPE has slices that are $\{0, 1\}^N$ for each $N \geq 0$. It allows length-preserving encryption of any binary string. **(6)** One can FPE rather unusual spaces. For example, slice \mathcal{X}_N could encode all N -vertex graphs. Or \mathcal{X}_N could

be all valid C-programs on N bytes. Designing an efficient FPE with this specification might be impossible. All of the examples just given are unique-format FPEs. The following example is not.

INTEGER FPEs. The specification for a particularly handy kind of FPE is the following. The slices are $\mathcal{X}_N = \mathbb{Z}_N$, for $N \in \mathcal{N} \subseteq \mathbb{N}$. This allows enciphering natural numbers with respect to any permitted modulus N . Assuming the tweak space is similarly rich, say $\mathcal{T} = \{0, 1\}^*$, we call such scheme an *integer FPE*. When used within the rank-then-encipher paradigm, integer FPEs enable the construction of FPEs with quite complex specifications.

3 FPE Security Notions

GAMES. Our definitions and proofs use *code-based games* [1], so we first review that material. A game has an **Initialize** procedure, an optional **Finalize** procedure, and any number of additional procedures. A game G is executed with an adversary \mathcal{A} as follows. First, **Initialize** executes, possibly returning an output s , and then $\mathcal{A}(\text{run}, s)$ is run ($s = \varepsilon$ if **Initialize** returns no string). As \mathcal{A} executes it may call any procedure G (but not **Initialize** or **Finalize**) provided by G . If there is no **Finalize** procedure, the output of \mathcal{A} is the output of the game. If the game does specify a **Finalize**, then, when \mathcal{A} terminates, \mathcal{A} 's output is **Finalize**'s input and the game's output is that of **Finalize**. Game procedures may call $\mathcal{A}(\text{identifier}, [x])$, which invokes an instance of the caller with distinct coins for each distinct **identifier**. Conceptually, then, each **identifier** thus names a separate adversarial algorithm. State is not shared among them. Let $G^{\mathcal{A}} \Rightarrow y$ denote the event that the game outputs y . We write $S \stackrel{\cup}{\leftarrow} x$ as shorthand for $S \leftarrow S \cup \{x\}$. Later we write $c \stackrel{\pm}{\leftarrow} d$ for $c \leftarrow c + d$.

Boolean variables, including *bad*, are silently initialized to **FALSE**, set variables to \emptyset , integer variables to 0. Games G and H are said to be *identical-until-bad* if their code differs only in the sequel of statements that first set *bad* to true. We say that “ $G^{\mathcal{A}}$ sets *bad*” for the event that game G , when executed with adversary \mathcal{A} , sets *bad* to true. If G, H are *identical-until-bad* and \mathcal{A} is an adversary then $\Pr [G^{\mathcal{A}} \text{ sets } bad] = \Pr [H^{\mathcal{A}} \text{ sets } bad]$. It is also standard (“the fundamental lemma”) that if G, H are *identical-until-bad* then $\Pr [G^{\mathcal{A}} \Rightarrow y] - \Pr [H^{\mathcal{A}} \Rightarrow y] \leq \Pr [G^{\mathcal{A}} \text{ sets } bad]$.

SECURITY NOTIONS. We will extend the standard PRP notion to our setting, but we will also describe notions weaker than it, because they can be achieved with better proven concrete security for the same efficiency and, at the same time, they suffice for typical applications. Coming at it from the latter perspective, the most basic and often sufficient requirement is security against message recovery (MR), under either an adaptive or nonadaptive attack. We define this as well as a stronger notion of message privacy (MP) that requires that partial information about the message is not leaked by the ciphertext. We also consider a weakening of the PRP notion that we call SPI. The reason for considering this notion is that it is simpler than MP and MR to work with yet implies them; at the same time, it can be achieved with better concrete security bounds than we currently know how to get for the ordinary PRP notion.

In the following let $E: \mathcal{K} \times \mathcal{N} \times \mathcal{T} \times \mathcal{X} \rightarrow \mathcal{X} \cup \{\perp\}$ be an FPE scheme. We consider the games in Figure 1. It is assumed that any query of the form (N, T, X) satisfies $N \in \mathcal{N}$, $X \in \mathcal{X}_N$, and $T \in \mathcal{T}$.

PRP security. The standard notion of PRP security is extended to FPE schemes via game PRP_E and the corresponding adversary advantage is

$$\text{Adv}_E^{\text{PRP}}(\mathcal{A}) = 2 \cdot \Pr [\text{PRP}_E^{\mathcal{A}} \Rightarrow \text{true}] - 1 .$$

In the game $\text{Perm}(\mathcal{X}_N)$ is the set of all permutations on \mathcal{X}_N .

SPI security. Single-point indistinguishability (SPI) requires that the adversary be unable to distinguish between the encryption of a single chosen message or a random range point, even when

<p>Initialize // Game PRP_E $b \xleftarrow{\\$} \{0, 1\}; K \xleftarrow{\\$} \mathcal{K}$ for $(N, T) \in \mathcal{N} \times \mathcal{T}$ do $\pi_{N,T} \xleftarrow{\\$} \text{Perm}(\mathcal{X}_N)$</p> <p>Enc($N, T, X$) if $b = 1$ then ret $E_K^{N,T}(X)$ if $b = 0$ then ret $\pi_{N,T}(X)$</p> <p>Finalize(b') ret ($b = b'$)</p>	<p>Initialize // Game SPI_E $b \xleftarrow{\\$} \{0, 1\}; K \xleftarrow{\\$} \mathcal{K}$</p> <p>Enc(N, T, X) if $(N, T, X) \in S$ then ret \perp $S \xleftarrow{\cup} (N, T, X)$ ret $E_K^{N,T}(X)$</p> <p>Test(N^*, T^*, X^*) if $(N^*, T^*, X^*) \in S$ then ret \perp $S \xleftarrow{\cup} (N^*, T^*, X^*)$ if $b = 1$ then $Y^* \leftarrow E_K^{N^*,T^*}(X^*)$ else $Y^* \xleftarrow{\\$} \mathcal{X}_{N^*}$ ret Y^*</p> <p>Finalize(b') ret ($b = b'$)</p>	<p>Initialize // Game MP_E $K \xleftarrow{\\$} \mathcal{K}$ $(N^*, T^*, X^*) \xleftarrow{\\$} \mathcal{A}(\text{dist})$ $Y^* \leftarrow E_K^{N^*,T^*}(X^*)$ ret (N^*, T^*)</p> <p>Enc(N, T, X) ret $E_K^{N,T}(X)$</p> <p>Eq(X) ret ($X = X^*$)</p> <p>Test ret Y^*</p> <p>Finalize(Z) ret ($Z = \mathcal{A}(\text{func}, X^*)$)</p>	<p>Initialize // Game MR_E $K \xleftarrow{\\$} \mathcal{K}$ $(N^*, T^*, X^*) \xleftarrow{\\$} \mathcal{A}(\text{dist})$ $Y^* \leftarrow E_K^{N^*,T^*}(X^*)$ ret (N^*, T^*)</p> <p>Enc(N, T, X) ret $E_K^{N,T}(X)$</p> <p>Eq(X) ret ($X = X^*$)</p> <p>Test ret Y^*</p> <p>Finalize(X) ret ($X = X^*$)</p>
--	--	---	---

Fig. 1. Games used for defining FPE security notions SPRP, PRP, SPI, MP, and MR. Procedure \mathcal{A} , invoked by games MP and MR, denotes the caller of the game.

given adaptive access to a true encryption oracle. The formalization is based on game SPI_E. An adversary \mathcal{A} is allowed to make only a single **Test** query, and this must be its first oracle query. Its associated advantage is

$$\mathbf{Adv}_E^{\text{spi}}(\mathcal{A}) = 2 \cdot \Pr[\text{SPI}_E^{\mathcal{A}} \Rightarrow \text{true}] - 1.$$

Let game SPI1_E (resp. SPI0_E) be the same as SPI_E except that b is set to 1 (resp. 0) in **Initialize**. A standard argument gives that

$$\mathbf{Adv}_E^{\text{spi}}(\mathcal{A}) = \Pr[\text{SPI1}_E^{\mathcal{A}} \Rightarrow \text{true}] - \Pr[\text{SPI0}_E^{\mathcal{A}} \Rightarrow \text{FALSE}].$$

The SPI notion is closely related to (and inspired by) a definition originally from [13], variants of which were also considered in [10, 27]. It is easy to see that PRP security implies SPI security, but there is an additive loss of q/M in the advantage bound, where q is the number of queries by the adversary and M is the minimum size of \mathcal{X}_N over all $N \in \mathcal{N}$. This is perhaps unfortunate, but SPI is only used as a tool anyway. A hybrid argument following [10, 13] shows that SPI security likewise implies PRP security. Here, $\mathbf{Adv}_E^{\text{spi}}(\mathcal{A}) \leq q \cdot \mathbf{Adv}_E^{\text{prp}}(B) + q^2/M$ where q is the number of **Enc** queries of starting prp adversary \mathcal{A} , and constructed spi adversary B makes $q - 1$ **Enc** queries.

Message recovery. An FPE scheme secure against message recovery is one for which an adversary is unable to recover plaintexts from ciphertexts, even given an encryption oracle and a favorable distribution of plaintexts, formats, and tweaks. If the encryption were randomized we would require that the target ciphertext Y^* and encryption oracle E_K be of *no* use in recovering the plaintext, but this is too much to ask for with a deterministic encryption scheme, as an adversary can always encrypt candidate messages X_1, \dots, X_q to ciphertexts Y_1, \dots, Y_q and, if $Y_i = Y^*$ for some i , it will know that the target plaintext is $X^* = X_i$. Our security definition will formalize that this attack is (up to the adversary's advantage) the best one possible.

The idea is formalized as game MR_E in Figure 1. An MR-adversary \mathcal{A} must begin with a **Test** query and have $Q_{\text{Test}}(\mathcal{A}) = 1$ and $Q_{\text{Eq}}(\mathcal{A}) = 0$, while a simulator \mathcal{S} for \mathcal{A} is an adversary that has $\mathcal{S}(\text{dist}) = \mathcal{A}(\text{dist})$, $Q_{\text{Test}}(\mathcal{S}) = Q_{\text{Enc}}(\mathcal{S}) = 0$ and $Q_{\text{Eq}}(\mathcal{S}) = Q_{\text{Enc}}(\mathcal{A})$. Here $Q_{\text{Proc}}(\mathcal{C})$ is the maximum number of calls that adversary \mathcal{C} might make to procedure **Proc**, the maximum over all coins of \mathcal{C}

and all possible oracle responses. The MR-advantage of adversary \mathcal{A} is then defined as

$$\mathbf{Adv}_E^{\text{mr}}(\mathcal{A}) = \Pr [\text{MR}_E^{\mathcal{A}} \Rightarrow \text{true}] - p_{\mathcal{A}}$$

where $p_{\mathcal{A}} = \max_{\mathcal{S}} \Pr [\text{MR}_E^{\mathcal{S}} \Rightarrow \text{true}]$ with the maximum over all simulators \mathcal{S} for \mathcal{A} . Translating our formalism into English, an adversary making a **Test** query and some number of **Enc**-queries could do just as well forgoing its **Test** query and trading its **Enc** queries for **Eq** queries.

In our experiment defining $p_{\mathcal{A}}$ it is easy to see *what* strategy an optimal \mathcal{S} should use: it makes q **Eq**-queries, X_1, \dots, X_q , where X_1 is a most likely point output by $\mathcal{A}(\text{dist})$ for the known (N^*, T^*) ; X_2 is a second most likely point ($X_2 \neq X_1$); X_3 is a third most likely point ($X_3 \notin \{X_1, X_2\}$); and so on. If the **Eq**-oracle returns true for some X_i then \mathcal{S} calls **Finalize**(X_i); otherwise, it calls **Finalize**(X_{q+1}) where $X_{q+1} \notin \{X_1, \dots, X_q\}$ is the next most likely point after X_q . In this way \mathcal{S} will win with probability $p_{\mathcal{A}} = \sum_{i=1}^{q+1} p_i$ where $p_i = \Pr[\mathcal{A}(\text{dist}) \Rightarrow (N, T, X_i) \mid (N, T) = (N^*, T^*)]$.

Message privacy. In message privacy we are trying to measure the ability of an adversary with an encryption oracle to compute some function of a challenge plaintext X^* from its encryption C^* . If the encryption is randomized we would require that the challenge ciphertext C^* is of no use in such an attack. The formalization of this is semantic security [14]. For deterministic encryption, the intuition we aim to capture is that the adversary should do no better than it could if the encryption were ideal. In this case, the encryption oracle provides no more than the capability of testing whether a message of the adversary's choice equals the challenge message.

Our formalization closely resembles that for MR. A difference is that \mathcal{A} is asked not only to come up with the distribution on plaintexts, but also the function on which it hopes to do well. See game MP in Figure 1. An MP-adversary \mathcal{A} must begin with a **Test** query and have $Q_{\text{Test}}(\mathcal{A}) = 1$ and $Q_{\text{Eq}}(\mathcal{A}) = 0$, while a simulator \mathcal{S} for \mathcal{A} is an adversary that has $\mathcal{S}(\text{dist}) = \mathcal{A}(\text{dist})$, $Q_{\text{Test}}(\mathcal{S}) = Q_{\text{Enc}}(\mathcal{S}) = 0$, $Q_{\text{Eq}}(\mathcal{S}) = Q_{\text{Enc}}(\mathcal{A})$ and $\mathcal{S}(\text{func}) = \mathcal{A}(\text{func})$. The advantage of \mathcal{A} is defined as

$$\mathbf{Adv}_E^{\text{mp}}(\mathcal{A}) = \Pr [\text{MP}_E^{\mathcal{A}} \Rightarrow \text{true}] - p_{\mathcal{A}}$$

where $p_{\mathcal{A}} = \max_{\mathcal{S}} \Pr [\text{MP}_E^{\mathcal{S}} \Rightarrow \text{true}]$ with the maximum over all simulators \mathcal{S} for \mathcal{A} . Translating our formalism into English, an adversary making a **Test** query and some number of **Enc**-queries could do just as well in guessing $Z = \mathcal{A}(\text{func}, X^*)$ forgoing its **Test** query and trading its **Enc** queries for **Eq** queries. Note that MR-security amounts to a special case of MP-security where the function $\mathcal{A}(\text{func}, \cdot)$ is the identity function.

RELATIONS BETWEEN NOTIONS. One can pictorially describe the relationships between our four security notions like this:

$$\text{PRP} \dashrightarrow \text{SPI} \xrightarrow{\text{solid}} \text{MP} \xrightarrow{\text{solid}} \text{MR}$$

The solid arrows indicate tight implications and the broken arrows indicate lossy ones. We already noted the implications between PRP and SPI above. These can be shown to be the best possible, with the counter-example in the first case being a perfect FPE scheme and in the second case following [10]. We also noted that MP tightly implies MR. The non-obvious implication is that SPI tightly implies MP, and is proved below. Finally, MP does not imply SPI, and MR does not imply MP. For the former separation, consider an FPE scheme that has a fixed point for all keys; for the latter separation, consider an FPE that always leaks a single bit of the plaintext. We now prove the implication SPI \rightarrow MP.

Proposition 1. [SPI \Rightarrow MP] Let $E: \mathcal{K} \times \mathcal{N} \times \mathcal{T} \times \mathcal{X} \rightarrow \mathcal{X} \cup \{\perp\}$ be an FPE scheme and let \mathcal{A} be an MP adversary. Then there is an SPI adversary \mathcal{B} such that

$$\mathbf{Adv}_E^{\text{mp}}(\mathcal{A}) \leq \mathbf{Adv}_E^{\text{spi}}(\mathcal{B}).$$

In addition, adversary \mathcal{B} runs in time that of \mathcal{A} and $Q_{\text{Enc}}(\mathcal{B}) = Q_{\text{Enc}}(\mathcal{A})$. \square

Proof. Let \mathcal{A} be an MP adversary. Without loss of generality we assume \mathcal{A} never repeats an oracle query. We first construct an SPI adversary \mathcal{B} using \mathcal{A} . First, \mathcal{B} runs $\mathcal{A}(\text{dist})$ to get triple (N^*, T^*, X^*) . Then, \mathcal{B} queries $\mathbf{Test}(N^*, T^*, X^*)$ to get ciphertext Y^* and runs $\mathcal{A}(N^*, T^*)$, returning Y^* in response to its \mathbf{Test} query. Upon query $\mathbf{Enc}(N, T, X)$, first \mathcal{B} checks if $(N, T, X) = (N^*, T^*, X^*)$. If so, then \mathcal{B} returns Y^* . Otherwise it queries its own oracle $\mathbf{Enc}(N, T, X)$ and returns the result. When \mathcal{A} halts without output Z , adversary \mathcal{B} returns 1 if $Z = \mathcal{A}(\text{func}, X^*)$ and returns 0 otherwise. In the case that \mathcal{B} is run in game SPI1_E , we have that \mathcal{B} simulates for \mathcal{A} exactly the environment of MP_E . Thus

$$\Pr[\text{SPI1}_E^{\mathcal{B}} \Rightarrow \text{true}] = \Pr[\text{MP}_E^{\mathcal{A}} \Rightarrow \text{true}]. \quad (1)$$

We now specify a simulator \mathcal{S} for \mathcal{A} . We let $\mathcal{S}(\text{dist}) = \mathcal{A}(\text{dist})$ and $\mathcal{S}(\text{func}) = \mathcal{A}(\text{func})$. When executed on inputs N^*, T^* , the simulator first selects a key $K \xleftarrow{\$} \mathcal{K}$ and a random value $Y^* \xleftarrow{\$} \mathcal{X}_{N^*}$. It then runs $\mathcal{A}(N^*, T^*)$. When \mathcal{A} queries \mathbf{Test} , \mathcal{S} returns Y^* . When \mathcal{A} queries $\mathbf{Enc}(N, T, X)$, \mathcal{S} first queries $\mathbf{Eq}(X)$. If the returned value is true and $(N, T) = (N^*, T^*)$ then \mathcal{S} returns Y^* to \mathcal{A} . Otherwise it computes and returns $E(K, N, T, X)$. Finally, when \mathcal{A} halts with output Z , \mathcal{S} outputs Z . Since \mathcal{S} is such that $Q_{\text{Enc}}(\mathcal{S}) = Q_{\text{Test}}(\mathcal{S}) = 0$ and $Q_{\text{Eq}}(\mathcal{S}) = Q_{\text{Enc}}(\mathcal{A})$ we have that $p_{\mathcal{A}} \geq \Pr[\text{MP}_E^{\mathcal{S}} \Rightarrow \text{true}]$. By design $\text{MP}_E^{\mathcal{S}}$ and $\text{SPI0}_E^{\mathcal{B}}$ simulate the same environment for \mathcal{A} . Therefore, $p_{\mathcal{A}} \geq \Pr[\text{SPI0}_E^{\mathcal{B}} \Rightarrow \text{FALSE}]$ and combining this with (1) and the definition of $\mathbf{Adv}_E^{\text{mp}}(\mathcal{A})$ yields the proposition statement. \blacksquare

NONADAPTIVE SECURITY, STRONG SECURITY. We expect that nonadaptive adversaries (the “static” security setting) are sufficient for many applications of FPE—the constructed scheme is not so much a tool as an end. We consider the class of static adversaries \mathcal{S} . An adversary $\mathcal{A} \in \mathcal{S}$, on input run , decides at the beginning of its execution the sequence of queries it will ask, their number and their kind being fixed. The relations between the non-adaptive notions of security remain the same as for their adaptive counterparts as described above.

In the other direction, the notions can be strengthened to require CCA-security. This is done by adding to the games a decryption procedure. In the PRP case, procedure $\mathbf{Dec}(N, T, Y)$ would return $D_K^{N, T}(Y)$ if $b = 1$ and $\pi_{N, T}^{-1}(Y)$ otherwise, where $D = E^{-1}$ denotes the inverse of E , as defined earlier. The resulting notion is the FPE analog of what is sometimes called strong-PRP (SPRP). In the games for SPI, MP and MR, $\mathbf{Dec}(N, T, Y)$ would return $D_K^{N, T}(Y)$. The adversary is not allowed to call it on inputs N^*, T^*, Y^* and the simulator is not allowed to call it at all.

ASYMPTOTIC NOTIONS. We can adapt our definitions to the asymptotic setting. We illustrate this for PRP-security. Recall first that, in speaking of complexity, we assume that \mathcal{K} , E , and D are all given by algorithms. Also, algorithm \mathcal{K} took no input. We must slightly adjust the syntax of our FPE schemes. In particular, we provide \mathcal{K} an input of the form 1^k . The algorithm must run in probabilistic polynomial time. Algorithm E and its inverse D must run in deterministic polynomial time in the sum of their input lengths. We then say that E is *PRP-secure* if, for any PPT adversary \mathcal{A} , the function $\varepsilon(k) = \mathbf{Adv}_E^{\text{prp}}(\mathcal{A}(1^k))$ is negligible, meaning $\varepsilon(k) \in k^{-\omega(1)}$. We emphasize that it is the key K output by \mathcal{K} that, presumably, grows with the security parameter k ; the specification $\mathcal{X} = \{\mathcal{X}_N\}$ does not grow with or otherwise depend on the security parameter.

4 The Rank-then-Encipher Approach

THE IDEA. Suppose we want to build an FPE scheme \mathcal{E} the slices of which may be quite complex. As an example, we might want to do length-preserving encryption of credit cards of various lengths, the CCNs of each length having a particular checksum and satisfying specified constraints on allowable substrings. It would be undesirable to design an encryption schemes whose internal workings were tailored to the specialized task in hand. Instead, what one can do is this. First, arbitrarily order and then number the points in each slice, $\mathcal{X}_N = \{X_0, X_1, \dots, X_{n-1}\}$ where $n = |\mathcal{X}_N|$. Then, to encipher $X \in \mathcal{X}_N$, find its index i in the enumeration, encipher i to j in \mathbb{Z}_n using an integer FPE scheme, and then return X_j as the encryption of X . We call this strategy the *rank-then-encipher* approach. The method will be efficient if there is an efficient way to map each point X to its index i , to encipher i to j , and to map j back to the corresponding point X_j . Details now follow, attending more closely to formats and tweaks, and also allowing the enumeration used for mapping j to X_j to differ from that used for ranking.

DEFINITIONS. To formalize RtE encryption, we first define a *ranking* and an *unranking* function for a specification $\mathcal{X} = \{\mathcal{X}_N\}$. A ranking function is a map $rank: \mathcal{N} \times \mathcal{X} \rightarrow \mathbb{N} \cup \{\perp\}$ for which $rank_N(\cdot) = rank(N, \cdot)$ is a bijection from \mathcal{X}_N to $\mathbb{Z}_{|\mathcal{X}_N|}$. In addition, $rank_N(X) = \perp$ if $N \notin \mathcal{N}$ or $X \notin \mathcal{X}_N$. An unranking function is a map $unrank: \mathcal{N} \times \mathbb{N} \rightarrow \mathcal{X} \cup \{\perp\}$ for which $unrank_N(\cdot) = unrank(N, \cdot)$ is a bijection from $\mathbb{Z}_{|\mathcal{X}_N|}$ to \mathcal{X}_N . In addition, $unrank_N(i) = \perp$ if $i \notin \mathbb{Z}_{|\mathcal{X}_N|}$.

For the asymptotic tradition, we say that a specification $\mathcal{X} = \{\mathcal{X}_N\}$ can be *efficiently ranked* if there are (deterministic) polynomial-time computable ranking and unranking functions for $\mathcal{X} = \{\mathcal{X}_N\}$. Polynomiality is in the sum of the input lengths. Note that the security parameter is not an input to the ranking or unranking functions, but it is already built in that larger slices may take more time to rank and unrank, as the input to these functions includes the format N .

THE SCHEME. Suppose one aims to create an FPE scheme \mathcal{E} with specification $\mathcal{X} = \{\mathcal{X}_N\}_{N \in \mathcal{N}}$. Let the desired tweak space for \mathcal{E} be the set \mathcal{T} . Let $\mathbb{N}_0 = \{|\mathcal{X}_N| : N \in \mathcal{N}\} \subseteq \mathbb{N}$ be the sizes of the different slices. Then we can construct our desired FPE scheme \mathcal{E} if we have in hand: **(1)** an integer FPE scheme $E: \mathcal{K} \times \mathbb{N}_0 \times \{0, 1\}^* \rightarrow \mathbb{N}$ (it enciphers points in \mathbb{Z}_n for each $n \in \mathbb{N}_0$), and **(2)** a ranking function $rank$ and an unranking function $unrank$ for $\mathcal{X} = \{\mathcal{X}_N\}_{N \in \mathcal{N}}$. Given such objects, define $\mathcal{E} = \text{RtE}[E, rank, unrank]$ as the map $\mathcal{E}: \mathcal{K} \times \mathcal{N} \times \mathcal{T} \times \mathcal{X} \rightarrow \mathcal{X} \cup \{\perp\}$ with

$$\mathcal{E}_K^{N,T}(X) = unrank_N(E_K^{|\mathcal{X}_N|, \langle N, T \rangle}(rank_N(X)))$$

when $X \in \mathcal{X}_N$, and $\mathcal{E}_K^{N,T}(X) = \perp$ otherwise. We call this *rank-then-encipher* approach. In words: convert the N -formatted string X to its corresponding number i ; encipher $i \in \mathbb{Z}_{|\mathcal{X}_N|}$ to some $j \in \mathbb{Z}_{|\mathcal{X}_N|}$, employing a tweak that encodes both the format N of X and the tweak of \mathcal{E} ; finally, convert j back to a domain point in $Y \in \mathcal{X}_N$ using a possibly unrelated enumeration of points.

We will omit formalizing and proving the rather obvious statements that, if E is secure with respect to the strong-PRP, PRP, SPI, MP, or MR notion of security, then so too will be the FPE scheme $\mathcal{E} = \text{RtE}[E, rank, unrank]$, the reduction being tight and having time complexity that is approximately the sum of the times to perform the ranking and unranking.

By way of the rank-then-encipher approach, one can take an integer FPE (based, e.g., on the techniques described in [4]) and create from it an FPE with a quite intricate specification $\mathcal{X} = \{\mathcal{X}_N\}_{N \in \mathcal{N}}$.

For many specifications the needed ranking and unranking functions are simple to design and fast to compute: an *ad hoc* approach will work fine. But what can one say in general about the power of the rank-then-encipher FPE approach? We now turn our attention to this.

<pre> algorithm BuildTable(n) for $q \in Q$ do if $q \in F$ then $T[q, 0] \leftarrow 1$ for $i \leftarrow 1, \dots, n$ do for $q \in Q$ do for $a \in \Sigma$ do $T[q, i] \leftarrow^{\pm} T[\delta(q, a), i - 1]$ </pre>	<pre> algorithm $rank(X)$ $q \leftarrow q_0; c \leftarrow 0; n \leftarrow X$ for $i \leftarrow 1, \dots, n$ do for $j \leftarrow 1, \dots, \mathbf{ord}(X[i]) - 1$ do $c \leftarrow^{\pm} T[\delta(q, a_j), n - i]$ $q \leftarrow \delta(q, X[i])$ ret c </pre>	<pre> algorithm $unrank(c)$ $X \leftarrow \varepsilon; q \leftarrow q_0; j \leftarrow 1$ for $i \leftarrow 1, \dots, n$ do while $c \geq T[\delta(q, a_j), n - i]$ do $c \leftarrow c - T[\delta(q, a_j), n - i]; j \leftarrow^{\pm} 1$ $X[i] \leftarrow a_j; q \leftarrow \delta(q, X[i]); j \leftarrow 1$ ret X </pre>
---	--	---

Fig. 2. Middle: Algorithm for computing the rank of a word in the regular language L of a DFA $M = (Q, \Sigma, \delta, q_0, F)$. **Left:** Initializing the table T . Each $T[\cdot, \cdot]$ starts at zero. **Right:** How to compute the inverse of the ranking function.

5 FPE for Arbitrary Regular Languages

THE PROBLEM. Let Σ be a (finite) alphabet and let $L \subseteq \Sigma^*$ be a language over it. We say that an FPE scheme $E: \mathcal{X} \times \mathcal{N} \times \mathcal{T} \times \mathcal{X} \rightarrow \mathcal{X} \cup \{\perp\}$ is an FPE scheme *for* L if $\mathcal{X} = L$, $\mathcal{N} = \mathbb{N}$, and the slices are $\mathcal{X}_n = L_n = L \cap \Sigma^n$ for all $n \in \mathbb{N}$. In this section we show how to build an FPE for an arbitrary regular language L by describing how to compute a corresponding ranking and unranking function.

Why attend to regular languages? Many FPE specifications can be cast as asking for an FPE for a regular language. This is trivially true when the domain is finite. Some important domains are finite and without an easily summarized structure; a domain like “a valid postal address” is likely to be defined by a database such as the US Address Information System (AIS) and, given such a database, ranking is easy. Other finite domains are large but have a concise description as a regular language, either in terms of a regular expression or a DFA. For example, a US social security number is a string in the regular language $(0 \cup 1 \cup \dots \cup 9)^9$. Alternatively, one may subtract from this any set of numbers that have not been assigned, such as those starting with an 8 or 9, having 0000 as the last four digits, or having 00 as the preceding two digits, but the resulting set will again have a concise description. For credit card numbers, a simple 20-state DFA M recognizes the language $Luhn^R$ of strings that are the reversals of numbers with a valid Luhn checksum [20]. Namely, the DFA is $M = (Q, \Sigma, \delta, q_0, F)$ with states $Q = \mathbb{Z}_{10} \times \mathbb{Z}_2$, final states $F = \{0\} \times \mathbb{Z}_2$, start state $q_0 = (0, 0)$, and transition rule $\delta((a, b), d) = (a + 2d + a[d/5] \bmod 10, 1 - b)$. We will continue to use the $M = (Q, \Sigma, \delta, q_0, F)$ syntax below, following the convention of Sipser’s book [40].

RANK COMPUTATION FOR REGULAR LANGUAGES. We will describe efficient ranking and unranking functions for the specification $\mathcal{X} = \{\mathcal{X}_M\}$ where M is a DFA and $\mathcal{X}_M = L(M)$ is its language. First impose a total order $a_1 \prec \dots \prec a_{|\Sigma|}$ on the elements of the alphabet $\Sigma = \{a_1, \dots, a_{|\Sigma|}\}$ and extend this to the lexicographic order \prec on each Σ^n . For $a \in \Sigma$ let $\mathbf{ord}(a)$ be the index i such that $a = a_i$ and for every $n \in \mathbb{N}$ let the ranking function be given by $rank_L(X) = |\{Y \in L: |X| = |Y| = n \text{ and } Y \prec X\}|$. We omit the argument $n = |X|$ because it is determined by X . Assume we have an integer FPE scheme E . Provided that we can efficiently compute each $rank_L(\cdot)$ and its inverse $unrank_L(\cdot)$, applying the RtE paradigm gives a practical FPE $\mathcal{E} = \text{RtE}[E, rank_L, unrank_L]$ with $\mathcal{E}: \mathcal{K} \times \mathcal{N} \times \mathcal{T} \times L \rightarrow L \cup \{\perp\}$.

Let $M = (Q, \Sigma, \delta, q_0, F)$ be a DFA recognizing the regular language $L \subseteq \Sigma^*$. Let $X[i]$ denote the i -th character of $X \in \Sigma^*$ (numbering from the left and starting at 1). Extend δ to $Q \times \Sigma^*$ so that $\delta(q, X)$ is the state we end up in by starting from q and following $X \in \Sigma^*$. Formally, set $\delta(q, \varepsilon) = q$ for all $q \in Q$ and recursively define $\delta(q, x) = \delta(\delta(q, X[1] \dots X[n-1]), X[n])$ for all $q \in Q$ and all $X \in \Sigma^*$ with $n = |X| \geq 1$.

We compute the ranking function for M by dynamic programming, following [12]. Let $T[q, n]$ be the number of strings $X \in \Sigma^n$ such that $\delta(q, X) \in F$. The first algorithm of Figure 2, on input n , uses dynamic programming to compute, for all $q \in Q$ and $j \in [1..n]$, the number $T[q, j]$ of

accepting paths of length j that start at q . The rank of a word in L can be computed based on T as shown by the second algorithm in Figure 2. The third algorithm in the figure computes the inverse, deriving a word in L by its rank. In the unit-cost model of computation, where arbitrary integer multiplications and additions are performed in unit time, $rank_M$ and $unrank_M$ can be computed in $O(|\Sigma| \cdot n)$ time, while the preprocessing step $BuildTable(n)$ takes time $O(|Q| \cdot |\Sigma| \cdot n)$ time.

We comment that ranking can be further sped up to require about n sums instead of $n|\Sigma|$ by precomputing the needed partial sums, adding a third coordinate to T . The unranking function would need a binary search, or some other method, to map a number into the corrected (precomputed) interval $[0, \beta_1), [\beta_1, \beta_2), \dots, [\beta_{\sigma-1}, \beta_\sigma)$ that contains it, where $\sigma = |\Sigma|$. Regardless, ranking and unranking are linear-time for any regular language L , with modest constants in terms of the DFA representation of L .

ON THE IMPORTANCE OF REPRESENTATIONS. It is important that we represented our regular language in terms of a DFA; had L been represented in terms of an NFA or a regular expression, we could not have efficiently computed the ranking and unranking functions. In particular, remember that it is NP-hard (even PSPACE-hard) to decide if the language of an NFA M (or a regular expression α) is Σ^* [11, #AL1], [15]. Consequently, if $P \neq NP$, we can't compute $unrank(2^n - 1)$ efficiently for all n , as such functionality would provide immediate means to decide if $L(M) = \Sigma^*$. Formally, if $P \neq NP$ then \mathcal{X}_M can't be efficiently ranked, where $\mathcal{X}_M = L(M)$ is the language of the NFA M . Note, however that this does not imply an inability to make an efficient FPE scheme for this specification—it only means that such a scheme could not use the RtE approach.

RANKING NON-REGULAR LANGUAGES. Beyond regular languages, we can also apply the RtE approach with Mäkinen's ranking algorithm for the language generated by an unambiguous context-free grammar [25]. Efficient ranking algorithms exist for various other classes of combinatorial objects. For example, if we wish to encrypt the domain $\mathcal{X}_{n!}$ consisting of the set of permutations on n elements, the Lucas-Lehmer encoding [18] provides an efficient ranking. Other examples are spanning trees of a graph [8], B-trees [21], and Dyck languages [19]. Efficient rankings have also been studied in coding theory, starting with [9].

6 FPE without Ranking

Given the ease of ranking regular languages and beyond, it is natural to ask if *every* language for which there is an efficient FPE scheme admits an RtE-style one. In this section we show that the answer is *no*. More specifically, we exhibit a specification $\mathcal{X} = \{\mathcal{X}_N\}_{N \in \mathcal{N}}$ where efficient FPE is possible but efficient ranking is not. This assumes the existence of a one-way function. We warn up front that our result is not practical, in the sense that nobody would ever want to do format-preserving encryption on our chosen specification $\{\mathcal{X}_N\}_{N \in \mathcal{N}}$. We leave it as an interesting open problem to find a practicable specification for which efficient FPE is possible but ranking is not.

A SPECIFICATION THAT CAN'T BE RANKED. We will let each format $G \in \mathcal{N}$ specify a (simple, undirected) graph $G = (V, E)$. Slice \mathcal{X}_G consists of all proper κ -colorings of G , where $\kappa = 2\Delta + 1$ and $\Delta = \max_{v \in V} d(v)$ is the maximum degree of any vertex. Recall that a coloring is an assignment of colors to vertices, and a coloring is *proper* if it uses only allowed colors and adjacent vertices never receive the same color. A coloring $color \in \mathcal{X}_G$ on the n -vertex graph G can be regarded as a map $color: \{1, \dots, n\} \rightarrow \{0, \dots, \kappa - 1\}$ where the vertices have names $1, \dots, n$ and the colors have names $0, 1, \dots, \kappa - 1$. The value of $color$ could be conveniently represented by a string $color \in \mathcal{C}^n$ where $\mathcal{C} = \{0, 1, \dots, \kappa - 1\}$.

Our definition of having an efficient ranking on \mathcal{X}_G is quite strict: in particular, via binary search, one can efficiently extract from the unranking function $unrank$ the cardinality of \mathcal{X}_G . In

other words, efficient ranking and unranking of $\mathcal{X} = \{\mathcal{X}_G\}$ is at least as hard counting the elements of each \mathcal{X}_X . But a result of Bubley, Dyer, Greenhill, and Jerrum [7, Section 6] says that it is $\#P$ -complete to count the number of proper κ -colorings of a maximum-degree- Δ graph, even for a fixed $\kappa, \Delta \geq 3$. As a consequence, one can't FPE via RtE on our specification $\mathcal{X} = \{\mathcal{X}_G\}$ assuming $P \neq \#P$. Note that $P \neq NP$, or the existence of a one-way function, already implies that $P \neq \#P$.

HOW TO FPE ON THIS SPECIFICATION. A classical paper by Jerrum [16] shows how to efficiently produce an almost-uniform proper κ -coloring $color$ of the n -vertex graph $G = (V, E)$ (where, as we have assumed already, that $\kappa \geq 2\Delta + 1$). Beginning with an arbitrary proper coloring $color$ of G , repeat the following, for iterations $t = 1, \dots, ntimes$:

- (1) Uniformly select $v \xleftarrow{\$} \{1, \dots, n\}$ and $c \xleftarrow{\$} \{0, \dots, \kappa - 1\}$.
- (2) Let $color^*$ be identical to $color$ except for setting $color^*(v) = c$.
- (3) If $color^*$ is a proper coloring then replace $color$ by $color^*$.

The final value of $color$ is the coloring produced. Jerrum shows that $ntimes = \kappa/(\kappa - 2\Delta)n \ln(n/\varepsilon)$ repetitions is enough so that no adversary can distinguish the resulting coloring from a uniform one with advantage exceeding ε (the original language is in terms of stopping times and total-variation distance). For our choice of $\kappa = 2\Delta + 1$, one needs at most $ntimes = 2n^2 \ln(n/\varepsilon)$ repetitions to ensure an advantage of at most ε .

The procedure above cannot directly be used to encrypt because it is not *computationally reversible*. What we mean by this is the following. Suppose that the transition function of a Markov chain is described by an (efficiently computable) function f : when the Markov chain is in state q and coins $c \xleftarrow{\$} \Omega$ are selected, the next state is $q' = f(q, c)$. Then f is computationally reversible if there exists an (efficiently computable) function g such that $q = g(q', c)$. In such a case, if the chain goes from an initial state of q_0 to a final state of q_t using coins c_1, \dots, c_t , then knowledge of the final state q_t and the coins c_1, \dots, c_t allows recovery of q_0 . Now to make our transition function efficiently recoverable for the Jerrum chain, just replace step (2), above, by

- (2) Let $color^*$ be identical to $color$ except for setting $color^*(v) = (color(v) + c) \bmod \kappa$.

Formally, the Markov chain is completely unchanged: one transitions from q to q' with the same probability as before. But with this alternative description, transitions can readily be reversed, by setting $\kappa(v) = (\kappa(v') - c) \bmod \kappa$. This provides the basis of a simple encryption scheme. First define a pseudorandom function F that, keyed by key K of length k , maps the numbers t, n, κ to an output in $\{1, \dots, n\} \times \{0, \dots, \kappa\}$. Use F to generate the needed coins in step t . Run the chain for $2n^2 k \lg n$ steps. It is easy to see that if F is indeed a PRF then, using the result of Jerrum, we have constructed a secure FPE for the specified domain. We summarize our finding as follows.

Theorem 1. *Suppose there exists a one-way function. Then there is a specification $\mathcal{X} = \{\mathcal{X}_N\}$ that admits a PRP-secure FPE scheme but for which $\{\mathcal{X}_N\}$ cannot be efficiently ranked.*

OTHER EXAMPLES. Our choice of an FPE specification involving proper graph colorings was not necessary: one could have selected other $\#P$ -complete problems. For example, it would have worked to select the space of perfect matchings of a bipartite graph [44]. Here again there is a Monte Carlo process that rapidly mixes the Markov chain, this due to Jerrum, Sinclair and Vigoda [17]. The process, as described by the authors, again fails to be computationally reversible, but it is once again possible to recast the process so that it is computationally reversible. It is somewhat more complex than with the example chosen. In fact, we rather expect that most computationally interesting Markov processes can be recast so as to make them computationally reversible.

An alternative approach to FPE without ranking is provided by cycle walking. Suppose, for example, that one defines a 1-bit pseudorandom function F on strings and declares that $\mathcal{X} = \{\mathcal{X}_n\}$ where $\mathcal{X}_n = \{x \in \{0, 1\}^n : F_K(x) = 1\}$. The set is easy to encrypt (in efficient *expected* time) but, one would expect, computationally inefficient to rank and unrank.

7 Feistel-Based Integer FPEs

We present two Feistel-based constructions of integer FPE schemes $E: \mathcal{K} \times \mathcal{N} \times \mathcal{T} \times \mathcal{X} \rightarrow \mathcal{X} \cup \{\perp\}$ with format space $\mathcal{N} = \mathbb{N} \times \mathbb{N}$ and \mathcal{X} such that $\mathcal{X}_N = \mathbb{Z}_{ab}$ for $N = (a, b)$ with $a \leq b$. Both are parameterized by the following: (1) a round function $F: \mathcal{K} \times \mathcal{N} \times \mathcal{T} \times \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$; and (2) a function $r: \mathcal{N} \rightarrow \mathbb{N}$ specifying the number of rounds.

Figure 3 defines encryption and decryption for the two integer FPE schemes FE1 and FE2. We refer to Feistel networks, such as FE1, that utilize the same kind of round function every round as type-1. Type-1 Feistel networks were previously treated in [28, 38] for the case of bit strings. We refer to Feistel networks, such as FE2, that alternate the kind of round function as type-2. Type-2 Feistel networks for the case of bit strings are due to Lucks [24]. Type-2 Feistel networks with modular arithmetic were first used in [4].

ROUND FUNCTIONS. The round functions should be PRFs. It is not clear what this means when the range is the infinite set \mathbb{N} . To specify a round function, we will first specify a range function $w: \mathcal{N} \rightarrow \mathbb{N}$ such that for all $N \in \mathcal{N}$ we have $w(N) \geq b$ where $N = (a, b)$. The PRF advantage of an adversary \mathcal{A} is then defined by

$$\mathbf{Adv}_F^{\text{prf}}(\mathcal{A}) = \Pr \left[\mathcal{A}^{F(K, \cdot, \cdot, \cdot)} \Rightarrow 1 \right] - \Pr \left[\mathcal{A}^{\mathfrak{S}(\cdot, \cdot, \cdot)} \Rightarrow 1 \right]$$

where \mathcal{A} 's oracle in the second case returns a random point in $\mathbb{Z}_{w(N)}$ in response to a query N, T, i, X . Adversary \mathcal{A} is not allowed to repeat an oracle query.

In cases of practical interest, we can build suitable round functions based on block ciphers (e.g. 3DES or AES) or cryptographic hash functions (e.g. SHA-256). See Appendix A for example instantiations. We also discuss there the use of precomputation for speed improvements (deriving from the fact that several of the inputs to F are the same across all rounds).

DISCUSSION. The round function takes as input the format and tweak, which effectively provides “separate” instances of the cipher for each format, tweak pair. To ensure independence between rounds, the round number is also input into the PRF.

FE1 and FE2 support domains of the form \mathbb{Z}_{ab} and only provide security when $a > 1$. To handle arbitrary \mathbb{Z}_n one can choose $N = (a, b)$ so that $ab > N$ and then utilize the cycle walking technique with FE1 or FE2 (see [4] for a treatment). Alternatively, one might utilize the off-by-one construction (see [4]) to avoid cycle-walking. But for typical applications like the encryption of credit card numbers, the requisite domains will be \mathbb{Z}_n for which $n = ab$ for a and b that are almost balanced.

8 Security of Feistel FPEs

We discuss FE1 and FE2 in turn, treating both attacks and proofs of security. We distinguish between when information-theoretic or computational security is compromised by an attack. The first means we are concerned only with the advantage as a function of the number of oracle queries, and the second means the running time of the adversary is also relevant. The distinction is important because there are some attacks, like Patarin’s [30], defending against which requires a lot of rounds for information-theoretic security but, due to the large running time of the attack, very few for

<pre> algorithm FE1$_{K}^{N,T}(X)$ $(a, b) \leftarrow N$; $X_0 \leftarrow X$ for $i = 1, \dots, r(N)$ do $L_{i-1} \leftarrow X_{i-1} \text{ div } b$ $R_{i-1} \leftarrow X_{i-1} \text{ mod } b$ $W_i \leftarrow (L_{i-1} + F_K(N, T, i, R_{i-1})) \text{ mod } a$ $X_i \leftarrow aR_{i-1} + W_i$ ret $X_{r(N)}$ </pre>	<pre> algorithm FE2$_{K}^{N,T}(X)$ $(a, b) \leftarrow N$ $L_0 \leftarrow X \text{ div } b$; $R_0 \leftarrow X \text{ mod } b$ for $i = 1, \dots, r(N)$ do If $i \text{ mod } 2 = 1$ then $s \leftarrow a$ else $s \leftarrow b$ $L_i \leftarrow R_{i-1}$ $R_i \leftarrow (L_{i-1} + F_K(N, T, i, R_{i-1})) \text{ mod } s$ ret $sL_{r(N)} + R_{r(N)}$ </pre>
<pre> algorithm FD1$_{K}^{N,T}(Y)$ $(a, b) \leftarrow N$; $Y_{r(N)} \leftarrow Y$ for $i = r(N), \dots, 1$ do $W_i \leftarrow Y_i \text{ mod } a$ $R_{i-1} \leftarrow Y_i \text{ div } a$ $L_{i-1} \leftarrow (W_i - F_K(N, T, i, R_{i-1})) \text{ mod } a$ $Y_{i-1} \leftarrow bL_{i-1} + R_{i-1}$ ret Y_0 </pre>	<pre> algorithm FD2$_{K}^{N,T}(Y)$ $(a, b) \leftarrow N$ If $r(N) \text{ mod } 2 = 1$ then $s \leftarrow a$ else $s \leftarrow b$ $R_{r(N)} \leftarrow Y \text{ mod } s$; $L_{r(N)} \leftarrow Y \text{ div } s$ for $i = r(N), \dots, 1$ do If $i \text{ mod } 2 = 1$ then $s \leftarrow a$ else $s \leftarrow b$ $R_{i-1} \leftarrow L_i$ $L_{i-1} \leftarrow (R_i - F_K(N, T, i, R_{i-1})) \text{ mod } s$ ret $bL_0 + R_0$ </pre>

Fig. 3. Top: Encryption and decryption algorithms for the integer FPE scheme FE1 (left) and FE2 (right), where $K \in \mathcal{K}$, $T \in \mathcal{T}$, $N \in \mathcal{N}$, and $X, Y \in \mathcal{X}_N$. Here $x \text{ div } y$ means $\lfloor x/y \rfloor$. **Bottom:** The corresponding decryption algorithms.




Fig. 4. Diagrams depicting three rounds of type-1 Feistel (**left**) or three rounds of type-2 Feistel (**right**). Here $N = (a, b) = (2^a, 2^b)$, input $X \in \mathbb{Z}_{ab}$, and F_K maps points in \mathbb{Z}_a to points in \mathbb{Z}_b .

computational security, so that it is not a practical threat. Computational security against an attack is what's relevant in practice. An information-theoretic attack is not an indication of a practical weakness but, due to the fact that proofs (so far) use information-theoretic techniques, is interesting because it tells us how good a bound we can hope to prove.

8.1 Security of FE1

MINIMAL NUMBER OF ROUNDS. Consider any $N = (a, b) \in \mathcal{N}$. Let $p(a, b) = \lceil \log_a ab \rceil = 1 + \lceil \log_a b \rceil$. When a, b are clear from context, we will sometimes write just p . This is the number of rounds

required to implement a full *pass*, as per the terminology in [27]. Security mandates that at the least $r((a, b)) \geq p(a, b) + 1$. Otherwise a simple PRP adversary can successfully distinguish with just a few queries. For example consider when $r((a, b)) = p(a, b)$. Let X, X' be such that $(X \operatorname{div} b) \neq (X' \operatorname{div} b)$ and $(X \operatorname{mod} b) = (X' \operatorname{mod} b)$. Then necessarily $(\operatorname{FE1}_K^{N,T}(X) \operatorname{div} b) - (\operatorname{FE1}_K^{N,T}(X') \operatorname{div} b) = (X \operatorname{div} b) - (X' \operatorname{div} b)$. But this relation only holds with probability a^{-1} when FE1 is replaced by a random function.

PATARIN'S ATTACK. In [30], Patarin gave a simple distinguishing attack against the PRP security of balanced Feistel networks using ideal round functions. In the attack, the adversary \mathcal{A} queries q distinct messages. Then \mathcal{A} checks if there exists a set of round functions that, when used within the Feistel network, map the q inputs to the q outputs. If no such set of rounds exists, then \mathcal{A} knows that its oracle is a random permutation. This attack easily generalizes to the type-1 Feistel network of FE1 and we discuss its implications for the choice of the number of rounds. The bottom line is that information-theoretic security against Patarin's attack requires rather a lot of rounds. But the attack is computationally expensive, and computational security requires few rounds. Additionally, the attack is a PRP one but not an MR one. For these reasons, the attack is not a practical threat. We now try to detail the analysis.

Fix some format $N = (a, b)$, let $n = ab$, and let $r = r(N)$. We let \mathcal{R} be the set of all possible round functions, so that $|\mathcal{R}| = a^b$. Each vector $\mathbf{r} \in \mathcal{R}^r$ of r round functions specifies a Feistel permutation that we denote $\operatorname{Fe}_{\mathbf{r}}$. The number of such vectors is a^{rb} . The adversary \mathcal{A} queries its oracle on $q \leq n$ distinct inputs $\mathbf{x}[1], \dots, \mathbf{x}[q]$ to get outputs $\mathbf{y}[1], \dots, \mathbf{y}[q]$. We say that \mathbf{y} is compatible with \mathbf{r} if $\operatorname{Fe}_{\mathbf{r}}(\mathbf{x}[i]) = \mathbf{y}[i]$ for all $1 \leq i \leq q$. If there is some $\mathbf{r} \in \mathcal{R}^r$ that is compatible with \mathbf{y} then \mathcal{A} returns 1, else 0.

The running time of the attack is at least a^{rb} . In practice we are safe if this is not more than, say, 2^{80} , meaning r needs to be at least $80/(b \lg a)$. Encrypting 2-digit numbers with $a = 10$, we need only three rounds to counter the attack, and this number does not increase if we encrypt more digits, regardless of the balance.

To reduce the running time, the adversary may mount the attack using only some subset of the vectors, say an ϵ fraction of them. In this case, it achieves advantage at most ϵ with time ϵa^{rb} . In such a setting the right metric of adversarial effort is work factor, meaning the ratio of its running time to its advantage, which continues to be a^{rb} . So this change in the attack has not bought the adversary anything.

That's the computational security. Now let's consider information-theoretic security, where the picture is quite different. The prp-advantage α of \mathcal{A} is $1 - p$ where p is the probability that, for a random vector \mathbf{y} with distinct entries, there exists $\mathbf{r} \in \mathcal{R}^r$ which is compatible with \mathbf{y} . We want to ensure that α is small, meaning p is close to 1. Let

$$p^* = \frac{a^{rb}}{n!/(n-q)!}. \quad (2)$$

This is a good approximation to p so we will endeavor to choose r to make it 1. (The quantity p^* is definitely an upper bound on p , but may not equal it because there may be different vectors \mathbf{r}, \mathbf{r}' for which $\operatorname{Fe}_{\mathbf{r}} = \operatorname{Fe}_{\mathbf{r}'}$. We are assuming this happens seldom, which is likely true if $q \geq a$, so let's assume this.) Let's say we want security up to $q = n$ queries. In this case, the equation $p^* = 1$ gives us $a^{rb} = n! \approx \sqrt{2\pi n}(n/e)^n$, whence $r \approx (n \ln n)/(b \ln a)$. If we are encrypting 9-digit numbers in an almost balanced way, meaning $a = 10^4$ and $b = 10^5$, this means we need almost 22,500 rounds. Things get better with more imbalance: if $a = 10$ and $b = 10^8$ then $r = 90$.

ATTACKS FROM [34]. Patarin, Nachev, and Berbain investigate the security of type-1 Feistel networks [34] in the case that $a = 2^m$ and $b = 2^{(p-1)m}$ for some m and $p \geq 3$. When $p = 3$, they give

PRP attacks requiring $q \approx 2^{m/2}$ for 4 rounds, $q \approx 2^m$ for 5 rounds, and $q \approx 2^{2m}$ for 6 rounds. When $p \geq 4$ they give PRP attacks requiring $q \approx 2^{(2i-1)m/2}$ for $p+i$ rounds with $1 \leq i < p$. These attacks require time approximately q^2 and aren't even applicable after r moves beyond $2p$.

PRIOR SECURITY BOUNDS. Naor and Reingold gave results regarding the security of type-1 Feistel networks in the case³ where $\alpha = \log_2 a \in \mathbb{N}$ and $\beta = \log_2 b \in \mathbb{N}$ [28]. They showed that when (say) $r = 3p$ any PRP adversary's advantage when round functions are random functions is upper bounded by

$$\frac{\alpha + \beta}{2\alpha} \cdot \frac{q^2}{b} + \frac{q^2}{n}.$$

In the extreme case when $a = 2$ security up to almost \sqrt{n} queries is proven. More recently, Morris, Stegers and Rogaway [27] considered the case $a = 2$ and $\log_2 b \in \mathbb{N}$, corresponding to the Thorpe shuffle, and proved non-adaptive PRP security up to about $n/8 \log_2(n)$ queries, where $n = ab = 2b$, if one uses $r \geq 2p \log_2(n) = 2 \log_2^2 n$ rounds. (Adaptive security requires twice as many rounds.) As an example, this means 1,800 rounds when encrypting a 30-bit string. (30 is about the number of bits in a 9-digit number but here we need b a power of 2 so can't work directly over digits.) Patarin's attacks says that information-theoretic security here (which is what is proved) would require 60 rounds, so the truth is somewhere in between. Computational security requires at least 30 rounds in this case by our first attack.

Proofs in general mean greater assurance, but lack of a proof for fewer rounds does not mean there is an attack in that case. Right now, there is a fairly large gap between the security indicated by attacks and proofs. Which you follow depends on how conservative you want to be and how much you are willing to pay in processing.

NON-ADAPTIVE SPI SECURITY. We provide a new security bound for FE1 by proving non-adaptive SPI security whenever $p(a, b)$ is a whole number (i.e. $b = a^k$ for $k \in \mathbb{N}$ and $k = p(a, b) - 1$) and when $r((a, b)) \geq p(a, b) + 1$. As discussed above, this number of rounds is minimal, and so our theorem shows that this number of rounds also achieves SPI security up to the indicated bounds.

Theorem 2. *Let the FE1 FPE scheme be parameterized by $F: \mathcal{K} \times \mathcal{N} \times \mathcal{T} \times \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ and $r: \mathcal{N} \rightarrow \mathbb{N}$ where $r(N) = p(N) + 1$ for all $N \in \mathcal{N}$. Let \mathcal{A} be a non-adaptive SPI adversary that makes q encrypt queries with formats $N = (a, b)$ with $2 \leq a \leq b$ and where $p(a, b) = \log_a ab$ is a whole number. Let $N_m = (a_m, b_m)$ be the format $N = (a, b)$ queried by \mathcal{A} with minimal a . Let $p_m = p(a_m, b_m)$. Then there exists a PRF adversary \mathcal{B} such that*

$$\mathbf{Adv}_{\text{FE1}}^{\text{spi}}(\mathcal{A}) \leq 2 \cdot \mathbf{Adv}_F^{\text{prf}}(\mathcal{B}) + \frac{p_m q}{a_m}.$$

\mathcal{B} runs in time that of \mathcal{A} and makes $\sum_{1 \leq i \leq q} p(N_i)$ queries where N_i is the format of the i -th query.

Proof. We assume without loss of generality that \mathcal{A} never makes a query for which \perp is returned. Our proof proceeds in steps, using intermediate games G0 and G1 with behaviors between that of SPI1_{FE1} and SPI0_{FE1}. Game G0 (shown in Figure 5, boxed statements included) is the same as SPI1_{FE1} except that F is replaced by random functions. Game G1 (shown in Figure 5, boxed statements excluded) is the same as G0 except that the random values used in responding to the **Test** query are no longer used when responding to **Enc** queries. The result of this change is that **Test** responds with a random, independent point. To see this, note that the values W_2^*, \dots, W_r^* are uniformly distributed in \mathbb{Z}_a due to the random choices of Z_2^*, \dots, Z_r^* . By assumption, we have

³ In fact, they treat a slightly more general case, see [28] for details. Moreover, their result utilizes pairwise-independent permutations for the first and last round; we assume these to be implemented via more rounds.




Fig. 5. Games used in the proof of Theorem 2. Each function \mathcal{S}_x maps every input to an independently and uniformly sampled point in \mathbb{Z}_x .

that $ab = a^p$ for some number p . (Equivalently, that $p = \log_a ab$ is a whole number.) The returned value X_r^* can therefore be written as $a^{p-1}W_2^* + \dots + a \cdot W_{r-1}^* + W_r^*$ which is distributed uniformly over \mathbb{Z}_{ab} .

We have that

$$\begin{aligned}
\mathbf{Adv}_{\text{FE1}}^{\text{spi}}(\mathcal{A}) &\leq \Pr[\text{SPI1}_{\text{FE1}}^{\mathcal{A}} \Rightarrow \text{true}] - \Pr[\text{G0}^{\mathcal{A}} \Rightarrow 1] \\
&\quad + \Pr[\text{G0}^{\mathcal{A}} \Rightarrow 1] - \Pr[\text{G1}^{\mathcal{A}} \Rightarrow 1] \\
&\quad + \Pr[\text{G1}^{\mathcal{A}} \Rightarrow 1] - \Pr[\text{SPI1}_{\text{FE1}}^{\mathcal{A}} \Rightarrow \text{FALSE}].
\end{aligned}$$

A standard reduction bounds the first and last differences each by $\mathbf{Adv}_F^{\text{prf}}(\mathcal{B})$ for an adversary \mathcal{B} using the resources indicated in the theorem statement. To bound the second difference, we use that game G0 and G1 are identical-until-*bad* with the fundamental lemma of game playing [1] to get that

$$\Pr[\text{G0}^{\mathcal{A}} \Rightarrow 1] - \Pr[\text{G1}^{\mathcal{A}} \Rightarrow 1] \leq \Pr[\text{G1}^{\mathcal{A}} \text{ sets } bad]$$

where “G1^A sets *bad*” is the event that *bad* is set during execution of G1^A. Recall that \mathcal{A} is non-adaptive, meaning has fixed queries independent of all other random choices in the game. The flag *bad* is set if, for some $i \in [2..r]$ and some $j \in [1..q]$ it is the case that $R_{i-1}^* = R_{i-1}^j$. We split the argument into two cases, depending on whether the flag *bad* is set when $i = 2$ or $i > 2$. For the $i = 2$ case, note that if $R_0^* = R_0^j$ for some j then $R_1^* \neq R_1^j$ (we have disallowed pointless queries). If $R_0^* \neq R_0^j$, then for *bad* to be set it must be that

$$R_1^* = (aR_0^* + Z_1^*) \text{ mod } b = (aR_0^j + Z_1^j) \text{ mod } b = R_1^j$$

where (a, b) is the format of the i -th **Enc** query. But this equality holds with probability at most a_m^{-1} (taken over the coins used to choose Z_i^*) because $a \geq a_m$ by assumption. For the case that $i > 2$, we have that $R_{i-1}^* = (aR_{i-2}^* + W_{i-1}^*) \bmod b$ where W_k^* is uniformly and independently distributed for all $k \in [2..r]$. Thus the probability that $R_{i-1}^* = R_{i-1}^j$ is at most a_m^{-1} .

Taking a union bound over all i, j we have that $\Pr[\text{G1}^{\mathcal{A}} \text{ sets bad}] \leq p_m q / a_m$. ■

8.2 Security of FE2

As one would expect, Patarin's attack also applies to type-2 unbalanced Feistel networks. Its success and running time can be easily adapted from the discussion for FE1 given above. In the following, we describe a novel PRP attack for type-2 Feistel networks, a variant of it that allows message recovery, and then provable security bounds for FE2.

A NEW DISTINGUISHING ATTACK. Highly unbalanced Feistel networks are susceptible to highly efficient attacks that succeed with exponentially vanishing probability as the number of rounds increases. Still, for small, fixed round number, the attacks could be dangerous. We present a PRP adversary \mathcal{A} against FE2 for some format $N = (a, b)$ and arbitrary tweak T . Denote $r(N)$ by r and assume r is even (the attack easily extends to the case that r is odd). Then adversary \mathcal{A} works as described below.

adversary $\mathcal{A}^{\text{Enc}(\cdot, \cdot)}$
 $L_0 \xleftarrow{\$} \mathbb{Z}_a; L'_0 \leftarrow L_0; R_0 \xleftarrow{\$} \mathbb{Z}_b; R'_0 \xleftarrow{\$} \mathbb{Z}_b \setminus \{R_0\}$
 $Y \leftarrow \mathbf{Enc}(N, T, b \cdot L_0 + R_0); Y' \leftarrow \mathbf{Enc}(N, T, b \cdot L'_0 + R'_0)$
 $D \leftarrow Y \bmod b; D' \leftarrow Y' \bmod b$
 if $D - R_0 \equiv D' - R'_0 \pmod{b}$ then ret 1 else ret 0

Let game PRP0 (resp. PRP1) be the same as game PRP (Figure 1) except b in **Initialize** is assigned 0 (resp. 1). First we analyze $\Pr[\text{PRP}^{\mathcal{A}} \Rightarrow 1]$. Adversary \mathcal{A} outputs 1 exactly when $D - R_0 + R'_0 \equiv D' \pmod{b}$. Thus $\Pr[\text{PRP0}^{\mathcal{A}} \Rightarrow 1] = \frac{1}{a(b-1)} + \frac{a-1}{ab}$. We next analyze $\Pr[\text{PRP1}^{\mathcal{A}} \Rightarrow 1]$. Let $d = r/2$. This is the number of times a value R_i is assigned in FE2 for $i > 0$ and even. (Refer to Figure 3.) Let Z_1, \dots, Z_r be the outputs of the round function F for rounds 1 to r (respectively) when evaluating FE2 on point $b \cdot L_0 + R_0$ in response to \mathcal{A} 's first query. Similarly let Z'_1, \dots, Z'_r be the outputs of the round function F for rounds 1 to r (respectively) when evaluating FE2 on point $b \cdot L'_0 + R'_0$ in response to \mathcal{A} 's second query. Consider the situation in which $Z_i = Z'_i$ for all $i > 0$ and i even. This occurs with probability at least a^{-d} . (This is true because the inputs to each of the relevant d round function applications will collide with probability $1/a$.) Then in this case it holds with probability one that $D - R_0 \equiv D' - R'_0 \pmod{b}$ since

$$D \equiv R_r \equiv R_0 + \sum_{\substack{i \leq r \\ i \text{ even}}} Z_i \pmod{b} \quad \text{and} \quad D' \equiv R'_r \equiv R'_0 + \sum_{\substack{i \leq r \\ i \text{ even}}} Z'_i \pmod{b}.$$

Therefore we have $\Pr[\text{PRP1}^{\mathcal{A}} \Rightarrow 1] \geq a^{-d}$. Combining this with the upper bound on $\Pr[\text{PRP0}^{\mathcal{A}} \Rightarrow 1]$ given above we get that the PRP advantage of \mathcal{A} is

$$\mathbf{Adv}_{\text{FE2}}^{\text{PRP}}(\mathcal{A}) \geq \frac{1}{a^d} - \frac{1}{a(b-1)} - \frac{a-1}{ab}.$$

For certain values of a, b, r this is large. Say $r = 7$, $a = 2$ and b is large. Then \mathcal{A} 's advantage is $1/3 - 1/(2b-2) - 1/(2b)$.

MESSAGE RECOVERY ATTACK. We can extend the attack above to enable message recovery. The distinguishing attack establishes a relationship $D - R_0 \equiv D' - R'_0 \pmod{b}$ for distinct messages with high probability. If R_0 is unknown, one can recover it if D, D' , and R'_0 are known. This requires a

<p>Initialize</p> <p>$i \leftarrow 0$</p> <p>Enc(N, T, X)</p> <p>$i \leftarrow i + 1; (a, b) \leftarrow N; L_0^i \leftarrow X \text{ div } b; R_0^i \leftarrow X \text{ mod } b$ $L_1^i \leftarrow R_0^i; Z_1^i \leftarrow \mathcal{S}_a(N, T, 1, R_0^i); R_1^i \leftarrow Z_1^i + L_0^i \text{ mod } a$ $L_2^i \leftarrow R_1^i; Z_2^i \leftarrow \mathcal{S}_b(N, T, 2, R_1^i)$ If $(N, T, R_1^i) = (N^*, T^*, R_1^*)$ then $bad1 \leftarrow \text{true}$; $Z_2^i \leftarrow Z_2^*$ $R_2^i \leftarrow Z_2^i + L_1^i \text{ mod } b$ $L_3^i \leftarrow R_2^i; Z_3^i \leftarrow \mathcal{S}_a(N, T, 3, R_2^i)$ If $(N, T, R_2^i) = (N^*, T^*, R_2^*)$ then $bad2 \leftarrow \text{true}$; $Z_3^i \leftarrow Z_3^*$ $R_3^i \leftarrow Z_3^i + L_2^i \text{ mod } a$ Ret $a \cdot L_3^i + R_3^i$</p> <p>Test(N, T, X)</p> <p>$(a, b) \leftarrow N; N^* \leftarrow N; T \leftarrow T^*; L_0^* \leftarrow X \text{ div } b; R_0^* \leftarrow X \text{ mod } b$ $L_1^* \leftarrow R_0^*; Z_1^* \leftarrow \mathcal{S}_a(N, T, 1, R_0^*); R_1^* \leftarrow Z_1^* + L_0^* \text{ mod } a$ $L_2^* \leftarrow R_1^*; Z_2^* \xleftarrow{\mathcal{S}} \mathbb{Z}_b; R_2^* \leftarrow Z_2^* + L_1^* \text{ mod } b$ $L_3^* \leftarrow R_2^*; Z_3^* \xleftarrow{\mathcal{S}} \mathbb{Z}_a; R_3^* \leftarrow Z_3^* + L_2^* \text{ mod } a$ Ret $a \cdot L_3^* + R_3^*$</p>	<p>Game G0, G1</p>
--	--

Fig. 6. Games used in the proof of Theorem 3. Each function \mathcal{S}_x maps every input to an independently and uniformly sampled point in \mathbb{Z}_x .

single known-plaintext and its associated ciphertext, which will have the desired collisions with the unknown plaintext with probability a^{-d} . From the known plaintext, ciphertext pair one can recover the unknown plaintext portion R_0 . Then L_0 can be guessed (with probability of success $1/a$) or the adversary can try all a potential values of L_0 by appropriate queries to **Enc**. The attack then runs in time that to compute R_0 (two modular additions), uses at most a **Enc** queries, and succeeds in recovering the full plaintext with probability at least a^{-d} .

NON-ADAPTIVE SPI SECURITY. We prove the following theorem, which is analogous to Theorem 2. It establishes non-adaptive SPI security of FE2 up to $q \approx a$ when using just 3 rounds. Note that it is easy to give attackers against 2 rounds, and so this is the minimal number of rounds for which one can expect security.

Theorem 3. *Let the FE2 FPE scheme be parameterized by $F: \mathcal{K} \times \mathcal{N} \times \mathcal{T} \times \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ and $r: \mathcal{N} \rightarrow \mathbb{N}$ where $r(N) = 3$ for all $N \in \mathcal{N}$. Let \mathcal{A} be a non-adaptive SPI adversary that makes q encrypt queries with formats $N = (a, b)$ with $2 \leq a \leq b$. Let a_m be the maximal value such that for all formats $N = (a, b)$ queried by \mathcal{A} it holds that $a_m \leq a$. Then there exists a PRF adversary \mathcal{B} such that*

$$\text{Adv}_{\text{FE2}}^{\text{spi}}(\mathcal{A}) \leq 2 \cdot \text{Adv}_F^{\text{prf}}(\mathcal{B}) + \frac{2q}{a_m}.$$

\mathcal{B} runs in time that of \mathcal{A} and makes $3q$ queries.

Proof. We assume without loss of generality that \mathcal{A} never makes a query for which \perp is returned. Our proof proceeds in steps, using intermediate games G0 and G1 with behaviors between that of SPII_{FE2} and SPI0_{FE2} . Game G0 (shown in Figure 5, boxed statements included) is the same

as SPI1_{FE2} except that F is replaced by random functions. Game G1 (shown in Figure 5, boxed statements excluded) is the same as G0 except that the random values used in responding to the **Test** query are no longer used when responding to **Enc** queries. The result of this change is that that **Test** responds with a random point independent of all handling of **Enc** queries.

We have that

$$\begin{aligned} \text{Adv}_{\text{FE2}}^{\text{spi}}(\mathcal{A}) &\leq \Pr[\text{SPI1}_{\text{FE2}}^{\mathcal{A}} \Rightarrow \text{true}] - \Pr[\text{G0}^{\mathcal{A}} \Rightarrow 1] \\ &\quad + \Pr[\text{G0}^{\mathcal{A}} \Rightarrow 1] - \Pr[\text{G1}^{\mathcal{A}} \Rightarrow 1] \\ &\quad + \Pr[\text{G1}^{\mathcal{A}} \Rightarrow 1] - \Pr[\text{SPI1}_{\text{FE2}}^{\mathcal{A}} \Rightarrow \text{FALSE}]. \end{aligned}$$

A standard reduction bounds the first and last differences each by $\text{Adv}_F^{\text{prf}}(\mathcal{B})$ for an adversary \mathcal{B} using the resources indicated in the theorem statement. To bound the second difference, we use that game G0 and G1 are identical-until-*bad* with the fundamental lemma of game playing to get that

$$\Pr[\text{G0}^{\mathcal{A}} \Rightarrow 1] - \Pr[\text{G1}^{\mathcal{A}} \Rightarrow 1] \leq \Pr[\text{G1}^{\mathcal{A}} \text{ sets } \textit{bad}]$$

where “G1^A sets *bad*” is the event that either *bad1* or *bad2* are set during execution of G1^A.

To conclude, then, we must bound the probability that *bad1* or *bad2* is set in game G1. We use a union bound to treat each case separately, starting with *bad2* which is easier. Recall that we are in the non-adaptive setting. The random choice of Z_2^* is therefore independent of all other values R_2^i for $i \in [1..q]$. Thus $\Pr[\text{G1}^{\mathcal{A}} \text{ sets } \textit{bad2}] \leq q/b_m$.

For *bad1*, consider a particular $i \in [1..q]$. Suppose that $R_0^i = R_0^*$, meaning also that $Z_1^i = Z_1^*$. But then having $R_2^i = R_2^*$ implies that $L_0^i = L_0^*$, yielding a contradiction since \mathcal{A} does not make such queries. Now consider the case that $R_0^i \neq R_0^*$. The probability that $R_1^i = R_1^*$ is at most $1/a_m$ due to the choice of Z_1^* (and because of non-adaptivity). A union bound over all i gives $\Pr[\text{G1}^{\mathcal{A}} \text{ sets } \textit{bad1}] \leq q/a_m$. ■

9 Cycle Walking Doesn't Give Rise to Timing Attacks

Suppose $E: \mathcal{K} \times \mathcal{N} \times \mathcal{T} \times \mathcal{X} \rightarrow \mathcal{X} \cup \{\perp\}$ is an FPE scheme and $\mathcal{X}_N \subseteq \bar{\mathcal{X}}_N$ for all $N \in \mathcal{N}$. Suppose FPE scheme $\mathcal{E}: \mathcal{K} \times \mathcal{N} \times \mathcal{T} \times \bar{\mathcal{X}} \rightarrow \bar{\mathcal{X}} \cup \{\perp\}$ is obtained from FPE scheme E by cycle walking, a folklore technique used, in a similar context, in [4]. Specifically, suppose we define

```
algorithm  $\mathcal{E}_K^{N,T}(X)$ 
 $c \leftarrow 0$ 
repeat  $c \leftarrow c + 1$ ;  $X \leftarrow E_K^{N,T}(X)$  until  $(X \in \mathcal{X}_N)$ 
ret  $X$ 
```

We refer to the value c at the end of the computation as the *cycle-length* associated to $\mathcal{E}_K^{N,T}(X)$. So far, our security models have provided the adversary with an oracle $\mathcal{E}_K(\cdot, \cdot, \cdot)$, for a hidden key K . But in implementations, it would not be surprising to have c leaked, as well as $\mathcal{E}_K^{N,T}(X)$, by way of timing information: for example, responses to \mathcal{E}_K should take twice as long if $c = 2$ than if $c = 1$, assuming that E itself runs in data-independent time across \mathcal{X}_N . A more conservative model, then, would provide to the adversary the cycle length c in addition to handing it $\mathcal{E}_K^{N,T}(X)$. In such a case, one wonders if leaking the timing information negatively impacts security. Here we will show that, in a formally specified model, it does not adversely impact PRP security.

Game $\text{Real}_{\mathcal{E}}$ of Figure 7 returns in response to query N, T, X not only $E_K^{N,T}(X)$ but also the associated cycle-length c . Game $\text{Sim}_{\mathcal{E}}$ implements $\text{Enc}(N, T, \cdot)$ as a random permutation over \mathcal{X}_N

<p>Initialize $K \leftarrow \mathcal{K}$</p> <p>Enc(N, T, X) $c \leftarrow 0$ repeat $c \leftarrow c + 1$ $X \leftarrow \mathcal{E}_K^{N,T}(X)$ until ($X \in \mathcal{X}_N$) ret (X, c)</p>	<p>// Game Real\mathcal{E}</p>	<p>Initialize $K \leftarrow \mathcal{K}$</p> <p>Enc(N, T, X) $Y \xleftarrow{\\$} \mathcal{X}_N \setminus R[N, T]; R[N, T] \stackrel{\leftarrow}{\cup} Y$ $c \leftarrow \text{Sample}(N, T)$ $Q[N, T] \leftarrow Q[N, T] + 1$ ret (Y, c)</p> <p>algorithm <i>Sample</i>(N, T) in $\leftarrow \mathcal{X}_N - Q[N, T]$ out $\leftarrow \tilde{\mathcal{X}}_N - \mathcal{X}_N - St[N, T]$ $c \leftarrow 0$ repeat $b \xleftarrow{\\$} \mathbf{B}(\text{in}, \text{out} - c)$ $c \leftarrow c + 1$ until $b = 1$ $St[N, T] \leftarrow St[N, T] + c - 1$ ret c</p>	<p>// Game Sim\mathcal{E}</p>	<p>Enc(N, T, X) $c \leftarrow 0$ repeat $c \leftarrow c + 1$ $X \xleftarrow{\\$} \tilde{\mathcal{X}}_N \setminus \bar{R}[N, T]; \bar{R}[N, T] \stackrel{\leftarrow}{\cup} X$ until ($X \in \mathcal{X}_N$) ret (X, c)</p>	<p>// Game G</p>
--	---	--	--	--	------------------

Fig. 7. Games used in evaluating risk of timing attacks. Integer-valued tables are silently initialized to 0 for each entry. Distribution $\mathbf{B}(\alpha, \beta)$ is Bernoulli with probability of returning 1 being $p = \alpha/(\alpha + \beta)$.

and, in addition, returns for any query a value of c computed by the subroutine call to *Sample*. The computation of *Sample* depends on some global arrays Q and St , and on N, T , but it does not depend on X . Game *Sample* thus provides a reference experiment defining a random permutation with an *irrelevant* side channel. The notation $b \xleftarrow{\$} \mathbf{B}(\alpha, \beta)$ means that the bit b is chosen randomly with bias $\Pr[b = 1] = \frac{\alpha}{\alpha + \beta}$. The following says that two games we have described are indistinguishable if \mathcal{E} is a good PRP.

Theorem 4. *Let $E: \mathcal{K} \times \mathcal{N} \times \mathcal{T} \times \mathcal{X} \rightarrow \mathcal{X} \cup \{\perp\}$ be an FPE scheme. Suppose $\mathcal{X}_N \subseteq \tilde{\mathcal{X}}_N$ for all $N \in \mathcal{N}$ and let $\mathcal{E}: \mathcal{K} \times \mathcal{N} \times \mathcal{T} \times \tilde{\mathcal{X}} \rightarrow \tilde{\mathcal{X}} \cup \{\perp\}$ be obtained from E by cycle-walking. Then for any \mathcal{A} making $q_{\mathcal{A}}$ queries and having running time $T_{\mathcal{A}}$ there is an adversary \mathcal{B} such that*

$$\Pr[\text{Real}_{\mathcal{E}}^{\mathcal{A}} \Rightarrow 1] - \Pr[\text{Sim}_{\mathcal{E}}^{\mathcal{A}} \Rightarrow 1] \leq \mathbf{Adv}_E^{\text{PRP}}(\mathcal{B}).$$

Let $\lambda = \max_{N \in \mathcal{N}} (|\tilde{\mathcal{X}}_N| - q_{\mathcal{A}}) / (|\mathcal{X}_N| - q_{\mathcal{A}})$, or 1 if this is less than 1. Then the expected number of queries made by \mathcal{B} and the expected running time are at most $\lambda q_{\mathcal{A}}$ and $\lambda T_{\mathcal{A}}$ respectively. \square

Proof. Consider game G in Figure 7. We design \mathcal{B} so that

$$\Pr[\text{Real}_{\mathcal{E}}^{\mathcal{A}} \Rightarrow 1] - \Pr[\text{G}^{\mathcal{A}} \Rightarrow 1] = \mathbf{Adv}_E^{\text{PRP}}(\mathcal{B}). \quad (3)$$

Adversary \mathcal{B} runs \mathcal{A} . When \mathcal{A} makes a query N, T, X to its **Enc** oracle, adversary \mathcal{B} responds, using its own **Enc** oracle, via

```

 $c \leftarrow 0$ 
repeat
   $c \leftarrow c + 1; X \leftarrow \mathbf{Enc}(N, T, X)$ 
until ( $X \in \mathcal{X}_N$ )
ret ( $X, c$ )

```

Equation (3) is clear. The proof is concluded by noting that

$$\Pr[\text{G}^{\mathcal{A}} \Rightarrow 1] = \Pr[\text{Sim}_E^{\mathcal{A}} \Rightarrow 1]. \blacksquare$$

Acknowledgments

Rogaway thanks Terence Spies for many useful discussions, and for sparking his interest in this topic. Rogaway and Stegers were supported by NSF grant CNS 0904380. Bellare and Ristenpart thank Clay Mueller, Lance Nakamura and Joachim Vance for useful discussions, feedback, corrections and support.

References

1. M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. *Advances in Cryptology – EUROCRYPT '06*, LNCS vol. 4004, pp. 409–426, Springer, 2006.
2. M. Bellare, M. Fischlin, A. O’Neill, and T. Ristenpart. Deterministic encryption: definitional equivalences and constructions without random oracles. *Advances in Cryptology – CRYPTO '08*, LNCS vol. 5157, pp. 360–378, Springer, 2008.
3. M. Bellare, T. Ristenpart, P. Rogaway, and T. Stegers. Format-preserving encryption. Full version of this paper. 2009.
4. J. Black and P. Rogaway. Ciphers with arbitrary finite domains. *Topics in Cryptology – CT-RSA '02*, LNCS vol. 2271, Springer, pp. 114–130, 2002.
5. A. Boldyreva, S. Fehr, and A. O’Neill. On notions of security for deterministic encryption, and efficient constructions without random oracles. *Advances in Cryptology – CRYPTO '08*, LNCS vol. 5157, pp. 335–359, Springer, 2008.
6. M. Brightwell and H. Smith. Using datatype-preserving encryption to enhance data warehouse security. *20th NISSC Proceedings*, pp. 141–149, 1997. Available at <http://csrc.nist.gov/nissc/1997>.
7. R. Buble, M. Dyer, C. Greenhill, and M. Jerrum. On approximately counting colorings of small degree graphs. *SIAM J. Computing*, 29(2), pp. 387–400, 1999.
8. C. Colbourn, R. Day, and L. Nel. Unranking and ranking spanning trees of a graph. *Journal of Algorithms*, 10(2), pp. 271–286, 1989.
9. T. Cover. Enumerative source encoding, *IEEE Transactions on Information Theory*, 19(1), pp. 73–77, 1977.
10. A. Desai and S. Miner. Concrete security characterizations of PRFs and PRPs: reductions and applications. *Advances in Cryptology – ASIACRYPT '00*, LNCS vol. 1976, pp. 503–516, Springer, 2000.
11. M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, 1979.
12. A. Goldberg and M. Sipser. Compression and Ranking. *17th Annual ACM Symposium on the Theory of Computing (STOC '85)*, ACM Press, pp. 440–448, 1985.
13. O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. *Journal of the ACM*, 33(4), pp. 792–807, 1986.
14. S. Goldwasser and S. Micali. Probabilistic encryption. *J. Comput. Syst. Sci.*, 28(2), pp. 270–299, 1984.
15. J. Hopcroft and J. Ullman. *Formal Languages and their Relation to Automata*. Addison-Wesley, 1969.
16. M. Jerrum. A very simple algorithm for estimating the number of k -colorings of a low-degree graph. *Random Structures and Algorithms*, 7(2), pp. 157–165, 1995.
17. M. Jerrum, A. Sinclair, and E. Vigoda. A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries. *Journal of the ACM*, 51(4), pp. 671–697, 2004.
18. D. Knuth. *The Art of Computer Programming, Vol. 2, Seminumerical Algorithms, 3rd ed.*, Addison-Wesley, 1997.
19. J. Liebehenschel. Ranking and unranking of a generalized Dyck language and the application to the generation of random trees, *Séminaire Lotharingien de Combinatoire*, 43, 2000.
20. ISO/IEC 7812-1:2006. Identification cards – Identification of issuers – Part 1: Numbering system.
21. P. Kelsen. Ranking and unranking trees using regular reductions. *STACS 1996*, LNCS vol. 1046, pp. 581–592, Springer, 1996.
22. M. Liskov, R. Rivest, and D. Wagner. Tweakable block ciphers. *Advances in Cryptology – CRYPTO 2002*, LNCS vol. 2442, Springer, pp. 31–46, 2002.
23. M. Luby and C. Rackoff. How to construct pseudorandom permutations from pseudorandom functions. *SIAM Journal of Computing*, vol. 17, no. 2, pp. 373–386, 1988.
24. S. Lucks. Faster Luby-Rackoff ciphers. *Fast Software Encryption 1996*, LNCS vol. 1039, Springer, pp. 189–203, 1996.
25. E. Mäkinen. Ranking and unranking left Szilard languages. Report A-1997-2, Department of Computer Science, University of Tampere, 1997.
26. U. Maurer and K. Pietrzak. The security of many-round Luby-Rackoff pseudo-random permutations. *Advances in Cryptology – EUROCRYPT '03*, LNCS vol. 2656, pp. 544–561, Springer, 2003.

27. B. Morris, P. Rogaway, and T. Stegers. How to encipher messages on a small domain: deterministic encryption and the Thorp shuffle. *Advances in Cryptology – CRYPTO ’09*, LNCS vol. 5677, Springer, 2009.
28. M. Naor and O. Reingold. On the construction of pseudorandom permutations: Luby-Rackoff revisited. *Journal of Cryptology*, 12(1), pp. 29–66, 1999.
29. National Bureau of Standards. FIPS PUB 74. Guidelines for Implementing and Using the NBS Data Encryption Standard. April 1, 1981.
30. J. Patarin. New results on pseudorandom permutation generators based on the DES Scheme. *Advances in Cryptology – CRYPTO ’91*, LNCS vol. 576, Springer, pp. 301–312, 1991.
31. J. Patarin. Generic attacks on Feistel schemes. *Advances in Cryptology – ASIACRYPT ’01*, LNCS vol. 2248, Springer, pp. 222–238, 2001.
32. J. Patarin. Luby-Rackoff: 7 rounds are enough for $2^{n(1-\epsilon)}$ security. *Advances in Cryptology – CRYPTO ’03*, LNCS vol. 2729, Springer, pp. 513–529, 2003.
33. J. Patarin. Security of random Feistel schemes with 5 or more rounds. *Advances in Cryptology – CRYPTO ’04*, LNCS vol. 3152, Springer, pp. 106–122, 2004.
34. J. Patarin, V. Nachev, and C. Berbain. Generic attacks on unbalanced Feistel schemes with contracting functions. *Advances in Cryptology – ASIACRYPT ’06*, LNCS vol. 4284, Springer, pp. 396–411, 2006.
35. S. Patel, Z. Ramzan, and G. Sundaram. Efficient constructions of variable-input-length block ciphers. *Selected Areas in Cryptography 2004*, LNCS vol. 3357, pp. 326–340, 2004.
36. PCI Security Standards Council. Payment Card Industry Data Security Standard Version 1.2. Available at https://www.pcisecuritystandards.org/security_standards/pci_dss.shtml
37. E. Petrank and C. Rackoff. CBC MAC for real-time data sources. *J. of Cryptology*, 13(3), pp. 315–338, 2000.
38. B. Schneier and J. Kelsey. Unbalanced Feistel networks and block cipher design. *Fast Software Encryption 1996*, LNCS vol. 1039, Springer, pp. 121–144, 1996.
39. R. Schroepel. Personal communication, approximately 2001.
40. M. Sipser. *Introduction to the Theory of Computation*, 2nd ed. Thomson Press, 2006.
41. T. Spies. Format preserving encryption. Unpublished white paper, available at www.voltage.com. See also: Format preserving encryption: www.voltage.com. *Database and Network Journal*, Dec. 2008.
42. T. Spies. Personal communications, Feb 2009.
43. T. Spies. Feistel finite set encryption mode. Manuscript, posted on NIST’s website on February 6, 2008. Available at <http://csrc.nist.gov/groups/ST/toolkit/BCM/documents/proposedmodes/ffsem/ffsem-spec.pdf>
44. L. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8, pp. 189–201, 1979.

A Building Efficient Round Functions

In this section we detail some sample methods for instantiating an efficient tweakable round function $F: \mathcal{K} \times D \rightarrow \mathbb{N}$ for $D = \mathcal{T} \times \mathbb{N}^3$ from a blockcipher $E: \mathcal{K} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$. We first describe building round functions that output n -bit strings, and then we consider what happens when one take the result mod b .

VIL PRF. A simple and secure approach is to first build a variable-input-length PRF F from E using any of the many well-known construction, such as CMAC. Then map the format N , tweak T , and round number i to a bit string via some canonical encoding, apply F , and interpret the resulting n -bit string as an integer in \mathbb{Z}_{2^n} . We detail two particular realizations of this approach below.

PREFIX-FREE CBC-MAC. Recall that CBC-MAC is secure for variable-input-length inputs if a prefix-free encoding of messages is used [37], so we can encode N, T, i in a prefix-free manner and apply the CBC-MAC. Let $\ell \geq \lceil \log N \rceil$ for any N to be used, and let $L \geq \lceil \log |T| \rceil$ for any T to be used. In practice there will always be such maximums. Define

$$\text{pad}(N, T) = \langle N \rangle_\ell \parallel \langle |T| \rangle_L \parallel T \parallel 0^p$$

where p is the minimum number of bits needed to ensure that $\ell + L + |T| + p$ is a multiple of n and StN (string to number) maps a string Y to the integer y , $0 \leq y < 2^{|Y|}$, that it represents as an unsigned binary value (for example, $\text{StN}(000011) = 3$). The round function $\text{RdF1}: \mathcal{K} \times D \rightarrow \mathbb{N}$ is then defined by

$$\text{RdF1}_K(N, T, i, X) = \text{StN}(\text{CBC-MAC}_K(\text{pad}(N, T) \parallel \langle i \rangle_8 \parallel \langle X \rangle_{n-8}))$$

for any $K \in \mathcal{K}$, $T \in \mathcal{T}$, and $X \in \mathbb{Z}_{2^{n-8}}$. Since typically one will have $n = 64$ or $n = 128$, that $X \in \mathbb{Z}_{2^{n-8}}$ is not a serious restriction, nor is it to assume that $i < 256$.

The encoding ensures that one can do efficient precomputation, since N, T do not change between rounds. One therefore precomputes $\tau \leftarrow \text{CBC-MAC}_K(\text{pad}(N, T))$ first. During round i , applying the PRF to number X can then be accomplished via a single call to E , $Z \leftarrow \text{StN}(E_K(\tau \oplus (\langle i \rangle_8 \parallel \langle X \rangle_{n-8})))$.

REKEYING BY TWEAKS. We modify the above construction to give one that uses rekeying. Let $E: \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a blockcipher with keylength $k \geq n$. Round function $\text{RdF2}: \mathcal{K} \times D \rightarrow \mathbb{N}$ is defined by

$$\text{RdF2}_K(T, \ell, i, X) = \text{StN}(E_{K'}(\langle i \rangle_8 \parallel \langle X \rangle_{n-8})) \text{ where } K' = \text{CBC-MAC}_K(\text{pad}(N, T))[1..k].$$

That is, the CBC MAC is applied to N and T to derive a new key K' , and this new key is used for a final application of E on the round number and message. Again precomputation is straightforward, so that each round requires a single blockcipher call.

EFFECT OF MODULAR ARITHMETIC. Our PRFs are built from traditional ones that output bit strings, and are proven to be indistinguishable from random functions outputting bit strings. But we will use these PRFs to output numbers that are taken modulo numbers smaller than 2^n within FE1 and FE2. Here we discuss how security is affected by this extra mod operation.

We begin by asking the following. Consider, on the one hand, the uniform distribution on \mathbb{Z}_M . Consider, on the other hand, the distribution on \mathbb{Z}_M that is obtained by picking a random point x in \mathbb{Z}_N and returning $x \bmod M$. What is the statistical difference between these distributions? To answer this, let INTDIV denote the integer division algorithm, which on inputs N, M returns a quotient q and remainder r satisfying $N = Mq + r$ and $0 \leq r < M$. Then, we claim the following.

Lemma 1. *Let $N \geq M \geq 1$ be integers, and let $(q, r) \leftarrow \text{INTDIV}(N, M)$. For $z \in \mathbb{Z}_M$ let*

$$P_{N,M}(z) = \Pr[x \bmod M = z : x \leftarrow \mathbb{Z}_N].$$

Then for any $z \in \mathbb{Z}_M$,

$$P_{N,M}(z) = \begin{cases} \frac{q+1}{N} & \text{if } 0 \leq z < r \\ \frac{q}{N} & \text{if } r \leq z < M. \end{cases}$$

Proof (Lemma 1). Let the random variable X be uniformly distributed over \mathbb{Z}_N . Then

$$\begin{aligned} P_{N,M}(z) &= \Pr[X \bmod M = z] \\ &= \Pr[X < Mq] \cdot \Pr[X \bmod M = z \mid X < Mq] \\ &\quad + \Pr[Mq \leq X < N] \cdot \Pr[X \bmod M = z \mid Mq \leq X < N] \\ &= \frac{Mq}{N} \cdot \frac{1}{M} + \frac{N - Mq}{N} \cdot \begin{cases} \frac{1}{N - Mq} & \text{if } 0 \leq z < N - Mq \\ 0 & \text{if } N - Mq \leq z < M. \end{cases} \\ &= \frac{q}{N} + \frac{r}{N} \cdot \begin{cases} \frac{1}{r} & \text{if } 0 \leq z < r \\ 0 & \text{if } r \leq z < M. \end{cases} \end{aligned}$$

Simplifying yields the claimed equation. **■**

As a result of the above, the statistical distance between the uniform distribution on \mathbb{Z}_M and the distribution obtained by picking a random point x in \mathbb{Z}_N and returning $x \bmod M$ is

$$\frac{1}{2} \sum_{z=0}^{r-1} \left| \frac{q+1}{N} - \frac{1}{M} \right| + \frac{1}{2} \sum_{z=r}^{M-1} \left| \frac{q}{N} - \frac{1}{M} \right| = \frac{r(M-r)}{NM} \leq \frac{1}{4} \frac{M}{N}.$$

Suppose that the maximum number of digits in a plaintext is 20. In this case, we have $M = 10^{10}$, so the above statistical distance is at most $10^{10}/2^{64} \approx 2^{-31}$. This is reasonably small, indicating that the mod operation does not dramatically affect the distribution.