

Sometimes-Recurse Shuffle

Almost-Random Permutations in Logarithmic Expected Time

Ben Morris¹ and Phillip Rogaway²

¹ Dept. of Mathematics, University of California, Davis, USA

² Dept. of Computer Science, University of California, Davis, USA

Abstract. We describe a security-preserving construction of a random permutation of domain size N from a random function, the construction tolerating adversaries asking all N plaintexts, yet employing just $\Theta(\lg N)$ calls, on average, to the one-bit-output random function. The approach is based on card shuffling. The basic idea is to use the *sometimes-recurse* transformation: lightly shuffle the deck (with some other shuffle), cut the deck, and then recursively shuffle one of the two halves. Our work builds on a recent paper of Ristenpart and Yilek.

Keywords: Card shuffling, format-preserving encryption, PRF-to-PRP conversion, mix-and-cut shuffle, pseudorandom permutations, sometimes-recurse shuffle, swap-or-not shuffle.

1 Introduction

FORMAT-PRESERVING ENCRYPTION. Suppose you are given a blockcipher, say AES, and want to use it to efficiently construct a cipher on a smaller domain, say the set of $N = 10^{16}$ sixteen-digit credit card numbers. You could, for example, use AES as the round function for several rounds of a Feistel network, the approach taken by emerging standards [1, 7]. But information-theoretic security will vanish by the time the adversary asks \sqrt{N} queries, which is a problem on small-sized domains. (It is a problem from the point of view of having a satisfying provable-security claim; likely it is not a problem with respect to their being a feasible attack.) Alternatively, you could precompute a random permutation on N points, but spending $\Omega(N)$ time in computation will become undesirable before \sqrt{N} adversarial queries becomes infeasible.

This paper provides a new solution to this problem of *format-preserving encryption*, where we aim to build ciphers with an arbitrary finite domain [3, 4, 8, 5], frequently $[N] = \{0, 1, \dots, N-1\}$ for some N . Our solution lets you encipher a sixteen-digit credit card with about 1000 expected AES calls, getting an essentially ideal provable-security claim. (One thousand AES calls comes to about 80K clock cycles, or 25 μ sec, on a recent Intel processor.) In particular, the adversary can ask any number of queries—including all N of them—and its advantage in distinguishing the constructed cipher from a random permutation

will be insignificantly more than its ability to break the underlying primitive (in our example, AES) with a like number of queries.

Cast in more general language, this paper is about constructing ciphers—meaning information theoretic or complexity theoretic PRPs—on an arbitrary domain $[N]$, starting from a PRF. (If starting from AES, only a single bit of each 128-bit output will be used. A random permutation on 128 bits that gets truncated to a single bit is extremely close to a random function [2].) As in other recent work [14, 11, 9], our ideas are motivated by card shuffling and its cryptographic interpretation. This connection was first observed by Naor [15, p. 62], [17, p. 17], who explained that when a card shuffle is *oblivious*—meaning that you can trace the trajectory of a card without attending to the trajectories of *other* cards in the deck—then it determines a computationally plausible cipher. We will move back and forth between the language of encryption and that of card shuffling: a PRP/cipher is a shuffle; a plaintext x encrypts to ciphertext y if the card initially at position x ends up at position y ; the PRP’s key is the randomness underlying the shuffle.

THE SWAP-OR-NOT AND MIX-AND-CUT SHUFFLES. Hoang, Morris, and Rogaway describe an oblivious shuffle well-suited for enciphering on a small domain [11]. In the binary-string setting ($N = 2^n$), round i of their *swap-or-not* shuffle employs a random string $K_i \in \{0, 1\}^n$ and replaces X by $K_i \oplus X$ if $F(i, \hat{X}) = 1$, where F is a random function to bits and $\hat{X} = \max(X, X \oplus K_i)$. If $F(i, \hat{X}) = 0$, then X is left alone. After all rounds are complete, the final value of input X is the result of the shuffle. The authors show that $O(\lg N)$ rounds suffice to get a cipher that will look uniform to an adversary that makes $q < (1 - \epsilon)N$ queries. But as q approaches N , one would need more and more rounds and, eventually, one gets a non-result.

Ristenpart and Yilek were looking for practical ways to tolerate adversaries asking all $q = N$ queries, a goal they called *full* security. Assume again that we want to shuffle $N = 2^n$ cards. Then Ristenpart and Yilek’s *Icicle* construction first mixes the cards using some given (we’ll call it the *inner*) shuffle. Then they cut the deck into two piles and recursively shuffle each. The authors explain that if the inner shuffle is a good *pseudorandom separator* (PRS), then the constructed shuffle will achieve full security. A shuffle is a good PRS if, after shuffling, the (unordered) *set* of cards ending up in each of the two piles is indistinguishable from a uniform partitioning of the cards into two equal-sized sets.

Ristenpart and Yilek apply the *Icicle* construction to the *swap-or-not* shuffle, a combination they call *mix-and-cut*. The combination achieves full security in $\Theta(\lg^2 N)$ rounds. When the underlying round function is realized by an AES call, *mix-and-cut* constructs a cipher on N points, achieving full security, with $\Theta(\lg^2 N)$ AES calls. While full security is directly achieved by other oblivious shuffles [9, 13, 18], *mix-and-cut* would seem to be much faster.

CONTRIBUTIONS. We reconceptualize what is going on in Ristenpart and Yilek’s *mix-and-cut*. Instead of thinking of the underlying transformation as turning a PRS into a PRP, we think of it as turning a mediocre PRP into a better one.

If the inner shuffle is good enough to mix half the cards—in the inverse shuffle, any $N/2$ cards end up in almost uniform positions—then the constructed shuffle will achieve full security.

After this shift in viewpoint, we make a simple change to mix-and-cut that dramatically improves its speed. As before, one begins by applying the inner shuffle to the N cards. Then one splits the deck and recursively shuffles *one* (rather than both) of the two halves. Using swap-or-not (SN) for the inner shuffle we now get a PRP over $[N]$ enjoying full security and computable in $\Theta(\lg N)$ expected time. We call the SN-based construction **SR**, for *sometimes-recurse*. The underlying transformation we call **SR** (in bold font).

Our definitions and results apply to an arbitrary domain size N (it need not be a power of two). We emphasize that the adversary may query *all* points in the domain. We give numerical examples to illustrate that the improvement over mix-and-cut is large. We also explain why, with SR, having the running time depend on the key and plaintext does *not* give rise to side-channel attacks. Finally, we explain how to cheaply tweak [12] the construction, degrading neither the run-time nor the security bound compared to the untweaked counterpart. (Ristenpart and Yilek likewise support tweaks [16], but their quantitative bounds give up more, and each round key needs to depend on the tweak.)

ADDITIONAL RELATED WORK. Granboulan and Pornin [9] also give a shuffle achieving full security, and Ristenpart and Yilek’s paper [16] can likewise be seen as building on it, reconceptualizing their work as the application of the Icicle construction to a particular PRS. But the chosen PRS is computationally expensive to realize, involving extensive use of arbitrary-precision floating-point arithmetic to do approximate sampling from a hypergeometric distribution. The mix-and-cut and sometimes-recurse shuffles are much more practical.

For realistic domain sizes N , both mix-and-cut and sometimes-recurse are also much faster than the method of Stefanov and Shi [18], which spends $\tilde{\Theta}(N)$ time to preprocess the key into a table of size $\tilde{\Theta}(\sqrt{N})$ that supports $\tilde{\Theta}(\sqrt{N})$ -time evaluation of the constructed cipher.

2 Preliminaries

SHUFFLES AS FORMAL OBJECTS. A shuffle SH_N on $N \geq 1$ cards is a distribution on permutations of $[N]$. We are only interested in distributions that can be described by efficient probabilistic algorithms, so one can alternatively consider a shuffle SH_N on N cards to be a probabilistic algorithm that bijectively maps each $x \in [N]$ to a value $\text{SH}_N(x) \in [N]$. The algorithm may be thought of as keyed, the key coinciding with the algorithm’s coins. A shuffle SH (now on an arbitrary number of cards) is a family of shuffles on N cards, one for each number $N \geq 1$. One can regard SH as taking two arguments, with $\text{SH}_N(x) \in [N]$ being the image of $x \in [N]$ under the random permutation on $[N]$. If we write $\text{SH}(x)$ for some shuffle SH we mean $\text{SH}_N(x)$ for some understood N .

As suggested already, we may refer to points $x \in [N]$ as *cards*. We then think of $\text{SH}_N(x)$ as the location that card x landed at following the shuffle of

these N cards. Locations are indexed 0 to $N - 1$. We think of 0 as the leftmost position and $N - 1$ as the rightmost position. If we shuffle a deck with an even number N of cards, the lefthand pile would be positions $\{0, \dots, N/2 - 1\}$ and the righthand pile would be positions $\{N/2, \dots, N - 1\}$. The card that landed at position $y \in [N]$ is card $\text{SH}_N^{-1}(y)$.

We are interested in operators that transform one shuffle into another. Such an operator \mathbf{OP} takes a shuffle SH and produces a shuffle $\text{SH}' = \mathbf{OP}[\text{SH}]$. The definition of $\text{SH}'_N(x)$ may depend on $\text{SH}_{N'}(x')$ values with $N' \neq N$.

PROBABILITY. For distributions μ and ν on a finite set V , define the total variation distance

$$\|\mu - \nu\| = \frac{1}{2} \sum_{x \in V} |\mu(x) - \nu(x)|.$$

If V_1, \dots, V_k are finite sets and τ is a probability distribution on $V_1 \times \dots \times V_k$, then for l with $0 \leq l \leq k - 1$ define

$$\tau(\cdot \mid x_1, \dots, x_l) = \mathbf{P}(X_{l+1} = \cdot \mid X_1 = x_1, \dots, X_l = x_l),$$

where $(X_1, \dots, X_k) \sim \tau$.

Lemma 1. *Let V_1, \dots, V_n be finite sets and let μ and ν be probability distributions on $V_1 \times \dots \times V_n$. Suppose that $(Z_1, \dots, Z_n) \sim \mu$. Then*

$$\|\mu - \nu\| \leq \sum_{l=0}^{n-1} \mathbf{E}(\|\mu(\cdot \mid Z_1, \dots, Z_l) - \nu(\cdot \mid Z_1, \dots, Z_l)\|).$$

We defer the proof of Lemma 1 to Appendix A. The lemma immediately gives us the following.

Corollary 2. *Suppose that for every l with $1 \leq l \leq n$ there is an $\epsilon_l > 0$ such that for any z_1, z_2, \dots, z_l we have $\|\mu(\cdot \mid z_1, \dots, z_l) - \nu(\cdot \mid z_1, \dots, z_l)\| \leq \epsilon_l$. Then $\|\mu - \nu\| \leq \epsilon_1 + \dots + \epsilon_n$.*

Let us explain part of the utility of this fact. Consider a random permutation π on $\{0, 1, \dots, N - 1\}$, which we view as a random ordering of cards arranged from left to right. Suppose N_1, \dots, N_n are positive integers with $N_1 + N_2 + \dots + N_n = N$. Let Z_1 be the configuration of cards in the rightmost N_1 positions, let Z_2 be the configuration of cards in the N_2 positions to the immediate left of these, and so on. Applying Corollary 2 to (Z_1, \dots, Z_n) shows that if the distribution of the rightmost N_1 cards is within ϵ_1 of uniform, and regardless of the values of these cards the conditional distribution of the N_2 cards to their immediate left is within ϵ_2 of uniform, and so on, then the whole deck is within distance $\epsilon = \epsilon_1 + \epsilon_2 + \dots + \epsilon_n$ of a uniform random permutation.

3 Mix-and-Cut Shuffle

This section reviews and reframes the prior work of Ristenpart and Yilek [16].

The mix-and-cut transformation can be described recursively as follows. Assume we want to shuffle $N = 2^n$ cards. If $N = 1$ then we are done; a single card is already shuffled. Otherwise, to mix-and-cut shuffle $N \geq 2$ cards,

1. shuffle the N cards using some other, inner shuffle; and then
2. cut the deck into two halves (that is, the cards in positions $0, \dots, \frac{N}{2} - 1$ and the cards in positions $\frac{N}{2}, \dots, N - 1$) and, recursively, shuffle each half.

The method can be seen as an operator, \mathbf{MC} , that maps a shuffle SH on a power-of-two number of cards to a shuffle $\text{SH}' = \mathbf{MC}[\text{SH}]$ on the same number of cards. A sufficient condition for SH' to achieve full security is for SH to *lightly shuffle* the deck. Informally, to lightly shuffle the deck means that if one identifies some $N/2$ positions of the deck, then the cards that land in these positions should be nearly uniform, that is, like $N/2$ samples without replacement from the N cards. More formally, we say that SH ε -lightly shuffles if for any $N/2$ positions the distribution of the *unordered set* of cards in those positions is within distance ε of a uniform random subset of cards of size $N/2$. Note that if the shuffle SH is swap-or-not (SN) then it is equivalent to ask that SH itself send $N/2$ cards to something ε -close to uniform, as SN is identical in its forward and backward direction, up to the naming of keys.

Let's consider the speed of \mathbf{MC} with SN as the underlying shuffle, a combination we'll write as $\text{MC} = \mathbf{MC}[\text{SN}]$. First some preliminaries. For a round-parameterized shuffle SH that approaches the uniform distribution, let $\tau_q^r(N)$ be the induced distribution after r rounds on some q distinct cards $(x_1, \dots, x_q) \in [N]^q$ from a deck of size N , and let $\pi_q(N)$ be the distribution of q samples, without replacement, from $[N]$. Let $\Delta_{\text{SH}}(N, q, r) = \|\tau_q^r(N) - \pi_q(N)\|$ be the total variation distance between these two distributions. Hoang, Morris, and Rogaway show that, for the swap-or-not shuffle, SN,

$$\Delta_{\text{SN}}(N, q, r) \leq \frac{2N^{3/2}}{r+2} \left(\frac{q+N}{2N} \right)^{r/2+1} = \Delta_{\text{SN}}^{\text{ub}}(N, q, r). \tag{1}$$

Assuming even N , setting $q = N/2$ in this equation gives

$$\Delta_{\text{SN}}(N, N/2, r) \leq N^{3/2} \left(\frac{3}{4} \right)^{r/2}$$

and so $\Delta_{\text{SN}}(N, N/2, r) \leq \varepsilon$ if

$$\frac{3}{2} \lg N + \frac{r}{2} \lg(3/4) \leq \lg \varepsilon,$$

which occurs if

$$\begin{aligned} r &\geq \frac{\lg \varepsilon - (3/2) \lg N}{(1/2) \lg(3/4)} \\ &\geq 7.23 \lg N - 4.82 \lg \varepsilon \\ &\in \Theta(\lg N - \lg \varepsilon). \end{aligned} \tag{2}$$

Let SH be a round-based shuffle approaching the uniform distribution and let $T_{\text{SH}}(N, q, \varepsilon)$ be the minimum number r such that $\Delta_{\text{SH}}(N, q, r) \leq \varepsilon$. Let $T_{\text{SH}}(N, \varepsilon) = T_{\text{SH}}(N, N, \varepsilon)$ be the time to mix all the cards to within ε . For $\text{MC} = \mathbf{MC}[\text{SN}]$ to mix all $N = 2^n$ cards to within ε it will suffice if we arrange that each invocation of SN mixes half the cards to within ε/n . Assuming this strategy, the total number of needed rounds will be

$$\begin{aligned} T_{\text{MC}}(2^n, \varepsilon) &\leq \sum_{\ell=1}^n T_{\text{SN}}(2^\ell, 2^{\ell-1}, \varepsilon/n) \\ &\leq \sum_{\ell=1}^n (7.23 \ell - 4.82 \lg(\varepsilon/n)) \quad (\text{from (2)}) \\ &\leq 14.46 n^2 + 4.82 n \lg n - 4.82 n \lg \varepsilon \\ &\in \Theta(\lg^2 N - \lg N \lg \varepsilon) \end{aligned}$$

Interpreting, the MC construction can encipher n -bit strings, getting to within any fixed total variation distance ε of uniform, by using $\Theta(n)$ stages of $\Theta(n)$ rounds, so $\Theta(n^2)$ total rounds. The round functions here are assumed uniform and independent. Replacing them by a complexity-theoretic PRF, we are converting a PRF into a PRP on domain $\{0, 1\}^n$ with $\Theta(n^2)$ calls, achieving tight provable security and no limit on the number of adversarial queries.

4 Sometimes-Recurse Shuffle

The SN shuffle has a stronger mixing property than light shuffling: namely, the SN shuffle randomizes the *sequence* of cards in any $N/2$ positions of the deck (as made precise by equation (1)). Therefore, after shuffling the deck with SN and cutting it in half, there is no need to recurse on one of the two halves. Either pile can be declared finished and in the next stage we recursively shuffle only the other pile. Assuming that the first stage brings the distribution of the cards in the rightmost $N/2$ positions to within distance ε_1 of uniform, and the next stage brings the conditional distribution of the cards in the prior $N/4$ positions to within distance ε_2 of uniform, and so on, the final permutation is with distance $\varepsilon_1 + \dots + \varepsilon_n$ of a uniform random permutation, where n is the number of stages. This follows by the remark that immediately followed Corollary 2.

POWER-OF-TWO DOMAINS. The *sometimes-recurse* (**SR**) transform can thus be described as follows. Assume for now that want to shuffle $N = 2^n$ cards. (We will generalize afterward.) If $N = 1$ then we are done; a single card is already shuffled. Otherwise, to **SR** shuffle $N \geq 2$ cards,

1. shuffle the N cards using some other, inner shuffle; and then
2. cut the deck into two halves and, recursively, shuffle the first half.

The method can be seen as an operator, **SR**, that maps a shuffle SH on any power-of-two cards to a shuffle $\text{SH}' = \mathbf{SR}[\text{SH}]$ on any power-of-two cards.

Recasting the method into more cryptographic language, suppose you are given a variable-input-length PRP $E : \mathcal{K} \times \{0, 1\}^* \rightarrow \{0, 1\}^*$. Write $E_K(\cdot)$ for $E(K, \cdot)$. Each $E_K(\cdot)$ is a length-preserving permutation. We construct from E a PRP $E' = \mathbf{SR}[E]$ as follows. First, assert that $E'_K(\epsilon) = \epsilon$, where ϵ is the empty string. Otherwise, let $E'_K(X) = Y$ if $Y = E_K(X) = 1 \parallel Y'$ begins with a 1-bit, and let $E'_K(X) = 0 \parallel E_K(Y')$ if $Y = E_K(X) = 0 \parallel Y'$ begins with a 0-bit.

THE SR TRANSFORMATION. The description above assumes a power-of-two number of cards and an even cut of the deck. The first assumption runs contrary to our intended applications, and dropping this assumption necessitates dropping the second assumption as well. Here then is the **SR** transform stated more broadly. Assume an inner shuffle, **SH**, that can mix an arbitrary number of cards. Let $p : \mathbb{N} \rightarrow \mathbb{N}$, the *split*, be a function with $1 \leq p(N) < N$. We'll sometimes write p_N for $p(N)$. We construct a shuffle $\mathbf{SH}' = \mathbf{SR}_p[\mathbf{SH}]$. Namely, if $N = 1$, we are done; a single card is shuffled. Otherwise,

1. shuffle the N cards using the inner shuffle, **SH**; and then
2. cut the deck into a first pile having p_N cards and a second pile having $q_N = N - p_N$ cards. Recursively, shuffle the first pile.

INITIAL AND GENERATED N -VALUES. A potential point of confusion is that, above, the name “ N ” effectively has two different meanings: it is used for both the *initial* N , call it N_0 , that specifies the domain $[N_0]$ on which we seek to encipher; and it is used as a generic name for any of the N -values that can arise in recursive calls that begin with the initial N . These are the *generated* N -values, a set of numbers $\mathcal{G}_p(N_0) = \mathcal{G}(N_0)$. Note that we count the initial N among the generated N -values $\mathcal{G}_g(N_0)$. As an example, if the initial N is $N_0 = 10^{16}$ and $p_N = \lfloor N/2 \rfloor$, then there are 54 generated N -values, which are $\mathcal{G}_p(10^{16}) = \{10^{16}, 10^{16}/2, 10^{16}/4, \dots, 71, 35, 17, 8, 4, 2, 1\}$. In general, $\mathcal{G}_p(N_0)$ is the set $\{N_0, N_1, \dots, N_n\}$ where $N_i = p(N_{i-1})$ and $N_n = 1$. We call n the number of *stages*.

THE TRANSFORMATION WORKS. Let $q : \mathbb{N} \rightarrow \mathbb{N}$ and let $\varepsilon : \mathbb{N} \rightarrow [0, 1]$ be functions, $1 \leq q(N) \leq N$. We may write $q(N)$ and ε_N for $q(N)$ and $\varepsilon(N)$. Let **SH** be a shuffle that can mix any number of cards. We say that **SH** is (q, ε) -good if for all $N \in \mathbb{N}$, for any distinct $y_1, \dots, y_{q(N)} \in [N]$, the total-variation distance between $(\mathbf{SH}^{-1}(y_1), \dots, \mathbf{SH}^{-1}(y_{q(N)}))$ and the uniform distribution on $q(N)$ distinct points from $[N]$ is at most $\varepsilon(N)$. A shuffle is ε -good if it is (q, ε) -good for $q(N) = N$. We have the following:

Theorem 3. *Let $p, q : \mathbb{N} \rightarrow \mathbb{N}$ and $\varepsilon : \mathbb{N} \rightarrow [0, 1]$ be functions, $p(N) + q(N) = N$, and fix $N_0 \in \mathbb{N}$. Suppose that **SH** is a (q, ε) -good shuffle. Then $\mathbf{SR}_p[\mathbf{SH}]$ is a δ -good shuffle where $\delta = \sum_{N \in \mathcal{G}_g(N_0)} \varepsilon_N$.*

Proof. Consider the indicated shuffle π on domain $[N_0]$. Enumerate the elements of $\mathcal{G}_p(N_0)$ as $\{N_0, N_1, \dots, N_n\}$ where $N_0 > N_1 > \dots > N_n$. The first stage of the shuffle brings the distribution of the rightmost q_{N_0} cards to within a distance

```

10 procedure  $E_{KF}^N(X)$  //invariant:  $X \in [N]$ 
11 if  $N = 1$  then return  $X$  //a single card is already shuffled

20 for  $i \leftarrow 1$  to  $t_N$  do //SN, for  $t_N$ -rounds
21  $X' \leftarrow K_i - X \pmod{N}$  // $X'$  is the "partner" of  $X$ 
22  $\hat{X} \leftarrow \max(X, X')$  //canonical name for  $\{X, X'\}$ 
23 if  $F(i, \hat{X}) = 1$  then  $X \leftarrow X'$  //maybe swap  $X$  and  $X'$ 

30 if  $X < p_N$  then return  $E_{KF}^{p_N}(X)$  //recursively shuffle the first pile
31 if  $X \geq p_N$  then return  $X$  //but second pile is done

```

Fig. 1. Construction $SR = \mathbf{SR}[SN]$. The method enciphers on $[N_0]$ (the initial value of N), each stage (recursive invocation) employing t_N -rounds of SN (lines 20–23). The split values, p_N , are a second parameter on which SR depends. The randomness for SN is determined by $F: \mathbb{N} \times \mathbb{N} \rightarrow \{0, 1\}$ and $K: \mathbb{N} \rightarrow \mathbb{N}$.

ε_{N_0} of uniform. Regardless of the values of these cards the second stage brings the conditional distribution of the preceding q_{N_1} cards to within distance ε_{N_1} of uniform, and so on. Therefore, applying Corollary 2 (as explained in the argument immediately following the statement of Corollary 2) shows that the final permutation is within δ of a uniform random permutation, where $\delta = \varepsilon_{N_0} + \varepsilon_{N_1} + \dots + \varepsilon_{N_n}$. \square

USING SN AS THE INNER SHUFFLE. We'll write SR (no bold) for $\mathbf{SR}[SN]$, the sometimes-recurse transformation applied to the swap-or-not shuffle. The algorithm is shown in Fig. 1, now written out in the manner of a cipher, where the trajectory of a single card X is followed. Of course $SN = SN_t$ depends on the round count and $\mathbf{SR} = \mathbf{SR}_p$ depends on the split, so $SR = SR_{t,p}$ depends on both. The canonical choice for the split p_N is $p_N = \lfloor N/2 \rfloor$; when no mention of p_N is made, this is assumed. There is no default for the round counts t_N ; we must select these values with care.

We proceed to analyze SR, for the canonical split, with the help of Proposition 3 and equation (2). We aim to shuffle N cards to within a target distance ε . Assume we run each stage (that is, each SN shuffle) with t_N adequate to achieve error ε/n for any half, rounded up, of the cards. When N is a power of 2, the expected total number of rounds to encipher a point will then be

$$\begin{aligned}
\mathbf{E}[T_{SR}(N, \varepsilon)] &\leq T_{SN}(N, \frac{N}{2}, \frac{\varepsilon}{\lg N}) + \frac{T_{SN}(\frac{N}{2}, \frac{N}{4}, \frac{\varepsilon}{\lg N})}{2} + \frac{T_{SN}(\frac{N}{4}, \frac{N}{8}, \frac{\varepsilon}{\lg N})}{4} + \dots \\
&\leq 2(7.23 \lg N + 4.82 \lg \lg N - 4.82 \log \varepsilon) \quad \text{from (2)}
\end{aligned}$$

For arbitrary N (not necessarily a power of two), simply replace N by $2N$ in the equation just given to get an upper bound. This is valid because the sequence of generated N -values for N_0 are bounded above by the sequence of generated

N -values for N'_0 the next higher power of two, and, additionally, the bound $\Delta_{\text{SN}}^{\text{ub}}(N, N/2, r)$ is increasing in N . Thus, for any N ,

$$\begin{aligned} \mathbf{E}[T_{\text{SR}}(N, \varepsilon)] &\leq 14.46 \lg N + 4.82 \lg \lg 2N - 4.82 \lg \varepsilon + 14.46 \\ &\in \Theta(\lg N - \lg \varepsilon) \end{aligned} \tag{3}$$

The worst-case number of rounds is similarly bounded. We summarize the result as follows.

Theorem 4. *For any $N \geq 1$ and $\varepsilon \in (0, 1)$, the SR construction enciphers points on $[N]$ in $\Theta(\lg N - \lg \varepsilon)$ expected rounds and $\Theta(\lg^2 - \lg N \lg \varepsilon)$ rounds in the worst case. No adversary can distinguish the construction from a uniform permutation on $[N]$ with advantage exceeding ε . This assumes uniformly random round keys and round functions for SN, appropriate round counts t_N , and the canonical split.*

As a numerical example, equation (3) gives $\mathbf{E}[T_{\text{SR}}(10^{16}, 10^{-10})] \leq 1159$. In the next section we will do better than this—but not by much—by doing calculations directly from equation (1) and by partitioning the error ε so as to give a larger portion to earlier (that is, larger) generated N .

5 Parameter Optimization

ROUND COUNTS. Let us continue to assume the canonical split of $p_N = \lfloor N/2 \rfloor$ and look at the optimization of round counts t_N under this assumption.

In speaking below of the number p of nontrivial stages of SR, we only count generated N -values with $N \geq 3$. This is because we will always select $t_2 = 1$, as this choice already contributes zero error, and the degenerate SR stage with $N = 1$ contributes no error and needs no t_1 value (let $t_1 = 0$). Corresponding to this convention for counting the number of nontrivial stages, we let $\mathcal{G}'(N_0) = \mathcal{G}(N_0) \setminus \{1, 2\}$ be the generated N -values when starting with N_0 but excluding $N = 1$ and $N = 2$.

Given an initial N_0 and a target ε , we consider two strategies for computing the round counts t_N for $N \in \mathcal{G}'(N_0)$. Both use the upper bound $\Delta_{\text{SN}}^{\text{ub}}(N, q, r) = (2N^{3/2}/(r+2)) \cdot ((q+N)/(2N))^{r/2+1}$ on $\Delta_{\text{SN}}(N, q, r)$ given by equation (1).

1. *Split the error equally.* Let $n = |\mathcal{G}'(N_0)| \approx \lg N_0$ be the number of nontrivial stages. For each $N \in \mathcal{G}'(N_0)$ let t_N be smallest number r for which $\Delta_{\text{SN}}^{\text{ub}}(N, \lfloor N/2 \rfloor, r) \leq \varepsilon/n$. This will result in rounds counts t_N that diminish with diminishing N , each stage contributing about the same portion to the error.
2. *Constant round count.* Let r_0 be the smallest number r for which the sum $\sum_{N \in \mathcal{G}'(N_0)} \Delta_{\text{SN}}^{\text{ub}}(N, \lfloor N/2 \rfloor, r) < \varepsilon$, and let $t_N = r_0$ for all $N \in \mathcal{G}'(N_0)$. This will result in stages that contribute a diminishing amount to the error.

The table of Fig. 2 illustrates the expected and worst-case number of rounds that result from these two strategies if we encipher on a domain of $N_0 = 10^d$

d	2	4	6	8	10	12	14	15	16	18	20	30
min-1	187	239	289	337	386	435	483	507	531	580	628	869
mean-1	359	464	563	660	758	856	952	1000	1048	1145	1242	1723
max-1	1110	2442	4411	6402	8885	11842	14790	16639	18239	22158	26069	51453
min-2	218	225	272	318	365	413	460	484	507	555	602	840
mean-2	427	450	544	636	730	826	920	968	1014	1110	1204	1680
max-2	1308	2701	5168	7951	11681	16107	20701	23716	26365	32745	39131	83160

Fig. 2. Speed of SR shuffle. Minimum, mean (rounded to nearest integer), and maximum number of rounds to SR-encipher a d -digit decimal string with error $\varepsilon \leq 10^{-10}$ and round counts t_N selected by strategy 1 or strategy 2, as marked. The split is $p_N = \lfloor N/2 \rfloor$. Round-counts for MC always coincide with the max-labeled rows.

points and cap the error at $\varepsilon = 10^{-10}$. The pronounced differences between mean and max round counts (a factor exceeding 17 when $n = 16$) coincides with the saving of SR over MC. In contrast, there is only a modest difference in mean round-counts between the two round-count selection strategies.

In numerical experiments, more complex strategies for determining the round counts did not work better.

NON-EQUAL SPLITS. Besides the split of $p_N = \lfloor N/2 \rfloor$, we considered splits of $p_N = \lfloor \alpha N \rfloor$ for $\alpha \in (0, 1)$. For example, if the input is a decimal string then a selection of $\alpha = 0.1$ corresponds to using SN until a 90% fraction of the cards are (almost) properly distributed, at which point there would be only a 10% chance of needing to recurse. When a recursive call is made, it would be on a string of length one digit less than before. But splits this uneven turn out to be inefficient; see Fig. 3. On the other hand, when the split $p_N = \lfloor \alpha N \rfloor$ has α close to $1/2$, the expected number of rounds is not very sensitive to α ; again see the figure. Small α make each SN stage slower, but there will be fewer of them; large α make each SN stage faster, but there will be more.

Given the similar mean round counts for strategies 1 and 2, the similar mean round counts all α near $1/2$, the implementation simplicity of dividing by 2, and the better maximum rounds counts of strategy 1, the choice of strategy 1 and $\alpha = 1/2$ seems best.

6 Incorporating Tweaks

The possibly-small domain for FPE makes it important, in applications, to have the constructed cipher be *tweaked*: an additional argument T , the tweak, names the desired permutation in a family of keyed permutations [12]. In the reference experiment that defines security one asks for indistinguishability (complexity theoretic or information theoretic) from a family of tweak-indexed, uniformly random permutations, each tweak naming an independent permutation from




Fig. 3. Selecting the split. Expected number of rounds (the y -coordinate) to encipher $N = 10^{16}$ points using SR and a split of $p_N = \lfloor \alpha N \rfloor$ for various α (the x -axis). The total variation distance is capped at $\varepsilon = 10^{-10}$. The top (blue) curve is with round counts t_N determined by for strategy 1; the bottom (red) curve for strategy 2. In both cases the smallest expected number of rounds occurs with a non-canonical split: 1048 rounds ($\alpha = 0.5$) reduced to 1043 rounds ($\alpha = 0.53$) for strategy 1; and 1014 rounds ($\alpha = 0.5$) reduced to 1010 rounds ($\alpha = 0.52$) for strategy 2.

the collection. As an example of a tweak’s use, in the context of enciphering a credit card number, one might encipher only the middle six digits, using the first six and last four digits as the tweak.

The obvious way to incorporate a tweak in SR is to make the round constants K_i (line 21 of Fig. 1) depend on it, and to make the round functions $F(i, \hat{X})$ (line 23 of Fig. 1) depend on it. Note, however, that an inefficiency emerges when the former is done: if there is a large space of possible tweaks, it will no longer be possible to precompute the round constants K_i . In addition, we do not want to get a security bound that gives up a factor corresponding to the number of tweaks used, which would be a potentially major loss in quantitative security.

As it turns out, neither price need be paid. In particular, it is fine to leave the round constants independent of the tweak T , and, even when doing so, there need be no quantitative security loss in the bound from making this change. What we call tweaked-SR, then, is identical to Fig. 1 except that the tweak T is added to the scope of F at line 23.

To establish security for this scheme, obtaining the same bounds as before, we go back to the swap-or-not shuffle and show that, in that context, if the round constants are left untweaked but the round function is tweaked, then equation (1) continues to hold. The result is as follows.

Theorem 5. Fix q_1, \dots, q_l with $\sum_{i=1}^l q_i = q$. Let $X_t^1, X_t^2, \dots, X_t^l$ be SN shuffles on G driven by the same round constants K_1, \dots, K_r , but independent round functions. Let $X_t = (X_t^1, \dots, X_t^l)$. For i with $1 \leq i \leq l$, let π^i be the uniform distribution on q_i samples without replacement from G , and let

$\pi = \pi^1 \times \pi^2 \cdots \times \pi^l$. That is, π is the distribution of l independent samples, one each from $\pi^1, \pi^2, \dots, \pi^l$. Let τ be the distribution of X_r . Then

$$\|\tau - \pi\| \leq \frac{2N^{3/2}}{r+2} \left(\frac{q+N}{2N}\right)^{r/2+1}. \tag{4}$$

Proof. Let

$$\Delta(j) = \sum_{m=0}^{j-1} \frac{\sqrt{N}}{2} \left(\frac{m+N}{2N}\right)^{r/2}.$$

We show that

$$\|\tau - \pi\| \leq \Delta(q)$$

from which (4) follows by way of

$$\begin{aligned} \|\tau - \pi\| &\leq \sum_{m=0}^{q-1} \frac{\sqrt{N}}{2} \left(\frac{m+N}{2N}\right)^{r/2} \\ &\leq N^{3/2} \int_0^{q/2N} (1/2+x)^{r/2} dx \\ &\leq \frac{2N^{3/2}}{r+2} \left(\frac{q+N}{2N}\right)^{r/2+1}. \end{aligned}$$

For random variables W_1, W_2, \dots, W_j , we write $\tau^i(\cdot | W_1, W_2, \dots, W_j)$ for the conditional distribution of X_r^i given W_1, W_2, \dots, W_j . Then Lemma 1 implies that

$$\|\tau - \pi\| \leq \sum_{i=1}^l \mathbf{E} (\|\tau^i(\cdot | X_r^1, \dots, X_r^{i-1}) - \pi^i\|). \tag{5}$$

We claim that

$$\mathbf{E} (\|\tau^i(\cdot | X_r^1, \dots, X_r^{i-1}) - \pi^i\|) \leq \Delta(q_i). \tag{6}$$

For distributions μ and ν the total variation distance $\|\mu - \nu\|$ is half the L^1 -norm of $\mu - \nu$. Since the L^1 -norm is convex, to verify the claim it is enough to show that

$$\mathbf{E} (\|\tau^i(\cdot | X_r^1, \dots, X_r^{i-1}, K_1, \dots, K_r) - \pi^i\|) \leq \Delta(q_i).$$

But the X_r^i are conditionally independent given K_1, K_2, \dots, K_r , so

$$\tau^i(\cdot | X_r^1, \dots, X_r^{i-1}, K_1, \dots, K_r) = \tau^i(\cdot | K_1, \dots, K_r).$$

Thus it remains to show that

$$\mathbf{E} (\|\tau^i(\cdot | K_1, \dots, K_r) - \pi^i\|) \leq \Delta(q_i) = \sum_{m=0}^{q_i-1} \frac{\sqrt{N}}{2} \left(\frac{m+N}{2N}\right)^{r/2},$$

but this inequality is shown on page 8 of [11]. This verifies (6), and combining this with (5) gives

$$\begin{aligned} \|\tau - \pi\| &\leq \sum_{i=1}^l \Delta(q_i) \\ &\leq \Delta(q), \end{aligned}$$

where the second inequality holds because the summands in the definition of $\Delta(j)$ are increasing. This completes the proof. \square

Theorem 5 plays the same role in establishing the security for tweaked-SR as equation (1) played for establishing the security of the basic version. The values in the table of Fig. 2, for example, apply equally well to the tweakable-SR.

We comment that in the the tweakable version of SR, the round constants do depend on the generated N -values. This dependency can also be eliminated, but we do not pursue this for now.

7 Absence of Timing Attacks

With SR (and, more generally, with **SR**), the total number of rounds t^* used to encipher a plaintext $X \in [N_0]$ to a ciphertext $Y \in [N_0]$ will depend on X and the key $\mathbf{K} = KF$. This suggests that an adversary’s acquiring t^* , perhaps by measuring the running time of the algorithm, could be damaging. But this is not the case—not in the typical setting, where the adversary knows the ciphertext—for, knowing Y , one can determine the corresponding t^* value.

It is easiest to describe this when $N_0 = 2^n$ is a power of two, whence the generated N -values are $2^n, 2^{n-1}, \dots, 4, 2, 1$. Let $t'_0, t'_1, \dots, t'_{n-2}, t'_{n-1}, t'_n$ be the corresponding round counts (the last two values are 1 and 0, respectively). Let $t^*_j = \sum_{i \leq j} t'_i$ be the cumulative round counts: the total number of SN rounds if we run for $j + 1$ stages. Then t^* is simply t^*_ℓ where ℓ is the number of leading 0-bits in the n -bit binary representation of Y . The adversary holding a ciphertext of $Y = 0^z 1Z$, knows that it was produced using $t^* = t^*_z$ rounds of SN. Ciphertext 0^n is the slowest to produce, needing t^*_n rounds.

The observation generalizes when N_0 is not a power of 2: the set $[N_0]$ is partitioned into easily-calculated intervals and the number of SN rounds that a ciphertext Y was subjected to is determined by the interval containing it.

8 Discussion

ALTERNATIVE DESCRIPTION. It is easy to eliminate the tail recursion of Fig. 1; no stack is needed. This and other changes are made to the alternative description of tweaked-SR given in Fig. 4. While the algorithm looks rather different from before, it is equivalent.

WHICH PILE TO RECURSE ON? The convention that **SR** recurses on the first (left) pile of cards, rather than on the second (right) pile of cards, simplifies

```

50 procedure  $E_{KF}^{T, N_0}(X)$  //Encipher  $X \in [N_0]$  with tweak  $T$ , key  $KF$ 
51  $N \leftarrow N_0$  //initial- $N$ 
52 for  $j \leftarrow 0$  to  $\infty$  do //for each stage, until we return
53   for  $i \leftarrow 1$  to  $t_N$  do //SN, for as many rounds as needed for this stage
54      $X' \leftarrow K_i - X \pmod{N}$  //  $X'$  is the partner of  $X$ 
55      $\hat{X} \leftarrow \max(X, X')$  //canonical name for  $\{X, X'\}$ 
56     if  $F(i, \hat{X}, T) = 1$  then  $X \leftarrow X'$  //maybe swap  $X$  and  $X'$ 
57   if  $X \geq \lfloor N/2 \rfloor$  then return  $X$  //right pile is done
58    $N \leftarrow \lfloor N/2 \rfloor$  //left pile is new domain to shuffle

```

Fig. 4. Alternative description of the tweaked construction. We eliminate the recursion and assume the canonical split. The values t_N again parameterize the algorithm, influencing the mechanism’s speed and the quality of enciphering.

bookkeeping: in this way, we will always be following a card $X \in [N]$ for decreasing values of N . Had we recursed on the second pile we would be following a card $X \in [N_0 - N + 1 .. N_0 - 1]$ for decreasing values of N . Concretely, the code in Figures 1 and 4 would become more complex with the recurse-right convention.

MULTIPLE CONCURRENT DOMAINS. Our assumption has been that the domain for the constructed cipher is $[N_0]$ for some N_0 . As with variable-input-length (VIL) PRFs, it makes sense to seek security against adversaries that can simultaneously encipher points from any number of domains $\{[N_0] : N_0 \in \mathcal{N}\}$, as previously formalized [3]. This can be handled by having the round-function and round-keys depend on the description of the domain N_0 . Once again it seems unnecessary to reflect the N_0 dependency in the round-keys. To prove the conjecture will take a generalization of Theorem 5.

OPEN QUESTION. The outstanding open question in this domain is whether there is an oblivious shuffle on N cards where a card can be tracked through the shuffle in *worst-case* $\Theta(\lg N)$ -time. Equivalently, can we do information-theoretic PRF to PRP conversion with $\Theta(\lg N)$ calls, always, to a constant-output-length PRF?

Acknowledgments. This work was made possible by Tom Ristenpart and Scott Yilek generously sharing an early draft of their work [16]. Thanks also to Tom and Scott for their comments and interaction. Thanks to Terence Spies and Voltage Security, whose interest in FPE has motivated this line of work. Our work was supported under NSF grants CNS-0904380, CNS-1228828 and DMS-1007739.

References

1. Accredited Standards Committee X9, Incorporated (ANSI X9): X9.124: Symmetric Key Cryptography for the Financial Services Industry — Format Preserving Encryption (2011) (manuscript)

2. Bellare, M., Impagliazzo, R.: A Tool for Obtaining Tighter Security Analyses of Pseudorandom Function Based Constructions, with Applications to PRP to PRF Conversion. ePrint report 1999/024 (1999)
3. Bellare, M., Ristenpart, T., Rogaway, P., Stegers, T.: Format-Preserving Encryption. In: Jacobson Jr., M.J., Rijmen, V., Safavi-Naini, R. (eds.) SAC 2009. LNCS, vol. 5867, pp. 295–312. Springer, Heidelberg (2009)
4. Black, J.A., Rogaway, P.: Ciphers with Arbitrary Finite Domains. In: Preneel, B. (ed.) CT-RSA 2002. LNCS, vol. 2271, pp. 114–130. Springer, Heidelberg (2002)
5. Brightwell, M., Smith, H.: Using Datatype-preserving Encryption to Enhance Data Warehouse Security. In: 20th National Information Systems Security Conference Proceedings (NISSC), pp. 141–149 (1997)
6. Did, user profile <http://math.stackexchange.com/users/6179/did>: Total Variation Inequality for the Product Measure. Mathematics Stack Exchange, <http://math.stackexchange.com/q/72322> (2011) (last visited June 2, 2014)
7. Dworkin, M.: NIST Special Publication 800-38G: Draft. Recommendation for Block Cipher Modes of Operation: Methods for Format-Preserving Encryption (July 2013)
8. FIPS 74: Guidelines for Implementing and Using the NBS Data Encryption Standard. U.S. National Bureau of Standards, U.S. Dept. of Commerce (1981)
9. Granboulan, L., Pornin, T.: Perfect Block Ciphers with Small Blocks. In: Biryukov, A. (ed.) FSE 2007. LNCS, vol. 4593, pp. 452–465. Springer, Heidelberg (2007)
10. Håstad, J.: The Square Lattice Shuffle. *Random Structures and Algorithms* 29(4), 466–474 (2006)
11. Hoang, V.T., Morris, B., Rogaway, P.: An Enciphering Scheme Based on a Card Shuffle. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 1–13. Springer, Heidelberg (2012)
12. Liskov, M., Rivest, R., Wagner, D.: Tweakable Block Ciphers. *J. of Cryptology* 24(3), 588–613 (2011)
13. Morris, B.: The Mixing Time of the Thorp Shuffle. *SIAM J. on Computing* 38(2), 484–504 (2008)
14. Morris, B., Rogaway, P., Stegers, T.: How to Encipher Messages on a Small Domain: Deterministic Encryption and the Thorp Shuffle. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 286–302. Springer, Heidelberg (2009)
15. Naor, M., Reingold, O.: On the Construction of Pseudo-Random Permutations: Luby-Rackoff Revisited. *J. of Cryptology* 12(1), 29–66 (1999)
16. Ristenpart, T., Yilek, S.: The Mix-and-Cut Shuffle: Small-Domain Encryption Secure against N Queries. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 392–409. Springer, Heidelberg (2013)
17. Rudich, S.: Limits on the Provable Consequences of One-Way Functions. Ph.D. Thesis, UC Berkeley (1989)
18. Stefanov, E., Shi, E.: FastPRP: Fast Pseudo-Random Permutations for Small Domains. *Cryptology ePrint Report* 2012/254 (2012)
19. Thorp, E.: Nonrandom Shuffling with Applications to the Game of Faro. *J. of the American Statistical Association* 68, 842–847 (1973)

A Proof of Lemma 1

We follow the approach outlined in [6] for bounding the total variation distance between two product measures. Define $V = V_1 \times V_2 \times \dots \times V_n$. Note that

$$2 \|\mu - \nu\| = \sum_{x \in V} |\mu(x) - \nu(x)| \tag{7}$$

$$= \sum_{x \in V} |\mu_1(x)\mu_2(x) \cdots \mu_n(x) - \nu_1(x)\nu_2(x) \cdots \nu_n(x)|, \tag{8}$$

where, for j with $1 \leq j \leq n$, we define $\mu_j(x)$ to be $\mu(x_j \mid x_1, \dots, x_{j-1})$, with a similar definition for $\nu_j(x)$. For $x \in V$, define $s_j(x)$ as

$$\mu_1(x)\mu_2(x) \cdots \mu_j(x)\nu_{j+1}(x) \cdots \nu_n(x).$$

Then

$$s_0(x) = \nu_1(x)\nu_2(x) \cdots \nu_n(x) \text{ and} \\ s_n(x) = \mu_1(x)\mu_2(x) \cdots \mu_n(x),$$

and hence by the triangle inequality the quantity (8) is at most

$$\sum_{x \in V} \sum_{j=0}^{n-1} |s_{j+1}(x) - s_j(x)| \tag{9}$$

$$= \sum_{l=0}^{n-1} \sum_{x \in V} |\mu_{l+1}(x) - \nu_{l+1}(x)| \mu_1(x)\mu_2(x) \cdots \mu_l(x)\nu_{l+2}(x) \cdots \nu_n(x). \tag{10}$$

If we sum the terms over all $x \in V$ whose first l components are x_1, x_2, \dots, x_l we get

$$\mu(x_1, x_2, \dots, x_l) \sum_{v \in V_{l+1}} \left| \mu(v \mid x_1, x_2, \dots, x_l) - \nu(v \mid x_1, x_2, \dots, x_l) \right| \\ = 2 \mu(x_1, x_2, \dots, x_l) \|\mu(\cdot \mid x_1, \dots, x_l) - \nu(\cdot \mid x_1, \dots, x_l)\|.$$

Summing this over x_1, \dots, x_l gives

$$2 \mathbf{E} \left(\|\mu(\cdot \mid Z_1, \dots, Z_l) - \nu(\cdot \mid Z_1, \dots, Z_l)\| \right)$$

where $(Z_1, \dots, Z_n) \sim \mu$, and now summing this over l proves the lemma.