On Generalized Feistel Networks

Viet Tung Hoang and Phillip Rogaway

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

Abstract. We prove beyond-birthday-bound security for most of the well-known types of generalized Feistel networks: (1) unbalanced Feistel networks, where the n-bit to m-bit round functions may have $n \neq m$; (2) alternating Feistel networks, where the round functions alternate between contracting and expanding; (3) type-1, type-2, and type-3 Feistel networks, where n-bit to n-bit round functions are used to encipher kn-bit strings for some $k \geq 2$; and (4) numeric variants of any of the above, where one enciphers numbers in some given range rather than strings of some given size. Using a unified analytic framework, we show that, in any of these settings, for any $\varepsilon > 0$, with enough rounds, the subject scheme can tolerate CCA attacks of up to $q \sim N^{1-\varepsilon}$ adversarial queries, where N is the size of the round functions' domain (the larger domain for alternating Feistel). Prior analyses for most generalized Feistel networks established security to only $q \sim N^{0.5}$ queries.

Keywords: Block ciphers, coupling, Feistel networks, generalized Feistel networks, modes of operation, provable security, symmetric techniques.

1 Introduction

Background. Feistel-like ciphers come in several flavors beyond the "classical" one used in DES [31,7]. In speaking of *generalized* Feistel networks we mean to encompass most all of them; see Fig. 1. In particular, we include: *unbalanced* Feistel networks with either expanding or contracting round functions, as described by Schneier and Kelsey [30]; *alternating* Feistel networks, where the rounds alternate between contracting and expanding steps, as described by Anderson and Biham [1] and by Lucks [11]; *type-1*, *type-2*, and *type-3* Feistel networks, as described by Zheng, Matsumoto, and Imai [35], each of which uses an n-bit to n-bit round function to create a kn-bit blockcipher for some $k \geq 2$; and *numeric* variants of any of the above, where one enciphers numbers in $\mathbb{Z}_N$, for some $N \in \mathbb{N}$, instead of enciphering binary strings. Well-known blockciphers that use generalized Feistel networks include Skipjack (an unbalanced Feistel network), BEAR/LION (alternating), CAST-256 (type-1), RC6 (type-2), and MARS (type-3).

The provable-security analysis of Feistel networks begins with the seminal work of Luby and Rackoff [10]. The r round functions used are assumed to be selected uniformly and independently at random ($r = 3$ or $r = 4$ in [10]). One then considers how close to a random permutation the constructed cipher is. Subsequent

T. Rabin (Ed.): CRYPTO 2010, LNCS 6223, pp. 613–630, 2010.

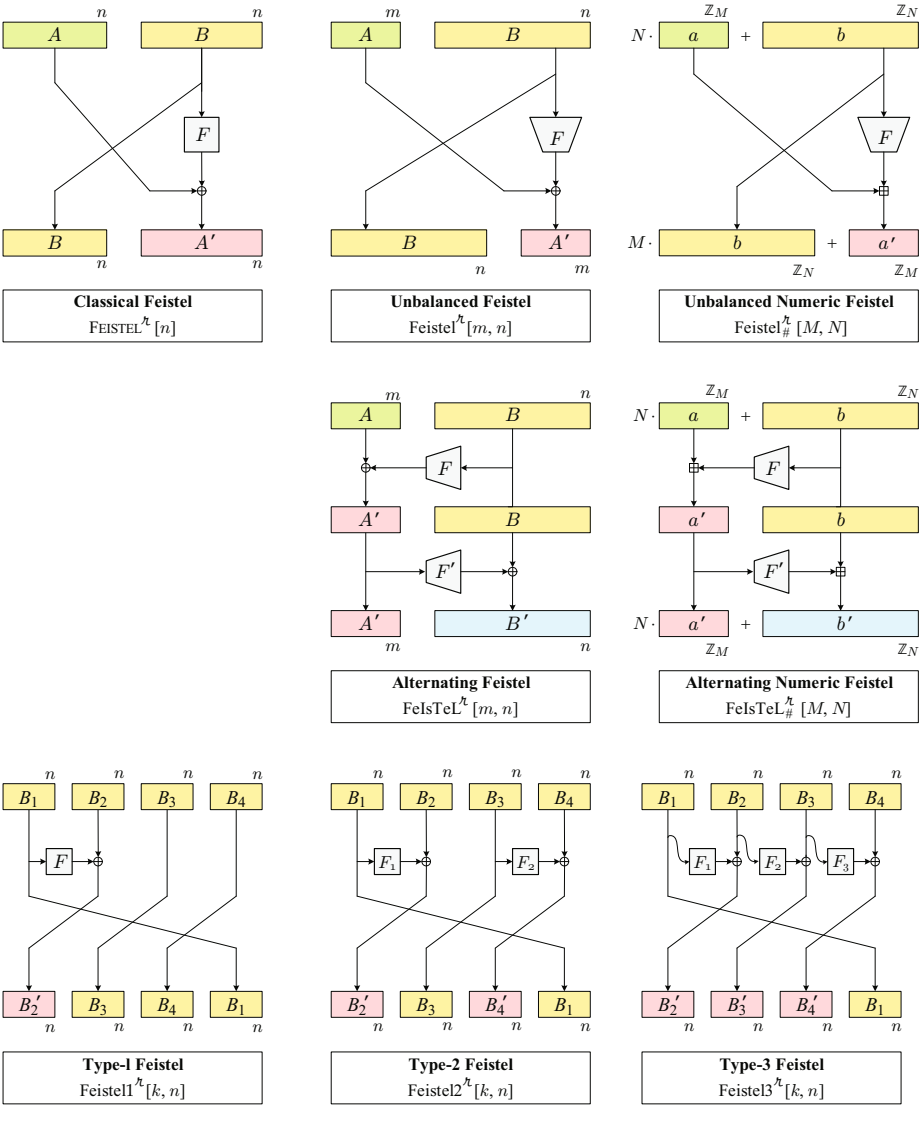

Fig. 1. Generalized Feistel networks. The superscript $\mathcal{R}$ is the number of rounds. The illustrations show a single round $\mathcal{R} = 1$ except for the alternating schemes, where $\mathcal{R} = 2$ rounds are shown. Scheme FEISTEL is the classical balanced-Feistel scheme; all remaining schemes are generalizations of it. Schemes Feistel$_\sharp$ and FeIsTeL$_\sharp$ are numeric variants of Feistel (unbalanced Feistel) and FeIsTeL (alternating Feistel); they encipher a number $x = aN + b \in \mathbb{Z}_{MN}$ ($a \in \mathbb{Z}_M$, $b \in \mathbb{Z}_N$) instead of a string $X \in \{0,1\}^{m+n}$. Schemes Feistel1, Feistel2, and Feistel3 are the so-called type-1, type-2, and type-3 Feistel networks. They are used in modern blockciphers like CAST-256, RC6, and MARS, respectively. Variable k refers to the number of n-bit input blocks $B_1, \ldots, B_k$. The illustrations are for $k = 4$.

scheme	$E =$	$\mathbf{Adv}^{\mathrm{cca}}_{E}(q) \leq$	**where** $\nu =$
classical	$\textsc{Feistel}^{\nu}[n]$	$\frac{2q}{r+1}\left(4q \,/\, 2^n\right)^r$	$6r-1$
unbalanced	$\textsc{Feistel}^{\nu}[m,n]$		
	with $n > m$	$\frac{2q}{r+1}\left((3\lceil n/m\rceil+3)q \,/\, 2^n\right)^r$	$r(4\lceil n/m\rceil+4)$
	with $n \leq m$	$\frac{2q}{r+1}\left(4\lceil m/n\rceil q \,/\, 2^n\right)^r$	$r(2\lceil m/n\rceil+4)$
unbalanced$\sharp$	$\textsc{Feistel}^{\nu}_{\sharp}[M,N]$		
	with $N > M$	$\frac{2q}{r+1}\left((9\lceil \log_M N\rceil+5)q \,/\, N\right)^r$	$r(6\lceil \log_M N\rceil+4)$
	with $N \leq M$	$\frac{2q}{r+1}\left((7\lceil \log_N M\rceil+7)q \,/\, N\right)^r$	$r(2\lceil \log_N M\rceil+6)$
alternating	$\textsc{FeIsTeL}^{\nu}[m,n]$	$\frac{2q}{r+1}\left((6\lceil n/m\rceil+3)q \,/\, 2^n\right)^r$	$r\,(12\lceil n/m\rceil+8)$
alternating$\sharp$	$\textsc{FeIsTeL}^{\nu}_{\sharp}[M,N]$	$\frac{2q}{r+1}\left((6\lceil \log_M N\rceil+3)q \,/\, N\right)^r$	$r\,(12\lceil \log_M N\rceil+8)$
type-1	$\textsc{Feistel1}^{\nu}[k,n]$	$\frac{2q}{r+1}\left(2k(k-1)q \,/\, 2^n\right)^r$	$r(4k-2)$
type-2	$\textsc{Feistel2}^{\nu}[k,n]$	$\frac{2q}{r+1}\left(2k(k-1)q \,/\, 2^n\right)^r$	$r(2k+2)$
type-3	$\textsc{Feistel3}^{\nu}[k,n]$	$\frac{2q}{r+1}\left(4(k-1)^2q \,/\, 2^n\right)^r$	$r(2k+2)$

Fig. 2. Summary of CCA bounds in this paper. The rows correspond to the generalized Feistel networks pictured in Fig. 1. Unbalanced schemes are distinguished by their using contracting ($n > m$) or expanding ($n \leq m$) round functions. Parameters k, m, n, M, N describe the scheme and $r \geq 1$ determines the number of rounds ν. The specified results appear as Theorems 6–10.

work in this information-theoretic framework (still analyzing the classical Feistel construction) includes Maurer [12], Naor and Reingold [19], Vaudenay [33], Maurer and Pietrzak [13], and a sequence of papers by Patarin [26,21,23,24,22]. The last culminates with the claim that six rounds of (classical) Feistel on a $2n$-bit string is enough to defeat (meaning the advantage goes to 0 as $n \to \infty$) adaptive chosen-ciphertext attacks of $2^{n(1-\varepsilon)}$ queries, for any $\varepsilon > 0$.

Information-theoretic analysis of *generalized* Feistel schemes is less mature. We postpone describing the known results except to say that they are either absent (alternating Feistel with highly-imbalanced round functions), quantitatively weak (birthday bounds that generalize Luby and Rackoff's 25-year-old work), or highly specialized (unbalanced Feistel networks with maximally unbalanced contracting round functions).

Contributions. Our CCA-security bounds for generalized Feistel networks are described in Fig. 2. Proofs omitted due to lack of space appear in the full version of this paper [8]. Let us briefly describe each result and how it compares with prior work.

For the classical Feistel network on $2n$ bits, our results are comparable to those of Maurer and Pietrzak (henceforth "MP") [13]. As with that work, the bounds get better as one increases the number of rounds ν. Asymptotically, for any $\varepsilon > 0$, there is a corresponding number of rounds ν (about $6/\varepsilon$) such

that any CCA-adversary has vanishing advantage if it asks at most $q = 2^{n(1-\varepsilon)}$ forwards or backwards queries. Our actual results are concrete, and are a little sharper than MP's bounds; see Fig. 3 for a graphical comparison. Our proof is much simpler than those of MP or Patarin. One reason for this is just that we employ the lovely result of Maurer, Pietrzak, and Renner for passing from NCPA-security to CCA-security [14]. The more important reason stems from our use of *coupling*, a well-known technique from the theory of Markov chains.

Next we look at unbalanced Feistel networks; the round functions are maps $F_i \colon \{0,1\}^n \to \{0,1\}^m$. For the contracting case ($n > m$) we prove CCA-security to $2^{n(1-\varepsilon)}$ queries. Earlier work by Naor and Reingold provided bounds that topped out at $2^{n/2}$ adversarial queries. Interpreting our result, if one holds fixed the block length $\ell = m + n$, bounds improve with increasing imbalance, the best bounds at $m = 1$, the setting earlier studied by Morris, Rogaway, and Stegers ("MRS") [17]. In effect, we "connect up" MP's bounds on balanced Feistel with MRS's bounds on maximally unbalanced Feistel, demonstrating a smooth increase in security with increasing imbalance. This behavior is not an artifact of the analysis; corresponding information-theoretic attacks exist [22,27].

For unbalanced Feistel networks with expanding random round functions our concrete-security results (again see Fig. 2) can similarly be interpreted asymptotically to show CCA security to $2^{n(1-\varepsilon)}$ queries. But note that as imbalance increases in an expanding round functions the value of n goes down, so provable security is effectively vanishing. Again this is no artifact; there are corresponding information-theoretic attacks [22,28].

We next treat unbalanced Feistel networks that acts on numbers instead of strings, the blockcipher we denote $\text{Feistel}_{\sharp}^{r}[M, N]$. This situation is seen in the card-shuffling technique of Thorp [32] (where $M = 2$) and is defined explicitly in the work of Bellare *et al.* [4]. While one might expect unbalanced Feistel schemes to behave similarly in the number-based and string-based settings, being able to show this is something else: the number-based setting is considerably more complex. We note that MRS only managed to deal with the case $M = 2$ and $N = 2^n$, leaving the generalization open. We show security to $q \sim N^{1-\varepsilon}$ queries.

Unbalanced Feistel networks are unpleasant in requiring a "repartitioning" of each round's output before it can be treated as the next round's input. An alternative is suggested by the "ladder" way of drawing DES (the way that avoids wire-crossings, as in our illustration of FeIsTeL). Information-theoretic security bounds for alternating Feistel networks [1,3,4,11] were weak in two ways: quantitatively, they top out at the birthday-bound; qualitatively, they depend on the domain size of the round function with *smaller* domain, leading to a non-result for the highly imbalanced setting. We overcome both issues. Our results cover the numeric as well as the string-based settings.

Finally, we consider type-1, type-2, and type-3 Feistel networks [35], as used in several modern blockciphers. We prove information-theoretically optimal bounds (as the number of rounds becomes large). The proofs here are straightforward compared to those for unbalanced and alternating Feistel, highlighting a strength of the coupling-based approach.

Unmentioned in all of the above is that our string-based results also work when the alphabet is non-binary. This turns out to be useful; for example, one could encipher a 16-digit credit card number (CCN) (the ciphertext again being a 16-digit number) using a scheme $\text{FEISTEL}^{\curvearrowright}_{10}[8]$ just like $\text{FEISTEL}^{\curvearrowright}[8]$ but over the decimal alphabet instead of the binary one [2] (re-interpret the xor operator as, say, modular addition). Our security bounds for schemes with non-binary alphabets are as given in Fig. 2 but with 2^n replaced by d^n, where d is the radix of the alphabet.

In general, finding a unified framework with which to analyze Feistel-like schemes—one that gives concrete, asymptotically optimal, humanly-verifiable bounds—is a contribution we see as being at least as important as all the improved bounds.

Additional related work. In work just subsequent to our own, Patarin provides a concrete security bound for the classical Feistel construction $\text{FEISTEL}^6[n]$ [25]. He goes on to claim beyond-birthday-bound security for the unbalanced scheme $\text{Feistel}^8[n, 2n]$. Earlier versions of our paper confessed an inability to extract concrete security bounds from Patarin's body of work.

Nachef attacks a Feistel variant that she calls an alternating unbalanced Feistel scheme [18], but the scheme is different from the more classical one that we study here. The specific rotation operation used in Nachef's scheme makes this Feistel variant highly insecure.

The first use of a coupling argument in cryptography that we know is due to Mirinov, who used the technique to gave a lovely (even if slightly heuristic) analysis of RC4 [15]. As mentioned earlier, Morris, Rogaway, and Stegers go on to use coupling to analyze the security of a maximally-unbalanced (contracting round function) Feistel network. Our work builds on theirs, but our use of coupling becomes considerably more complex.

Beyond their use in making conventional blockciphers, generalized Feistel networks have been proposed as blockcipher modes-of-operation for format-preserving encryption (FPE) [5,3,4]. Here one usually aims to encipher points within some arbitrary string-valued domain Σ^n, or within some arbitrary numeric domain $\mathbb{Z}_N$. Commercial interest in doing this has been spurred by PCI regulations [29] that require vendors to encipher CCNs they store; an architecturally clean way to do this is to encipher a column in a database without making any modification to the database's schema. There is now a NIST proposal for an FPE-providing mode of operation, FFX [2], that employs an unbalanced or alternating Feistel network over a possibly non-binary alphabet.

2 Preliminaries

Notation. For finite nonempty sets A and B, let $\text{Func}(A, B)$ be the set of all functions from A to B and let $\text{Perm}(A)$ be the set of all permutations on A. For numbers $a, b \geq 1$, let $\text{Func}(a, b)$ be the set of all functions from $\{0,1\}^a$ to $\{0,1\}^b$.

BLOCKCIPHERS. Let $E\colon \mathcal{K} \times \mathcal{M} \to \mathcal{M}$ be a blockcipher, meaning that each $E_K(\cdot) = E(K, \cdot)$ is a permutation on the finite nonempty set $\mathcal{M}$. We emphasize that $\mathcal{M}$ (and also $\mathcal{K}$) need not consist of binary strings of some particular length, as is often assumed to be the case. For any blockcipher E, we let E^{-1} be its inverse blockcipher. For blockcipher $E\colon \mathcal{K}\times\mathcal{M} \to \mathcal{M}$ and adversary A the *advantage* of A in carrying out an (adaptive) chosen-ciphertext attack (CCA) on E is $\mathbf{Adv}^{\mathrm{cca}}_E(A) = \Pr[K \xleftarrow{\$} \mathcal{K}\colon A^{E_K(\cdot),E_K^{-1}(\cdot)} \Rightarrow 1] - \Pr[\pi \xleftarrow{\$} \mathrm{Perm}(\mathcal{M})\colon A^{\pi(\cdot),\pi^{-1}(\cdot)} \Rightarrow 1]$. We say that A carries out an (adaptive) chosen-plaintext attack (CPA) if it asks no queries to its second oracle. Adversary A is *non-adaptive* if it asks the same queries on every run. Let $\mathbf{Adv}^{\mathrm{cca}}_E(q)$ be the maximum advantage of any (adaptive) CCA adversary against E subject to the adversary asking at most q total oracle queries. Similarly define $\mathbf{Adv}^{\mathrm{ncpa}}_E(q)$ for nonadaptive CPA attacks (NCPA).

For blockciphers $F, G\colon \mathcal{K} \times \mathcal{M} \to \mathcal{M}$ let $F \circ G$ denote their cascade, with F's output fed into G's input; formally, $F \circ G\colon \mathcal{K}^2 \times \mathcal{M} \to \mathcal{M}$ is defined by $(F \circ G)_{(K,K')} = G_{K'}(F_K(X))$. To be consistent with this left-to-right convention for composing blockciphers we define composition of permutations by $(f \circ g)(x) = g(f(x))$. (This won't be used often and should not cause confusion for those used to the opposite convention.)

COUPLING ARGUMENTS. The high-level idea for a coupling argument can be explained like this. We have a Markov chain X_t that we want to analyze. For example, the Markov chain may consist of the image of the distinct, fixed strings $(x_1, \ldots, x_q) \in (\{0,1\}^{2n})^q$ as each point is enciphered for t rounds according to the classical Feistel network on $2n$ bits. We would like to show that, after $t = \varkappa$ rounds, the tuple of points X_t is pretty close to being uniformly distributed. For this purpose, we introduce a *second* Markov chain U_t that, after any number of rounds t, is indisputably uniform. We arrange so that X_t and U_t can be viewed as co-evolving on a common probability space; formally, we create a joint distribution that yields the correct marginal distributions. We try to arrange our joint distribution so that, usually, X_t and U_t quickly *couple*: for *most* random choices, it does not take long until $X_t = U_t$. After X_t and U_t come together, they should remain so. The basic observation underlying coupling is that the statistical distance between the distributions associated to X_t and U_t is upperbounded by the probability that $X_t \neq U_t$.

More formally, let μ and ν be probability distributions on a finite event space Ω. The *total variation distance* between distributions μ and ν is defined as $\|\mu - \nu\| = \frac{1}{2}\sum_{x\in\Omega} |\,\mu(x) - \nu(x)\,| = \max_{S\subset\Omega}\{\mu(S) - \nu(S)\}$. A *coupling* of μ and ν is a pair of random variables $X, Y\colon \Omega \to R$ (the set R is arbitrary) such that $X \sim \mu$ and $Y \sim \nu$, that is, variables X and Y have marginal distributions μ and ν, respectively. The *coupling lemma* we will use is as follows.

Lemma 1 (Coupling lemma). *Let μ and ν be probability distributions on a finite event space Ω and let (X, Y) be a coupling of μ and ν. Then $\|\mu - \nu\| \leq \Pr[X \neq Y]$.*

From coupling to ncpa-security. Suppose that an adversary asks some non-adaptive distinct queries. The adversary's NCPA advantage cannot exceed the total variation distance between the distribution of the outputs from her queries and the uniform distribution. The uniform distribution itself can be viewed as the distribution of outputs from a uniformly random choice of distinct queries. Think of a coupling argument as a computer program that accepts as its input either the actual adversarial queries or a pool of uniformly random, distinct queries. On each input, the program implements a Feistel network and gives a random output. The program tries to produce the same output on its two possible inputs. Hence the total variation distance between the distributions of the program's outputs is upperbounded by the program's probability of failure (that is, its failure to produce the same output in the two cases).

To ease the design of such a program, a hybrid argument is employed and a chain of inputs is created—the first being the adversarial queries and the last being the pool of uniformly random, distinct ones. The purpose of this hybrid argument is to reduce the difference between any pair of adjacent inputs in the chain. Given an arbitrary pair of adjacent inputs, our goal now is to design a coupling program that produces identical output on those two inputs with high probability. The program runs both inputs, one after another. When the program starts running the second input, it has finished the operations on the first input and now knows all the random choices of the first Feistel network. It then uses this knowledge in implementing the second Feistel network. For example, if at some step the second network needs a uniformly random string then the program may reuse the corresponding string from the first network. The random choices in the second network are geared toward the first output, but they are subject to the restriction that the round functions in the second network must be independent and uniformly random.

From ncpa to cca-security. We bound the CCA-security of a Feistel network from its NCPA-security by using the following result of Maurer, Pietrzak, and Renner [14, Corollary 5]. It is key to our approach, effectively letting us assume that our adversaries are of the simple, NCPA breed. Recall that in writing $F \circ G$, the blockciphers are, in effect, independently keyed.

Lemma 2 (Maurer-Pietrzak-Renner). *If F and G are blockciphers on the same message space then, for any q, $\mathbf{Adv}^{\mathrm{cca}}_{F\circ G^{-1}}(q) \leq \mathbf{Adv}^{\mathrm{ncpa}}_{F}(q) + \mathbf{Adv}^{\mathrm{ncpa}}_{G}(q)$.*

3 Classical Feistel

This section provides a strong, concrete security bound for conventional, balanced Feistel networks. It also serves as a pedagogical example for proving security of a Feistel network using coupling; some later examples get much more complex.

Defining the scheme. Fix $n \geq 1$ and let $F\colon \{0,1\}^n \to \{0,1\}^n$ be a function. Define from F the permutation $\Psi_F\colon \{0,1\}^{2n} \to \{0,1\}^{2n}$ by way of

$\Psi_F(A,B) = (B, A \oplus F(B))$ where $|A| = |B| = n$, and $\oplus$ denotes xor. Blockcipher $\text{FEISTEL}^{\nu}[n]\colon \mathcal{K} \times \{0,1\}^{2n} \to \{0,1\}^{2n}$ has key space $\mathcal{K} = (\text{Func}(n,n))^{\nu}$ and a key $(F_1, \ldots, F_{\nu}) \in \mathcal{K}$ names the permutation $\Psi_{F_1} \circ \cdots \circ \Psi_{F_{\nu}}$ on $\{0,1\}^{2n}$. Each F_i is called the round function at round i. For an illustration, see Fig. 1.

INITIAL NOTATION. Given a query X to $E = \text{FEISTEL}^{\nu}[n]$, define its round-0 output to be X itself, while the round-t output is $(\Psi_{F_1} \circ \cdots \circ \Psi_{F_t})(X)$. The *coin* of the query X at round t is the string $A \oplus F(B)$, where F is the round function at round t and (A,B) is the round-$(t-1)$ output, with $|A| = |B| = n$. Two queries *collide* at time t if their round-t outputs have the same final n bits.

NCPA-SECURITY. We will now prove the NCPA-security of E by way of coupling, afterwards lifting this to show CCA-security using the result of [14] from Lemma 2. The lemma below will help us bound the probability that we *fail* to couple.

Lemma 3. *For the blockcipher $E = \text{FEISTEL}^{\nu}[n]$, the chance that two distinct non-adaptive queries collide at time $t \geq 1$ is at most 2^{-n}.*

Proof. Suppose that the Feistel network receives distinct nonadaptive queries X_1 and X_2. For each $i \in \{1,2\}$, let (A_i, B_i) be the output at round $t-1$ of X_i, where $|A_i| = |B_i| = n$. The queries X_1 and X_2 collide at time t if and only if $A_1 \oplus F(B_1) = A_2 \oplus F(B_2)$, with F being the round function at round t. This occurs with probability 2^{-n} if B_1 and B_2 differ, because F is uniformly random. If $B_1 = B_2$ then so are A_1 and A_2, which contradicts the hypothesis that X_1 and X_2 are distinct. □

Theorem 4. *Let $E = \text{FEISTEL}^{\nu}[n]$, $\nu = 3r$. Then $\mathbf{Adv}^{\text{ncpa}}_E(q) \leq \frac{q}{r+1}\,(4q/2^n)^r$.*

Proof. Suppose that E receives non-adaptive distinct queries $X_1, \ldots, X_q$. For each $\ell \leq q$, consider a vector of queries $(Z_1, \ldots, Z_q)$ such that Z_i is X_i if $i \leq \ell$ and Z_i is chosen uniformly from $\{0,1\}^{2n} \backslash \{Z_1, \ldots, Z_{i-1}\}$ otherwise. Let μ_ℓ be the distribution of the vector of q outputs when E receives queries $Z_1, \ldots, Z_q$. We will show in a moment that the total variation distance between μ_ℓ and $\mu_{\ell+1}$ is at most $(4\ell \,/\, 2^n)^r$ for every $\ell \leq q-1$. Assuming this, we have, by hybrid argument,

$$\mathbf{Adv}^{\text{ncpa}}_E(q) \leq \sum_{\ell=0}^{q-1} \|\mu_\ell - \mu_{\ell+1}\| \leq \sum_{\ell=0}^{q-1} (4\ell \,/\, 2^n)^r \leq 2^{r(2-n)} \int_0^q x^r\, dx,$$

which is $\frac{q}{r+1}\,(4q \,/\, 2^n)^r$. Now we show the claim. Fix a value $\ell \leq q-1$. We must bound the total variation distance between μ_ℓ and $\mu_{\ell+1}$, each of them is a distribution of a vector of q outputs. However, only the first $\ell+1$ components of the vector matter, because of the uniform sampling of the other. Consider a $3r$-round balanced Feistel network on n bits that receives queries $X_1, \ldots, X_{\ell+1}$. Let $X_i(t)$ be the output at round t from the query X_i.

THE COUPLING. We construct another $3r$-round balanced Feistel network on n bits with its non-adaptive distinct queries $U_1, \ldots, U_{\ell+1}$. Let $U_i(t)$ be the output at round t of the new Feistel network on input U_i. The construction of the new Feistel network will satisfy the following conditions:

- Query U_j equals to X_j for every $j \leq \ell$, and $U_{\ell+1}$ is uniformly chosen over $\{0,1\}^{2n} \backslash \{U_1, \ldots, U_\ell\}$.
- If for all $i \leq \ell + 1$, the outputs at round t of X_i and U_i are identical then so are their outputs in any subsequent round.

Let T be the smallest round for which X_i and U_i have identical outputs for every $i \leq \ell + 1$. From the second condition above and from Lemma 1, we have that

$$\|\mu_\ell - \mu_{\ell+1}\| \leq \Pr[X_i(3r) \neq U_i(3r) \text{ for some } i \leq \ell + 1] = \Pr[T > 3r] \ .$$

The first condition above describes how to initialize $U_1(0), \ldots, U_{\ell+1}(0)$. As the coin of U_i at round $t+1$ dictates how to update $U_i(t+1)$ from $U_i(t)$, it suffices to show how to construct just that coin.

- If U_i collides with some previous query U_j at time t then the coin at round $t+1$ of U_i is defined so as to ensure consistency with the earlier query.
- Suppose that, in the new Feistel network, U_i does not collide with any previous query at time t. If the query X_i collides with some previous query X_j at time t then we choose a string uniformly from $\{0,1\}^n$ to be the coin of U_i at round $t+1$. Otherwise, the coin of X_i at round $t+1$ is uniformly distributed over $\{0,1\}^n$ and U_i will use exactly the same coin at round $t+1$.

Note that U_i and X_i always have the same output at round t, for every $i \leq \ell$ and every t. Consider the event Coll that in either Feistel networks, the $(\ell+1)$-th query collides with some previous query at some time $t \in \{1,2\}$. From Corollary 3, each such collision occurs with probability at most 2^{-n}. Summing over the two Feistel networks, two rounds, and ℓ previous queries shows that the probability Coll occurs is at most $4\ell \ / \ 2^n$. Unless Coll occurs, $U_{\ell+1}$ and $X_{\ell+1}$ will share the coins at the second and third rounds, and then have identical outputs at the third round. Hence $\Pr[T > 3] \leq \Pr[\mathsf{Coll}]$, which is at most $4\ell \ / \ 2^n$.

Now imagine that we run a sequence of trials. In each trial, we observe the outputs of $X_{\ell+1}$ and $U_{\ell+1}$ for an additional three rounds. The probability that $X_{\ell+1}$ and $U_{\ell+1}$ have different outputs after the first trial is at most $4\ell \ / \ 2^n$. Since the round functions of both Feistel networks in each trial are independent with those in previous trials, the conditional probability that $X_{\ell+1}$ and $U_{\ell+1}$ have different outputs after the r-th trial, given that their outputs remain different after the first $r-1$ trials, is again at most $4\ell \ / \ 2^n$. Hence $\Pr[T > 3r] \leq (4\ell \ / \ 2^n)^r$. □

CCA-SECURITY. Let $\mathtt{Rev}$ denote the permutation on $\{0,1\}^{2n}$ where $\mathtt{Rev}(A, B) = (B, A)$, for $|A| = |B| = n$. The following observation is standard; see [13] for proof.

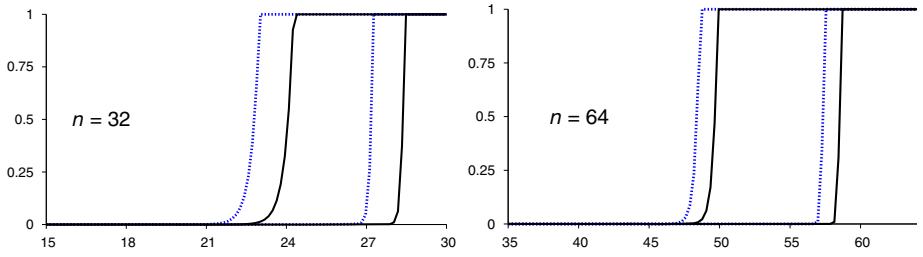

Fig. 3. Proven CCA-security for the classical Feistel network: our own bounds and MP's. The x-axis gives the log base-2 of the number of adversarial queries and the y-axis gives upper bounds on an adversary's CCA advantage. In the left-hand plot (64-bit inputs), the dashed lines depict MP's bounds for $\text{FEISTEL}^{24}[32]$ (left) and $\text{FEISTEL}^{96}[32]$ (right); the solid lines depict our own bounds. In the right-hand plot (128-bit inputs), the dashed lines likewise depict MP's bounds for $\text{FEISTEL}^{24}[64]$ (left) and $\text{FEISTEL}^{96}[64]$ (right); the solid lines depict our own bounds.

Lemma 5. *If F and G are the blockcipher* $\text{FEISTEL}^{\varkappa}[n]$ *then $F \circ G^{-1} \circ \texttt{Rev}$ is the blockcipher* $\text{FEISTEL}^{2\varkappa-1}[n]$. □

Employing Lemma 2 we conclude the following.

Theorem 6. *Let $E = \text{FEISTEL}^{\varkappa}[n]$, $\varkappa = 6r-1$. Then* $\mathbf{Adv}^{\text{cca}}_{E}(q) \le \frac{2q}{r+1}\,(4q/2^n)^r$.

Asymptotic interpretation. For an asymptotic interpretation of Theorem 6, fix $r > 0$. Suppose that $q = 2^{n(1-1/r)}$. Let E_n be the blockcipher $\text{FEISTEL}^{6r-1}[n]$. Then

$$\mathbf{Adv}^{\text{cca}}_{E_n}(q) \le \frac{2q}{r+1}\,(4q \,/\, 2^n)^r = \frac{2^{2r+1}}{r+1} \,/\, 2^{n/r},$$

which goes to 0 as $n \to \infty$. Translating into English, CCA security is guaranteed to about $q = 2^{n(1-\varepsilon)}$ adversarial queries as long as one employs $\varkappa \ge 6/\varepsilon - 1$ rounds. At a higher level still, ignoring the $1-\varepsilon$ multiplier in the exponent, an appropriate number of rounds lets one tolerate nearly $q = 2^n$ adversarial queries.

Comparisons. Maurer and Pietrzak's earlier work proves a security bound of $\mathbf{Adv}^{\text{cca}}_{E}(q) \le 4q^2 \,/\, 2^{2n} + 2q\,(8q \,/\, 2^n)^r$ for $E = \text{FEISTEL}^{6r-1}[n]$. Our own bound is always tighter than this; see Fig. 3 for a comparison of Theorem 6 and MP's bound. Earlier versions of our paper explained that we were unable to plot Patarin's latest bounds [26] due to the absence of a concrete security statement. In very recent work [25] (subsequent to our own), Patarin bounds the security of $E = \text{FEISTEL}^{6}[n]$ by $\mathbf{Adv}^{\text{cca}}_{E}(q) \le 8q \,/\, 2^n + q^2 \,/\, 2^{2n+1}$ (assuming $q \le 2^n \,/\, 128n$).

4 Unbalanced Feistel

Defining the scheme. Fix $n, m \ge 1$ and let $F\colon \{0,1\}^n \to \{0,1\}^m$ be a function. Define from F the permutation $\Psi_F\colon \{0,1\}^{m+n} \to \{0,1\}^{m+n}$ by way

of $\Psi_F(A,B) = (B, A \oplus F(B))$ where $|A| = m$ and $|B| = n$, and $\oplus$ denotes xor. We call Ψ_F a Feistel (m,n)-permutation and F its round function. Blockcipher $\text{Feistel}^{\nu}[m,n]\colon \mathcal{K} \times \{0,1\}^{m+n} \to \{0,1\}^{m+n}$ has key space $\mathcal{K} = (\text{Func}(m,n))^{\nu}$ and a key $(F_1,\ldots,F_{\nu}) \in \mathcal{K}$ names the permutation $\Psi_{F_1} \circ \cdots \circ \Psi_{F_{\nu}}$ on $\{0,1\}^{m+n}$. For an illustration, see Fig. 1.

Security of unbalanced feistel schemes. The theorem below shows the CCA-security of $\text{Feistel}^{\nu}[m,n]$. The proof can be found in Appendix A. Interpreted asymptotically, the result says that, with an adequate number of rounds, CCA security is guaranteed to about 2^n adversarial queries. Note that for *expanding* round functions this guarantee eventually becomes meaningless. This is as it should be; expanding round functions with small domains give rise to information-theoretically insecure schemes.

Theorem 7. *Fix integers $m,n,r \geq 1$.*

1) *Let $E = \text{Feistel}^{\nu}[m,n]$ where $n > m$ and $\nu = r(4\lceil n/m \rceil + 4)$. Then $\mathbf{Adv}^{\text{cca}}_E(q) \leq \frac{2q}{r+1}\big((3\lceil n/m \rceil + 3)q/\ 2^n\big)^r$.*

2) *Let $E = \text{Feistel}^{\nu}[m,n]$ where $n \leq m$ and $\nu = r(2\lceil m/n \rceil + 4)$. Then $\mathbf{Adv}^{\text{cca}}_E(q) \leq \frac{2q}{r+1}\big(4\lceil m/n \rceil q\ /\ 2^n\big)^r$.*

Non-binary alphabets. We can replace the binary alphabet $\{0,1\}$ in an unbalanced Feistel scheme with an arbitrary alphabet Σ where $d = |\Sigma| \geq 2$. Regard the characters as numbers $\{0,1,\ldots,d-1\}$ and reinterpret $\oplus$ either as integer addition modulo d^m or as characterwise addition modulo d. The analysis associated to Theorem 7 is trivially lifted to this setting; for example, if $E = \text{Feistel}^{\nu}_d[m,n]$, the radix of the alphabet indicated by the subscript, with $n > m$ and $\nu = r(4\lceil n/m \rceil + 4)$, then $\mathbf{Adv}^{\text{cca}}_E(q) \leq \frac{2q}{r+1}\big((3\lceil n/m \rceil + 3)q\ /\ d^n\big)^r$. We comment that our proof for part (1) of Theorem 7 works for any group operator on Σ^m, but our proof for part (2) does not.

Graphical illustration. Fig. 4 illustrates our CCA-security bounds for $\text{Feistel}^{\nu}[32,96]$ versus $\text{Feistel}^{\nu}[64,64]$. Given an adequate number of rounds, imbalance provably helps.

Unbalanced numeric feistel. We now go on to show security for the numeric variant of the unbalanced Feistel scheme. We begin by defining this. Let $M \geq 2$ and $N \geq 2$ be numbers and let F have signature $F\colon \mathbb{Z}_N \to \mathbb{Z}_M$. Let $\boxplus\colon \mathbb{Z}_M \times \mathbb{Z}_M \to \mathbb{Z}_M$ represent addition modulo M, that is, $a \boxplus b = (a+b) \bmod M$. Consider the permutation $\Psi_F\colon \mathbb{Z}_{MN} \to \mathbb{Z}_{MN}$ that maps $Na+b$ to $Mb + (a \boxplus F(b))$ for every $(a,b) \in \mathbb{Z}_M \times \mathbb{Z}_N$. We call Ψ_F a numeric Feistel (M,N)-permutation and F its round function. Blockcipher $\text{Feistel}^{\nu}_{\sharp}[M,N]\colon \mathcal{K} \times \mathbb{Z}_{MN} \to \mathbb{Z}_{MN}$ has key space $(\text{Func}(\mathbb{Z}_N, Z_M))^{\nu}$. A key $(F_1,\ldots,F_{\nu}) \in \mathcal{K}$ names the permutation $\Psi_{F_1} \circ \cdots \circ \Psi_{F_{\nu}}$ on $\mathbb{Z}_{MN}$, permutations composing from the left. For an illustration, see Fig. 1.

Security of numeric feistel schemes. The following theorem establishes CCA-security for $\text{Feistel}_{\sharp}$. Interpreted asymptotically, the result implies that,

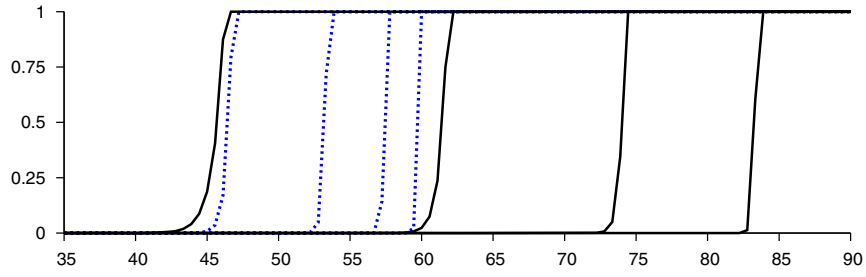

Fig. 4. Unbalanced Feistel versus classical Feistel on a 128-bit string. Proven CCA-security of $\text{Feistel}^{\curvearrowright}[32, 96]$ (bold lines) versus $\text{Feistel}^{\curvearrowright}[64, 64] = \text{FEISTEL}^{\curvearrowright}[64]$ (dashed lines) when $\curvearrowright$ is 18, 36, 72, and 144 (the curves from left to right). The x-axis gives the log base-2 of the number of queries; the y-axis gives an upper bound on an adversary's CCA advantage by Theorems 6 and 7.

with an adequate number of rounds, unbalanced numeric Feistel with a $\mathbb{Z}_N \to \mathbb{Z}_M$ round function withstands a chosen-ciphertext attack to nearly N queries.

Theorem 8. *Fix $M, N \geq 2$, $r \geq 1$.*

1) *Let $E = \text{Feistel}^{\curvearrowright}_{\sharp}[M, N]$ where $N > M$ and $\curvearrowright = r(6\lceil \log_M N \rceil + 4)$. Then $\mathbf{Adv}^{\text{cca}}_E(q) \leq \frac{2q}{r+1}\big((9\lceil \log_M N \rceil + 5)q \;/\; N\big)^r$.*
2) *Let $E = \text{Feistel}^{\curvearrowright}_{\sharp}[M, N]$ where $N \leq M$ and $\curvearrowright = r(2\lceil \log_N M \rceil + 6)$. Then $\mathbf{Adv}^{\text{cca}}_E(q) \leq \frac{2q}{r+1}\big((7\lceil \log_N M \rceil + 7)q \;/\; N\big)^r$.*

Proof ideas. Let us briefly give an overview of the proof; see the full version of this paper [8, Appendix B] for the complete proof. We begin by extending the concepts of *coin* and *collision* of Section 3. The coupling method in Section 3 requires that every pair of queries share coins at each round, if possible. But this does not work here because if M and N are relatively prime, we may find two deterministic queries that *never* yield the same output under such a coupling strategy. Instead, think of coupling as a computer program trying to produce the same output for two different inputs by manipulating the coins. The program first creates a rule for coin-renaming. For example, suppose that each Feistel network is programmed to create a sequence of uniformly random, independent coins. The rule will map each possible value of the random sequence in the first network to a *unique* value of the corresponding sequence in the second network. The program then runs the first input. Now, knowing the exact value of the sequence of coins in the first network, it runs the second input and uses the rule above to specify how the coins of the second network are created. The uniqueness property is to ensure that the round functions in the second network are independent and uniformly random.

5 Alternating Feistel

Defining the schemes. Let m and n be positive integers such that $m \le n$. The blockcipher $\text{FeIsTeL}^{\curvearrowright}[m,n]\colon \mathcal{K} \times \{0,1\}^{m+n} \to \{0,1\}^{m+n}$ consists of $\curvearrowright$ rounds in which the odd rounds are Feistel (m,n)-permutations (contracting) and the even rounds are Feistel (n,m)-permutations (expanding). For simplicity, we assume that $\curvearrowright$ is even. The key space of $\text{FeIsTeL}^{\curvearrowright}[m,n]$ is then $\mathcal{K} = (\text{Func}(n,m) \times \text{Func}(m,n))^{\curvearrowright/2}$. Given integers M and N such that $2 \le M \le N$, we define the blockcipher $\text{FeIsTeL}^{\curvearrowright}_{\sharp}[M,N]\colon \mathcal{K} \times \mathbb{Z}_{MN} \to \mathbb{Z}_{MN}$, with numeric Feistel (M,N) permutations at odd rounds and numeric Feistel (N,M) permutations at even rounds. See Fig. 1 for illustration. We comment that it does not much matter whether one starts with a contracting or expanding round because a security bound with respect to one notion implies the same security bound with respect to the other after one additional round.

Security of alternating feistel. The information-theoretic security of blockciphers FeIsTeL and $\text{FeIsTeL}_{\sharp}$ are established by the following results. Interpreted asymptotically, the result says that, with an adequate number of rounds, alternating Feistel can withstand a chosen-ciphertext attack to nearly N adversarial queries.

Theorem 9. *Fix $r > 0$, $1 \le m \le n$, and $2 \le M \le N$.*

1) *Let $E = \text{FeIsTeL}^{\curvearrowright}[m,n]$ where $\curvearrowright = r\,(12\lceil n/m\rceil + 8)$. Then $\mathbf{Adv}^{\text{cca}}_{E}(q) \le \frac{2q}{r+1}\big((6\lceil n/m\rceil + 3)q \,/\, 2^n\big)^r$.*

2) *Let $E = \text{FeIsTeL}^{\curvearrowright}_{\sharp}[M,N]$ where $\curvearrowright = r\,(12\lceil \log_M N\rceil + 8)$. Then $\mathbf{Adv}^{\text{cca}}_{E}(q) \le \frac{2q}{r+1}\big((6\lceil \log_M N\rceil + 3)q \,/\, N\big)^r$.*

Proof ideas. We give an overview; see the full version of this paper for all details [8, Appendix C]. We consider the generalization of $\text{FeIsTeL}_{\sharp}$ in which the operator $\boxplus$ is replaced by any two group operators on $\mathbb{Z}_M$ and $\mathbb{Z}_N$, regarding FeIsTeL as a special case. While we still follow the framework of Section 3, extending the concepts of *coin* and *collision* is tricky. Following the birthday-bound proof of Black and Rogaway [3] and using the simple coupling method for classical Feistel, one may be tempted to define two types of coins, one for odd rounds and one for even rounds; and, likewise, two types of collisions. This will indeed give rise to a bound, which however falls off with $\min(N,M)$ queries instead of $\max(N,M)$ queries; that is, the approach is only good in the nearly-balanced setting. Instead, we define coins only at odd rounds, and collisions only at even rounds.

We are left with the task of coupling two pools of queries. Coins alone cannot completely determine the outputs, because they dictate only the randomness at odd rounds. However, if we require that the two pools use the same expanding round functions (that control the randomness at even rounds), it suffices to specify how coins evolve. While some specific choice of expanding round functions may give us a poor chance of coupling, the expected value of the success probability is good when those functions are uniformly chosen.

6 Type-1, Type-2, and Type-3 Feistel

Defining the schemes. For illustrations, refer again to Fig. 1.

1) Fix $k \geq 2$ and $n \geq 1$, and let $F\colon \{0,1\}^n \to \{0,1\}^n$ name a permutation $\Psi_F\colon \{0,1\}^{kn} \to \{0,1\}^{kn}$ by way of setting $\Psi_F\big(B_1, \cdots, B_k\big) = \big(B_2 \oplus F(B_1), B_3, \ldots, B_k, B_1\big)$, where $|B_i| = n$. Then Feistel1$^{\varkappa}[k,n]\colon \mathcal{K}\times\{0,1\}^{kn} \to \{0,1\}^{kn}$ is the blockcipher obtained by the $\varkappa$-fold composition of Ψ_F permutations, the key space being $\mathcal{K} = (\mathrm{Func}(n,n))^{\varkappa}$.

2) Assume $k \geq 2$ is even, $n \geq 1$, and $f_i\colon \{0,1\}^n \to \{0,1\}^n$ for every $i \leq k/2$. Let $F = (f_1, \ldots, f_{k/2})$ name a permutation $\Psi_F\colon \{0,1\}^{kn} \to \{0,1\}^{kn}$ by $\Psi_F(B_1, \ldots, B_k) = \big(B_2 \oplus f_1(B_1), B_3, B_4 \oplus f_2(B_3), B_5, \ldots, B_k \oplus f_{k/2}(B_{k-1}), B_1\big)$ where $|B_i| = n$. Then the blockcipher Feistel2$^{\varkappa}[k,n]\colon \mathcal{K}\times\{0,1\}^{kn} \to \{0,1\}^{kn}$ is obtained by the $\varkappa$-fold composition of Ψ_F permutations, the key space being $\mathcal{K} = (\mathrm{Func}(n,n))^{k\varkappa/2}$.

3) Finally, with $k \geq 2$ and $n \geq 1$, consider $f_i\colon \{0,1\}^n \to \{0,1\}^n$ for every $i \leq k-1$. Let $F = (f_1, \ldots, f_{k-1})$ name a permutation $\Psi_F\colon \{0,1\}^{kn} \to \{0,1\}^{kn}$ by $\Psi_F(B_1, \cdots, B_k) = \big(B_2 \oplus f_1(B_1), B_3 \oplus f_2(B_2), \ldots, B_k \oplus f_{k-1}(B_{k-1}), B_1\big)$, where $|B_i| = n$. Then Feistel3$^{\varkappa}[k,n]\colon \mathcal{K} \times \{0,1\}^{kn} \to \{0,1\}^{kn}$ is the blockcipher obtained by the $\varkappa$-fold composition of Ψ_F permutations, the key space being $\mathcal{K} = (\mathrm{Func}(n,n))^{(k-1)\varkappa}$.

Security results. The following results show CCA-security of type-1, type-2, type-3 Feistel variants to $2^{n(1-\varepsilon)}$ queries. Of course this may be a disappointing bound when n is small—and the type-i Feistel variants are in part motivated by a desire to keep n small despite a long block length. But the bound is the best possible, up to the asymptotic behavior, and substantially improves the prior bound in the literature [35].

Theorem 10. *Fix $k, r \geq 1$. Then:*

1) $E = \mathrm{Feistel1}^{\varkappa}[k,n]$, $\varkappa = r(4k-2) \Rightarrow \mathbf{Adv}^{\mathrm{cca}}_{E}(q) \leq \frac{2q}{r+1}\big(2k(k-1)q/2^n\big)^r$.

2) $E = \mathrm{Feistel2}^{\varkappa}[k,n]$, $\varkappa = r(2k+2) \Rightarrow \mathbf{Adv}^{\mathrm{cca}}_{E}(q) \leq \frac{2q}{r+1}\big(2k(k-1)q/2^n\big)^r$.

3) $E = \mathrm{Feistel3}^{\varkappa}[k,n]$, $\varkappa = r(2k+2) \Rightarrow \mathbf{Adv}^{\mathrm{cca}}_{E}(q) \leq \frac{2q}{r+1}\big(4(k-1)^2q/2^n\big)^r$.

The proofs for the results above can be found in the full version of this paper [8, Appendix D].

Acknowledgments

The authors gratefully acknowledge the support of NSF grant 0904380. Thanks particularly to program directors Richard Beigel and Lenore Zuck.

References

1. Anderson, R., Biham, E.: Two practical and provably secure block ciphers: BEAR and LION. In: Gollmann, D. (ed.) FSE 1996. LNCS, vol. 1039, pp. 113–120. Springer, Heidelberg (1996)
2. Bellare, M., Rogaway, P., Spies, T.: The FFX mode of operation for format-preserving encryption (draft 1.1). NIST submission (February 2010), http://csrc.nist.gov/groups/ST/toolkit/BCM/modes_development.html
3. Black, J., Rogaway, P.: Ciphers with arbitrary finite domains. In: Preneel, B. (ed.) CT-RSA 2002. LNCS, vol. 2271, pp. 114–130. Springer, Heidelberg (2002)
4. 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)
5. Brightwell, M., Smith, H.: Using datatype-preserving encryption to enhance data warehouse security. In: 20th NISSC Proceedings, pp. 141–149 (1997), http://csrc.nist.gov/nissc/1997
6. Coppersmith, D.: Luby-Rackoff: four rounds is not enough. Technical Report RC 20674, IBM (December 1996)
7. Feistel, H., Notz, W., Smith, J.: Some cryptographic techniques for machine-to-machine data communications. Proc. of the IEEE 63, 1545–1554 (1975)
8. Hoang, V., Rogaway, P.: On generalized Feistel networks. Full version of this paper. Cryptology ePrint report 2010/301, May26 (2010)
9. Jutla, C.: Generalized birthday attacks on unbalanced Feistel networks. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 186–199. Springer, Heidelberg (1998)
10. Luby, M., Rackoff, C.: How to construct pseudorandom permutations from pseudorandom functions. SIAM Journal on Computing 17(2), 373–386 (1988); Earlier version in CRYPTO 1985
11. Lucks, S.: Faster Luby-Rackoff ciphers. In: Gollmann, D. (ed.) FSE 1996. LNCS, vol. 1039, pp. 189–203. Springer, Heidelberg (1996)
12. Maurer, U.: A simplified and generalized treatment of Luby-Rackoff pseudorandom permutation generator. In: Rueppel, R.A. (ed.) EUROCRYPT 1992. LNCS, vol. 658, pp. 239–255. Springer, Heidelberg (1993)
13. Maurer, U., Pietrzak, K.: The security of many-round Luby-Rackoff pseudo-random permutations. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 544–561. Springer, Heidelberg (2003)
14. Maurer, U., Pietrzak, K., Renner, R.: Indistinguishability amplification. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 130–149. Springer, Heidelberg (2007)
15. Mirinov, I. (Not so) random shuffles of RC4. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 304–319. Springer, Heidelberg (2002)
16. Moriai, S., Vaudenay, S.: On the pseudorandomness of top-level schemes of block ciphers. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 289–302. Springer, Heidelberg (2000)
17. 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)
18. Nachef, V.: Generic attacks on alternating unbalanced Feistel schemes. Cryptology ePrint report 2009/287, June 16 (2009)

19. Naor, M., Reingold, O.: On the construction of pseudo-random permutations: Luby-Rackoff revisited. Journal of Cryptology 12(1), 29–66 (1997)
20. Nyberg, K.: Generalized Feistel networks. In: Kim, K.-c., Matsumoto, T. (eds.) ASIACRYPT 1996. LNCS, vol. 1163, pp. 91–104. Springer, Heidelberg (1996)
21. Patarin, J.: About Feistel schemes with six (or more) rounds. In: Vaudenay, S. (ed.) FSE 1998. LNCS, vol. 1372, pp. 103–121. Springer, Heidelberg (1998)
22. Patarin, J.: Generic attacks on Feistel schemes. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 222–238. Springer, Heidelberg (2001)
23. Patarin, J.: Luby-Rackoff: 7 Rounds are enough for $2^{n-\varepsilon}$ security. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 513–529. Springer, Heidelberg (2003)
24. Patarin, J.: New results on pseudorandom permutation generators based on the DES scheme. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 301–312. Springer, Heidelberg (1992)
25. Patarin, J.: Security of balanced and unbalanced Feistel schemes with linear non equalities. Cryptology ePrint report 2010/293. May 17 (2010)
26. Patarin, J.: Security of random Feistel schemes with 5 or more rounds. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 106–122. Springer, Heidelberg (2004)
27. Patarin, J., Nachef, V., Berbain, C.: Generic attacks on unbalanced Feistel schemes with contracting functions. In: Lai, X., Chen, K. (eds.) ASIACRYPT 2006. LNCS, vol. 4284, pp. 396–411. Springer, Heidelberg (2006)
28. Patarin, J., Nachef, V., Berbain, C.: Generic attacks on unbalanced Feistel schemes with expanding functions. In: Kurosawa, K. (ed.) ASIACRYPT 2007. LNCS, vol. 4833, pp. 325–341. Springer, Heidelberg (2007)
29. PCI Security Standards Council. Payment Card Industry (PCI) Data Security Standard: Requirements and Security Assessment Procedures, version 1.2.1 (July 2009), `www.pcisecuritystandards.org`
30. Schneier, B., Kelsey, J.: Unbalanced Feistel networks and block cipher design. In: Gollmann, D. (ed.) FSE 1996. LNCS, vol. 1039, pp. 121–144. Springer, Heidelberg (1996)
31. Smith, J.: The design of Lucifer: a cryptographic device for data communications. IBM Research Report RC 3326. IBM T.J. Watson Research Center, Yorktown Heights, New York, USA (April 15, 1971)
32. Thorp, E.: Nonrandom shuffling with applications to the game of Faro. Journal of the American Statistical Association 68, 842–847 (1973)
33. Vaudenay, S.: Provable security for block ciphers by decorrelation. In: Meinel, C., Morvan, M. (eds.) STACS 1998. LNCS, vol. 1373, pp. 249–275. Springer, Heidelberg (1998)
34. Yun, A., Park, J., Lee, J.: On Lai-Massey and quasi-Feistel ciphers. In: Designs, Codes and Cryptography, Online First (2010)
35. Zheng, Y., Matsumoto, T., Imai, H.: On the construction of block ciphers provably secure and not relying on any unproved hypotheses. In: Brassard, G. (ed.) CRYPTO 1989. LNCS, vol. 435, pp. 461–480. Springer, Heidelberg (1990)

A Proof for Unbalanced Feistel — Theorem 7

Given a query X to $\text{Feistel}^{\mathcal{n}}[m,n]$, its *coin* at round t is the string $A \oplus F(B)$, where F is the round function at round t and (A,B) is the round-$(t-1)$ output, with $|A| = m$ and $|B| = n$. We say that two queries *collide* at time t if their outputs at round t have the same last n bits. We begin with the following.

Lemma 11. *In the blockcipher* Feistel$^{\nu}[m,n]$*, the chance that two distinct non-adaptive queries have the same coin at round* $t \geq 1$ *is at most* 2^{-m}*.*

Proof. Suppose that the Feistel network receives distinct non-adaptive queries X_1 and X_2. For each $i \in \{1,2\}$, let (A_i, B_i) be the output at round $t-1$ of X_i, where $|A_i| = m$ and $|B_i| = n$. The queries X_1 and X_2 collide at time t if and only if $A_1 \oplus F(B_1) = A_2 \oplus F(B_2)$, with F being the round function at round t. This occurs with probability 2^{-m} if B_1 and B_2 differ, because F is uniformly random. If $B_1 = B_2$ then so are A_1 and A_2, which contradicts the hypothesis that the two queries are distinct. □

Contracting round functions. We first consider the security of the blockcipher Feistel$^{\nu}[m,n]$ with $n > m$ (that is, the round functions are contracting). Later we show how to deal with expanding round functions.

Lemma 12. *In the blockcipher* Feistel$^{\nu}[m,n]$ *with* $n > m$*, the chance that two distinct non-adaptive queries collide at time* $t > \lceil n/m \rceil$ *is at most* $3/2^{n+1}$*.*

Proof. Suppose that the Feistel network receives distinct non-adaptive queries X_1 and X_2. We shall prove by induction on b that for any $b \leq n$, the probability that outputs at round $t > \lceil b/m \rceil$ of the two queries have the same last b bits is at most $3/2^{b+1}$. The claim of this lemma corresponds to the special case $b = n$.

First consider the base case $b < m$. For each $i \in \{1,2\}$, let (A_i, B_i) be the output at round $t-1$ of X_i, where $|A_i| = m$ and $|B_i| = n$. The last m-bit substring of the round-t output of X_i is $A_i \oplus F(B_i)$, with F being the round function at round t. If B_1 and B_2 differ then the probability that outputs at round t of the two queries have the same last b bits is at most 2^{-b}, because F is uniformly random. If $B_1 = B_2$ then the two queries have the same coin at round $t-1$, which by Lemma 11 occurs with probability at most 2^{-m}. Hence, by union bound, the chance that the two queries have the same last b bits is at most $2^{-b} + 2^{-m} \leq 3/2^{b+1}$.

Next consider $b \geq m$ and assume that the chance round-$(t-1)$ outputs of the two queries have the same last $b-m$ bits is at most $3/2^{b-m+1}$. The outputs at round t of the two queries have the same last b bits if and only if (i) they have the same coin at round t, which by Lemma 11 occurs with probability at most 2^{-m}, and (ii) their output at round $t-1$ have the same lat $b-m$ bits, which occurs with probability at most $3/2^{b-m+1}$ by induction hypothesis. As the round functions in the network are independent, the chance that both (i) and (ii) occur is at most $2^{-m} \cdot 3 \,/\, 2^{b-m+1} = 3/2^{b+1}$. □

We now prove NCPA-security of Feistel$^{r(2\lceil n/m \rceil+2)}[m,n]$. Employing Lemma 2 then yields the desired result. Let $b = \lceil n/m \rceil + 1$. Suppose that the network receives nonadaptive distinct queries $X_1, \ldots, X_q$. We shall use a similar strategy as in the proof of Theorem 4. Fix an integer $\ell \leq q-1$. For every $i \leq \ell$, let $U_i = X_i$ and let $U_{\ell+1}$ be chosen uniformly from $\{0,1\}^{n+m} \setminus \{U_1, \ldots, U_\ell\}$. We shall construct another Feistel$^{2rb}[m,n]$ for the queries $U_1, \ldots, U_\ell$. Let $X_i(t)$ and $U_i(t)$ be the outputs at round t of X_i and U_i respectively. It suffices to define

the coupling in the first $2b$ rounds, and then show that the probability that $X_i(2b) \neq U_i(2b)$ for some $i \leq \ell + 1$ is at most $3b\ell \,/\, 2^n$.

THE COUPLING. In the first b rounds, for every $i \leq \ell$, we use the same coin to update $X_i(t)$ and $U_i(t)$, and couple $X_{\ell+1}(t)$ and $U_{\ell+1}(t)$ in an arbitrary way. In the next b rounds, we couple as follows.

- If U_i collides with some previous query U_j at time t then the coin at round $t+1$ of U_i is defined so as to ensure consistency with the earlier query.
- Suppose that, in the new Feistel network, U_i does not collide with any previous query at time t. If the query X_i collides with some previous query X_j at time t then we choose a string uniformly from $\{0,1\}^{n+m}$ to be the coin of U_i at round $t+1$. Otherwise, the coin of X_i at round $t+1$ is uniformly distributed over $\{0,1\}^{n+m}$ and U_i will use exactly the same coin at round $t+1$.

Note that U_i and X_i always have the same output at round t, for every $i \leq \ell$ and every t. Consider the event Coll that in either Feistel networks, the $(\ell+1)$-th query collides with some previous query at some time $t \in \{b, \ldots, 2b-1\}$. From Lemma 12, each such collision occurs with probability at most $3/2^{n+1}$. Summing over the two Feistel networks, b rounds, and ℓ previous queries shows that the probability Coll occurs is at most $3b\ell \,/\, 2^n$. Unless Coll occurs, $U_{\ell+1}$ and $X_{\ell+1}$ will share the coins at the rounds $b+1, \ldots, 2b$, and then have identical outputs at the round $2b$. Hence the chance that we fail to couple at round $2b$ cannot exceed $3b\ell \,/\, 2^n$.

EXPANDING ROUND FUNCTIONS. We follow the same proof as before, but Lemma 12 is replaced by the following result.

Lemma 13. *In the blockcipher* $\mathrm{Feistel}^{\nu}[m,n]$ *with* $n \leq m$*, the chance that two distinct non-adaptive queries collide at time* $t \geq \lceil m/n \rceil$ *is at most* $\lceil m/n \rceil \,/\, 2^n$.

Proof. Suppose that the Feistel network receives distinct non-adaptive queries X_1 and X_2. For each $i \in \{1,2\}$, let (A_i, B_i) be the output at round $t-1$ of X_i, where $|A_i| = m$ and $|B_i| = n$. The queries X_1 and X_2 collide at time t if and only if the two strings $A_1 \oplus F(B_1)$ and $A_2 \oplus F(B_2)$ have the same last n bits, with F being the round function at round t. This occurs with probability 2^{-n} if B_1 and B_2 differ, because F is uniformly random. If $B_1 = B_2$ then A_1 and A_2 must have the same last n bits. In other words, the round-$(t-1)$ outputs of the two queries must agree at the last $2n$ bits. Repeating this argument leads us to examine the case that for every $j < \lceil m/n \rceil$ the round-$(t-j)$ outputs of the two queries must agree at the last $(j+1)n$ bits. When this chain of reasoning stops at round $t - \lceil m/n \rceil + 1$, the outputs at that round must have the same last m bits. In other words, the queries have the same coin at that round, which by Lemma 11 occurs with probability at most $2^{-m} \leq 2^{-n}$. Hence by union bound, the chance that the two queries collide at time t is at most $\lceil m/n \rceil \,/\, 2^n$. □