On Tweaking Luby-Rackoff Blockciphers

David Goldenberg[1], Susan Hohenberger[2*], Moses Liskov[1],
Elizabeth Crump Schwartz[1], and Hakan Seyalioglu[1**]

[1] The College of William and Mary,
{dcgold,mliskov,eacrum}@cs.wm.edu, hakan.seyalioglu@gmail.com
[2] The Johns Hopkins University, susan@cs.jhu.edu

Abstract. Tweakable blockciphers, first formalized by Liskov, Rivest, and Wagner [12], are blockciphers with an additional input, the *tweak*, which allows for variability. An open problem proposed by Liskov et al. is how to construct tweakable blockciphers without using a pre-existing blockcipher. There are many natural questions in this area: is it significantly more efficient to incorporate a tweak directly? How do direct constructions compare to existing techniques? Are these direct constructions *optimal* and for what levels of security? How large of a tweak can be securely added? In this work, we explore these questions for Luby-Rackoff blockciphers. We show that tweakable blockciphers can be created directly from Luby-Rackoff ciphers, and in some cases show that direct constructions of tweakable blockciphers are more efficient than previously known constructions.

1 Introduction

A *blockcipher*, also known as a *pseudorandom permutation*, is a pair of algorithms E and D. The encryption algorithm E takes two inputs – a key K and a message block M, and produces a ciphertext block C of the same length as M, while the decryption algorithm D reverses this process. A blockcipher is considered secure if, for a random secret key K, the cipher is indistinguishable from a random permutation.

A *tweakable blockcipher* takes an extra input, the *tweak* ,(T), that is used only to provide variation and is not kept secret. Unlike changing the key, changing the tweak should involve minimal extra cost. A tweakable blockcipher is considered secure if it is indistinguishable from a family of random permutations indexed by the tweak. The Hasty Pudding Cipher by Schroeppel [21] was the first to introduce an auxiliary blockcipher input called a "spice" and Liskov, Rivest, and Wagner [12] later formalized the notion of tweakable blockciphers. Liskov et al. describe two levels of

* Supported by an NDSEG Fellowship and NSF grant CT-0716142.
** Partially supported by a Monroe Grant and a Cummings Grant.

security: a secure (CPA) tweakable blockcipher is one that is indistinguishable from a random permutation family to any adversary that may make chosen plaintext queries, while a strongly secure (CCA) tweakable blockcipher is pseudorandom even to an adversary that may also make chosen ciphertext queries.

Tweakable blockciphers have many practical applications. Liskov et al. describe how they can be used to implement secure symmetric encryption and authenticated encryption. Halevi and Rogaway [9, 10] suggest an immediate application to private storage where the tweak is set to be the memory address of an enciphered block; and thus, the encryptions of two blocks with the same plaintext are not likely to look the same and yet decryption remains straightforward. Tweakable blockciphers have also been studied in a variety of other contexts [1, 11, 20, 2].

Feistel Blockciphers. Feistel blockciphers [6] have been an actively studied class of constructions since Horst Feistel invented them in 1973. In particular, Luby and Rackoff showed how to construct a pseudorandom permutation from a pseudorandom function by composing three (or four in the case of CCA security) Feistel permutations [13]. We call this construction the Luby-Rackoff blockcipher. In 1996, Lucks [14] described an optimization for the secure 3-round version by replacing the first round with a universal hash function. Shortly afterwards, Naor and Reingold [15] provided the analogous optimization for the strongly secure 4-round cipher, replacing both the first and last rounds with a more general type of function. In 2001, Ramzan [18] formally studied many variations on the Luby-Rackoff cipher. Patarin gave proofs of security for certain constructions against unbounded adversaries with access to exponentially many queries, albeit assuming the individual round functions are random functions rather than pseudorandom. Specifically, Patarin proved security for 7 rounds against $q \ll 2^k$ queries, where the blockcipher input is of size $2k$ [16], and later improved this to show that 5 rounds is sufficient, both for chosen-plaintext and chosen-ciphertext attacks [17], which remains the best proven security level for Feistel ciphers. Dodis and Puniya recently provided a combinatorial understanding of Feistel networks when the round functions are *unpredictable* rather than pseudorandom [5].

Our Work. Liskov, Rivest, and Wagner [12] give two constructions for tweakable blockciphers, each one constructed from an underlying blockcipher. Subsequent work has also taken this approach; Halevi and Rogaway's EMD and EME modes [9, 10] and Rogaway's XEX mode [20] were all blockcipher modes of operation. The only examples of specific tweakable blockciphers are the Hasty Pudding [21] and the Mercy [4] ciphers.

One open problem proposed by Liskov et al. was to study how to incorporate tweaks into existing blockciphers, or design tweakable blockciphers directly. In this work, we perform a systematic study of issues relating to directly tweaking Luby-Rackoff blockciphers. We analyze the approach of including a tweak by XOR-ing the tweak value into one or more places in the dataflow. This natural model for adding a tweak changes the cipher minimally. Also, approaches involving more direct cryptographic processing of the tweak (e.g. hashing the tweak) have a significant additional cost associated with changing the tweak.

Our Contributions. We present tweakable Luby-Rackoff blockciphers, for both CPA and CCA security, and against both polynomial-time adversaries, and against unbounded adversaries with $q \ll 2^k$ queries[1], where k is half the size of the input (matching the best result for ordinary blockciphers [17]). Specifically, we construct tweakable blockciphers:

– CPA-secure against polynomial adversaries in 4 rounds (Theorem 3)
– CCA-secure against polynomial adversaries in 6 rounds (Theorem 8)
– CPA-secure against $q \ll 2^k$ queries in 7 rounds (Theorem 4)
– CCA-secure against $q \ll 2^k$ queries in 10 rounds (Theorem 9)

Recall that for polynomial adversaries CPA-security requires 3 rounds whereas CCA-security requires 4. It is thus natural to wonder if our constructions are optimal. We prove our constructions against polynomial adversaries are indeed round-optimal in our model (Theorems 1 and 7). Furthermore, we show that any construction of 6 or fewer rounds in our model can be attacked with $O(2^{k/2})$ queries (Table 1), so our construction of Theorem 4 is also round-optimal. In addition, the attacks used to prove the round-optimality of our constructions, as well as our extension of the proof methods of Naor and Reingold, help to form the theoretical foundation necessary for the secure design of tweakable blockciphers regardless of construction, as well as shedding light on the difficulties in adding a tweak to Feistel-based blockciphers such as RC6 [19] and MARS [3].

We also explicitly address the problem of incorporating tweaks of arbitrary length, an important issue not addressed in the literature.[2] We show that our CPA-secure constructions can incorporate additional blocks

[1] That is, any non-negative $q < 2^k$ such that $q2^{-k}$ is negligible.

[2] Using tweaks of arbitrary length has been considered for tweakable symmetric encryption [8], but not for one-block constructions. Certain applications require different, specific tweak sizes, and one may want to allow longer tweaks to include more

of tweak at the cost of 1 round per block (Theorems 11 and 14), and that our CCA-secure constructions may be similarly extended at the cost of 2 rounds per block of tweak (Theorems 12 and 15).

2 Definitions

A *tweakable blockcipher* is a triple of algorithms $(\widetilde{G}, \widetilde{E}, \widetilde{D})$ for key generation, encryption, and decryption, respectively. We restrict our attention to tweakable blockciphers where $\widetilde{G}(\cdot)$, $\widetilde{E}_K(\cdot,\cdot)$, and $\widetilde{D}_K(\cdot,\cdot)$ are all efficiently computable algorithms; and where the correctness property holds; that is, for all M, T, and for all keys $K \in \widetilde{G}(1^k)$, $\widetilde{D}_K(\widetilde{E}_K(M,T),T) = M$. We also generally assume that $\widetilde{G}(1^k)$ draws keys uniformly at random from $\{0,1\}^{p(k)}$ for some polynomial p.

We have two notions of security: (1) chosen-plaintext secure (CPA) and (2) chosen-ciphertext secure (CCA). Security is defined in terms of both a polynomial and an exponential adversary; polynomial adversaries are limited to a number of queries and computations polynomial in the message size, whereas an exponential adversary is allowed unlimited computation, but is bounded by an exponential number of queries relative to the message size.

Definition 1. *Over all adversaries with access to an encryption oracle, the maximum advantage is defined as:*

$$\mathsf{ADV\text{-}TBC}_{\mathsf{K}}(\widetilde{E}, \widetilde{D}, q, t) = \max_{\mathcal{A}} : |\Pr[\mathcal{A}^{\widetilde{E}_K(\cdot,\cdot)}(1^k) = 1] - \Pr[\mathcal{A}^{\Pi}(1^k) = 1]|$$

where (1) *for all* k*,* K *is generated by* $\widetilde{G}(1^k)$*,* (2) Π *is a random permutation family parameterized by its second input, and* (3) $\mathcal{A}$ *is allowed to run for* t *steps and make at most* q *oracle queries.*

Definition 2. *Over all adversaries with access to an encryption and decryption oracle, the maximum advantage is defined as:*

$$\mathsf{ADV\text{-}STBC}_{\mathsf{K}}(\widetilde{E}, \widetilde{D}, q, t) = \max_{\mathcal{A}} : |\Pr[\mathcal{A}^{\widetilde{E}_K(\cdot,\cdot),\widetilde{D}_K(\cdot,\cdot)}(1^k) = 1] - \Pr[\mathcal{A}^{\Pi,\Pi^{-1}}(1^k) = 1]|$$

where (1) *for all* k*,* K *is generated by* $\widetilde{G}(1^k)$*,* (2) Π, Π^{-1} *are a pseudorandom permutation family and its inverse, and* (3) $\mathcal{A}$ *is allowed to run for* t *steps and make at most* q *oracle queries.*

A tweakable blockcipher is CPA secure if for all k, for q queries and time t, $\mathsf{ADV\text{-}TBC}_{\mathsf{K}}(\widetilde{E}, \widetilde{D}, q, t)$ is negligible in k. A tweakable cipher is said to be polynomially-secure if q and t are polynomial in k. If t is unspecified, then it may be unbounded. We define CCA security in the same manner.

information. Indeed, this was the motivation for Schroeppel to allow spice values of 512 bits in the Hasty Pudding Cipher [21].

3 The Feistel Blockcipher

Recall the formula for the Feistel blockcipher [6] on input $M = (L^0, R^0)$:

$$L^{i+1} = R^i$$
$$R^{i+1} = f_{i+1}(R^i) \oplus L^i$$

where the output after n rounds is (L^n, R^n), and each f_i is a pseudorandom function specified by the key. Further recall that the 3-round Feistel construction is secure against chosen plaintext attacks, and the 4-round construction is secure against chosen ciphertext attack [13].

3.1 Notation

In order to talk about where to add a tweak, we must first establish some notation. Unless otherwise specified, the tweaks we refer to are a *half-block* in length; that is, on input M of size $2k$, the tweak is of size k. As we will later see, a blockcipher may allow for longer tweaks; we think of these as "multiple tweaks," as conceptually, the longer tweak can be thought of as being composed of multiple tweaks, each of the same size.

For an n-round Luby-Rackoff construction, a single half-block of tweak can conceivably be XOR-ed in at any of the following unique locations: $\mathcal{L}_0$, $\mathcal{L}_1$, …, $\mathcal{L}_n$, $\mathcal{R}_0, \mathcal{R}_{0.5}$, $\mathcal{R}_1$, $\ldots, \mathcal{R}_{n-0.5}, \mathcal{R}_n$. Let this set be denoted by Λ_n. We illustrate the Λ_3 (3-round) locations in Figure 1.

Let T^λ be the XOR of all the tweaks used at location $\lambda \in \Lambda_n$. The formula for our construction is:

$$L^{i+1} = R^i \oplus T^{\mathcal{R}_i}$$
$$R^{i+1} = f_{i+1}(R^i \oplus T^{\mathcal{R}_i} \oplus T^{\mathcal{R}_{i+0.5}}) \oplus L^i \oplus T^{\mathcal{L}_i}$$

We use "$\mathcal{BC}(n, \lambda)$" to refer to the tweakable blockcipher construction where the number of Luby-Rackoff rounds is n and a tweak T^λ is XOR-ed in at some location $\lambda \in \Lambda_n$. To denote adding multiple tweaks, we write "$\mathcal{BC}(n, \lambda_1, \ldots, \lambda_t)$", where $T^{\lambda_i} = T_i$ is the tweak for location λ_i and different locations each have their own independent tweak. Thus, in such a construction, the tweak size is tk.

We might also want to denote adding the *same* tweak value at two or more locations. We write this as "$\mathcal{BC}(n, \lambda_1 + \lambda_2)$", where the implication of using the *compound* location $\lambda_1 + \lambda_2$ is that $T^{\lambda_1} = T^{\lambda_2}$. Of course, we may also consider a construction with multiple tweaks, each of which may be a compound location; we use the obvious notation for this. We use the symbol Γ to denote a (possibly) compound tweak location.

In Λ_n, we have listed all tweaks at ".5" locations, i.e., $\mathcal{R}_{l+0.5}$ for some l. However, we do not have to consider these locations.

Lemma 1. *For all m, $\mathcal{R}_{m+0.5}$ is equivalent to $\mathcal{R}_m + \mathcal{L}_{m+1}$.*

Lemma 2. *For all $0 \leq m < n$, $\mathcal{L}_m$ is equivalent to $\mathcal{R}_{m+1}$.*

Since $\mathcal{L}_m$ and $\mathcal{R}_{m+1}$ are equivalent, we will use them interchangeably. This starts us off with a reduced set of tweakable constructions to study including tweaks at locations $\mathcal{L}_n$, $\mathcal{R}_0$, …, $\mathcal{R}_n$ and all combinations thereof.

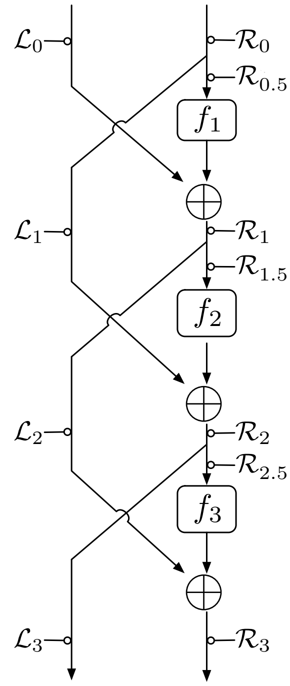

Fig. 1. An illustration of Λ_3; the locations at which to XOR a tweak of length $|M|/2$ for 3-round LR.

4 Tweakable Blockciphers With CPA Security

In this section, we focus on achieving CPA security. In the next section, we will discuss the stronger CCA notion of security.

We begin by presenting some general results that hold for an arbitrary number of rounds. These results will help us to narrow down the possibilities for secure constructions and to prove the optimality of our final construction. As stated in Section 3, the set of possibly secure constructions includes those with tweaks at locations $\mathcal{L}_n$, $\mathcal{R}_0$, …, $\mathcal{R}_n$ and all combinations thereof. However, we remark in Lemma 3 that we do not need to consider all possible locations, and that some locations can be simulated without directly tweaking the blockcipher; this important observation is used frequently throughout the paper.

Lemma 3. *For all n, without loss of generality, we can consider only constructions that never use the tweak locations $\mathcal{L}_n$ $\mathcal{R}_n$, $\mathcal{R}_0$, or $\mathcal{R}_1$, even in compound locations, and even when considering CCA security.*

Proof. We can simulate oracle queries with or without the tweaks in $\mathcal{L}_n$ $\mathcal{R}_n$, $\mathcal{R}_0$, or $\mathcal{R}_1$. To simulate a query $(L^0, R^0, T_1, \ldots, T_t)$ to a construction with these tweaks, we make a query $(L^0 \oplus T^{\mathcal{R}_1}, R^0 \oplus T^{\mathcal{R}_0}, T_1, \ldots, T_t)$ to the construction without these tweaks to obtain (L^n, R^n), and we return $(L^n \oplus T^{\mathcal{L}_n}, R^n \oplus T^{\mathcal{R}_n})$. Decryption queries can be simulated similarly. □

The set of tweak locations we need to consider is thus reduced to $\{\mathcal{R}_2 \ldots, \mathcal{R}_{n-1}\}$. From here on, we consider Λ_n to be $\{\mathcal{R}_2, \ldots, \mathcal{R}_{n-1}\}$.

Lemma 4. *For all n, $\mathcal{BC}(n, \mathcal{R}_{n-1})$ is not CPA-secure.*

Proof. We use a 2-query attack. If we query (L, R, T) to get (L_1^n, R_1^n), and then query (L, R, T') to get (L_2^n, R_2^n), then $L_1^n \oplus L_2^n = T \oplus T'$. $\square$

Thus, we arrive at our first round-specific conclusion.

Theorem 1 (No Tweakable 3-Round Constructions). *For all $n < 4$ and all compound locations Γ of elements in Λ_n, $\mathcal{BC}(n, \Gamma)$ is not CPA-secure.*

Proof. This follows from Lemmas 3 and 4, and the set $\{\mathcal{R}_2, \ldots, \mathcal{R}_{n-2}\}$ being empty for $n = 3$. $\square$

4.1 Secure Locations

We have reduced the set of possible secure single tweak locations to $\{\mathcal{R}_2, \ldots, \mathcal{R}_{n-2}\}$. We now show that each of these locations are secure for $n \geq 4$. However, first we must define $\epsilon - \mathsf{ARCU}_2$ hash functions and introduce some related work.

Definition 3. *An $\epsilon - \mathsf{ARCU}_2$ ("Almost Right-Collision-avoiding Universal") hash function family is a hash function family given a range of $\{0,1\}^{2k}$ with the property that for all $x \neq y$, the probability that $h_R(x) = h_R(y)$ is at most $2^{-k} + \epsilon$, over the choice of h, where h_R denotes the right half of the output of h.*

Naor and Reingold [15] create a secure blockcipher using two Luby-Rackoff rounds in combination with a potentially less expensive function.

Theorem 2 (Naor-Reingold). *If E denotes two Luby-Rackoff rounds with truly random round functions, and h is drawn from an $\epsilon - \mathsf{ARCU}_2$ hash function family, then $E \circ h$ is indistinguishable (in a CPA attack) from a random function.*

Using Definition 3 and Theorem 2, we are able construct CPA-secure tweakable blockciphers.

Theorem 3 (Several Tweakable n-Round Constructions (for $n \geq 4$)). *For all $n \geq 4$ and $m \in \{2, \ldots, n-2\}$, $\mathcal{BC}(n, \mathcal{R}_m)$ is CPA-secure against polynomially bounded adversaries.*

Proof. We can capitalize on Theorem 2 as follows. We will prove that when we let $h(L, R, T) = (L \oplus f_{m-1}(R) || R \oplus T \oplus f_m(L \oplus f_{m-1}(R)))$ over random choice of f_{m-1} and f_m, these conditions hold. Here, h is comprised of the last two rounds of the construction before the tweak, including the tweak. Once we prove this, the result will follow: the first $m-2$ rounds are a permutation, so if h' is comprised of the first m rounds, it will be $\epsilon - \mathsf{ARCU}_2$ if h is. Furthermore, since $m \leq n-2$, there are at least 2 more rounds to follow; any further rounds are another permutation and pseudorandomness will be maintained.

Lemma 5. *The family* $h(L, R, T) = (L \oplus f_1(R) || R \oplus T \oplus f_2(L \oplus f_1(R)))$*, where* f_1 *and* f_2 *are randomly chosen over the domain of all functions from* k *bits to* k *bits, is* $\epsilon - \mathsf{ARCU}_2$*, for* $\epsilon = 2^{-k} + 2^{-2k}$*.*

Proof. Let $x = (L, R, T)$ and $y = (L', R', T')$, where $x \neq y$. Note that if $R \neq R'$ then the probability that $L \oplus f_1(R) = L' \oplus f_1(R')$ is the probability that $f_1(R) = L \oplus L' \oplus f_1(R')$ which is 2^{-k}. Similarly, if $R = R'$ but $L \neq L'$ then $L \oplus f_1(R) \neq L' \oplus f_1(R')$. In either case, the probability that $L \oplus f_1(R) = L' \oplus f_1(R')$ is at most 2^{-k}. Finally, if $R = R'$ and $L = L'$ then $T \neq T'$ so $h_R(L, R, T) = h_R(L, R, T') \oplus T \oplus T' \neq h_R(L, R, T')$.

The probability that $h_R(L, R, T) = h_R(L', R', T')$ given that $L \oplus f_1(R) \neq L' \oplus f_1(R')$ is the probability that $f_2(L \oplus f_1(R)) = R \oplus R' \oplus f_2(L' \oplus f_1(R')$, which is 2^{-k}, so the probability we hit a collision is at most $(1 - 2^{-k})(2^{-k}) + 2^{-k} = 2^{-k} + 2^{-2k} + 2^{-k} = 2^{-k} + \epsilon$. □

From the Lemma, if all the round functions are random, then the h we are interested in is $\epsilon - \mathsf{ARCU}_2$. By Theorem 2, $\mathcal{BC}(n, \mathcal{R}_m)$ is indistinguishable from a random function if all round functions are random. Therefore, $\mathcal{BC}(n, \mathcal{R}_m)$ must be CPA secure if its round functions are pseudorandom (since random functions are indistinguishable from random permutation families). This completes the proof of Theorem 3. □

Corollary 1 (CPA Security In 4 Rounds). $\mathcal{BC}(4, \mathcal{R}_2)$ *is CPA-secure and round-optimal.*

Proof. This follows directly from Theorems 1 and 3. □

4.2 Exponential Attacks

In this section, we investigate the security of tweakable blockcipher constructions against an adversary who is capable of making an exponential number of queries. We provide general attacks against several types of tweakable constructions built from Luby-Rackoff permutations. In this section, we assume all round functions are ideal, in other words, that they are uniform random functions.[3] We consider a construction secure against exponentially many queries if the probability of any computationally unbounded adversary allowed $q \ll 2^k$ queries to distinguish the construction from a random permutation family is negligible in k. These attacks appertain to constructions with both single and compound tweak locations (where the same tweak value is XOR-ed in multiple locations) and are used to prove that all constructions of less than 7 rounds can be distinguished from a random permutation family in $O(2^{\frac{k}{2}})$ queries.

Lemma 6. *For any $0 \leq r < n$, $\mathcal{BC}(n, \mathcal{R}_{r+0.5})$ is insecure against $O(2^{\frac{k}{2}})$ queries.*

Proof. The attack is as follows: fix the message and query with $2^{\frac{k}{2}}$ different tweaks. The probability that two different queries lead to the same output is negligible for a random permutation family. However, the probability that two queries lead to a collision in this construction is not negligible. On each query, the internal values stay constant until the input to f_{r+1}. Since we have made $2^{\frac{k}{2}}$ queries to an ideal round function, we can expect with non-negligible probability to get a collision on the output of f_{r+1} for two distinct queries. If we get such a collision, notice the entire output ciphertext will collide. □

Corollary 2. *For any $0 \leq r < n$, $\mathcal{BC}(n, \mathcal{R}_{r+0.5} + \mathcal{R}_{r+1})$ is insecure against $O(2^{\frac{k}{2}})$ queries.*

Proof. The attack is identical to that used in Lemma 6, except that instead of expecting a collision of the type $f_{r+1}(R^r \oplus T) = f_{r+1}(R^r \oplus T')$, we expect a collision of the type $f_{r+1}(\mathcal{R}^r \oplus T) \oplus T = f_{r+1}(\mathcal{R}^r \oplus T') \oplus T'$. □

Lemma 7. *For any $0 \leq r < n$, $\mathcal{BC}(n, \mathcal{R}_{r+0.5} + \mathcal{R}_{n-1})$ is insecure against $O(2^{\frac{k}{2}})$ queries.*

Proof. For this proof we will first need a result from probability.

Lemma 8 (Strong Birthday Lemma). *For all $k > 1$, there exists an $m < 1.2 \times 2^{\frac{k}{2}}$ such that if p is the probability of picking an element twice when selecting m elements from a 2^k-element set with replacement uniformly at random, then p and $1 - p$ are both non-negligible in k.*

[3] This is the standard assumption when we want to prove security in a setting where the adversary has beyond-polynomial capabilities [16, 17].

Proof. For proof of the Strong Birthday Lemma, see full version [7]. □

The attack is as follows: Compute the m described in Lemma 8. Keep the message constant and query with m different tweaks. The probability that two ciphertexts are such that $L^n \oplus T = L'^n \oplus T'$ is significantly higher for the actual construction than for a random permutation family. Since $m \leq 1.2 \times 2^{\frac{k}{2}}$, this attack can be performed by an exponential adversary.

Notice that the internal values of any pair of queries are the same up to the input of f_{r+1}. For every query, f_{r+1} receives a different input (as the input is a fixed value XOR-ed by the tweak). Since the round functions are ideal, the event of getting a collision on two outputs of f_{r+1} with m different queries reduces to the event of picking the same element twice as described in Lemma 8; say that probability is p. Notice that if such a collision happens, we always get a collision of the type, $L^n \oplus T = L'^n \oplus T'$.

Assume that the outputs of f_{r+1} are distinct for each of the m queries. Notice that in order to have a collision of two R^{n-2} values, it must be true that the L^{n-2} values differ for both queries, because the intervening rounds act as a permutation. Therefore, we will get a collision on R^{n-2} if and only if we have a collision of the type:

$$f_{n-2}(L^{n-2}) \oplus L^{n-3} = f_{n-2}(L'^{n-2}) \oplus L'^{n-3}.$$

Since the probability of such a collision for any two queries is either 2^{-k} or 0 (in the case that the L^{n-2} values coincide), we can bound the probability of having such a collision above by $\frac{(1.2)^2 2^k}{2 \times 2^k} = .72$ since $m \leq 1.2 \times 2^{\frac{k}{2}}$. Therefore, in this case, with probability greater equal to .28, we can assume all R^{n-2} values are distinct. Notice:

$$L^n \oplus T = L'^n \oplus T' \Leftrightarrow f_{n-1}(R^{n-2}) \oplus L^{n-2} \oplus T = f_{n-1}(R'^{n-2}) \oplus L'^{n-2} \oplus T'.$$

The probability of such an event occurring over m queries with distinct R^{n-2} and ideal round functions is, again, p. Therefore, the overall probability of getting at least two ciphertexts with the described property is at least $p + (1-p)(.28p)$.

If the construction we are given is the random permutation family, the probability of getting the coincidence described is clearly p. Therefore the difference in probabilities of this event happening for the tweakable construction and the random permutation family is at least $p + .28p(1-p) - p = .28p(1-p)$. Since p and $1-p$ are non-negligible in k (by Lemma 8), this value is also non-negligible, and therefore our attack successfully distinguishes the two constructions. □

Corollary 3. *$\mathcal{BC}(n, \mathcal{R}_{r+0.5} + \mathcal{R}_{r+1} + \mathcal{R}_{n-1})$ is insecure against $O(2^{\frac{k}{2}})$ queries.*

Proof. The generalization of Lemma 7 to Lemma 3 is identical to the extension of Lemma 6 to Lemma 2. □

These four attacks can be used to attack every tweakable Luby-Rackoff blockcipher of 6 or fewer rounds. A rundown of which general attack applies for each construction can be found in Table 1. We do not include $\mathcal{L}_1, \mathcal{R}_1, \mathcal{L}_6$ or $\mathcal{R}_6$ in the possible locations, or their equivalent constructions of Table 1 since they can be simulated away by Lemma 3.

Tweak Locations		
Location	Equivalent	Attack
$\mathcal{R}_2$	$\mathcal{R}_{0.5}$	Lemma 6
$\mathcal{R}_3$	$\mathcal{R}_{1.5}$	Lemma 6
$\mathcal{R}_4$	$\mathcal{R}_{4.5}$	Lemma 6
$\mathcal{R}_5$	N/A	Lemma 4
$\mathcal{R}_2 + \mathcal{R}_3$	$\mathcal{R}_{1.5} + \mathcal{R}_2$	Corollary 2
$\mathcal{R}_2 + \mathcal{R}_4$	$\mathcal{R}_{2.5}$	Lemma 6
$\mathcal{R}_2 + \mathcal{R}_5$	$\mathcal{R}_{0.5} + \mathcal{R}_5$	Lemma 7
$\mathcal{R}_3 + \mathcal{R}_4$	$\mathcal{R}_{3.5} + \mathcal{R}_4 + \mathcal{R}_5$	Corollary 3
$\mathcal{R}_3 + \mathcal{R}_5$	$\mathcal{R}_{3.5}$	Lemma 6
$\mathcal{R}_4 + \mathcal{R}_5$	$\mathcal{R}_{4.5} + \mathcal{R}_5$	Corollary 2
$\mathcal{R}_2 + \mathcal{R}_3 + \mathcal{R}_4$	$\mathcal{R}_{2.5} + \mathcal{R}_3$	Corollary 2
$\mathcal{R}_2 + \mathcal{R}_3 + \mathcal{R}_5$	$\mathcal{R}_{1.5} + \mathcal{R}_2 + \mathcal{R}_5$	Corollary 3
$\mathcal{R}_2 + \mathcal{R}_4 + \mathcal{R}_5$	$\mathcal{R}_{2.5} + \mathcal{R}_5$	Lemma 7
$\mathcal{R}_3 + \mathcal{R}_4 + \mathcal{R}_5$	$\mathcal{R}_{3.5} + \mathcal{R}_4$	Corollary 2
$\mathcal{R}_2 + \mathcal{R}_3 + \mathcal{R}_4 + \mathcal{R}_5$	$\mathcal{R}_{2.5} + \mathcal{R}_3 + \mathcal{R}_5$	Corollary 3

Table 1. All possible 6-round tweakable blockcipher constructions and the corresponding lemmas that prove the constructions are insecure.

4.3 A Tweakable Construction Secure for $q \ll 2^k$ Queries

We now show a 7-round Luby-Rackoff construction that is secure against an adversary allowed $q \ll 2^k$ queries.

Theorem 4. *$\mathcal{BC}(7, \mathcal{R}_3 + \mathcal{L}_3)$ is CPA-secure for $q \ll 2^k$ queries.*

Proof. To prove that this construction is a secure tweakable blockcipher we utilize the following theorem from Patarin [16]:

Theorem 5 (Patarin). *Let F be a function from $2k$ bits to $2k$ bits. If F has the property that for $q \ll 2^k$ queries, the probability of having $l > O(k)$ indices such that $R_{i_1} = R_{i_2} = R_{i_3} = ...R_{i_l}$ is negligible, (where R_{i_j} is the right half of the j'th output of F), and on distinct inputs F has only a negligible probability of a full collision on its outputs, then $E \circ F$,(where E is a four-round Luby-Rackoff function), is indistinguishable from random for $q \ll 2^k$ input queries.*

We decompose our 7-round construction into two functions, F and E, where F is the first three rounds, including the XOR-ed tweak at both $\mathcal{L}_3$ and $\mathcal{R}_3$,[4] and E is the last four rounds. It is obvious that E is a four-round Luby-Rackoff function. To prove that F has the properties enumerated in Theorem 5, we need to prove the following two properties about F.

Lemma 9. *F is such that for any two distinct queries, the probability of the outputs being equal is $O(2^{-2k})$ and the probability of the right halves of the outputs being equal is $O(2^{-k})$.*

Proof. For proof see full version [7]. □

So long as the queries the adversary makes do not produce a full collision on F or a multi-collision on the right half of the output of F, the responses are indistinguishable from random. Therefore, the queries of the adversary are independent of the outputs of F so long as the required conditions hold. By Lemma 9, the probability of an overall collision in $q \ll 2^k$ queries is $O(q^2 2^{-2k})$ which is negligible. Similarly, the probability of an l-way multicollision on the right is $O(q^l 2^{-(l-1)k}) = O(2^k (q2^{-k})^l)$. Since $q < 2^{k(1-\epsilon)}$ for some ϵ, we know that $(q2^{-k})^l < (2^{-k\epsilon})^l = 2^{-kl\epsilon}$. If $l \geq k \geq 2/\epsilon$, which will be true for sufficiently large k, this probability is bounded by 2^{-k}. Thus, F satisfies the necessary properties with all but a negligible probability, which completes our proof of Theorem 4. □

5 Tweakable Blockciphers With CCA Security

In this section, we study the problem of achieving CCA security. An important observation to make in constructing a CCA-secure tweakable blockcipher is a distinguishing attack we will call the *four-message attack*, which is a type of Boomerang attack [22]. The attack can be performed by any adversary with access to encryption and decryption oracles, E and D respectively. To perform the attack, the adversary makes four queries:

1. For an arbitrary message M and tweak T, obtain $C = E(M, T)$.
2. For an arbitrary tweak $T' \neq T$, obtain $C' = E(M, T')$.
3. Obtain $M' = D(C', T)$.
4. Obtain $C'' = E(M', T')$. If $C = C''$; output 1, otherwise output 0.

A wide class of tweakable blockciphers fall to the four-message attack:

[4] Although $\mathcal{L}_3$ is equivalent to $\mathcal{R}_4$, we think of this construction as using $\mathcal{L}_3$, so that we can conceptually split the function this way.

Theorem 6 (Four Message Attack). *Suppose that* $g_1 : \{0,1\}^n \to \{0,1\}^l$ *is an injective function that is invertible on its domain, that* $g_2 : \{0,1\}^t \to \{0,1\}^l$ *is any deterministic function, and that* $g_3 : \{0,1\}^l \to \{0,1\}^n$ *is a function such that for all* C *and* T *there exists a unique* A *such that* $g_3(A \oplus g_2(T)) = C$*. Then the construction* $\widetilde{E}_K(M,T) = g_3(g_2(T) \oplus g_1(M))$ *is not CCA-secure.*

Proof. Note that $C = g_3(g_2(T) \oplus g_1(M))$, $C' = g_3(g_2(T') \oplus g_1(M))$. Now if we decrypt C' with tweak T, we obtain $M' = g_1^{-1}(g_2(T') \oplus g_2(T) \oplus g_1(M))$. When we encrypt M' under tweak T', we get $C'' = g_3(g_2(T') \oplus g_1(g_1^{-1}(g_2(T') \oplus g_2(T) \oplus g_1(M))) = g_3(g_2(T') \oplus g_2(T') \oplus g_2(T) \oplus g_1(M)) = g_3(g_2(T) \oplus g_1(M)) = C$. □

Note in particular that if both g_1 and g_3 are permutations, the conditions are satisfied. This has immediate consequences:

Corollary 4. *For all* $n, \mathcal{R}_m \in \Lambda_n$*, both* $\mathcal{BC}(n, \mathcal{R}_m)$ *and* $\mathcal{BC}(n, \mathcal{R}_m + \mathcal{R}_{m+1})$ *are not CCA-secure.*

Proof. Here, g_1 is the permutation described by the m rounds of Luby-Rackoff before the tweak, $g_2(T) = 0^k || T$ for $\mathcal{BC}(n, \mathcal{R}_m)$ and $g_2(T) = T||T$ for $\mathcal{BC}(n, \mathcal{R}_m + \mathcal{R}_{m+1})$, and g_3 is the remaining $n - m$ rounds. Clearly g_1 and g_3 are permutations, so the four message attack applies. □

This shows that if we are to be able to add a half-block of tweak to the construction anywhere, it must be used at multiple locations, and those locations must be separated by at least one round.[5] In fact, however, a one round distance will not suffice:

Lemma 10. *For all* $n, \mathcal{R}_m \in \Lambda_n$*,* $\mathcal{BC}(n, \mathcal{R}_m + \mathcal{R}_{m+2})$ *is not CCA-secure, and* $\mathcal{BC}(n, \mathcal{R}_m + \mathcal{R}_{m+1} + \mathcal{R}_{m+2})$ *is also not CCA-secure.*

Proof. To simplify, recall that $\mathcal{R}_m + \mathcal{R}_{m+2}$ is equivalent to $\mathcal{R}_{m+0.5}$ by Lemma 1. Noticing this makes it clear why this is unlikely to be secure, in light of the previous two corollaries, but we still have some work to do.

Here, we use the four-message attack again, but this time, suppose g_1 and g_3 are not permutations. Rather, if (L, R) is the output of the first m rounds of the Luby-Rackoff permutations, then $g_1(M)$ is the $3k$ bit response (L, R, R). Notice that $g_2(T)$ is $0^{2k}||T$, and $g_3(A, B, C)$ computes the remaining rounds, computing $L^{m+1} = B$ and $R^{m+1} = f_m(C) \oplus A$,

[5] This shows, along with Lemma 10, that an adversary making a CCA attack with XOR injection will be able to succeed, regardless of the location of the XOR.

and continuing from there. Note that $g_3(g_2(T) \oplus g_1(M))$ is the output we get from applying $\mathcal{BC}(n, \mathcal{R}_{m+0.5})$ to M with tweak T. For the $\mathcal{BC}(n, \mathcal{R}_m + \mathcal{R}_{m+1} + \mathcal{R}_{m+2})$ construction, this is just the same as $\mathcal{BC}(n, \mathcal{R}_{m+0.5} + \mathcal{L}_m)$, and change g_2 so that it produces $T||0^k||T$ rather than $0^{2k}||T$. Clearly g_1 is injective and invertible, and g_3 has unique inverses of the proper form, which we can find by inverting the tweakable blockcipher and noting the values in the proper place. Doing so requires the tweak T, but the answer is unique regardless, or we wouldn't have unique decryption. By Theorem 6, neither of these constructions are CCA-secure. □

Theorem 7. *For all $n < 6$ and all compound locations Γ of elements in Λ_n, $\mathcal{BC}(n, \Gamma)$ is not CCA-secure.*

Proof. In order to construct a CCA-secure tweakable blockcipher, we must use the tweak at (minimally) $\mathcal{R}_m$ and $\mathcal{R}_{m+d}$ for some $d \geq 3$. And naturally, m and $m + d$ must be in the range $2, \ldots, n - 1$ since all other locations can be simulated. For $n \leq 5$ no such pair of locations exists. □

Therefore, the first construction that can be CCA-secure is $\mathcal{BC}(6, \mathcal{R}_2 + \mathcal{R}_5)$, and is in fact a secure construction!

Theorem 8. *$\mathcal{BC}(6, \mathcal{R}_2 + \mathcal{R}_5)$ is a CCA-secure tweakable blockcipher.*

Proof. For proof, see full version [7]. □

5.1 CCA Security Against Exponential Attacks

Theorem 9. *$\mathcal{BC}(10, \mathcal{L}_3 + \mathcal{R}_3 + \mathcal{L}_7 + \mathcal{R}_7)$ is CCA-secure for $q \ll 2^k$ queries.*

Proof. In order to construct a tweakable blockcipher secure against CCA exponential attacks, we use a theorem of Patarin [17]:

Theorem 10 (Patarin). *Let F and F' be functions from $2k$ bits to $2k$ bits. If F and F'^{-1} each have the property that for $q \ll 2^k$ queries, the probability of having $l > O(k)$ indices such that $R_{i_1} = R_{i_2} = R_{i_3} = ...R_{i_l}$ is negligible, (where R_{i_j} is the right half of the j'th output of F or F'^{-1}), and on distinct inputs F (and F'^{-1}) has only a negligible probability of a full collision on its outputs, then $F' \circ E \circ F$,(where E is a four-round Luby-Rackoff function), is indistinguishable from random against chosen-ciphertext attack for $q \ll 2^k$ input queries.*

In our construction, the first three rounds, including the tweaks at $\mathcal{L}_3$ and $\mathcal{R}_3$, form F, and the last three rounds, including the tweaks at $\mathcal{L}_7$

and $\mathcal{R}_7$, form F'. F'^{-1} is just the same as F, except with distinct round functions. Both F and F'^{-1} meet the properties of Theorem 10, as we have shown in our proof of Lemma 9. $\mathcal{BC}(10, \mathcal{L}_3 + \mathcal{R}_3 + \mathcal{L}_7 + \mathcal{R}_7) = F' \circ E \circ F$, and is therefore CCA-secure against $q \ll 2^k$ queries. □

6 Allowing Longer Tweaks

In our previous results, all tweaks were assumed to be a half block in length. It may be desirable however, to have tweaks of arbitrary lengths. We can always lengthen a tweak that is less than a half block, by padding it in a deterministic way. However, increasing the length of a tweak beyond a half block in length does not follow easily. It may be useful to have constructions that are still secure with longer tweaks, as one usual way of choosing a tweak is to include data with it that makes it unique [21]. The longer the tweak, the more data can be included.

Tweakable Blockciphers with Longer Tweaks. For t half-blocks of tweak, we show how to construct a CPA-secure tweakable blockcipher in $t + 3$ rounds and a CCA-secure tweakble blockcipher in $2t + 4$ rounds.

Theorem 11. *For all n, one can use $n - 3$ half-blocks of tweak but no more. Specifically, $\mathcal{BC}(n, \mathcal{R}_2, \ldots, \mathcal{R}_{n-2})$ is secure, but any construction $\mathcal{BC}(n, \Gamma_1, \ldots, \Gamma_t)$ for $t > n - 3$ is not secure.*

Theorem 12. *For all n, the tweakable blockcipher $\mathcal{BC}(2n, \mathcal{R}_2 + \mathcal{R}_{2n-1}, \mathcal{R}_3 + \mathcal{R}_{2n-2}, \ldots, \mathcal{R}_{n-1} + \mathcal{R}_{n+2})$ is a CCA-secure tweakable blockcipher.*

Proof. For proof of Theorem 11 and Theorem 12 see full version [7].

Longer Tweaks with Exponential Security. Next, we focus on constructing Luby-Rackoff based tweakable blockciphers which are secure against an unbounded adversary with $q \ll 2^k$ queries. For t half-blocks of tweak, we show how to construct a CPA-secure tweakable blockcipher in $t + 6$ rounds and give a CCA-secure tweakable blockcipher in $2t + 8$ that meets this security goal. These constructions are based on a $t + 2$ round function F designed to meet the properties required by Patarin.

Theorem 13. *Let $\mu_i = \mathcal{L}_{i+2}$ if $i \equiv 1$ or $i \equiv 2 \bmod 4$, let $\mu_i = \mathcal{L}_{i+2} + \mathcal{L}_1$ if $i \equiv 3 \bmod 4$, and $\mu_i = \mathcal{L}_{i+2} + \mathcal{L}_2$ if $i \equiv 0 \bmod 4$. Let $\mu'_i = \mu_i + \mathcal{R}_i$ if $i \not\equiv 2 \bmod 4$, and $\mu'_i = \mu_i + \mathcal{R}_i + \mathcal{L}_1$ otherwise. Then let F be $\mathcal{BC}(n + 2, \mu_1, \ldots, \mu_{n-1}, \mu'_n)$. F is a function such that for $q \ll 2^k$ queries, the probability of having $l = O(k)$ indices such that $R_{i_1} = R_{i_2} = R_{i_3} = \ldots R_{i_l}$*

is negligible, (where R_{i_j} is the right half of the j'th output of F), and on q distinct inputs F has only a negligible probability of a full collision on its outputs.

Proof. For proof, see full paper [7]. □

Theorem 14. *$E \circ F$ is a tweakable blockcipher with t tweaks that is secure against any unbounded adversary with at most $q \ll 2^k$ queries, where E is a four-round Luby-Rackoff cipher.*

Proof. This follows from Theorem 13 and Theorem 5. Note that $E \circ F$ requires a total of $t + 6$ rounds.

Theorem 15. *$F' \circ E \circ F$ is a tweakable blockcipher with t tweaks that is CCA-secure against any unbounded adversary with at most $q \ll 2^k$ queries, where E is a four-round Luby-Rackoff cipher, F' is the inverse of the F described above, with new independent round functions.*

Proof. This follows from Theorem 13 and Theorem 10. Here, $F' \circ E \circ F$ requires $2(t + 2) + 4 = 2t + 8$ rounds.

7 Conclusion

Table 2 summarizes our constructions, compared to regular blockciphers and the second construction of Liskov et al. [12]. This table shows that our results are better for CPA constructions, equivalent for CCA against polynomial attacks, and worse for CCA against exponential ones.

Security Level	Blockciphers	Prior TBCs [12]	This paper
CPA with polynomial queries	3 rounds [13]	3 + 2 rounds/tweak	3 + 1 round/tweak
CPA with $\ll 2^k$ queries	5 rounds [17]	5 + 2 rounds/tweak	6 + 1 round/tweak
CCA with polynomial queries	4 rounds [13]	4 + 2 rounds/tweak	4 + 2 rounds/tweak
CCA with $\ll 2^k$ queries	5 rounds [17]	5 + 2 rounds/tweak	8 + 2 rounds/tweak

Table 2. Number of rounds required for each construction. The prior tweakable construction we consider is $\widetilde{E_{K,h}}(M, T) = h(T) \oplus E_K(M \oplus h(T))$, where h is an ϵ−AXU$_2$ hash function; subsequent constructions are similar. The natural way to realize the hash function would be to simply use two random functions on the tweak, one for each half of the data stream. Although Liskov et al. do not explicitly consider arbitrary tweak length, their construction and proof can be easily extended to do so.

We conclude with some open problems: (1) incorporating tweaks securely into other blockcipher structures, (2) direct, specific design of tweakable blockciphers (Luby-Rackoff or otherwise) and (3) improving the provable level of security for tweakable blockciphers in general.

Acknowledgments. We thank Ronald L. Rivest and several anonymous reviewers for their helpful comments.

References

1. M. Bellare and T. Kohno. A Theoretical Treatment of Related-Key Attacks: PKA-PRPs, RKA-PRFs, and Applications. In *Advances in Cryptology – EUROCRYPT '03*, volume 2656 of LNCS, pages 491–506, 2003.
2. J. Black, M. Cochran, and T. Shrimpton. On The Impossibility of Highly-Efficient Blockcipher-Based Hash Functions. In *Advances in Cryptology – EUROCRYPT '05*, volume 3494 of *LNCS*, pages 526–541, 2005.
3. C. Burwick, D. Coppersmith, E. D'Avignon, R. Gennaro, S. Halevi, C. Jutla, S. M. Matyas Jr., L. O'Connor, M. Peyravian, D. Safford, and N. Zunic. MARS - A Candidate Cipher for AES. In *NIST AES proposal*, June 1998.
4. P. Crowley. Mercy: A Fast Large Block Cipher for Disk Sector Encryption. In *Fast Software Encryption*, volume 1978 of LNCS, pages 49–63, 2000.
5. Y. Dodis and P. Puniya. Feistel networks made public, and applications. In *EUROCRYPT '07*, volume 4515 of LNCS, pages 534–554, 2007.
6. H. Feistel. Cryptography and Computer Privacy. *Scientific American*, 15–23, 1973.
7. D. Goldenberg, S. Hohenberger, M. Liskov, E. Crump Schwartz, and H. Seyalioglu. Full version of this paper, Cryptology ePrint Archive, Report 2007/350.
8. S. Halevi EME*: Extending EME to Handle Arbitrary-Length Messages with Associated Data. In *Proceedings of INDOCRYPT 2004*, pp. 315–327, 2004.
9. S. Halevi and P. Rogaway. A Tweakable Enciphering Mode. In *Advances in Cryptology – CRYPTO '03*, volume 2729 of LNCS, pages 482–499, 2003.
10. S. Halevi and P. Rogaway. A Parallelizable Enciphering Mode. In *Topics in Cryptology – CT-RSA '04*, volume 2964 of LNCS, pages 292–304, 2004.
11. A. Joux. Cryptanalysis of the EMD Mode of Operation. In *Advances in Cryptology – EUROCRYPT '03*, volume 2656 of LNCS, pages 1–16, 2003.
12. M. Liskov, R. Rivest, and D. Wagner. Tweakable Block Ciphers. In *Advances in Cryptology – CRYPTO '02*, volume 2442 of LNCS, pages 31–46, 2002.
13. M. Luby and C. Rackoff. How To Construct Pseudorandom Permutations from Pseudorandom Functions. *SIAM J. of Computing*, 17(2):373–386, 1988.
14. S. Lucks. Faster Luby-Rackoff Ciphers. In *Fast Software Encryption*, volume 1039 of LNCS, pages 189–203, 1996.
15. M. Naor and O. Reingold. On the Construction of Pseudorandom Permutations: Luby-Rackoff Revisited. *Journal of Cryptology*, 12(1):29–66, 1999.
16. J. Patarin. Luby-Rackoff: 7 Rounds are Enough for $2^{n(1-\varepsilon)}$ Security. In *Advances in Cryptology – CRYPTO*, volume 2729 of LNCS, pages 513–529, 2003.
17. J. Patarin. Security of Random Feistel Schemes with 5 or More Rounds. In *Advances in Cryptology – CRYPTO*, volume 3152 of LNCS, pages 106–122, 2004.
18. Z. Ramzan. *A Study of Luby-Rackoff Ciphers*. PhD thesis, MIT, 2001.
19. R. Rivest, M. Robshaw, R. Sidney, and Y. L. Yin. The RC6 Block Cipher. In *First AES conference*, August 1998.
20. P. Rogaway. Efficient Instantiations of Tweakable Blockciphers and Refinements to Mode OCB and PMAC. In *ASIACRYPT*, vol. 3329, pages 16–31, 2004.
21. R. Schroeppel. The Hasty Pudding Cipher. NIST AES proposal, available `http://www.cs.arizona.edu/~rcs/hpc`, 1998.
22. D. Wagner. The Boomerang Attack. In *Fast Software Encryption*, vol. 1636 of LNCS, pages 156–170, 1999.