

New Method for Upper Bounding the Maximum Average Linear Hull Probability for SPNs

Liam Keliher¹, Henk Meijer¹, and Stafford Tavares²

¹ Department of Computing and Information Science
Queen's University at Kingston, Ontario, Canada, K7L 3N6
{keliher,henk}@cs.queensu.ca

² Department of Electrical and Computer Engineering
Queen's University at Kingston, Ontario, Canada, K7L 3N6
tavares@ee.queensu.ca

Abstract. We present a new algorithm for upper bounding the maximum average linear hull probability for SPNs, a value required to determine provable security against linear cryptanalysis. The best previous result (Hong et al. [9]) applies only when the linear transformation branch number (\mathcal{B}) is M or $(M + 1)$ (maximal case), where M is the number of s-boxes per round. In contrast, our upper bound can be computed for any value of \mathcal{B} . Moreover, the new upper bound is a function of the number of rounds (other upper bounds known to the authors are not). When $\mathcal{B} = M$, our upper bound is consistently superior to [9]. When $\mathcal{B} = (M + 1)$, our upper bound does not appear to improve on [9]. On application to Rijndael (128-bit block size, 10 rounds), we obtain the upper bound $UB = 2^{-75}$, corresponding to a lower bound on the data complexity of $\frac{8}{UB} = 2^{78}$ (for 96.7% success rate). Note that this does not demonstrate the existence of a such an attack, but is, to our knowledge, the first such lower bound.

Keywords: substitution-permutation networks, linear cryptanalysis, maximum average linear hull probability, provable security

1 Introduction

The substitution-permutation network (SPN) [6] is a fundamental block cipher architecture designed to be a practical implementation of Shannon's principles of *confusion* and *diffusion* [15], through the use of substitution and linear transformation (LT), respectively. There has been a recent increase in interest in SPNs, in part because their simplicity lends itself to analysis, and, from an implementation viewpoint, because they tend to be highly parallelizable. This interest will no doubt be spurred on by the recent adoption of Rijndael (a straightforward SPN) as the U.S. Government Advanced Encryption Standard (AES)[5].

The two most powerful cryptanalytic attacks on block ciphers are generally considered to be linear cryptanalysis (LC) [11] and differential cryptanalysis (DC) [2]. There exists a strong duality between these two attacks which allows

certain results related to one of the attacks to be translated into the corresponding results for the other attack [1][12]. This duality applies to the work of this paper; for this reason we will limit our focus to LC.

In carrying out LC, an attacker typically computes a vector called the *best linear characteristic*, for which the associated *linear characteristic probability* (LCP) is maximal. This LCP allows the attacker to estimate the number of chosen plaintexts required to mount a successful attack. In [14], Nyberg showed that the use of linear characteristics underestimates the success of LC. In order to guarantee *provable security*, a block cipher designer needs to consider *approximate linear hulls* instead of linear characteristics, and the *maximum average linear hull probability* instead of the LCP of the best linear characteristic.

In this paper we present a new method for computing an upper bound on the maximum average linear hull probability for SPNs. The best previous result is that of Hong et al. [9], which applies only to SPNs with highly diffusive LTs. In contrast, our method can be applied to an SPN with any LT (computation time may vary). Moreover, the upper bound we compute is a function of the number of rounds of the SPN; all other upper bounds known to the authors do not depend on the number of rounds. When the diffusiveness of the LT is one less than maximum (the relevant definition is given in Section 5.3), our upper bound is consistently superior to that of [9]. For LTs with maximum diffusiveness, our upper bound does not appear to improve on [9].

Application of our method to Rijndael (128-bit block size, 10 rounds), which involved extensive computation, yielded the upper bound $UB = 2^{-75}$, for a corresponding lower bound on the data complexity of LC of $\frac{8}{UB} = 2^{78}$ (for 96.7% success rate—see Section 3). Note that this does not demonstrate the existence of a such an attack, but is, to our knowledge, the first such lower bound.

Conventions

In what follows, $\{0, 1\}^d$ denotes the set of all d -bit vectors, which we view as row vectors. For a vector or matrix \mathbf{w} , \mathbf{w}' denotes the transpose of \mathbf{w} . We adopt the convention that numbering of the bits of a binary vector proceeds from left to right, beginning at 1. The Hamming weight of a vector \mathbf{x} is written $wt(\mathbf{x})$. If \mathbf{Z} is a random variable (r.v.), $E[\mathbf{Z}]$ denotes the expected value of \mathbf{Z} . And we use $\#\mathcal{A}$ to indicate the number of elements in the set \mathcal{A} .

2 Substitution-Permutation Networks

A block cipher is a bijective mapping from N bits to N bits (N is called the *block size*) parameterized by a bitstring called a *key*, denoted \mathbf{k} . Common block sizes are 64 and 128 bits. The input to a block cipher is called a *plaintext*, and the output is called a *ciphertext*.

An SPN encrypts a plaintext through a series of R simpler encryption steps called *rounds*. The input to round r ($1 \leq r \leq R$) is first bitwise XOR'd with an N -bit *subkey*, denoted \mathbf{k}^r , which is typically derived from the key, \mathbf{k} , via a separate *key-scheduling algorithm*. The *substitution stage* then partitions the resulting vector into M subblocks of size n ($N = Mn$), which become the inputs

to a row of bijective $n \times n$ *substitution boxes (s-boxes)*—bijective mappings from $\{0, 1\}^n$ to $\{0, 1\}^n$. Finally, the *permutation stage* applies an invertible LT to the output of the s-boxes (classically, a bitwise permutation). Often the permutation stage is omitted from the last round. A final subkey, \mathbf{k}^{R+1} , is XOR'd with the output of round R to form the ciphertext. Figure 1 depicts an example SPN with $N = 16$, $M = n = 4$, and $R = 3$.

We assume the most general situation for the key, namely, that \mathbf{k} is an *independent key* [1], a concatenation of $(R + 1)$ independent subkeys—symbolically, $\mathbf{k} = \langle \mathbf{k}^1, \mathbf{k}^2, \dots, \mathbf{k}^{R+1} \rangle$. We use \mathcal{K} to denote the set of all independent keys.




Fig. 1. SPN with $N = 16$, $M = n = 4$, $R = 3$

3 Linear Cryptanalysis

Linear cryptanalysis (LC) was introduced by Matsui in 1993 [11]. The more powerful version is known as Algorithm 2 (Algorithm 1 extracts only a single subkey bit). As applied to SPNs, Algorithm 2 can be used to extract the first subkey, \mathbf{k}^1 . Once \mathbf{k}^1 is known, the first round can be stripped off, and LC can be reapplied to obtain \mathbf{k}^2 , and so on.

Let \mathbf{P} , \mathbf{C} , and \mathbf{X} be r.v.'s representing the plaintext, ciphertext, and intermediate input to round 2, respectively. The attacker attempts to identify the best correlation between the parity of a subset of the bits of \mathbf{X} and the parity of a subset of the bits of \mathbf{C} . Symbolically, the attacker wants *masks* $\mathbf{a}, \mathbf{b} \in \{0, 1\}^N \setminus \mathbf{0}$ which maximize the following *linear probability*:

$$LP_{\mathbf{k}}(\mathbf{a} \rightarrow \mathbf{b}) \stackrel{\text{def}}{=} (2 \cdot \text{Prob}\{\mathbf{a} \bullet \mathbf{X} = \mathbf{b} \bullet \mathbf{C}\} - 1)^2, \quad (1)$$

for a fixed key, \mathbf{k} (the symbol \bullet denotes the inner product over $\text{GF}(2)$). Note that $LP_{\mathbf{k}}(\mathbf{a} \rightarrow \mathbf{b}) \in [0, 1]$. Given \mathbf{a} and \mathbf{b} , the attack proceeds as in Figure 2.

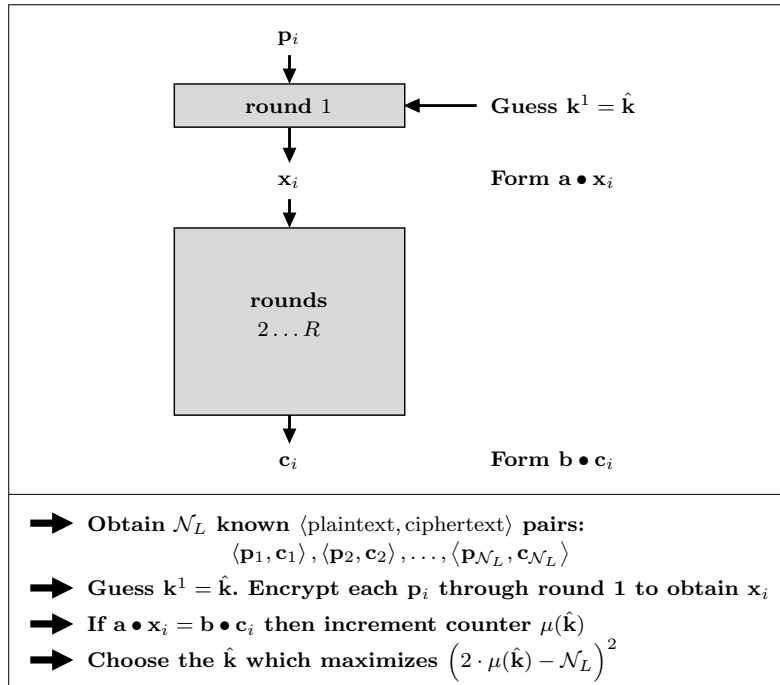


Fig. 2. Summary of linear cryptanalysis (Algorithm 2)

The probability that Algorithm 2 will determine the correct value of \mathbf{k}^1 increases as the number of known $\langle \text{plaintext, ciphertext} \rangle$ pairs, \mathcal{N}_L , is increased. The value \mathcal{N}_L is called the *data complexity* of the attack—this is what the attacker wants to minimize. Given an assumption about the behavior of round-1 output [11], Matsui shows that if $\mathcal{N}_L = \frac{c}{LP_{\mathbf{k}}(\mathbf{a} \rightarrow \mathbf{b})}$, then Algorithm 2 has the success rates in the following table, for various values of the constant, c .

c	2	4	8	16
Success rate	48.6%	78.5%	96.7%	99.9%

3.1 Notational Generalization

In describing Algorithm 2, we have discussed input and output masks (\mathbf{a} and \mathbf{b} , respectively) and the associated linear probability for rounds $2 \dots R$ of an R -round SPN. It is useful to consider these and other related concepts as applying to any $T \geq 2$ consecutive rounds of an SPN. Hereafter, unless specified otherwise, terms such as “first round” and “last round” are relative to the T rounds under consideration. For Algorithm 2, then, $T = R - 1$, and the “first round,” or “round 1,” is actually round 2 of the SPN. And for simplicity, we will always assume that the LT is absent from round T (this does not affect LC).

4 Linear Characteristics

For fixed $\mathbf{a}, \mathbf{b} \in \{0, 1\}^N$, direct computation of $LP_{\mathbf{k}}(\mathbf{a} \rightarrow \mathbf{b})$ is generally infeasible, first since it requires encrypting all N -bit vectors through rounds $1 \dots T$, and second because $LP_{\mathbf{k}}(\mathbf{a} \rightarrow \mathbf{b})$ depends on the unknown key, \mathbf{k} . The latter is usually handled by working not with $LP_{\mathbf{k}}(\mathbf{a} \rightarrow \mathbf{b})$, but with the *average* (expected) value of $LP_{\mathbf{k}}(\mathbf{a} \rightarrow \mathbf{b})$ over all independent keys $\mathbf{k} \in \mathcal{K}$, denoted $E_T[\mathbf{a}, \mathbf{b}]$:

$$E_T[\mathbf{a}, \mathbf{b}] \stackrel{\text{def}}{=} E[LP_{\mathbf{K}}(\mathbf{a} \rightarrow \mathbf{b})] \quad (2)$$

(\mathbf{K} is an r.v. uniformly distributed over \mathcal{K}). The implicit assumption is that $LP_{\mathbf{k}}(\mathbf{a} \rightarrow \mathbf{b})$ is approximately equal to $E_T[\mathbf{a}, \mathbf{b}]$ for most values of \mathbf{k} (Harper et al. refer to this as the *Hypothesis of fixed-key equivalence* [7]). The data complexity of Algorithm 2 for masks \mathbf{a} and \mathbf{b} is now taken to be $\mathcal{N}_L = \frac{c}{E_T[\mathbf{a}, \mathbf{b}]}$. The problem of computational complexity is usually treated by approximating $E_T[\mathbf{a}, \mathbf{b}]$ through the use of *linear characteristics* (or simply *characteristics*).

4.1 One-Round and Multi-round Linear Characteristics

Note that the linear probability in (1) can be defined for any binary mapping—in particular, for a bijective $n \times n$ s-box, S . Let $\alpha, \beta \in \{0, 1\}^n$, and let \mathbf{X} be an r.v. uniformly distributed over $\{0, 1\}^n$. Define

$$LP^S(\alpha \rightarrow \beta) \stackrel{\text{def}}{=} (2 \cdot \text{Prob}\{\alpha \bullet \mathbf{X} = \beta \bullet S(\mathbf{X})\} - 1)^2 \quad (3)$$

$$q \stackrel{\text{def}}{=} \max_{S \in \text{SPN}} \max_{\alpha, \beta \in \{0, 1\}^n \setminus \mathbf{0}} LP^S(\alpha \rightarrow \beta) . \quad (4)$$

A *one-round* characteristic for round t , $1 \leq t \leq T$, is a pair $\Omega^t = \langle \mathbf{a}^t, \mathbf{b}^t \rangle$ in which \mathbf{a}^t and \mathbf{b}^t are input and output masks, respectively, for round t , *excluding the permutation stage*. The *linear characteristic probability* of Ω^t , denoted $LCP^t(\Omega^t)$ or $LCP^t(\mathbf{a}^t \rightarrow \mathbf{b}^t)$, is simply the linear probability obtained by viewing round t (minus the permutation stage) as an $N \times N$ s-box:

$$LCP^t(\Omega^t) \stackrel{\text{def}}{=} (2 \cdot \text{Prob}\{\mathbf{a}^t \bullet \mathbf{X} = \mathbf{b}^t \bullet S^t(\mathbf{X} \oplus \mathbf{k}^t)\} - 1)^2 , \quad (5)$$

where $S^t(\cdot)$ denotes application of the s-boxes of round t , and \mathbf{X} is an r.v. uniformly distributed over $\{0, 1\}^N$. (Note: It can be shown that $LCP^t(\Omega^t)$ is independent of the (unknown) subkey \mathbf{k}^t , and therefore the operation $\oplus \mathbf{k}^t$ can be removed from (5).) Let the M s-boxes of round t be enumerated from left to right as $S_1^t, S_2^t, \dots, S_M^t$. Note that \mathbf{a}^t and \mathbf{b}^t determine input and output masks for each s-box in round t ; let the masks for S_i^t be denoted α_i^t and β_i^t , respectively. Then by Matsui's Piling-up Lemma [11],

$$LCP^t(\Omega^t) = \prod_{i=1}^M LP^{S_i^t}(\alpha_i^t \rightarrow \beta_i^t) . \quad (6)$$

Definition 1. Let \mathbf{L} denote the N -bit LT of the SPN represented as a binary $N \times N$ matrix, i.e., if $\mathbf{x}, \mathbf{y} \in \{0, 1\}^N$ are the input and output, respectively, for the LT, then $\mathbf{y} = (\mathbf{L}\mathbf{x})'$.

Lemma 1 ([3]). If $\mathbf{b} \in \{0, 1\}^N$ and $\mathbf{a} = (\mathbf{L}'\mathbf{b})'$, then $\mathbf{a} \bullet \mathbf{x} = \mathbf{b} \bullet \mathbf{y}$ for all N -bit inputs to the LT, \mathbf{x} , and corresponding outputs, \mathbf{y} (i.e., if \mathbf{b} is an output mask for the LT, then $\mathbf{a} = (\mathbf{L}'\mathbf{b})'$ is the (unique) corresponding input mask).

Now given one-round characteristics for each of rounds $1 \dots T$, $\Omega^1 = \langle \mathbf{a}^1, \mathbf{b}^1 \rangle$, $\Omega^2 = \langle \mathbf{a}^2, \mathbf{b}^2 \rangle$, \dots , $\Omega^T = \langle \mathbf{a}^T, \mathbf{b}^T \rangle$, these can be concatenated to form a single T -round characteristic if \mathbf{a}^{t+1} and \mathbf{b}^t are corresponding output and input masks for the LT, respectively, for $1 \leq t \leq (T-1)$ (see Lemma 1). The resulting T -round characteristic is the tuple $\Omega = \langle \mathbf{a}^1, \mathbf{a}^2, \dots, \mathbf{a}^T, \mathbf{b}^T \rangle$. The linear characteristic probability of Ω is again given by Matsui's Piling-up Lemma:

$$LCP(\Omega) = \prod_{t=1}^T LCP^t(\Omega^t) . \quad (7)$$

4.2 Choosing the Best Characteristic

In carrying out LC, the attacker typically runs an algorithm to find the T -round characteristic, Ω , for which $LCP(\Omega)$ is maximal; such a characteristic (not necessarily unique) is called the *best characteristic* [12]. If $\Omega = \langle \mathbf{a}^1, \mathbf{a}^2, \dots, \mathbf{a}^T, \mathbf{b}^T \rangle$, and if the input and output masks used in Algorithm 2 are taken to be $\mathbf{a} = \mathbf{a}^1$ and $\mathbf{b} = \mathbf{b}^T$, respectively, then $E_T[\mathbf{a}, \mathbf{b}]$ (used to determine $\mathcal{N}_L = \frac{c}{E_T[\mathbf{a}, \mathbf{b}]}$) is approximated by

$$E_T[\mathbf{a}, \mathbf{b}] \approx LCP(\Omega) . \quad (8)$$

5 Provable Security Against Linear Cryptanalysis

The approximation in (8) has been widely used to evaluate the security of block ciphers against LC [8]. Knudsen calls a block cipher *practically secure* if the data complexity determined by this method is prohibitive [10]. However, in 1994 Nyberg demonstrated that this approach underestimates the success of LC [14]. We state Nyberg's results in the context of SPNs.

5.1 Approximate Linear Hulls

Definition 2 (Nyberg). Given nonzero N -bit masks \mathbf{a}, \mathbf{b} , the approximate linear hull, $\text{ALH}(\mathbf{a}, \mathbf{b})$, is the set of all T -round characteristics, for the T rounds under consideration, having \mathbf{a} as the input mask for round 1 and \mathbf{b} as the output mask for round T , i.e., all characteristics of the form $\Omega = \langle \mathbf{a}, \mathbf{a}^2, \mathbf{a}^3, \dots, \mathbf{a}^T, \mathbf{b} \rangle$.

Remark: Recall that any characteristic $\Omega \in \text{ALH}(\mathbf{a}, \mathbf{b})$ determines an input and an output mask for each s-box in rounds $1 \dots T$. If this yields at least one s-box for which the input mask is zero and the output mask is nonzero, or vice versa, the linear probability associated with that s-box will be 0 (see (3)) and therefore $LCP(\Omega) = 0$ by (6) and (7). We exclude such characteristics from consideration.

Definition 3. For $\mathbf{a}, \mathbf{b} \in \{0, 1\}^N \setminus \mathbf{0}$, let $\text{ALH}(\mathbf{a}, \mathbf{b})^*$ consist of the elements $\Omega \in \text{ALH}(\mathbf{a}, \mathbf{b})$ such that for each s-box in rounds $1 \dots T$, the input and output masks determined by Ω for that s-box are either both zero or both nonzero.

Theorem 1 (Nyberg). Let \mathbf{a} and \mathbf{b} be fixed nonzero N -bit input and output masks, respectively, for T rounds of an SPN. Then

$$E_T[\mathbf{a}, \mathbf{b}] = \sum_{\Omega \in \text{ALH}(\mathbf{a}, \mathbf{b})^*} LCP(\Omega). \quad (9)$$

It follows immediately from Theorem 1 that (8) does not hold in general, since $E_T[\mathbf{a}, \mathbf{b}]$ is shown to be equal to the sum of terms $LCP(\Omega)$ over a (large) set of characteristics. Therefore, on average, the linear characteristic probability of the best characteristic will be strictly *less than* $E_T[\mathbf{a}, \mathbf{b}]$. An important implication of this is that the attacker will overestimate the number of (plaintext, ciphertext) pairs required for a given success rate. Indeed, Harpes et al. [7] comment that Matsui observed that his attacks performed better than expected.

5.2 Maximum Average Linear Hull Probability

An SPN is considered to be provably secure against LC if the *maximum average linear hull probability* (MALHP), $\max_{\mathbf{a}, \mathbf{b} \in \{0, 1\}^N \setminus \mathbf{0}} E_T[\mathbf{a}, \mathbf{b}]$, is sufficiently small that the resulting data complexity is prohibitive for any conceivable attacker. Note that this must hold for $T = R - 1$, because Algorithm 2 as presented attacks the first round. Since variations of LC can be used to attack the first and last rounds of an SPN simultaneously, it may also be important that the data complexity remain prohibitive for $T = R - 2$.

5.3 Best Previous Result

Since evaluation of the MALHP appears to be infeasible in general, researchers have adopted the approach of upper bounding this value. If such an upper bound is sufficiently small, provable security can be claimed. Hong et al. [9] give the best previously known result for the SPN architecture, stated in Theorem 2 below. First we need the following concepts.

Definition 4 ([1]). Any T -round characteristic, Ω , determines an input and an output mask for each s-box in rounds $1 \dots T$. Those s-boxes having nonzero input and output masks are called *active*.

Definition 5. Let $\Omega \in \text{ALH}(\mathbf{a}, \mathbf{b})^*$, and let \mathbf{v} be one of the masks in Ω . Then \mathbf{v} is either an input or an output mask for the substitution stage of some round of the SPN. By the definition of $\text{ALH}(\mathbf{a}, \mathbf{b})^*$ (Definition 3), the active s-boxes in this round can be determined from \mathbf{v} (without knowing the corresponding output/input mask). We define $\gamma_{\mathbf{v}}$ to be the M -bit vector which encodes this pattern of active s-boxes: $\gamma_{\mathbf{v}} = \gamma_1\gamma_2 \dots \gamma_M$, where $\gamma_i = 1$ if the i^{th} s-box is active, and $\gamma_i = 0$ otherwise, for $1 \leq i \leq M$.

Definition 6 ([4]). The branch number of the LT, denoted \mathcal{B} , is the minimum number of active s-boxes in any two consecutive rounds. It can be given by

$$\mathcal{B} = \min \{ wt(\gamma_{\mathbf{v}}) + wt(\gamma_{\mathbf{w}}) : \mathbf{w} \in \{0, 1\}^N \setminus \mathbf{0} \text{ and } \mathbf{v} = (\mathbf{L}'\mathbf{w}')' \} .$$

It is not hard to see that $2 \leq \mathcal{B} \leq (M + 1)$.

Theorem 2 (Hong et al.). If $\mathcal{B} = (M + 1)$, then $\max_{\mathbf{a}, \mathbf{b} \in \{0, 1\}^N \setminus \mathbf{0}} E_T[\mathbf{a}, \mathbf{b}] \leq q^M$, and if $\mathcal{B} = M$, then $\max_{\mathbf{a}, \mathbf{b} \in \{0, 1\}^N \setminus \mathbf{0}} E_T[\mathbf{a}, \mathbf{b}] \leq q^{M-1}$, where q is defined as in (4).

6 New Upper Bound for Maximum Average Linear Hull Probability

In this section we present a new method for upper bounding the maximum average linear hull probability. Our main results are Theorem 3 and Theorem 4. The upper bound we compute depends on:

- (a) q , the maximum linear probability over all SPN s-boxes (see (4))
- (b) T , the number of rounds being approximated by Algorithm 2
- (c) the structure of the SPN LT (via the $W[\]$ table in Definition 7 below)

6.1 Definition and Technical Lemmas

Definition 7. Let $\gamma, \hat{\gamma} \in \{0, 1\}^M$. Then

$$W[\gamma, \hat{\gamma}] \stackrel{\text{def}}{=} \# \{ \mathbf{y} \in \{0, 1\}^N : \gamma_{\mathbf{x}} = \gamma, \gamma_{\mathbf{y}} = \hat{\gamma}, \text{ where } \mathbf{x} = (\mathbf{L}'\mathbf{y}')' \} .$$

Remark: Informally, the value $W[\gamma, \hat{\gamma}]$ represents the number of ways the LT can “connect” a pattern of active s-boxes in one round (γ) to a pattern of active s-boxes in the next round ($\hat{\gamma}$).

Lemma 2. Let Ω be a one-round or T -round characteristic that makes A s-boxes active. Then $LCP(\Omega) \leq q^A$.

Proof. Follows directly from (4), (6), and (7).

Lemma 3. Let $1 \leq t \leq T$, and $\mathbf{a}, \mathbf{b}^t \in \{0, 1\}^N$. Then

$$\sum_{\mathbf{x} \in \{0, 1\}^N} E_T[\mathbf{a}, \mathbf{x}] = \sum_{\mathbf{x} \in \{0, 1\}^N} LCP^t(\mathbf{x} \rightarrow \mathbf{b}^t) = 1 .$$

Proof. The second sum equals 1 by application of Parseval's Theorem[13] to round t . To see that the first sum is equal to 1, apply Parseval's Theorem to the decryption function for rounds $1 \dots T$ (masked by \mathbf{a}), and take the expected value over the set of independent keys with uniform distribution.

Lemma 4. *Let $T \geq 2$, and let $\mathbf{a}, \mathbf{b} \in \{0, 1\}^N \setminus \mathbf{0}$. For any $\mathbf{x} \in \{0, 1\}^N$ viewed as an input mask for the LT, let \mathbf{y} denote the unique corresponding output mask (via the relationship given in Lemma 1). Then*

$$E_T[\mathbf{a}, \mathbf{b}] = \sum_{\mathbf{x} \in \{0, 1\}^N \setminus \mathbf{0}} E_{T-1}[\mathbf{a}, \mathbf{x}] \cdot LCP^T(\mathbf{y} \rightarrow \mathbf{b}) .$$

Proof. Follows immediately from (7) and (9).

Lemma 5. *Let $m \geq 2$, and suppose $\{c_i\}_{i=1}^m, \{d_i\}_{i=1}^m$ are sequences of nonnegative values. Let $\{\dot{c}_i\}_{i=1}^m, \{\dot{d}_i\}_{i=1}^m$ be the sequences obtained by sorting $\{c_i\}$ and $\{d_i\}$, respectively, in nonincreasing order. Then $\sum_{i=1}^m c_i d_i \leq \sum_{i=1}^m \dot{c}_i \dot{d}_i$.*

Proof. See Appendix A.

Lemma 6. *Suppose $\{\dot{c}_i\}_{i=1}^m, \{\ddot{c}_i\}_{i=1}^m$, and $\{\dot{d}_i\}_{i=1}^m$ are sequences of nonnegative values, with $\{\dot{d}_i\}$ sorted in nonincreasing order. Suppose there exists $\tilde{m}, 1 \leq \tilde{m} \leq m$, such that*

- (a) $\ddot{c}_i \geq \dot{c}_i$, for $1 \leq i \leq \tilde{m}$
- (b) $\ddot{c}_i \leq \dot{c}_i$, for $(\tilde{m} + 1) \leq i \leq m$
- (c) $\sum_{i=1}^m \dot{c}_i \leq \sum_{i=1}^m \ddot{c}_i$

Then $\sum_{i=1}^m \dot{c}_i \dot{d}_i \leq \sum_{i=1}^m \ddot{c}_i \dot{d}_i$.

Proof. See Appendix A.

6.2 Derivation of New Upper Bound

Our approach is to compute an upper bound for each nonzero pattern of active s-boxes in round 1 and round T ($T \geq 2$); that is, we compute $UB_T[\gamma, \hat{\gamma}]$, for $\gamma, \hat{\gamma} \in \{0, 1\}^M \setminus \mathbf{0}$, such that the following holds:

UB Property for T . For all $\mathbf{a}, \mathbf{b} \in \{0, 1\}^N \setminus \mathbf{0}$, $E_T[\mathbf{a}, \mathbf{b}] \leq UB_T[\gamma_{\mathbf{a}}, \gamma_{\mathbf{b}}]$.

If the *UB Property for T* holds, then an upper bound for the MALHP is given by $\max_{\gamma, \hat{\gamma} \in \{0, 1\}^M \setminus \mathbf{0}} UB_T[\gamma, \hat{\gamma}]$. We first handle the case $T = 2$ in Theorem 3, and then use a recursive technique for $T \geq 3$ in Theorem 4.

Theorem 3. *Let $\gamma, \hat{\gamma} \in \{0, 1\}^M \setminus \mathbf{0}$, $f = wt(\gamma)$, $\ell = wt(\hat{\gamma})$, and $W = W[\gamma, \hat{\gamma}]$. If*

$$UB_2[\gamma, \hat{\gamma}] \stackrel{\text{def}}{=} \begin{cases} \min \{q^f, q^\ell\} & \text{if } \max \{q^f, q^\ell\} \cdot W > 1 \\ q^{f+\ell} \cdot W & \text{if } \max \{q^f, q^\ell\} \cdot W \leq 1 \end{cases} \quad (10)$$

then the UB Property for 2 holds.

Proof. Let $\gamma, \hat{\gamma} \in \{0, 1\}^M \setminus \mathbf{0}$ be fixed, and let $\mathbf{a}, \mathbf{b} \in \{0, 1\}^N \setminus \mathbf{0}$ such that $\gamma_{\mathbf{a}} = \gamma$ and $\gamma_{\mathbf{b}} = \hat{\gamma}$. We want to show that $E_2[\mathbf{a}, \mathbf{b}] \leq UB_2[\gamma, \hat{\gamma}]$. There are $W = W[\gamma, \hat{\gamma}]$ ways that the LT can “connect” the f active s-boxes in round 1 to the ℓ active s-boxes in round 2. Let $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_W$ be the corresponding input masks for the LT, and let $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_W$ be the respective output masks (so $\gamma_{\mathbf{x}_i} = \gamma$ and $\gamma_{\mathbf{y}_i} = \hat{\gamma}$). Let $c_i = LCP^1(\mathbf{a} \rightarrow \mathbf{x}_i)$ and $d_i = LCP^2(\mathbf{y}_i \rightarrow \mathbf{b})$, for $1 \leq i \leq W$.

From Lemma 4 we have $E_2[\mathbf{a}, \mathbf{b}] = \sum_{i=1}^W c_i d_i$. We know that $0 \leq c_i \leq q^f$, $0 \leq d_i \leq q^\ell$ (by Lemma 2) and $\sum_{i=1}^W c_i \leq 1$, $\sum_{i=1}^W d_i \leq 1$ (by Lemma 3). Without loss of generality, assume that $f \geq \ell$, so $\min\{q^f, q^\ell\} = q^\ell$ and $\max\{q^f, q^\ell\} = q^f$ (since $0 \leq q \leq 1$). Note that q^f always upper bounds $E_2[\mathbf{a}, \mathbf{b}]$, since $E_2[\mathbf{a}, \mathbf{b}] = \sum_{i=1}^W c_i d_i \leq q^f \sum_{i=1}^W d_i \leq q^f$; we use this upper bound in the first case of (10). On the other hand, $q^{f+\ell} \cdot W$ also upper bounds $E_2[\mathbf{a}, \mathbf{b}]$, since $\sum_{i=1}^W c_i d_i \leq \sum_{i=1}^W q^f q^\ell = q^{f+\ell} \cdot W$. If $q^\ell \cdot W = 1$, the two upper bounds are identical. If $q^\ell \cdot W < 1$, then $q^{f+\ell} \cdot W < q^f$, so we use $q^{f+\ell} \cdot W$ as the upper bound in the second case of (10).

Theorem 4. *Let $T \geq 3$. Assume that values $UB_{T-1}[\gamma, \hat{\gamma}]$ have been computed for all $\gamma, \hat{\gamma} \in \{0, 1\}^M \setminus \mathbf{0}$ such that the UB Property for $(T-1)$ holds. Let values $UB_T[\gamma, \hat{\gamma}]$ be computed using the algorithm in Figure 3. Then the UB Property for T holds.*

Proof. Throughout this proof, “Line X ” refers to the X^{th} line in Figure 3. Let $\mathbf{a}, \mathbf{b} \in \{0, 1\}^N \setminus \mathbf{0}$. It suffices to show that if $\gamma = \gamma_{\mathbf{a}}$ in Line 1 and $\hat{\gamma} = \gamma_{\mathbf{b}}$ in Line 2, then the value $UB_T[\gamma, \hat{\gamma}]$ computed in Figure 3 satisfies $E_T[\mathbf{a}, \mathbf{b}] \leq UB_T[\gamma, \hat{\gamma}]$. Enumerate the nonzero output masks for the LT as $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_{2^N-1}$, and let the corresponding input masks be given by $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{2^N-1}$, respectively. From Lemma 4 we have $E_T[\mathbf{a}, \mathbf{b}] = \sum_{i=1}^{2^N-1} E_{T-1}[\mathbf{a}, \mathbf{x}_i] \cdot LCP^T(\mathbf{y}_i \rightarrow \mathbf{b})$. If $\gamma_{\mathbf{y}_i} \neq \gamma_{\mathbf{b}}$ ($= \hat{\gamma}$), then $LCP^T(\mathbf{y}_i \rightarrow \mathbf{b}) = 0$ (by the Piling-up Lemma), so these \mathbf{y}_i can be removed from consideration, leaving $\bar{\mathbf{y}}_1, \bar{\mathbf{y}}_2, \dots, \bar{\mathbf{y}}_L$, and corresponding input masks, $\bar{\mathbf{x}}_1, \bar{\mathbf{x}}_2, \dots, \bar{\mathbf{x}}_L$, respectively.

Let $c_i = E_{T-1}[\mathbf{a}, \bar{\mathbf{x}}_i]$ and $d_i = LCP^T(\bar{\mathbf{y}}_i \rightarrow \mathbf{b})$, for $1 \leq i \leq L$. Then $E_T[\mathbf{a}, \mathbf{b}] = \sum_{i=1}^L c_i d_i$. Let $\ell = wt(\hat{\gamma})$ (Line 3), and let $u_i = UB_{T-1}[\gamma, \gamma_{\bar{\mathbf{x}}_i}]$, for $1 \leq i \leq L$. Then $0 \leq c_i \leq u_i$, $0 \leq d_i \leq q^\ell$ (the latter by Lemma 2), and $\sum c_i \leq 1$, $\sum d_i \leq 1$ (by Lemma 3). It follows immediately that $E_T[\mathbf{a}, \mathbf{b}] \leq q^\ell \cdot \sum_{i=1}^L d_i \leq q^\ell$. We use this upper bound in **Case I** (Lines 19, 20).

Now note that some of the terms in $\{u_i\}$ are identical, since if $1 \leq i < j \leq L$ and $\gamma_{\bar{\mathbf{x}}_i} = \gamma_{\bar{\mathbf{x}}_j}$, then $u_i = u_j$. We use this to define an equivalence relation on $\{\bar{\mathbf{x}}_i\}$: $\bar{\mathbf{x}}_i \equiv \bar{\mathbf{x}}_j$ iff $\gamma_{\bar{\mathbf{x}}_i} = \gamma_{\bar{\mathbf{x}}_j}$. It can be seen that the number of elements in the equivalence class of $\bar{\mathbf{x}}_i$ is $W[\gamma_{\bar{\mathbf{x}}_i}, \hat{\gamma}]$.

Select indices j_1, j_2, \dots, j_H such that $\{\bar{\mathbf{x}}_{j_h}\}_{h=1}^H$ consists of one representative from each equivalence class. Let $\gamma_h = \gamma_{\bar{\mathbf{x}}_{j_h}}$, $U_h = u_{j_h} = UB_{T-1}[\gamma, \gamma_h]$, and $W_h = W[\gamma_h, \hat{\gamma}]$, for $1 \leq h \leq H$. Without loss of generality, assume that the indices are ordered such that $U_1 \geq U_2 \geq \dots \geq U_H$. It is an important observation

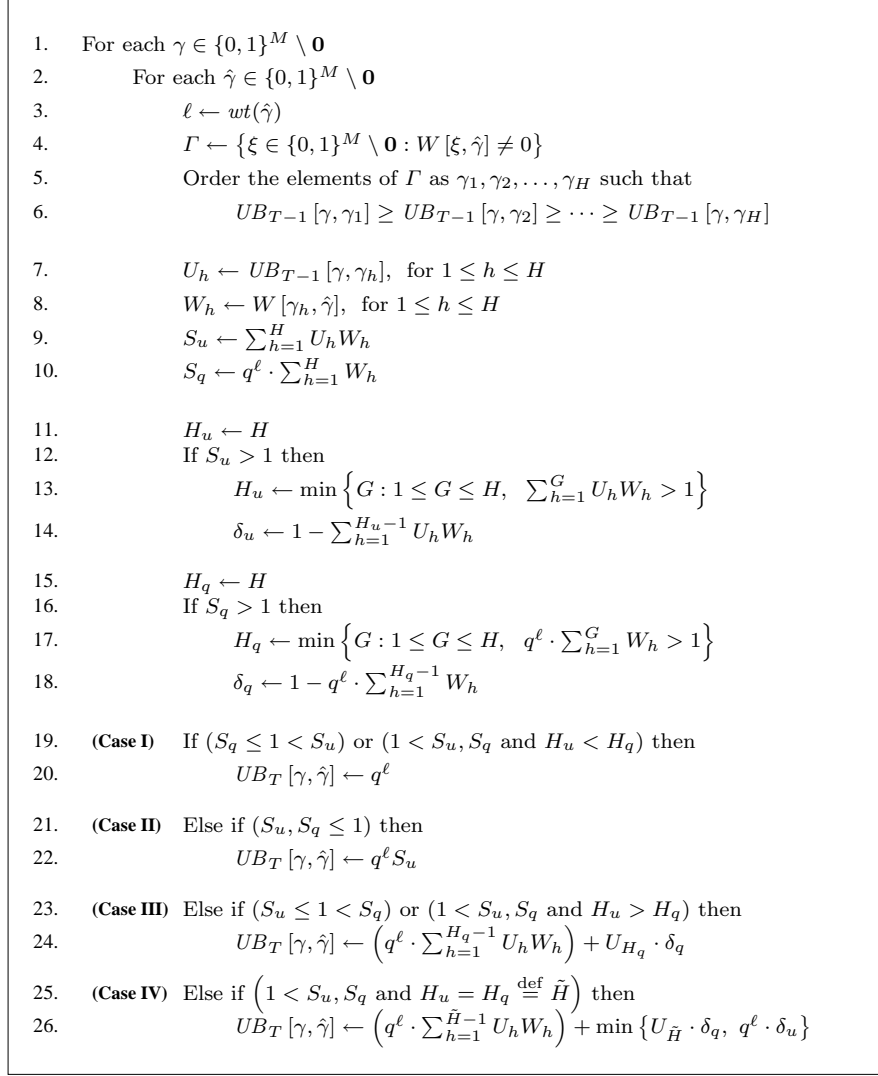


Fig. 3. Algorithm to compute $UB_T[\]$ for $T \geq 3$

that the values γ_h , U_h , and W_h are the same as those defined in Lines 5, 7, and 8. The following four facts are straightforward.

Fact 1 $\sum_{h=1}^H W_h = L$.

Fact 2 $\sum_{i=1}^L u_i = \sum_{h=1}^H U_h W_h = S_u$ (S_u is defined in Line 9).

Fact 3 $q^\ell L = q^\ell \cdot \sum_{h=1}^H W_h = S_q$ (S_q is defined in Line 10).

Using Fact 2, we get the upper bound $E_T[\mathbf{a}, \mathbf{b}] = \sum_{i=1}^L c_i d_i \leq q^\ell \cdot \sum_{i=1}^L u_i = q^\ell S_u$. If $S_u \leq 1$, this upper bound is no larger than that of **Case I**; if $S_u < 1$, it is strictly smaller. This is the upper bound we use in **Case II** (Lines 21, 22).

The proofs of **Cases I** and **II** have parallels to the proofs of the two cases in Theorem 3. For **Cases III** and **IV**, however, we require additional techniques, since the terms which upper bound the c_i (namely, the u_i) are not, in general, all the same (in the proof of Theorem 3, all the c_i are upper bounded by q^f). The intuition for what follows is this: Since $\sum_{i=1}^L d_i \leq 1$, it is not necessary to replace all the d_i by the value q^ℓ if the consequence is that $\sum_{i=1}^L q^\ell > 1$. Instead, certain of the d_i are replaced by q^ℓ and the rest by 0, so that the resulting summation is 1 (a residue term may be required). To ensure an upper bound, it is necessary that the q^ℓ terms be multiplied by the *largest* of the u_i terms. This is the reason for the sorting in Lines 5–6. (A “cutoff” of the u_i terms at the value 1 is also applied.)

Sort $\{c_i\}$, $\{d_i\}$, and $\{u_i\}$ in nonincreasing order to obtain the sequences $\{\dot{c}_i\}$, $\{\dot{d}_i\}$, and $\{\dot{u}_i\}$, respectively. Clearly $\dot{c}_i \leq \dot{u}_i$, for $1 \leq i \leq L$. Applying Lemma 5 we have $\sum_{i=1}^L c_i d_i \leq \sum_{i=1}^L \dot{c}_i \dot{d}_i$. If $S_u = \sum_{i=1}^L \dot{u}_i \leq 1$, let $\ddot{c}_i = \dot{u}_i$, for $1 \leq i \leq L$. If $S_u > 1$, let L_u ($1 \leq L_u \leq L$) be minimum such that $\sum_{i=1}^{L_u} \dot{u}_i > 1$, and let $\{\ddot{c}_i\}$ consist of the first L terms of $\dot{u}_1, \dot{u}_2, \dots, \dot{u}_{L_u-1}, \left(1 - \sum_{i=1}^{L_u-1} \dot{u}_i\right), 0, 0, 0, \dots$

If $S_q = q^\ell L \leq 1$, let $\ddot{d}_i = q^\ell$ for $1 \leq i \leq L$. Otherwise, if $S_q > 1$, let $L_q = \left\lfloor \frac{1}{q^\ell} \right\rfloor$, and let $\{\ddot{d}_i\}$ consist of the first L terms of $\underbrace{q^\ell, \dots, q^\ell}_{L_q \text{ terms}}, (1 - q^\ell L_q), 0, 0, 0, \dots$

Then $\{\dot{c}_i\}$, $\{\ddot{c}_i\}$, and $\{\dot{d}_i\}$ satisfy the conditions on the identically named sequences in the statement of Lemma 6, so $\sum_{i=1}^L \dot{c}_i \dot{d}_i \leq \sum_{i=1}^L \ddot{c}_i \dot{d}_i$. Also, $\{\dot{d}_i\}$, $\{\ddot{d}_i\}$, and $\{\ddot{c}_i\}$ satisfy the conditions on the three sequences in the statement of Lemma 6 (in that order), and therefore $\sum_{i=1}^L \ddot{c}_i \dot{d}_i \leq \sum_{i=1}^L \ddot{c}_i \ddot{d}_i$. Combining, we get $E_T[\mathbf{a}, \mathbf{b}] \leq \sum_{i=1}^L \ddot{c}_i \ddot{d}_i$, so it remains to show that $\sum_{i=1}^L \ddot{c}_i \ddot{d}_i \leq UB_T[\gamma, \hat{\gamma}]$.

Define the partial sums $P_0 = 0$ and $P_h = \sum_{j=1}^h W_j$, for $1 \leq h \leq H$ (so $P_H = L$).

Case III (Lines 23, 24) If either condition in Line 23 holds, then

- (a) $\ddot{c}_i = U_h$ and $\ddot{d}_i = q^\ell$, for $(P_{h-1} + 1) \leq i \leq P_h$, $1 \leq h \leq (H_q - 1)$
- (b) $\ddot{c}_i = U_{H_q}$, for $(P_{H_q-1} + 1) \leq i \leq P_{H_q}$
- (c) $\sum_{i=(P_{H_q-1}+1)}^{P_{H_q}} \ddot{d}_i = \delta_q$
- (d) $\ddot{d}_i = 0$, for $i \geq (P_{H_q} + 1)$

It follows that $\sum_{i=1}^L \ddot{c}_i \ddot{d}_i = \left(q^\ell \cdot \sum_{h=1}^{H_q-1} U_h W_h\right) + U_{H_q} \cdot \delta_q$, which is the upper bound used in **Case III**.

Case IV (Lines 25, 26) If the condition in Line 25 holds, then (using the definition of \tilde{H} in Line 25)

- (a) $\ddot{c}_i = U_h$ and $\ddot{d}_i = q^\ell$, for $(P_{h-1} + 1) \leq i \leq P_h$, $1 \leq h \leq (\tilde{H} - 1)$
- (b) $\sum_{i=(P_{\tilde{H}-1}+1)}^{P_{\tilde{H}}} \ddot{c}_i = \delta_u$ and $\sum_{i=(P_{\tilde{H}-1}+1)}^{P_{\tilde{H}}} \ddot{d}_i = \delta_q$
- (c) $\ddot{c}_i = \ddot{d}_i = 0$, for $i \geq (P_{\tilde{H}} + 1)$

Let $Y = q^\ell \cdot \sum_{h=1}^{\tilde{H}-1} U_h W_h$. For $(P_{\tilde{H}-1} + 1) \leq i \leq P_{\tilde{H}}$, replacing \ddot{c}_i by its upper bound $U_{\tilde{H}}$ gives $\sum_{i=1}^L \ddot{c}_i \ddot{d}_i \leq Y + U_{\tilde{H}} \cdot \delta_q$. For i in the same range, replacing \ddot{d}_i by its upper bound q^ℓ gives $\sum_{i=1}^L \ddot{c}_i \ddot{d}_i \leq Y + q^\ell \cdot \delta_u$. Combining the above, we get $\sum_{i=1}^L \ddot{c}_i \ddot{d}_i \leq Y + \min\{U_{\tilde{H}} \cdot \delta_q, q^\ell \cdot \delta_u\}$, the right-hand side of which is the upper bound used in **Case IV**.

7 Application of New Upper Bound to Rijndael

To test our new upper bound, we generated random invertible LTs for SPNs with various parameters. We found that for LTs with branch number $\mathcal{B} = M$, our upper bound was consistently superior to that of Hong et al. [9]. We give the results for one such LT in Appendix B. For LTs with $\mathcal{B} = (M + 1)$, our upper bound did not appear to improve on that of [9].

However, the bulk of our analysis we reserved for Rijndael with the following parameters: $N = 128$, $R = 10$, $M = 16$, $n = 8$, $q = 2^{-6}$. Note that the result of [9] does not apply to Rijndael, since for Rijndael, $\mathcal{B} = 5 < M = 16$) [5]. Tailoring our algorithm to any particular SPN involves computation of the values in $W[\]$ (Definition 7), which for Rijndael is a $2^{16} \times 2^{16}$ table. The Rijndael LT is depicted in Figure 4.




Fig. 4. Rijndael linear transformation

The 128-bit input block can be viewed as an array of 16 bytes. These bytes are first shuffled according to the figure, and then consecutive 4-byte sequences are fed into copies of the same highly diffusive 32-bit LT (based on maximum-distance-separable (MDS) codes). We first computed the $2^4 \times 2^4$ $W[\]$ table for the MDS LT, denoted $W_{\text{MDS}}[\]$, by transforming all 2^{32} output masks (see Definition 7). Given $\gamma \in \{0, 1\}^{16}$ representing a pattern of active s-boxes for the Rijndael LT *input*, a corresponding 4-bit input pattern is determined for each copy of the MDS LT simply by tracing through the byte “shuffle”: denote these $\gamma_1, \gamma_2, \gamma_3, \gamma_4 \in \{0, 1\}^4$, from left to right, respectively. Then given $\hat{\gamma} \in \{0, 1\}^{16}$ representing a pattern of active s-boxes for the Rijndael LT *output*, partition $\hat{\gamma}$ into consecutive 4-bit sequences representing output patterns for the MDS LT, denoted $\hat{\gamma}_1, \hat{\gamma}_2, \hat{\gamma}_3, \hat{\gamma}_4 \in \{0, 1\}^4$. Then $W[\gamma, \hat{\gamma}] = \prod_{i=1}^4 W_{\text{MDS}}[\gamma_i, \hat{\gamma}_i]$.

Since $W[\]$ turns out to be quite sparse (roughly 80,000,000 of the 2^{32} entries are nonzero, around 2%), we precompute it, and store the nonzero entries. By

doing this first for the $W_{\text{MDS}}[\]$ table, computation of $W[\]$ becomes fairly fast. Computing the upper bound in the case $T = 2$ using Theorem 4 is easy. The main work involves executing the algorithm in Figure 3 for $T = 3 \dots 10$. Lines 3–26 are executed $(2^{16} - 1)^2$ times for each value of T ($3 \leq T \leq 10$), a total of $\approx 2^{35}$ iterations. Once the values $\gamma_1, \gamma_2, \dots, \gamma_H$ in Line 5 are known, the time complexity of Lines 7–26 is $O(H)$. Since the values in Γ in Line 4 can be precomputed and stored during generation of $W[\]$, the sorting specified in Lines 5–6 is the most expensive ($O(H \log H)$). The average value for H is 1191, although individual values vary widely.

For a fixed value of γ , computing $UB_T[\gamma, \hat{\gamma}]$ for all $\hat{\gamma} \in \{0, 1\}^M \setminus \mathbf{0}$ and all T ($2 \leq T \leq 10$) takes approximately 40 minutes on a Sun Ultra 5, for a total running time in the range of 44,000 hours on that platform. We completed the computation by distributing it over roughly 60 CPUs for several weeks.

Our results for Rijndael are given in Figure 5. For $7 \leq T \leq 10$, the upper bound value is 2^{-75} , giving a corresponding lower bound on the data complexity of LC of 2^{78} , for a 96.7% success rate (see Section 3). Note that for Algorithm 2 as described in Section 3, $T = R - 1 = 9$.




Fig. 5. New upper bound applied to Rijndael

8 Conclusion

We have presented a new method for computing an upper bound on the maximum average linear hull probability for SPNs. Our method has the advantage that it can be computed for an SPN with any LT layer, whereas the best previous result (Hong et al. [9]) applies only to SPNs with highly diffusive LTs, i.e., those having branch number $\mathcal{B} = M$ or $\mathcal{B} = (M + 1)$, where M is the number

of s-boxes per round. In addition, our upper bound is a function of the number of rounds being approximated; other known upper bounds do not vary with the number of rounds. When applied to an SPN whose LT has branch number $\mathcal{B} = (M + 1)$ (the maximal case), our upper bound does not appear to improve on that of [9]. For SPNs whose LTs have branch number $\mathcal{B} = M$, our upper bound is consistently superior to that of [9].

A significant part of our work involved application of our method to Rijndael (with $N = 128$ and $R = 10$). This yielded the upper bound $UB = 2^{-75}$, for a corresponding lower bound on the data complexity of LC of $\frac{8}{UB} = 2^{78}$ (for a 96.7% success rate). Note that this does not demonstrate the existence of a such an attack, but is, to our knowledge, the first such lower bound.

Acknowledgments The authors are grateful to the reviewers for comments which significantly improved the presentation of this paper. For help in obtaining the required computing resources, we are grateful to: the High Performance Computing Virtual Laboratory (Canada), Tom Bradshaw, Randy Ellis, Howard Heys, Alex MacPherson, Andrew Pollard, Carolyn Small, and Amr Youssef.

References

1. E. Biham, *On Matsui's linear cryptanalysis*, Advances in Cryptology—EUROCRYPT'94, Springer-Verlag, pp. 341–355, 1995.
2. E. Biham and A. Shamir, *Differential cryptanalysis of DES-like cryptosystems*, Journal of Cryptology, Vol. 4, No. 1, pp. 3–72, 1991.
3. J. Daemen, R. Govaerts, and J. Vandewalle, *Correlation matrices*, Fast Software Encryption : Second International Workshop, Springer-Verlag, pp. 275–285, 1995.
4. J. Daemen, L. Knudsen, and V. Rijmen, *The block cipher SQUARE*, Fast Software Encryption—FSE'97, Springer-Verlag, pp. 149–165, 1997.
5. J. Daemen and V. Rijmen, *AES proposal: Rijndael*, <http://csrc.nist.gov/encryption/aes/round2/AESAlgs/Rijndael/Rijndael.pdf>, 1999.
6. H. Feistel, *Cryptography and computer privacy*, Scientific American, Vol. 228, No. 5, pp. 15–23, May 1973.
7. C. Harpes, G. Kramer, and J. Massey, *A generalization of linear cryptanalysis and the applicability of Matsui's piling-up lemma*, Advances in Cryptology—EUROCRYPT'95, Springer-Verlag, pp. 24–38, 1995.
8. H.M. Heys and S.E. Tavares, *Substitution-permutation networks resistant to differential and linear cryptanalysis*, Journal of Cryptology, Vol. 9, No. 1, pp. 1–19, 1996.
9. S. Hong, S. Lee, J. Lim, J. Sung, and D. Cheon, *Provable security against differential and linear cryptanalysis for the SPN structure*, Fast Software Encryption (FSE 2000), Proceedings to be published by Springer-Verlag.
10. L.R. Knudsen, *Practically secure Feistel ciphers*, Fast Software Encryption, Springer-Verlag, pp. 211–221, 1994.
11. M. Matsui, *Linear cryptanalysis method for DES cipher*, Advances in Cryptology—EUROCRYPT'93, Springer-Verlag, pp. 386–397, 1994.
12. M. Matsui, *On correlation between the order of s-boxes and the strength of DES*, Advances in Cryptology—EUROCRYPT'94, Springer-Verlag, pp. 366–375, 1995.

13. W. Meier and O. Staffelbach, *Nonlinearity criteria for cryptographic functions*, Advances in Cryptology—EUROCRYPT'89, Springer-Verlag, pp. 549–562, 1990.
14. K. Nyberg, *Linear approximation of block ciphers*, Advances in Cryptology—EUROCRYPT'94, Springer-Verlag, pp. 439–444, 1995.
15. C.E. Shannon, *Communication theory of secrecy systems*, Bell System Technical Journal, Vol. 28, no. 4, pp. 656–715, 1949.

Appendix A

Proof (Lemma 5). Without loss of generality, assume that $\{d_i\}$ is already sorted in nonincreasing order, so $\dot{d}_i = d_i$. If $m = 2$ and $\{c_i\}$ is not sorted, i.e., if $c_1 < c_2$, then $\dot{c}_1 = c_2$ and $\dot{c}_2 = c_1$, so

$$\begin{aligned} \sum_{i=1}^2 c_i d_i \leq \sum_{i=1}^2 \dot{c}_i \dot{d}_i &\iff c_1 d_1 + c_2 d_2 \leq c_2 \dot{d}_1 + c_1 \dot{d}_2 \\ &\iff (c_2 - c_1) d_2 \leq (c_2 - c_1) d_1 \\ &\iff d_2 \leq d_1, \end{aligned}$$

which is true since $\{d_i\}$ was assumed to be sorted. Let $m \geq 3$ and assume the lemma holds for $m - 1$. Let s be the index of a minimal element in $\{c_i\}$, and let $\{\hat{c}_i\}_{i=1}^m$ be the sequence obtained by exchanging c_s and c_m in $\{c_i\}$. Then $\dot{c}_m = \hat{c}_m$, and therefore sorting $\{\hat{c}_i\}_{i=1}^{m-1}$ in nonincreasing order gives $\{\dot{c}_i\}_{i=1}^{m-1}$. By an argument similar to that of the $m = 2$ case, we have $\sum_{i=1}^m c_i \dot{d}_i \leq \sum_{i=1}^m \hat{c}_i \dot{d}_i$. Applying the induction hypothesis to the first $m - 1$ terms of $\{\hat{c}_i\}$ and $\{\dot{d}_i\}$ gives $\sum_{i=1}^{m-1} \hat{c}_i \dot{d}_i \leq \sum_{i=1}^{m-1} \dot{c}_i \dot{d}_i$. Combining these facts, we get

$$\sum_{i=1}^m c_i d_i \leq \sum_{i=1}^m \hat{c}_i \dot{d}_i = \sum_{i=1}^{m-1} \hat{c}_i \dot{d}_i + \hat{c}_m \dot{d}_m \leq \sum_{i=1}^{m-1} \dot{c}_i \dot{d}_i + \dot{c}_m \dot{d}_m = \sum_{i=1}^m \dot{c}_i \dot{d}_i.$$

Proof (Lemma 6). Let

$$\begin{aligned} \dot{A} &= \sum_{i=1}^{\tilde{m}} \dot{c}_i & \dot{B} &= \sum_{i=\tilde{m}+1}^m \dot{c}_i & \dot{C} &= \sum_{i=1}^m \dot{c}_i \\ \ddot{A} &= \sum_{i=1}^{\tilde{m}} \ddot{c}_i & \ddot{B} &= \sum_{i=\tilde{m}+1}^m \ddot{c}_i & \ddot{C} &= \sum_{i=1}^m \ddot{c}_i \end{aligned}$$

By assumption, $\dot{A} \leq \ddot{A}$, $\dot{B} \geq \ddot{B}$, and $\dot{C} \leq \ddot{C}$. Let $\Delta A = \ddot{A} - \dot{A} \geq 0$ and $\Delta B = \dot{B} - \ddot{B} \geq 0$. Note that $\Delta A - \Delta B = \ddot{C} - \dot{C} \geq 0$. We have

$$\sum_{i=1}^{\tilde{m}} \ddot{c}_i \dot{d}_i \geq \sum_{i=1}^{\tilde{m}} \dot{c}_i \dot{d}_i + \Delta A \cdot \dot{d}_{\tilde{m}} \quad (11)$$

$$\sum_{i=\tilde{m}+1}^m \ddot{c}_i \dot{d}_i \geq \sum_{i=\tilde{m}+1}^m \dot{c}_i \dot{d}_i - \Delta B \cdot \dot{d}_{\tilde{m}+1} \quad (12)$$

Adding (11) and (12), we get

$$\begin{aligned} \sum_{i=1}^m \ddot{c}_i \dot{d}_i &\geq \sum_{i=1}^m \dot{c}_i \dot{d}_i + \Delta A \cdot \dot{d}_{\bar{m}} - \Delta B \cdot \dot{d}_{\bar{m}+1} \geq \sum_{i=1}^m \dot{c}_i \dot{d}_i + \Delta A \cdot \dot{d}_{\bar{m}+1} - \Delta B \cdot \dot{d}_{\bar{m}+1} \\ &= \sum_{i=1}^m \dot{c}_i \dot{d}_i + (\Delta A - \Delta B) \cdot \dot{d}_{\bar{m}+1} \\ &\geq \sum_{i=1}^m \dot{c}_i \dot{d}_i . \end{aligned}$$

Appendix B

Some of the LTs which we randomly generated were for SPNs with parameters $N = 24$, $M = 3$, and $n = 8$. For one example of such an LT for which $\mathcal{B} = M = 3$, we plot our upper bound against that of Hong et al. [9] in Figure 6, using a \log_2 scale on the y-axis. We also plot the value q^M (the upper bound of [9] for $\mathcal{B} = (M + 1) = 4$) for comparison purposes. On the x-axis we use *minimum nonlinearity*, NL_{\min} ; for $n = 8$, the relationship between NL_{\min} and q is given by $q = \left(1 - \frac{NL_{\min}}{128}\right)^2$. For this particular LT, it happened that our upper bound settled on a fixed value for $T = 2$, and did not decrease with an increasing number of rounds—this is the value we plot for each NL_{\min} .




Fig. 6. Comparison of new upper bound with that of Hong et al. [9]