An All-In-One Approach to Differential Cryptanalysis for Small Block Ciphers*

Martin R. Albrecht[1] and Gregor Leander[2]

[1] INRIA, Paris-Rocquencourt Center, POLSYS Project
UPMC Univ Paris 06, UMR 7606, LIP6, F-75005, Paris, France
CNRS, UMR 7606, LIP6, F-75005, Paris, France
[2] DTU Mathematics, Department of Mathematics, Technical University of Denmark,
2800 Kgs. Lyngby, Denmark
malb@lip6.fr, G.Leander@mat.dtu.dk

Abstract. We present a framework that unifies several standard differential techniques. This unified view allows us to consider many, potentially all, output differences for a given input difference and to combine the information derived from them in an optimal way. We then propose a new attack that implicitly mounts several standard, truncated, impossible, improbable and possible future variants of differential attacks in parallel and hence allows to significantly improve upon known differential attacks using the same input difference. To demonstrate the viability of our techniques, we apply them to KATAN-32. In particular, our attack allows us to break 115 rounds of KATAN-32. For this, our attack exploits the non-uniformity of the difference distribution after 91 rounds which is 20 rounds more than the previously best known differential characteristic.

Keywords: symmetric cryptography, block cipher, differential attack.

1 Introduction

Designing a secure block cipher that, at the same time, is very efficient is still challenging. In particular, lightweight cryptography which recently received considerable attention from the cryptographic community calls for block ciphers that can be efficiently implemented even in very resource constrained devices. Designing secure ciphers for such tiny devices – e.g., RFID tags or sensor networks – requires, on the one hand, innovative design strategies and, on the other hand, perhaps compromises in the security level. One such constraint is the block size used in block ciphers. As the block size, along with the key size, greatly influences the required circuit size, block ciphers tailored to be implemented in small devices have a strong tendency to feature smaller block sizes compared to modern block ciphers mainly focusing on software such as the AES. While modern block ciphers focusing on software usually have a block size of no less than 128 bits, most ciphers designed for efficient implementations in hardware have block

* This is an extended abstract of the work available as [1].

L.R. Knudsen and H. Wu (Eds.): SAC 2012, LNCS 7707, pp. 1–15, 2013.

sizes of 64 bits or less (see for example PRESENT [8] or HIGHT [12]). A block cipher with a particular small block size of 32-bit is KATAN-32 [10] presented at CHES 2009.

Block ciphers with very small block sizes have some interesting characteristics. From the point of view of the attacker, when using the block cipher in counter mode, it is possible to distinguish the output from a random sequences faster. Similarly, an attacker can build a complete code book faster and time-memory tradeoffs are a greater concern. From the perspective of the designer, most statistical attacks like differential or linear cryptanalysis seem at first glance to become more difficult as the amount of data available to the attacker is much more restricted.

Finally, from a theoretical point of view, small block sizes provide the opportunity to understand well-established attacks better since computations involving the entire code-book are feasible. In particular, for differential cryptanalysis, it becomes feasible to compute the exact expected probabilities for many (sometimes all) differentials. This data then allows to study the behaviour of (classical) differential cryptanalysis and related techniques more precisely.

Yet, it is not obvious a priori how to provide an optimal unified view on these differentials even if this data is available. To provide an answer to this question, this work investigates the probability distribution of output differences under one (or many) input difference and provides an optimal way to use the non-uniform distribution of differences in an attack.

Prior Work: *Differential cryptanalysis* was first proposed by Biham and Shamir [4] and since became one of the most prominent tools in the analysis of block ciphers. Many improvements and extensions have been proposed in the past, we mention some of the most influential ones. Knudsen [15] and later Biham, Biryukov and Shamir [3] proposed to use differentials with zero probability, that is *impossible differential* attacks. Based on the work of Lai [17] *High-order differentials* were introduced in [16] and are most effective against ciphers where the algebraic degree can be limited. *Truncated differentials*, first mentioned in [16] can be seen as a collection of differentials and in some cases allow to push differential attacks one or two rounds further. *Boomerang attacks* can be viewed as special cases of second order differentials and are most efficient when the probability of any differential drops rapidly with an increasing number of rounds. Recently, *improbable differentials* have been suggested [22] as a natural extension of impossible differentials and have been successfully applied to the block cipher CLEFIA. Also recently, differential cryptanalysis was extended to *multi-differential cryptanalysis* in [6]. Finally, our application of the log-likelihood can be seen in the framework of [21].

Our Contribution: Abstractly, differential cryptanalysis exposes a non-uniform distribution of output differences given one (or several) input differences. This is also the point of view from which our investigation sets out. Phrased in these terms, recovering key information using differential techniques becomes the task of distinguishing between distributions, one for the right key and one for the

wrong keys. However, usually the attacker does not have access to a full description of these distributions. In standard differential cryptanalysis only one output difference is considered and usually the probability of the best differential characteristic is considered in place of the probability of the output differential. Furthermore, for wrong keys it is assumed that the distribution is uniform.

In comparison the advantage of an attacker when dealing with small block-size ciphers become apparent. The attacker has, under mild assumptions, the ability to compute the parameters of those distributions precisely. Thus, the task is no longer to distinguish (essentially) unknown distributions, but distributions which are known completely. In particular, the usual hypotheses that wrong keys result in random permutations can be lifted. To this end, we first introduce a model to study and distinguish these distributions. As an important side effect, our framework unifies and generalises standard differential attacks, impossible differentials, improbable differentials and truncated differentials into one attack framework. Since our framework considers the distribution of all output differences it captures all techniques which exploit statistically significant subspaces of the output space.

We then propose a new attack based on this model that implicitly mounts several standard, truncated, impossible, improbable and possible future variants of differential attacks in parallel and hence allows to significantly improve upon known differential attacks using the same input difference. We stress that these "parallel applications" of various differential attacks are such that they are strictly better than those attacks considered independently. To demonstrate the viability of our model and attack, we apply our attack to two ciphers with small block sizes: the toy-cipher SmallPresent[4] and KATAN-32. For KATAN-32 we present the best known differential attack.[1] In particular, our attack allows us to break 115 rounds of KATAN-32, which is 37 rounds more than previous work [14], although we note that our attack requires considerably more resources than [14]. For this, our attack exploits the non-uniformity of the difference distribution after 91 rounds which is 20 rounds more than the previously best known differential characteristic. Since our results takes into account several standard techniques and still cover less than 1/2 of the cipher, they further strengthen our confidence in KATAN-32's resistance against differential attacks. For completeness, we also like to mention a recent preprint [13] using a meet-in-the-middle variant to recover the key for the full KATAN (slightly) faster than exhaustive search.

Furthermore, our model allows to combine many *input-* and output-differences which allows to reduce the data complexity compared to previous works significantly. This is mainly due to the fact that our approach almost naturally provides the optimal way of combining information from several input and output differences. This is the major difference between our work and [6].

We highlight that similar approaches have been independently developed by Blondeau, Gérard and Nyberg [7] and Murphy [20]. While these approaches also differ in some theoretical respects (such as using the likelihood instead of the

[1] Our attack is also the best known differential attack on SmallPresent[4].

likelihood ratio in the latter case), the main difference between these works and ours is that we put our model to practice and use it to improve upon known attacks.

2 Preliminaries and Notation

In this work we focus on block ciphers where the key is XORed to (parts of) the state. Let R_k denote one round function of a block cipher with (round)-key k, where without loss of generality the key is added in last. By R we denote the round function without the final key addition, that is $R_k(x) = R(x) \oplus k$. Moreover let $E_K : \mathbb{F}_2^n \to \mathbb{F}_2^n$ be the corresponding r round block cipher, where $K = (k_0, k_1, \ldots, k_r)$ consist of all round keys. More precisely $E_K(x) = R_{k_r} \circ R_{k_{r-1}} \circ \cdots \circ R_{k_1}(x \oplus k_0)$ where k_0 is the whitening key. For a function $F : \mathbb{F}_2^n \to \mathbb{F}_2^n$ given an input difference δ and an output difference γ we denote

$$P_F(\delta, \gamma) := \mathbf{Pr}(F(X) \oplus F(X \oplus \delta) = \gamma)$$

for randomly uniformly chosen X. That is, $P_F(\delta, \gamma)$ is the probability of the differential $\delta \to \gamma$. Using N (unordered) pairs, the number of pairs following the given differential is denoted by $D_F^{(N)}(\delta, \gamma)$. The expected value of $D_F^{(N)}(\delta, \gamma)$ is $NP_F(\delta, \gamma)$ and we discuss below more precisely how $D_F^{(N)}(\delta, \gamma)$ is distributed.

Note that in the following N always denotes the number of (unordered) plaintext/ciphertext pairs used. As we use unordered pairs, using the full code book corresponds to choosing $N = 2^{n-1}$.

We consider the case where we assume E is a Markov cipher. A cipher E is a Markov cipher when the transitional probabilities for the output differences of round $r + 1$ only depend on the output difference of round r. More precisely the round function has to satisfy [18]:

$$\mathbf{Pr}(R(X \oplus k) \oplus R(X \oplus \delta \oplus k) = \gamma \mid X = x_0) = P_R(\delta, \gamma)$$

for all choices of x_0 and uniformly random chosen subkeys k. If, furthermore, all round keys are independent, then one can compute the average value of $P_{E_K}(\delta, \gamma)$ over all possible keys by adding the probabilities for all differential characteristics included in the differential. This has first been formalised in [18] and is summarised in the next proposition.

Proposition 1. *For a function $E_K : \mathbb{F}_2^n \to \mathbb{F}_2^n = R_{k_r} \circ R_{k_{r-1}} \circ \cdots \circ R_{k_1}(x \oplus k_0)$ with input difference δ, output difference γ and $P_R(\gamma', \delta')$ the probability of the differential $\gamma' \to \delta'$ for the function R we have*

$$\begin{aligned} \tilde{P}_E(\delta, \gamma) &:= \frac{1}{\sharp K} \sum_K P_{E_K}(\delta, \gamma) \\ &= \sum_{\gamma_1, \ldots, \gamma_{r-1}} P_R(\delta, \gamma_1) \left(\prod P_R(\gamma_i, \gamma_{i+1}) \right) P_R(\gamma_{r-1}, \gamma) \end{aligned}$$

The *hypothesis of stochastic equivalence* states (cf. [18]) that for almost all keys we expect $P_{E_K}(\delta,\gamma) \approx \tilde{P}_E(\delta,\gamma)$ which implies that $D_K^{(N)}(\delta,\gamma) \approx N\tilde{P}_E(\delta,\gamma)$ for almost all keys.

This approximation has to be understood as expected value taken over all expanded keys. However, for our purpose, we are not only interested in the expected value of the counter $D_{E_K}^{(N)}(\delta,\gamma)$ but moreover how these values are distributed. This was analysed in [11] and more recently in [5]. It turns out, considering $D_{E_K}^{(N)}(\delta,\gamma)$ as the results of N independent Bernoulli trials with success probability $\tilde{P}_E(\delta,\gamma)$ leads to a good model of the actual distribution. More precisely, denoting by $\mathcal{B}(n,p)$ the Binomial distribution with n tries and success probability p, the following is a reasonable approximation for the distribution of $D_{E_K}^{(N)}(\delta,\gamma)$.

Assumption 1 (cf. Theorem 14 in [11]). *The counter $D_{E_K}^{(N)}(\delta,\gamma)$ is distributed according to the Binomial distribution $\mathcal{B}(N,\tilde{P}_E(\delta,\gamma))$, that is*

$$\mathbf{Pr}(D_{E_K}^{(N)}(\delta,\gamma)=c) = \binom{N}{c}\tilde{P}_E(\delta,\gamma)^c(1-\tilde{P}_E(\delta,\gamma))^{N-c}$$

where the probability is taken over random keys K.

We note that we experimentally validated this assumption for all ciphers considered in this work although these ciphers are not Markov ciphers.

2.1 $\tilde{P}_E(\delta,\gamma)$ in Differential Cryptanalysis

In standard differential cryptanalysis the attacker attempts to find an input difference and an output difference such that $\tilde{P}_E(\delta,\gamma)$ is "sufficiently high", i.e., bounded away from uniform. In this case we can expect that, with high probability, for each key K there exist sufficiently many right pairs to mount an attack, i.e., to detect the bias of $\tilde{P}_E(\delta,\gamma)$. Traditionally, in a 1R attack on the cipher $R_{k_{r+1}} \circ E_K$ one (partially) decrypts the last round with all possible (partial) round keys and increases a counter for the current round key guess iff the computed difference fits the expected output difference γ of round r. Afterwards, the keys are ranked according to their counters, that is, the attacker first tries the key with the highest counter, than the one with the second highest counter, etc.

The success probability of a differential attack is usually computed under the *Wrong-Key Randomization Hypotheses*. The Wrong-Key Randomization Hypotheses (see for example [18]) states that for a wrong key guess the corresponding counter is distributed as for a random permutation. Using the notation established above the Wrong-Key Randomization Hypotheses can be stated as follows

Assumption 2 (Wrong-Key Randomization Hypotheses, cf. [18]). *For a wrong key guess the corresponding counter is distributed as for a random permutation, that is $D^{(N)}_{R_{k'}^{-1}\circ R_{k_{r+1}}\circ E_K}(\delta,\gamma) \sim \mathcal{B}(N,2^{-n})$ for all $k' \neq k_{r+1}$.*

2.2 Distinguishing Distributions

Following the above discussion on the distribution of counter values, it is natural to view a differential attack as a technique to find the value k_{r+1} which maximises the likelihood function corresponding to the right-key distribution (Maximum Likelihood Estimation). This estimation may take two distributions into account. For the right key guess, according to Assumption 1 the counter is distributed according to $\mathcal{B}(N, \tilde{P}_E(\delta, \gamma))$ while the counter of a wrong key guess is assumed (cf. Assumption 2) to be distributed accordingly to $\mathcal{B}(N, 2^{-n})$.

In this setting, the maximum likelihood estimation is equivalent to maximising the log-likelihood ratio of the two distributions under consideration. Indeed, by the Neyman-Pearson Lemma the log-likelihood ratio is the most poweful test to determine whether a sample comes from one of two distributions. Denoting $p = \tilde{P}_E(\delta, \gamma)$ and $q = 2^{-n}$, if a key K resulted in a counter value c one computes

$$l_k(c) := \log\left(\frac{\binom{N}{c}p^c(1-p)^{N-c}}{\binom{N}{c}q^c(1-q)^{N-c}}\right) = c\log\left(\left(\frac{p(1-q)}{q(1-p)}\right)\right) + N \cdot \log\left(\frac{1-p}{1-q}\right).$$

The key guesses are ranked according to their $l_k(c)$ values, that is, the key with highest $l_k(c)$ value is tested first. To simplify the computation one can equivalently rank the keys according to

$$l'_k(c) = c \cdot w \text{ where } w = \log\left(\frac{p(1-q)}{q(1-p)}\right),$$

as we are only interested in the relative value of $l_k(c)$.[2]

We may write $l_{k'}$ and $l'_{k'}$ for $l_{k'}(c)$ and $l'_{k'}(c)$ respectively if it is clear from the context which c we are referring to.

Now, observe that $l'_k(c)$ is monotone increasing iff $p > q$ (as in this case $w > 0$). Thus, if the expected counter for the right key is higher than for wrong key guesses then l'_k has the same ranking and the rankings accordingly to $l'_k(c)$ and c is the same. However, if $p < q$ the function is monotone decreasing (as $w < 0$) and the ranks get reversed. This corresponds to *improbable differentials* as defined in [22]. The special case where $p = 0$ corresponds to *impossible differentials* (as introduced in [15] and later used in [3]), as in this case for each counter $c \neq 0$ the value $l_k(c)$ is formally minus infinity. In the latter case we use the convention $w = -\infty$ and $0 \cdot w = 0$. To conclude, we state the following observation.

Observation 1. *Ranking keys according to their maximum likelihood estimation as defined in Equation (1) unifies in a natural way standard differential attacks, impossible differentials and improbable differentials.*

As explained in the next section, it is this unified view that allows for a generalised attack that considers many (in principle all) counters $D^{(N)}_{E_K}(\delta, \gamma)$ simultaneously.

[2] As discussed below, this is actually equivalent to sorting according to the counters c in the case of $p > q$ and to $-c$ in the case of $p < q$.

3 The Attack Model

In this section, we present our attack in detail and provide formulas for computing the gain of our attack. In summary, we use many (or even all) counters $D_{E_K}^{(N)}(\delta, \gamma)$ for different δ and γ values simultaneously. We view those counters as samples from one out of two possible (this time multi-dimensional) distributions. One distribution corresponds to the correct round-key guess and the other to the wrong key guesses. Using many counters at the same time allows us to significantly improve the success probability (or – equivalently – reduce the data complexity) compared to standard differential attacks. Informally, and this is the major difference and biggest improvement over a related approach performed in [6], this allows us to perform several standard differential attacks and impossible (or more generally improbable) differential attacks at the same time. In our attacks these simultaneous differential attacks are weighted appropriately ensuring that we do not lose information compared to standard attacks. That is, considering more information never reduces the success probability but strictly improves it.

3.1 Multi-dimensional Distribution of $D_{E_K}^{(N)}(\delta, \gamma)$

While in general any subset of pairs of input/output differences could be considered, here we focus on the case where one input difference is fixed and we consider all possible output differences. In this case, we denote by

$$\mathcal{D}_{E_K}^{(N)}(\delta) = \left(D_{E_K}^{(N)}(\delta, 1), D_{E_K}^{(N)}(\delta, 2), \ldots, D_{E_K}^{(N)}(\delta, 2^n - 1)\right)$$

the vector of all corresponding counters. As discussed in Section 2, each individual counter is distributed according to a binomial distribution $\mathcal{B}(N, \tilde{P}_E(\delta, \gamma))$. As each pair of the N pairs with input difference δ results in exactly one output difference, we have that

$$\sum_{\gamma} D_{E_K}^{(N)}(\delta, \gamma) = N.$$

Thus, assuming that this is the only dependency between the counter values, the vector $\mathcal{D}_{E_K}^{(N)}(\delta)$ follows a multinomial distribution with parameters N and $\tilde{P}_E(\delta) := \left(\tilde{P}_E(\delta, 1), \ldots, \tilde{P}_E(\delta, 2^n - 1)\right)$, denoted by $\mathcal{D}_{E_K}^{(N)}(\delta, \gamma) \sim \text{Multi}(N, \tilde{P}_E(\delta))$.

Later in this work we present experimental evidence comparing the empirical and theoretical gain of the attack to justify this assumption for the ciphers considered in this work. We summarise our assumption on the behaviour below.

Assumption 3. *The vector of counters $\mathcal{D}_{E_K}^{(N)}(\delta)$ follows a multinomial distribution where each component is distributed according to Assumption 1 and*

$$\sum_{\gamma} D_{E_K}^{(N)}(\delta, \gamma) = N.$$

In contrast to previous works, we do not rely on the Wrong-Key Randomization Hypotheses (Assumption 2) for our attack. Before mounting our attack, the attacker has to compute the expected probability (or the expected counter value) for all possible output differences. If the attacker is able to do this, he is usually also able to compute the expected probability for wrong keys, that is compute the distribution of $D^{(N)}_{R^{-1}_{k'} \circ R_{k_{r+1}} \circ E_K}(\delta, \gamma)$ as this is essentially computing two more rounds. We note that even if k' differs from k in only a few bits, this affects at least one S-box and hence many output differences.

3.2 The Attack Algorithm

First, recall that the attack uses N plaintext/ciphertext pairs, to recover the secret key. Following the previous section, we assume that the attacker has – in an offline phase – computed the parameters of two distributions. Namely, vectors of parameters $p = (p_i)_i$ and $q = (q_i)_i$ such that

$$p_i = \tilde{P}_E(\delta, i) \tag{1}$$
$$q_i = \tilde{P}_{R^{-1} \circ R \circ E}(\delta, i). \tag{2}$$

That is, for a right key the vector of counters is a sample from the distribution $\text{Dist}_1 = \text{Multi}(N, p)$ and for the wrong keys sampled from the distribution $\text{Dist}_2 = \text{Multi}(N, q)$. After this pre-computation phase, the attack proceeds as follows . For all possible last round keys k', the attacker first computes the vector of difference counters $\mathcal{D}^{(N)}_{R^{-1}_{k'} \circ R_{k_{r+1}} \circ E_K}(\delta)$. That is, given the guess for the last round key, the attacker partially decrypts every ciphertext and for all output differences γ computes the number of pairs fulfilling the differential $\delta \to \gamma$. Next, the attacker estimates the likelihood that the vector was sampled from Dist_1. In our case, this is equivalent to computing the difference of the log-likelihood of the vector with respect to Dist_1 and with respect to Dist_2, i.e., to compute the log-likelihood-ratio.

Given that for a random variable X following a multinomial distribution $X \sim \text{Multi}(M, p)$ it holds that

$$\mathbf{Pr}((X_1, \ldots, X_n) = (x_1, \ldots, x_n)) = \begin{cases} \frac{n!}{x_1! x_2! \ldots x_n!} p_1^{x_1} \ldots p_n^{x_n} & \text{if } \sum x_i = M \\ 0 & \text{else} \end{cases},$$

the log-likelihood-ratio is given by $l_{k'} = \sum_i D^{(N)}_{R^{-1}_{k'} \circ R_{k_{r+1}} \circ E_K}(\delta, i) \log\left(\frac{p_i}{q_i}\right)$ Thus, denoting $w_i = \log\left(\frac{p_i}{q_i}\right)$ one computes

$$l_{k'} = \sum_i w_i \cdot D^{(N)}_{R^{-1}_{k'} \circ R_{k_{r+1}} \circ E_K}(\delta, i).$$

This is a weighted extension of the case where one considers only one counter. As before these weights naturally capture various types of differential attacks,

i.e., in each component one considers a standard differential, improbable or impossible differential attack. Furthermore, truncated differentials are captured in this model since these correspond to a sub-vector of $\mathcal{D}^{(N)}_{E_K}(\delta)$.[3]

The time complexity is $|K'| \cdot N$ where N is the number of pairs considered and $|K'|$ is the number of all last-round subkeys.

3.3 Computing the Gain of the Attack

What remains to be established is the efficiency of this attack. The key observation (cf. also [2]) is that the distribution of $l_{k'}$ can be well approximated by a normal distribution in the case where all values w_i are relatively close together. The case where all w_i are close to uniform is the most interesting case for our attack, as otherwise standard differential techniques, considering only one counter are sufficient to break the cipher. Recall that there are two distributions to be considered. First, there is a random variable (and a corresponding distribution) for the log-likelihood-ratio of the right key. We denote this random variable by $\mathcal{R}$ and it is defined as $\mathcal{R} = \sum_i w_i D^{(N)}_{E_K}(\delta, i)$. By Assumption 3 we expect $\mathcal{D}^{(N)}_{E_K}(\delta)$ to be multinomial distributed with parameters N and $(p_i)_i$, with p_i defined in Equation (1). Hence the expected value of $\mathcal{R}$ is given by $E(\mathcal{R}) = N \sum w_i p_i$. Using that the pairwise covariances for a multinomial distribution is known, the variance of $\mathcal{R}$ can be computed to be

$$\mathrm{Var}(\mathcal{R}) = N \left(\left(\sum_i w_i^2 p_i \right) - \left(\sum_i w_i p_i \right)^2 \right).$$

Therefore, denoting by $\mathcal{N}(E, V)$ the normal distribution with expected value E and variance V, we will use the following approximation

$$\mathcal{R} \sim \mathcal{N} \left(N \sum w_i p_i, N \left(\left(\sum_i w_i^2 p_i \right) - \left(\sum_i w_i p_i \right)^2 \right) \right)$$

which we will justify with experimental results later in this work.

For the wrong keys, we introduce a random variable $\mathcal{W}$ and, following the same lines of argumentation, we approximate the distribution of $\mathcal{W}$ with a normal distribution, as follows

$$\mathcal{W} \sim \mathcal{N} \left(N \sum w_i q_i, N \left(\left(\sum_i w_i^2 q_i \right) - \left(\sum_i w_i q_i \right)^2 \right) \right)$$

with q_i as defined in Equation (2). This enables to estimate the gain of the attack. The gain is related to the probability that a wrong key candidate is

[3] We note, however, that in the case of truncated differential attacks we might have to assume that Assumption 2 holds.

ranked higher than the right key candidate. More precisely, if the task is to recover an n bit key and the rank of the correct key is r on average the gain is defined as $-\log_2 \frac{2r-1}{2^n}$. Given the probability that a wrong key is ranked higher than the right key the expected number of wrong keys ranked higher than the right key can be computed. This corresponds in turn to the expected rank of the right key.

For analyzing this, we assume that the right key value is sampled according to $\mathcal{R}$. As the normal distribution is symmetric, with a probability of $1/2$, the result is larger or equal to $E(\mathcal{R})$. For the wrong keys values are sampled from $\mathcal{W}$. For 50% percent of the keys, computing the gain is now reduced to computing the probability that $\mathcal{W} \geq E(\mathcal{R})$, as this corresponds to an upper bound on to the probability that a wrong key is ranked above the right key. Using the density function of $\mathcal{W}$, defined as

$$f_W(x) = \frac{1}{\sqrt{2\pi \operatorname{Var}(\mathcal{W})}} e^{-\frac{1}{2\operatorname{Var}(\mathcal{W})}(x-E(\mathcal{W}))^2}$$

this probability of a wrong key being ranked higher than the right key is given by $\mathbf{Pr}(\mathcal{W} \geq E(\mathcal{R})) = \int_{E(\mathcal{R})}^{\infty} f_W(x)$. Using the relation of the standard Normal distribution and the Gaussian error function, this can be rewritten as

$$\mathbf{Pr}(\mathcal{W} \geq E(\mathcal{R})) = \frac{1}{2}\left(1 - \operatorname{erf}\left(\frac{E(\mathcal{R}) - E(\mathcal{W})}{\sqrt{2\operatorname{Var}(\mathcal{W})}}\right)\right). \tag{3}$$

Concluding this part, we have now at hand an expression that allows us two compute the gain of the attack. Moreover, compared to computing the values of p_i and q_i, the time for evaluating the above expressions is negligible. We will make use of this in Section 4 where the model is applied to SmallPRESENT-[4] and KATAN. The experimental data given there justifies in turn the model for those two ciphers.

More Input Differences. A straight-forward extension which does not require any change to the analysis above is to use a different subset of input- and output-differences. In particular, the attack might benefit from not only using one vector $\mathcal{D}^{(N)}_{E_K}(\delta)$ but several such vectors for several choices of δ. We followed this approach in our experiments for SmallPRESENT-[4].

4 Application

In this section, we apply our framework to two blockciphers with very small block sizes. First, we consider SmallPRESENT-[4] to demonstrate the idea and then we consider reduced round variants of KATAN-32 for which we present the currently best known differential attack.

4.1 Toy Example SmallPRESENT-[4]

SmallPRESENT-[s] [19] is a small-scale (toy) cipher designed to aid the development and verification of cryptanalysis techniques. The cipher is an SP-network with s parallel 4-bit S-box applications. Hence the block size is $4s$. The permutation layer is a simple permutation of wires. We focus on SmallPRESENT-[4], the version with 16 bit block size, as this allows us to derive sufficient experimental data rather quickly. The S-box S is the same as for PRESENT (cf. [8]) itself and the round keys are independent. For more details we refer to [19]. A standard differential attack, with one round of partial decryption, seems feasible for not more than 7 rounds. By looking at all the whole vector of output differences, we are able to break 9 rounds with a significant gain. Moreover, compared to standard differential attacks, the data complexity for 7 rounds is reduced by a factor of 2^5. We summarise our findings for attacking $7, 8$ and 9 rounds in Table 1. All attacks in Table 1 are 1R attacks. Hence, the length of the differentials is 6, 7 and 8 respectively. In Table 1 we give the number of input differences considered, the values for $E(\mathcal{R}), V(\mathcal{R}), E(\mathcal{W}), V(\mathcal{W})$ and the number of right-key ranks smaller a than given threshold observed in 100 experiments (except for the last column, see below) compared with the number of such ranks predicted by our model (given in brackets in Table 1).

Table 1. Experimental Results for SmallPRESENT-[4]

#rounds	7	7	8	8	9	9
Data used	2^{16}	2^{9}	2^{16}	2^{16}	2^{16}	2^{16}
$\#\Delta$	1	1	1	5	1	60
$E(\mathcal{R})$	53.8210	0.8409	2.2250	9.7636	0.0570	1.5181
$V(\mathcal{R})$	124.0870	1.9388	4.6110	20.1537	0.1130	3.0441
$E(\mathcal{W})$	−47.2890	−0.7389	−2.1490	−9.4631	−0.0560	−1.5141
$V(\mathcal{W})$	84.2370	1.3162	4.1520	18.3502	0.1120	3.0203
#ranks $< 2^0$	100 (10000.00)	4 (1.10)	1 (2.70)	57 (61.95)	0 (6.81E-3)	1 (0.56)
#ranks $< 2^1$		4 (1.50)	2 (3.90)	65 (67.64)	0 (0.01)	2 (0.84)
#ranks $< 2^2$		4 (2.10)	3 (5.40)	73 (73.13)	0 (0.02)	2 (1.26)
#ranks $< 2^3$		4 (2.90)	5 (7.40)	79 (78.30)	0 (0.05)	2 (1.96)
#ranks $< 2^4$		4 (4.10)	8 (10.20)	82 (83.03)	0 (0.09)	2 (2.87)
#ranks $< 2^5$		4 (5.70)	11 (13.70)	84 (87.21)	1 (0.16)	5 (4.27)
#ranks $< 2^6$		5 (7.80)	16 (18.30)	88 (90.78)	1 (0.30)	8 (6.23)
#ranks $< 2^7$		9 (10.70)	25 (24.20)	91 (93.69)	1 (0.56)	14 (8.96)
#ranks $< 2^8$		14 (14.50)	31 (31.30)	95 (95.95)	3 (1.04)	22 (12.67)
#ranks $< 2^9$		19 (19.60)	42 (39.80)	96 (97.59)	3 (1.92)	28 (17.57)

For comparison, the best *6 round differential* for one active S-box is $\delta =$ `0x0007`, $\gamma =$ `0x0404` where $\tilde{P}_E(\delta, \gamma) = 2^{-13.57}$ which is still sufficient to mount a standard differential attack. Consequently, our attack always succeeded as well. However, to go beyond standard differential attacks, even when using only

2^9 pairs, which for a standard differential attack would not be sufficient, we expect and observe a gain of more than 3.5 for 50% of the keys (cf. column 3 to Table 1). The best *7 round differential* for one active S-box is $\delta = \texttt{0x0007}, \gamma = \texttt{0x0505}$ where $\tilde{P}_E(\delta, \gamma) = 2^{-15.39}$ which is not sufficient to mount a standard differential attack while our attack provides a gain of 5.97 for 50% of the keys. Using $N = 2^{14}$, $N = 2^{13}$, $N = 2^{12}$ and $N = 2^{11}$ we get a gain of 3.954, 2.821, 2.159 and 1.758 respectively. Using more than one input difference and the full code book, namely `0x0007`, `0x000f`, `0x0700`, `0x0070` and `0x0f00` we expect and observe (cf. column 5 of Table 1) a gain of 18.03 for 50% of the keys. Finally, the best *8 round differential* for one active S-box is $\delta = \texttt{0x0007}, \gamma = \texttt{0x5055}$ where $\tilde{P}_E(\delta, \gamma) = 2^{-15.92}$.. Our attack has a gain of 1.44 for 50% of the keys(cf. column 6 of Table 1). Using all sixty input differences where one S-box is active in round one, we expect a gain of 4.625 which is better than exhaustive key search (over half the key space) by a factor of 3.625. Our experimental results for this case are presented in the last column of Table 1.

4.2 Application to KATAN-32

KATAN-32 is one member of a family of ciphers defined in [10]. It has a block-size of 32 bits, an 80 bit key and 254 rounds. The relatively small block-size of 32 bits makes it an interesting target for our technique. The plaintext is loaded into two registers of length 13 and 19, respectively. In each round, two bits of the registers are updated, involving one key bit each. We refer to [10] for more information. The currently best know differential attack on KATAN-32 is a conditional differential attack that can break up to 78 rounds given 2^{22} chosen plaintext/ciphertext pairs (see [14]). The best attack overall breaks the full cipher slightly faster than exhaustive key search [13]. Note that, for KTANTAN-32, which differs from KATAN-32 only in the key-scheduling, better attacks are known (see [9,23]) but they do not apply to KATAN-32.

Below, we always assume $\delta = \texttt{0x1006a880}$ which is the input difference for the best known differential characteristic which holds with probability 2^{-31} after 71 rounds disregarding any dependencies. Note, however, that the special structure of KATAN-32 means that in fact the first probabilistic difference only depends on the plaintext values and not on the key values.

We consider a $\ell = 24$R attack below to maximise the number of rounds. This implies a computational cost of $2^{32} \cdot 2^{2\ell} = 2^{32+48} = 2^{80}$ partial decryptions in the online phase of the attack. Exhaustive search over half the key space would have to perform 2^{79} full encryptions where one full encryption costs roughly 4 partial decryptions. Hence, our attacks are twice as fast as exhaustive search. However, we emphasise that compared to exhaustive search the gain in our attack is significantly smaller.

71 + 24 Rounds of KATAN-32. The best output difference $\gamma = \texttt{0x00000008}$ has probability $\tilde{P}_E(\delta, \gamma) \approx 2^{-29.52}$, the output difference with the lowest probability is $\tilde{\gamma} = \texttt{0x00000080}$ with $\tilde{P}_E(\delta, \gamma) \approx 2^{-32.10}$. We get $E(\mathcal{R}) \approx 2505.211$, $V(\mathcal{R}) \approx$

5096.661, $E(\mathcal{W}) \approx -2467.448$, $V(\mathcal{W}) \approx 4868.280$. Which gives an expected gain of > 50 for 50% of the keys. We verified this estimate by considering the 16 differences with the highest probability. We compared randomly chosen right keys with randomly chosen wrong keys and always recovered the right key as the key with the highest rank.

91 + 24 Rounds of KATAN-32. The best output difference $\gamma = \texttt{0x00400000}$ has probability $\tilde{P}_E(\delta, \gamma) \approx 2^{-31.98}$, the output difference with the lowest probability is $\tilde{\gamma} = \texttt{0x02000000}$ with $\tilde{P}_E(\delta, \tilde{\gamma}) \approx 2^{-32.00}$. We get $E(\mathcal{R}) \approx 0.3695390$, $V(\mathcal{R}) \approx 0.7390803$, $E(\mathcal{W}) \approx -0.3695384$, $V(\mathcal{W}) \approx 0.7390745$. Which gives an expected gain of 2.3586180 for 50% of the keys. The expected gain for 50% of the keys for 92 and 94 rounds is 1.9220367 and 1.2306869 respectively.

5 Conclusions and Further Work

In this work we presented a unifying framework for several standard differential attacks. This unified view allows to naturally consider multiple differentials and by that improving upon known results. Our framework always provides better success probabilities than any of the combined differential attacks alone; although at the potential cost of increased computation time and memory. We demonstrated the viability of our approach by extending the the best differential for SmallPRESENT-[4] by two rounds and the best known differential for KATAN-32 by 20 rounds.

However, for many ciphers computing the distribution of counter values, i.e., $\mathcal{D}^{(N)}_{E_K}(\delta)$, is prohibitively expensive. For example, computing the Markov model exactly for KATAN-48 would require $\mathcal{O}(2^{48})$ memory which is well beyond what is feasible today. Yet, starting from one difference computing one or two rounds is usually feasible since only few output differences are possible after such a small number of rounds. It is thus possible to extend a standard differential attack using techniques discussed in this work. Instead of considering $\mathcal{D}^{(N)}_{E_K}(\delta)$ the attacker would consider $\mathcal{D}^{(N)}_{R_{k_r}}(\delta)$. Then, in the online phase of the attack counters are weighted accordingly to their distribution. We leave the details of such an approach open for further investigation.

Acknowledgements. We would like to thank Sean Murphy for helpful discussions on the log-likelihood ratio and its relation to the maximum likelihood estimation. We would also like to thank Céline Blondeau, Benoit Gérard and Kaisa Nyberg for helpful discussions on an earlier draft of this work. Furthermore, we would like to thank William Stein for allowing us to use his computers purchased under National Science Foundation Grant No. DMS-0821725. The second author gratefully acknowledges the support from the Danish National Research Foundation and the National Science Foundation of China(Grant No.11061130539) for the Danish-Chinese Center for Applications of Algebraic Geometry in Coding Theory and Cryptography. Finally, we would like to thank Michael Hortmann for introducing the authors to cryptography.

References

1. Albrecht, M.R., Leander, G.: An all-in-one approach to differential cryptanalysis for small block ciphers. Cryptology ePrint Archive, Report 2012/401 (2012), http://eprint.iacr.org/
2. Baignères, T., Junod, P., Vaudenay, S.: How Far Can We Go Beyond Linear Cryptanalysis? In: Lee, P.J. (ed.) ASIACRYPT 2004. LNCS, vol. 3329, pp. 432–450. Springer, Heidelberg (2004)
3. Biham, E., Biryukov, A., Shamir, A.: Cryptanalysis of Skipjack Reduced to 31 Rounds Using Impossible Differentials. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592, pp. 12–23. Springer, Heidelberg (1999)
4. Biham, E., Shamir, A.: Differential Cryptanalysis of DES-like Cryptosystems. In: Menezes, A., Vanstone, S.A. (eds.) CRYPTO 1990. LNCS, vol. 537, pp. 2–21. Springer, Heidelberg (1991)
5. Blondeau, C., Gérard, B.: Links between theoretical and effective differential probabilities: Experiments on PRESENT. In: Ecrypt II Workshop on Tools for Cryptanalysis (2010)
6. Blondeau, C., Gérard, B.: Multiple Differential Cryptanalysis: Theory and Practice. In: Joux, A. (ed.) FSE 2011. LNCS, vol. 6733, pp. 35–54. Springer, Heidelberg (2011)
7. Blondeau, C., Gérard, B., Nyberg, K.: Multiple Differential Cryptanalysis using LLR and χ^2 Statistics. Cryptology ePrint Archive, Report 2012/360 (2012), http://eprint.iacr.org/
8. Bogdanov, A.A., Knudsen, L.R., Leander, G., Paar, C., Poschmann, A., Robshaw, M., Seurin, Y., Vikkelsoe, C.: PRESENT: An Ultra-Lightweight Block Cipher. In: Paillier, P., Verbauwhede, I. (eds.) CHES 2007. LNCS, vol. 4727, pp. 450–466. Springer, Heidelberg (2007)
9. Bogdanov, A., Rechberger, C.: A 3-Subset Meet-in-the-Middle Attack: Cryptanalysis of the Lightweight Block Cipher KTANTAN. In: Biryukov, A., Gong, G., Stinson, D.R. (eds.) SAC 2010. LNCS, vol. 6544, pp. 229–240. Springer, Heidelberg (2011)
10. De Cannière, C., Dunkelman, O., Knežević, M.: KATAN and KTANTAN — A Family of Small and Efficient Hardware-Oriented Block Ciphers. In: Clavier, C., Gaj, K. (eds.) CHES 2009. LNCS, vol. 5747, pp. 272–288. Springer, Heidelberg (2009)
11. Daemen, J., Rijmen, V.: Probability distributions of correlation and differentials in block ciphers. Cryptology ePrint Archive, Report 2005/212 (2005), http://eprint.iacr.org/
12. Hong, D., Sung, J., Hong, S.H., Lim, J.-I., Lee, S.-J., Koo, B.-S., Lee, C.-H., Chang, D., Lee, J., Jeong, K., Kim, H., Kim, J.-S., Chee, S.: HIGHT: A New Block Cipher Suitable for Low-Resource Device. In: Goubin, L., Matsui, M. (eds.) CHES 2006. LNCS, vol. 4249, pp. 46–59. Springer, Heidelberg (2006)
13. Knellwolf, S.: Meet-in-the-Middle cryptanalysis of KATAN. In: ECRYPT Workshop on Lightweight Cryptography 2011 (to appear)
14. Knellwolf, S., Meier, W., Naya-Plasencia, M.: Conditional Differential Cryptanalysis of NLFSR-Based Cryptosystems. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 130–145. Springer, Heidelberg (2010)
15. Knudsen, L.: DEAL – a 128-bit block cipher. Technical report, Department of Informatics, University of Bergen, Norway (1998)

16. Lars, R.: Truncated and Higher Order Differentials. In: Preneel, B. (ed.) FSE 1994. LNCS, vol. 1008, pp. 196–211. Springer, Heidelberg (1995)
17. Lai, X.: Higher order derivatives and differential cryptanalysis. In: Communications and Cryptography (1994)
18. Lai, X., Massey, J.L.: Markov Ciphers and Differential Cryptanalysis. In: Davies, D.W. (ed.) EUROCRYPT 1991. LNCS, vol. 547, pp. 17–38. Springer, Heidelberg (1991)
19. Leander, G.: Small scale variants of the block cipher PRESENT. Cryptology ePrint Archive, Report 2010/143 (2010), http://eprint.iacr.org/
20. Murphy, S.: The analysis of simultaneous differences in Differential Cryptanalysis (2011), http://www.isg.rhul.ac.uk/~sean/SimDiffA.pdf
21. Murphy, S., Piper, F., Walker, M., Wild, P.: Likelihood estimation for block cipher keys (1995), http://www.isg.rhul.ac.uk/~sean/maxlik.pdf
22. Tezcan, C.: The Improbable Differential Attack: Cryptanalysis of Reduced Round CLEFIA. In: Gong, G., Gupta, K.C. (eds.) INDOCRYPT 2010. LNCS, vol. 6498, pp. 197–209. Springer, Heidelberg (2010)
23. Wei, L., Rechberger, C., Guo, J., Wu, H., Wang, H., Ling, S.: Improved Meet-in-the-Middle cryptanalysis of KTANTAN. Cryptology ePrint Archive, Report 2011/201 (2011), http://eprint.iacr.org/