Blind Differential Cryptanalysis for Enhanced Power Attacks

Helena Handschuh[1] and Bart Preneel[2]

[1] Spansion,
7 Avenue Georges Pompidou,
92593 Levallois-Perret Cedex, France
helena.handschuh@spansion.com
[2] Katholieke Universiteit Leuven, Dept. Electrical Engineering-ESAT/COSIC,
Kasteelpark Arenberg 10,
B-3001 Leuven-Heverlee, Belgium
bart.preneel@esat.kuleuven.be

Abstract. At FSE 2003 and 2004, Akkar and Goubin presented several masking methods to protect iterated block ciphers such as DES against Differential Power Analysis and higher-order variations thereof. The underlying idea is to randomize the first few and last few rounds of the cipher with independent masks at each round until all intermediate values depend on a large number of secret key bits, thereby disabling power attacks on subsequent inner rounds. We show how to combine differential cryptanalysis applied to the first few rounds of the cipher with power attacks to extract the secret key from intermediate unmasked (unknown) values, even when these already depend on all secret key bits. We thus invalidate the widely believed claim that it is sufficient to protect the outer rounds of an iterated block cipher against side-channel attacks.

Keywords: differential cryptanalysis, power analysis, side channel attacks, Hamming weights, combined cryptanalysis, blind cryptanalysis.

1 Introduction

In 1998, Kocher et al. introduced Differential Power Attacks on block ciphers and digital signature algorithms [10]. These attacks allow to recover secrets used in cryptographic computations even if these are executed inside tamper-resistant devices such as smart cards. Kocher noted that these devices leak information which is directly correlated to the secret data being manipulated inside the device. The information may be recovered for example by measuring the power consumption of the device and the correlation of its variation with the secret data. Differential Power Analysis exploits the fact that computing a given output bit of a non-linear S-box requires different power consumption when this bit is set to zero or to one; correlation analysis extends this by correlating the power consumption with the key dependent power consumption predicted by a model. Since 1998, these techniques have been generalized to other side-channels such as timing information, electro-magnetic radiation, and even sound waves; research

E. Biham and A.M. Youssef (Eds.): SAC 2006, LNCS 4356, pp. 163–173, 2007.

has also focused on how to protect tamper-resistant computations against these attacks. Countermeasures are applied at the hardware, software and protocol level. At the hardware level, power consumption scramblers and ad-hoc noise introduction via random execution delays or random operation execution are the preferred methods. At the software level, the most useful technique against first order differential attacks are randomization techniques. In essence, the intermediate values are blinded using some randomized masks in order to decorrelate them from the actual values which would reveal information to the opponent. These techniques include random masking methods, randomized exponentiation techniques, and randomized execution paths or integer representation. At the protocol level, fast key refreshing has been a useful countermeasure. For further references on side channel attacks and countermeasures, see for example [13,5,9,23,8,20].

Some of these methods have since been shown to be vulnerable to higher order differential attacks (see Messerges [14]) in which an opponent can measure information at different places in a single power consumption curve. Kunz-Jacques et al. [11] have shown how to improve higher order attacks on DES by combining them with the Davies-Murphy attack. For block ciphers, new masking methods have been proposed in which an independent random mask is applied at each round, thus preventing an attacker to take advantage of the repeating mask in a higher-order differential attack. However, these protection methods require very large quantities of volatile memory and pre-computation time, which is typically cost-prohibitive in secure embedded devices. Therefore only a few rounds are eventually masked against power analysis attacks and the inner rounds of the cipher are left unmasked (see Akkar et al. [3,2]). In their paper on cache-based attacks, Osvik et al. [18] independently suggested that differential attacks could be used to bypass protection in the outer rounds, but no details are provided.

Our Contribution. In this paper we show how to attack secure implementations of iterated block ciphers which apply reduced-round masking methods to protect their secret keys against side-channel attacks such as power attacks. We are able to mount key recovery attacks based on differential cryptanalysis techniques [4] and power traces providing only the Hamming weight of the internal variables used throughout the computation. Compared to differential cryptanalysis, the main difference is that our technique is *blind* in the sense that we do not *see* the actual values at the output of the differential path since the path stops somewhere in the middle rounds of the cipher, but can only *derive* them from their measured Hamming weight. As an example we explain how the technique works on the Unified Masking Method applied to the DES.

Organization of the Paper. Section 2 explains the unique masking method and its extensions which formed the inspiration of this attack. In Sect. 3 we explain our blind differential attack for the specific case of DES with the outer four rounds masked. In Sect. 4 we present our simulation results. Section 5 discusses improvements and generalizations and Sect. 6 concludes the paper.

2 Extended Unique Masking Method

The Unique Masking Method (UMM) described below was proposed by Akkar and Goubin [3] and applies to Feistel ciphers such as DES [15] and Substitution Permutation Networks such as AES [16]. We use here the first type as an example. In Feistel ciphers, the plaintext M is split into two halves L_0 and R_0 such that $M = L_0||R_0$ and a round function f is applied to the right half of its input before the result is exored to the left half. Next, both halves are swapped and the procedure iterates for r rounds:

$$L_{i+1} := R_i \quad \text{and} \quad R_{i+1} := L_i \oplus f(R_i) \quad 0 \leq i \leq r-1 \ .$$

The ciphertext is equal to $C = R_r||L_r$ (in the last round, the halves are not swapped). The DES round function comprises a key addition operation, followed by an expansion operation E and substitution through a layer of 8 tables or S-boxes S (each mapping 6 bits to 4 bits); next a bit-level permutation P is applied to the result. UMM proposes to mask the outer rounds of a Feistel block cipher such as DES with different independent masks at each round for at least four rounds in order to decorrelate the calculations from the actual intermediate data. In order to achieve this, the S-boxes S are replaced with different S-boxes, the input and output of which are masked with random data. Since these S-boxes are the only non-linear part of the cipher, new S-boxes need to be generated dynamically at each execution of the algorithm to account for the different input and output translations introduced by the random masks. UMM uses two sets of S-boxes.

Let S_1 and S_2 denote the following two new functions based on the original DES S-boxes S, where α represents the 32-bit input and output mask used during one execution of the algorithm:

$$\begin{cases} \forall\, x \in \{0,1\}^{48}, & S_1(x) = S(x \oplus E(\alpha)) \\ \forall\, x \in \{0,1\}^{48}, & S_2(x) = S(x) \oplus P^{-1}(\alpha) \ . \end{cases}$$

These two new functions are logically combined such that one output mask synchronizes with the input mask of the next round automatically as shown in Fig. 1. Akkar et al. show in [2] that their initial UMM method does not achieve the desired goal as the second round output remains unmasked. Therefore they propose to use a third independent set of S-boxes S_3 such that $\forall\, x \in \{0,1\}^{48}$, $S_3(x \oplus E(\alpha)) = S(x) \oplus P^{-1}(\alpha)$ and completely mask all intermediate data up to the fourth round of DES as shown in Fig. 2. A different value for α should be used at each execution of the algorithm. This scheme uses S_3 in both round 2 and 3; one could avoid this by introducing an additional mask β and by defining S'_3 that transforms a mask α into a mask β and S''_3 that does the opposite. Since these masking techniques require the generation of an independent translated S-box for each round, which represents a large cost in terms of volatile memory (256 bytes of RAM each), only four rounds are masked until each intermediate data bit depends on all the key bits of the block cipher. By combining Power Attacks with Differential Cryptanalysis, we will show that this scheme is still not secure even if four or more independent S-box layers are chosen in each run of the algorithm.

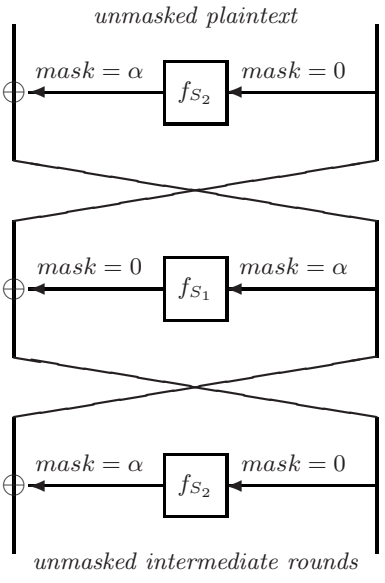

Fig. 1. Example application of Akkar and Goubin's initial masked f-function chaining method with three masked rounds [3]

3 Mounting a Blind Differential Attack on 4-Round DES

Differential attacks as first described by Biham and Shamir [4] are chosen plaintext attacks in which an adversary chooses pairs of plaintexts with given differences and tries to deduce information from the corresponding pairs of ciphertexts. With a certain probability, a given plaintext difference follows a pre-determined path throughout the encryption operations and results in a given ciphertext difference. An input pair that results in the correct intermediate and output differences is called a right pair. When the adversary finds a right pair, he can deduce information on the last round key from the pre-determined differential path. For more details we refer the reader to the original paper describing this technique [4]. We apply differential cryptanalysis to four-round DES using Biham's original four-round differential characteristics. The innovation in our attack is that we cannot observe the differences directly in the ciphertext pair at the output of the cipher as these differences only appear in internal rounds of the encryption process. Therefore we call our method 'blind differential cryptanalysis.'

3.1 Enhanced Power Attacks

Now Power Attacks come into play. In power attacks, a generally admitted model is that the adversary can measure side-channel information which leaks a linear function of the individual data bits, for example the Hamming weight of the data (see for example [1,6,7] for a discussion of this model). In this setting, we can

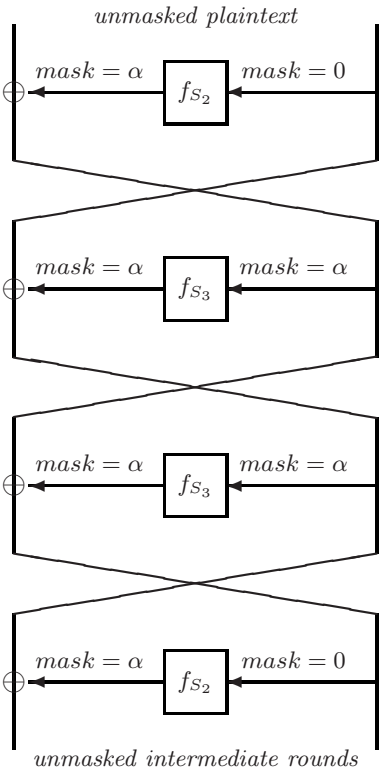

Fig. 2. Example application of Akkar, Bévan and Goubin's improved masked f-function chaining method with four masked rounds [2]

combine the expected value of the difference at round four of the DES and the Hamming weight observations of the power attack. In our four-round differential represented in Fig. 3, several difference bytes are equal to zero, meaning that the corresponding data bytes are equal. This in turn implies that the Hamming weights are equal. Note that the converse is not necessarily true. Thus, our power measurements will both enable us to *see* when two bytes are potentially equal and will provide a large number of false alarms at the same time, since equal Hamming weights do not imply equal data bytes. However, filtering out the pairs which *do* reveal the same Hamming weight on the required data bytes, we can now apply our blind differential key recovery attack to recover part of the secret key of the fourth round. Note that a collision technique (i.e., searching for identical values and thus Hamming weights) has been used by Schramm et al. to improve power analysis attacks [21]. Ledig et al. [12] showed that this attack can be further enhanced by exploiting the slow increase of Hamming distances in the rounds following a collision. However, in our attack we explicitly make use of right pairs for arbitrary characteristics and we show how the key can be recovered even if no collisions occur at all.

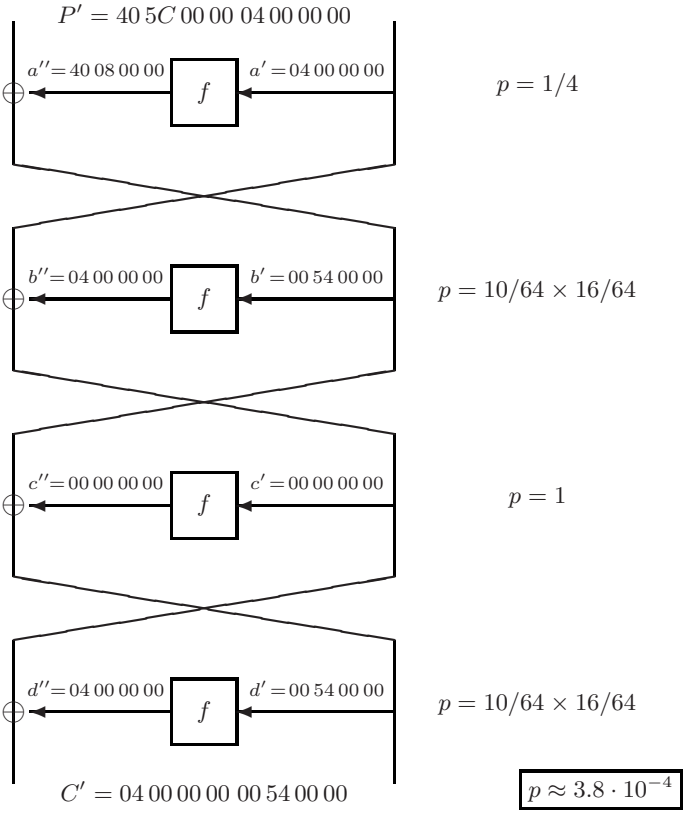

Fig. 3. Biham and Shamir's 4-round differential characteristic for DES and the associated probabilities [4]

3.2 Blind Key Recovery

Once we have filtered out the right pairs using equal Hamming weights on the 6-bit S-box inputs, we need to find a technique which allows to recover the secret key bits involved in this round without knowing the actual intermediate values. Recall that we also have access to the absolute value of the Hamming weight of the data before key addition. For every DES S-box it is easy to construct the difference distribution table of DES and to determine which input pairs yield the right output difference. These actual input values will now be used in the following way.

Consider an active S-box in the fourth round, that is, an S-box with a non-zero input difference δ and output difference Δ. Denote the 6-bit input of the fourth round corresponding to this S-box by x_i, the 6-bit key by k, the 6-bit input of the S-box by y_i and the 4-bit output of the S-box by z_i. Clearly one has $y_i = x_i \oplus k$. In classical differential cryptanalysis applied to a Feistel cipher, one finds a number of right pairs and then deduces candidate values for k from the

values (y_i, y_i') that correspond to the characteristic and from the known values of (x_i, x_i') (note that $x_i \oplus x_i' = y_i \oplus y_i' = \delta$). Ohta and Matsui [17] and Preneel et al. [19] have extended this attack to the case where not all bits of the x_i and y_i are known in order to attack CBC-MAC and the CFB mode of DES (or reduced-round variants of DES).

In our new attack only the Hamming weights of the intermediate plaintexts (x_i, x_i') are known. At first sight it seems rather easy to deduce candidate values for k by generating all the 6-bit values with the correct Hamming weight and eliminating those which are not compatible with the characteristic. The remaining candidate intermediate plaintexts then suggest several values for the key k; by considering multiple right pairs, the correct value of k should appear. Unfortunately this attack does not work, since all or almost all 64 key values are suggested by each right pair. Therefore we have developed a new approach.

Consider the set $Y = \{(y_i, y_i') \mid y_i \oplus y_i' = \delta \text{ and } z_i \oplus z_i' = \Delta\}$. Consider a fixed value of k. Define the set X_k as follows:

$$X_k = \{(x_i, x_i') = (y_i \oplus k, y_i' \oplus k) \mid (y_i, y_i') \in Y\} \ .$$

This set can be partitioned according to the value $(\mathrm{hwt}(x_i), \mathrm{hwt}(x_i'))$. Note that due to the constraint $x_i \oplus x_i' = \delta$ not all combinations of these integers can occur. It is easy to see the following cases:

$\mathrm{hwt}(\delta) = 1$**:** then $\mathrm{hwt}(x_i) = \mathrm{hwt}(x_i') \pm 1$
$\mathrm{hwt}(\delta) = 2$**:** then $\mathrm{hwt}(x_i) = \mathrm{hwt}(x_i') \pm 2$ or $\mathrm{hwt}(x_i) = \mathrm{hwt}(x_i')$
$\mathrm{hwt}(\delta) = 3$**:** then $\mathrm{hwt}(x_i) = \mathrm{hwt}(x_i') \pm 1$ or $\mathrm{hwt}(x_i) = \mathrm{hwt}(x_i') \pm 3$

We now define the Hamming weight profile PP_k of the key k as follows:

$$\mathrm{PP}_k[i,j] = |\{(x_i, x_i') \mid (x_i, x_i') \in X_k \text{ and } \mathrm{hwt}(x_i) = i, \mathrm{hwt}(x_i') = j\}| \ ,$$

or in words: $\mathrm{PP}_k[i,j]$ is the number of input pairs to an active S-box for which the round inputs have Hamming weight i and j respectively.

The attack proceeds as follows:

1. Collect a sufficient number of right pairs; note that since our filtering mechanism based on Hammming weights is not perfect, not all the retained pairs will be right pairs.
2. Compute an estimate for the Hamming weight profile $\hat{\mathrm{PP}}_k[i,j]$ based on the measurements.
3. Perform a matching between the observed profile and the profiles of all the values for k. We propose the use of a mean square error (MSE) as a matching criterion
$$\mathrm{MSE}_{k^*} = \sum_{[i,j]} \left(\mathrm{PP}_{k*}[i,j] - \hat{\mathrm{PP}}_k[i,j]\right)^2 \ .$$
The idea of the attack is that for the correct value of k the Euclidean distance between the two profiles will be very small (and typically smaller than that for other keys).

Note that due to the symmetry property, we have that $\mathrm{PP}_k[i,j] = \mathrm{PP}_{k\oplus\delta}[i,j]$, hence we cannot distinguish between k and $k \oplus \delta$.

Overall, a few dozen right pairs and the corresponding power measurements provide enough information about the Hamming weights of intermediate data before the S-boxes to recover two candidates for the secret 6-bit key. Note however that this technique only allows to recover key elements corresponding to the active S-boxes. Therefore we need to use several *different* four-round characteristics to recover all 6-bit elements of the secret key. Fortunately, there are many four-round characteristics available for differential cryptanalysis of DES, and we do not have to use only the best one in our attack. Once sufficient key bits have been obtained, the remaining bits can be recovered by exhaustive search.

4 Simulations

We have performed some experiments in software on a PC to validate the analysis in Sect. 3. We assume that we can measure Hamming weights in a reliable way, that is, our simulated measurements are noise free. This assumption may not hold in practice, in particular if additional noise is added as a countermeasure. However, we believe that our methods are sufficiently robust to also work (with an increased number of measurements) under noisy conditions.

The probability p of the characteristic we use is $3.8 \cdot 10^{-4}$ (see Fig. 3); it has two active S-boxes. We use the Hamming weights of the left half output of the fourth round (which is also the input to the next round) to filter out right pairs. There are seven passive S-boxes in the fifth round which allows to almost uniquely identify right pairs as being those which follow the differential path. A very rough estimate shows that the probability p_f to obtain equal Hamming weights on seven 6-bit elements in the pair is approximately $3 \cdot 10^{-5}$; it can be computed as follows:

$$p_f = (p^*)^7 \quad \text{with } p^* = \frac{1}{2^{12}} \cdot \sum_{i=0}^{6} \binom{6}{i}^2 .$$

This filtering function has roughly the same probability of success as regular differential filters, and as noted in Sect. 3.1, we can guarantee that all right pairs are correctly identified, and few false alarms appear. A right pair in the sense of differential cryptanalysis will automatically yield a right pair in the sense of Hamming weights. The Hamming weight difference for a wrong pair is not uniformly distributed – it is more likely to be equal to that of a right pair. In our simulations we have noted that our Hamming weight filter is about a factor of two worse than the above rough estimate, but this is still more than sufficient for the attack to work. As an example, for about 2^{14} random plaintext pairs, we obtain 3 right pairs for differential cryptanalysis and one false alarm. For 2^{20} random plaintext pairs, we obtain 408 right pairs for differential cryptanalysis and 69 false alarms, i.e., about 15%; note that the rough estimate suggests $3 \cdot 10^{-5}/3.8 \cdot 10^{-4} \approx 7.8\%$.

Next, for every possible 6-bit key entering S-box S3 in the fourth round, we computed the Hamming weight profile of the key according to the difference distribution table for that S-box. There are 10 differential pairs which follow our characteristic for S-box S3, and thus the profile distributes these 10 possible input pairs according to the key value as described in the previous section. Note that the profiles only depend on the S-box (namely its difference distribution table and the associated differential pairs) and the key value.

Now taking our real data, we try to match the observed distribution of Hamming weight profiles at the input of S-box S3 in round four with the theoretical profiles we have using a mean square error matching criterion. With 3 right pairs and one false alarm, the right key ends up in second position. However there are many indistinguishable keys at this stage. With as few as 26 right pairs and 4 false alarms (derived from 2^{16} plaintext pairs) the right key ends up in first position and there are only 2 (indistinguishable) keys left. Thus a 20-30 right pairs suffice to recover the 6-bit key element corresponding to S-box S3. Note that if four keys survive at this stage, the overall attack only requires a small extra exhaustive search step. The experiments have been repeated for several keys, with similar or better results.

Now that we recover a few candidates for the first 6-bit key element, we continue with the adjacent S-box S4. Next we change our differential characteristic to get different active S-boxes in round four and recover the whole key piece by piece.

Since the probability of the four-round differential is about $p \approx 3.8 \cdot 10^{-4}$ and there are approximately 15% false alarms, the whole attack requires $\mathcal{O}(1/p) \approx 30\,000$ plaintext pairs with their associated power traces and Hamming weight measurements. Blind differential cryptanalysis completely bypasses any type of random data masking on the four first rounds. The computational complexity of the attack is negligible; it requires only a few seconds on a regular PC.

5 Improvements and Generalizations

So far we have only explained a rather straightforward approach and we have illustrated it with an example as a proof of concept. There are several ways in which our attack can be further optimized and improved. First, some keys are clearly easier to recover than others; we need to analyze this phenomenon in more detail in order to assess the entropy reduction of the key that can be obtained. Second we can optimize the differential characteristics for this type of attack – the example we have used is a good characteristic for a regular differential attack, but it is plausible that we can find characteristics that are better suited for a blind differential attack. Third, the attack could be expanded to taking into account measurement noise and other side channel leakage models (such as leakage of the Hamming weight transitions in the registers). Finally, the attack is independent of the masking technique and only builds on the difference distribution tables of the cipher as well as the Hamming weight profiles of the differential pairs for a given key. Hence it is clear that it can be extended to other

Feistel ciphers (including 2-key and 3-key triple-DES) but also to substitution permutation network ciphers such as the AES.

The technique we describe applies to four or more initial masked rounds, as long as high probability characteristics can be found for this reduced number of rounds of the cipher. Since for all modern block ciphers, resistance against differential cryptanalysis is achieved only after sufficiently many rounds, independent masks should be applied to just as many internal rounds. Power attacks provide enough information about Hamming weights of the intermediate values to put external round-masking techniques at risk when side-channel analysis is combined with differential cryptanalysis. Note that our attacks simply start measuring side-channel information on the first unmasked internal round. Again we stress that our techniques are completely independent from the underlying masking technique and apply to any block cipher for which well-chosen reduced-round differentials with high probability and with several colliding bytes can be found.

6 Conclusion

We have introduced the notion of blind differential cryptanalysis where an attacker uses internal differentials to by-pass outer round masking against power attacks. This technique retrieves the secret key of the target block cipher given only Hamming weight measurements on selected internal values. We therefore invalidate the widely claimed belief that only outer rounds need be protected from power attacks. Our method easily generalizes to different block cipher structures and requires a reasonable amount of plaintext-ciphertext pairs and power measurements. We believe other powerful combinations of side-channel attacks and traditional cryptanalysis will provide for interesting future developments in the area of secure embedded tokens.

Acknowledgements. We would like to thank the anonymous referees for the constructive comments.

References

1. Akkar, M.-L., Bevan, R., Dischamp, P., Moyart, D.: Power Analysis, What Is Now Possible.... In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 489–502. Springer, Heidelberg (2000)
2. Akkar, M.-L., Bevan, R., Goubin, L.: Two Power Analysis Attacks against One-Mask Methods. In: Roy, B.K., Meier, W. (eds.) FSE 2004. LNCS, vol. 3017, pp. 332–347. Springer, Heidelberg (2004)
3. Akkar, M.-L., Goubin, L.: A Generic Protection against High-Order Differential Power Analysis.. In: Johansson, T. (ed.) FSE 2003. LNCS, vol. 2887, pp. 192–205. Springer, Heidelberg (2003)
4. Biham, E., Shamir, A.: Differential Cryptanalysis of DES-like Cryptosystems. J. Cryptology 4(1), 3–72 (1991)

5. Koç, Ç.K., Naccache, D., Paar, C. (eds.): CHES 2001. LNCS, vol. 2162. Springer, Heidelberg (2001)
6. Chari, S., Jutla, C.S., Rao, J.R., Rohatgi, P.: Towards Sound Approaches to Counteract Power-Analysis Attacks. In: Wiener [24], pp. 398–412
7. Goubin, L., Patarin, J.: DES and Differential Power Analysis (The "Duplication" Method). In: Koç, Ç.K., Paar, C. (eds.) CHES 1999. LNCS, vol. 1717, pp. 158–172. Springer, Heidelberg (1999)
8. Joye, M., Quisquater, J.-J. (eds.): CHES 2004. LNCS, vol. 3156. Springer, Heidelberg (2004)
9. Kaliski Jr., B.S., Koç, Ç.K., Paar, C. (eds.): CHES 2002. LNCS, vol. 2523. Springer, Heidelberg (2003)
10. Kocher, P.C., Jaffe, J., Jun, B.: Differential Power Analysis. In: Wiener [24], pp. 388–397
11. Kunz-Jacques, S., Muller, F., Valette, F.: The Davies-Murphy Power Attack. In: Lee, P.J. (ed.) ASIACRYPT 2004. LNCS, vol. 3329, pp. 451–467. Springer, Heidelberg (2004)
12. Ledig, H., Muller, F., Valette, F.: Enhancing collision attacks. In: Joye and Quisquater [8], pp. 176–190
13. Messerges, T.S.: Securing the AES Finalists Against Power Analysis Attacks. In: Schneier, B. (ed.) FSE 2000. LNCS, vol. 1978, pp. 150–164. Springer, Heidelberg (2001)
14. Messerges, T.S.: Using Second-Order Power Analysis to Attack DPA Resistant Software. In: Paar, C., Koç, Ç.K. (eds.) CHES 2000. LNCS, vol. 1965, pp. 238–251. Springer, Heidelberg (2000)
15. National Institute of Standards and Technology (NIST) FIPS Publication 46-3: Data Encryption Standard (1999)
16. National Institute of Standards and Technology (NIST). FIPS Publication 197: Advanced Encryption Standard (AES) (2001)
17. Ohta, K., Matsui, M.: Differential Attack on Message Authentication Codes. In: Stinson [22], pp. 200–211
18. Osvik, D.A., Shamir, A., Tromer, E.: Cache attacks and countermeasures: The case of AES. In: Pointcheval, D. (ed.) CT-RSA 2006. LNCS, vol. 3860, pp. 1–20. Springer, Heidelberg (2006)
19. Preneel, B., Nuttin, M., Rijmen, V., Buelens, J.: Cryptanalysis of the CFB Mode of the DES with a Reduced Number of Rounds. In: Stinson [22], pp. 212–223
20. Rao, J.R., Sunar, B. (eds.): CHES 2005. LNCS, vol. 3659. Springer, Heidelberg (2005)
21. Schramm, K., Leander, G., Felke, P., Paar, C.: A Collision-Attack on AES: Combining Side Channel- and Differential-Attack. In: Joye and Quisquater [8], pp. 163–175
22. Stinson, D.R. (ed.): CRYPTO 1993. LNCS, vol. 773. Springer, Heidelberg (1994)
23. Walter, C.D., Koç, Ç.K., Paar, C. (eds.): CHES 2003. LNCS, vol. 2779. Springer, Heidelberg (2003)
24. Wiener, M.J. (ed.): CRYPTO 1999. LNCS, vol. 1666. Springer, Heidelberg (1999)