

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Alfred Kobsa

University of California, Irvine, CA, USA

Friedemann Mattern

ETH Zurich, Switzerland

John C. Mitchell

Stanford University, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

Oscar Nierstrasz

University of Bern, Switzerland

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

TU Dortmund University, Germany

Madhu Sudan

Microsoft Research, Cambridge, MA, USA

Demetri Terzopoulos

University of California, Los Angeles, CA, USA

Doug Tygar

University of California, Berkeley, CA, USA

Gerhard Weikum

Max Planck Institute for Informatics, Saarbruecken, Germany

Anne Canteaut (Ed.)

Fast Software Encryption

19th International Workshop, FSE 2012
Washington, DC, USA, March 19-21, 2012
Revised Selected Papers

 Springer

Volume Editor

Anne Canteaut
INRIA Paris-Rocquencourt
B.P. 105
78153 Le Chesnay, France
E-mail: anne.canteaut@inria.fr

ISSN 0302-9743 e-ISSN 1611-3349
ISBN 978-3-642-34046-8 e-ISBN 978-3-642-34047-5
DOI 10.1007/978-3-642-34047-5
Springer Heidelberg Dordrecht London New York

Library of Congress Control Number: 2012948466

CR Subject Classification (1998): E.3, K.6.5, D.4.6, C.2, J.1, G.2.1

LNCS Sublibrary: SL 4 – Security and Cryptology

© International Association for Cryptologic Research 2012

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

These are the proceedings of FSE 2012, the 19th International Workshop on Fast Software Encryption. The workshop, organized in cooperation with the International Association for Cryptologic Research, was held March 19–21, 2012, in Washington DC. The General Chair was Bruce Schneier, from British Telecom, USA.

This year, a total of 89 papers were submitted to the workshop. Each submission was reviewed by at least three Program Committee (PC) members, while submissions co-authored by PC members were reviewed by at least five PC members. After the reviews were submitted, the committee deliberated online in depth and we eventually selected 24 submissions for presentation. The authors of the accepted papers were then given more than one month to revise their manuscript and to take into account the comments from the reviewers. This revision process allowed some interactions between the authors and the PC, and I am grateful to the PC members who spent a lot of time on this and contributed to guaranteeing the high standards of these papers. At the workshop, the papers were made available to the audience in electronic form. Then, the authors prepared the final versions which are included in these proceedings. Since these final versions were not checked again before publication, the authors bear the responsibility for the contents of their papers.

The PC selected three papers for invitation to the *Journal of Cryptology*: “Improved Rebound Attack on the Finalist Grøstl” by J  r  my Jean, Mar  a Naya-Plasencia, and Thomas Peyrin, “Recursive Diffusion Layers for Block Ciphers and Hash Functions” by Mahdi Sajadieh, Mohammad Dakhilalian, Hamid Mala, and Pouyan Sepehrdad, and “New attacks on Keccak-224 and Keccak-256” by Itai Dinur, Orr Dunkelman, and Adi Shamir.

In addition to the papers included in this volume, we were fortunate to have in the program two invited talks: one by Kaisa Nyberg on “Provable” Security against Differential and Linear Cryptanalysis” and the other by Mitsuru Matsui on “The History of Linear Cryptanalysis.” An invited paper corresponding to Kaisa Nyberg’s talk is included in the proceedings. The conference also featured a rump session with short informal presentations. Dan Bernstein and Tanja Lange served as the Chairs of the rump session.

I wish to thank all the authors who submitted their work to the conference. I am very grateful to the PC members for their hard and generous work. In addition, I gratefully acknowledge the help of a number of colleagues who provided reviews for the PC. I am also indebted to Andrei Voronkov for his very nice EasyChair conference management system that helped me compile this volume.

Finally, I would like to say that being the Program Chair for FSE 2012 has been a great honor and an exciting task.

Conference Organization

General Chair

Bruce Schneier British Telecom, USA

Program Chair

Anne Canteaut INRIA Paris-Rocquencourt, France

Program Committee

Alex Biryukov	University of Luxembourg, Luxembourg
Guang Gong	University of Waterloo, Canada
Martin Hell	Lund University, Sweden
Antoine Joux	Université de Versailles Saint-Quentin-en-Yvelines and DGA, France
Pascal Junod	HEIG-VD, Switzerland
John Kelsey	NIST, USA
Dmitry Khovratovich	Microsoft Research, USA
Lars Knudsen	Technical University of Denmark, Denmark
Gregor Leander	Technical University of Denmark, Denmark
Stefan Lucks	Bauhaus-Universität Weimar, Germany
Subhamoy Maitra	ISI Kolkata, India
Willi Meier	FHNW, Switzerland
Shiho Moriai	Sony Corporation, Japan
María Naya-Plasencia	Université de Versailles Saint-Quentin-en-Yvelines, France
Elisabeth Oswald	University of Bristol, UK
Vincent Rijmen	K.U. Leuven, Belgium and TU Graz, Austria
Matt Robshaw	Orange Labs, France
Yu Sasaki	NTT Corporation, Japan
François-Xavier Standaert	Université catholique de Louvain, Belgium
Gilles Van Assche	STMicroelectronics, Belgium
Serge Vaudenay	EPFL, Switzerland

External Reviewers

Mohamed Ahmed Abdelraheem
Toru Akishita
Kazumaro Aoki
Jean-Philippe Aumasson
Subhadeep Banik
Ash Bay
Guido Bertoni
Rishiraj Bhattacharyya
Céline Blondeau
Andrey Bogdanov
Julia Borghoff
Ioana Boureanu
Qi Chai
Anupam Chattopadhyay
Jiazhe Chen
Baudoin Collard
Joan Daemen
Xinxin Fan
Matthieu Finiasz
Ewan Fleischmann
Christian Forler
Thomas Fuhr
Praveen Gauravaram
Benedikt Gierlichs
Kishan Gupta
Benoît Gérard
Honggang Hu
Takanori Isobe
Tetsu Iwata
Selçuk Kavut
Shahram Khazaei
Simon Knellwolf
Yuichi Komano
Gaëtan Leurent
Marco Macchetti
Atefeh Mashatan

Marcel Medwed
Florian Mendel
Mridul Nandi
Svetla Nikova
Kaisa Nyberg
Khaled Ouafi
Goutam Paul
Emmanuel Prouff
Christian Rechberger
Jean-René Reinhard
Arnab Roy
Santanu Sarkar
Martin Schläffer
Sourav Sen Gupta
Pouyan Sepehrdad
Yannick Seurin
Kyoji Shibusaki
Taizo Shirai
Paul Stankovski
Fatih Sulak
Petr Sušil
Soren S. Thomsen
Stefan Tillich
Elmar Tischhauser
Deniz Toz
Michael Tunstall
Kerem Varici
Lei Wang
Ralf-Philipp Weinmann
Jakob Wenzel
Carolyn Whitnall
Teng Wu
Kan Yasuda
Bo Zhu
Martin Ågren

Table of Contents

Invited Talk

- “Provable” Security against Differential and Linear Cryptanalysis 1
Kaisa Nyberg

Block Ciphers

- Improved Attacks on Full GOST 9
Itai Dinur, Orr Dunkelman, and Adi Shamir
- Zero Correlation Linear Cryptanalysis with Reduced Data
Complexity 29
Andrey Bogdanov and Meiqin Wang

Differential Cryptanalysis

- A Model for Structure Attacks, with Applications to PRESENT
and Serpent 49
Meiqin Wang, Yue Sun, Elmar Tischhauser, and Bart Preneel
- A Methodology for Differential-Linear Cryptanalysis and
Its Applications 69
Jiqiang Lu
- New Observations on Impossible Differential Cryptanalysis
of Reduced-Round Camellia 90
*Ya Liu, Leibo Li, Dawu Gu, Xiaoyun Wang, Zhiqiang Liu,
Jiazhe Chen, and Wei Li*

Hash Functions I

- Improved Rebound Attack on the Finalist Grøstl 110
Jérémy Jean, María Naya-Plasencia, and Thomas Peyrin
- (Pseudo) Preimage Attack on Round-Reduced Grøstl Hash Function
and Others 127
*Shuang Wu, Dengguo Feng, Wenling Wu, Jian Guo, Le Dong, and
Jian Zou*
- Practical Cryptanalysis of ARMADILLO2 146
María Naya-Plasencia and Thomas Peyrin

On the (In)Security of IDEA in Various Hashing Modes 163
*Lei Wei, Thomas Peyrin, Przemysław Sokołowski, San Ling,
 Josef Pieprzyk, and Huaxiong Wang*

Modes of Operation

The Security of Ciphertext Stealing 180
Phillip Rogaway, Mark Wooding, and Haibin Zhang

McOE: A Family of Almost Foolproof On-Line Authenticated
 Encryption Schemes 196
Ewan Fleischmann, Christian Forler, and Stefan Lucks

Cycling Attacks on GCM, GHASH and Other Polynomial MACs
 and Hashes 216
Markku-Juhani Olavi Saarinen

Hash Functions II

Collision Attacks on the Reduced Dual-Stream Hash Function
 RIPEMD-128 226
Florian Mendel, Tomislav Nad, and Martin Schl affer

Bicliques for Preimages: Attacks on Skein-512 and the SHA-2 Family . . . 244
Dmitry Khovratovich, Christian Rechberger, and Alexandra Savelieva

Converting Meet-In-The-Middle Preimage Attack into Pseudo Collision
 Attack: Application to SHA-2 264
Ji Li, Takanori Isobe, and Kyoji Shibutani

New Tools for Cryptanalysis

UNAF: A Special Set of Additive Differences with Application
 to the Differential Analysis of ARX 287
*Vesselin Velichkov, Nicky Mouha, Christophe De Canni ere, and
 Bart Preneel*

ElimLin Algorithm Revisited 306
*Nicolas T. Courtois, Pouyan Sepehrdad, Petr Su il, and
 Serge Vaudenay*

New Designs

Short-Output Universal Hash Functions and Their Use in Fast
 and Secure Data Authentication 326
Long Hoang Nguyen and A.W. Roscoe

Lapin: An Efficient Authentication Protocol Based on Ring-LPN	346
<i>Stefan Heyse, Eike Kiltz, Vadim Lyubashevsky, Christof Paar, and Krzysztof Pietrzak</i>	
Higher-Order Masking Schemes for S-Boxes	366
<i>Claude Carlet, Louis Goubin, Emmanuel Prouff, Michael Quisquater, and Matthieu Rivain</i>	
Recursive Diffusion Layers for Block Ciphers and Hash Functions	385
<i>Mahdi Sajadieh, Mohammad Dakhilalian, Hamid Mala, and Pouyan Sepehrdad</i>	
Keccak	
Unaligned Rebound Attack: Application to Keccak	402
<i>Alexandre Duc, Jian Guo, Thomas Peyrin, and Lei Wei</i>	
Differential Propagation Analysis of Keccak	422
<i>Joan Daemen and Gilles Van Assche</i>	
New Attacks on Keccak-224 and Keccak-256	442
<i>Itai Dinur, Orr Dunkelman, and Adi Shamir</i>	
Author Index	463

“Provable” Security against Differential and Linear Cryptanalysis

Kaisa Nyberg

Aalto University School of Science and Nokia, Finland
kaisa.nyberg@aalto.fi

Abstract. In this invited talk, a brief survey on the developments of countermeasures against differential and linear cryptanalysis methods is presented.

1 Nonlinearity of S-Boxes

Throughout the eighties the unpublished design criteria of the DES had inspired various authors to invent formal nonlinearity criteria for S-boxes such as the *strict avalanche criterion* [30] and the *propagation criterion* [27]. At the same time, correlation attacks on combination generators inspired definitions of *correlation immunity* [29] and *perfect nonlinearity* [21] of Boolean functions. W. Meier and O. Staffelbach realized that perfect nonlinear Boolean functions had been invented before under the name *bent functions* [28,12]. Then the discovery of the differential cryptanalysis method [4] led to the notion of *perfect nonlinear S-boxes* [22], with the property that for any non-zero input difference the output differences are uniformly distributed. In particular, the output difference zero would occur with the same probability as the non-zero output differences and would significantly improve the probability of the two-round iterative characteristic for a Feistel cipher as pointed out to the author by E. Biham at Eurocrypt 1991. It also means that perfect nonlinear S-boxes cannot be bijective, even worse, the number of input bits must be at least twice the number of output bits [22].

It was clear that the requirement of perfect nonlinearity must be relaxed. But it was not sufficient to take care that the output bits were highly nonlinear Boolean functions as in [26], but also all non-zero linear combinations of the output bits should be highly nonlinear as noted in [23], where the definition of nonlinearity of a vector Boolean function was formulated. The importance of nonlinearity as a cryptographic criterion was highlighted even more as the linear cryptanalysis method was presented by M. Matsui in 1993 [20]. The relationship between nonlinearity (resistance against linear cryptanalysis) and differential uniformity (resistance against differential cryptanalysis) was established in [8]. Since then H. Dobbertin and C. Carlet followed by many other authors have contributed with combinatorial designs and constructions that are almost perfect nonlinear (APN) or satisfy other nonlinearity criteria of S-boxes.

2 CRADIC

We observed that if the differential probabilities of a round function of a Feistel cipher are bounded from above, then also the differential probabilities over four rounds of the cipher are bounded by a significantly smaller bound. There is a penalty of allowing zero output difference as noted by E. Biham, but it takes only one more round to achieve the same security level. In [25] we formulated and proved the following theorem.

Theorem 1. (KN Theorem) *It is assumed that in a DES-like cipher with $f : \mathbb{F}_2^m \rightarrow \mathbb{F}_2^n$ the round keys are independent and uniformly random. Then the probability of an s -round differential, $s \geq 4$, is less than or equal to $2p_{max}^2$.*

Here

$$\begin{aligned} p_{max} &= \max_{\beta} \max_{\alpha_R \neq 0} \Pr[\alpha_L + f(E(X + \alpha_R)) + K) + f(E(X) + K) = \beta_R] \\ &\leq p_f = \max_b \max_{a \neq 0} \Pr[f(Y + a) + f(Y) = b] \end{aligned}$$

If f bijective, then the claim of the KN Theorem holds for $s \geq 3$, in which case the multiplier 2 can be removed [1].

The high nonlinearity of the Cube function $f(x) = x^3$ in \mathbb{F}_{2^n} had been observed already in [26]. It is bijective for odd n only, so we made one-bit adjustments to it, so that it was possible to fit it into a balanced $2(n-1)$ -bit Feistel cipher as a round function. We called this cipher CRADIC, as Cipher Resistant Against Differential Cryptanalysis, but in public it became known as KN Cipher. The cipher was later broken using algebraic cryptanalysis making use of the low degree of the Cube monomial.

Since then, designers of block ciphers continue using small nonlinear S-boxes in the spirit of C. Shannon. Would it be possible to use a monolithic algebraic construction? Recently, the Discrete Logarithm function was proved to achieve optimum algebraic immunity [7]. Let α generator of the multiplicative group $\mathbb{F}_{2^n}^*$ and set

$$f(x) = \begin{cases} \log_{\alpha}(x), & \text{for } x \neq 0 \\ (1, 1, \dots, 1,) & \text{for } x = 0. \end{cases}$$

Then f gives an n -bit S-box. Previously, it is known that any single output bit of f exhibits asymptotically low correlation with linear functions [6]. The correlations are bounded from above by

$$\mathcal{O}(n2^{-n/2}).$$

But no useful general upper bound is known to the linearity of combinations of output bits. The known bounds increase exponentially as the length of the linear mask grows [7,14]. Later we managed to establish a smaller bound where the increase is exponential with respect to the number of output bits involved, that is, the Hamming weight of the mask. In experiments, however, it seems that the linearity does not grow exponentially but essentially slower. It remains an open question, whether CRADIC would be secure if the Cube function were replaced by the Discrete Logarithm function.

3 Linear Hulls

The essential notion in the KN Theorem is *differential* first introduced in [18]. The approach taken in this work was to model an iterated block cipher as a stochastic process and assume that the rounds are independent. This can be achieved for a key-alternating cipher by selecting the round keys to be statistically independent and then taking the average over all keys. Under the *hypothesis of stochastic equivalence* it is then possible to draw conclusions about the behaviour of the cipher for a fixed unknown key. We adopted the same stochastic model and introduced in [24] the concept of linear hull and proved the following result for the expected squared correlation.

Theorem 2. (Linear Hull Theorem) *Let X , K and Y be random variables in \mathbb{F}_2^m , \mathbb{F}_2^ℓ , and \mathbb{F}_2^n , resp. where $Y = F(X, K)$ and X and K are independent. If K is uniformly distributed, then for all $a \in \mathbb{F}_2^m$ and $b \in \mathbb{F}_2^n$,*

$$\text{Exp}_K \text{corr}(a \cdot X + b \cdot Y)^2 = \sum_{c \in \mathbb{F}_2^\ell} \text{corr}(a \cdot X + b \cdot Y + c \cdot K)^2.$$

Here, for random variable Z in \mathcal{Z} (binary strings) we defined

$$\text{corr}(u \cdot Z) = \frac{1}{|\mathcal{Z}|} \sum_{z \in \mathcal{Z}} \Pr[z] (-1)^{u \cdot z}. \quad (1)$$

Then the linear hull (originally called as approximate linear hull) was defined as the set of all linear approximations

$$ALH(a, b) = \{a \cdot X + b \cdot Y + c \cdot K \mid c \in \mathbb{F}_2^\ell\}$$

of plaintext, ciphertext and key, with fixed input and output masks a and b , but letting the key mask vary. Thus taking squares of the correlations and summing over c gives the average correlation over the cipher with plaintext mask a and ciphertext mask b .

J. Daemen abandoned the Markov cipher model and took the fixed key approach [11]. He investigated correlations of linear approximations over a key alternating block cipher E , with round functions $x \mapsto f_i(x + K_i)$, and fixed set of round keys K_0, \dots, K_r . Given vector Boolean function: $f : \mathbb{F}_2^n \mapsto \mathbb{F}_2^m$ with $f = (f_1, \dots, f_m)$, where $b \cdot f$ are Boolean functions, for all $b \in \mathbb{F}_2^m$, the correlation between $b \cdot f(x)$ and $a \cdot x$ is defined by

$$c_f(a, b) = \frac{1}{2^n} (\#\{x \in \mathbb{F}_2^n \mid b \cdot f(x) = a \cdot x\} - \#\{x \in \mathbb{F}_2^n \mid b \cdot f(x) \neq a \cdot x\})$$

Then the correlation of a composed function computed as the matrix product is

$$c_{f \circ g}(a, b) = \sum_u c_g(a, u) c_f(u, b),$$

from where we obtain

$$c_E(u_0, u_r) = \sum_{u_1, \dots, u_{r-1}} (-1)^{u_0 \cdot K_0 + \dots + u_r \cdot K_r} \prod_{i=1}^r c_{f_i}(u_{i-1}, u_i),$$

where u_0 and u_r are the linear masks of data after 0 and r rounds of encryption, respectively. This result holds for all fixed keys. By taking the squares and averaging over uniformly distributed and independent keys we get as a corollary

$$\text{Average}_{K_0, \dots, K_r} c_E(u_0, u_r)^2 = \sum_{u_1, \dots, u_{r-1}} \prod_{i=1}^r c_{f_i}(u_{i-1}, u_i)^2.$$

This result was given in [24] for the special case of DES. Related to this, let us also observe that the correlation of a single trail of a linear hull, taken over plaintext, ciphertext and key, gives another presentation of the piling-up lemma

$$\text{corr}(a \cdot X + b \cdot Y + c \cdot K) = \prod_{i=1}^r c_{f_i}(u_{i-1}, u_i),$$

where $a = u_0$, $b = u_r$, and c is in unique correspondence with the trail masks u_1, \dots, u_{r-1} .

Finally let us make an observation of the effect of key scheduling, which should be designed in such a way that the magnitudes of the correlations

$$c_E(u_0, u_r) = \sum_{u_1, \dots, u_{r-1}} (-1)^{u_0 \cdot K_0 + \dots + u_r \cdot K_r} \prod_{i=1}^r c_{f_i}(u_{i-1}, u_i)$$

do not vary too much with the key. This can be achieved if all dominating trail correlations are of about equal magnitude and the map:

$$(u_1, \dots, u_{r-1}) \mapsto \text{sign} \left(\prod_{i=1}^r c_{f_i}(u_{i-1}, u_i) \right)$$

is highly nonlinear. Then the correlations $|c_E(u_0, u_r)|$ are bounded by the small linearity bound. Known examples of mappings with highly nonlinear correlation sign functions are bent functions and the Cube function. For bent functions the sign function is also bent. For the Cube function, correlations are zero in a half space while restricted in the other half space the sign function is bent.

4 Provable Security in Practice

It would be easier to achieve security guarantees against differential and linear attacks for round functions composed of a highly nonlinear monolithic design. In case of substitution permutation networks and similar designs such as AES, cryptographers must work harder. The basic approach is to design the diffusion layer in such a way that the minimum number of active S-boxes involved in the attack is large enough to make the linear trail correlations and differential

characteristic probabilities sufficiently small. To achieve this goal, the designers of the AES used MDS matrices for creating larger S-boxes and the Wide-Trail Strategy for ensuring diffusion of trails over the entire width of the cipher [10]. Then obtaining any useful upper bounds to linear correlations and differential probabilities becomes hard. The best known upperbounds for 4 and more rounds are due to L. Keliher [16].

The block cipher PRESENT makes use of bit permutations between rounds for optimal diffusion [5]. Its hardware optimized S-box exhibits, however, strong linear correlations for single-bit masks. Consequently, fairly accurate estimates of correlations can be obtained using single-bit linear approximation trails. As demonstrated in [9], linear hull effect is significant and therefore linear attacks are more powerful than initially estimated by the designers. The other side of the coin is that now we have better estimates of resistance of PRESENT against linear attacks. Can the linear hull effect for PRESENT be computed with sufficient accuracy using single-bit trails only is an interesting question.

5 Linear Approximations and Distributions

The correspondence between correlations of linear projections and probability distributions has been well-known for cryptographers since at least [2] but not exploited in cryptanalysis until in the multidimensional linear cryptanalysis [15]. It allows to use a number of linear approximations simultaneously. More generally, let Z be a vector of (binary) random variables over domain \mathcal{Z} . By applying the inverse Walsh-Hadamard transform to (1) we get

$$p_z = \Pr[z] = \sum_{u \in \mathcal{Z}} \text{corr}(u \cdot Z)(-1)^{u \cdot z}.$$

In cryptanalysis, Z is a random variable, which can be sampled from cipher data, such as multidimensional linear approximation, difference, or ciphertext from chosen biased plaintext, anything expected to have non-random behaviour. In this sense, linear approximations, that is, linear projection $z \mapsto u \cdot z$ gives a universal tool for analyzing probability distributions. For example, G. Leander used it to prove that the statistical saturation attack averaged over the fixations and the multidimensional linear cryptanalysis attack are essentially the same [19].

This approach is not restricted to binary variables but can be extended to any finite group. For example, projections $x \mapsto ux \bmod p$, for p prime, have been used in cryptanalysis of block ciphers with non-binary diffusion layer [3]. This leads to the following generalized notion of correlation

$$c_f(u, w) = \frac{1}{q} \sum_{x \in \mathbb{Z}_q} e^{\frac{2\pi i}{p} w f(x)} e^{-\frac{2\pi i}{q} ux}$$

for a function $f : \mathbb{Z}_q \rightarrow \mathbb{Z}_p$ and positive integers p and q . The generalized bent functions achieve the smallest linearity with respect to generalized correlation [17]. The Discrete Logarithm function for integers is another example

with known asymptotic upper bound of linearity [13]. This upperbound is of the same magnitude than the bound conjectured to the binary Discrete Logarithm function.

Given such Z related to a cipher, how many samples of Z is needed to distinguish it from a true random variable? If the distribution of Z is close to uniform, then the answer can be given in terms of the capacity of the distribution Z defined as follows:

$$C(Z) = M \sum_{z \in \mathcal{Z}} \left(p_z - \frac{1}{M}\right)^2,$$

where $M = |\mathcal{Z}|$. Using the relationship between the distribution and correlations we obtain

$$C(Z) = \sum_{u \neq 0} |\text{corr}(u \cdot Z)|^2.$$

Let us summarize the known upper bounds of data complexities for two commonly used distinguishers.

The strongest distinguisher based on the log-likelihood ratio (LLR) requires good knowledge of the probability distribution of Z . If it is available, then the data requirement of the LLR distinguisher can be given as:

$$N_{\text{LLR}} = \frac{\lambda}{C(Z)},$$

where the constant λ depends only on the success probability.

In cryptanalysis, the variable Z and its probability distribution typically depend on the unknown key. While the χ^2 distinguisher is less optimal than the LLR, it can be used also in this case, as it does not require knowledge of the distribution of Z . Its data requirement is

$$N_{\chi^2} = \frac{\lambda' \sqrt{M}}{C(Z)}, \quad \text{where}$$

$$\lambda' \approx (\sqrt{2} + 2)\Phi^{-1}(P_S) \approx \lambda.$$

Cryptanalysts aim at minimizing the data complexity. To be able to use the LLR bound, they must make convincing arguments that LLR works. Else they are left with the higher value given by the χ^2 complexity bound. Cryptographers want to work in the opposite direction and claim as high values as possible for the data complexity. In general, provable security may be difficult to achieve given only such upper bounds of average data complexities. It takes practical experiments and other evidence to see what the actual distinguishing data complexities are and how much they vary with the keys.

Acknowledgement. Thanks to Céline and Risto for their help in the final editing of the paper.

References

1. Aoki, K.: On Maximum Non-averaged Differential Probability. In: Tavares, S., Meijer, H. (eds.) SAC 1998. LNCS, vol. 1556, pp. 118–130. Springer, Heidelberg (1999)
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. Baignères, T., Stern, J., Vaudenay, S.: Linear Cryptanalysis of Non Binary Ciphers. In: Adams, C., Miri, A., Wiener, M. (eds.) SAC 2007. LNCS, vol. 4876, pp. 184–211. Springer, Heidelberg (2007)
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. Bogdanov, 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)
6. Brandstätter, N., Lange, T., Winterhof, A.: On the Non-linearity and Sparsity of Boolean Functions Related to the Discrete Logarithm in Finite Fields of Characteristic Two. In: Ytrehus, Ø. (ed.) WCC 2005. LNCS, vol. 3969, pp. 135–143. Springer, Heidelberg (2006)
7. Carlet, C., Feng, K.: An Infinite Class of Balanced Vectorial Boolean Functions with Optimum Algebraic Immunity and Good Nonlinearity. In: Chee, Y.M., Li, C., Ling, S., Wang, H., Xing, C. (eds.) IWCC 2009. LNCS, vol. 5557, pp. 1–11. Springer, Heidelberg (2009)
8. Chabaud, F., Vaudenay, S.: Links between Differential and Linear Cryptanalysis. In: De Santis, A. (ed.) EUROCRYPT 1994. LNCS, vol. 950, pp. 356–365. Springer, Heidelberg (1995)
9. Cho, J.Y.: Linear Cryptanalysis of Reduced-Round PRESENT. In: Pieprzyk, J. (ed.) CT-RSA 2010. LNCS, vol. 5985, pp. 302–317. Springer, Heidelberg (2010)
10. Daemen, J., Rijmen, V.: The Design of Rijndael – AES, the Advanced Encryption Standard. Springer (2002)
11. Daemen, J., Govaerts, R., Vandewalle, J.: Correlation Matrices. In: Preneel, B. (ed.) FSE 1994. LNCS, vol. 1008, pp. 275–285. Springer, Heidelberg (1995)
12. Dillon, J.F.: Elementary Hadamard difference sets. In: Proceedings of the Sixth Southeastern Conference on Combinatorics, Graph Theory and Computing, Boca Raton, Florida. Congressus Numerantium, vol. XIV, pp. 237–249. Utilitas Math., Winnipeg, Manitoba (1975)
13. Hakala, R.M.: An upper bound for the linearity of Exponential Welch Costas functions. Finite Fields and Their Applications (to appear, 2012), <http://dx.doi.org/10.1016/j.ffa.05.001>
14. Hakala, R.M., Nyberg, K.: On the Nonlinearity of Discrete Logarithm in \mathbb{F}_{2^n} . In: Carlet, C., Pott, A. (eds.) SETA 2010. LNCS, vol. 6338, pp. 333–345. Springer, Heidelberg (2010)
15. Hermelin, M., Cho, J.Y., Nyberg, K.: Multidimensional Linear Cryptanalysis of Reduced Round Serpent. In: Mu, Y., Susilo, W., Seberry, J. (eds.) ACISP 2008. LNCS, vol. 5107, pp. 203–215. Springer, Heidelberg (2008)
16. Keliher, L.: Refined Analysis of Bounds Related to Linear and Differential Cryptanalysis for the AES. In: Dobbertin, H., Rijmen, V., Sowa, A. (eds.) AES 2005. LNCS, vol. 3373, pp. 42–57. Springer, Heidelberg (2005)

17. Kumar, P.V., Scholtz, R.A., Welch, L.R.: Generalized bent functions and their properties. *J. Combin. Theory Ser. A* 40(1), 90–107 (1985)
18. Preneel, B., Govaerts, R., Vandewalle, J.: Boolean Functions Satisfying Higher Order Propagation Criteria. In: Davies, D.W. (ed.) EUROCRYPT 1991. LNCS, vol. 547, pp. 141–152. Springer, Heidelberg (1991)
19. Leander, G.: On Linear Hulls, Statistical Saturation Attacks, PRESENT and a Cryptanalysis of PUFFIN. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 303–322. Springer, Heidelberg (2011)
20. Matsui, M.: Linear Cryptanalysis Method for DES Cipher. In: Helleseth, T. (ed.) EUROCRYPT 1993. LNCS, vol. 765, pp. 386–397. Springer, Heidelberg (1994)
21. Meier, W., Staffelbach, O.: Nonlinearity Criteria for Cryptographic Functions. In: Quisquater, J.-J., Vandewalle, J. (eds.) EUROCRYPT 1989. LNCS, vol. 434, pp. 549–562. Springer, Heidelberg (1990)
22. Nyberg, K.: Perfect Nonlinear S-Boxes. In: Davies, D.W. (ed.) EUROCRYPT 1991. LNCS, vol. 547, pp. 378–386. Springer, Heidelberg (1991)
23. Nyberg, K.: On the Construction of Highly Nonlinear Permutations. In: Rueppel, R.A. (ed.) EUROCRYPT 1992. LNCS, vol. 658, pp. 92–98. Springer, Heidelberg (1993)
24. Nyberg, K.: Linear Approximation of Block Ciphers. In: De Santis, A. (ed.) EUROCRYPT 1994. LNCS, vol. 950, pp. 439–444. Springer, Heidelberg (1995)
25. Nyberg, K., Knudsen, L.R.: Provable security against a differential attack. *Journal of Cryptology* 8(1), 27–37 (1995)
26. Pieprzyk, J.: On bent permutations. Tech. rep., The University of South Wales, Department of Computer Science. Presented at the International Conference on Finite Fields, Coding Theory and Advances in Communications and Computing, Las Vegas (1991)
27. Preneel, B., Van Leekwijck, W., Van Linden, L., Govaerts, R., Vandewalle, J.: Propagation Characteristics of Boolean Functions. In: Damgård, I.B. (ed.) EUROCRYPT 1990. LNCS, vol. 473, pp. 161–173. Springer, Heidelberg (1991)
28. Rothaus, O.S.: On “bent” functions. *J. Combinatorial Theory Ser. A*(20), 300–305 (1976)
29. Siegenthaler, T.: Correlation-immunity of nonlinear combining functions for cryptographic applications. *IEEE Transactions on Information Theory* 30(5), 776–780 (1984)
30. Webster, A.F., Tavares, S.: On the Design of S-boxes. In: Williams, H.C. (ed.) CRYPTO 1985. LNCS, vol. 218, pp. 523–534. Springer, Heidelberg (1986)

Improved Attacks on Full GOST

Itai Dinur¹, Orr Dunkelman^{1,2}, and Adi Shamir¹

¹ Computer Science Department, The Weizmann Institute, Rehovot, Israel

² Computer Science Department, University of Haifa, Israel

Abstract. GOST is a well known block cipher which was developed in the Soviet Union during the 1970's as an alternative to the US-developed DES. In spite of considerable cryptanalytic effort, until very recently there were no published single key attacks against its full 32-round version which were faster than the 2^{256} time complexity of exhaustive search. In February 2011, Isobe used the previously discovered reflection property in order to develop the first such attack, which requires 2^{32} data, 2^{64} memory and 2^{224} time. In this paper we introduce a new fixed point property and a better way to attack 8-round GOST in order to find improved attacks on full GOST: Given 2^{32} data we can reduce the memory complexity from an impractical 2^{64} to a practical 2^{36} without changing the 2^{224} time complexity, and given 2^{64} data we can simultaneously reduce the time complexity to 2^{192} and the memory complexity to 2^{36} .

Keywords: Block cipher, cryptanalysis, GOST, reflection property, fixed point property, 2D meet in the middle attack.

1 Introduction

During the 1970's, the US decided to publicly develop the Data Encryption Standard (DES), which was the first standardized block cipher intended for civilian applications. At roughly the same time, the Soviet Union decided to secretly develop GOST [14], which was supposed to be used in civilian applications as well but in a more controlled way. The general design of GOST was finally published in 1994, but even today some of the crucial elements (e.g., the choice of Sboxes) do not appear in the public description, and a different choice can be made for each application.

GOST is a Feistel structure over 64-bit blocks. The round function consists of adding (modulo 2^{32}) a 32-bit round key to the right half of the block, and then applying the function f described in Figure 1. This function has an Sbox layer consisting of eight different 4×4 Sboxes, followed by a rotation of the 32-bit result by 11 bits to the left using the little-endian format (i.e. the LSB of the 32-bit word enters the rightmost entry of the first Sbox).

The full GOST has 32 rounds, and its key schedule is extremely simple: the 256-bit key is divided into eight 32-bit words (K_1, K_2, \dots, K_8). Each round of GOST uses one of these words as a round key in the following order: in the first 24 rounds, the keys are used in their cyclic order (i.e. K_1 in rounds 1,9,17, K_2




Fig. 1. One round of GOST

in rounds 2,10,18, and so forth). In the final 8 rounds (25–32), the round keys are used in reverse order (K_8 in round 25, K_7 in round 26, and so forth).

A major difference between the design philosophies of DES and GOST was that the publicly available DES was intentionally chosen with marginal parameters (16 rounds, 56-bit keys), whereas the secretive GOST used larger parameters (32 rounds, 256-bit keys) which seemed to offer an extra margin of security. As a result, DES was broken theoretically (by using differential and linear techniques) and practically (by using special purpose hardware) about 20 years ago, whereas in the case of GOST, all the single key attacks [1,9,17] published before 2011 were only applicable to reduced-round versions of the cipher.¹

The first single key attack on the full 32-round version of GOST was published by Isobe at FSE’11 [8]. It exploited a surprising reflection property which was first pointed out by Kara [9] in 2008: Whenever the left and right halves of the state after 24 rounds are equal (which happens with probability 2^{-32}), the last 16 rounds become the identity mapping, and thus the effective number of rounds is reduced from 32 to 16. Isobe developed a new key-extraction algorithm for the remaining 16 rounds of GOST which required 2^{192} time and 2^{64} memory, and used it 2^{32} times for different plaintext/ciphertext pairs in order to get the full 256-bit key using a total of 2^{32} data, 2^{64} memory, and 2^{224} time. This is much faster than exhaustive search, but neither the time complexity nor the memory complexity are even close to being practical.

Shortly afterwards, Courtois [4] published on ePrint a new attack on the full GOST. It uses a very different algebraic approach, but had an inferior complexity of 2^{64} data, 2^{64} memory, and 2^{248} time. Later, Courtois and Misztal [5] described a differential attack which again used 2^{64} data and memory, but reduced the time complexity to 2^{226} .

In this paper we improve several aspects of these previously published attacks. We describe a new *fixed point property*, and show how to use either the previous reflection property or the new fixed point property in order to reduce the general cryptanalytic problem of attacking the full 32-round GOST into an attack on

¹ Attacks on full GOST in the stronger related-key model are known for about a decade, see [7,10,11,16,17].

8-round GOST with two known input-output pairs. We then develop a new way to extract all the 2^{128} possible values of the full 256-bit key given only two known 64-bit input-output pairs of 8-round GOST, which requires 2^{128} time and 2^{36} memory² (all the previously published attacks on 8-round GOST have an impractical memory complexity of at least 2^{64}). By combining these improved elements, we can get the best known attacks on GOST for the two previously considered data complexities of 2^{32} and 2^{64} .

Our new results on GOST (including the fixed point based attack) use well known and easy to analyze cryptanalytic techniques such as “Guess and Determine” and “meet-in-the-middle”. A month after this paper appeared on eprint [6] (and more than four months after its results were publicly disclosed in a public talk by Adi Shamir at MIT), Courtois posted to ePrint his independently discovered attacks [3], which use several different algebraic techniques. Some of his attacks are also based on the fixed point property, but all of them have higher claimed complexities: Given 2^{32} data, the best attack in [3] has a time complexity of 2^{224} and a memory complexity of 2^{128} , and given 2^{64} data, the best attack in [3] has a time complexity of 2^{216} and a negligible memory complexity. We include the results of [3] in Table 1 (which summarizes all the previously known single-key attacks on the full GOST, our new results, and Courtois’ subsequent results) for the sake of completeness.

An important observation about Isobe’s attack is that it uses in an essential way the assumption that the Sboxes are invertible. Since the GOST standard does not specify the Sboxes, and there is no need to make them invertible in a Feistel structure, Isobe’s attack might not be applicable to some valid incarnations of this standard. A similar problem occurs in most of Courtois’ attacks [3,4,5], as their complexities are only estimated for one particular choice of Sboxes described in [15] which is used in the Russian banking system, and it is possible that for other choices of Sboxes the complexities will be different. Our new attacks do not suffer from these limitations, since they can be applied with the same complexity to any given set of Sboxes.

2 Overview of Our New Attacks on the Full GOST

The 32 rounds of GOST can be described using only two closely related 8-round encryption functions. Let $G_{K_{i_1}, \dots, K_{i_j}}$ be j rounds of GOST under the subkeys K_{i_1}, \dots, K_{i_j} (where $i_1, \dots, i_j \in \{1, 2, \dots, 8\}$), and let (P_L, P_R) be a 64-bit plaintext, such its right half, P_R , enters the first round. Then $GOST_K(P_L, P_R) = G_{K_8, \dots, K_1}(G_{K_1, \dots, K_8}(G_{K_1, \dots, K_8}(G_{K_1, \dots, K_8}(P_L, P_R))))$.

Our new attacks on the full GOST exploit its high degree of self-similarity using a general framework which is shared by other attacks: the algorithm of

² We can reduce the memory complexity by an additional factor of 2^{17} (to 2^{19}) if we are willing to increase the time by a factor of 2^{12} (to 2^{140}). This may seem like an unattractive tradeoff since the 2^{36} memory complexity is already practical, but one can argue that 2^{19} words fit into the cache whereas 2^{36} do not, which may result in a big performance penalty.

Table 1. Single-key Attacks on the Full GOST

Reference	Data (KP) ^{††}	Memory	Time	Self-Similarity Property	8-Round Attack	Sboxes
[8]	2^{32}	2^{64}	2^{224}	Reflection	-	Bijective
[4]	2^{64}	2^{64}	2^{248}	Other (unnamed)	Algebraic	Russian Banks [15]
[5]	2^{64}	2^{64}	2^{226}	None (differential attack)	-	Russian Banks [15]
[3] ^{†††}	2^{32}	2^{128}	2^{224}	Reflection	-	any
[3] ^{†††}	2^{64}	Negligible	2^{216}	fixed point	Algebraic	Russian Banks [15]
This paper	2^{64}	2^{36}	2^{192} [†]	fixed point	2DMITM	any
This paper	2^{64}	2^{19}	2^{204} [†]	fixed point	low-memory	any
This paper	2^{32}	2^{36}	2^{224} [†]	Reflection	2DMITM	any
This paper	2^{32}	2^{19}	2^{236} [†]	Reflection	low-memory	any

[†] The time complexity can be slightly reduced by exploiting GOST's complementation properties (as described in the full version of the paper [6])

^{††} Known plaintext

^{†††} Published on ePrint after the original version of this paper [6].

each attack consists of an outer loop which iterates over the given 32-round plaintext-ciphertext pairs, and uses each one of them to obtain suggestions for two input-output pairs for G_{K_1, \dots, K_s} . For each suggestion of the 8-round input-output pairs, we apply an 8-round attack which gives suggestions for the 256-bit GOST key. We then verify the key suggestions by using some of the other plaintext-ciphertext pairs. The self-similarity properties of GOST ensure that the 8-round attack needs to be applied a relatively small number of times, leading to attacks which are much faster than exhaustive search.

We describe several attacks on the full GOST which belong to this common framework but differ according to the property and the type of 8-round attack we use. The two self-similarity properties are:

1. The *reflection property* which was first described in [9], where it was used to attack 30 rounds of GOST (and 2^{224} weak keys of the full GOST). This property was later exploited in [8] to attack the full GOST for all keys. We describe this property again in Section 3.1 for the sake of completeness.
2. A new *fixed point property* which is described in Section 3.2.

The two properties differ according to the amount of data required to satisfy them, and thus offer different points on a time/data tradeoff curve.

Given two 8-round input-output pairs, we describe in this paper several possible attacks of increasing sophistication:

1. A very basic meet-in-the-middle (MITM) attack [2], which is described in Section 4.1.

2. An improved MITM attack, described in Section 4.2, which uses the idea of equivalent keys (first described by Isobe in [8]).
3. A low-memory attack, described in Section 5, which requires 2^{19} memory and 2^{140} time.
4. A new *2-dimensional meet-in-the-middle* (2DMITM) attack, described in Section 6, which requires 2^{36} memory and 2^{128} time.

In order to attack the full GOST, we first select one of the two self-similarity properties to use in the outer loop of the attack according to the number of plaintext-ciphertext pairs available: in case we have 2^{64} pairs available, we select the fixed point property, and if we only have 2^{32} pairs, we select the reflection property. We then select one of last two 8-round attacks according to the amount of available memory: in case we have 2^{36} memory available, we select the 2DMITM attack, and if we only have 2^{19} memory, we select the low-memory attack. The outcome of this selection is an attack algorithm of the form:

1. For each plaintext-ciphertext pair (P, C) :
 - (a) Assuming that (P, C) satisfies the conditions of the self-similarity property, derive suggestions for two 8-round input-output pairs (I, O) and (I^*, O^*) .
 - (b) For each suggestion for (I, O) and (I^*, O^*) :
 - i. Execute the 8-round attack on (I, O) and (I^*, O^*) in order to derive suggestions for the key, and test each suggestion by performing trial encryptions on the remaining plaintext-ciphertext pairs.

The total time complexity of our attacks is calculated by multiplying the complexity of the 8-round attack by the expected number of times it needs to be applied according to the self-similarity property: An arbitrary (P, C) pair satisfies the fixed point property with probability of about 2^{-64} . Thus, it requires about 2^{64} known (P, C) pairs to succeed with high probability, and since we do not know in advance which pair satisfies the property, we need to repeat step 1 of the attack 2^{64} times. For each (P, C) pair, the fixed point property immediately suggests two 8-round input-output pairs (which are correct if the pair indeed satisfies the property). Hence, we need to perform step 1.(b) of the attack only once per (P, C) pair. In total, we need to execute the 8-round attack about 2^{64} times. On the other hand, an arbitrary (P, C) pair satisfies the reflection property with a much higher probability of about 2^{-32} . Thus, it requires about 2^{32} known (P, C) pairs, and we need to repeat the attack only 2^{32} times. However, for each (P, C) pair, the reflection property suggests a large number of 2^{64} values for (I, O) and (I^*, O^*) (out of which only one is correct if the pair indeed satisfies the property). Hence, we need to perform step 1.(b) of the attack 2^{64} times per (P, C) pair. In total, we need to execute the 8-round attack about $2^{32+64} = 2^{96}$ times.

Altogether, we obtain four new attacks on the full GOST. In three out of the four cases, we obtain better combinations of complexities than in all the previously published attacks. In the remaining case, we use the reflection property

and the low-memory 8-round attack to significantly reduce the memory requirements of Isobe’s attack [8], at the expense of a small time complexity penalty. We note that the computation required by each one of our attacks can be easily parallelized, and thus using x CPUs reduces the expected running time of the attack by a factor of x .

As described in the full version of this paper [6], the time complexity of all these attacks can be slightly reduced by exploiting GOST’s complementation properties. However, in some of these improved attacks we have to use chosen rather than known plaintexts, which reduces their attractiveness.

3 Obtaining Two 8-Round Input-Output Pairs for GOST

In this section, we describe the two self-similarity properties of GOST which we exploit in order to obtain two 8-round input-output pairs.

3.1 The Reflection Property [8,9]

Assume that the encryption of a plaintext P after 24 rounds of GOST results in a 64-bit value Y , such that the 32-bit right and left halves of Y are equal (i.e. $Y_R = Y_L$). Thus, exchanging the two halves of Y at the end of round 24 does not change the intermediate encryption value. In rounds 25–32, the round keys K_1 – K_8 are applied in the reverse order, and Y undergoes the same operations as in rounds 17–24, but in the reverse order. As a result, the encryption of P after 32 rounds, which is the ciphertext C , is equal to its encryption after 16 rounds (see Figure 2). By guessing the state of the encryption of P after 8 rounds, denoted by the 64-bit value X , we obtain two 8-round input-output pairs (P, X) and (X, C) . For an arbitrary key, the probability that a random plaintext gives such a symmetric value Y after 24 rounds is 2^{-32} , implying that we have to try about 2^{32} known plaintexts (in addition to guessing X) in order to obtain the two pairs. Note that the reflection property actually gives us another “half pair” (\widehat{C}, Y) , where the 64-bit word \widehat{C} is obtained from C by exchanging the right and left 32-bit halves of C , and the 32-bit right and left halves of Y are equal.³ However, it is not clear how to exploit this additional knowledge in order to significantly improve the running time of our attacks on the full GOST which are based on the reflection property.

³ In our attacks, we use 8-round input-output pairs whose encryption starts with K_1 and thus need to apply the Feistel structure in the reverse order (starting from round 32) for input-output pairs obtained for rounds 25–32. Since in Feistel structures the right and left halves of the block are exchanged at the end (rather than at the beginning) of the round function, we exchange the right and left sides of the input and the output of the input-output pairs obtained for rounds 25–32. We call (\widehat{C}, Y) a “half pair” since we have to guess only 32 additional bits in order to find it, once (P, C) is known.




Fig. 2. The Reflection Property of GOST

3.2 The Fixed Point Property

Assume that for a plaintext P , $G_{K_8, \dots, K_1}(P) = P$. Since rounds 9–16 and 17–24 are identical to rounds 1–8, we obtain P after 16 and 24 rounds as well. In rounds 25–32, the round keys K_1 – K_8 are applied in the reverse order, and we obtain some arbitrary ciphertext C (see Figure 3). The knowledge of P and C immediately gives us the 8-round input-output pairs (P, P) and (\hat{C}, \hat{P}) (in which the right and left 32-bit halves of P and C are exchanged).

For an arbitrary key, the probability that a random plaintext is a fixed point is about 2^{-64} , implying that we need about 2^{64} known plaintexts to have a single fixed point, from which we obtain the two input-output pairs needed in our attack. If we have only $c \cdot 2^{64}$ known plaintexts for some fraction c , we expect this fixed point to occur among the given plaintexts with probability c , and thus the time complexity, the data complexity, and the success probability are all reduced by the same linear factor c . Consequently, it makes sense to try the fixed point based attack even when we are given only a small fraction of the entire code book of GOST. Such a graceful degradation when we are given fewer plaintexts (which also occurs for the reflection property) should be contrasted with other attacks such as slide attacks, in which we have to wait for some random birthday phenomenon to occur among the given data points. Since the existence of birthdays has a much sharper threshold, the probability of finding an appropriate pair of points goes down quadratically rather than linearly in c , and thus they are much more likely to fail in such situations.

We note that our fixed point property is closely related to a previously published property which (in addition to the assumption the P is an 8-round fixed point) also assumes that the right and left halves of P are equal. Such a plaintext exists for an arbitrary key with probability 2^{-32} and thus was used in [9] to attack 2^{224} weak keys of the full GOST. The same property was also used later in [13] in cryptanalysis of the GOST hash function.




Fig. 3. The fixed point property of GOST

4 Simple Meet-In-The-Middle Attacks on 8 Rounds of GOST

Meet-in-the-middle (MITM) attacks can be efficiently applied to block ciphers in which some intermediate encryption variables (bits, or combinations of bits) depend only on a subset of key bits from the encryption side and on another subset of key bits from the decryption side: the attacker guesses the relevant key bits from the encryption and the decryption sides independently, and tries only keys in which the values suggested by the computed intermediate variables match. While the full 32-round GOST resists such attacks, 8-round GOST uses completely independent round keys. Thus, the full 64-bit value after 4 encryption rounds depends only on round keys K_1 – K_4 from the encryption side and on round keys K_5 – K_8 from the decryption side.

4.1 The Basic Meet-In-The-Middle Attack

We describe how to mount a simple meet-in-the-middle attack on 8 rounds of GOST given two 8-round input-output pairs and several additional 32-round plaintext-ciphertext pairs:

1. For each of the 2^{128} possible values of K_1 – K_4 , encrypt both inputs and obtain two 64-bit intermediate encryption values after 4 rounds of GOST (i.e., 2^{128} intermediate values of 128 bits each). Store the intermediate values in a list, sorted according to these 128 bits, along with the corresponding value of K_1 – K_4 .
2. For each of the 2^{128} possible values of K_5 – K_8 , decrypt both outputs, obtain two 64-bit intermediate values and search the sorted list for these two values.
3. For each match, obtain the corresponding value of K_1 – K_4 from the sorted list and derive a full 256-bit key by concatenating the value of value of K_1 – K_4 with the value of K_5 – K_8 of the previous step. Using the full key, perform a trial encryption of several plaintexts and return the full key, i.e., the one that remains after successfully testing the given 32-round pairs.

We expect to try about $2^{128+128-128} = 2^{128}$ full keys in step 3 of the attack, out of which only the correct key is expected to pass the exhaustive search of step 3. Including the 2^{128} 8-round encryptions which are performed in each of the first two steps of the attack, the total time complexity of the attack is slightly more than 2^{128} GOST encryptions. The memory complexity of the attack is about 2^{128} words of 256 bits.⁴

⁴ Note that it is possible obtain a time-memory tradeoff: we partition the 2^{128} possible values of K_1 – K_4 into 2^x sets of size 2^{128-x} (for $0 \leq x \leq 128$), and run the second and third steps of the attack independently for each set. Thus, the memory complexity decreases by a factor 2^x to 2^{128-x} , and the time complexity increases by a factor of 2^x to 2^{128+x} .

4.2 An Improved Meet-In-The-Middle Attack Using Equivalent Keys

In this section, we use a more general variant of Isobe’s equivalent keys idea [8] to significantly improve the memory complexity of the attack. Both our and Isobe’s MITM attacks are based on a 4-round attack that uses one 4-round input-output pair to find all the 2^{64} possible values of subkeys K_1 – K_4 that yield this pair. However, our MITM attack is more general since we can attack all possible incarnations of the GOST standard, whereas Isobe’s attack works only on those which use bijective Sboxes.⁵ An additional advantage of our MITM attack over Isobe’s one, is that our attack can use any two input-output pairs for 8-round GOST, regardless of how they are obtained. We can thus use the same algorithm to exploit both the reflection and the fixed point properties. On the other hand, Isobe’s attack is restricted to the case of a single input-output pair obtained for the first 16 rounds of GOST (by guessing the intermediate values obtained after 4 and 12 rounds) and thus can be combined with the reflection property, but cannot be directly applied to the two input-output pairs produced by the fixed point property.

We now describe Isobe’s 4-round attack procedure: Denote the 4-round input (divided into two 32-bit words) by (X_L, X_R) and the output by (Y_L, Y_R) . Denote the middle values (after the second round) by (Z_L, Z_R) (see Figure 4). Then:

$$Z_L = X_L \oplus f(X_R \boxplus K_1)$$

$$Z_R = Y_R \oplus f(Y_L \boxplus K_4)$$

$$Y_L \oplus Z_L = f(Z_R \boxplus K_3)$$

$$X_R \oplus Z_R = f(Z_L \boxplus K_2)$$

Isobe’s attack assumes bijective Sboxes (making f invertible), and finds the equivalent keys as follows:⁶ for each value of K_1, K_2 , compute Z_L from the first equation and Z_R from the fourth equation. From the second equation we have: $K_4 = f^{-1}(Z_R \oplus Y_R) \boxplus Y_L$ and from the third equation: $K_3 = f^{-1}(Z_L \oplus Y_L) \boxplus Y_R$.

Our 8-round attack is a variant of Isobe’s MITM attack, given two 8-round input-output pairs (I, O) and (I^*, O^*) :

1. For each possible value of the 64-bit word $Y = (Y_L, Y_R)$ obtained after 4 encryption rounds of I :
 - (a) Apply the 4-round attack on (I, Y) to obtain 2^{64} candidates for K_1 – K_4 .
 - (b) Partially encrypt I^* using the 2^{64} candidates and store $Y^* = (Y_L^*, Y_R^*)$ in a list with K_1 – K_4 .

⁵ The Feistel structure of GOST does not require bijective Sboxes and the published standard does not discuss this issue, but all the known choices of Sboxes happen to be bijective (perhaps due to the weakness of non-bijective Sboxes against differential cryptanalysis).

⁶ In case f is not bijective, then for a random (X_L, X_R) and (Y_L, Y_R) there exist an average of 2^{64} equivalent keys which can be found using a simple preprocessing MITM algorithm that requires about 2^{64} time and memory.

- (c) Apply the 4-round attack on (Y,O) to obtain 2^{64} candidates for K_5-K_8 .
- (d) Partially decrypt O^* using each one of the 2^{64} candidates and obtain $Y^* = (Y_L^*, Y_R^*)$.
- (e) Search the list obtained in step (b) for Y^* , and test the full 256-bit keys for which there is a match.

The expected time complexity of steps (a–d) is about 2^{64} (regardless of the algorithm that is used to find the equivalent keys). The time complexity of step (e) is also about 2^{64} since we expect to try about $2^{64+64-64} = 2^{64}$ full keys. Steps (a–e) are performed 2^{64} times, hence the total time complexity of the attack is about 2^{128} GOST encryptions, which is similar to the first attack. However, the memory complexity is significantly reduced from 2^{128} to slightly more than 2^{64} words of 64 bits.




Fig. 4. Four Rounds of GOST

5 A New Attack on 8 Rounds of GOST with Lower Memory Complexity

Simple meet-in-the-middle attacks, such as the ones described in Sections 4.1 and 4.2 are much faster than exhaustive search for the entire 256-bit key. However, they do not fully exploit the slow diffusion of the key bits in 4 rounds of GOST. As a result, these MITM attacks use a large amount of memory to store the many intermediate encryption values obtained for all the possible values of large sets of key bits. In this section, we describe an improved 8-round attack which exploits the slow diffusion properties of 4 rounds of GOST in order to reduce the memory complexity from the impractical value of 2^{64} to the very practical value of 2^{19} words of memory, with a very small time complexity penalty. The main idea of this attack is to guess the 4 round keys K_5-K_8 and apply an optimized “Guess and Determine” attack on the remaining 4 rounds using two input-output pairs.

In the 4-round attacks we have 128-bits of unknown key and 128 bits of input-output pairs. Thus, we expect that only one value for K_1 – K_4 exists (although there are likely to be input-output pairs for which the encryptions of the inputs does not match the outputs for any of the keys, and input-output pairs for which the encryptions of the inputs matches the outputs for several values of K_1 – K_4).

In the rest of this section we describe the algorithm for deriving the 32 bits of K_1 and the 32 bits of K_4 . Afterwards, deriving the values of K_2 and K_3 is immediate using the third and fourth equations of Section 4.2 (Z_L and Z_R are known from the first and second equations).

5.1 Overview of the “Guess and Determine” Attack on 4-Round GOST

Now that we deal with 4-round GOST, we apply a typical “Guess and Determine” attack which traverses a tree of partial guesses for the round keys K_1 and K_4 and intermediate encryption values. The tree is composed of layers of nodes ℓ_i for integral $0 \leq i \leq k$, where each layer contains nodes that specify the potential values (i.e. guesses) for a certain subset of key and intermediate encryption values. In each layer we expand each node by guessing the values of a small number of additional key bits and state bits that are needed to calculate some intermediate encryption bits, both from the encryption and the decryption sides. We then calculate the bits by evaluating the Feistel structure from both sides on a small number of bits, compare the values obtained, and discard guesses in which the values do not match (i.e., we discard child nodes that do not satisfy a predicate which checks the consistency of intermediate encryption values).

We traverse the partial guess tree starting from the root using DFS (which requires only a small amount of memory). In our attack, the nodes of the last layer of the tree ℓ_k contain guesses for the full key, which can be verified using trial encryptions.

The total number of operations performed during the traversal is proportional to the total number of nodes in the tree. However, the operations performed when expanding a single node work only on a few bits (rather than on full words). At the same time, when expanding a full path of nodes in the tree from the root to the last layer, we work on the full-size Feistel structure to obtain a guess for the full key. Hence, we estimate the time complexity of expanding a full path by a single Feistel structure evaluation on a full 64-bit input. Using this estimation, we can upper bound the time complexity of the tree traversal (in terms of Feistel structure evaluations) as the width of the tree, or the size of the layer which contains the highest number of nodes. Note that when counting the number of nodes in a layer for the time complexity analysis, we must also include nodes that were expanded and discarded since they do not satisfy the predicate of the previous layer.

5.2 Notations

Assume that we have two input-output pairs for 4 encryption rounds of GOST under the subkeys K_1, K_2, K_3, K_4 . Similarly to Section 4.2, denote the input, output and middle values (after using K_2) for the first pair by (X_L, X_R) , (Y_L, Y_R) and (Z_L, Z_R) , respectively. For the second pair, denote these values by (X_L^*, X_R^*) , (Y_L^*, Y_R^*) and (Z_L^*, Z_R^*) respectively.

Since our attack analyzes 4-bit words (which are outputs of single Sboxes), we introduce additional notations: Define the functions f^0, f^1, \dots, f^7 where each f^i takes a 4-bit word as an input, and outputs a 4-bit word by applying Sbox i to the input. Denote by W^i the i 'th bit of the 32-bit word W , and by $W^{i,j}$ the $(j - i + 1)$ -bit word composed of consecutive bits of W starting from bit i and ending at bit j . We treat W as a cyclic word, and thus $W^{24,3}$ contains 12 bits which are bits 24 to 31 and 0 to 3 of W .

5.3 An Attack on 4 Rounds of Simplified GOST

We start by describing an attack on 4 rounds of a simplified variant of GOST (which we call S-GOST), in which the round-key addition is replaced by XOR, and the 11-bit rotation is replaced by 12-bit rotation. The simplified variant is easier to analyze since it provides much slower diffusion of the key bits compared to full GOST: unlike addition, the XOR operation does not produce carries, and since 12 is a multiple of 4, rotating by 12 bits implies that the output of any Sbox effects the input of only a single Sbox in the next round.

We now describe the basic procedure performed by a node in layer 0 of our guess tree for S-GOST. The procedure requires the value of $K_1^{0,3}$ (whose value we guess before executing the procedure), and expands nodes in the next layer, which suggest a value for the additional 4 bits of $K_4^{20,23}$. The steps of this procedure can be easily verified using a variant of Figure 4 where the addition is replaced by XOR.

1. Given $K_1^{0,3}$ and $X_R^{0,3}$, compute $Z_L^{12,15} \equiv f^0(X_R^{0,3} \oplus K_1^{0,3})$ for both pairs (i.e., given $K_1^{0,3}$ and $X_R^{*0,3}$, compute $Z_L^{*12,15} \equiv f^0(X_R^{*0,3} \oplus K_1^{0,3})$).
2. Obtain $f^0(Z_R^{0,3} \oplus K_3^{0,3}) \equiv Z_L^{12,15} \oplus Y_L^{12,15}$ for both pairs. Then, invert⁷ f^0 to obtain $Z_R^{0,3} \oplus K_3^{0,3}$ and $Z_R^{*0,3} \oplus K_3^{0,3}$.
3. XOR the two expressions calculated in step 2, to eliminate $K_3^{0,3}$, and obtain the value of $Z_R^{0,3} \oplus Z_R^{*0,3}$.
4. XOR the 4-bit difference obtained in step 3 to the difference $Y_R^{0,3} \oplus Y_R^{*0,3}$ and obtain the value of $T = Z_R^{0,3} \oplus Y_R^{0,3} \oplus Z_R^{*0,3} \oplus Y_R^{*0,3} \equiv (f(Y_L \oplus K_4) \oplus f(Y_L^* \oplus K_4))^{0,3}$ (from the encryption side).
5. For each of the 2^4 possible values of $K_4^{20,23}$:
 - (a) Allocate a node in the next layer.

⁷ We expect one solution on average. However, in case the inversion has more than one solution, we need to try each one. In case the inversion has no solution, we can discard the node.

- (b) Evaluate the expression $f^5(Y_L^{20,23} \oplus K_4^{20,23}) \oplus f^5(Y_L^{*20,23} \oplus K_4^{20,23})$ from the decryption side by plugging the current value of $K_4^{20,23}$ into the expression. Discard nodes which do not agree with the value T .

Note that given $K_1^{0,3}$, we expect the procedure above to process a single child in the next layer: in step 5 we have a 4-bit condition on 4 bits of the key $K_4^{20,23}$, and thus we expect one node to satisfy the predicate. Moreover, step 5 can be optimized by using a small amount of precomputation and memory in order to calculate in advance the solutions to the 4-bit condition (as described in the full version of this paper [6]).

We now generalize the procedure above in order to derive more key bits in a similar way:

- Since encryption and decryption are completely symmetric (except the order of the subkeys), steps 1–5 can also be performed from the decryption side: in steps 1–5 we use the value of $K_1^{0,3}$ in order to obtain the value of $K_4^{20,23}$, and thus we define the symmetric steps 6–10 which use the value of $K_4^{20,23}$ in order to obtain the value of $K_1^{20+20,23+20}$, i.e. $K_1^{8,11}$.
- Given any integer $0 \leq i \leq 7$, we can rotate the indices of all the 32-bit words in steps 1–10 by $4i$ bits. Namely, given i , we define analogues steps 1–10 which use the value of $K_1^{4i,4i+3}$ to obtain the value of $K_4^{4i+20,4i+23}$ and $K_1^{4i+8,4i+11}$.

In order to derive the full 32-bit values of K_1 and K_4 , we define a tree which contains 9 layers $\ell_0, \ell_1, \dots, \ell_8$ (and an additional root node). The nodes of each layer are expanded using the generalized procedure which uses 4 bits of K_1 in order to derive 4 additional bits of K_1 and 4 additional bits of K_4 . Since the 10 steps of the procedure for expanding the nodes of layers 0–7 are basically the same, we call this procedure an *iteration*, and index it according to the value of i (which determines the 4-bit chunks that we work on).

5.4 Extending the Attack to 4 Rounds of the Real GOST

In order to extend the iteration procedure from S-GOST to full GOST, we need to make several adjustments. The most significant adjustments are given below:

- Since the round keys are added (and not XORed) to the state, we have to guess the carry bits into the LSBs of several addition operations of 4-bit words. For example, in the expression $f^5(Y_L^{20,23} \boxplus K_4^{20,23}) \oplus f^5(Y_L^{*20,23} \boxplus K_4^{20,23})$ evaluated in step 5, we have to guess two carry bits (one for $Y_L^{20,23}$ and one for $Y_L^{*20,23}$).
- GOST uses 11-bit rotation (instead of 12-bit rotation), and thus the 4-bit chunks that we work on in each iteration are not aligned. Consequently, we have to guess additional state bits in order to compare the evaluation of the 4-bit predicates from both sides. For example, since $20 + 11 = 31$, in step 5 of the iteration we actually calculate $(f(Y_L \oplus K_4) \oplus f(Y_L^* \oplus K_4))^{31,2}$ from the decryption side. Thus, we additionally guess bit 31 of this expression from the encryption side.

These adjustments create strong dependencies between iterations with consecutive indexes (i.e., i and $i + 1$), namely:

- The carry bits required by iteration $i + 1$ are known after iteration i . For example, iteration 1 requires the carry into bit 24 of the addition operation $Y_L \boxplus K_4$ (in order to calculate $f^6(Y_L^{24,27} \boxplus K_4^{24,27}) \oplus f^6(Y_L^{*24,27} \boxplus K_4^{24,27})$ in step 5). This bit can be calculated after step 5 of iteration 0, where the 4-bit value of $Y_L^{20,23} \boxplus K_4^{20,23}$ is calculated in order to evaluate the predicate.
- The state bits required by iteration $i + 1$ are known after iteration i . For example, iteration 1 requires calculation of bit 3 of the expression $f(Y_L \boxplus K_4) \oplus f(Y_L^* \boxplus K_4)$ from the encryption side. However, this bit is already guessed in step 4 of iteration 0.

This suggests that we perform the iterations in their natural order, namely assign layer ℓ_i iteration i for $0 \leq i \leq 7$. As a result, we need to guess carry and state bits only in the first iteration. Afterwards, the required carry and state bits for each iteration can be calculated by the knowledge from the previous one. On the other hand, we pay a (relatively small) penalty on key bit guesses since key bits required by iteration $i + 2$ are derived in iteration i (and not in iteration $i + 1$). Since iteration i requires key bits $K_1^{4i,4i+3}$, we need to guess 4 key bits in both iterations 0 and 1 ($K_1^{0,3}$ and $K_1^{4,7}$). For iterations $i \geq 2$, the required key bits are already derived in previous iterations (as shown in Table 2).

We note that since there is no carry into the LSBs of addition operations, starting the process with iteration 0 has the advantage that we do not need to guess the carries for all the addition operations (e.g., we do not need to guess the carry into the addition $f^0(X_R^{0,3} \boxplus K_1^{0,3})$ in step 1).

The full details and analysis of the “Guess and Determine” attack are given in the full version of this paper [6], most of which is not required in order to understand the rest of this paper. It shows that the expected number of nodes in the widest layer of the partial guess tree is 2^{14} , and it is obtained at iterations 1 to 5 (this was also verified using simulations performed on a PC). Basically, the number 2^{14} is obtained due to the 8 key-bit guesses ($K_1^{0,3}$ and $K_1^{4,7}$) and 6 additional carry and state bit guesses in iteration 0. This gives an expected time complexity of about 2^{14} 4-round Feistel structure evaluations for two input-output pairs, which is equivalent to about 2^{12} full GOST evaluations. Since we apply this 4-round attack 2^{128} times, the time complexity of the 8-round attack is about $2^{128+12} = 2^{140}$ GOST evaluations. In terms of memory, the attack has a completely practical complexity of 2^{25} bits, which is equivalent to 2^{19} 64-bit words.

6 A New 2-Dimensional Meet-In-The-Middle Attack on 8 Rounds of GOST

In this section, we present a new attack on 8 rounds of GOST given two input-output pairs, which combines the ideas of the “Guess and Determine” attack

Table 2. The key bits derived in each iteration

Iteration	0	1	2	3	4	5	6	7
K_1 bits derived	0–3	4–7	<u>8–11</u>	<u>12–15</u>	<u>16–19</u>	<u>20–23</u>	<u>24–27</u>	<u>28–31</u>
	8–11	12–15	16–19	20–23	24–27	28–31	<u>0–3</u>	<u>4–7</u>
K_4 bits derived	20–23	24–27	28–31	0–3	4–7	8–11	12–15	16–19

The key bits which are already known from previous iterations are underlined.

(which progresses horizontally across the state) and the MITM attack (which progresses vertically across the rounds). Unlike the attack of the previous section, we do not guess the last 4 round keys in advance. Instead, we divide the 8-round Feistel structure horizontally by splitting it into a *top part*, which uses round keys K_1 – K_4 , and a *bottom part*, which uses round keys K_5 – K_8 .

Our main observation is that due to the slow diffusion of the data bits into the state, we can run a substantial part of the “Guess and Determine” attack of Section 5 with very partial knowledge of Y and Y^* (obtained after 4 rounds of encryption). This allows us to split the “Guess and Determine” attack into two partial 4-round attacks which we run a relatively small number of times (once for each value of the bits of Y and Y^* that it requires). Our full 4-round attacks on the top and bottom parts combine the suggestions of the partial attacks in order to suggest values for the 4-round keys. Finally, we use an 8-round attack which joins the suggestions of the two partial attacks in order to obtain suggestions for the full 256-bit key.

Schematically, we split the top and bottom parts of the block cipher vertically into two (potentially overlapping) cells, such that on each cell we execute an independent partial attack to obtain suggestions for some subset of key bits. We then join all the suggestions to obtain suggestions for the full key using three MITM attacks. This can be visualized using a 2×2 matrix (as shown in Figure 5), where the horizontal line separates the four initial and final rounds of the 8-round block cipher, and the dashed vertical line separates the left and right cells in each one of the top and bottom parts.

After the MITM attacks on the top and bottom parts of the Feistel structure, we obtain 2^{128} suggestions for K_1 – K_4 and 2^{128} suggestions for K_5 – K_8 , each accompanied by corresponding 128-bit values of Y and Y^* . Note that so far we did not filter out any possible keys, and thus the final MITM attack, which compares the 128-bit values of Y and Y^* to obtain about 2^{128} suggestions for the full key, is essentially the basic MITM attack of Section 4.1, which would normally require 2^{128} memory.

To reduce the memory consumption, we guess many of the 128 bits of Y and Y^* in advance (in the outer loop of the 8-round attack). For each possible value of those bits, we execute the 2DMITM (2-dimensional MITM) attack described above, but obtain fewer suggestions for the key which we have to store. This increases the number of times that we execute the partial 4-round attacks and could potentially increase the overall time complexity of the full 8-round attack.




Fig. 5. The general framework of the 2-dimensional meet-in-the-middle attack

However, this is not the case, as the partial 4-round attacks are relatively efficient (the time complexity of each one is at most 2^{18}) and is executed only 2^{82} times. Thus, the partial 4-round attacks are not the bottleneck of the time complexity of the attack.⁸

6.1 Details of the 8-Round Attack

Formally, we define the following sets which contain bits of Y and Y^* :

- S_1 is the set of bits that we guess in the outer loop of the 8-round attack.
- S_2 is chosen such that $S_1 \cap S_2 = \emptyset$, and $S_1 \cup S_2$ is the minimal set that contains all the bits of Y and Y^* which are required by the partial 4-round attack on the left cell of the top part.
- S_3 is chosen such that $S_1 \cap S_3 = \emptyset$, and $S_1 \cup S_3$ is the minimal set of bits which are required by the partial 4-round attack on the right cell of the top part.

For the bottom MITM attack, we define S_4 and S_5 in a similar way to S_2 and S_3 , respectively. Note that since the 4-round attacks on both the top and bottom parts require all the 128 intermediate bits, $S_2 \cup S_3 = S_4 \cup S_5$.

The details of the 4-round attacks are given in the next section. We now refer to them as black boxes, and give the algorithm of the full 8-round attack:

1. For each value of the bits of the set S_1 :
 - (a) Perform the 4-round attack on the top part of the Feistel structure, and obtain a list with values of K_1-K_4 , sorted according to the value of the bits of $S_2 \cup S_3$.
 - (b) Perform the 4-round attack on the bottom part of the Feistel structure. For each value of $S_4 \cup S_5 = S_2 \cup S_3$ (given along with the value of K_5-K_8), search the list obtained in the previous step of matches. For each match test the full key K_1-K_8 .

⁸ Note again that we expect about 2^{128} keys to fulfill the filtering conditions of the two input-output pairs. Thus, the time required for the attack to list all of them cannot be reduced below 2^{128} (without exploiting additional filtering conditions).

6.2 Details of the 4-Round Attacks

We concentrate first on the top part of the 8-round Feistel structure: each one of the two partial 4-round attacks on the top part sequentially executes a subset of the iterations defined in Section 5, and is called an iteration *batch*. The first (left) iteration batch executes iterations 0–3, and the second (right) executes iterations 4–7.

After performing iteration batches 0–3 and 4–7 independently, we get suggestions for the values of some key bits, along with some carry and state bits. We then discard inconsistent suggestions by comparing the values of the common bits that are derived by batches. We partition these bits into three groups (which are fully specified in the full version of this paper [6]):

- G_1 contains 16 key bits which are derived by both of the left and right batches.
- G_2 contains 6 carry and state input bits that we guess in iteration 0. These bits are also contained in the set of output bits of iteration 7 (of the right batch), and can thus be used to discard inconsistent suggestions made by the two batches.
- G_3 contains 10 carry and state input bits that we guess in iteration 4. This bits are also contained in the set of iteration output bits of iteration 3 (of the left batch), and can thus be used to discard inconsistent suggestions made by the two batches.

Assume that the values of all the bits of S_1 are known. We now give the algorithm of the MITM attack performed on the top part of the 8-round Feistel structure:

1. For each value of the bits of S_2 , perform the left batch. Save all the nodes of the final layer in a list. These nodes contain the values 40 bits of K_1 and K_4 (including the values of the bits of G_1), and also the values of the bits of G_3 . In addition to the information obtained by each node, save the value of the initial guess of the bits of G_2 , and the value of the bits of S_2 per node. Sort the list according to the values of G_1, G_2 and G_3 .
2. For each value of the bits of S_3 , perform the right batch. For each node in the final layer obtain the value of the bits of G_1, G_2 and G_3 and search the list obtained in the first step for their value. For each match, save the value of the full K_1-K_4 in a sorted list according to the value of the bits of $S_2 \cup S_3$.

The iteration batches of the MITM attack on the bottom part of the Feistel structure are performed from the decryption side and are completely analogous to the iteration batches on the top part (i.e. in iteration 0, we start by guessing $K_8^{0,3}$, and derive $K_5^{20,23}$ and $K_8^{8,11}$). We also define analogous sets to G_1, G_2 and G_3 for the bottom part.

The specific choices of S_1-S_5 are given in the full version of this paper [6]. This choice of sets satisfies $|S_1| = 92$ and $|S_2| = |S_3| = |S_4| = |S_5| = 18$.

We now analyze the complexity of the MITM attack on the top part of the Feistel structure: as specified in the full version of this paper [6], when starting

the iteration batch from iteration 0, the expected maximal size of the tree is 2^{14} . It is obtained after iteration 1, and is maintained until the end of iteration 5 (even though we do not perform 5 consecutive iterations in this attack). The time complexity of the first step of the attack is thus about $2^{|S_2|+14} = 2^{14+18} = 2^{32}$, and this is also the size of the sorted list at the end of the first step. The maximal size of the tree of the iteration batch 4–7 is $2^{14+4} = 2^{18}$ (as described in the full version of this paper [6], we have to guess 4 more carry bits compared to iterations 0–3). Thus, the time complexity of expanding the tree in the second step is $2^{|S_3|+18} = 2^{36}$. The time and memory complexities of the remainder of step 2 (in which we match the batches) are $2^{|S_2|+|S_3|+14+18-(|G_1|+|G_2|+|G_3|)} = 2^{|S_2|+|S_3|+14+18-(16+6+10)} = 2^{|S_2|+|S_3|} = 2^{36}$. Note that it is not surprising that the time and memory complexities of the matching part of the attack reduce to $2^{|S_2|+|S_3|}$, since given the full 128-bit intermediate value, we expect that only one key survives the filtering conditions. Altogether, the memory complexity of the top MITM attack is about 2^{36} 64-bit words. The time complexity is dominated by step 2 and is equivalent to about 2^{36} 4-round Feistel structure evaluations, which is equivalent to about 2^{33} evaluations of the full GOST cryptosystem. For the bottom MITM attack, we obtain the same time and memory complexities, since the sizes of S_4 and S_5 are equal to the sizes of S_2 and S_3 , and the sets corresponding to G_1 , G_2 and G_3 are completely symmetrical.

6.3 The Complexity of the 8-Round Attack on GOST

We can now analyze the complexity of the attack described in Section 6.1: The time complexities of each of the MITM attacks on the bottom and top parts in steps (a) and (b) are equivalent to about 2^{36} 4-round Feistel structure evaluations, as calculated above. The number of expected matches for which we run the full cipher in step (b) is $2^{36+36-36} = 2^{36}$. Hence, the time complexity of these steps is equivalent to a bit more than 2^{36} full GOST evaluations. Since $|S_1| = 92$, the total time complexity of the attack is equivalent to about $2^{92+36} = 2^{128}$ GOST evaluations. The total memory complexity of the attack is about 2^{36} 64-bit words, and is dominated by the sorted list calculated in step (a).

7 Conclusions and Open Problem

In this paper we introduced several new techniques such as the fixed point property and two dimensional meet in the middle attacks, and used them to greatly improve the best known attacks on the full 32-round GOST. In particular, we reduced the memory complexity of the attacks from an impractical 2^{64} to a practical 2^{36} (and to an even more practical 2^{19} complexity, which can fit into the cache of modern microprocessors, with a small penalty in the running time). The lowest time complexity of our attacks is 2^{192} , which is 2^{32} times better than previously published attacks but still very far from being practical. Consequently, we are concerned about the demonstrated weaknesses in the design of GOST

(especially in its simplistic key schedule), but do not advocate that its current users should stop using it right away.

The main open problems left in this paper are whether it is possible to find faster attacks, and how to better exploit other amounts of available data (in addition to the 2^{32} and 2^{64} complexities considered in this paper, which are the natural thresholds for our techniques).

Acknowledgements. The authors thank Nathan Keller, Pierre-Alain Fouque and Charles Bouillaguet for useful discussions on this work, and the anonymous referees for their helpful comments on this paper which greatly improved the presentation of our results.

References

1. Biham, E., Dunkelman, O., Keller, N.: Improved Slide Attacks. In: Biryukov, A. (ed.) FSE 2007. LNCS, vol. 4593, pp. 153–166. Springer, Heidelberg (2007)
2. Chaum, D., Evertse, J.-H.: Cryptanalysis of DES with a Reduced Number of Rounds. In: Williams, H.C. (ed.) CRYPTO 1985. LNCS, vol. 218, pp. 192–211. Springer, Heidelberg (1986)
3. Courtois, N.T.: Algebraic Complexity Reduction and Cryptanalysis of GOST. Cryptology ePrint Archive, Report 2011/626 (2011), <http://eprint.iacr.org/>
4. Courtois, N.T.: Security Evaluation of GOST 28147-89 in View of International Standardisation. Cryptology ePrint Archive, Report 2011/211 (2011), <http://eprint.iacr.org/>
5. Courtois, N.T., Misztal, M.: Differential Cryptanalysis of GOST. Cryptology ePrint Archive, Report 2011/312 (2011), <http://eprint.iacr.org/>
6. Dinur, I., Dunkelman, O., Shamir, A.: Improved Attacks on Full GOST. Cryptology ePrint Archive, Report 2011/558 (2011), <http://eprint.iacr.org/>
7. Fleischmann, E., Gorski, M., Huehne, J.-H., Lucks, S.: Key Recovery Attack on full GOST Block Cipher with Negligible Time and Memory. Presented at Western European Workshop on Research in Cryptology (WEWoRC) (2009)
8. Isobe, T.: A Single-Key Attack on the Full GOST Block Cipher. In: Joux, A. (ed.) FSE 2011. LNCS, vol. 6733, pp. 290–305. Springer, Heidelberg (2011)
9. Kara, O.: Reflection Cryptanalysis of Some Ciphers. In: Chowdhury, D.R., Rijmen, V., Das, A. (eds.) INDOCRYPT 2008. LNCS, vol. 5365, pp. 294–307. Springer, Heidelberg (2008)
10. Kelsey, J., Schneier, B., Wagner, D.: Key-Schedule Cryptanalysis of IDEA, G-DES, GOST, SAFER, and Triple-DES. In: Kobitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 237–251. Springer, Heidelberg (1996)
11. Ko, Y., Hong, S., Lee, W., Lee, S., Kang, J.-S.: Related Key Differential Attacks on 27 Rounds of XTEA and Full-Round GOST. In: Roy, B., Meier, W. (eds.) FSE 2004. LNCS, vol. 3017, pp. 299–316. Springer, Heidelberg (2004)
12. Mendel, F., Pramstaller, N., Rechberger, C.: A (Second) Preimage Attack on the GOST Hash Function. In: Nyberg, K. (ed.) FSE 2008. LNCS, vol. 5086, pp. 224–234. Springer, Heidelberg (2008)
13. Mendel, F., Pramstaller, N., Rechberger, C., Kontak, M., Szmids, J.: Cryptanalysis of the GOST Hash Function. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 162–178. Springer, Heidelberg (2008)

14. National Bureau of Standards. Federal Information Processing Standard-Cryptographic Protection - Cryptographic Algorithm. GOST 28147-89 (1989)
15. OpenSSL. A Reference Implementation of GOST, <http://www.openssl.org/source/>
16. Rudskoy, V.: On Zero Practical Significance of Key Recovery Attack on Full GOST Block Cipher with Zero Time and Memory. Cryptology ePrint Archive, Report 2010/111 (2010), <http://eprint.iacr.org/>
17. Seki, H., Kaneko, T.: Differential Cryptanalysis of Reduced Rounds of GOST. In: Stinson, D.R., Tavares, S. (eds.) SAC 2000. LNCS, vol. 2012, pp. 315–323. Springer, Heidelberg (2001)

Zero Correlation Linear Cryptanalysis with Reduced Data Complexity

Andrey Bogdanov^{1,*} and Meiqin Wang^{1,2,*}

¹ KU Leuven, ESAT/COSIC and IBBT, Belgium

² Shandong University, Key Laboratory of Cryptologic Technology and Information Security, Ministry of Education, Shandong University, Jinan 250100, China

Abstract. Zero correlation linear cryptanalysis is a novel key recovery technique for block ciphers proposed in [5]. It is based on linear approximations with probability of exactly $1/2$ (which corresponds to the zero correlation). Some block ciphers turn out to have multiple linear approximations with correlation zero for each key over a considerable number of rounds. Zero correlation linear cryptanalysis is the counterpart of impossible differential cryptanalysis in the domain of linear cryptanalysis, though having many technical distinctions and sometimes resulting in stronger attacks.

In this paper, we propose a statistical technique to significantly reduce the data complexity using the high number of zero correlation linear approximations available. We also identify zero correlation linear approximations for 14 and 15 rounds of TEA and XTEA. Those result in key-recovery attacks for 21-round TEA and 25-round XTEA, while requiring less data than the full code book. In the single secret key setting, these are structural attacks breaking the highest number of rounds for both ciphers.

The findings of this paper demonstrate that the prohibitive data complexity requirements are not inherent in the zero correlation linear cryptanalysis and can be overcome. Moreover, our results suggest that zero correlation linear cryptanalysis can actually break more rounds than the best known impossible differential cryptanalysis does for relevant block ciphers. This might make a security re-evaluation of some ciphers necessary in the view of the new attack.

Keywords: block ciphers, key recovery, linear cryptanalysis, zero correlation linear cryptanalysis, data complexity, TEA, XTEA.

1 Introduction

1.1 Motivation

Differential and linear cryptanalyses [3, 31] are the two basic tools for evaluating the security of block ciphers such as the former U.S. encryption standard DES as well as its successor AES. While DES was developed at the time when differential

* Corresponding authors.

and linear cryptanalyses were not publicly known, the design of AES provably addresses these attacks.

Design strategies have been proposed such as the wide-trail design strategy [14] or decorrelation theory [43] to make ciphers resistant to the basic flavours of differential and linear cryptanalysis. However, a proof of resistance according to these strategies does not necessarily imply resistance to the extensions of these techniques such as impossible differential cryptanalysis [1, 7] and the recently proposed zero correlation linear cryptanalysis [5].

Standard differential cryptanalysis uses differentials with probabilities significantly higher than those expected for a randomly drawn permutation. Similarly, basic linear cryptanalysis uses linear approximations whose probabilities detectably deviate from $1/2$. At the same time, impossible differential cryptanalysis and zero correlation linear cryptanalysis are based on structural deviations of another kind: Differentials with zero probability are targeted in impossible differential cryptanalysis and linear approximations with probability of exactly $1/2$ correlation are exploited in zero correlation linear cryptanalysis. Thus, zero correlation linear cryptanalysis can be seen as the counterpart of impossible differential cryptanalysis in the domain of linear cryptanalysis.

The name of the attack originated from the notion of *correlation* [12,35]: If $\frac{1+c}{2}$ is the probability for a linear approximation to hold, c is called the correlation of this linear approximation. Clearly, putting $c = 0$ yields an unbiased linear approximation of probability $1/2$, or a *zero correlation linear approximation*.

Impossible differential cryptanalysis has been known to the cryptographic community since over a decade now. It has turned out a highly useful tool of attacking block ciphers [2, 16, 28–30, 42]. In fact, among meet-in-the-middle [15] and multiset-type attacks [19], it is the impossible differential cryptanalysis [29] that breaks the highest numbers of rounds of AES-128 and AES-256 in the classical single-key attack model as to date, the recent biclique cryptanalysis [4] being the notable exception though.

Zero correlation linear cryptanalysis is a novel promising attack technique that bears some technical similarities to impossible differential cryptanalysis but has its theoretical foundation in a different mathematical theory. Despite its newness, it has already been demonstrated to successfully apply to round-reduced AES and CLEFIA even in its basic form [5], which is highly motivating for further studies.

In this paper, we show how to remove the data requirement of the full codebook which was the major limitation of basic zero correlation linear cryptanalysis [5]. As an application of zero correlation linear cryptanalysis and this data complexity reduction technique, we propose attacks against round-reduced TEA and XTEA. For both ciphers, we can cryptanalyze more rounds than it was previously possible using less than the full codebook.

1.2 Contributions

The work at hand has two major contributions.

Data complexity reduction for zero correlation linear cryptanalysis.

The data requirements of the full codebook have been a crucial limitation for the recent zero correlation linear cryptanalysis to become a major cryptanalytic technique, though the length of the fundamental property (the length of the zero correlation linear approximation) was demonstrated to be comparable to that of impossible differentials for several cipher structures [5]. Overcoming this annoying limitation, a statistical technique of data complexity reduction for zero correlation linear cryptanalysis is the first contribution of this paper.

The data complexity reduction technique is based on the fact that, like any exploitable impossible differential, a typical zero correlation linear approximation is *truncated*: That is, once a zero correlation linear approximation has been identified that holds for all keys, it will as a rule imply an entire class of similar zero correlation linear approximations to exist. Those can be typically obtained by just changing several bits of the input mask, output mask or both. In other words, in most practical cases, there will be *multiple* zero correlation linear approximations available to the adversary which has been ignored by the previous analysis.

However, unlike in impossible differential cryptanalysis, the actual value of the correlation has to be estimated in zero correlation linear cryptanalysis and it is not enough to just wait for the impossible event to occur. In fact, the idea we use for zero correlation linear cryptanalysis is more similar to that of multiple linear cryptanalysis: We estimate the correlation of each individual linear approximation using a limited number of texts. Then, for a group of zero correlation linear approximations (i.e. for the right key), we expect the cumulative deviation of those estimations from 0 to be lower than that for a group of randomly chosen linear approximations (i.e. for a wrong key). Given the statistical behaviour of correlation for a randomly drawn permutation [13, 36], this consideration results in a χ^2 statistic and allows for a theoretical analysis of the complexity and error probabilities of a zero correlation linear attack that are confirmed by experiments.

Zero correlation linear cryptanalysis of round-reduced TEA and XTEA. TEA (Tiny Encryption Algorithm) is one of the first lightweight block ciphers. It is a 64-bit block cipher based on a balanced Feistel-type network with a simple ARX round function. TEA has 64 rounds and accepts a key of 128 bits. It favours both efficient hardware [23] and software implementations. TEA was designed by Wheeler and Needham and proposed at FSE'94 [44]. It was used in Microsoft's Xbox gaming console for checking software authenticity until its weakness as a hash function was used [41] to compromise the chain of trust. The block cipher XTEA [34] is the fixed version of TEA eliminating this property (having the same number rounds, block size, and key size). TEA and XTEA being rather popular ciphers, both are implemented in the Linux kernel.

Similarly to the complementation property of DES, TEA has an equivalent key property and its effective key size is 126 bits (compared to 128 bits suggested by the nominal key input size) [24]. Kelsey, Scheier and Wagner [25] proposed a practical related-key attack on the full TEA. Using complementation

Table 1. Summary of cryptanalytic results on round-reduced TEA* and XTEA in the single-key setting

attack	#rounds	data	comp. compl.	memory	Pr[success]	ref.
TEA						
impossible differential	11	$2^{52.5}$ CP	2^{84}	NA	NA	[33]
truncated differential	17	1920 CP	$2^{123.37}$	NA	NA	[21]
impossible differential	17	2^{57} CP	$2^{106.6}$	2^{49}	NA	[9]
zero correlation linear	21	$2^{62.62}$ KP	$2^{121.52}$	negligible	0.846	this paper
zero correlation linear	23	2^{64}	$2^{119.64}$	negligible	1	this paper
XTEA						
impossible differential	14	$2^{62.5}$ CP	2^{85}	NA	NA	[33]
truncated differential	23	$2^{20.55}$ CP	$2^{120.65}$	NA	0.969	[21]
meet-in-the-middle	23	18 KP	2^{117}		$1 - 2^{-1025}$	[38]
impossible differential	23	$2^{62.3}$ CP	$2^{114.9}$	$2^{94.3}$	NA	[9]
impossible differential	23	2^{63}	2^{101} MA + $2^{105.6}$	2^{103}	NA	[9]
zero correlation linear	25	$2^{62.62}$ KP	$2^{124.53}$	2^{80}	0.846	this paper
zero correlation linear	27	2^{64}	$2^{120.71}$	negligible	1	this paper

CP: Chosen Plaintexts, KP: Known Plaintexts.

Memory: the number of 32-bit words.

*The effective key length for TEA is 126 bit

cryptanalysis [8], up to 36 rounds of XTEA can be attacked with related keys for all keys. The work [8] also contains related-key attacks for up to 50 rounds of XTEA working for a weak key class.

In the classical single-key setting, however, by far not all rounds of TEA are broken by structural attacks (whereas the effective key size is 126 bits for the full cipher). The truncated differential result on 17 rounds remains the best cryptanalysis of TEA [21]. Impossible differential cryptanalysis [9] has yielded a faster attack against 17 rounds of TEA. Similarly, 23 rounds of XTEA have been cryptanalyzed so far using truncated differential [21], impossible differential [9] and well as meet-in-the-middle attacks [38]. That is, for both TEA and XTEA, there has been no progress in terms of the number of attacked rounds since 2003.

In this paper, using zero correlation linear cryptanalysis, we cryptanalyze 21 rounds of TEA and 25 rounds of XTEA with $2^{62.62}$ *known* plaintexts (in contrast to *chosen* texts required in impossible differential cryptanalysis). Certainly, zero correlation linear cryptanalysis for lower number of rounds yields a lower data complexity for both TEA and XTEA. Moreover, unlike most impossible differential attacks including those on TEA and XTEA [9], zero correlation linear cryptanalysis is able to profit from the full code available. If all 2^{64} texts are available to the adversary, we propose zero correlation linear cryptanalysis for 23 rounds of TEA and 27 rounds of XTEA. Our cryptanalytic results are summarized and compared to previous cryptanalysis in Table 1.

As opposed to the initial intuition expressed in [5], both major contributions of this work — the data complexity reduction and the new attacks on more rounds of TEA and XTEA — demonstrate that zero correlation linear cryptanalysis can actually perform better than impossible differential cryptanalysis. Moreover, we expect the security of more ciphers to be reevaluated under the consideration of zero correlation linear cryptanalysis.

1.3 Outline

We start with a review of the basic zero correlation linear cryptanalysis for block ciphers in Section 2. In Section 3, we introduce a χ^2 statistical technique for reducing the data requirements of zero correlation linear cryptanalysis and thoroughly investigate its complexity. In Section 4, the 14- and 15-round zero correlation linear approximations are demonstrated for block ciphers TEA and XTEA. Section 5 gives several zero correlation key recoveries for round-reduced TEA and XTEA. The full version [6] of this paper is available online and contains proofs of some technical statements as well as further zero correlation linear attacks on round-reduced TEA and XTEA.

2 Basic Zero Correlation Linear Cryptanalysis

Zero correlation linear cryptanalysis has been introduced in [5]. Below we briefly review its basic ideas and methods.

2.1 Linear Approximations with Correlation Zero

Consider an n -bit block cipher f_K with key K . Let P denote a plaintext which is mapped to ciphertext C under key K , $C = f_K(P)$. If Γ_P and Γ_C are nonzero plaintext and ciphertext linear masks of n bit each, we denote by $\Gamma_P \rightarrow \Gamma_C$ the linear approximation

$$\Gamma_P^T P \oplus \Gamma_C^T C = 0.$$

Here, $\Gamma_A^T A$ denotes the multiplication of the transposed bit vector Γ_A (linear mask for A) by a column bit vector A over \mathbb{F}_2 . The linear approximation $\Gamma_P \rightarrow \Gamma_C$ has probability

$$p_{\Gamma_P, \Gamma_C} = \Pr_{P \in \mathbb{F}_2^n} \{ \Gamma_P^T P \oplus \Gamma_C^T C = 0 \}. \quad (1)$$

The value

$$c_{\Gamma_P, \Gamma_C} = 2p_{\Gamma_P, \Gamma_C} - 1 \quad (2)$$

is called the *correlation (or bias)* of linear approximation $\Gamma_P \rightarrow \Gamma_C$. Note that $p_{\Gamma_P, \Gamma_C} = 1/2$ is equivalent to *zero correlation* $c_{\Gamma_P, \Gamma_C} = 0$:

$$p_{\Gamma_P, \Gamma_C} = \Pr_{P \in \mathbb{F}_2^n} \{ \Gamma_P^T P \oplus \Gamma_C^T C = 0 \} = 1/2. \quad (3)$$

In fact, for a randomly drawn permutation of sufficiently large bit size n , zero is the most frequent single value of correlation for a nontrivial linear approximation. Correlation goes to small values for increasing n , the probability to get exactly zero decreases as a function of n though. More precisely, the probability of the linear approximation $\Gamma_P \rightarrow \Gamma_C$ with $\Gamma_P, \Gamma_C \neq 0$ to have zero correlation has been shown [5, Proposition 2] to be approximated by

$$\frac{1}{\sqrt{2\pi}} 2^{\frac{4-n}{2}}. \quad (4)$$

2.2 Two Examples

Given a randomly chosen permutation, however, it is hard to tell a priori which of its nontrivial linear approximations in particular has zero correlation. At the same time, it is often possible to identify groups of zero correlation linear approximations for a block cipher f_K once it has compact description with a distinct structure. Moreover, in many interesting cases, these linear approximations will have zero correlation *for any key* K . Here are two examples provided in [5]:

- **AES:** The data transform of AES has a set of zero correlation linear approximations over 4 rounds (3 full rounds appended by 1 incomplete rounds with MixColumns omitted). If Γ and Γ' are 4-byte column linear masks with exactly one nonzero byte, then each of the linear approximations $(\Gamma, 0, 0, 0) \rightarrow (\Gamma', 0, 0, 0)$ over 4 AES rounds has zero correlation [5, Theorem 2].
- **CLEFIA-type GFNs:** CLEFIA-type generalized Feistel networks [40] (also known as type-2 GFNs with 4 lines [45]) have zero correlation linear approximations over 9 rounds, if the underlying F-functions of the Feistel construction are invertible. For $a \neq 0$, the linear approximations $(a, 0, 0, 0) \rightarrow (0, 0, 0, a)$ and $(0, 0, a, 0) \rightarrow (0, a, 0, 0)$ over 9 rounds have zero correlation [5, Theorem 1].

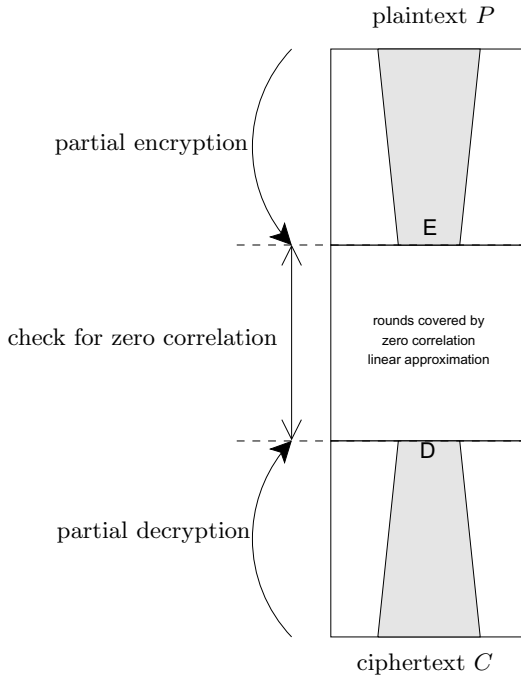


Fig. 1. High-level view of key recovery in zero correlation linear cryptanalysis

2.3 Key Recovery with Zero Correlation Linear Approximations

Based on linear approximations of correlation zero, a technique similar to Matsui's Algorithm 2 [31] can be used for key recovery. Let the adversary have N known plaintext-ciphertexts and ℓ zero correlation linear approximations $\{\Gamma_E \rightarrow \Gamma_D\}$ for a part of the cipher, with $\ell = |\{\Gamma_E \rightarrow \Gamma_D\}|$. The linear approximations $\{\Gamma_E \rightarrow \Gamma_D\}$ are placed in the middle of the attacked cipher. Let E and D be the partial intermediate states of the data transform at the boundaries of the linear approximations.

Then the key can be recovered using the following approach (see also Figure 1):

1. Guess the bits of the key needed to compute E and D . For each guess:
 - (a) Partially encrypt the plaintexts and partially decrypt the ciphertexts up to the boundaries of the zero correlation linear approximation $\Gamma_E \rightarrow \Gamma_D$.
 - (b) Estimate the correlations $\{\hat{c}_{\Gamma_E, \Gamma_D}\}$ of all linear approximations in $\{\Gamma_E \rightarrow \Gamma_D\}$ for the key guess using the partially encrypted and decrypted values E and D by counting how many times $\Gamma_E^T E \oplus \Gamma_D^T D$ is zero over N input/output pairs, see (1) and (2).
 - (c) Perform a test on the estimated correlations $\{\hat{c}_{\Gamma_E, \Gamma_D}\}$ for $\{\Gamma_E \rightarrow \Gamma_D\}$ to tell of the estimated values of $\{\hat{c}_{\Gamma_E, \Gamma_D}\}$ are compatible with the hypothesis that all of the actual values of $\{c_{\Gamma_E, \Gamma_D}\}$ are zero.
2. Test the surviving key candidates against a necessary number of plaintext-ciphertext pairs according to the unicity distance for the attacked cipher.

Step 1(c) of the technique above relies on an efficient test distinguishing between the hypothesis that $\{c_{\Gamma_E, \Gamma_D}\}$ are all zero and the alternative hypothesis. The work [5] requires the exact evaluation of the correlation value (defined by the probability of a linear approximation) and the data complexity is restricted to $N = 2^n$ in [5]. Thus, a small number ℓ of linear approximations is usually enough in [5] and $\hat{c}_{\Gamma_E, \Gamma_D} = c_{\Gamma_E, \Gamma_D}$, though the data complexity of the full codebook is too restrictive.

For most ciphers (including the examples of Subsection 2.2), however, a large number ℓ of zero correlation linear approximations is available. This freedom is not used in [5]. At the same time, it has been shown in the experimental work [10] that any value of correlation can be used for key recovery in a linear attack with reduced data complexity, once enough linear approximations are available. Despite its convincing experimental evidence, [10] gives no theoretical data complexity estimations and does not provide any ways of constructing linear approximations with certain properties.

In the next section of this paper, we provide a framework for reducing the data complexity N if many zero correlation linear approximations are known.

3 Reduction of Data Complexity with Many Approximations

3.1 Distinguishing between Two Normal Distributions

Consider two normal distributions: $\mathcal{N}(\mu_0, \sigma_0)$ with mean μ_0 and standard deviation σ_0 , and $\mathcal{N}(\mu_1, \sigma_1)$ with mean μ_1 and standard deviation σ_1 . A sample s is

drawn from either $\mathcal{N}(\mu_0, \sigma_0)$ or $\mathcal{N}(\mu_1, \sigma_1)$. It has to be decided if this sample is from $\mathcal{N}(\mu_0, \sigma_0)$ or from $\mathcal{N}(\mu_1, \sigma_1)$. The test is performed by comparing the value s to some threshold value t . Without loss of generality, assume that $\mu_0 < \mu_1$. If $s \leq t$, the test returns " $s \in \mathcal{N}(\mu_0, \sigma_0)$ ". Otherwise, if $s > t$, the test returns " $s \in \mathcal{N}(\mu_1, \sigma_1)$ ". There will be error probabilities of two types:

$$\begin{aligned}\beta_0 &= \Pr\{ "s \in \mathcal{N}(\mu_1, \sigma_1)" | s \in \mathcal{N}(\mu_0, \sigma_0) \}, \\ \beta_1 &= \Pr\{ "s \in \mathcal{N}(\mu_0, \sigma_0)" | s \in \mathcal{N}(\mu_1, \sigma_1) \}.\end{aligned}$$

Here a condition is given on μ_0, μ_1, σ_0 , and σ_1 such that the error probabilities are β_0 and β_1 . The proof immediately follows from the basics of probability theory (see e.g. [18, 20]) and is given in the full version [6] of the paper for completeness.

Proposition 1. *For the test to have error probabilities of at most β_0 and β_1 , the parameters of the normal distributions $\mathcal{N}(\mu_0, \sigma_0)$ and $\mathcal{N}(\mu_1, \sigma_1)$ with $\mu_0 \neq \mu_1$ have to be such that*

$$\frac{z_{1-\beta_1}\sigma_1 + z_{1-\beta_0}\sigma_0}{|\mu_1 - \mu_0|} = 1,$$

where $z_{1-\beta_1}$ and $z_{1-\beta_0}$ are the quantiles of the standard normal distribution.

3.2 A Known Plaintext Distinguisher with Many Zero Correlation Linear Approximations

Let the adversary be given N known plaintext-ciphertext pairs and ℓ zero correlation linear approximations for an n -bit block cipher. The adversary aims to distinguish between this cipher and a randomly drawn permutation.

The procedure is as follows. For each of the ℓ given linear approximations, the adversary computes the number T_i of times the linear approximations are fulfilled on N plaintexts, $i \in \{1, \dots, \ell\}$. Each T_i suggests an empirical correlation value $\hat{c}_i = 2\frac{T_i}{N} - 1$. Then, the adversary evaluates the statistic:

$$\sum_{i=1}^{\ell} \hat{c}_i^2 = \sum_{i=1}^{\ell} \left(2\frac{T_i}{N} - 1 \right)^2. \quad (5)$$

It is expected that for the cipher with ℓ known zero correlation linear approximations, the value of statistic (5) will be lower than that for ℓ linear approximations of a randomly drawn permutation. In a key-recovery setting, the right key will result in statistic (5) being among the lowest values for all candidate keys if ℓ is high enough. In the sequel, we treat this more formally.

3.3 Correlation under Right and Wrong Keys

Consider the key recovery procedure outlined in Subsection 2.3 given N known plaintext-ciphertext pairs. There will be two cases:

- *Right key guess:* Each of the values \hat{c}_i in (5) approximately follows the normal distribution with zero mean and standard deviation $1/\sqrt{N}$ with good precision (c.f. e.g. [22, 39]) for sufficiently large N :

$$\hat{c}_i \sim \mathcal{N}(0, 1/\sqrt{N}).$$

- *Wrong key guess:* Each of the values \hat{c}_i in (5) approximately follows the normal distribution with mean c_i and standard deviation $1/\sqrt{N}$ for sufficiently large N :

$$\hat{c}_i \sim \mathcal{N}(c_i, 1/\sqrt{N}) \text{ with } c_i \sim \mathcal{N}(0, 2^{-n/2}),$$

where c_i is the exact value of the correlation which is itself distributed as $\mathcal{N}(0, 2^{-n/2})$ over random permutations with $n \geq 5$ — a result due to [13, 36]. Thus, our wrong key hypothesis is that for each wrong key, the adversary obtains a permutation with linear properties close to those of a randomly chosen permutation.

3.4 Distribution of the Statistic

Based on these distributions of \hat{c}_i , we now derive the distributions of statistic (5) in these two cases.

Right Key Guess. In this case, we deal with ℓ zero correlation linear approximations:

$$\sum_{i=1}^{\ell} \hat{c}_i^2 \sim \sum_{i=1}^{\ell} \mathcal{N}^2(0, 1/\sqrt{N}) = \frac{1}{N} \sum_{i=1}^{\ell} \mathcal{N}^2(0, 1) = \frac{1}{N} \chi_{\ell}^2,$$

where χ_{ℓ}^2 is the χ^2 -distribution with ℓ degrees of freedom which has mean ℓ and standard deviation $\sqrt{2\ell}$, assuming the independency of underlying distributions. For sufficiently large ℓ , χ_{ℓ}^2 converges to the normal distribution. That is, χ_{ℓ}^2 approximately follows $\mathcal{N}(\ell, \sqrt{2\ell})$, and:

$$\sum_{i=1}^{\ell} \hat{c}_i^2 \sim \frac{1}{N} \chi_{\ell}^2 \approx \frac{1}{N} \mathcal{N}(\ell, \sqrt{2\ell}) = \mathcal{N}\left(\frac{\ell}{N}, \frac{\sqrt{2\ell}}{N}\right). \quad (6)$$

Proposition 2. *Consider ℓ nontrivial zero correlation linear approximations for a block cipher with a fixed key. If N is the number of known plaintext-ciphertext pairs, T_i is the number of times such a linear approximation is fulfilled for $i \in \{1, \dots, \ell\}$, and ℓ is high enough, then, assuming the counters T_i are independent, the following approximate distribution holds for sufficiently large N and n :*

$$\sum_{i=1}^{\ell} \left(2\frac{T_i}{N} - 1\right)^2 \sim \mathcal{N}\left(\frac{\ell}{N}, \frac{\sqrt{2\ell}}{N}\right).$$

Wrong Key Guess. The wrong key hypothesis is that we deal with pick a permutation at random for each wrong key. Therefore, the ℓ given linear approximations will have randomly drawn correlations, under this hypothesis. Thus, as mentioned above:

$$\sum_{i=1}^{\ell} \hat{c}_i^2 \sim \sum_{i=1}^{\ell} \mathcal{N}^2 \left(c_i, 1/\sqrt{N} \right), \text{ where } c_i \sim \mathcal{N} \left(0, 2^{-n/2} \right).$$

First, we show that the underlying distribution of \hat{c}_i is actually normal with mean 0. Then we show that the sum approximately follows χ^2 -distribution assuming the independency of underlying distributions, and can be approximated by another normal distribution.

Since

$$\begin{aligned} \mathcal{N} \left(c_i, 1/\sqrt{N} \right) &= c_i + \mathcal{N} \left(0, 1/\sqrt{N} \right) \\ &= \mathcal{N} \left(0, 1/\sqrt{2^n} \right) + \mathcal{N} \left(0, 1/\sqrt{N} \right) \\ &= \mathcal{N} \left(0, \sqrt{1/N + 1/2^n} \right), \end{aligned}$$

the distribution above is a χ^2 -distribution with ℓ degrees of freedom up to a factor, under the independency assumption:

$$\begin{aligned} \sum_{i=1}^{\ell} \mathcal{N}^2 \left(c_i, 1/\sqrt{N} \right) &= \sum_{i=1}^{\ell} \mathcal{N}^2 \left(0, \sqrt{\frac{1}{N} + \frac{1}{2^n}} \right) \\ &= \left(\frac{1}{N} + \frac{1}{2^n} \right) \sum_{i=1}^{\ell} \mathcal{N}^2 \left(0, 1 \right) \\ &= \left(\frac{1}{N} + \frac{1}{2^n} \right) \chi_{\ell}^2. \end{aligned}$$

As for the right keys, for sufficiently large ℓ , χ_{ℓ}^2 can be approximated by the normal distribution with mean ℓ and standard deviation $\sqrt{2\ell}$. Thus:

$$\begin{aligned} \sum_{i=1}^{\ell} \hat{c}_i^2 &\sim \left(\frac{1}{N} + \frac{1}{2^n} \right) \chi_{\ell}^2 \approx \left(\frac{1}{N} + \frac{1}{2^n} \right) \mathcal{N} \left(\ell, \sqrt{2\ell} \right) \\ &= \mathcal{N} \left(\frac{\ell}{N} + \frac{\ell}{2^n}, \frac{\sqrt{2\ell}}{N} + \frac{\sqrt{2\ell}}{2^n} \right). \end{aligned}$$

Proposition 3. *Consider ℓ nontrivial linear approximations for a randomly drawn permutation. If N is the number of known plaintext-ciphertext pairs, T_i is the number of times a linear approximation is fulfilled for $i \in \{1, \dots, \ell\}$, and ℓ is high enough, then, assuming the independency of T_i , the following approximate distribution holds for sufficiently large N and n :*

$$\sum_{i=1}^{\ell} \left(2 \frac{T_i}{N} - 1 \right)^2 \sim \mathcal{N} \left(\frac{\ell}{N} + \frac{\ell}{2^n}, \frac{\sqrt{2\ell}}{N} + \frac{\sqrt{2\ell}}{2^n} \right).$$

3.5 Data Complexity of the Distinguisher

Combining Propositions 2 and 3 with Proposition 1, one obtains the condition:

$$\frac{z_{1-\beta_1} \left(\frac{\sqrt{2\ell}}{N} + \frac{\sqrt{2\ell}}{2^n} \right) + z_{1-\beta_0} \frac{\sqrt{2\ell}}{N}}{\left(\frac{\ell}{N} + \frac{\ell}{2^n} \right) - \frac{\ell}{N}} = 1.$$

The left part of this equation can be simplified to

$$\frac{2^{n+0.5}}{N\sqrt{\ell}} (z_{1-\beta_0} + z_{1-\beta_1}) + \frac{z_{1-\beta_1}\sqrt{2}}{\sqrt{\ell}},$$

which yields

Theorem 1. *With the assumptions of Propositions 1 to 3, using ℓ nontrivial zero correlation linear approximations, to distinguish between a wrong key and a right key with probability β_1 of false positives and probability β_0 of false negatives, a number N of known plaintext-ciphertext pairs is sufficient if the following condition is fulfilled:*

$$\frac{2^{n+0.5}}{N\sqrt{\ell}} (z_{1-\beta_0} + z_{1-\beta_1}) + \frac{z_{1-\beta_1}\sqrt{2}}{\sqrt{\ell}} = 1.$$

The success probability of an attack is defined by the probability β_0 of false negatives. The probability β_1 of false positives determines the number of surviving key candidates and, thus, influences the computational complexity of the key recovery.

4 Linear Approximations with Correlation Zero for TEA and XTEA

In [5], a sufficient condition is given for a linear approximation to have a correlation of zero. Namely, if for a linear approximation there exist no linear characteristics with non-zero correlation contributions, then the correlation of the linear approximation is exactly zero.

4.1 The Block Ciphers TEA and XTEA

TEA is a 64-round iterated block cipher with 64-bit block size and 128-bit key which consist of four 32-bit words $K[0]$, $K[1]$, $K[2]$ and $K[3]$. TEA does not have any iterative key schedule algorithm. Instead, the key words are used directly in round functions. The round constant is derived from the constant $\delta = 9e3779b9_x$ and the round number. We denote the input and the output of the r -th round for $1 \leq r \leq 64$ by (L_r, R_r) and (L_{r+1}, R_{r+1}) , respectively. $L_{r+1} = R_r$ and R_{r+1} is computed as follows:

$$R_{r+1} = \begin{cases} L_r + (((R_r \ll 4) + K[0]) \oplus (R_r + i \cdot \delta) \oplus (R_r \gg 5 + K[1])) & r = 2i - 1, \\ L_r + (((R_r \ll 4) + K[2]) \oplus (R_r + i \cdot \delta) \oplus (R_r \gg 5 + K[3])) & r = 2i, 1 \leq i \leq 32. \end{cases}$$

Like TEA, XTEA is also a 64-round Feistel cipher with 64-bit block and 128-bit key. Its 128-bit secret key K is represented by four 32-bit words $K[0]$, $K[1]$, $K[2]$ and $K[3]$ as well. The derivation of the subkey word number is slightly more complex though. The input of the r -th round is (L_r, R_r) and the output is (L_{r+1}, R_{r+1}) . Again, $L_{r+1} = R_r$ and R_{r+1} is derived as:

$$R_{r+1} = \begin{cases} L_r + (((R_r \ll 4 \oplus R_r \gg 5) + R_r) \oplus ((i-1) \cdot \delta + K[((i-1) \cdot \delta \ll 11) \& 3])) & r = 2i - 1, \\ L_r + (((R_r \ll 4 \oplus R_r \gg 5) + R_r) \oplus (i \cdot \delta + K[(i \cdot \delta \ll 11) \& 3])) & r = 2i, 1 \leq i \leq 32. \end{cases}$$

These round functions of TEA and XTEA are illustrated in Figure 2.

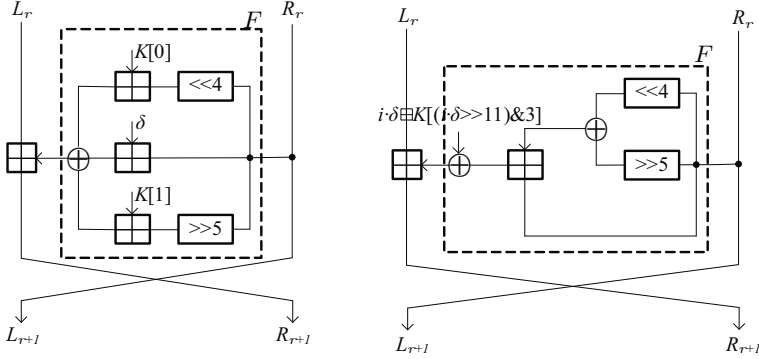


Fig. 2. Round function for TEA(left) and XTEA(right)

4.2 Notations

To demonstrate zero correlation linear approximations for TEA and XTEA, we will need the following notations (the least significant bit of a word has number 0):

- $e_{i,\sim}$ is a 32-bit word that has zeros in bits $(i + 1)$ to 31, one in bit i and undefined values in bits 0 to $(i - 1)$,
- $e_{i\sim j}$ is a 32-bit word that has zeros in bits $(i + 1)$ to 31 and bits 0 to $(j - 1)$, a one in bit i and undefined values in bits j to $(i - 1)$ for $j < i$,
- $\bar{e}_{i,\sim}$ is a 32-bit word that has zeros in bits $(i + 1)$ to 31, undefined values in bits 0 to i ,
- $?$ is an undefined value,
- $X^{i\sim j}$ is bits from j to i of the value X , $j < i$, and
- X^i is the value of bit i of X .

4.3 Linear Approximation of Modular Addition

Here, we first demonstrate the properties of linear approximations with non-zero correlation over the modular addition, which is the only nonlinear part of the TEA and XTEA transformation (summarized as Property 1). Then we use it to show a condition for linear approximation with non-zero correlation for one round of TEA and XTEA (stated as Property 2).

For the modular addition of two n -bit inputs x and y , the output z can be computed as:

$$z = (x + y) \pmod{2^n}.$$

We denote the mask values for x , y and z as $\Gamma x, \Gamma y$ and Γz , respectively ($x, y, z, \Gamma x, \Gamma y$, and $\Gamma z \in \mathbb{F}_2^n$). The linear approximation for the modular addition is then $\Gamma x^T \cdot x \oplus \Gamma y^T \cdot y = \Gamma z^T \cdot z$ and is referred to as

$$+ : (\Gamma x | \Gamma y) \rightarrow \Gamma z.$$

Property 1 (Modular addition). In any linear approximation $(\Gamma x|\Gamma y) \rightarrow \Gamma z$ of the modular addition with a non-zero correlation, the most significant non-zero mask bit for $\Gamma x, \Gamma y$ and Γz is the same.

Property 1 is proven in the full version [6] of the paper.

4.4 Linear Approximation of One TEA/XTEA Round

Using Property 1 for modular addition, as all other operations in TEA and XTEA are linear, we can derive conditions on a special class of approximations with non-zero correlation for the round function of TEA and XTEA. See Figures 4 and 3 for an illustration.

As in Subsection 4.1, the input and output of round r in TEA and XTEA are $(L_r|R_r)$ and $(L_{r+1}|R_{r+1})$, respectively. Correspondingly, $(\Gamma_r^L|\Gamma_r^R)$ and $(\Gamma_{r+1}^L|\Gamma_{r+1}^R)$ are input and output linear masks of the round. So the linear approximation over the round is

$$(X)TEA \text{ round } r : (\Gamma_r^L|\Gamma_r^R) \rightarrow (\Gamma_{r+1}^L|\Gamma_{r+1}^R)$$

and has the following

Property 2 (One round). If $\Gamma_r^L = e_{i,\sim}$ and $\Gamma_r^R = e_{j,\sim}$, ($j < i$), then one needs $\Gamma_{r+1}^R = e_{i,\sim}$ and $\Gamma_{r+1}^L = e_{i,\sim} \oplus e_{i+5\sim 5}$ for the approximation to have a non-zero correlation. Similarly, for the decryption round function of TEA, if the input mask and the output mask for round r are $(\Gamma_r^L|\Gamma_r^R)$ and $(\Gamma_{r+1}^L|\Gamma_{r+1}^R)$, respectively. If $\Gamma_r^R = e_{i,\sim}$ and $\Gamma_r^L = e_{j,\sim}$, ($j < i$), then we have $\Gamma_{r+1}^L = e_{i,\sim}$ and $\Gamma_{r+1}^R = e_{i,\sim} \oplus e_{i+5\sim 5}$.

4.5 Zero Correlation Approximations for 14 and 15 Rounds of TEA/XTEA

With the one-round property of linear approximation in TEA and XTEA derived in the previous subsection, we can identify zero correlation approximations over 14 and 15 rounds of both TEA and XTEA.

Proposition 4. Over 15 rounds of TEA and XTEA, any linear approximation with input mask $(\Gamma_1^R|\Gamma_1^L) = (1|0)$ and output mask $(\Gamma_{15}^R|\Gamma_{15}^L) = (0|e_{1,\sim})$ has a correlation of exactly zero. Moreover, over 14 rounds of TEA and XTEA, any linear approximation with input mask $(\Gamma_1^R|\Gamma_1^L) = (1|0)$ and output mask $(\Gamma_{14}^R|\Gamma_{14}^L) = (e_{1,\sim}|e_{5,\sim})$ has zero correlation.

Proof. First, we follow the linear approximation in the forward direction. From $\Gamma_1^L = 0$ and $\Gamma_1^R = 1$, it is obtained that $\Gamma_2^L = 0$ and $\Gamma_2^R = 1$, then we get $\Gamma_3^L = 1 \oplus (1 \ll 5)$ and $\Gamma_3^R = 1$. From Property 2, $\Gamma_3^L = 1 \oplus (1 \ll 5)$ and $\Gamma_3^R = 1$, then we have $\Gamma_4^R = e_{5,\sim}$ and $\Gamma_4^L = e_{5,\sim} \oplus e_{5+5\sim 5} \oplus 1 = e_{10,\sim}$. Similarly, we get $(\Gamma_5^R|\Gamma_5^L) = (e_{10,\sim}|e_{15,\sim})$, $(\Gamma_6^R|\Gamma_6^L) = (e_{15,\sim}|e_{20,\sim})$, $(\Gamma_7^R|\Gamma_7^L) = (e_{20,\sim}|e_{25,\sim})$, $(\Gamma_8^R|\Gamma_8^L) = (e_{25,\sim}|e_{30,\sim})$ and $(\Gamma_9^R|\Gamma_9^L) = (e_{30,\sim}|?)$.




Fig. 3. Zero correlation linear approximation for 14-round TEA and XTEA (grey – undefined bits, black – bits set to 1)




Fig. 4. Zero correlation linear approximation for 15-round TEA and XTEA (grey – undefined bits, black – bits set to 1)

Second, we follow the 7-round linear approximation in the backward direction. From $\Gamma_{16}^L = e_{1,\sim}$ and $\Gamma_{16}^R = 0$, we can derive that $(\Gamma_{15}^R|\Gamma_{15}^L) = (e_{1,\sim}|0)$, $(\Gamma_{14}^R|\Gamma_{14}^L) = (e_{1,\sim}\oplus e_{6\sim 5}|e_{1,\sim})$, $(\Gamma_{13}^R|\Gamma_{13}^L) = (e_{11,\sim}|e_{6,\sim})$, $(\Gamma_{12}^R|\Gamma_{12}^L) = (e_{16,\sim}|e_{11,\sim})$, $(\Gamma_{11}^R|\Gamma_{11}^L) = (e_{21,\sim}|e_{16,\sim})$, $(\Gamma_{10}^R|\Gamma_{10}^L) = (e_{26,\sim}|e_{21,\sim})$ and $(\Gamma_9^R|\Gamma_9^L) = (e_{31,\sim}|e_{26,\sim})$.

From the forward direction, the most significant bit of Γ_9^R has to be zero, and from the backward direction, the most significant bit of Γ_9^R has to be one. This yields a contradiction and shows that there are no characteristics for this linear approximation. By the sufficient condition of [5] for constructing zero correlation linear approximations, this is enough for the approximation to have correlation zero. So the linear approximation for 15-round TEA and XTEA with the input mask $(1|0)$ and the output mask $(0|e_{1,\sim})$ has zero correlation. By restricting this linear approximation to 14 rounds and adding several undefined bits to the output mask, one gets all the claims of the proposition. \square

There are only 2 zero correlation linear approximations of this form over 15 rounds. We note however that there are 2^7 different zero correlation linear approximations over 14 rounds of both TEA and XTEA. They can be generated by setting the undefined bits (depicted in gray in Figure 3 and Figure 4) to different values.

5 Zero Correlation Linear Cryptanalysis of Round-Reduced (X)TEA




Fig. 5. Key recovery for 21 rounds of TEA. For the estimation of correlation, grey and black bits need to be computed and white bits are irrelevant. Uses the zero correlation approximation of Figure 3.




Fig. 6. Key recovery for 25 rounds of XTEA. For the estimation of correlation, grey and black bits need to be computed and white bits are irrelevant. Uses the zero correlation approximation of Figure 3.

5.1 Key Recovery for 21 Rounds of TEA

For the cryptanalysis of 21-round TEA, we use the 14-round zero correlation approximations of the type depicted in Figure 3 of Subsection 4.5. The availability of many such approximations allows us to use the data complexity reduction technique of Section 3.

We place the 14-round zero correlation linear approximations in the middle of the 21-round TEA. It covers rounds 5 to 18. Following the procedure outlined in

Subsection 2.3, up to the boundaries of the linear approximations, we partially encrypt over the 4 first rounds 1 to 4 and partially decrypt over the 3 last rounds 19 to 21. The attack is illustrated in Figure 5.

The linear approximations involve 9 state bits: R_5^0 , $R_{19}^{1\sim 0}$, and $L_{19}^{5\sim 0}$. In the corresponding 9 bits of the input and output masks, only 7 can take on 0 and 1 values: Γ_{19}^R and $\Gamma_{19}^{L^{5\sim 0}}$. For the evaluation of the linear approximations from a plaintext-ciphertext pair, we guess 54 key bits $K_0^{15\sim 0}$, $K_1^{15\sim 0}$, $K_2^{10\sim 0}$, and $K_3^{10\sim 0}$. The attack flow is as follows given N known plaintext-ciphertext pairs.

For each possible guess of the 54-bit subkey $\kappa = (K_0^{15\sim 0}|K_1^{15\sim 0}|K_2^{10\sim 0}|K_3^{10\sim 0})$:

1. Allocate a 128-bit counter W and set it to zero. W will contain the χ^2 statistic for the subkey guess κ .
2. Allocate a 64-bit counter $V[x]$ for each of 2^9 possible values of

$$x = (R_5^0|R_{19}^{1\sim 0}|L_{19}^{5\sim 0})$$

and set it to 0. $V[x]$ will contain the number of times the partial state value x occurs for N texts.

3. For each of N plaintext-ciphertext pairs: partially encrypt 4 rounds and partially decrypt 3 rounds, obtain the 9-bit value for $x = (R_5^0|R_{19}^{1\sim 0}|L_{19}^{5\sim 0})$ and add one to the counter $V[x]$.
4. For each of 2^7 zero correlation linear approximations:
 - (a) Set the 64-bit counter U to zero.
 - (b) For 2^9 values of x , verify if the linear approximation holds. If so, add $V[x]$ to U .
 - (c) $W = W + (2 \cdot U/N - 1)^2$.
5. If $W < t$, then κ is a possible subkey candidate and all cipher keys it is compatible with are tested exhaustively against a maximum of 3 plaintext-ciphertext pairs.

The correct 54-bit subkey κ is likely to be among the candidates with the χ^2 statistic W lower than the threshold $t = \sigma_0 \cdot z_{1-\beta_0} + \mu_0 = \frac{\sqrt{2\ell}}{N} \cdot z_{1-\beta_0} + \frac{\ell}{N} = \frac{\sqrt{2 \cdot 2^7}}{N} \cdot z_{1-\beta_0} + \frac{2^7}{N}$, see Subsection 3.1 with its Proposition 1 as well as Theorem 1.

In this attack, we set $\beta_0 = 2^{-2.7}$, $\beta_1 = 2^{-4.49}$ and get $z_{1-\beta_0} = 1$, $z_{1-\beta_1} = 1.7$. Note once again that $n = 64$ and $\ell = 2^7$. Theorem 1 suggests the data complexity of $N = 2^{62.62}$ known plaintext-ciphertexts with those parameters. The decision threshold is $t = 2^{-55.56}$.

The computational complexity is dominated by Steps 3 and 5. The computational complexity T_3 of Step 3 is 2^{54} times 7 half-round encryptions for each of N texts. This gives $T_3 = 2^{54} \cdot 2^{62.62} \cdot 7 \cdot 0.5/21 = 2^{114.03}$ 21-round TEA encryptions.

One in $1/\beta_1 = 2^{4.49}$ keys is expected to survive the test against zero correlation. The remaining key space is covered by exhaustive search which is performed in Step 5. The computational complexity T_5 of Step 5 is about $T_5 = 2^{126-4.49} = 2^{121.51}$ 21-round encryptions using the equivalent key property. T_5 dominates the total computational complexity.

Summarizing the attack, its computational complexity is about $2^{121.51}$, data complexity is about $2^{62.62}$ known plaintext-ciphertext pairs, and the memory complexity is negligible. The success probability is about 0.846.

5.2 Key Recovery for 25-Round XTEA

Similarly to the attack on 21 rounds of TEA provided in the previous subsection, we use the 14-round zero correlation linear approximation depicted in Figure 3 to attack 25-round XTEA. Note that the attack covers rounds 8 to 32. It is illustrated in Figure 6. The linear approximations are placed in rounds 14 to 27. We partially encrypt 6 rounds (8 to 13) and partially decrypt 5 rounds (28 to 32) to evaluate the parity of approximations.

The linear approximations involve 9 bits and in the corresponding 9 bits of the input and output masks, again only 7 can take on 0 and 1 values: $\Gamma_{28}^{R_0}$ and $\Gamma_{28}^{L_{5\sim 0}}$. For the evaluation of the linear approximations from a plaintext-ciphertext pair, we guess altogether 74 key bits $K_0^{25\sim 0}$, $K_1^{10\sim 0}$, $K_2^{10\sim 0}$, and $K_3^{25\sim 0}$. The attack itself is similar to that on 21-round TEA.

For each possible 63-bit value of $(K_0^{25\sim 0}|K_1^{10\sim 0}|K_3^{25\sim 0})$:

1. Allocate and set to zero the 32-bit counter $V_1[x]$ for each of 2^{30} possible values of

$$x = (R_{13}^0|R_{13}^5|L_{13}^0|R_{30}^{10\sim 0}|L_{30}^{15\sim 0}).$$

2. For each of N plaintext-ciphertext pairs: partially encrypt 5 rounds and partially decrypt 3 rounds, obtain 30-bit $x = (R_{13}^0|R_{13}^5|L_{13}^0|R_{30}^{10\sim 0}|L_{30}^{15\sim 0})$, and add one to $V_1[x]$.
3. For each possible 11 bits value of $K_2^{10\sim 0}$:
 - (a) Allocate and set to zero a 128-bit counter W .
 - (b) Allocate and set to zero a 64-bit counter $V_2[y]$ for each of 2^9 possible values of

$$y = (R_{14}^0|L_{28}^{5\sim 0}|R_{28}^{1\sim 0}).$$

- (c) Encrypt one round and decrypt two rounds for 2^{30} values for x to get 9 bits of y and add $V_1[x]$ to $V_2[y]$.
- (d) For each of 2^7 zero correlation linear approximations:
 - i. Set the 64-bit counter U to zero.
 - ii. For 2^9 values of y , verify if the linear approximation holds. If so, add $V_2[y]$ to the counter U .
 - iii. $W = W + (2 \cdot U/N - 1)^2$.
- (e) If $W < t$, then κ is a possible subkey candidate and all cipher keys it is compatible with are tested exhaustively against a maximum of 3 plaintext-ciphertext pairs.

The correct 74-bit subkey is likely to be among the candidates with the χ^2 statistic W lower than the threshold t . As we again set $\beta_0 = 2^{-2.7}$ and $\beta_1 = 2^{-4.49}$, we obtain $N = 2^{62.62}$ and $t = 2^{-55.56}$.

The computational complexity is dominated by Step 2 and checking for false positives in Step 3(e). T_2 of Step 2 is constituted by $2^{63}N$ computations of 5 rounds

of 25-round XTEA and by $2^{63}N$ increments in the memory of 2^{30} 32-bit counters. Assuming that one increment of a memory cell costs one XTEA round, we obtain $T_2 = 2^{63} \cdot 2^{62.62} \cdot (5/25 + 1/25) = 2^{123.56}$. In Step 3(e), the remaining $T_{3(e)} = 2^{128-4.49} = 2^{123.51}$ keys can be checked exhaustively by the same number of 25-round XTEA encryptions. Thus, the overall computational complexity is about $T_2 + T_{3(e)} = 2^{123.56} + 2^{123.51} = 2^{124.53}$ 25-round XTEA encryptions. The memory complexity is 2^{30} 32-bit words. Again, the data complexity is about $2^{62.62}$ known plaintext-ciphertext pairs, and the success probability is about 0.846.

5.3 Attacking More Rounds with the Full Codebook

The attacks in the previous subsections use 14-round zero correlation linear approximations to enable data complexity reduction. As we only identified 2 15-round approximations, we cannot use this longer property to attack more rounds and still get a non-negligible decrease in the number of texts required. By taking advantage of the full codebook, we are however able to perform key recovery for up to 23 rounds of TEA and up to 27 rounds of XTEA, see the full version [6] of this paper.

Acknowledgements. We would like to thank Vincent Rijmen and Gregor Leander for insightful discussions. Andrey Bogdanov is postdoctoral fellow of the Fund for Scientific Research - Flanders (FWO). This work has been supported in part by the IAP Programme P6/26 BCRYPT of the Belgian State, by the European Commission under contract number ICT-2007-216676 ECRYPT NoE phase II, by KU Leuven-BOF (OT/08/027), by the Research Council KU Leuven (GOA TENSE), by NSFC Projects (No.61133013, No.61070244 and No.60931160442) as well as Outstanding Young Scientists Foundation Grant of Shandong Province (No.BS2009DX030).

References

1. 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)
2. Biham, E., Dunkelman, O., Keller, N.: Related-Key Impossible Differential Attacks on 8-Round AES-192. In: Pointcheval, D. (ed.) CT-RSA 2006. LNCS, vol. 3860, pp. 21–33. Springer, Heidelberg (2006)
3. 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)
4. Bogdanov, A., Khovratovich, D., Rechberger, C.: Biclique Cryptanalysis of the Full AES. In: Lee, D.H. (ed.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 344–371. Springer, Heidelberg (2011)
5. Bogdanov, A., Rijmen, V.: Zero Correlation Linear Cryptanalysis of Block Ciphers. IACR Eprint Archive Report 2011/123 (March 2011)
6. Bogdanov, A., Wang, M.: Zero Correlation Linear Cryptanalysis with Reduced Data Complexity. IACR Eprint Archive Report (2012)
7. Borst, J., Knudsen, L.R., Rijmen, V.: Two Attacks on Reduced IDEA. In: Fumy, W. (ed.) EUROCRYPT 1997. LNCS, vol. 1233, pp. 1–13. Springer, Heidelberg (1997)

8. Bouillaguet, C., Dunkelman, O., Leurent, G., Fouque, P.-A.: Another Look at Complementation Properties. In: Hong, S., Iwata, T. (eds.) FSE 2010. LNCS, vol. 6147, pp. 347–364. Springer, Heidelberg (2010)
9. Chen, J., Wang, M., Preneel, B.: Impossible Differential Cryptanalysis of Lightweight Block Ciphers TEA, XTEA and HIGHT. IACR Eprint Archive Report 2011/616 (2011)
10. Collard, B., Standaert, F.-X.: Experimenting Linear Cryptanalysis. In: Junod, P., Canteaut, A. (eds.) Advanced Linear Cryptanalysis of Block and Stream Ciphers. Cryptology and Information Security Series, vol. 7. IOS Press (2011)
11. Collard, B., Standaert, F.-X., Quisquater, J.-J.: Improving the Time Complexity of Matsui’s Linear Cryptanalysis. In: Nam, K.-H., Rhee, G. (eds.) ICISC 2007. LNCS, vol. 4817, pp. 77–88. Springer, Heidelberg (2007)
12. Daemen, J., Govaerts, R., Vandewalle, J.: Correlation Matrices. In: Preneel, B. (ed.) FSE 1994. LNCS, vol. 1008, pp. 275–285. Springer, Heidelberg (1995)
13. Daemen, J., Rijmen, V.: Probability distributions of correlations and differentials in block ciphers. *Journal on Mathematical Cryptology* 1(3), 221–242 (2007)
14. Daemen, J., Rijmen, V.: The Design of Rijndael: AES – The Advanced Encryption Standard. Springer (2002)
15. Demirci, H., Selçuk, A.A.: A Meet-in-the-Middle Attack on 8-Round AES. In: Nyberg, K. (ed.) FSE 2008. LNCS, vol. 5086, pp. 116–126. Springer, Heidelberg (2008)
16. Dunkelman, O., Keller, N.: An Improved Impossible Differential Attack on MISTY1. In: Pieprzyk, J. (ed.) ASIACRYPT 2008. LNCS, vol. 5350, pp. 441–454. Springer, Heidelberg (2008)
17. Etrog, J., Robshaw, M.J.B.: On Unbiased Linear Approximations. In: Steinfeld, R., Hawkes, P. (eds.) ACISP 2010. LNCS, vol. 6168, pp. 74–86. Springer, Heidelberg (2010)
18. Feller, W.: An Introduction to Probability Theory and Its Applications, vol. 1. Wiley & Sons (1968)
19. Dunkelman, O., Keller, N., Shamir, A.: Improved Single-Key Attacks on 8-Round AES-192 and AES-256. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 158–176. Springer, Heidelberg (2010)
20. Hoel, P., Port, S., Stone, C.: Introduction to Probability Theory. Brooks Cole (1972)
21. Hong, S., Hong, D., Ko, Y., Chang, D., Lee, W., Lee, S.: Differential Cryptanalysis of TEA and XTEA. In: Lim, J.-I., Lee, D.-H. (eds.) ICISC 2003. LNCS, vol. 2971, pp. 402–417. Springer, Heidelberg (2004)
22. Junod, P.: On the Complexity of Matsui’s Attack. In: Vaudenay, S., Youssef, A.M. (eds.) SAC 2001. LNCS, vol. 2259, pp. 199–211. Springer, Heidelberg (2001)
23. Kaps, J.-P.: Chai-Tea, Cryptographic Hardware Implementations of xTEA. In: Chowdhury, D.R., Rijmen, V., Das, A. (eds.) INDOCRYPT 2008. LNCS, vol. 5365, pp. 363–375. Springer, Heidelberg (2008)
24. Kelsey, J., Schneier, B., Wagner, D.: Key-Schedule Cryptanalysis of IDEA, G-DES, GOST, SAFER, and Triple-DES. In: Kobitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 237–251. Springer, Heidelberg (1996)
25. Kelsey, J., Schneier, B., Wagner, D.: Related-key Cryptanalysis of 3-WAY, Biham-DES, CAST, DES-X, NewDES, RC2, and TEA. In: Han, Y., Quing, S. (eds.) ICISC 1997. LNCS, vol. 1334, pp. 233–246. Springer, Heidelberg (1997)
26. Lee, E., Hong, D., Chang, D., Hong, S., Lim, J.: A Weak Key Class of XTEA for a Related-Key Rectangle Attack. In: Nguyèn, P.Q. (ed.) VIETCRYPT 2006. LNCS, vol. 4341, pp. 286–297. Springer, Heidelberg (2006)

27. Lu, J.: Related-key rectangle attack on 36 rounds of the XTEA block cipher. *International Journal of Information Security* 8(1), 1–11 (2009)
28. Lu, J., Kim, J.-S., Keller, N., Dunkelman, O.: Improving the Efficiency of Impossible Differential Cryptanalysis of Reduced Camellia and MISTY1. In: Malkin, T. (ed.) *CT-RSA 2008*. LNCS, vol. 4964, pp. 370–386. Springer, Heidelberg (2008)
29. Lu, J., Dunkelman, O., Keller, N., Kim, J.: New Impossible Differential Attacks on AES. In: Chowdhury, D.R., Rijmen, V., Das, A. (eds.) *INDOCRYPT 2008*. LNCS, vol. 5365, pp. 279–293. Springer, Heidelberg (2008)
30. Mala, H., Dakhilalian, M., Rijmen, V., Modarres-Hashemi, M.: Improved Impossible Differential Cryptanalysis of 7-Round AES-128. In: Gong, G., Gupta, K.C. (eds.) *INDOCRYPT 2010*. LNCS, vol. 6498, pp. 282–291. Springer, Heidelberg (2010)
31. Matsui, M.: Linear Cryptanalysis Method for DES Cipher. In: Helleseth, T. (ed.) *EUROCRYPT 1993*. LNCS, vol. 765, pp. 386–397. Springer, Heidelberg (1994)
32. Matsui, M.: The First Experimental Cryptanalysis of the Data Encryption Standard. In: Desmedt, Y.G. (ed.) *CRYPTO 1994*. LNCS, vol. 839, pp. 1–11. Springer, Heidelberg (1994)
33. Moon, D., Hwang, K., Lee, W., Lee, S., Lim, J.: Impossible Differential Cryptanalysis of Reduced Round XTEA and TEA. In: Daemen, J., Rijmen, V. (eds.) *FSE 2002*. LNCS, vol. 2365, pp. 49–60. Springer, Heidelberg (2002)
34. Needham, R.M., Wheeler, D.J.: Tea extensions. Technical report, Computer Laboratory, University of Cambridge (October 1997), <http://www.cix.co.uk/~klockstone/xtea.pdf>
35. Nyberg, K.: Correlation theorems in cryptanalysis. *Discrete Applied Mathematics* 111(1-2), 177–188 (2001)
36. O’Connor, L.: Properties of Linear Approximation Tables. In: Preneel, B. (ed.) *FSE 1994*. LNCS, vol. 1008, pp. 131–136. Springer, Heidelberg (1995)
37. Röck, A., Nyberg, K.: Exploiting Linear Hull in Matsui’s Algorithm 1. In: *WCC 2011* (2011)
38. Sekar, G., Mouha, N., Velichkov, V., Preneel, B.: Meet-in-the-Middle Attacks on Reduced-Round XTEA. In: Kiayias, A. (ed.) *CT-RSA 2011*. LNCS, vol. 6558, pp. 250–267. Springer, Heidelberg (2011)
39. Selçuk, A.A.: On Probability of Success in Linear and Differential Cryptanalysis. *Journal of Cryptology* 21(1), 131–147 (2008)
40. Shirai, T., Shibutani, K., Akishita, T., Moriai, S., Iwata, T.: The 128-Bit Block-cipher CLEFIA (Extended Abstract). In: Biryukov, A. (ed.) *FSE 2007*. LNCS, vol. 4593, pp. 181–195. Springer, Heidelberg (2007)
41. Steil, M.: 17 Mistakes Microsoft Made in the Xbox Security System. *Chaos Communication Congress* (2005), <http://events.ccc.de/congress/2005/fahrplan/events/559.en.html>
42. Tsunoo, Y., Tsujihara, E., Shigeri, M., Saito, T., Suzaki, T., Kubo, H.: Impossible Differential Cryptanalysis of CLEFIA. In: Nyberg, K. (ed.) *FSE 2008*. LNCS, vol. 5086, pp. 398–411. Springer, Heidelberg (2008)
43. Vaudenay, S.: Decorrelation: A Theory for Block Cipher Security. *J. Cryptology* 16(4), 249–286 (2003)
44. Wheeler, D.J., Needham, R.M.: TEA, a Tiny Encryption Algorithm. In: Preneel, B. (ed.) *FSE 1994*. LNCS, vol. 1008, pp. 363–366. Springer, Heidelberg (1995)
45. Zheng, Y., Matsumoto, T., Imai, H.: On the Construction of Block Ciphers Provably Secure and Not Relying on Any Unproved Hypotheses. In: Brassard, G. (ed.) *CRYPTO 1989*. LNCS, vol. 435, pp. 461–480. Springer, Heidelberg (1990)

A Model for Structure Attacks, with Applications to PRESENT and Serpent

Meiqin Wang^{1,2,3,*}, Yue Sun⁴, Elmar Tischhauser^{2,3,**}, and Bart Preneel^{2,3}

¹ Key Laboratory of Cryptologic Technology and Information Security,
Ministry of Education, Shandong University, Jinan 250100, China

² Department of Electrical Engineering ESAT/SCD-COSIC,
Katholieke Universiteit Leuven, Kasteelpark Arenberg 10, B-3001 Heverlee, Belgium

³ Interdisciplinary Institute for BroadBand Technology (IBBT), Belgium

⁴ Institute for Advanced Study, Tsinghua University, Beijing 100084, China
mqwang@sdu.edu.cn

Abstract. As a classic cryptanalytic method for block ciphers, hash functions and stream ciphers, many extensions and refinements of differential cryptanalysis have been developed. In this paper, we focus on the use of so-called structures in differential attacks, *i.e.* the use of multiple input and one output difference. We give a general model and complexity analysis for structure attacks and show how to choose the set of differentials to minimize the time and data complexities. Being a subclass of multiple differential attacks in general, structure attacks can also be analyzed in the model of Blondeau *et al.* from FSE 2011. In this very general model, a restrictive condition on the set of input differences is required for the complexity analysis. We demonstrate that in our dedicated model for structure attacks, this condition can be relaxed, which allows us to consider a wider range of differentials. Finally, we point out an inconsistency in the FSE 2011 attack on 18 rounds of the block cipher PRESENT and use our model for structure attacks to attack 18-round PRESENT and improve the previous structure attacks on 7-round and 8-round Serpent. To the best of our knowledge, those attacks are the best known differential attacks on these two block ciphers.

Keywords: Structure Attack, Block Cipher, Differential, PRESENT, Serpent.

1 Introduction

Differential cryptanalysis [2] is a classic cryptanalytic method that has been successfully applied to block ciphers, hash functions and stream ciphers. The key

* This author is supported by 973 Project (No.2007CB807902), National Natural Science Foundation of China (Grant No.61133013, No.61070244 and No.60931160442), Outstanding Young Scientists Foundation Grant of Shandong Province (No.BS2009DX030), Shandong University Initiative Scientific Research Program (2009TS087)

** Elmar Tischhauser is a research assistant of the F.W.O., Fund for Scientific Research — Flanders.

step for a differential attack is to identify a differential characteristic with high probability as a distinguisher, then use it to recover (part of) the key. Lai *et al.* propose the notion of *differential* which encompasses the collection of all possible differential characteristics [13] for one fixed input and output difference. A lower bound for the probability of a differential (and thus, an upper bound for the complexity of the attack) can be obtained by combining the probabilities of a number of differential characteristics belonging to the differential. Therefore, differentials give a better estimation of the actual attack complexity than characteristics, since the distinguisher can exploit any characteristic belonging to the differential. In order to further improve differential attacks, multiple differentials with a single output difference but multiple input differences can be used. This can reduce the data complexity provided that the set of input differences for the differentials can be combined in a so-called *structure*. Therefore, we call this type of differential attacks *structure attacks*. The structure technique in differential cryptanalysis was originally introduced in a more restrictive way as *quartets* to attack DES [2], and multiple differential characteristics with multiple input differences and a single output difference have been used to attack DES. In addition, Biham *et al.* use the structure technique to attack reduced-round versions of the Serpent block cipher [3].

At FSE 2011, Blondeau *et al.* proposed multiple differential cryptanalysis with multiple input differences and multiple output differences [4] and gave an explicit formula to compute the success probability of multiple differential cryptanalysis. Traditionally, a normal approximation to the binomial distribution was used to evaluate the success probability of a differential attack [18,19]. The approach of [4] provides a more accurate estimation of the success probability. Since structure attacks are a special case of multiple differential cryptanalysis, those results also apply to our structure attacks.

However, in order to ensure that one pair of ciphertexts can be only counted once, the model of [4] requires a certain condition to be met (see Definition 1), which severely restricts the set of input difference values that can be used in an attack. In this paper, we demonstrate that this condition on the set of the input difference values is so strong that many valuable differentials may be excluded. We show that in the structure technique, this condition can be relaxed without counting ciphertexts more than once. This enables us to choose our differentials more freely, leading to improved attack complexities.

We stress that this condition and the general model of [4] are still necessary for the analysis of the general case where one has multiple input and multiple output differences. What we propose in this paper, is a tailored model for structure attacks, which are an important and often particularly efficient subclass of multiple differential cryptanalysis.

Furthermore, the multiple differential attack on 18-round PRESENT [4] uses 561 differentials with 17 input differences and 33 output differences [5]. It turns out that the sum of the probabilities of those 561 differentials is not correct in [4]. When calculated correctly, however, the obtained probability is lower than the random probability, implying that this set of 561 differentials cannot be used

in an attack. Finally, we compare our attack to the corrected version [6] of the attack of [4].

In order to evaluate the resistance of a block cipher to differential cryptanalysis, it is crucial to take into account the effect of combining multiple differentials. However, it is often not clear a priori which choice of differentials can actually lead to an improvement. Compared to classic differential cryptanalysis with one differential, a structure attack can obviously reduce the data complexity. In order to reduce the overall time complexity, however, the differentials have to be chosen carefully.

In this paper, we first present a general model for structure attacks, providing guidance on how to choose the differentials to minimize the time complexity. Secondly, we demonstrate structure attacks for 18-round PRESENT-80 with a data complexity of 2^{64} chosen plaintexts and time complexity of 2^{76} 18-round encryptions. We find that the properties of differentials in PRESENT cause structure attacks to be more efficient than the multiple differential cryptanalysis proposed in [4]. Thirdly, we improve the differential cryptanalytic result for the block cipher Serpent. In [3], Biham *et al.* describe a differential attack for 7-round Serpent with a data complexity of 2^{84} chosen plaintexts and a time complexity of 2^{85} memory accesses. Biham *et al.* also give a differential attack on 8-round Serpent-256 with 2^{213} memory accesses and 2^{84} chosen plaintexts. In our attack for 7-round Serpent, the data complexity is reduced to 2^{71} chosen plaintexts and the time complexity is $2^{74.99}$ encryptions. The attack can be further extended to 8-round Serpent-256. The time complexity is then increased to $2^{203.81}$ encryptions, with the data complexity remaining at 2^{71} chosen plaintexts.

For PRESENT-80, the best known attack is the linear hull cryptanalysis of 26-round PRESENT [8]. For Serpent-128, the best known cryptanalytic result is the differential-linear cryptanalysis on 12 rounds [12]. Although our attacks do not improve on those results for PRESENT and Serpent, to the best of our knowledge, they are the best *differential* attacks for PRESENT and Serpent. Moreover, our proposed attack model can be used to improve differential cryptanalytic results on other block ciphers as well.

This paper is organized as follows. Section 2 briefly describes the method for computing the success probability with multiple differentials. Section 3 introduces the structure attack model and the probability distribution of the key under multiple differentials. In Sect. 4, we demonstrate the attack for 18-round PRESENT. In Sect. 5, the improved attacks on 7-round and 8-round Serpent are presented. Section 6 concludes the paper.

2 Brief Description of Blondeau *et al.*'s Multiple Differential Cryptanalysis

In [4], Blondeau *et al.* propose multiple differential cryptanalysis using multiple differentials with different input differences and different output differences and give a precise analytical model to compute the success probability. In [18], Selçuk uses a Gaussian approximation of the binomial distribution to derive a formula

for the success probability for differential cryptanalysis. Since then, his formula has been used in many papers on differential cryptanalysis. Blondeau *et al.* demonstrate that Selçuk's method cannot be applied to multiple differential cryptanalysis and express the distribution of key counters instead in terms of a hybrid distribution including the Kullback-Leibler divergence and a Poisson distribution [4]. Blondeau *et al.* obtain the following formula for the success probability P_S :

$$P_S \approx 1 - G_*[G^{-1}(1 - \frac{l-1}{2^{n_k-2}}) - 1/N_s], \quad (1)$$

where n_k is the number of key candidates, l is the size of the list to keep, G is defined by $G^{-1}(y) = \min\{x|G(x) \geq y\}$, and N_s is the number of samples. Note that (1) corrects a typo in [4] by dividing by N_s for normalization. The functions G and G_* are defined as $G_*(\tau) \stackrel{\text{def}}{=} G(\tau, p_*)$ and $G(\tau) \stackrel{\text{def}}{=} G(\tau, p)$, where $p_* = \frac{\sum_{i,j} p_*^{(i,j)}}{|\Delta_0|}$ and $p = \frac{|\Delta|}{2^m |\Delta_0|} \cdot p_*^{(i,j)}$ is the probability for the differential with the i -th input difference value and the j -th output difference value, m is the block size, $|\Delta_0|$ is the number of input difference values and $|\Delta|$ is the number of differentials. $G(\tau, p_*)$ and $G(\tau, p)$ can be calculated with the following equations:

$$G(\tau, q) \stackrel{\text{def}}{=} \begin{cases} G_-(\tau, q) & \text{if } \tau < q - 3 \cdot \sqrt{q/N_s}, \\ 1 - G_+(\tau, q) & \text{if } \tau > q + 3 \cdot \sqrt{q/N_s}, \\ G_{\mathcal{P}}(\tau, q) & \text{otherwise,} \end{cases} \quad (2)$$

where $G_{\mathcal{P}}(\tau, q)$ is the cumulative distribution function of the Poisson distribution with parameter qN_s . $G_-(\tau, q)$ and $G_+(\tau, q)$ are defined as follows:

$$\begin{aligned} G_-(\tau, q) &\stackrel{\text{def}}{=} e^{-N_s D(\tau||q)} \cdot \left[\frac{q\sqrt{1-\tau}}{(q-\tau)\sqrt{2\pi\tau N_s}} + \frac{1}{\sqrt{8\pi\tau N_s}} \right], \\ G_+(\tau, q) &\stackrel{\text{def}}{=} e^{-N_s D(\tau||q)} \cdot \left[\frac{(1-q)\sqrt{\tau}}{(\tau-q)\sqrt{2\pi N_s(1-\tau)}} + \frac{1}{\sqrt{8\pi\tau N_s}} \right], \end{aligned} \quad (3)$$

where $D(\tau||q)$ is the Kullback-Leibler divergence defined by $D(\tau||q) \stackrel{\text{def}}{=} \tau \ln\left(\frac{\tau}{q}\right) + (1-\tau) \ln\left(\frac{1-\tau}{1-q}\right)$.

On the Assumptions for This Analysis. In order to guarantee that each pair is counted only once, Blondeau *et al.* give Definition 1 as a necessary condition for the set of the input differences Δ_0 .

Definition 1. *The set of input differences Δ_0 is admissible if there exists a set χ of $N/2$ plaintexts that fulfils the condition:*

$$\forall \delta_0^{(i)} \in \Delta_0, \forall x \in \chi, x \oplus \delta_0^{(i)} \notin \chi, \quad (4)$$

where N is the number of chosen plaintexts. However, this condition is so strong that many differentials will be excluded. For example, independent of the algorithm under consideration, the set of input differences $\Delta_0 = \{1_x, 2_x, 3_x\}$ is

never admissible in any substitution-permutation network (SPN) because of this condition, since the overlapping bits of $3_x = 1_x \oplus 2_x$ will always result in double-counting.

By contrast, in the structure technique, we can use a hash table to exclude the duplicate pair arising from the violation of Definition 1. In fact, making use of hash tables, structure attacks can use more differentials while still ensuring that each pair is counted only once. Since we only have one possible output difference, this also enables the use of the complexity analysis of [4] for sets of plaintexts not satisfying Def. 1: This condition is only necessary to avoid counting both x and $x \oplus \delta_0^{(i)}$ for any $\delta_0^{(i)} \in \Delta_0$, i.e. guarantee $N_s = N|\Delta_0|/2$. This is satisfied in our approach, since each hash table will produce $N/2$ plaintext pairs with one input difference from N plaintexts, in total therefore $N_s = |\Delta_0|N/2$ plaintext pairs with $|\Delta_0|$ input difference values. For structure attacks, the complexity analysis of [4] is therefore applicable independent of Def. 1.

This has additionally been verified by experiments on SMALLPRESENT with block length of 24 bits, 12 rounds, and a set of 11 differentials with input differences violating Definition 1 and a single output difference.

On Previous Attacks on 18-Round PRESENT. There are two previously published differential attacks on 18-round PRESENT [4,6]. In this section, we point out two inconsistencies in both attacks, and demonstrate that our attack compares favourably to them.

In [4], a multiple differential attack for 18-round PRESENT is presented. They identify 561 differentials¹ including 17 input differences and 33 output differences using a branch-and-bound algorithm. In [4], the probabilities p_* and p are calculated as $p_* = \frac{\sum_{i,j} p_*^{(i,j)}}{|\Delta_0|} = 2^{-58.50}$ and $p = \frac{|\Delta|}{2^m |\Delta_0|} = 2^{-64} \cdot 33 = 2^{-58.96}$. However, the value of p_* is not correct; it should be $p_* = 2^{-60.39}$, which is less than the random probability for 33 output differences $p = 2^{-58.96}$. We found that even when one chooses an optimal subset of these 561 differentials, this attack compares unfavourably to our structure attack.

In [6], another multiple differential attack on 18-round PRESENT is presented. It can be seen from Table 4 of [6], that $|\Delta_0| = 17$ (and not 16 as assumed in the paper). This results in $p_* = 2^{-62.6765}$ (instead of $2^{-62.59}$) and $p = 2^{-63.56}$ (instead of $p = 2^{-63.47}$). Based on these values, we compare this attack to our attack from Sect. 4 for different values of the number ℓ of remaining key candidates:

Attack of [6]		Attack of Sect. 4		N	time complexity
ℓ	P_S	ℓ	P_S		
2^{38}	65.27%	2^{36}	85.94%	2^{64}	2^{76}
2^{39}	79.68%	2^{37}	92.30%	2^{64}	2^{77}
2^{41}	94.62%	2^{39}	98.36%	2^{64}	2^{79}

¹ These differentials have been obtained through private communication with Blondeau *et al.*

One can see that for the same data and time complexities, the structure attack performs consistently better than multiple differential cryptanalysis with multiple input differences and multiple output differences. This implies that PRESENT is not a good example to show the efficiency of multiple differential cryptanalysis with different input differences and different output differences.

3 Structure Attack

3.1 Principle of the Attack

The structure attack is a form of differential cryptanalysis which uses multiple input differences and a single output difference. Structure attacks are a special case of multiple differential cryptanalysis, but their form allows for a dedicated attack procedure, which we describe in this section.

A structure attack is performed in three phases:

1. **Data Collection Phase:** Collect a large number of ciphertext pairs with the differences produced from the output difference of the differentials and the corresponding plaintext differences belong to the set of the input differences.
2. **Data Analysis Phase:** Derive the list of the best candidates for some key bits from the collected ciphertext pairs.
3. **Key Search Phase:** Search the list of candidates and all the corresponding master keys (*i.e.*, the unexpanded key from which the round subkeys are derived).

The idea of the structure attack is to use more differentials with multiple input differences and a single output difference to reduce the data complexity. However, the set of the input differences must be controlled in order to reduce the time complexity. This is done by organizing the plaintext in so-called *structures*:

Definition 2. Let $\{\Delta_0^1, \dots, \Delta_0^t\}$ be a set of t input differences. A collection of plaintexts of the form

$$\bigcup_x \{x \oplus \Delta \mid \Delta \in \text{span}\{\Delta_0^1, \dots, \Delta_0^t\}\}, \quad (5)$$

with span denoting the linear span operator, is called a structure.

In this way, we can construct structures to produce the expected number of right pairs with lower data complexity compared with a single differential. Now we will give a model to choose the differentials to reduce the complexity. For clarity of exposition, we describe the model for the case of a substitution-permutation network (SPN); however, the concept can analogously be applied to other block cipher constructions, most importantly Feistel ciphers.

If we attack an R -round block cipher with $|\Delta_0|$ r -round differentials with a single output difference and multiple input differences, we denote these differentials as follows:

$$\Delta_0^i \xrightarrow{r} \Delta_r, \text{ Probability} = p_i, (1 \leq i \leq |\Delta_0|),$$

where Δ_0^i and Δ_r are the i -th input difference and the output difference, respectively. The following notations are related with the attack:

- m : the block size of the block cipher.
- k : the key size of the block cipher.
- $|\Delta_0|$: the number of differentials.
- p_i : the probability of the differential with input difference Δ_0^i .
- N_{st} : the number of structures is $2^{N_{st}}$.
- N_p : the number of plaintexts bits involved in the active S-boxes in the first round for all differentials.
- N_c : the number of ciphertexts bits involved in the non-active S-boxes in the last round deriving from Δ_r .
- β : the filtering probability for the ciphertext pairs.
- p_f : the filtering probability for the ciphertext pairs according to active S-boxes, $p_f = \beta \cdot 2^{N_c}$.
- l : the size of the candidate list.
- n_k : the number of guessed subkey bits in the last $R - r$ rounds.

In the attack, $2^{N_{st}}$ structures are constructed. In each structure, all the input bits to non-active S-boxes in the first round are fixed to some random value, while N_p input bits of all active S-boxes take all 2^{N_p} possible values. There are $2^{N_{st}} \cdot 2^{N_p-1} = 2^{N_{st}+N_p-1}$ pairs for each differential. We expect that about $2^{N_{st}+N_p-1} \cdot \sum_{i=1}^{|\Delta_0|} p_i$ pairs produce the output difference Δ_r . These pairs are right pairs.

The attack is described as follows.

1. For each structure:
 - (a) Insert all the ciphertexts into a hash table indexed by N_c bits of the non-active S-boxes in the last round.
 - (b) For each entry with the same N_c bits value, check whether the input difference is any one of the total $|\Delta_0|$ possible input differences. If a pair satisfies one input difference, then go to the next step.
 - (c) For the pairs in each entry, check whether the output differences of active S-boxes in the last round can be caused by the input differences according to the differential distribution table. If the pair passes the test, then go to the next step.
 - (d) Guess n_k bits subkeys to decrypt the ciphertext pairs to round r and check whether the obtained output difference at round r is equal to Δ_r . If so, add one to the corresponding counter.
2. Choose the list of the l best key candidates from the counters.
3. Search the list of candidates and all the corresponding master key.

Obviously the time complexity in step 2 is negligible, so we denote T_a , T_b , T_c , T_d and T_3 as the time complexity in step (a), (b), (c), (d) and 3, respectively, which are listed in following:

$$\begin{cases} T_a : 2^{N_{st}+N_p} \text{ memory accesses;} \\ T_b : 2^{N_{st}+2N_p-N_c} \text{ memory accesses;} \\ T_c : |\Delta_0| \cdot 2^{N_{st}+N_p-N_c} \text{ memory accesses;} \\ T_d : |\Delta_0| \cdot 2^{N_{st}+N_p-N_c} \cdot p_f \cdot 2^{n_k} \text{ partial decryptions;} \\ T_3 : l \cdot 2^{k-n_k}. \end{cases}$$

This assumes that there are n_k independent subkey bits from the key schedule. In general, T_d can be approximated by $|\Delta_0| \cdot 2^{N_{st} + N_p - N_c} = T_c$. Since $|\Delta_0| < 2^{N_p}$, we have $T_c < T_b$. Then the whole time complexity can be expressed as follows:

$$T_a + T_b + T_c + T_d + T_3 \simeq \begin{cases} T_a + T_3 & \text{if } N_p < N_c, \\ T_b + T_3 & \text{if } N_p > N_c, \\ 2T_a + T_3 = 2T_b + T_3 & \text{if } N_p = N_c. \end{cases}$$

If the time complexity in the key searching process T_3 is much smaller than the time complexity of the data collection process and the data analysis process, we will take $N_p = N_c$ to minimise the whole time complexity as the minimum value $2T_a$. Otherwise, we can try to take a larger value for N_p to increase the sum of the probabilities for differentials to further reduce the data complexity.

It is worth noting that in the structure attack, any pair of plaintexts with the given input difference is only counted once. In this way, the number of input differences can be increased compared with the condition in Definition 1, improving the efficiency of the attacks. This is especially applicable in an attack scenario where the probability of many differentials are close to 2^{-m} , implying a low success rate P_S . Therefore, a large value for l has to be chosen, which causes the complexity T_3 of step 3 to increase. In this case, increasing the number of input differences can help improving the attack, whereas increasing the number of output differences does not have this effect in the case of multiple differential cryptanalysis.

In the case of reduced-round PRESENT, we have the above-mentioned scenario (many differentials with probability close to 2^{-64}), so that when choosing our set of differentials, we only include a limited number of high-probability differentials to maintain a good success probability P_S . For reduced-round Serpent, the probabilities of the differentials are much larger than 2^{-128} (the inverse of the block size), so that we can choose more differentials here without affecting the success probability. In order to minimize the time complexity, we choose $N_p = N_c$ according to our model.

3.2 Ratio of Weak Keys for Multiple Differentials

In general, the differential probability is related to the value of the key. As we use multiple differentials in the structure attack, we need to consider the ratio of keys which can produce the expected number of right pairs. We call those keys *weak keys* since the attacks are only expected to work for those.

A cipher is called *key-alternating* if it consists of an alternating sequence of unkeyed rounds and simple bitwise key additions. Note that most block cipher proposals, including PRESENT and Serpent, are key-alternating ciphers. The fixed-key cardinality of a differential $N[K](a, b)$ is the number of pairs with input difference a and output difference b where the key K is fixed to a specific value. In [11,10], Daemen and Rijmen give the following theorem.

Theorem 1. *Assuming that the set of pairs following a characteristic for a given key can be modeled by a sampling process, the fixed-key cardinality of a differential in a key-alternating cipher is a stochastic variable with the following distribution:*

$$\Pr(N[K](a, b) = i) \approx \text{Poisson} \left(i, 2^{m-1} EDP(a, b) \right),$$

where m is the block size, $EDP(a, b)$ denotes the expected differential probability of the differential (a, b) , and the distribution function measures the probability over all possible values of the key and all possible choices of the key schedule.

For multiple differentials with multiple input differences and a single output difference, we have $p_j = EDP(a_j, b), 1 \leq j \leq |\Delta_0|$. We denote the fixed-key cardinality of multiple differentials (a_j, b) with a single output difference b by $N[K]\{(a_j, b)\}_j$. Based on Theorem 1, we can now derive Theorem 2.

Theorem 2. *Under the assumptions of Theorem 1, in a key-alternating cipher, the fixed-key cardinality of multiple differentials is a stochastic variable with the following distribution:*

$$\Pr \left(N[K]\{(a_j, b)\}_j = i \right) \approx \text{Poisson} \left(i, 2^{m-1} \sum_j EDP(a_j, b) \right).$$

Proof. The cardinality of multiple differentials equals the sum of the cardinalities of each differential (a_j, b) for the iterative cipher, so we have

$$N[K]\{(a_j, b)\}_j = \sum_j N[K](a_j, b).$$

From Theorem 1, the cardinality for each differential (a_j, b) has Poisson distribution. Making the standard assumption that the cardinalities of the differentials are independent random variables, the sum still is Poisson distributed with as λ -parameter the sum of the λ -parameters of the terms:

$$\lambda = \sum_j 2^{m-1} EDP(a_j, b). \quad \square$$

From Theorem 2, in the structure attack based on the differentials $\Delta_0^i \xrightarrow{r} \Delta_r$, $Probability = p_i, (1 \leq i \leq |\Delta_0|)$, the ratio of the weak keys r_w that can produce more than or equal to μ right pairs can be computed as follows:

$$r_w = 1 - \sum_{x=0}^{\mu-1} \text{Poisson} \left(x, 2^{m-1} \sum_{j=1}^{|\Delta_0|} p_i \right).$$

Note that when evaluating the ratio of weak keys, we have a different setting than when dealing with the distribution of the counters in a (multiple) differential attack. While approximating the distribution of the counters with either

normal or Poisson distributions was shown to be problematic for accurately estimating the tails [19,4], the distribution of the weak keys instead depends on the *cardinality* of the multiple differentials. In this setting, using the Poisson distribution as in Theorem 2 also yields a good approximation for the tails. This was also experimentally verified with small-scale variants of the block cipher PRESENT [14], with block lengths ranging from 8 to 24 bits.

Additionally, the accuracy of the weak key ratio r_w based on Theorem 2 has been verified by experiments on SMALLPRESENT with a block length of 24 bits, 12 rounds and an master key with 8 bit entropy. 7 differentials with 7 different input and a single output difference were used. It was found that the experimental results very closely follow the theoretical estimate.

4 Attack on 18-Round PRESENT

The block cipher PRESENT is designed as a very lightweight cipher. It has a 31-round SPN structure in which the S-box layer has 16 parallel 4-bit S-boxes and the diffusion layer is a bit permutation [7]. The block size is 64 bits and the key size can be 80 bits or 128 bits. One round of PRESENT is illustrated in Fig. 1.




Fig. 1. One round of the PRESENT block cipher

PRESENT has been extensively analyzed. Wang presents a differential attack on 16-round PRESENT [20]. Collard *et al.* give a statistical saturation attack for 24-round PRESENT [9]. There are three papers about attacks based on linear hulls for PRESENT [8,17,16], leading to linear attacks for up to 26 rounds. Since the S-box of PRESENT admits linear approximations with single-bit linear masks, the attacker can exploit linear hulls containing many single-bit linear trails over an arbitrary number of rounds. However, for differential attacks, we have to use paths in which two active S-boxes appear per round. Hence, a linear attack will typically be more efficient than differential attacks.

In order to identify a differential with high probability, we must collect more differential paths with high probability for a differential. The differential paths with two active S-boxes in every round have a much bigger contribution to the differential, so we will focus on differential paths with only two active S-boxes in each round. Then we can choose more differentials to improve the attack according to the formulas for the overall time complexity described in Sect. 3, .

4.1 Searching Differential Paths for PRESENT

We now give a method to search all differential characteristics with two active S-boxes in each round which have higher probability compared with other differential paths.

First, we introduce some notation. The block size of PRESENT is 64 bits and we can divide 16 nibbles into four groups, in each of which there are four nibbles. We define G as the index of a group, so the four least significant nibbles belong to the group $G = 0$ and the four most significant nibbles belong to the group $G = 3$. Analogously, we denote the index of a nibble in a group as N , $N = 0, \dots, 3$, and B as the B -th bit in a nibble, from $B = 0$ to $B = 3$. In this way, the position of any bit can be denoted by a triple (G, N, B) , as also illustrated in Fig. 1. The permutation layer P is computed as follows,

$$P(16 \cdot G + 4 \cdot N + B) = 16 \cdot B + 4 \cdot G + N, 0 \leq G, N, B \leq 3.$$

After the permutation layer P , the bit (G, N, B) will be transferred to the bit (B, G, N) . Here we also give another triple (G, N, V) where G and N are the group index and nibble index, respectively, while V is the difference of the nibble. We will write $(G_{r,k}, N_{r,k}, B_{r,k})$ for the position of the k -th ($k = 1, 2, 3, 4$) output bit for S-box in round r , and $(G_{r,k}, N_{r,k}, V_{r,k})$ for the output difference value of the k -th ($k = 1, 2$) active S-box for nibble $(G_{r,k}, N_{r,k})$ in round r .

We focus on finding differential characteristics with two active S-boxes in each round. The foundation for this search is formulated in Theorem 3.

Theorem 3. *For the PRESENT block cipher, differential characteristics with only two active S-boxes per round must have the following pattern:*

1. *If two active S-boxes are in the same group in round r , their output difference will be equal and must have two non-zero bits to ensure that only two active S-boxes appear in the $(r+2)$ -nd round, and two active S-boxes in round $r+1$ will be in the different groups;*
2. *If two active S-boxes are in different groups in round r , their output difference will be equal and must have only one non-zero bit to ensure that only two active S-boxes appear in the $(r+1)$ -st round, and two active S-boxes in round $r+1$ will be in the same group.*

Proof. The output differences for the two active S-boxes are $(G_{r,1}, N_{r,1}, V_{r,1})$ and $(G_{r,2}, N_{r,2}, V_{r,2})$. First, we will prove the case for two active S-boxes in the same group in round r . We have $G_{r,1} = G_{r,2}$ and $N_{r,1} \neq N_{r,2}$.

1. $V_{r,1} \in \{1, 2, 4, 8\}$: If $V_{r,2} \in \{1, 2, 4, 8\}$, we denote their two non-zero bits as $\{(G_{r,1}, N_{r,1}, B_{r,1}), (G_{r,1}, N_{r,2}, B_{r,2})\}$.

We have

$$\begin{aligned} & \{(G_{r,1}, N_{r,1}, B_{r,1}), (G_{r,1}, N_{r,2}, B_{r,2})\} \xrightarrow{P} \{(B_{r,1}, G_{r,1}, N_{r,1}), (B_{r,2}, G_{r,1}, N_{r,2})\} \xrightarrow{S} \\ & \{(B_{r,1}, G_{r,1}, N_{r+1,1}), (B_{r,1}, G_{r,1}, N_{r+1,2}), (B_{r,2}, G_{r,1}, N_{r+1,3}), (B_{r,2}, G_{r,1}, N_{r+1,4})\} \\ & \xrightarrow{P} \{(N_{r+1,1}, B_{r,1}, G_{r,1}), (N_{r+1,2}, B_{r,1}, G_{r,1}), (N_{r+1,3}, B_{r,2}, G_{r,1}), (N_{r+1,4}, B_{r,2}, \\ & G_{r,1})\}. \end{aligned}$$

As there are two active S-boxes in round $r+1$, we have $B_{r,1} \neq B_{r,2}$. Because bit $N_{r+1,1}$ and bit $N_{r+1,2}$ are from the same S-box, we have $N_{r+1,1} \neq N_{r+1,2}$.

Similarly, we have $N_{r+1,3} \neq N_{r+1,4}$. There will be four active S-boxes in the $(r+2)$ -nd round. If $V_{r,2} \in \{3, 5, 6, 9, 10, 12\}$, we denote the three non-zero bits as $\{(G_{r,1}, N_{r,1}, B_{r,1}), (G_{r,1}, N_{r,2}, B_{r,2}), (G_{r,1}, N_{r,2}, B_{r,3})\}$.

We have

$$\begin{aligned} & \{(G_{r,1}, N_{r,1}, B_{r,1}), (G_{r,1}, N_{r,2}, B_{r,2}), (G_{r,1}, N_{r,2}, B_{r,3})\} \\ & \xrightarrow{P} \{(B_{r,1}, G_{r,1}, N_{r,1}), (B_{r,2}, G_{r,1}, N_{r,2}), (B_{r,3}, G_{r,1}, N_{r,2}) | B_{r,1} = B_{r,2} \neq B_{r,3}\} \\ & \xrightarrow{S} \{(B_{r,1}, G_{r,1}, N_{r+1,1}), (B_{r,3}, G_{r,1}, N_{r+1,2}), (B_{r,3}, G_{r,1}, N_{r+1,3}) | N_{r+1,2} \neq N_{r+1,3}\} \\ & \xrightarrow{P} \{(N_{r+1,1}, B_{r,1}, G_{r,1}), (N_{r+1,2}, B_{r,3}, G_{r,1}), (N_{r+1,3}, B_{r,3}, G_{r,1})\}. \end{aligned}$$

There will be three active S-boxes in round $r+2$.

2. $V_{r,1} \in \{7, 11, 13, 14, 15\}$ **or** $V_{r,2} \in \{7, 11, 13, 14, 15\}$: There will be at least three active S-boxes in round $r+1$.
3. $V_{r,1}, V_{r,2} \in \{3, 5, 6, 9, 10, 12\}$: We denote the four non-zero bits as $\{(G_{r,1}, N_{r,1}, B_{r,1}), (G_{r,1}, N_{r,1}, B_{r,2}), (G_{r,1}, N_{r,2}, B_{r,3}), (G_{r,1}, N_{r,2}, B_{r,4})\}$.

We have

$$\begin{aligned} & \{(G_{r,1}, N_{r,1}, B_{r,1}), (G_{r,1}, N_{r,1}, B_{r,2}), (G_{r,1}, N_{r,2}, B_{r,3}), (G_{r,1}, N_{r,2}, B_{r,4})\} \\ & \xrightarrow{P} \{(B_{r,1}, G_{r,1}, N_{r,1}), (B_{r,2}, G_{r,1}, N_{r,1}), (B_{r,3}, G_{r,1}, N_{r,2}), (B_{r,4}, G_{r,1}, N_{r,2})\}. \end{aligned}$$

Only if $B_{r,1} = B_{r,3}$ and $B_{r,2} = B_{r,4}$, there will be 2 active S-boxes in round $r+1$, so we have $V_{r,1} = V_{r,2}$. For $B_{r,1} \neq B_{r,2}$, the two active S-boxes in round $r+1$ will be in different groups.

Next, we will prove the case for two active S-boxes in different groups in round r . We have $G_{r,1} \neq G_{r,2}$.

1. $V_{r,1} \in \{7, 11, 13, 14, 15\}$ **or** $V_{r,2} \in \{7, 11, 13, 14, 15\}$: There will be at least three active S-boxes in round $r+1$.
2. $V_{r,1} \in \{3, 5, 6, 9, 10, 12\}$: There are at least three non-zero bits, namely $(G_{r,1}, N_{r,1}, B_{r,1}), (G_{r,1}, N_{r,1}, B_{r,2})$ and $(G_{r,2}, N_{r,2}, B_{r,3})$.

We have

$$\begin{aligned} & \{(G_{r,1}, N_{r,1}, B_{r,1}), (G_{r,1}, N_{r,1}, B_{r,2}), (G_{r,2}, N_{r,2}, B_{r,3})\} \\ & \xrightarrow{P} \{(B_{r,1}, G_{r,1}, N_{r,1}), (B_{r,2}, G_{r,1}, N_{r,1}), (B_{r,3}, G_{r,2}, N_{r,2})\}. \end{aligned}$$

For $B_{r,1} \neq B_{r,2}$ and $G_{r,1} \neq G_{r,2}$, there are three active S-boxes in round $r+1$.

3. $V_{r,1} \in \{1, 2, 4, 8\}$: From the above proof, we have $V_{r,2} \in \{1, 2, 4, 8\}$. There are two non-zero bits $\{(G_{r,1}, N_{r,1}, B_{r,1}), (G_{r,2}, N_{r,2}, B_{r,2})\}$. We have

$$\begin{aligned} & \{(G_{r,1}, N_{r,1}, B_{r,1}), (G_{r,2}, N_{r,2}, B_{r,2})\} \xrightarrow{P} \{(B_{r,1}, G_{r,1}, N_{r,1}), (B_{r,2}, G_{r,2}, N_{r,2})\} \\ & \xrightarrow{S} \{(B_{r,1}, G_{r,1}, N_{r+1,1}), (B_{r,1}, G_{r,1}, N_{r+1,2}), (B_{r,2}, G_{r,2}, N_{r+1,3}), (B_{r,2}, G_{r,2}, N_{r+1,4})\} \\ & \xrightarrow{P} \{(N_{r+1,1}, B_{r,1}, G_{r,1}), (N_{r+1,2}, B_{r,1}, G_{r,1}), (N_{r+1,3}, B_{r,2}, G_{r,2}), (N_{r+1,4}, B_{r,2}, G_{r,2})\}. \end{aligned}$$

In order to ensure that there are two active S-boxes in round $r+2$, $N_{r+1,1} = N_{r+1,3}$, $N_{r+1,2} = N_{r+1,4}$ and $B_{r,1} = B_{r,2}$. So we have $V_{r,1} = V_{r,2}$ and the two active S-boxes in round $r+1$ are in the same group. \square

Based on Theorem 3, a branch-and-bound search algorithm for differential paths can be devised.

Using this algorithm, we search for 16-round differential paths (characteristics) with two active S-boxes in each round having a probability greater than 2^{-92} . In total, we find 139 *differentials* with probability greater than 2^{-64} , among

Table 1. Filter probability for the structure attack on 18-round PRESENT

N_a	$(Y_{17,2}, Y_{17,10})$	$p_f^{(a)}$
2	$\{(1, 1), (1, 4), (4, 1), (4, 4)\}$	$2^{-24} \cdot (\frac{7}{16})^2 \cdot 4 = 2^{-24.83}$
3	$\{(1, 9), (1, 10), (1, 12), (4, 9), (4, 10), (4, 12), (9, 1), (9, 4), (10, 1), (10, 4), (12, 1), (12, 4)\}$	$2^{-20} \cdot (\frac{7}{16})^3 \cdot 12 = 2^{-19.99}$
4	$\{(9, 9), (9, 10), (9, 12), (10, 9), (10, 10), (10, 12), (12, 9), (12, 10), (12, 12), (1, 11), (1, 13), (4, 11), (4, 13), (11, 1), (11, 4), (13, 1), (13, 4)\}$	$2^{-16} \cdot (\frac{7}{16})^4 \cdot 17 = 2^{-16.68}$
5	$\{(9, 11), (9, 13), (10, 11), (10, 13), (12, 11), (12, 13), (11, 9), (11, 10), (11, 12), (13, 9), (13, 10), (13, 12)\}$	$2^{-12} \cdot (\frac{7}{16})^5 \cdot 12 = 2^{-14.38}$
6	$\{(11, 11), (11, 13), (13, 11), (13, 13)\}$	$2^{-8} \cdot (\frac{7}{16})^6 \cdot 4 = 2^{-13.16}$

Table 2. Differentials for 16-round PRESENT with $\Delta_{16} = 00000900_x || 00000900_x$

i	Δ_0^i	$\log_2^{p_i}$	i	Δ_0^i	$\log_2^{p_i}$
1	000f0000 _x 0000000f _x	-62.98	10	000f0000 _x 00000005 _x	-63.98
2	00070000 _x 00000007 _x	-63.42	11	000f0000 _x 0000000b _x	-63.98
3	0f000000 _x 00000f00 _x	-63.68	12	000f0000 _x 0000000d _x	-63.98
4	000f0000 _x 00000007 _x	-63.69	13	00030000 _x 0000000f _x	-63.98
5	00070000 _x 0000000f _x	-63.69	14	00050000 _x 0000000f _x	-63.98
6	000d0000 _x 0000000d _x	-63.72	15	000b0000 _x 0000000f _x	-63.98
7	00f00000 _x 000000f0 _x	-63.92	16	000d0000 _x 0000000f _x	-63.98
8	00090000 _x 00000009 _x	-63.94	17	f0000000 _x 0000000f _x	-63.98
9	000f0000 _x 00000003 _x	-63.98	18	000f0000 _x 0000f000 _x	-63.98

which 91 differentials have output difference $\Delta_{16} = 00000500_x || 00000500_x$ and 18 differentials have output difference $\Delta_{16} = 00000900_x || 00000900_x$. We list them in Table 3 and Table 2, respectively. The differentials have been ordered according to their probabilities in these two tables. In both Table 3 and Table 2, the first column i contains the number of the differential, Δ_0^i is the input difference and p_i is the probability for each differential. Moreover, we present the number of differential paths ordered by probability for Table 3 in Table 4 and Table 5. In Table 4, the first column denotes the index number in the first column of Table 3. For example, the differentials with number 19 and 20 consist of differential trails with the same probabilities. Columns 2, 3, ..., 12 denote the number of differential paths with probability $2^{-71}, 2^{-73}, \dots, 2^{-91}$, respectively. In Table 5, the first column denotes the index number in the first column of Table 3. Column 2, 3, ..., 13 denote the number of differential paths with probability $2^{-70}, 2^{-72}, \dots, 2^{-92}$, respectively. There is no differential path with probability greater than 2^{-70} or less than 2^{-92} for the 91 differentials.

Table 3. Differentials for 16-round PRESENT with output difference $00000500_x || 00000500_x$

i	Δ_0^i	$\log_2^{p_i}$	i	Δ_0^i	$\log_2^{p_i}$
1	000f0000 _x 0000000f _x	-62.13	47	000f0000 _x 00000f00 _x	-63.79
2	00070000 _x 00000007 _x	-62.57	48	0f000000 _x 0000000f _x	-63.79
3	0f000000 _x 00000f00 _x	-62.79	49	0f000000 _x 00000d00 _x	-63.79
4	000f0000 _x 00000007 _x	-62.84	50	0f000000 _x 00000b00 _x	-63.79
5	00070000 _x 0000000f _x	-62.84	51	0f000000 _x 00000300 _x	-63.79
6	000d0000 _x 0000000d _x	-62.88	52	0f000000 _x 00000500 _x	-63.79
7	00f00000 _x 00000f0 _x	-62.95	53	03000000 _x 00000f00 _x	-63.79
8	00090000 _x 00000009 _x	-63.10	54	05000000 _x 00000f00 _x	-63.79
9	000f0000 _x 00000003 _x	-63.13	55	0d000000 _x 00000f00 _x	-63.79
10	000f0000 _x 00000005 _x	-63.13	56	0b000000 _x 00000f00 _x	-63.79
11	000f0000 _x 0000000b _x	-63.13	57	00070000 _x 00000003 _x	-63.84
12	000f0000 _x 0000000d _x	-63.13	58	00070000 _x 00000005 _x	-63.84
13	00030000 _x 0000000f _x	-63.13	59	00030000 _x 00000007 _x	-63.84
14	00050000 _x 0000000f _x	-63.13	60	00050000 _x 00000007 _x	-63.84
15	000b0000 _x 0000000f _x	-63.13	61	f0000000 _x 00000007 _x	-63.84
16	000d0000 _x 0000000f _x	-63.13	62	70000000 _x 0000000f _x	-63.84
17	f0000000 _x 0000000f _x	-63.13	63	000f0000 _x 00007000 _x	-63.84
18	000f0000 _x 0000f000 _x	-63.13	64	00070000 _x 0000f000 _x	-63.84
19	000d0000 _x 00000007 _x	-63.19	65	0d000000 _x 00000700 _x	-63.85
20	00070000 _x 0000000d _x	-63.19	66	07000000 _x 00000d00 _x	-63.85
21	0f000000 _x 000000f0 _x	-63.21	67	00000f00 _x 00000f00 _x	-63.87
22	0f000000 _x 00000f00 _x	-63.21	68	00000000 _x 0f000f00 _x	-63.87
23	00000000 _x 000f000f _x	-63.21	69	d0000000 _x 0000000d _x	-63.88
24	0000000f _x 0000000f _x	-63.21	70	000d0000 _x 0000d000 _x	-63.88
25	07000000 _x 00000700 _x	-63.23	71	00000000 _x 000f0007 _x	-63.91
26	00700000 _x 00000070 _x	-63.39	72	00000000 _x 0007000f _x	-63.91
27	000b0000 _x 0000000b _x	-63.44	73	0000000f _x 00000007 _x	-63.91
28	000f0000 _x 00000009 _x	-63.50	74	00000007 _x 0000000f _x	-63.91
29	00090000 _x 0000000f _x	-63.50	75	00900000 _x 00000090 _x	-63.92
30	0f000000 _x 00000700 _x	-63.50	76	0f000000 _x 00000070 _x	-63.92
31	07000000 _x 00000f00 _x	-63.50	77	07000000 _x 000000f0 _x	-63.92
32	000b0000 _x 00000007 _x	-63.52	78	00f00000 _x 00000700 _x	-63.92
33	00070000 _x 0000000b _x	-63.52	79	00700000 _x 00000f00 _x	-63.92
34	0d000000 _x 00000d00 _x	-63.54	80	00f00000 _x 00000030 _x	-63.95
35	70000000 _x 00000007 _x	-63.57	81	00f00000 _x 00000050 _x	-63.95
36	00070000 _x 00007000 _x	-63.57	82	00f00000 _x 000000b0 _x	-63.95
37	000d0000 _x 00000009 _x	-63.58	83	00f00000 _x 000000d0 _x	-63.95
38	00090000 _x 0000000d _x	-63.58	84	00300000 _x 000000f0 _x	-63.95
39	00000000 _x 00070007 _x	-63.64	85	00500000 _x 000000f0 _x	-63.95
40	00000007 _x 00000007 _x	-63.64	86	00b00000 _x 000000f0 _x	-63.95
41	07000000 _x 00000070 _x	-63.65	87	00d00000 _x 000000f0 _x	-63.95
42	00700000 _x 00000700 _x	-63.65	88	0d000000 _x 000000d0 _x	-63.95
43	00700000 _x 000000f0 _x	-63.66	89	00d00000 _x 00000d00 _x	-63.95
44	00f00000 _x 00000070 _x	-63.66	90	00000000 _x 000d000d _x	-63.95
45	00d00000 _x 000000d0 _x	-63.70	91	0000000d _x 0000000d _x	-63.95
46	09000000 _x 00000900 _x	-63.76			

Table 4. Number of differential paths for differentials in Table 3 (first part)

i	2^{-71}	2^{-73}	2^{-75}	2^{-77}	2^{-79}	2^{-81}	2^{-83}	2^{-85}	2^{-87}	2^{-89}	2^{-91}
9,10,...,18	12	160	986	3744	9654	17440	21988	18536	9280	1920	0
19,20	12	157	952	3567	9092	16264	20348	17068	8520	1760	0
32,33	9	123	769	2913	7350	12692	14780	10980	4600	800	0
35,36	9	117	707	2669	7056	13858	20936	24568	21248	11520	2560
37,38	6	89	628	2795	8562	18504	27976	28004	16200	3680	0
47,48	8	104	628	2348	5976	10676	13340	11160	5568	1152	0
49,50,...,56	4	64	486	2336	7838	19064	33976	43600	38368	20736	4608
57,58,...,64	9	114	655	2258	5092	7600	7180	3800	800	0	0
65,66	4	63	472	2243	7448	17942	31704	40376	35344	19040	4224
69,70	3	55	457	2295	7744	18318	30608	35268	26256	11040	1920
80,81,...,87	4	60	438	2066	6886	16766	30064	38908	34584	18880	4224

Table 5. Number of differential paths for differentials in Table 3 (second part)

i	2^{-70}	2^{-72}	2^{-74}	2^{-76}	2^{-78}	2^{-80}	2^{-82}	2^{-84}	2^{-86}	2^{-88}	2^{-90}	2^{-92}
1	12	160	986	3744	9654	17440	21988	18536	9280	1920	0	0
2	9	117	707	2669	7056	13858	20936	24568	21248	11520	2560	0
3	4	64	486	2336	7838	19064	33976	43600	38368	20736	4608	0
4,5	9	114	655	2258	5092	7600	7180	3800	800	0	0	0
6	3	55	457	2295	7744	18318	30608	35256	26256	11040	1920	0
7	4	60	438	2066	6886	16766	30064	38908	34584	18880	4224	0
8	3	49	383	1897	6526	16098	28564	35504	28928	13440	2560	0
21,22	4	56	382	1708	5490	13088	23300	30260	27208	15168	3456	0
23,24	0	48	472	2112	5724	10404	13104	11336	6400	1920	0	0
25	3	47	351	1673	5650	14212	27472	41472	48928	43520	25600	6144
26	3	44	316	1480	4971	12516	24286	36824	43656	39168	23296	5632
27	0	21	274	1641	6002	14746	25040	29168	22336	10080	1920	0
28,29,30,31	3	46	331	1486	4562	9840	14808	14736	8480	1920	0	0
34	1	21	205	1243	5222	15940	35960	59616	70464	55488	24960	4608
39,40	0	36	342	1496	4090	8128	12572	14936	12928	7680	2560	0
41,42	3	41	275	1223	3976	9836	18950	28680	34008	30720	18688	4608
43,44	3	43	297	1309	4000	8664	13168	13268	7720	1760	0	0
45	1	20	188	1112	4609	14004	31658	52832	63048	50160	22752	4224
46	1	19	175	1037	4364	13596	31832	55600	70336	60416	30208	6144
67,68	0	16	200	1184	4420	11276	20280	26080	23392	13824	4608	0
71,72,73,74	0	36	330	1330	3072	4480	4280	2600	800	0	0	0
75	1	18	160	928	3857	11954	28014	49196	62800	54528	27520	5632
76,77,78,79	3	40	257	1070	3152	6706	10188	10412	6200	1440	0	0
88,89	1	19	169	949	3768	11100	24650	40920	49128	39696	18336	3456
90,91	0	12	178	1160	4430	10944	18260	20952	16416	8160	1920	0

4.2 Key Recovery Attack on 18-Round PRESENT-80

In this section, we show how to use the 16-round differentials listed in Table 3 to attack 18-round PRESENT-80. The first step is to choose the set of differentials. From the output difference $00000500_x || 00000500_x$ at round 16, we can derive that the number of recovered subkey bits in round 17 and round 18 is $8 + 32 = 40$. Those 40 subkey bits are independent according to the key schedule. In this attack, we will use the whole codebook and set the size of the candidates of subkey counters l to 2^{36} . In our structure attack, we will use Blondeau *et al.*'s method (see Sect. 2) to compute the success rate. With Equation (1), we have $n_k = 40$, $l = 2^{36}$ and $N = 2^{64}$. We gradually increase the number of differentials with higher probability from Table 3 to compute the success probability for every case. As a result, we found that the success rate will increase as $|\Delta_0| = i$ increases if $1 \leq i \leq 36$. The success probability is 85.95% as $|\Delta_0| = 36$. If we add the i -th ($37 \leq i \leq 91$) differential to the set, the success probability will be reduced. This implies that the i -th ($37 \leq i \leq 91$) differential has no contribution to reduce the data complexity since its probability is too low. Therefore, in our attack, we will only use the first 36 differentials in Table 3.

If we use multiple differentials cryptanalysis for PRESENT following Blondeau *et al.*, we can choose more output difference values. We can add the 18 differentials in Table 2 to the set of 36 differentials. The input difference values for the 18 differentials belong to the set of the input difference values for the 36 differentials, so we have $|\Delta_0| = 36$ and $|\Delta_{16}| = 2$. Then we get $p_* = 2^{-62.74}$ and $p = 2^{-63}$. As τ ($p < \tau < p_*$) increases, $G(\tau, p)$ will decrease. Even if we take $\tau = p_*$, $G(\tau, p)$ is still larger than $(1 - \frac{l-1}{2^{n_k}-2})$, so the attack will not work for $l = 2^{36}$. Therefore, our structure attack works better for PRESENT than the multiple differential cryptanalysis presented in [4].

Moreover, we have identified the differential trails with two active S-boxes per round but more than two active S-boxes in the last round. As a result, those differentials have no advantage compared with the differentials in Table 3. Therefore, these differentials do not contribute to improving multiple differential cryptanalysis for PRESENT.

We will use the structure attack for 18-round PRESENT-80 with the first 36 differentials with $p_* = 2^{-63.14}$ and $p = 2^{-64}$. For the 36 input differences, there are 10 active S-boxes in the first round which are nibbles 0, 1, 2, 3, 4, 8, 12, 13, 14 and 15, so the S-boxes for the nibbles 5, 6, 7, 9, 10 and 11 are all non-active.

We construct 2^{24} structures of 2^{40} chosen plaintexts each. In each structure, all the inputs to the 6 non-active S-boxes in the first round take a fixed random value, while 40 bits of input to 10 active S-boxes take 2^{40} possible values. In all structures, there are $2^{24} \cdot 2^{39} = 2^{63}$ pairs for each possible differential. The sum of the probabilities for all 36 differentials is $2^{-57.97}$, so the number of right pairs is $2^{63} \cdot 2^{-57.97} = 2^{5.03}$.

According to the output difference of 16-round differentials, there are two active S-boxes in round 17 in nibble 2 and 10 whose input difference is 5 and the possible output differences will be 1, 4, 9, 10, 11, 12 or 13. After the bit permutation, 8 output bits from the two active S-boxes in round 17 will be one

input bit to 8 different S-boxes in round 18 respectively. As the number of non-zero bits among the 8 output bits is at most 6, the maximum number of active S-boxes for round 18 is 6 and the minimum number of active S-boxes for round 18 is 2. We denote the number of active S-boxes in round 18 as N_a ($2 \leq N_a \leq 6$), the output difference for the j -th S-box in round i as $Y_{i,j}$, the filter probability with N_a active S-boxes in round 18 as $p_f^{(a)}$. We present the filter probability for different values of N_a in Table 1. The filter probability for the ciphertext pairs β according to active S-boxes can be computed with the sum of column 3 in Table 1, and we get $\beta = 2^{-12.55}$.

We now describe in detail the attack procedure of Sect. 3 for 18-round PRESENT-80. We have $|\Delta_0| = 36$, $\sum_{i=1}^{|\Delta_0|} p_i = 2^{-57.97}$, $N_{st} = 24$, $N_p = 40$, $N_c = 32$, $\beta = 2^{-12.55}$, $p_f = 2^{-44.55}$, $n_k = 40$ and $l = 2^{36}$. We denote T_a , T_b , T_c , T_d and T_3 as the time complexity in step (a), (b), (c) (d) and 3, respectively, which are as follows: $T_a = 2^{64}$ memory accesses, $T_b = 2^{72}$ memory accesses, $T_c = 36 \cdot 2^{32}$ memory accesses, $T_d = 36 \cdot 2^{31} \cdot 2^{-12.55} \cdot 2^{40} \cdot (\frac{1}{2} + \frac{1}{8}) \cdot 2 = 2^{65.20}$ 1-round encryptions and $T_3 = 2^{36} \cdot 2^{40} = 2^{76}$ 18-round encryptions. Therefore, the total time complexity will be 2^{76} 18-round encryptions. The data complexity is 2^{64} chosen plaintexts and the memory requirements are 2^{40} 128-bit cells for the hash table, which can be reused for the 2^{40} counters. The success probability is 85.95%.

The ratio of weak key satisfying the sum of the probabilities of the 36 differentials is computed as follows:

$$r_w = 1 - \sum_{x=0}^{\mu-1} \text{Poisson} \left(x, 2^{n-1} \sum_{j=1}^{N_d} p_i \right) = 1 - \sum_{x=0}^{2^{5.03}-1} \text{Poisson} (x, 2^{63} \cdot 2^{-57.97}) = 0.57.$$

This means that the number of weak keys for which our attack can succeed is $2^{80} \cdot 0.57 = 2^{79.19}$ for PRESENT-80. A comparison with the attack of [6] can be found in Section 2.

5 Attack on Reduced-Round Serpent

Serpent was one of the five AES candidates in the final round; it is an SPN block cipher with 32 rounds [1]. In our attacks, we consider Serpent from rounds 4 to 11.

In the previous differential cryptanalysis of Serpent in [3], Biham *et al.* used the structure attack for Serpent. They identify a differential characteristic for $\frac{1}{2} + 5$ rounds starting from the linear transformation with fewer active S-boxes (13 active S-boxes) in the first half round, then extend it backwards to 6 rounds. Moreover, there is only one differential characteristic in each differential due to the strong avalanche characteristics of Serpent. Biham *et al.* claim that 2^{14} differential characteristics with probability 2^{-93} have been found. However, it can be shown that there are only 2^{13} differential characteristics with probability 2^{-93} . The proof has been omitted due to space constraints.

For the differential characteristics, the output difference of S-boxes in the first round is $\{0906b010_x || 00000080_x || 13000226_x || 06040030_x\}$. We will use all the possible non-zero input differences according to the output differences for the S-boxes (S_4) in the first round. According to the differential distribution table of S_4 , we have $|\Delta_0| = 2^{35.32}$ and $\sum_{i=1}^{|\Delta_0|} p_i = 2^{-65}$ which is equal to the probability of the differential characteristic from round 2 to round 6.

We now apply the structure attack described in Sect. 3. We construct 2^{19} structures of 2^{52} chosen plaintexts each. In each structure, all the inputs to non-active S-boxes in the first round are fixed to some random value, while the 52 bits of input to all the active S-boxes take all the 2^{52} possible values. There are $2^{19} \cdot 2^{51} = 2^{70}$ pairs for each differential characteristic. We expect that about $2^{70} \cdot 2^{-65} = 2^5$ pairs produce the output difference Δ_6 . In order to reduce the time complexity and ensure a higher success probability, 52 bits subkey are guessed after the data collection process. After retrieving 52 bits of the subkey, we can use the right pairs to recover the remaining 24 bits of the subkey.

The success probability P_S can be computed with Equation (1). Here $N = 2^{71}$, $|\Delta_0| = 2^{35.32}$, $p_* = 2^{-65} \cdot 2^{-35.32} \cdot 2^{52} = 2^{-48.32}$, $N_s = 2^{70} \cdot 2^{35.32} \cdot 2^{-52} = 2^{53.32}$, $p = 2^{-52}$, $n_k = 52$, $l = 2$, $\beta = 2^{-26.22}$, hence we get $P_S = 89.87\%$.

The time complexity is $2^{27.10} \cdot 2^{52} \cdot 13/32 = 2^{77.81}$ one-round encryptions which is equivalent to $2^{74.99}$ 7-round encryptions, the data complexity is 2^{71} chosen plaintexts and the memory requirements are 2^{52} hash cells of 256 bits and 2^{52} 32-bit counters storing 2^5 pairs each, hence using about 2^{57} 256-bit words. This attack consequently applies to Serpent with all key sizes of 128,192 and 256 bits.

The attack can be further extended to 8-round Serpent-256. By exhaustively searching the 128-bit subkey in the last round to decrypt to round 7, the above attack for 7 rounds can be applied. The time complexity is $2^{203.81}$ 8-round encryptions, the data complexity is 2^{71} chosen plaintexts and the memory requirements are the same as for the 7-round attack. This attack therefore applies only to Serpent with a 256-bit key.

In comparison, the previous differential attack for 7-round Serpent described in [3] has a time complexity of 2^{85} memory accesses and a data complexity of 2^{84} chosen plaintexts. For the previous differential attack on 8-round Serpent, the time complexity is 2^{213} memory accesses and the data complexity is 2^{84} chosen plaintexts. This implies that our attacks require much less chosen plaintexts and improve the time complexity.

It is possible to further reduce the data requirements at the expense of the time complexity. We have identified another set of differentials for 5.5 rounds which have 16 instead of 13 active S-boxes in the first round (the sequence of active S-Boxes is 16–10–6–2–1–5, and there are $2^{41.49}$ input differences). The combined probability of these differentials is $2^{-62.85}$, leading to a total time complexity greater than the previously described attack.

The ratio of weak keys satisfying the probability of the multiple differentials is computed as follows:

$$r_w = 1 - \sum_{x=0}^{\mu-1} \text{Poisson} \left(x, 2^{n-1} \sum_{j=1}^{N_d} p_j \right) = 1 - \sum_{x=0}^{2^5-1} \text{Poisson} (x, 2^{70} \cdot 2^{-65}) = 0.52.$$

This means that this attack is expected to work with about half of all possible keys, independent of the key size.

6 Conclusion

In this paper, we give a general model for the structure attack, providing guidance on how to choose the set of differentials to minimize the time complexity. As concrete applications of our model, we present structure attacks on 18-round PRESENT and improve the previous differential cryptanalytic results for the Serpent block cipher. To the best of our knowledge, those attacks are the best known differential attacks on these two block ciphers.

Comparing our model for structure attacks against the general model for multiple differential cryptanalysis proposed in [4], we conclude that the limitation for the set of input differences imposed by the model of [4] excludes many valuable differentials. We show that in structure attacks, a very important – and often particularly efficient – subclass of multiple differential attacks, this restriction can be relaxed. In our model presented in Sect. 3, the analysis of an attack can be carried out without this assumption.

The relevance of the limitation imposed by the condition of Definition 1 is additionally supported by our concrete application of the structure attack to PRESENT, which is more efficient than the multiple differential cryptanalysis with different output differences described in [4] and [6] where this condition was necessary. By removing this limitation, we have identified new sets of differentials that improve on the previous analysis.

It remains an interesting open question to find a block cipher other than PRESENT for which multiple differential cryptanalysis with multiple output differences produces superior results to the structure attack. Furthermore, our attack model can be used as a guidance to improve differential attacks for other algorithms. Applying it to other block ciphers than PRESENT or Serpent will be subject of future work.

Acknowledgements. We would like to thank Vincent Rijmen for valuable advice and Kerem Varici for his support for the work on PRESENT. The work in this paper was supported in part by the European Commission through the ICT programme under contract ICT-2007-216676 ECRYPT II and in part by the Concerted Research Action (GOA) TENSE 2011 of the Flemish Government, the IAP Program P6/26 BCrypt of the Belgian State (Belgian Science Policy) and by the Research Fund KU Leuven (project OT/08/027).

References

1. Anderson, R., Biham, E., Knudsen, L.R.: A Proposal for the Advanced Encryption Standard. NIST AES proposal (1998)
2. Biham, E., Shamir, A.: Differential Cryptanalysis of DES-like Cryptosystems. *Journal of Cryptology* 4(1), 3–72 (1991)
3. Biham, E., Dunkelman, O., Keller, N.: The Rectangle Attack - Rectangling the Serpent. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 340–357. Springer, Heidelberg (2001)
4. 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)
5. Blondeau, C., Gérard, B.: Private communication: The 561 Differentials (2011)
6. Blondeau, C., Gérard, B.: Multiple Differential Cryptanalysis: Theory and Practice (Corrected). *Cryptology ePrint Archive: Report 2011/115*
7. Bogdanov, A., Knudsen, L.R., Leander, G., Paar, C., Poschmann, A., Robshaw, M.J.B., Seurin, Y., Vikkelsoe, C.: PRESENT: An Ultra-Lightweight Block Cipher. In: Paillier, P., Verbaauwhede, I. (eds.) CHES 2007. LNCS, vol. 4727, pp. 450–466. Springer, Heidelberg (2007)
8. Cho, J.: Linear Cryptanalysis of Reduced-Round PRESENT. In: Pieprzyk, J. (ed.) CT-RSA 2010. LNCS, vol. 5985, pp. 302–317. Springer, Heidelberg (2010)
9. Collard, B., Standaert, F.-X.: A Statistical Saturation Attack against the Block Cipher PRESENT. In: Fischlin, M. (ed.) CT-RSA 2009. LNCS, vol. 5473, pp. 195–210. Springer, Heidelberg (2009)
10. Daemen, J., Rijmen, V.: Probability distributions of correlations and differentials in block ciphers. *Journal of Mathematical Cryptology* 1(3), 221–242 (2007)
11. Daemen, J., Rijmen, V.: Probability distributions of Correlation and Differentials in Block Ciphers (2005), <http://eprint.iacr.org/2005/212>
12. Dunkelman, O., Indestege, S., Keller, N.: A Differential-Linear Attack on 12-Round Serpent. In: Chowdhury, D.R., Rijmen, V., Das, A. (eds.) INDOCRYPT 2008. LNCS, vol. 5365, pp. 308–321. Springer, Heidelberg (2008)
13. 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)
14. Leander, G.: Small scale variants of the block cipher PRESENT. *Cryptology ePrint Archive, Report 2010/143* (2010)
15. Matsui, M., Nakajima, J.: On the Power of Bitslice Implementation on Intel Core2 Processor. In: Paillier, P., Verbaauwhede, I. (eds.) CHES 2007. LNCS, vol. 4727, pp. 121–134. Springer, Heidelberg (2007)
16. Nakahara Jr., J., Sepehrdad, P., Zhang, B., Wang, M.: Linear (Hull) and Algebraic Cryptanalysis of the Block Cipher PRESENT. In: Garay, J.A., Miyaji, A., Otsuka, A. (eds.) CANS 2009. LNCS, vol. 5888, pp. 58–75. Springer, Heidelberg (2009)
17. Ohkuma, K.: Weak Keys of Reduced-Round PRESENT for Linear Cryptanalysis. In: Jacobson Jr., M.J., Rijmen, V., Safavi-Naini, R. (eds.) SAC 2009. LNCS, vol. 5867, pp. 249–265. Springer, Heidelberg (2009)
18. Selçuk, A.A., Biçak, A.: On Probability of Success in Linear and Differential Cryptanalysis. In: Cimato, S., Galdi, C., Persiano, G. (eds.) SCN 2002. LNCS, vol. 2576, pp. 174–185. Springer, Heidelberg (2003)
19. Selçuk, A.A.: On Probability of Success in Linear and Differential Cryptanalysis. *Journal of Cryptology* 21(1), 131–147 (2008)
20. Wang, M.: Differential Cryptanalysis of Reduced-Round PRESENT. In: Vaudenay, S. (ed.) AFRICACRYPT 2008. LNCS, vol. 5023, pp. 40–49. Springer, Heidelberg (2008)

A Methodology for Differential-Linear Cryptanalysis and Its Applications^{*}

(Extended Abstract)

Jiqiang Lu

Institute for Infocomm Research,
Agency for Science, Technology and Research
1 Fusionopolis Way, #19-01 Connexis, Singapore 138632
lvjiqiang@hotmail.com, jlu@i2r.a-star.edu.sg

Abstract. In 1994 Langford and Hellman introduced a combination of differential and linear cryptanalysis under two default independence assumptions, known as differential-linear cryptanalysis, which is based on the use of a differential-linear distinguisher constructed by concatenating a linear approximation with a (truncated) differential with probability 1. In 2002, by using an additional assumption, Biham, Dunkelman and Keller gave an enhanced version that can be applicable to the case when a differential with a probability of smaller than 1 is used to construct a differential-linear distinguisher. In this paper, we present a new methodology for differential-linear cryptanalysis under the original two assumptions implicitly used by Langford and Hellman, without using the additional assumption of Biham et al. The new methodology is more reasonable and more general than Biham et al.'s methodology, and apart from this advantage it can lead to some better differential-linear cryptanalytic results than Biham et al.'s and Langford and Hellman's methodologies. As examples, we apply it to attack 10 rounds of the CTC2 block cipher with a 255-bit block size and key, 13 rounds of the DES block cipher, and 12 rounds of the Serpent block cipher. The new methodology can be used to cryptanalyse other block ciphers, and block cipher designers should pay attention to this new methodology when designing a block cipher.

Keywords: Block cipher, CTC2, DES, Serpent, Differential cryptanalysis, Linear cryptanalysis, Differential-linear cryptanalysis.

1 Introduction

Differential cryptanalysis was introduced in 1990 by Biham and Shamir [8]. Linear cryptanalysis was introduced in 1992 by Matsui and Yamagishi [31]. A differential cryptanalysis attack is based on the use of one or more so-called differentials,

^{*} An earlier version of this work appeared in 2010 as part of Cryptology ePrint Archive Report 2010/025 [28], which was done when the author was with Eindhoven University of Technology (The Netherlands) under the support of the Dutch Sentinels project PINPASJC (No. TIF.6687).

Table 1. Our and previous main cryptanalytic results on CTC2 and DES

Cipher	Attack Technique	Rounds	Data	Time	Success Rate	Source
CTC2 (255-bit version)	Algebraic [12]	6	4CP	2^{253} Enc.	not specified	[11]
	Differential	7 [†]	2^{15} CP	2^{15} Enc.	not specified	[15]
	Differential-linear	8 [†]	2^{37} CP	2^{37} Enc.	61.8%	[15]
		10	2^{142} CP	2^{207} Enc.	99.9%	Section 5.4
DES	Differential	full	$2^{47.2}$ CP	2^{37} Enc.	not specified	[10]
	Linear	full	2^{43} KP	2^{47} Enc.	85%	[29]
	Davis's attack [13]	full	2^{50} KP	2^{50} Enc.	51%	[4]
	Differential-linear	8	768CP	2^{40} Enc.	95%	[26]
		9	$2^{15.75}$ CP	2^{38} Enc.	88.8%	[14]
10		$2^{29.66}$ CP	2^{44} Enc.	97%	Section 4.2	
13		$2^{52.1}$ CP	$2^{54.2}$ Enc.	99%	Section 4.2	

†: There is a flaw; see Section 5.2 for detail.

and a linear cryptanalysis attack is based on the use of one or more so-called linear approximations. Both the cryptanalytic methods were used to attack the full Data Encryption Standard (DES) [32] algorithm faster than exhaustive key search [10, 29].

In 1994 Langford and Hellman [26] introduced a combination of differential and linear cryptanalysis under two default independence assumptions, known as differential-linear cryptanalysis, and they applied it to break 8-round DES. Such an attack is constructed on a so-called differential-linear distinguisher; a differential-linear distinguisher treats a block cipher as a cascade of two sub-ciphers, and it uses a linear approximation for a sub-cipher and, for the other sub-cipher it uses a differential (or a truncated differential [22]) with a one probability that does not affect the bit(s) concerned by the input mask of the linear approximation. In 2002, by using an additional assumption Biham, Dunkelman and Keller [5] introduced an enhanced version of differential-linear cryptanalysis, which is applicable to the case when a differential with a smaller probability is used to construct a differential-linear distinguisher; and they applied the enhanced version to break 9-round DES. Differential-linear cryptanalysis has been used to yield the best currently published cryptanalytic results for a number of state-of-the-art block ciphers [5, 6, 15, 16].

In this paper, we present a new methodology for differential-linear cryptanalysis under the two default assumptions implicitly used by Langford and Hellman, without using the additional assumption due to Biham et al. The new methodology is more reasonable and more general than Biham et al.'s methodology, and it can lead to some better differential-linear cryptanalytic results than Biham et al.'s and Langford and Hellman's methodologies. As examples, we apply the new methodology to mount differential-linear attacks on 10 rounds of the CTC2 [11] block cipher with a 255-bit block size and key, 13 rounds of DES, and 12 rounds

of the Serpent [1,2] block cipher. In terms of the numbers of attacked rounds: The 10-round CTC2 attack is the first published cryptanalytic attack on the version of CTC2; the 13-round DES attack is much better than any previously published differential-linear cryptanalytic results for DES, though it is inferior to the best previously published cryptanalytic results for DES; and the 12-round Serpent attack matches the best previously published cryptanalytic result for Serpent, that was obtained under Biham et al.’s methodology. Due to page constraints, we will only present the attacks on CTC2 and DES in this paper, and give the attack on Serpent in the full version of this paper (which contains more material). Table 1 summarises both our and previous main cryptanalytic results on CTC2 and DES, where CP and KP refer respectively to the required numbers of chosen plaintexts and known plaintexts, and Enc. refers to the required number of encryption operations of the relevant version of CTC2 or DES.

The remainder of the paper is organised as follows. In the next section we give the notation used throughout the paper and briefly describe differential and linear cryptanalysis. In Section 3 we review Langford and Hellman’s and Biham et al.’s methodologies and give our methodology for differential-linear cryptanalysis. In Sections 4–5 we present our cryptanalytic results on DES and CTC2, respectively. We discuss a few possible extensions to our methodology in Section 6. Section 7 concludes this paper.

2 Preliminaries

In this section we describe the notation, differential and linear cryptanalysis.

2.1 Notation

In the following descriptions, we assume that a number without a prefix is in decimal notation, and a number with prefix $0x$ is in hexadecimal notation, unless otherwise stated. The bits of a value are numbered from right to left, the leftmost bit is the most significant bit, and the rightmost bit is the least significant bit, except in the case of DES, where we use the same numbering notation as in FIPS-46 [32]. We use the following notation.

- \oplus bitwise logical exclusive OR (XOR) of two bit strings of the same length
- \odot dot product of two bit strings of the same length
- \parallel string concatenation
- \lll left rotation of a bit string
- \circ functional composition. When composing functions X and Y, $X \circ Y$ denotes the function obtained by first applying X and then applying Y
- e_j a 255-bit value with zeros everywhere except for bit position j , ($0 \leq j \leq 254$)
- e_{i_0, \dots, i_j} the 255-bit value equal to $e_{i_0} \oplus \dots \oplus e_{i_j}$, ($0 \leq i_0, \dots, i_j \leq 254$)
- \mathbb{E} an n -bit block cipher when used with a specific user key

2.2 Differential Cryptanalysis

Differential cryptanalysis [8] takes advantage of how a specific difference in a pair of inputs of a cipher can affect a difference in the pair of outputs of the cipher, where the pair of outputs are obtained by encrypting the pair of inputs using the same key. The notion of difference can be defined in several ways; the most widely discussed is with respect to the XOR operation. The difference between the inputs is called the input difference, and the difference between the outputs of a function is called the output difference. The combination of the input difference and the output difference is called a differential. The probability of a differential is defined as follows.

Definition 1 (from [27]). *If α and β are n -bit blocks, then the probability of the differential (α, β) for \mathbb{E} , written $\Delta\alpha \rightarrow \Delta\beta$, is defined to be*

$$\Pr_{\mathbb{E}}(\Delta\alpha \rightarrow \Delta\beta) = \Pr_{P \in \{0,1\}^n} (\mathbb{E}(P) \oplus \mathbb{E}(P \oplus \alpha) = \beta).$$

The following result follows trivially from Definition 1:

Proposition 1 (from [27]). *If α and β are n -bit blocks, then*

$$\Pr_{\mathbb{E}}(\Delta\alpha \rightarrow \Delta\beta) = \frac{|\{x | \mathbb{E}(x) \oplus \mathbb{E}(x \oplus \alpha) = \beta, x \in \{0,1\}^n\}|}{2^n}.$$

For a random function, the expected probability of a differential for any pair (α, β) is 2^{-n} . Therefore, if $\Pr_{\mathbb{E}}(\Delta\alpha \rightarrow \Delta\beta)$ is larger than 2^{-n} , we can use the differential to distinguish \mathbb{E} from a random function, given a sufficient number of chosen plaintext pairs.

Sometimes, we simply write $\Delta\alpha \xrightarrow{\mathbb{E}} \Delta\beta$ to denote the differential $\Delta\alpha \rightarrow \Delta\beta$ for \mathbb{E} in this paper.

2.3 Linear Cryptanalysis

Linear cryptanalysis [29, 31] exploits correlations between a particular linear function of the input blocks and a second linear function of the output blocks. The combination of the two linear functions is called a linear approximation. The most widely used linear function involves computing the bitwise dot product operation of the block with a specific binary vector (the specific value combined with the input blocks may be different from the value applied to the output blocks). The value combined with the input blocks is called the input mask, and the value applied to the output blocks is called the output mask. The probability of a linear approximation is defined as follows.

Definition 2 (from [27]). *If α and β are n -bit blocks, then the probability of the linear approximation (α, β) for \mathbb{E} , written $\Gamma\alpha \rightarrow \Gamma\beta$, is defined to be*

$$\Pr_{\mathbb{E}}(\Gamma\alpha \rightarrow \Gamma\beta) = \Pr_{P \in \{0,1\}^n} (P \odot \alpha = \mathbb{E}(P) \odot \beta).$$

We refer to below the dot product $P \odot \alpha$ as the input parity, and the dot product $\mathbb{E}(P) \odot \beta$ as the output parity. The following result follows trivially from Definition 2:

Proposition 2 (from [27]). *If α and β are n -bit blocks, then*

$$\Pr_{\mathbb{E}}(\Gamma\alpha \rightarrow \Gamma\beta) = \frac{|\{x|x \odot \alpha = \mathbb{E}(x) \odot \beta, x \in \{0, 1\}^n\}|}{2^n}.$$

For a random function, the expected probability of a linear approximation for any pair (α, β) is $\frac{1}{2}$. The bias of a linear approximation $\Gamma\alpha \rightarrow \Gamma\beta$, denoted by ϵ , is defined to be $\epsilon = |\Pr_{\mathbb{E}}(\Gamma\alpha \rightarrow \Gamma\beta) - \frac{1}{2}|$. Thus, if the bias ϵ is sufficiently large, we can use the linear approximation to distinguish \mathbb{E} from a random function, given a sufficient number of matching plaintext-ciphertext pairs.

2.4 General Assumptions Used in Practice

Propositions 1 and 2 give the accurate probability values of a differential and a linear approximation from a theoretical point of view. However, it is usually hard to apply them in practice to a block cipher with a large block size, for example, $n = 64$ or 128 which is currently being widely used in reality, and even harder when the differential or linear approximation operates on many rounds of the cipher. In practice, for a Markov block cipher [24], a multi-round differential (or linear approximation) is usually obtained by concatenating a few one-round differential characteristics (respectively, linear approximations), and the probability of the multi-round differential (or linear approximation) is regarded as the product (respectively, the piling-up function [29]) of the probabilities of the one-round differential characteristics (respectively, linear approximations) under the following Assumption 1.

Assumption 1. *The involved round functions behave independently.*

We note that one may argue the correctness of Assumption 1 and may use a different assumption, for example, many people would like to use the assumption that the round keys are independent and uniformly distributed; however, it is not accurate, either, for generally the round keys are actually dependent, being generated from a global user key under the key schedule algorithm of the cipher. Anyway, all such assumptions require us to treat the involved rounds as independent. As mentioned in [17], this is “most often not exactly the case, but as often it is a good approximation”.

Differential and linear cryptanalyses generally treat a basic unit of input (i.e. a chosen-plaintext pair for differential cryptanalysis; a known-plaintext for linear cryptanalysis) as a random variable, and assume that given a set of inputs of the basic unit, the inputs that satisfy the required property can be approximated by an independent distribution, as followed in [9, 29].

3 Differential-Linear Cryptanalysis: Previous and Our Methodologies

In this section we first review previous methodologies on differential-linear cryptanalysis, namely Langford and Hellman's and Biham et al.'s methodologies, and then give our new methodology, followed by a few implications. First observe that for simplicity we assume that the probability for a linear approximation with bias ϵ is $\frac{1}{2} + \epsilon$ in all the following descriptions; but the same results can be obtained when the probability is $\frac{1}{2} - \epsilon$.

3.1 Langford and Hellman's Methodology

In 1994 Langford and Hellman [26] introduced differential-linear cryptanalysis as a combination of differential and linear cryptanalysis, which is based on the use of a differential-linear distinguisher. To construct a differential-linear distinguisher, they treated \mathbb{E} as a cascade of two sub-ciphers \mathbb{E}_0 and \mathbb{E}_1 , where $\mathbb{E} = \mathbb{E}_0 \circ \mathbb{E}_1$. A differential-linear distinguisher uses a (truncated) differential $\Delta\alpha \rightarrow \Delta\beta$ with probability 1 for \mathbb{E}_0 and a linear approximation $\Gamma\gamma \rightarrow \Gamma\delta$ with bias ϵ for \mathbb{E}_1 , where the output difference β of the (truncated) differential has a zero value in the bit positions concerned by the input mask of the linear approximation (thus $\beta \odot \gamma = 0$ holds). Let P be a plaintext chosen uniformly at random from $\{0, 1\}^n$. Thus, we have $\mathbb{E}_0(P) \odot \gamma = \mathbb{E}_0(P \oplus \alpha) \odot \gamma$ with probability 1. The differential-linear distinguisher is concerned with the event $\delta \odot \mathbb{E}(P) = \delta \odot \mathbb{E}(P \oplus \alpha)$; and under Assumption 1 and the following Assumption 2 it has a probability of $\Pr(\delta \odot \mathbb{E}(P) = \delta \odot \mathbb{E}(P \oplus \alpha)) = (\frac{1}{2} + \epsilon) \times (\frac{1}{2} + \epsilon) + (\frac{1}{2} - \epsilon) \times (\frac{1}{2} - \epsilon) = \frac{1}{2} + 2\epsilon^2$.

Assumption 2. *The two inputs $\mathbb{E}_0(P)$ and $\mathbb{E}_0(P \oplus \alpha)$ of the linear approximation for \mathbb{E}_1 behave as independent inputs with respect to the linear approximation.*

Note that $\mathbb{E}(P) = \mathbb{E}_1(\mathbb{E}_0(P))$ and $\mathbb{E}(P \oplus \alpha) = \mathbb{E}_1(\mathbb{E}_0(P \oplus \alpha))$ in the above descriptions. Assumption 2 is somewhat like assuming an independent distribution for plaintext pairs generated from a particular plaintext structure with certain property in differential cryptanalysis.

By contrast, for a random function, the expected probability of a differential-linear distinguisher is $\frac{1}{2}$. Therefore, if the bias $|\Pr(\delta \odot \mathbb{E}(P) = \delta \odot \mathbb{E}(P \oplus \alpha)) - \frac{1}{2}| = 2\epsilon^2$ is sufficiently large, we can distinguish \mathbb{E} from a random function.

3.2 Biham et al.'s Methodology

A differential-linear distinguisher plays a fundamental role in a differential-linear cryptanalysis attack. In 2002 Biham, Dunkelman and Keller [5] presented an enhanced version to make a differential-linear distinguisher cover more rounds of a block cipher, so that an attacker can potentially break more rounds of the cipher. Biham et al.'s enhanced version includes the case when the (truncated) differential $\Delta\alpha \rightarrow \Delta\beta$ has a smaller probability than 1, p say, with β meeting the

condition $\beta \odot \gamma = 0$.¹ A slightly revised version was given in [14]. They applied Langford and Hellman’s analysis described above when $\mathbb{E}_0(P) \oplus \mathbb{E}_0(P \oplus \alpha) = \beta$, and used the following Assumption 3 for the cases where $\mathbb{E}_0(P) \oplus \mathbb{E}_0(P \oplus \alpha) \neq \beta$:²

Assumption 3. *The output parities $\delta \odot \mathbb{E}(P)$ and $\delta \odot \mathbb{E}(P \oplus \alpha)$ have a uniform and independent distribution in $\{0, 1\}$ for the cases where $\mathbb{E}_0(P) \oplus \mathbb{E}_0(P \oplus \alpha) \neq \beta$.*

As a result, under Assumptions 1, 2 and 3, Biham et al. got $\Pr(\delta \odot \mathbb{E}(P) = \delta \odot \mathbb{E}(P \oplus \alpha)) = p \times (\frac{1}{2} + 2\epsilon^2) + (1 - p) \times \frac{1}{2} = \frac{1}{2} + 2p\epsilon^2$.

Finally, they concluded that if the bias $2p\epsilon^2$ is sufficiently large, the distinguisher can be used as the basis of a differential-linear attack to distinguish \mathbb{E} from a random function. Roughly, the attack has a data complexity of about $O(p^{-2}\epsilon^{-4})$.

Note. We learnt from the comments of an anonymous reviewer that the same methodology appeared earlier in Langford’s PhD thesis [25], (which seems to be not publicly accessible). For simplicity, in this paper we use the phrase “Biham et al.’s methodology” to express this methodology, but hope the reader to keep in mind that Langford proposed the same methodology a few years earlier.

3.3 Our Methodology

In summary, the differential-linear distinguishers described above are concerned with the correlation between a pair of output parities, where the pair of output parities are obtained by applying a linear function (e.g. bitwise dot product with δ) to the outputs of a pair of input blocks with difference α (under the same key). The combination of the input difference and the linear function is called a differential-linear distinguisher. More formally, we define the probability of the differential-linear distinguisher as follows.

Definition 3. *If α and δ are n -bit blocks, then the probability of the differential-linear distinguisher (α, δ) for \mathbb{E} , written $\Delta\alpha \rightarrow \Gamma\delta$, is defined to be*

$$\Pr_{\mathbb{E}}(\Delta\alpha \rightarrow \Gamma\delta) = \Pr_{P \in \{0,1\}^n} (\mathbb{E}(P) \odot \delta = \mathbb{E}(P \oplus \alpha) \odot \delta).$$

The following result follows trivially from Definition 3:

Proposition 3. *If α and δ are n -bit blocks, then*

$$\Pr_{\mathbb{E}}(\Delta\alpha \rightarrow \Gamma\delta) = \frac{|\{x | \mathbb{E}(x) \odot \delta = \mathbb{E}(x \oplus \alpha) \odot \delta, x \in \{0, 1\}^n\}|}{2^n}.$$

¹ A more general condition is $\beta \odot \gamma = c$, where $c \in \{0, 1\}$ is a constant. Without loss of generality, we consider the case with $c = 0$ throughout this paper.

² We note that Biham et al. used a different assumption when reviewing the enhanced version in a few other papers, [7] say, where they assumed that $\mathbb{E}_0(P) \odot \gamma = \mathbb{E}_0(P \oplus \alpha) \odot \gamma$ holds with half a chance for the cases where $\mathbb{E}_0(P) \oplus \mathbb{E}_0(P \oplus \alpha) \neq \beta$, yielding the same probability value $\frac{1}{2} + 2p\epsilon^2$ as under Assumption 3. We treat this assumption as Assumption 3, though they are different.

For a random function, the expected probability of a differential-linear distinguisher for any combination (α, δ) is $\frac{1}{2}$. Similarly, the bias of the differential-linear distinguisher $\Delta\alpha \rightarrow \Gamma\delta$ is defined to be $|\Pr_{\mathbb{E}}(\Delta\alpha \rightarrow \Gamma\delta) - \frac{1}{2}|$. Thus, if the bias is sufficiently large, we can use the differential-linear distinguisher to distinguish \mathbb{E} from a random function, given a sufficient number of chosen plaintext pairs.

In practice, it is usually infeasible to compute the accurate probability of a differential-linear distinguisher $\Delta\alpha \rightarrow \Gamma\delta$ by Proposition 3, and we have to make use of some assumptions to approximate it, like Biham et al.'s methodology described in Section 3.2. However, Biham et al.'s methodology uses the three assumptions as hypotheses and works only when Assumption 3 holds; otherwise it may give probability values that are highly inaccurate in some situations; for example, let's intuitively consider the naive situation where the differential $\Delta\alpha \rightarrow \Delta\beta$ has probability $\frac{1}{2}$ and meets $\beta \odot \gamma = 0$, and all the other possible differentials $\{\Delta\alpha \rightarrow \Delta\hat{\beta}\}$ meet $\hat{\beta} \odot \gamma = 1$. Such an example can be easily built for a practical block cipher, DES say. The differential $\Delta\alpha \rightarrow \Delta\beta$ contributes $\frac{1}{2}[(\frac{1}{2} + \epsilon) \times (\frac{1}{2} + \epsilon) + (\frac{1}{2} - \epsilon) \times (\frac{1}{2} - \epsilon)] = \frac{1}{4} + \epsilon^2$ to the probability of the distinguisher, and the other differentials $\{\Delta\alpha \rightarrow \Delta\hat{\beta}\}$ contribute $\frac{1}{2}[(\frac{1}{2} + \epsilon) \times (\frac{1}{2} - \epsilon) + (\frac{1}{2} - \epsilon) \times (\frac{1}{2} + \epsilon)] = \frac{1}{4} - \epsilon^2$, which also cause a bias, but in a negative way, canceling the bias due to $\Delta\alpha \rightarrow \Delta\beta$. So the real bias of the distinguisher is 0, that is, the distinguisher has no cryptanalytic significance. But if we applied Biham et al.'s methodology in this situation, the distinguisher would have a bias of $2 \times \frac{1}{2} \times \epsilon^2 = \epsilon^2$, and thus the distinguisher would be useful (if ϵ^2 is large enough); but nevertheless it is useless in fact. Notice that this case is not truly a counterexample to Biham et al.'s methodology, for it is clear that Assumption 3 does not hold for it, but it suggests that we should be cautious about using Assumption 3 and actually, we should be careful with using any assumption, and it is preferable to use as few assumptions as possible.

Biham, Dunkelman and Keller used a heuristic way to approximate the probability of a differential-linear distinguisher. We make an analysis for the probability of a differential-linear distinguisher from a mathematical point, and obtain a new methodology under only Assumptions 1 and 2. Our result is given as Theorem 1, followed by a proof.

Theorem 1. *An n -bit block cipher \mathbb{E} is represented as a cascade of two sub-ciphers \mathbb{E}_0 and \mathbb{E}_1 , where $\mathbb{E} = \mathbb{E}_0 \circ \mathbb{E}_1$. If $\alpha (\neq 0)$ is an input difference for \mathbb{E}_0 , $\Gamma\gamma \rightarrow \Gamma\delta$ is a linear approximation with bias ϵ for \mathbb{E}_1 , and the sum of the probabilities for the differentials $\{\Delta\alpha \rightarrow \Delta\beta \mid \Pr_{\mathbb{E}_0}(\Delta\alpha \rightarrow \Delta\beta) > 0, \gamma \odot \beta = 0, \beta \in \{0, 1\}^n\}$ is \hat{p} ($= \sum_{\gamma \odot \beta = 0} \Pr_{\mathbb{E}_0}(\Delta\alpha \rightarrow \Delta\beta)$), then under Assumptions 1 and 2 the probability of the differential-linear distinguisher $\Delta\alpha \rightarrow \Gamma\delta$ is*

$$\Pr_{P \in \{0,1\}^n} (\mathbb{E}(P) \odot \delta = \mathbb{E}(P \oplus \alpha) \odot \delta) = \frac{1}{2} + 2(2\hat{p} - 1)\epsilon^2.$$

Proof. Given the input difference α for \mathbb{E}_0 , there are one or more possible output differences $\{\beta \mid \Pr_{\mathbb{E}_0}(\Delta\alpha \rightarrow \Delta\beta) > 0, \beta \in \{0, 1\}^n\}$; these output differences can be classified into two sets: one is $\{\beta \mid \gamma \odot \beta = 0, \Pr_{\mathbb{E}_0}(\Delta\alpha \rightarrow \Delta\beta) > 0, \beta \in \{0, 1\}^n\}$, and the other is $\{\beta \mid \gamma \odot \beta = 1, \Pr_{\mathbb{E}_0}(\Delta\alpha \rightarrow \Delta\beta) > 0, \beta \in \{0, 1\}^n\}$.

Let P be a plaintext chosen uniformly at random from $\{0,1\}^n$. Then, under Assumptions 1 and 2 we have

$$\begin{aligned}
& \Pr(\mathbb{E}(P) \odot \delta = \mathbb{E}(P \oplus \alpha) \odot \delta | \mathbb{E}_0(P) \oplus \mathbb{E}_0(P \oplus \alpha) = \beta, \gamma \odot \beta = 0) \\
&= \Pr(\mathbb{E}_0(P) \odot \gamma = \mathbb{E}(P) \odot \delta, \mathbb{E}_0(P \oplus \alpha) \odot \gamma = \mathbb{E}(P \oplus \alpha) \odot \delta | \\
&\quad \mathbb{E}_0(P) \oplus \mathbb{E}_0(P \oplus \alpha) = \beta, \gamma \odot \beta = 0) + \\
&\quad \Pr(\mathbb{E}_0(P) \odot \gamma \neq \mathbb{E}(P) \odot \delta, \mathbb{E}_0(P \oplus \alpha) \odot \gamma \neq \mathbb{E}(P \oplus \alpha) \odot \delta | \\
&\quad \mathbb{E}_0(P) \oplus \mathbb{E}_0(P \oplus \alpha) = \beta, \gamma \odot \beta = 0) \\
&= \left(\frac{1}{2} + \epsilon\right) \times \left(\frac{1}{2} + \epsilon\right) + \left[1 - \left(\frac{1}{2} + \epsilon\right)\right] \times \left[1 - \left(\frac{1}{2} + \epsilon\right)\right] \\
&= \frac{1}{2} + 2\epsilon^2,
\end{aligned}$$

and

$$\begin{aligned}
& \Pr(\mathbb{E}(P) \odot \delta = \mathbb{E}(P \oplus \alpha) \odot \delta | \mathbb{E}_0(P) \oplus \mathbb{E}_0(P \oplus \alpha) = \beta, \gamma \odot \beta = 1) \\
&= \Pr(\mathbb{E}_0(P) \odot \gamma = \mathbb{E}(P) \odot \delta, \mathbb{E}_0(P \oplus \alpha) \odot \gamma \neq \mathbb{E}(P \oplus \alpha) \odot \delta | \\
&\quad \mathbb{E}_0(P) \oplus \mathbb{E}_0(P \oplus \alpha) = \beta, \gamma \odot \beta = 1) + \\
&\quad \Pr(\mathbb{E}_0(P) \odot \gamma \neq \mathbb{E}(P) \odot \delta, \mathbb{E}_0(P \oplus \alpha) \odot \gamma = \mathbb{E}(P \oplus \alpha) \odot \delta | \\
&\quad \mathbb{E}_0(P) \oplus \mathbb{E}_0(P \oplus \alpha) = \beta, \gamma \odot \beta = 1) \\
&= \left(\frac{1}{2} + \epsilon\right) \times \left[1 - \left(\frac{1}{2} + \epsilon\right)\right] + \left[1 - \left(\frac{1}{2} + \epsilon\right)\right] \times \left(\frac{1}{2} + \epsilon\right) \\
&= \frac{1}{2} - 2\epsilon^2.
\end{aligned}$$

Next, under Assumptions 1 and 2 we can compute the probability of the differential-linear distinguisher as follows.

$$\begin{aligned}
& \Pr(\mathbb{E}(P) \odot \delta = \mathbb{E}(P \oplus \alpha) \odot \delta) \\
&= \sum_{\beta \in \{0,1\}^n, Y \in \{0,1\}} \Pr(\mathbb{E}(P) \odot \delta = \mathbb{E}(P \oplus \alpha) \odot \delta, \mathbb{E}_0(P) \odot \gamma \oplus \mathbb{E}_0(P \oplus \alpha) \odot \gamma = Y, \\
&\quad \mathbb{E}_0(P) \oplus \mathbb{E}_0(P \oplus \alpha) = \beta) \\
&= \sum_{\beta \in \{0,1\}^n, Y \in \{0,1\}} \Pr(\mathbb{E}(P) \odot \delta = \mathbb{E}(P \oplus \alpha) \odot \delta | \mathbb{E}_0(P) \odot \gamma \oplus \mathbb{E}_0(P \oplus \alpha) \odot \gamma = Y, \\
&\quad \mathbb{E}_0(P) \oplus \mathbb{E}_0(P \oplus \alpha) = \beta) \times \\
&\quad \Pr(\mathbb{E}_0(P) \odot \gamma \oplus \mathbb{E}_0(P \oplus \alpha) \odot \gamma = Y, \mathbb{E}_0(P) \oplus \mathbb{E}_0(P \oplus \alpha) = \beta) \\
&= \sum_{\beta \in \{0,1\}^n} \Pr(\mathbb{E}(P) \odot \delta = \mathbb{E}(P \oplus \alpha) \odot \delta | \mathbb{E}_0(P) \odot \gamma \oplus \mathbb{E}_0(P \oplus \alpha) \odot \gamma = 0, \\
&\quad \mathbb{E}_0(P) \oplus \mathbb{E}_0(P \oplus \alpha) = \beta) \times \Pr(\mathbb{E}_0(P) \odot \gamma \oplus \mathbb{E}_0(P \oplus \alpha) \odot \gamma = 0, \\
&\quad \mathbb{E}_0(P) \oplus \mathbb{E}_0(P \oplus \alpha) = \beta) + \\
&\quad \sum_{\beta \in \{0,1\}^n} \Pr(\mathbb{E}(P) \odot \delta = \mathbb{E}(P \oplus \alpha) \odot \delta | \mathbb{E}_0(P) \odot \gamma \oplus \mathbb{E}_0(P \oplus \alpha) \odot \gamma = 1,
\end{aligned}$$

$$\begin{aligned}
& \mathbb{E}_0(P) \oplus \mathbb{E}_0(P \oplus \alpha) = \beta) \times \Pr(\mathbb{E}_0(P) \odot \gamma \oplus \mathbb{E}_0(P \oplus \alpha) \odot \gamma = 1, \\
& \mathbb{E}_0(P) \oplus \mathbb{E}_0(P \oplus \alpha) = \beta) \\
&= \left(\frac{1}{2} + 2\epsilon^2\right) \times \sum_{\beta \in \{0,1\}^n, \gamma \odot \beta = 0} \Pr(\mathbb{E}_0(P) \oplus \mathbb{E}_0(P \oplus \alpha) = \beta) + \\
& \left(\frac{1}{2} - 2\epsilon^2\right) \times \sum_{\beta \in \{0,1\}^n, \gamma \odot \beta = 1} \Pr(\mathbb{E}_0(P) \oplus \mathbb{E}_0(P \oplus \alpha) = \beta) \\
&= \frac{1}{2} + 2(2\hat{p} - 1)\epsilon^2. \quad \square
\end{aligned} \tag{1}$$

Consequently, the bias of the differential-linear distinguisher $\Delta\alpha \rightarrow \Gamma\delta$ is

$$\left| \Pr_{P \in \{0,1\}^n} (\mathbb{E}(P) \odot \delta = \mathbb{E}(P \oplus \alpha) \odot \delta) - \frac{1}{2} \right| = 2|2\hat{p} - 1|\epsilon^2.$$

3.4 Implications

Biham et al.'s methodology requires Assumptions 1, 2 and 3, while our methodology requires only Assumptions 1 and 2. Thus, our methodology is more reasonable than Biham et al.'s methodology.

Biham et al.'s methodology holds only when Assumption 3 holds, and under the situation we have $\hat{p} = p + (1 - p)\frac{1}{2} = \frac{1}{2} + \frac{p}{2}$, meaning that the probability value obtained using Biham et al.'s methodology equals that obtained using our methodology. Thus, when Biham et al.'s methodology holds, our methodology always holds. However, our methodology holds under some situations where Biham et al.'s methodology does not hold, for example, it works for the naive situation discussed in Section 3.3 where $\hat{p} = p = \frac{1}{2}$. Therefore, our methodology is more general than Biham et al.'s methodology. (When Langford and Hellman's methodology holds, our methodology always holds as well.)

Our methodology still requires Assumptions 1 and 2. Assumption 1 is extensively used in and is commonly regarded as necessary for differential and linear cryptanalysis in practice. Assumption 2 seems irremovable to get such a simple and practical probability formula; otherwise, the formula could not be so simple, but a more accurate version can be easily obtained from our above reasonings, for instance, from Eq. (1), though it is complicated and appears to be hardly applicable in practice. The assumptions mean that, in some cases, the probability of a differential-linear distinguisher may be overestimated or underestimated, and so is the success probability of the attack; however, computer experiments [6, 16, 23, 26, 29, 30] have shown that the assumptions work well in practice for some block ciphers. Anyway, it seems reasonable to take the worst case assumption from the point of the user of a cipher. We suggest that if possible an attacker should check the validity of these assumptions when applying them to a specific cipher.

Our result shows that using only one (truncated) differential satisfying $\beta \odot \gamma = 0$ is not sufficient in most situations, and it is likely to be not sufficient in the general situation; we should use all the differentials satisfying $\beta \odot \gamma = 0$ instead. This makes

the distinguisher harder and even impossible to construct in practice, due to a large number of possible output differences. Anyway, we should use at least those differentials with a significant contribution to reduce the deviation if we are able to do so. Biham et al.'s methodology suggests that if the bias of the linear approximation keeps constant, the larger p is, the bigger is the bias of the distinguisher. Now, we know that may be not true in the general situation: A differential with a bigger probability will not necessarily result in a distinguisher with a bigger bias.

When constructing a differential-linear distinguisher, in Biham et al.'s methodology the attacker first chooses a (truncated) differential that meets the condition (as followed in [5, 6, 15, 16], in practice the output difference of the differential has zeros in the bit positions concerned by the input mask of the linear approximation), then calculates the probability of the differential, and finally takes this probability as the value of p . Our new methodology suggests a different format, that is, computing \hat{p} . Once the linear approximation and the input difference of the differentials are chosen, that how many rounds can be constructed for a distinguisher depends to some extent on the computational power available for the attacker.

Our new methodology can lead to some better differential-linear cryptanalytic results than Biham et al.'s and Langford and Hellman's methodologies, as to be demonstrated by its applications to the block ciphers DES and CTC2 in the following two sections. Before further proceeding, observe that DES is a Markov cipher under the XOR difference notion [24], and similarly we can learn that CTC2 as well as Serpent is a Markov cipher under the XOR difference notion.

At last, to be conservative, we would like to suggest that one should pay attention to all these methodologies, for a real situation is usually hard to predict, and it may make the Assumption 3 for Biham et al.'s methodology hold.

4 Application to the DES Block Cipher

The DES block cipher is well known to both academia and industry, which has a 64-bit block size, a 56-bit user key, and a total of 16 rounds. We refer the reader to [32] for the specifications of DES.

In 1994, under the two default Assumptions 1 and 2 Langford and Hellman [26] used their methodology to obtain a 6-round differential-linear distinguisher of DES, and finally applied it to break 8-round DES; the attack recovers 16 key bits with a time complexity of $2^{14.6}$ 8-round DES encryptions, so it would take 2^{40} encryptions to recover the remaining 40 key bits with an exhaustive search, meaning that a total of approximately 2^{40} 8-round DES encryptions are required to recover the whole 56 key bits (Note that there might exist an efficient way to obtain the remaining key bits). In 2002, under Assumptions 1, 2 and 3, Biham, Dunkelman and Keller [5] described a 7-round differential-linear distinguisher of DES using their enhanced methodology, and finally gave differential-linear attacks on 8 and 9-round DES; and an improved version of the 9-round attack appeared in pages 108–111 of [14]. Their attack recovers 18 key bits with a time complexity of $2^{29.17}$ 9-round DES encryptions, the remaining 38 key bits would take 2^{38} encryptions to recover with a key exhaustion, and thus it has a total of approximately 2^{38} 9-round DES encryptions to recover the whole 56 key bits.

Nevertheless, we find that our new methodology enables us to construct 7 and 8-round differential-linear distinguishers of DES based on the same 3-round linear approximation as used in the previous differential-linear cryptanalysis of DES [5,26]; the 8-round distinguisher can allow us to break 10-round DES. More importantly, we are able to construct a 11-round differential-linear distinguisher of DES, and finally use it as the basis of a differential-linear attack on 13-round DES. Below we describe the 11-round differential-linear distinguisher and our attack on 13-round DES. We write the subkey used in the S_l S-box of Round m as $K_{m,l}$, where $1 \leq m \leq 16, 1 \leq l \leq 8$.

4.1 A 11-Round Differential-Linear Distinguisher with Bias $2^{-24.05}$

The 11-round differential-linear distinguisher is made up of a 6-round linear approximation $\Gamma\gamma \rightarrow \Gamma\delta$ with bias $1.95 \times 2^{-9} \approx 2^{-8.04}$ and all the 5-round differentials $\{\Delta\alpha \rightarrow \Delta\beta\}$ with $\Delta\alpha = 0x4000000000000000$. The 6-round linear approximation $\Gamma\gamma \rightarrow \Gamma\delta$ is $0x0000000001040080 \rightarrow 0x2104008000000800$, (which is the best 6-round linear approximation given in [29]). Let's compute the probability of the 11-round differential-linear distinguisher using our new methodology.

We first consider the 5-round differentials $\{\Delta\alpha \rightarrow \Delta\beta\}$. There is a one probability in the first round, meaning that the first round is bypassed by the differential characteristic with probability 1. After the **E** expansion operation of the second round, $0x4$ in $\Delta\alpha$ becomes $0x8$, which enters the S_1 S-box of the second round and generates 11 differences after the S-box: $\{\omega|\omega = 0x3, 0x5, 0x6, 0x7, 0x9, 0xA, 0xB, 0xC, 0xD, 0xE, 0xF\}$; the probabilities for the output differences are given in the second column of Table 2. We represent ω as a concatenation of four one-bit variables $a||b||c||d$, where $a, b, c, d \in \{0, 1\}$. Thus, the right half of the third round has the input difference $00000000a0000000b000000c0000000d0$ in binary notation, and this input difference can make at most 6 S-boxes of the third round active: $S_2, S_3, S_4, S_5, S_6, S_8$.

In the third round, the S_2 S-box has an input difference $00000a$ in binary notation, the S_3 S-box has an input difference $0a0000$ in binary notation, the S_4 S-box has an input difference $00000b$ in binary notation, the S_5 S-box has an input difference $0b0000$ in binary notation, the S_6 S-box has an input difference $000c00$ in binary notation, and the S_8 S-box has an input difference $000d00$ in binary notation. We denote respectively by x_0, x_1, x_2 the most significant bit, the second most significant bit and the second least significant bit of the output difference of the S_2 S-box, by $x_3||x_4||x_5||x_6$ the output difference of the S_3 S-box, by x_7, x_8, x_9 the second most significant bit, the second least significant bit and the least significant bit of the output difference of the S_4 S-box, by $x_{10}||x_{11}||x_{12}||x_{13}$ the output difference of the S_5 S-box, by x_{14}, x_{15}, x_{16} the most significant bit, the second most significant bit and the second least significant bit of the output difference of the S_6 S-box, and by x_{17}, x_{18}, x_{19} the most significant bit, the second least significant bit and the least significant bit of the output difference of the S_8 S-box.

Table 2. Probabilities for the eleven output differences in $\{\omega\}$

ω	$\Pr_{S_1}(\Delta 0x8 \rightarrow \Delta\omega)$	$\Pr(\Delta\beta_\omega \odot \Gamma\gamma = 0 \Delta 0x8 \rightarrow \Delta\omega)$
$0x3$	$\frac{12}{64}$	0.49779944866895676
$0x5$	$\frac{8}{64}$	0.49595199525356293
$0x6$	$\frac{8}{64}$	0.50433863041689619
$0x7$	$\frac{4}{64}$	0.50256029706542904
$0x9$	$\frac{6}{64}$	0.50855094581311278
$0xA$	$\frac{2}{64}$	0.50591027818154544
$0xB$	$\frac{8}{64}$	0.50239421910760029
$0xC$	$\frac{8}{64}$	0.49929085310759547
$0xD$	$\frac{2}{64}$	0.49968796220765910
$0xE$	$\frac{2}{64}$	0.50061782109781916
$0xF$	$\frac{4}{64}$	0.50005227406592345

In the fourth round, the S_1 S-box has the input difference $0||x_9||x_2 \oplus 1||x_{13}||x_{14}||x_{17}$, and we denote by y_0 the second most significant bit of its output difference; the S_2 S-box has the input difference $x_{14}||x_{17}||x_6||0||x_{10}||0$, and we denote by y_1 the least significant bit of its output difference; the S_3 S-box has the input difference $x_{10}||0||x_8||x_{16}||0||x_0$, and we denote by y_2 the second most significant bit of its output difference; the S_4 S-box has the input difference $0||x_0||x_{11}||x_{18}||x_4||0$, and we denote by y_3 the second most significant bit of its output difference; the S_6 S-box has the input difference $x_7||x_{19}||0||0||x_3||x_{12}$, and we denote by y_4 the least significant bit of its output difference; the S_8 S-box has the input difference $x_1||x_{15}||x_5||0||0||x_9$, and we denote by y_5 the least significant bit of its output difference. Thus we have that the input difference of the S_5 S-box of the fifth round is $y_2||y_0 \oplus b||y_1||y_4||y_3||y_5$.

A simple analysis reveals that the three bits concerned by the input mask $\Gamma\gamma$ depend on: (1) x_{10} , x_{11} and x_{12} ; and (2) The three most significant bits of the output difference of the S_5 S-box of the fifth round; and we denote the XOR of the three bits by z .

For each difference ω , we denote by β_ω the output difference(s) of the 5-round DES. Now, by the differential distribution tables of the S-boxes (see [9]) we can compute the probability that the XOR of the concerned three bits of β_ω (i.e., $x_{10} \oplus x_{11} \oplus x_{12} \oplus z$) is zero by performing a computer program over all the possible (truncated) differential characteristics. These probabilities are given in the third column of Table 2. The largest number of possible differential characteristics happens when $\omega = 0xF$, which is $7 \times 10 \times 4 \times 10 \times 6 \times 7 \times 2^6 \times 2 \approx 2^{23.9}$; and it takes a few seconds to check on a personal computer.

Finally, by Theorem 1 we have that the probability of the 11-round distinguisher $\Delta\alpha \rightarrow \Gamma\delta$ is $\frac{1}{2} + 2 \times [2 \times \sum_\omega \Pr_{S_1}(\Delta 0x8 \rightarrow \Delta\omega) \times \Pr(\Delta\beta_\omega \odot \Gamma\gamma = 0 | \Delta 0x8 \rightarrow \Delta\omega) - 1] \times (2^{-8.04})^2 \approx \frac{1}{2} + 2 \times 2^{-8.97} \times (2^{-8.04})^2 = \frac{1}{2} + 2^{-24.05}$. Therefore, the 11-round distinguisher has a bias of $2^{-24.05}$.

4.2 Differential-Linear Attack on 13-Round DES

The 11-round distinguisher $\Delta\alpha \rightarrow \Gamma\delta$ can be used to break 13-round DES. We assume the attacked rounds are the first thirteen rounds from Rounds 1 to 13. A simple analysis on the key schedule of DES reveals that $K_{1,1}$ and $K_{13,1}$ overlap in 2 bits (i.e. bits 17 and 34 of the user key), and thus given $K_{1,1}$ we know 2 bits of $K_{13,1}$. The attack procedure is as follows.

1. Choose $2^{47.1}$ structures \mathcal{S}_i , ($i = 1, 2, \dots, 2^{47.1}$), where a structure is defined to be a set of 2^4 plaintexts $P_{i,j}$ with bits (9,17,23, 31) of the left half taking all the possible values, bit (2) of the right half fixed to 0 and the other 59 bits fixed, ($j = 1, 2, \dots, 2^4$). In a chosen-plaintext attack scenario, obtain all the ciphertexts for the 2^4 plaintexts in each of the $2^{47.1}$ structures; we denote by $C_{i,j}$ the ciphertext for plaintext $P_{i,j}$.
2. Choose $2^{47.1}$ structures $\widehat{\mathcal{S}}_i$, ($i = 1, \dots, 2^{47.1}$), where a structure $\widehat{\mathcal{S}}_i$ is obtained by setting 1 to bit (2) of the right half of all the plaintexts $P_{i,j}$ in \mathcal{S}_i . In a chosen-plaintext attack scenario, obtain all the ciphertexts for the 2^4 plaintexts in each $\widehat{\mathcal{S}}_i$.
3. Guess a value for $K_{1,1}$, and do as follows.
 - (a) Initialize 2^{20} counters to zero, which correspond to the 2^{20} possible pairs consisting of the possible values for a couple of the 10 ciphertext bits: bit (17) of the left half and bits (1,2,3,4,5,8,14,25,32) of the right half.
 - (b) Partially encrypt every (remaining) plaintext $P_{i,j}$ with the guessed $K_{1,1}$ to get its intermediate value immediately after Round 1; we denote it by $\varepsilon_{i,j}$.
 - (c) Partially decrypt $\varepsilon_{i,j} \oplus 0x4000000000000000$ with the guessed $K_{1,1}$ to get its plaintext, and find the plaintext in $\widehat{\mathcal{S}}_i$; we denote it by $\widehat{P}_{i,j}$, and denote by $\widehat{C}_{i,j}$ the corresponding ciphertext for $\widehat{P}_{i,j}$. Store $(C_{i,j}, \widehat{C}_{i,j})$ in a table.
 - (d) For every ciphertext pair $(C_{i,j}, \widehat{C}_{i,j})$, add 1 to the counter corresponding to the pair of the 10 ciphertext bits specified by $(C_{i,j}, \widehat{C}_{i,j})$.
 - (e) Guess a value for the unknown 4 bits of $K_{13,1}$, and do as follows.
 - i. For each of the 2^{20} pairs of the concerned 10 ciphertext bits, partially decrypt it with the guessed $K_{13,1}$ to get the pair of the 5 bits concerned by the output mask $\Gamma\delta$, and compute the XOR of the pair of the 5 bits (concerned by the output mask).
 - ii. Count the number of the ciphertext pairs $(C_{i,j}, \widehat{C}_{i,j})$ such that the XOR of the pair of the 5 bits concerned by $\Gamma\delta$ is zero, and compute its deviation from $2^{50.1}$.
 - iii. If the guess for $(K_{1,1}, K_{13,1})$ is the first guess for $(K_{1,1}, K_{13,1})$, then record the guess and the deviation computed in Step 3(e)(ii); otherwise, record the guess and its deviation only when the deviation is larger than that of the previously recorded guess, and remove the guess with the smaller deviation.
4. For the $(K_{1,1}, K_{13,1})$ recorded in Step 3(e)(iii), exhaustively search for the remaining 46 key bits with two known plaintext-ciphertext pairs. If a 56-bit key is suggested, output it as the user key of the 13-round DES.

The attack requires $2^{52.1}$ chosen plaintexts. The required memory for the attack is dominated by the storage of the plaintexts and ciphertexts, which is $2^{52.1} \times 16 = 2^{56.1}$ bytes. Steps 1 and 2 have a time complexity of $2^{52.1}$ 13-round DES encryptions. Steps 3(b) and 3(c) have a time complexity of $2 \times 2^{51.1} \times 2^6 \times \frac{1}{8 \times 13} \approx 2^{51.4}$ 13-round DES encryptions. Step 3(d) has a time complexity of $2^{51.1} \times 2^6 = 2^{57.1}$ memory accesses. Roughly, an extremely conservative estimate is: 13 memory accesses equal a 13-round DES encryption in terms of time, assuming that the 13-round DES is implemented with 8 parallel S-box lookups per round and one round is equivalent to one memory access. So the time complexity of Step 3(d) is equivalent to $\frac{2^{57.1}}{13} \approx 2^{53.4}$ 13-round DES encryptions. The time complexity of Step 3(e) is dominated by the time complexity of Step 3(e)(i), which is $2 \times 2^6 \times 2^4 \times 2^{20} \times \frac{1}{8 \times 13} \approx 2^{24.3}$ 13-round DES encryptions. Step 4 has a time complexity of 2^{46} 13-round DES encryptions. Therefore, the attack has a total time complexity of approximately $2^{54.2}$ 13-round DES encryptions, faster than exhaustive key search. There are $2^{51.1}$ plaintext pairs $(P_{i,j}, \bar{P}_{i,j})$ for a guess of $(K_{1,1}, K_{13,1})$, and thus following Theorem 2 of [33], we can know that the attack has a success probability of about 99%.

This shows that our new methodology enables us to break more rounds of DES than Biham et al.'s or Langford and Hellman's methodology. Since our attack works under only two assumptions, it is more reasonable than Biham et al.'s attack.

Note. Using the new methodology we can obtain a few differential-linear distinguishers operating on a smaller number of rounds, for example, a 7-round distinguisher ($\Delta\alpha = 0x4000000000000000, \Gamma\delta = 0x2104008000008000$) with bias $2^{-7.94}$ and an 8-round distinguisher ($\Delta\alpha = 0x4000000000000000, \Gamma\delta = 0x2104008000008000$) with bias $2^{-12.83}$, both using the same 3-round linear approximation as used in Biham et al.'s and Langford and Hellman's differential-linear cryptanalysis of DES. These distinguishers can allow us to break DES with a smaller number of rounds at a smaller complexity, for example, the 8-round distinguisher can similarly be used to break 10-round DES with a data complexity of $2^{29.66}$ chosen plaintexts and a time complexity of 2^{44} 10-round DES encryptions at a success rate of about 97%.

5 Application to the CTC2 Block Cipher

The CTC2 [11] cipher was designed to show the strength of algebraic cryptanalysis [12] on block ciphers by the proposer of algebraic cryptanalysis, who described an algebraic attack on 6 rounds of the version of CTC2 that uses a 255-bit block size and a 255-bit key. Using Biham et al.'s methodology, in 2009 Dunkelman and Keller [15] described 6 and 7-round differential-linear distinguishers for the version of CTC2, and finally presented differential-linear attacks on 7 and 8 rounds of CTC2 (with a 255-bit block size and key). The 8-round attack is known as the best previously published cryptanalytic result on the version of CTC2 in terms of the numbers of attacked rounds.

In this section, we first describe a flaw in the previous differential-linear cryptanalysis of CTC2. Then, under the new methodology we present an 8.5-round differential-linear distinguisher with bias 2^{-68} for the CTC2 with a 255-bit block size and key, and finally give a differential-linear attack on 10-round CTC2 (with a 255-bit block size and a key). First we briefly describe the CTC2 cipher.

5.1 The CTC2 Block Cipher

The CTC2 [11] block cipher has a variable block size, a variable length key and a variable number of rounds. There are many combinations for the block size, key size and round number. As in [15], we only consider the version of CTC2 that uses a 255-bit block size and a 255-bit key. CTC2 uses the following two elementary operations to construct its round function.

- **S** is a non-linear substitution operation constructed by applying the same 3×3 -bit bijective S-box 85 times in parallel to an input.
- **D** is a linear diffusion operation, which takes a 255-bit block $Y = (Y_{254}, \dots, Y_1, Y_0)$ as input, and outputs a 255-bit block $Z = (Z_{254}, \dots, Z_1, Z_0)$, computed as defined below.

$$\begin{cases} Z_{151} = Y_2 \oplus Y_{139} \oplus Y_{21} \\ Z_{(i \times 202 + 2) \bmod 255} = Y_i \oplus Y_{(i+137) \bmod 255} & i = 0, 1, 3, 4, \dots, 254 \end{cases}$$

CTC2 takes as input a 255-bit plaintext block P , and its encryption procedure for N_r rounds is, where $Z_0, X_i, Y_i, Z_i, X_{N_r}, Y_{N_r}, Z_{N_r}$ are 255-bit variables, and K_0, K_i, K_{N_r} are round keys generated from a user key K as $K_j = K \lll j$ in our notation, ($0 \leq j \leq N_r$).

1. $Z_0 = P$.
2. For $i = 1$ to $N_r - 1$:
 - $X_i = Z_{i-1} \oplus K_{i-1}$,
 - $Y_i = \mathbf{S}(X_i)$,
 - $Z_i = \mathbf{D}(Y_i)$.
3. $X_{N_r} = Z_{N_r-1} \oplus K_{N_r-1}$, $Y_{N_r} = \mathbf{S}(X_i)$, $Z_{N_r} = \mathbf{D}(Y_{N_r})$.
4. Ciphertext = $Z_{N_r} \oplus K_{N_r}$.

To keep in accordance with [11], the i th iteration of Step 2 in the above description is referred to as Round i , ($1 \leq i \leq N_r - 1$), and the transformations in Steps 3 and 4 are referred to as Round N_r . We number the 85 S-boxes in a round from 0 to 84 from right to left.

5.2 A Flaw in Previous Differential-Linear Cryptanalysis of CTC2

Observe that Dunkelman and Keller used the 0.5-round differential $e_{30,151} \xrightarrow{\mathbf{D}} e_2$ with probability 1 in their differential-linear attacks presented in [15]. However, we find that this differential is not correct: For the **D** operation, given the input difference $e_{30,151}$, we cannot get the output difference e_2 ; and the correct output

difference should be $e_{25,63,159,197}$. On the other hand, for the **D** operation, given the output difference e_2 , the input difference has over fifty non-zero bits, much more than the number two in $e_{30,151}$. As a consequence, the differential-linear cryptanalytic results are flawed.

Note that Dunkelman and Keller also described differential attacks on 5, 6 and 7-round CTC2 in [15], and the 0.5-round differential $e_{30,151} \xrightarrow{\mathbf{D}} e_2$ with probability 1 was also used and played a very important role in the differential results; thus they are flawed, too. It seems very hard to correct these differential and differential-linear cryptanalytic results to break that many rounds of CTC2.

5.3 An 8.5-Round Differential-Linear Distinguisher with Bias 2^{-68}

The 8.5-round differential-linear distinguisher with bias 2^{-68} is made up of a 5.5-round linear expression $\Gamma\gamma \rightarrow \Gamma\delta$ with bias 2^{-33} and all the 3-round differentials $\{\Delta\alpha \rightarrow \Delta\beta\}$ with $\Delta\alpha = e_0$. The 5.5-round linear expression $\Gamma\gamma \rightarrow \Gamma\delta$ is $e_{5,33,49,54,101,112,131,138,155,168,188,193,217,247,251} \rightarrow e_{32,151}$. Using the new methodology we can compute that the 8.5-round distinguisher $\Delta\alpha \rightarrow \Gamma\delta$ has a bias of 2^{-68} , in a manner similar to that for the above 11-round DES distinguisher.

5.4 Differential-Linear Attack on 10-Round CTC2 with a 255-Bit Block Size and Key

The above 8.5-round distinguisher can be used as the basis for a differential-linear attack breaking the version of CTC2 that has a 255-bit block size, a 255-bit key and a total of 10 rounds.

We assume the attacked rounds are the first ten rounds from Rounds 1 to 10; and we use the distinguisher from Rounds 2 until before the **D** operation of Round 10. We can learn that the input difference α propagates to 16 bit positions after the inverse of the **D** operation of Round 1: Bits 17, 21, 40, 59, 78, 97, 116, 135, 139, 154, 158, 177, 196, 215, 234 and 253. The 16 active bits correspond to 16 S-boxes of Round 0: S-boxes 5, 7, 13, 19, 26, 32, 38, 45, 46, 51, 52, 59, 65, 71, 78 and 84; let Θ be the set of the 16 S-boxes, and K_Θ be the 48 bits of K_0 corresponding to the 16 S-boxes in Θ . Another observation is that we do not need to guess the subkey bits from K_{10} , because the output mask $\Gamma\delta$ of the 8.5-round distinguisher concerns the intermediate value immediately after the **S** operation of Round 10, and for a pair of ciphertexts (C, \hat{C}) the value of $\delta \odot \mathbf{D}^{-1}(C) \oplus \delta \odot \mathbf{D}^{-1}(\hat{C})$ equals to $\delta \odot \mathbf{D}^{-1}(C \oplus \hat{C})$, which is independent of K_{10} . The attack procedure is as follows.

1. Choose 2^{94} structures \mathcal{S}_i , ($i = 0, 1, \dots, 2^{94} - 1$), where a structure is defined to be a set of 2^{48} plaintexts $P_{i,j}$ with the 48 bits for the S-boxes in Θ taking all the possible values and the other 207 bits fixed, ($j = 0, 1, \dots, 2^{48} - 1$). In a chosen-plaintext attack scenario, obtain all the ciphertexts for the 2^{48} plaintexts in each of the 2^{94} structures; we denote by $C_{i,j}$ the ciphertext for plaintext $P_{i,j}$.

2. Initialize 2^{48} counters to zero, which correspond to all the possible values for K_Θ .
3. For every structure \mathcal{S}_i , guess a value for K_Θ , and do as follows.
 - (a) Partially encrypt every (remaining) plaintext $P_{i,j}$ with the guessed K_Θ to get its intermediate value immediately after the **S** operation of Round 1; we denote it by $\varepsilon_{i,j}$.
 - (b) Take bitwise complements to bits (17, 21, 40, 59, 78, 97, 116, 135, 139, 154, 158, 177, 196, 215, 234, 253) of $\varepsilon_{i,j}$, and keep the other bits of $\varepsilon_{i,j}$ invariant; we denote the resulting value by $\widehat{\varepsilon}_{i,j}$.
 - (c) Partially decrypt $\widehat{\varepsilon}_{i,j}$ with the guessed K_Θ to get its plaintext, and find the plaintext in \mathcal{S}_i ; we denote it by $\widehat{P}_{i,j}$, and denote by $\widehat{C}_{i,j}$ the corresponding ciphertext for $\widehat{P}_{i,j}$.
 - (d) For $(C_{i,j}, \widehat{C}_{i,j})$, compute the XOR of bits 32 and 151 of $\mathbf{D}^{-1}(C_{i,j} \oplus \widehat{C}_{i,j})$. If the XOR is zero, add 1 to the counter corresponding to the guessed K_Θ .
4. For the K_Θ with the highest deviation from 2^{140} , exhaustively search for the remaining 207 key bits with a known plaintext-ciphertext pair. If a 255-bit key is suggested, output it as the user key of the version of CTC2.

The attack requires 2^{142} chosen plaintexts. Note that we start to collect another structure of plaintexts only after testing a structure of plaintexts, so that we can reuse the memory for storing the structure of plaintexts, hence the required memory of the attack is dominated by the storage of the 2^{48} counters and a structure of 2^{48} plaintext-ciphertext pairs, which is $2^{48} \times \frac{48}{8} + 2 \times 2^{48} \times \frac{255}{8} \approx 2^{54.2}$ bytes of memory. The time complexity of Step 3 is dominated by the time complexity of Steps 3(a), 3(c) and 3(d), which is approximately $2 \times 2^{141} \times 2^{48} \times \frac{16}{85 \times 10} + 2^{141} \times 2^{48} \times \frac{1}{10} \approx 2^{186.2}$ 10-round CTC2 encryptions. Step 4 has a time complexity of 2^{207} 10-round CTC2 encryptions. Therefore, the attack has a total time complexity of 2^{207} 10-round CTC2 encryptions to find the 255-bit key. There are 2^{141} plaintext pairs $(P_{i,j}, \widehat{P}_{i,j})$ for a guess of K_Θ . Following Theorem 2 of [33], we can learn that the probability that the correct guess for K_Θ has the highest deviation is about 99.9%. Thus, the attack has a success probability of about 99.9%.

6 Possible Extensions of Our Methodology

In this section we briefly discuss several possible extensions of our methodology, although particulars should be noticed.

The first possible extension is to consider the case when using two different values for the output mask δ in Definition 3, say δ_1, δ_2 ; that is, we might consider the event $\mathbb{E}(P) \odot \delta_1 = \mathbb{E}(P \oplus \alpha) \odot \delta_2$ for a randomly chosen $P \in \{0, 1\}^n$. The resulting differential-linear distinguisher would have a bias of $2(2\widehat{p} - 1)\epsilon_1\epsilon_2$ for some ϵ_1 and ϵ_2 denoting the respective bias of the two linear approximations. From a theoretical point of view, there seems no need to use two different output masks, for we can always choose the output mask with a bigger bias,

and a key-recovery attack based on a differential-linear distinguisher with two different output masks requires us to guess no less key bits than that based on a differential-linear distinguisher with one output mask; however, the case with two different output masks may depend on Assumption 2 to a lesser degree than the discussed case with one output mask, for the two linear approximations can be independent somewhat, instead of two identical linear approximations used in the case with one output mask, and thus it may potentially be particularly helpful when making a practicable attack in reality.

The second possible extension is to consider the case when applying our methodology in a related-key [3,19,21] attack scenario. The notion of the related-key differential-linear analysis appeared in [18], and later Kim [20] described an enhanced version based on Biham et al.'s enhanced methodology. Likewise, we can get a more reasonable and general version based on our new methodology.

Other possible extensions are to obtain new methodologies, in a way similar to the above new methodology for differential-linear cryptanalysis, for the high-order differential-linear attack, the differential-bilinear attack and the differential-bilinear-boomerang attack, which were proposed in [7]. At present, however, these attack techniques appear to be hard to apply to obtain good cryptanalytic results in practice.

7 Conclusions

In this paper we have given a new methodology for differential-linear cryptanalysis under only the two assumptions implicitly used in the very first published paper on this technique. The new methodology is more reasonable and more general than Biham et al.'s methodology, and it can lead to some better differential-linear cryptanalytic results for some block ciphers than the previously known methodologies.

Using the new methodology, we have presented differential-linear attacks on 13-round DES and 10-round CTC2 with a 255-bit block size and key. In terms of the numbers of attacked rounds, the 10-round CTC2 attack is the first published cryptanalytic attack on the version of CTC2; and the 13-round DES attack is much better than any previously published differential-linear cryptanalytic results for DES, though it is inferior to the best previously published cryptanalytic results for DES. In addition, an important merit for these new differential-linear cryptanalytic results is that they are obtained under only two assumptions and thus are more reasonable than those obtained using Biham et al.'s methodology. Like most cryptanalytic results on block ciphers, most of these attacks are far less than practical at present, but they provide a comprehensive understanding of the security of the block ciphers.

The new methodology is a general cryptanalysis technique and can be potentially used to cryptanalyse other block ciphers; and block cipher designers should pay attention to this new methodology when designing ciphers.

The new methodology still requires Assumptions 1 and 2. As a direction for future research on differential-linear cryptanalysis, it would be interesting to

investigate how to further reduce the number of assumptions used, making a more reasonable and more general methodology that could be used in practice.

Acknowledgments. The author is very grateful to Dr. Orr Dunkelman and Dr. Nathan Keller for their discussions on the flaw about CTC2 and to the anonymous referees for their comments on earlier versions of the paper.

References

1. Biham, E., Anderson, R., Knudsen, L.R.: Serpent: A New Block Cipher Proposal. In: Vaudenay, S. (ed.) FSE 1998. LNCS, vol. 1372, pp. 222–238. Springer, Heidelberg (1998)
2. Anderson, R., Biham, E., Knudsen, L.R.: Serpent: a proposal for the Advanced Encryption Standard (1998)
3. Biham, E.: New types of cryptanalytic attacks using related keys. *Journal of Cryptology* 7(4), 229–246 (1994)
4. Biham, E., Biryukov, A.: An improvement of Davies’ attack on DES. *Journal of Cryptology* 10(3), 195–206 (1997)
5. Biham, E., Dunkelman, O., Keller, N.: Enhancing Differential-Linear Cryptanalysis. In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 254–266. Springer, Heidelberg (2002)
6. Biham, E., Dunkelman, O., Keller, N.: Differential-Linear Cryptanalysis of Serpent. In: Johansson, T. (ed.) FSE 2003. LNCS, vol. 2887, pp. 9–21. Springer, Heidelberg (2003)
7. Biham, E., Dunkelman, O., Keller, N.: New Combined Attacks on Block Ciphers. In: Gilbert, H., Handschuh, H. (eds.) FSE 2005. LNCS, vol. 3557, pp. 126–144. Springer, Heidelberg (2005)
8. 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)
9. Biham, E., Shamir, A.: Differential cryptanalysis of DES-like cryptosystems. *Journal of Cryptology* 4(1), 3–72 (1991)
10. Biham, E., Shamir, A.: Differential Cryptanalysis of the Full 16-Round DES. In: Brickell, E.F. (ed.) CRYPTO 1992. LNCS, vol. 740, pp. 487–496. Springer, Heidelberg (1993)
11. Courtois, N.T.: CTC2 and fast algebraic attacks on block ciphers revisited. IACR ePrint report 2007/152 (2007)
12. Courtois, N.T., Pieprzyk, J.: Cryptanalysis of Block Ciphers with Overdefined Systems of Equations. In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 267–287. Springer, Heidelberg (2002)
13. Davies, D.: Investigation of a potential weakness in the DES algorithm (1987)
14. Dunkelman, O.: Techniques for cryptanalysis of block ciphers. Ph.D. thesis, Technion — Israel Institute of Technology, Israel (2006)
15. Dunkelman, O., Keller, N.: Cryptanalysis of CTC2. In: Fischlin, M. (ed.) CT-RSA 2009. LNCS, vol. 5473, pp. 226–239. Springer, Heidelberg (2009)
16. Dunkelman, O., Indestege, S., Keller, N.: A Differential-Linear Attack on 12-Round Serpent. In: Chowdhury, D.R., Rijmen, V., Das, A. (eds.) INDOCRYPT 2008. LNCS, vol. 5365, pp. 308–321. Springer, Heidelberg (2008)

17. Handschuh, H., Naccache, D.: SHACAL. In: Proceedings of the First Open NESSIE Workshop (2000)
18. Hawkes, P.: Differential-Linear Weak Key Classes of IDEA. In: Nyberg, K. (ed.) EUROCRYPT 1998. LNCS, vol. 1403, pp. 112–126. Springer, Heidelberg (1998)
19. Kelsey, J., Schneier, B., Wagner, D.: Key-Schedule Cryptanalysis of IDEA, G-DES, GOST, SAFER, and Triple-DES. In: Koblitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 237–251. Springer, Heidelberg (1996)
20. Kim, J.: Combined differential, linear and related-key attacks on block ciphers and MAC algorithms. Ph.D. thesis, Katholieke Universiteit Leuven, Belgium (2006)
21. Knudsen, L.R.: Cryptanalysis of LOKI91. In: Zheng, Y., Seberry, J. (eds.) AUSCRYPT 1992. LNCS, vol. 718, pp. 196–208. Springer, Heidelberg (1993)
22. Knudsen, L.R.: Truncated and Higher Order Differentials. In: Preneel, B. (ed.) FSE 1994. LNCS, vol. 1008, pp. 196–211. Springer, Heidelberg (1995)
23. Knudsen, L.R., Mathiassen, J.E.: A Chosen-Plaintext Linear Attack on DES. In: Schneier, B. (ed.) FSE 2000. LNCS, vol. 1978, pp. 262–272. Springer, Heidelberg (2001)
24. 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)
25. Langford, S.K.: Differential-linear cryptanalysis and threshold signatures. Ph.D. thesis, Stanford University, USA (1995)
26. Langford, S.K., Hellman, M.E.: Differential-Linear Cryptanalysis. In: Desmedt, Y.G. (ed.) CRYPTO 1994. LNCS, vol. 839, pp. 17–25. Springer, Heidelberg (1994)
27. Lu, J.: Cryptanalysis of block ciphers. Ph.D. thesis, University of London, UK (2008)
28. Lu, J.: New methodologies for differential-linear cryptanalysis and its extensions. Cryptology ePrint Archive, Report 2010/025 (2010), <http://eprint.iacr.org/2010/025>
29. Matsui, M.: Linear Cryptanalysis Method for DES Cipher. In: Helleseht, T. (ed.) EUROCRYPT 1993. LNCS, vol. 765, pp. 386–397. Springer, Heidelberg (1994)
30. Matsui, M.: The First Experimental Cryptanalysis of the Data Encryption Standard. In: Desmedt, Y.G. (ed.) CRYPTO 1994. LNCS, vol. 839, pp. 1–11. Springer, Heidelberg (1994)
31. Matsui, M., Yamagishi, A.: A New Method for Known Plaintext Attack of FEAL Cipher. In: Rueppel, R.A. (ed.) EUROCRYPT 1992. LNCS, vol. 658, pp. 81–91. Springer, Heidelberg (1993)
32. National Bureau of Standards (NBS), Data Encryption Standard (DES), FIPS-46 (1977)
33. Selçuk, A.A.: On probability of success in linear and differential cryptanalysis. *Journal of Cryptology* 21(1), 131–147 (2008)

New Observations on Impossible Differential Cryptanalysis of Reduced-Round Camellia

Ya Liu¹, Leibo Li^{2,3,*}, Dawu Gu¹, Xiaoyun Wang^{2,3,4},
Zhiqiang Liu¹, Jiazhe Chen^{2,3}, and Wei Li^{5,6,7}

¹ Department of Computer Science and Engineering,
Shanghai Jiao Tong University, Shanghai 200240, China
{liyua0611,dwgu,ilu_zq}@sjtu.edu.cn

² Key Laboratory of Cryptologic Technology and Information Security,
Ministry of Education, Shandong University, Jinan 250100, China

³ School of Mathematics, Shandong University,
Jinan 250100, China
{lileibo,jiazhechen}@mail.sdu.edu.cn

⁴ Institute for Advanced Study, Tsinghua University, Beijing 100084, China
xiaoyunwang@mail.tsinghua.edu.cn

⁵ School of Computer Science and Technology, Donghua University,
Shanghai 201620, China

⁶ Shanghai Key Laboratory of Integrate Administration Technologies
for Information Security, Shanghai 200240, China
liwei.cs.cn@gmail.com

⁷ State Key Laboratory of Information Security, Institute of Software,
Chinese Academy of Sciences, Beijing 100190, China

Abstract. Camellia is one of the widely used block ciphers, which has been selected as an international standard by ISO/IEC. In this paper, by exploiting some interesting properties of the key-dependent layer, we improve previous results on impossible differential cryptanalysis of reduced-round Camellia and gain some new observations. First, we introduce some new 7-round impossible differentials of Camellia for weak keys. These weak keys that work for the impossible differential take 3/4 of the whole key space, therefore, we further get rid of the weak-key assumption and leverage the attacks on reduced-round Camellia to all keys by utilizing the multiplied method. Second, we build a set of differentials which contains at least one 8-round impossible differential of Camellia with two FL/FL^{-1} layers. Following this new result, we show that the key-dependent transformations inserted in Camellia cannot resist impossible differential cryptanalysis effectively. Based on this set of differentials, we present a new cryptanalytic strategy to mount impossible differential attacks on reduced-round Camellia.

Keywords: Block Cipher, Camellia, Impossible Differential Cryptanalysis.

* Corresponding author.

1 Introduction

The block cipher Camellia was jointly proposed by NTT and Mitsubishi in 2000 [1]. It was selected as one of the CRYPTREC e-government recommended ciphers in 2002 [4] and as a member of the NESSIE block cipher portfolio in 2003 [20]. In 2005, it was adopted as the international standard by ISO/IEC [6]. Camellia is a 128-bit block cipher. It supports variable key sizes and the number of the rounds depends on the key size, i.e., 18 rounds for a 128-bit key size and 24 rounds for 192/256-bit key sizes. For simplicity, they can be usually denoted as Camellia-128, Camellia-192 and Camellia-256, respectively. Camellia adopts the basic Feistel structure with some key-dependent functions FL/FL^{-1} inserted every six rounds, where these key-dependent transformations must be linear and reversible for any fixed key. The goals for such a design are to provide non-regularity across rounds and to thwart future unknown attacks.

Up to now, many cryptanalytic methods were used to evaluate the security of reduced-round Camellia such as linear cryptanalysis, differential cryptanalysis, higher order differential attack, truncated differential attack, collision attack, square attack and impossible differential attack. Before 2011, most attacks focused on the security of simplified versions of Camellia, which did not take the FL/FL^{-1} and whitening layers into account [9–11, 16, 19, 21–24]. Recently, some attacks involved in the study of the original structure of Camellia. For instance, Chen *et al.* constructed a 6-round impossible differential with the FL/FL^{-1} layer to attack 10-round Camellia-192 and 11-round Camellia-256 [3], Lu, Liu and Li independently improved Chen’s results to attack on reduced-round Camellia [12, 14, 17], Lu *et al.* proposed higher order meet-in-the-middle attacks on 10-round Camellia-128, 11-round Camellia-192 and 12-round Camellia-256 [18].

Impossible differential cryptanalysis was independently proposed by Knudsen [7] and Biham [2]. Its main idea is to use impossible differentials that hold with probability zero to discard the wrong keys until only one key is left. So far, impossible differential cryptanalysis has received much attention and been used to attack a variety of well-known block ciphers such as AES, ARIA, CLEFIA, MISTY1 and so on.

In this paper, we reevaluate the security of reduced-round Camellia with FL/FL^{-1} and whitening layers against impossible differential cryptanalysis from two aspects. On the one hand, we construct some new 7-round impossible differentials of Camellia for weak keys, which work for 75% of the keys. Based on one of them, we mount impossible differential attacks on reduced-round Camellia in the weak-key setting. Then we further propose a multiplied method to extend our attacks for the whole key space. The basic idea is that if the correct key belongs to the set of weak keys, then it will never satisfy the impossible differential. While if the correct key is not a weak key, we get 2-bit conditions about the key. In fact, for the whole key space, we attack 10-round Camellia-128 with about $2^{113.8}$ chosen plaintexts and 2^{120} 10-round encryptions, 11-round Camellia-192 with about $2^{114.64}$ chosen plaintexts and 2^{184} 11-round encryptions as well as 12-round Camellia-256 with about $2^{116.17}$ chosen plaintexts or chosen ciphertexts and 2^{240} 12-round encryptions, respectively. Meanwhile, we can also

extend these attacks to 12-round Camellia-192 and 14-round Camellia-256 with two FL/FL^{-1} layers. On the other hand, by studying some properties of key-dependent functions FL/FL^{-1} , we build a set of differentials which contains at least one 8-round impossible differential of Camellia with two FL/FL^{-1} layers. The length of this impossible differential with two FL/FL^{-1} layers is the same as the length of the longest known impossible differential of Camellia without the FL/FL^{-1} layer given by Wu and Zhang [24]. Consequently, we show that the key-dependent transformations inserted in Camellia cannot resist impossible differential cryptanalysis effectively. Based on this set of differentials, we propose a new cryptanalytic strategy to attack 11-round Camellia-128 with about 2^{122} chosen plaintexts and 2^{122} 11-round encryptions, 12-round Camellia-192 with approximately 2^{123} chosen plaintexts and $2^{187.2}$ 12-round encryptions as well as 13-round Camellia-256 with about 2^{123} chosen plaintexts and $2^{251.1}$ 13-round encryptions (not from the first round but with the whitening layers), respectively. All attacks adopt the early abort technique [15]. In table 1, we summarize our results along with the former known ones on reduced-round Camellia.

Table 1. Summary of the attacks on Reduced-Round Camellia

Key Size	Rounds	Attack Type	Data	Time(Enc)	Memory	Source
128 bits	9†	Square	2^{48} CP	2^{122}	2^{53} Bytes	[10]
	10†	Impossible DC	2^{118} CP	2^{118}	2^{93} Bytes	[17]
	10†	HO-MitM	2^{93} CP	$2^{118.6}$	2^{109} Bytes	[18]
	10†	Impossible DC	$2^{118.5}$ CP	$2^{123.5}$	2^{127} Bytes	[12]
	10(WK)	Impossible DC	$2^{111.8}$ CP	$2^{111.8}$	$2^{84.8}$ Bytes	Section 3.2
	10	Impossible DC	$2^{113.8}$ CP	2^{120}	$2^{84.8}$ Bytes	Section 3.2
	11	Impossible DC	2^{122} CP	2^{122}	2^{102} Bytes	Section 4.4
192 bits	10	Impossible DC	2^{121} CP	$2^{175.3}$	$2^{155.2}$ Bytes	[3]
	10	Impossible DC	$2^{118.7}$ CP	$2^{130.4}$	2^{135} Bytes	[12]
	11†	Impossible DC	2^{118} CP	$2^{163.1}$	2^{141} Bytes	[17]
	11†	HO-MitM	2^{94} CP	$2^{180.2}$	2^{174} Bytes	[18]
	11(WK)	Impossible DC	$2^{112.64}$ CP	$2^{146.54}$	$2^{141.64}$ Bytes	Section 3.3
	11	Impossible DC	$2^{114.64}$ CP	2^{184}	$2^{141.64}$ Bytes	Section 3.3
	12	Impossible DC	2^{123} CP	$2^{187.2}$	2^{160} Bytes	Section 4.3
12†	Impossible DC	$2^{120.1}$ CP	2^{184}	$2^{124.1}$ Bytes	Section 3.5	
256 bits	11	High Order DC	2^{93} CP	$2^{255.6}$	2^{98} Bytes	[5]
	11	Impossible DC	2^{121} CP	$2^{206.8}$	2^{166} Bytes	[3]
	11	Impossible DC	$2^{119.6}$ CP	$2^{194.5}$	2^{135} Bytes	[12]
	12†	HO-MitM	2^{94} CP	$2^{237.3}$	2^{174} Bytes	[18]
	12(WK)	Impossible DC	$2^{121.12}$ CP	$2^{202.55}$	$2^{142.12}$ Bytes	Section 3.4
	12	Impossible DC	$2^{116.17}$ CP/CC	2^{240}	$2^{150.17}$ Bytes	Section 3.4
	13	Impossible DC	2^{123} CP	$2^{251.1}$	2^{208} Bytes	Section 4.2
	14†	Impossible DC	2^{120} CC	$2^{250.5}$	2^{125} Bytes	Section 3.5

DC: Differential Cryptanalysis; CP/CC: Chosen Plaintexts/Chosen Ciphertexts; Enc: Encryptions; †: The attack doesn't include the whitening layers; WK: Weak Key; HO-MitM: Higher Order Meet-in-the-Middle Attack.

The remainder of this paper is organized as follows. Section 2 gives some notations and a brief introduction of Camellia. Section 3 presents several 7-round impossible differentials of Camellia for weak keys. Based on one of them, impossible differential attacks on 10-round Camellia-128, 11-round Camellia-192 and 12-round Camellia-256 are elaborated. Section 4 first constructs a set of differentials which contains at least one 8-round impossible differential of Camellia with two FL/FL^{-1} layers, and then proposes impossible differential attacks on 11-round Camellia-128, 12-round Camellia-192 and 13-round Camellia-256, respectively. Section 5 summarizes this paper.

2 Preliminaries

2.1 Some Notations

- P, C : the plaintext and the ciphertext;
- L_{i-1}, R_{i-1} : the left half and the right half of the i -th round input;
- $\Delta L_{i-1}, \Delta R_{i-1}$: the left half and the right half of the input difference in the i -th round;
- $X | Y$: the concatenation of X and Y ;
- $kw_1 | kw_2, kw_3 | kw_4$: the pre-whitening key and the post-whitening key;
- k_i : the subkey used in the i -th round;
- $kl_i (1 \leq i \leq 6)$: 64-bit keys used in the functions FL/FL^{-1} ;
- $S_r, \Delta S_r$: the output and the output difference of the S-boxes in the r -th round;
- $X \lll j$: left rotation of X by j bits;
- $X_{L(\frac{n}{2})}, X_{R(\frac{n}{2})}$: the left half and the right half of a n -bit word X ;
- $X_i, \bar{X}_{\{i,j\}}, \bar{X}_{\{i \sim j\}}$: the i -th byte, the i -th and j -th bytes and the i -th to the j -th bytes of X ;
- $X^i, X^{(i,j)}, X^{(i \sim j)}$: the i -th bit, the i -th and j -th bits and the i -th to j -th bits of X ;
- \oplus, \cap, \cup : bitwise exclusive-OR (XOR), AND, and OR operations, respectively;
- $0_{(i)}, 1_{(i)}$: consecutive i bits are zero or one.

2.2 Overview of Camellia

Camellia [1] is a 128-bit block cipher. It adopts the basic Feistel structure with keyed functions FL/FL^{-1} inserted every 6 rounds. Camellia uses variable key sizes and the number of rounds depends on the key size, i.e., 18 rounds for a 128-bit key size and 24 rounds for 192/256-bit key sizes. Its round function uses a SPN structure, including the XOR operation with the round subkey, the non-linear transformation S and the linear permutation P . Please refer to [1] for detailed information.

The key schedule algorithm of Camellia applies a 6-round Feistel structure to derive two 128-bit intermediate variables K_A and K_B from K_L and K_R , and then all round subkeys can be generated by K_L, K_R, K_A and K_B . For Camellia-128, the 128-bit key K is used as K_L and K_R is 0. For Camellia-192, the left

128-bit of the key K is used as K_L , and the concatenation of the right 64-bit of the key K and the complement of the right 64-bit of the key K is used as K_R . For Camellia-256, the main key K is separated into two 128-bit variables K_L and K_R , i.e., $K = K_L | K_R$.

3 7-Round Impossible Differentials of Camellia for Weak Keys and Their Applications¹

In this section, we construct some 7-round impossible differentials of Camellia in weak-key setting. Based on one of them, we present impossible differential attacks on 10-round Camellia-128, 11-round Camellia-192 and 12-round Camellia-256 which start from the first round. In addition, we also extend these attacks to 12-round Camellia-192 and 14-round Camellia-256 with two FL/FL^{-1} layers.

3.1 7-Round Impossible Differentials of Camellia for Weak Keys

This section introduces 7-round impossible differentials of Camellia in weak-key setting, which is based on the following lemmas and propositions.

Lemma 1 ([8]). *Let X, X', K be l -bit values, and $\Delta X = X \oplus X'$, then the differential properties of AND and OR operations are:*

$$(X \cap K) \oplus (X' \cap K) = (X \oplus X') \cap K = \Delta X \cap K,$$

$$(X \cup K) \oplus (X' \cup K) = (X \oplus K \oplus (X \cap K)) \oplus (X' \oplus K \oplus (X' \cap K)) = \Delta X \oplus (\Delta X \cap K).$$

Lemma 2 ([3]). *Let ΔX and ΔY be the input and output differences of FL . Then $\Delta Y_R = ((\Delta X_L \cap kl_L) \lll 1) \oplus \Delta X_R, \Delta Y_L = \Delta X_L \oplus \Delta Y_R \oplus (\Delta Y_R \cap kl_R); \Delta X_L = \Delta Y_L \oplus \Delta Y_R \oplus (\Delta Y_R \cap kl_R), \Delta X_R = ((\Delta X_L \cap kl_L) \lll 1) \oplus \Delta Y_R$.*

Proposition 1. *If the output difference of FL is $\Delta Y = (0|0|0|0|d|0|0|0)$, where $d \neq 0$ and $d^{(1)} = 0$, then the input difference of FL should satisfy $\Delta X_{\{2,3,4,6,7,8\}} = 0$.*

Proposition 2. *If the output difference of FL^{-1} is $\Delta X = (0|e|e|e|0|e|e|e)$, and the subkeys of FL^{-1} satisfy that $KL_L^{(9)}$ is 0 or $KL_R^{(8)}$ is 1, then the first byte of input difference ΔY should be zero, where e is a non-zero byte.*

Proposition 3. *Given a 7-round Camellia encryption and a FL/FL^{-1} layer inserted between the fifth and sixth round. If the input difference of the first round is $(0|0|0|0|0|0|0|0, a|0|0|0|c|0|0|0)$, and the subkeys of FL^{-1} satisfy $KL_L^{(9)} = 0$ or $KL_R^{(8)} = 1$, then the output difference $(0|0|0|0|d|0|0|0, 0|0|0|0|0|0|0|0)$ with $d^{(1)} = 0$ is impossible, where a and d are non-zero bytes, c is an arbitrary value (see Fig. 1).*

We also obtain three other impossible differentials under different weak-key assumptions:

$$- (0|0|0|0|0|0|0|0, 0|a|0|0|0|c|0|0) \rightarrow (0|0|0|0|0|d|0|0, 0|0|0|0|0|0|0|0) \text{ with conditions } KL_L^{(17)} = 0 \text{ or } KL_R^{(16)} = 1, \text{ and } d^{(1)} = 0,$$

¹ By Leibo Li, Xiaoyun Wang and Jiazhe Chen. See [13] for more details.

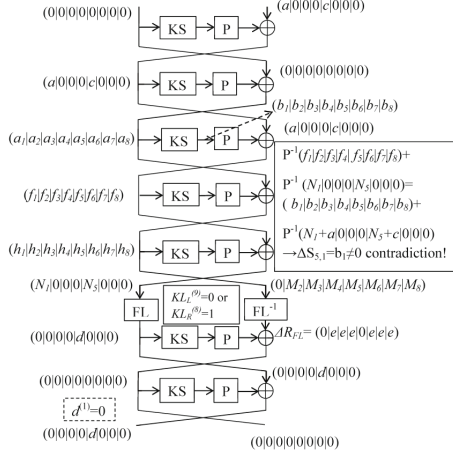


Fig. 1. A 7-Round Impossible Differential for Weak Keys

- $(0|0|0|0|0|0|0|0, 0|0|a|0|0|0|c|0) \rightarrow (0|0|0|0|0|0|d|0, 0|0|0|0|0|0|0|0)$ with conditions $KL_L^{(25)} = 0$ or $KL_R^{(24)} = 1$, and $d^{(1)} = 0$,
- $(0|0|0|0|0|0|0|0, 0|0|0|a|0|0|0|c) \rightarrow (0|0|0|0|0|0|0|d, 0|0|0|0|0|0|0|0)$ with conditions $KL_L^{(1)} = 0$ or $KL_R^{(32)} = 1$, and $d^{(1)} = 0$.

We denote this type of impossible differentials above as **5+2 WKID** (weak-key impossible differentials). Due to the feature of Feistel structure, we also deduce another type of 7-round impossible differentials with the FL/FL^{-1} layer inserted between the second and the third rounds. We call them **2+5 WKID**, which are depicted as follows.

- $(0|0|0|0|0|0|0|0, 0|0|0|0|d|0|0|0) \rightarrow (a|0|0|0|c|0|0|0, 0|0|0|0|0|0|0|0)$ with conditions $KL'_L^{(9)} = 0$ or $KL'_R^{(8)} = 1$, and $d^{(1)} = 0$,
- $(0|0|0|0|0|0|0|0, 0|0|0|0|0|d|0|0) \rightarrow (0|a|0|0|0|c|0|0, 0|0|0|0|0|0|0|0)$ with conditions $KL'_L^{(17)} = 0$ or $KL'_R^{(16)} = 1$, and $d^{(1)} = 0$,
- $(0|0|0|0|0|0|0|0, 0|0|0|0|0|0|d|0) \rightarrow (0|0|a|0|0|0|c|0, 0|0|0|0|0|0|0|0)$ with conditions $KL'_L^{(25)} = 0$ or $KL'_R^{(24)} = 1$, and $d^{(1)} = 0$,
- $(0|0|0|0|0|0|0|0, 0|0|0|0|0|0|0|d) \rightarrow (0|0|0|a|0|0|0|c, 0|0|0|0|0|0|0|0)$ with conditions $KL'_L^{(1)} = 0$ or $KL'_R^{(32)} = 1$, and $d^{(1)} = 0$,

where KL' represents the subkey used in FL -function.

3.2 Impossible Differential Attack on 10-Round Camellia-128

We first propose an attack that works for $3 \times 2^{126} (= \frac{3}{4} \times 2^{128})$ keys, which is mounted by adding one round on the top and two rounds on the bottom of the **5+2 WKID** (See Fig. 2).




Fig. 2. Impossible Differential Attack on 10-Round Camellia-128 for Weak Keys

Data Collection

1. Choose 2^n structures of plaintexts, and each structure contains 2^{32} plaintexts $(L_0, R_0) = (\alpha_1|x_1|x_2|x_3|\alpha_2|x_4|x_5|x_6, P(\beta_1|y_1|y_2|y_3|\beta_2|y_4|y_5|y_6))$, where x_i and y_i ($i = 1, \dots, 6$) are fixed values in each structure, while α_j and β_j ($j = 1, 2$) take all the possible values.
2. For each structure, ask for the encryption of the plaintexts and get 2^{32} ciphertexts. Store them in a hash table H indexed by $C_{R,\{1,5\}}$, the XOR of $C_{R,2}$ and $C_{R,3}$, the XOR of $C_{R,2}$ and $C_{R,4}$, the XOR of $C_{R,2}$ and $C_{R,6}$, the XOR of $C_{R,2}$ and $C_{R,7}$, the XOR of $C_{R,2}$ and $C_{R,8}$. Then by birthday paradox, we get 2^{n+7} pairs of ciphertexts with the differences $(\Delta C_L, \Delta C_R) = (g_1|g_2|g_3|g_4|g_5|g_6|g_7|g_8, 0|f|f|f|0|f|f|f)$, and the differences of corresponding plaintext pairs satisfy $(\Delta L_0, \Delta R_0) = (a|0|0|0|c|0|0|0, P(b_1|0|0|0|b_2|0|0|0))$, where a, c, f and b_i ($i = 1, 2$) are non-zero bytes, and g_i are unknown bytes. For every pair, compute $P^{-1}(\Delta C_L) = P^{-1}(g_1|g_2|g_3|g_4|g_5|g_6|g_7|g_8) = (g'_1|g'_2|g'_3|g'_4|g'_5|g'_6|g'_7|g'_8)$. Keep only the pairs whose ciphertexts satisfy $g'_1 = 0$. The probability of this event is 2^{-8} , thus the expected number of remaining pairs is 2^{n-1} .

Key Recovery

1. For each pair obtained in the data collection phase, guess the 16-bit value $K_{1,\{1,5\}}$, partially encrypt its plaintext $(L_{0,\{1,5\}}, L'_{0,\{1,5\}})$ to get the intermediate value $(S_{1,\{1,5\}}, S'_{1,\{1,5\}})$ and the difference $\Delta S_{1,\{1,5\}}$. Then discard the pairs whose intermediate values do not satisfy $\Delta S_{1,1} = b_1$ and $\Delta S_{1,5} = b_2$. The probability of a pair being kept is 2^{-16} , so the expected number of remaining pairs is 2^{n-17} .
2. In this step, the ciphertext of every remaining pair is considered.
 - (a) Guess the 8-bit value $K_{10,8}$ for every remaining pair, partially decrypt the ciphertext $(C_{R,8}, C'_{R,8})$ to get the intermediate value $(S_{10,8}, S'_{10,8})$ and

the difference $\Delta S_{10,8}$, and discard the pairs whose intermediate values do not satisfy $\Delta S_{10,8} = g'_8$. The expected number of remaining pairs is 2^{n-25} .

- (b) For $l = 2, 3, 4, 6, 7$, guess the 8-bit value $K_{10,l}$. For every remaining pair, partially decrypt the ciphertext $(C_{R,l}, C'_{R,l})$ to get the intermediate value $(S_{10,l}, S'_{10,l})$ and the difference $\Delta S_{10,l}$, and keep only the pairs whose intermediate values satisfy $\Delta S_{10,l} = g'_l \oplus g'_5$. Since each pair will remain with probability 2^{-40} , the expected number of remaining pairs is 2^{n-65} .
 - (c) Guess the 8-bit value $K_{10,1}$, partially decrypt the ciphertext $C_{R,1}$ of every remaining pair to get the intermediate value $S_{10,1}$, which is also the value of $S'_{10,1}$.
 - (d) Partially decrypt (S_{10}, S'_{10}) to get the intermediate values $(R_{9,5}, R'_{9,5})$, and discard the pairs whose intermediate values do not satisfy $\Delta R_{9,5}^{(1)} = 0$. As the probability of a pair being discarded is 0.5, the expected number of remaining pairs is 2^{n-66} .
3. For every remaining pair, guess the 8-bit value $K_{9,5}$, partially decrypt the output value $(R_{9,5}, R'_{9,5})$ to get the intermediate value $(S_{9,5}, S'_{9,5})$ and the difference $\Delta S_{9,5}$. If there is a pair satisfying $\Delta S_{9,5} = \Delta C_{R,2}$, we discard the guessed key and try another one. Otherwise we exhaustively search for the remaining 48 bits of the key under this guessed key, if the correct key is obtained, we halt the attack; otherwise, another key guess should be tried.

Complexity. Since the probability of the event $\Delta S_{9,5} = \Delta C_{R,2}$ in step 3 of key recovery phase is 2^{-8} , the expected number of remaining guesses for 72-bit target subkeys is about $\epsilon = 2^{80} \times (1 - 2^{-8})^{2^{n-66}}$. If we choose $\epsilon = 1$, then n is 79.8, and the proposed attack requires $2^{n+32} = 2^{111.8}$ chosen plaintexts. The time and memory complexities are dominated by step 2 of data collection phase, which are about $2^{111.8}$ 10-round encryptions and $2^{n-1} \times 4 \times 2^4 = 2^{84.8}$ bytes.

Extending the Attack to the Whole Key Space. On the basis of the above impossible differential attack for weak keys, we construct a multiplied attack on 10-Round Camellia-128.

- **Phase 1.** Perform an impossible differential attack by using the **5+2 WKID** $(0|0|0|0|0|0|0|0, a|0|0|0|c|0|0|0) \rightarrow (0|0|0|0|d|0|0|0, 0|0|0|0|0|0|0|0)$. This phase is extremely similar to the weak-key attack that is described above. However, it is slightly different when the attack is finished. That is, if there is a key kept, then the key is the correct key, and we halt the procedure of the attack. Otherwise, we conclude that the correct key does not belong to this set of weak keys, which means that $kl_1^{(9)} = 1$ and $kl_2^{(8)} = 0$. In this case, we get 2-bit information of the key and perform the next phase.
- **Phases 2 to 4.** Perform an impossible differential attack by using each **5+2 WKID** in the following:

$$(0|0|0|0|0|0|0|0, 0|a|0|0|0|c|0|0) \rightarrow (0|0|0|0|0|d|0|0, 0|0|0|0|0|0|0|0),$$

$$(0|0|0|0|0|0|0|0, 0|0|a|0|0|0|c|0) \rightarrow (0|0|0|0|0|0|d|0, 0|0|0|0|0|0|0|0),$$

$$(0|0|0|0|0|0|0|0|0, 0|0|0|a|0|0|0|c) \leftrightarrow (0|0|0|0|0|0|0|d, 0|0|0|0|0|0|0|0).$$

The procedure is similar to Phase 1, and either recover the correct key or get another 2-bit information about the key and execute the next phase.

- **Phase 5.** Announce the intermediate key $K_A^{(95,103,111,119)} = 0$ and $K_A^{(6,14,22,30)} = 1$, then exhaustively search for the remaining 120-bit value of K_A and recover the key K_L .

The upper bound of the time complexity is $2^{111.8} \times 4 + 2^{120} \approx 2^{120}$. The data complexity is about $2^{113.8}$. The memory could be reused in different phase, so the memory requirement is about $2^{84.8}$ bytes.

3.3 Attack on 11-Round Camellia-192

We add one round on the bottom of 10-round attack and give an attack on 11-round Camellia-192.

Data Collection. Choose $2^{80.64}$ structures of plaintexts. Each structure contains 2^{32} plaintexts satisfying $(L_0, R_0) = (\alpha_1|x_1|x_2|x_3|\alpha_2|x_4|x_5|x_6, P(\beta_1|y_1|y_2|y_3|\beta_2|y_4|y_5|y_6))$, where x_i and y_i ($i = 1, \dots, 6$) are fixed values in each structure, while α_j and β_j ($j = 1, 2$) take all the possible values. Ask for the encryption of the corresponding ciphertext for each plaintext, compute $P^{-1}(C_R)$ and store the plaintext-ciphertext pairs (L_0, R_0, C_L, C_R) in a hash table indexed by 8-bit value $(P^{-1}(C_R))_1$. By birthday paradox, we get $2^{135.64}$ pairs whose ciphertext differences satisfy $P^{-1}(\Delta C_L) = (h'_1|h'_2|h'_3|h'_4|h'_5|h'_6|h'_7|h'_8)$ and $P^{-1}(\Delta C_R) = (0|g'_2|g'_3|g'_4|g'_5|g'_6|g'_7|g'_8)$, where h'_i and g'_i are unknown values.

Key Recovery

1. For $l = 1, 5$, guess the 8-bit value of $K_{1,l}$, partially encrypt their plaintext $(L_{0,l}, L'_{0,l})$ and discard the pairs whose intermediate value do not satisfy $\Delta S_{1,l} = (P^{-1}(\Delta R_0))_l$. The expected number of remaining pairs is $2^{119.64}$.
2. In this step, we consider the ciphertext of each remaining pair.
 - (a) For $l = 1, 2, 3, 4, 6, 7, 8$, guess the 8-bit value of $K_{11,l}$. Partially decrypt the ciphertext $(C_{R,l}, C'_{R,l})$ and keep only the pairs which satisfy $\Delta S_{11,l} = h'_l$. The expected number of remaining pairs is $2^{63.64}$.
 - (b) Guess the 8-bit value $K_{11,5}$. Partially decrypt the ciphertext $(C_{R,5}, C'_{R,5})$, then compute the intermediate value (R_{10}, R'_{10}) , where $\Delta R_{10} = (0|f|f|f|0|f|f|f)$ and $f = \Delta S_{11,5} \oplus h'_5$.
3. Application of the 10-round attack.
 - (a) Guess the 8-bit value $K_{10,8}$, partially decrypt $(R_{10,8}, R'_{10,8})$ and discard the pairs whose intermediate values do not satisfy $\Delta S_{10,8} = g'_8$. The expected number of remaining pairs is $2^{63.64} \times 2^{-8} = 2^{55.64}$.
 - (b) For $l = 2, 3, 4, 6, 7$, guess the 8-bit value $K_{10,l}$. Partially decrypt the intermediate value $(R_{10,l}, R'_{10,l})$ and keep only the pairs whose intermediate values satisfy $\Delta S_{10,l} = g'_l \oplus g'_5$. The expected number of remaining pairs is $2^{15.64}$.

- (c) Guess the 8-bit value $K_{10,1}$, partially decrypt the intermediate value $R_{10,1}$ and calculate the intermediate values $(R_{9,5}, R'_{9,5})$. Discard the pairs whose intermediate values do not satisfy $\Delta R_{9,5}^{(1)} = 0$. Then the expected number of remaining pairs is $2^{14.64}$.
- (d) Guess the 8-bit value $K_{9,5}$, partially decrypt the intermediate value $(R_{9,5}, R'_{9,5})$ to get the difference $\Delta S_{9,5}$. If there is a pair satisfies $\Delta S_{9,5} = \Delta R_{10,2}$, we discard the guessed key and try another one. Otherwise we exhaustively search for the remaining 48 bits of K_L and K_R under this key, if the correct key is obtained, we halt the attack; otherwise, another key should be tried.

Complexity. The data complexity of the attack is $2^{112.64}$ chosen plaintexts. The time complexity is dominated by step 3 (d) which requires about $2^{144} \times (1 + (1 - 2^{-8}) + (1 - 2^{-8})^2 + \dots + (1 - 2^{-8})^{2^{13.7}-1}) \times 2 \times \frac{1}{11} \times \frac{1}{8} \approx 2^{146.54}$ 11-round encryptions. The memory complexity is about $2^{133.56} \times 4 \times 2^4 = 2^{141.64}$ bytes.

Reduce the Time Complexity to $2^{138.54}$. Assume 16-bit value α_2 and β_2 are fixed in data collection phase of above attack, then we can collect $2^{n+31} \times 2^{-8} = 2^{n+23}$ pairs, where n represents the number of structures. Nevertheless, it is unnecessary for us to guess 8-bit subkey $K_{1,5}$ in this case. Then there are totally 136-bit values of subkey to be guessed in the attack, therefore, the expected number of remaining guesses of target subkey is about $\epsilon = 2^{136} \times (1 - 2^{-8})^{2^n - 90}$ after the attack. If we chose $\epsilon = 1$, n is 104.56. Then the data complexity increases to $2^{n+16} = 2^{120.56}$, but the time complexity reduces to $2^{138.54}$, the memory requirement reduces to $2^{133.56}$ bytes.

Extending the Attack to the Whole Key Space. Similar to 10-round attack on Camellia-128, we mount a multiplied attack on Camellia-192 for the whole key space. The time complexity is about $4 \times 2^{146.54} + 2^{192} \times (1 - \frac{3}{4})^4 = 2^{184}$ 10-round encryptions. The data and memory complexities are approximately $2^{114.64}$ chosen plaintexts and $2^{141.64}$ bytes, respectively.

3.4 The Attack on 12-Round Camellia-256

We add one round on the bottom of 11-round attack, and present a 12-round attack on Camellia-256. The attack procedure is similar to the 11-round attack. First choose $2^{81.17}$ structures and collect $2^{144.17}$ plaintext-ciphertext pairs in data collection phase. After guessing the subkey $K_{1,\{1,5\}}$, we guess the 64-bit value K_{12} and compute the intermediate value (R_{11}, R'_{11}) , then apply the 11-round attack to perform the remaining steps. In summary, the proposed attack requires $2^{81.17+32} = 2^{113.17}$ chosen plaintexts. The time complexity is about $2^{210.55}$ 12-round encryptions, and the memory requirement is about $2^{150.17}$ bytes. Similar to the above subsection, the time complexity and memory requirement can also reduce to $2^{202.55}$ and $2^{142.12}$, respectively, but data complexity increases to $2^{121.12}$ in this case.

We also construct another type of impossible differential attack of Camellia-256, which adds four rounds on the top and one round on the bottom of the **2+5 WKID** (see section 3.1). The attack is performed under the chosen ciphertext attack scenario. Similar to the attack based on the **5+2 WKID**, the data and time complexity are about $2^{113.17}$ and $2^{216.3}$, respectively.

Extending the Attack to the Whole Key Space. On the basis of two types of impossible differential attacks for weak keys, we mount a multiplied attack on 12-round Camellia-256 for the whole key space as below.

- **Phases 1 to 8.** Perform impossible differential attacks by using of all conditional impossible differentials **2+5 WKID** list in section 3.1. For each phase, if success, output the actual key, else perform the next phase.
- **Phase 9.** Announce 16-bit value of the master key $K_R^{(31,39,47,55,95,103,111,119)} = 0$ and $K_R^{(6,14,22,30,70,78,86,94)} = 1$, then exhaustively search for the remaining 240-bit value of K_R , K_L and recover the actual key.

The expected time of the attack is $2^{216.3} \times 8 + 2^{256} \times (\frac{1}{4})^8 \approx 2^{240}$ encryptions, and the expected data complexity is about $2^{116.17}$.

3.5 The Attacks Including Two FL/FL^{-1} Layers

If we do not start from the first round, we can take the attacks that include two FL/FL^{-1} layers into account. By exploiting some new properties of FL and FL^{-1} , we mount impossible differential attacks on variants of 14-round Camellia-256 and 12-round Camellia-192. Specifically, we attack 14-round Camellia-256 from round 10 to round 23 with about 2^{120} chosen ciphertexts, $2^{250.5}$ 14-round encryptions and 2^{125} bytes of memory, and 12-round Camellia-192 from round 3 to round 14 with about $2^{120.1}$ chosen plaintexts, $2^{180.1}$ 12-round encryptions and $2^{124.1}$ bytes of memory. The detailed information can be found in [13].

4 8-Round Impossible Differentials of Camellia and Their Applications²

In this section, we first present a method to construct a set of differentials, which contains at least one 8-round impossible differential of Camellia with two FL/FL^{-1} layers for any fixed key. Based on this set of differentials, we propose a new strategy to attack on reduced-round Camellia-128/192/256 with the whitening and FL/FL^{-1} layers.

4.1 The Construction of 8-Round Impossible Differentials of Camellia

We first illustrate some properties of FL/FL^{-1} .

² By Ya Liu, Dawu Gu, Zhiqiang Liu and Wei Li.

Proposition 4. *If the input difference of FL is $(a|0|0|0|a'|0|0|0)$, where $a^{(1)} = a^{(8)} = 0$ and*

$$a^{(i)} = \begin{cases} 0, & kl_L^{(i+1)} = 0; \\ a^{(i+1)}, & kl_L^{(i+1)} = 1; \end{cases} \text{ for } 1 \leq i \leq 7,$$

then the output difference of FL is $(a|0|0|0|0|0|0|0)$.

By Propositions 4, we construct an 8-round impossible differential of Camellia with two FL/FL^{-1} layers for any fixed subkey.

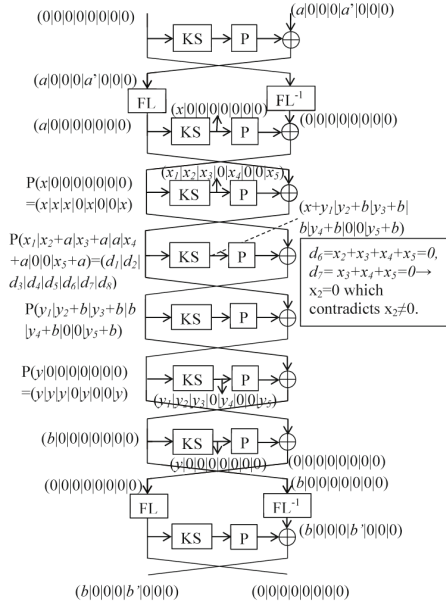


Fig. 3. The Structure of 8-Round Impossible Differential of Camellia

Proposition 5. *For an 8-round Camellia encryption with two FL/FL^{-1} layers inserted after the first and seventh rounds, the input difference of the first round is $(0|0|0|0|0|0|0|0, a|0|0|0|a'|0|0|0)$ and the output difference of the eighth round is $(b|0|0|0|b'|0|0|0, 0|0|0|0|0|0|0|0)$ with a and b being nonzero bytes and $a^{(1)} = b^{(1)} = a^{(8)} = a'^{(8)} = 0$. Four subkeys $kl_i (i = 1, \dots, 4)$ are used in two FL/FL^{-1} layers. If a' and b' satisfy the following equations:*

$$a^{(i)} = \begin{cases} 0, & \text{if } kl_1^{(i+1)} = 0; \\ a^{(i+1)}, & \text{if } kl_1^{(i+1)} = 1; \end{cases} \quad b^{(i)} = \begin{cases} 0, & \text{if } kl_4^{(i+1)} = 0; \\ b^{(i+1)}, & \text{if } kl_4^{(i+1)} = 1; \end{cases} \text{ for } 1 \leq i \leq 7,$$

then $(0|0|0|0|0|0|0|0, a|0|0|0|a'|0|0|0) \rightarrow_8 (b|0|0|0|b'|0|0|0, 0|0|0|0|0|0|0|0)$ is an 8-round impossible differential of Camellia with two FL/FL^{-1} layers (See Fig. 3).

For any fixed subkey, an 8-round impossible differential with two FL/FL^{-1} layers can be constructed. Each possible value of $kl_1^{(2\sim 8)} \mid kl_4^{(2\sim 8)}$ corresponds to the existence of an 8-round impossible differential. All possible values of $kl_1^{(2\sim 8)} \mid kl_4^{(2\sim 8)}$ are from $0_{(14)}$ to $1_{(14)}$. Denote their corresponding impossible differentials by Δ_i for $0 \leq i \leq 2^{14} - 1$. Let A be a set including all differentials $\Delta_i (0 \leq i \leq 2^{14} - 1)$, i.e., $A = \{\Delta_i \mid 0 \leq i \leq 2^{14} - 1\}$. According to Proposition 5, 8-round differentials of A must have the form: $\Delta = (0|0|0|0|0|0|0|0, a|0|0|0|a'|0|0|0) \rightarrow_8 (b|0|0|0|b'|0|0|0, 0|0|0|0|0|0|0|0)$ with a and b being nonzero bytes and $a^{(1)} = b^{(1)} = a^{(8)} = b^{(8)} = 0$. Among them, a' and b' are either zero or nonzero bytes. We divide all differentials of A into three cases: (1) $a' = b' = 0$, (2) $a' = 0$ and $b' \neq 0$, or $a' \neq 0$ and $b' = 0$, (3) $a' \neq 0$ and $b' \neq 0$.

By proposition 5, we only know the existence of an 8-round impossible differential of Camellia with two FL/FL^{-1} layers for any fixed key, but cannot distinguish it from other differentials of A . Therefore, we require to propose a new attack strategy to recover the correct key based on this set of differentials.

The Attack Strategy. Select a differential Δ_i from A . Based on it, we mount an impossible differential attack on reduced-round Camellia given enough plaintext pairs.

1. If one subkey will be kept, we recover the secret key by the key schedule and verify whether it is correct by some plaintext-ciphertext pairs. If success, halt this attack. Otherwise, try another differential $\Delta_j (j \neq i)$ of A and perform a new impossible differential attack.
2. If no subkeys or more than one subkeys are left, select another differential of A to execute a new impossible differential attack.

Our attack strategy can really recover the correct key. As a matter of fact, if Δ_i is an impossible differential, we make sure the expected number of remaining wrong keys will be almost zero given enough chosen plaintexts. Therefore, we only consider those differentials which result in one subkey remaining. By Proposition 5, we know the set A contains at least one impossible differential. So we try each differential of A until the correct key is recovered. The worst scenario is that the correct key is retrieved from the last try.

4.2 Impossible Differential Attack on 13-Round Camellia-256

Based on three scenarios of differentials in A , we present an impossible differential attack on 13-round Camellia-256 with the FL/FL^{-1} and whitening layers from rounds 4 to 16. Let $k_a \triangleq kw_1 \oplus k_4, k_b \triangleq kw_2 \oplus k_5, k_c \triangleq kw_4 \oplus k_{16}, k_d \triangleq kw_3 \oplus k_{15}, k_e \triangleq kw_4 \oplus k_{14}$. We use these equivalent subkeys k_a, k_b, k_c, k_d and k_e instead of the round subkeys k_4, k_5, k_{14}, k_{15} and k_{16} so as to remove the whitening layers. In the following, we will illustrate this attack.

Case 1 $a' = b' = 0$: The differential $\Delta = (0|0|0|0|0|0|0|0, a|0|0|0|0|0|0|0) \rightarrow_8 (b|0|0|0|0|0|0|0, 0|0|0|0|0|0|0|0)$, where a and b are nonzero bytes and $a^{(1)} = b^{(1)} = 0$ (See Fig. 4).




Fig. 4. Impossible Differential Attack on 13-round Camellia-256 for Case 1

Data Collection. Select a structure of plaintexts, which contains 2^{55} plaintexts with the following form:

$$(P(\alpha_1|x_1|x_2|x_3|x_4|x_5|x_6|x_7), P(\alpha_2|\alpha_3|\alpha_4|\alpha_5|\alpha_6|x_8|x_9|\alpha_7)), \quad (1)$$

where $\alpha_5^{(1)}, x_i (1 \leq i \leq 9)$ are fixed and $\alpha_j (1 \leq j \leq 7, i \neq 5), \alpha_5^{(2 \sim 8)}$ takes all possible values. Clearly, each structure forms 2^{109} plaintext pairs, the differences of which have the form: $(P(g_1|0|0|0|0|0|0|0), P(g_2|g_3 \oplus a|g_4 \oplus a|g_5 \oplus a|0|g_6 \oplus a))$ with a and $g_i (1 \leq i \leq 6)$ being nonzero bytes and $a^{(1)}=0$. We take all possible values of $(\alpha_5^{(1)}, x_4, x_8, x_9)$ and 2^{43} different values of $x_i (1 \leq i \leq 7, i \neq 4)$ to derive 2^{68} special structures. In total, there are 2^{123} chosen plaintexts which form 2^{177} plaintext pairs. Encrypt these plaintext pairs to obtain the corresponding ciphertext pairs. If the right halves of their ciphertexts differences have the form: $P(h_1|h_2 \oplus b|h_3 \oplus b|h_4 \oplus b|0|h_5 \oplus b|0|h_6 \oplus b)$ with $b^{(1)} = 0$, then these pairs will be kept. The expected number of remaining pairs is about 2^{160} .

Key Recovery

1. Guess $k_{a,1}$. For each remaining pair, check whether the equation $\Delta S_{4,1} = (P^{-1}(\Delta P_R))_1$ holds. If $\Delta S_{4,1} \neq (P^{-1}(\Delta P_R))_1$ for some pair, then this pair will be discarded. Next guess each possible value of $k_{a,l}$ for $l = 2, 3, 5, 8$. Keep only the pairs satisfying $\Delta S_{4,l} = (P^{-1}(\Delta P_R))_l \oplus (P^{-1}(\Delta P_R))_4$. The expected number of remaining pairs is about 2^{120} . Finally, guess $k_{a,\{4,6,7\}}$ and compute the inputs of the fifth round for each remaining pair.
2. Guess $k_{b,1}$ and test whether $\Delta S_{5,1}$ is equal to $(P^{-1}(\Delta P_L))_1$ for each remaining pair. If $\Delta S_{5,1} \neq (P^{-1}(\Delta P_L))_1$ for one pair, then this pair will be removed. Finally, about 2^{112} pairs will be kept.

3. Guess $k_{c,l}$ for $2 \leq l \leq 8$. Verify whether $\Delta S_{16,l}$ is equal to $(P^{-1}(\Delta C_L))_l$ for every remaining pair. If $\Delta S_{16,l} \neq (P^{-1}(\Delta C_L))_l$ for some pair, then this pair is discarded. The expected number of remaining pairs is about 2^{56} . Next guess $k_{c,1}$ and compute the outputs of the 15-th round for each remaining pair.
4. Guess $k_{d,l}$ for $l = 1, 2, 3, 5, 8$. For each remaining pair, verify whether the equations $\Delta S_{15,1} = (P^{-1}(\Delta C_R))_1$ and $\Delta S_{15,j} = (P^{-1}(\Delta C_R))_j \oplus (P^{-1}(\Delta C_R))_4$ ($j = 2, 3, 5, 8$) hold. The probability that to happen is about 2^{-40} . Thus about 2^{16} pairs will be kept. Next guess other bytes of k_d and calculate the outputs of the 14-th round.
5. Guess $k_{e,1}$ and compute the output difference of the S-Boxes in the 14-th round. If $\Delta S_{14,1}$ is equal to $(P^{-1}(\Delta L_{14}))_1$, then we remove this value of $k_{e,1}$ with $(k_a, k_{b,1}, k_c, k_d)$. The probability of this event is about 2^{-8} . After trying all possible values of $(k_a, k_{b,1}, k_c, k_d, k_{e,1})$, if only one joint subkey remains, then Δ is likely to be an impossible differential. At this time, we recover the secret key by the key schedule and verify whether it is correct by some plaintext-ciphertext pairs. If no subkeys or more than one subkeys are left, then Δ is possible to exist. At this time, try another differential of A . As a matter of fact, if Δ is an impossible differential, the expected number of remaining wrong subkeys is about $2^{208} \times (1 - 2^{-8})^{2^{16}} \approx 2^{-161.4}$. We consider that all wrong subkeys are removed and only the correct subkey is left. Therefore, we require to perform the following step only if one subkey will be kept.
6. According to the key schedule of Camellia-256, we can recover the secret key from this unique 208-bit subkey $(k_a, k_{b,1}, k_c, k_d, k_{e,1})$. As a matter of fact, we guess K_B and K_R , and then calculate K_L and K_A by property 4 of [18]. Finally, the number of remaining main keys is approximately 2^{48} . By about 2^{48} trail encryptions, if some key is correct, stop the attack. Otherwise, try another differential of A .

Case 2 $a' = 0$ and $b' \neq 0$, or $a' \neq 0$ and $b' = 0$: We only attack a special scenario, i.e., $a' = 0$ and $b'(1 \sim 7) = b(2 \sim 8)$. Others can be attacked in the similar way. At this time, the differential is $\Delta' = (0|0|0|0|0|0|0|0, a|0|0|0|0|0|0|0) \rightarrow_8 (b|0|0|0|b'|0|0|0, 0|0|0|0|0|0|0|0)$, where a , b and b' are non-zero bytes, $b'(1 \sim 7) = b(2 \sim 8)$ and $a^{(1)} = b^{(1)} = b'^{(8)} = 0$.

Data Collection. We apply 2^{68} special structures of above Case 1. Totally, there are 2^{123} chosen plaintexts which form 2^{177} pairs.

Key Recovery

1. Guess $k_{c,l}$ for $2 \leq l \leq 8$ and $l \neq 5$. Verify whether the equation $\Delta S_{16,l} = (P^{-1}(\Delta C_L))_l$ holds for every remaining pair. If $\Delta S_{16,l} \neq (P^{-1}(\Delta C_L))_l$ for some pair, then this pair is discarded. The expected number of remaining pairs is about 2^{129} . Next guess $k_{c,\{1,5\}}$ and compute the outputs of the 15-th round for each remaining pair.

2. We first guess $k_{d,1}$ and check whether the equation $\Delta S_{15,1} = (P^{-1}(\Delta C_R))_1$ holds for each remaining pair. If $\Delta S_{15,1} \neq (P^{-1}(\Delta C_R))_1$ for one pair, then this pair will be removed. Next guess $k_{d,8}$ and keep only the pairs satisfying $\Delta S_{15,8}^{(1)} = (P^{-1}(\Delta C_R))_8^{(1)}$. Finally, guess $k_{d,\{2\sim 7\}}$. Test whether $\Delta S_{15,l} = (P^{-1}(\Delta C_R))_l \oplus (((P^{-1}(\Delta C_R))_8 \oplus \Delta S_{15,8})^{(2\sim 8)}|0)$ for $l = 6, 7$ and $\Delta S_{15,l} = (P^{-1}(\Delta C_R))_l \oplus (P^{-1}(\Delta C_R))_8 \oplus \Delta S_{15,8} \oplus (P^{-1}(\Delta C_R))_7 \oplus \Delta S_{15,7}$ for $l = 2, 3, 4, 5$. The total probability of this step is about 2^{-57} . So the expected number of remaining pairs is approximately 2^{72} . Compute the outputs of the 14-th round for each remaining pair.
3. Guess $k_{e,l}$ for $l = 1, 5$. Verify whether the equation $\Delta S_{14,l} = (P^{-1}(\Delta L_{14}))_l$ holds for each remaining pair. If this equation is correct for some pair, then this pair will be kept. The probability of this event is about 2^{-16} . About 2^{56} pairs will be kept.
4. Guess each possible value of k_a as like Case 1. The expected number of remaining pairs is about 2^{16} . Calculate the inputs of the fifth round.
5. Guess $k_{b,1}$. This step is similar to Step 5 of Case 1. If only one joint subkey is left, then we consider Δ' is an impossible differential and recover the secret key by the key schedule. Otherwise try another differential of A . In fact, the expected number of remaining wrong subkeys is approximately $2^{216} \times (1 - 2^{-8})^{2^{16}} \approx 2^{-153.4}$ if Δ' is an impossible differential.
6. This step is similar to Step 6 of Case 1. Finally, about 2^{40} keys will be left. By about 2^{40} trail encryptions, if some key is correct, stop the attack. Otherwise, try another differential of A .

Case 3 $a' \neq 0$ and $b' \neq 0$: We only discuss an example, i.e., $a'^{(1\sim 7)} = a^{(2\sim 8)}$ and $b'^{(1\sim 7)} = b^{(2\sim 8)}$. The differential is $\Delta'' = (0|0|0|0|0|0|0|0, a|0|0|0|a'|0|0|0) \rightarrow_8 (b|0|0|0|b'|0|0|0, 0|0|0|0|0|0|0|0)$, where a, b, a' and b' are nonzero bytes and $a^{(1)} = b^{(1)} = a'^{(8)} = b'^{(8)} = 0$.

Data Collection. Continue to adopt 2^{123} chosen plaintexts of Case 1. Because each structure of Case 1 takes all possible values of $\alpha_5^{(1)}, x_4, x_8$ and x_9 , 2^{123} chosen plaintexts of Case 1 are equivalent to 2^{43} structures, each of which contains 2^{80} plaintexts with the form: $(P(\beta_1|y_1|y_2|y_3|\beta_2|y_4|y_5|y_6), \beta_3|\beta_4|\beta_5|\beta_6|\beta_7|\beta_8|\beta_9|\beta_{10})$, where $y_i (1 \leq i \leq 6)$ are fixed and $\beta_j (1 \leq j \leq 10)$ takes all possible values. It is obvious that one structure generates 2^{159} pairs. Totally, there are approximately 2^{202} plaintext pairs satisfying the input differences.

Key Recovery

1. Guess each byte of $k_c, k_d, k_{e,\{1,5\}}$. This step is similar to above Case 2. After this step, about 2^{81} pairs will be kept.
2. Guess $k_{a,1}, k_{a,8}, k_{a,\{6,7\}}, k_{a,\{2\sim 5\}}$ and $k_{b,5}$ in turn. The expected number of remaining pairs is about 2^{16} . Compute the inputs of the 5-th round for each remaining pair.
3. Guess $k_{b,1}$ and test whether $\Delta S_{5,1}$ is equal to $(P^{-1}(\Delta P_L))_1$ for each remaining pair. If $\Delta S_{5,1} = (P^{-1}(\Delta P_L))_1$ for some pair, then this guessed key are

removed. After guessing all possible subkeys, if only one joint subkey is left, then we consider Δ'' is an impossible differential. At this moment, we execute the following step. Otherwise try another differential of A . As a matter of fact, the expected number of remaining wrong subkeys is approximately $2^{224} \times (1 - 2^{-8})^{2^{16}} \approx 2^{-145.4}$ if Δ'' is an impossible differential.

4. Similarly, we recover the secret key from this subkey. The number of remaining main keys is approximately 2^{32} . By about 2^{32} trail encryptions, if some key is correct, stop the attack. Otherwise, try another differential of A .

Complexity. We calculate that the total time complexities of Cases 1 to 3 are about 2^{216} 1-round encryptions, 2^{224} 1-round encryptions and $2^{240.8}$ 1-round encryptions, respectively. Thus the total time complexity is at most $2^{14} \times 2^{240.8} \times \frac{1}{13} \approx 2^{251.1}$ 13-round encryptions. Furthermore, the total data and memory complexities are 2^{123} chosen plaintexts and 2^{208} bytes, respectively.

4.3 Impossible Differential Attack on 12-Round Camellia-192

In this section, we attack 12-round Camellia-192 from rounds 4 to 15 with the 8-round differentials inserted rounds 6 to 13. Some equivalent subkeys k_a and k_b are defined as before. In addition, let $k'_d = kw_4 \oplus k_{15}$ and $k'_e = kw_3 \oplus k_{14}$.

Case 1 $a' = b' = 0$: The differential is Δ .

We select the same plaintexts of Case 1 mentioned in section 4.2, i.e., 2^{123} chosen plaintexts and 2^{177} pairs. Encrypt them and keep those pairs whose ciphertext differences have the form: $(P(h_2|h_3 \oplus b|h_4 \oplus b|b|h_5 \oplus b|0|0|h_6 \oplus b), P(h_1|0|0|0|0|0|0|0))$, where b and $h_i (1 \leq i \leq 6)$ are nonzero bytes and $b^{(1)} = 0$. The expected number of remaining pairs is about 2^{104} .

Guess all possible values $(k_a, k_{b,1}, k'_d, k'_{e,1})$ and discard those subkeys which acquire the input and output differences of Δ . This step is similar to section 4.2. If Δ is an impossible differential, about $2^{144} \times (1 - 2^{-8})^{2^{16}} \approx 2^{-225.4}$ wrong subkeys are expected to remain. Therefore, we will recover the secret key by the key schedule of Camellia-192 only if one subkey is left. Otherwise, try another differential of A . By the key schedule of Camellia-192, we derive 2^{48} candidates of the secret key from the 144-bit subkey $(k_a, k_{b,1}, k'_d, k'_{e,1})$. By about 2^{48} trail encryptions, if the correct key is retrieved, halt the attack. Otherwise, try another differential of A .

Case 2 $a' = 0, b' \neq 0$ or $a' \neq 0, b' = 0$: For simplicity, we consider a special differential Δ' .

We still select 2^{123} plaintexts of above Case 1. In total, there are 2^{68} special structures, each of which contains 2^{55} plaintexts. Encrypt these plaintext pairs. If the right halves of their ciphertexts differences have the form: $P(h|0|0|0|h'|0|0|0)$ with h and h' being nonzero bytes, then these pairs will be kept. Consequently, the expected number of remaining pairs is about 2^{129} . Similarly, we can remove some subkeys $(k_a, k_{b,1}, k'_d, k'_{e,\{1,5\}})$ which obtain the input and output differences of Δ' for some pair. If only one subkey is left, we recover the secret key by the key

schedule. Otherwise, try another differential of A . In fact, if Δ' is an impossible differential, about $2^{-217.4} (\approx 2^{152} \times (1 - 2^{-8})^{2^{16}})$ wrong subkeys will be left.

Case 3 $a' \neq 0, b' \neq 0$: A special differential Δ'' will be considered.

The similar attacking procedure can be performed as before. We select 2^{43} structure, each of which contains 2^{80} plaintexts. Totally, they can form 2^{202} pairs. After filtering some pairs by the ciphertext differences, about 2^{154} pairs are expected to remain. The following steps can be performed in the similar way.

We found that the time complexity of Case 3 is maximal. Therefore, the total time complexity is at most $2^{14} \times 2^{173.2} \approx 2^{187.2}$ 12-round encryptions. The data and memory complexities are 2^{123} chosen plaintexts and 2^{160} bytes, respectively.

4.4 Impossible Differential Attack on 11-Round Camellia-128

For Camellia-128, we put two additional rounds on the top and one additional round on the bottom of 8-round differentials. Based on it, we attack 11-round Camellia-128 from rounds 4 to 14. Similarly, we divide all possible differentials into three different cases as before. For Case 1, we take 2^{67} special structures (1). Totally, the data complexity is 2^{122} chosen plaintexts which form 2^{176} pairs. Encrypt these pairs to acquire the corresponding ciphertext pairs. Then we discard some pairs whose ciphertext differences don't satisfy this form: $(P(h|0|0|0|0|0|0|0), b|0|0|0|0|0|0|0)$ with b and h being non-zero bytes and $b^{(1)} = 0$. The number of remaining pairs after this test is about 2^{63} . Guess $k_{e,1}, k_a$ and $k_{b,1}$ in turn and operate the similar steps. If only one subkey is left, we retrieve the secret key by the key schedule. Otherwise, try another differential of A . As a matter of fact, if Δ is an impossible differential, the expected number of remaining pairs is about $2^{80} \times (1 - 2^{-8})^{15} \approx 2^{-104.7}$. For other two cases, we execute the similar attack procedure.

We find that the dominant time complexity of all steps in three cases is the data collection. Therefore, the total data, time and memory complexities are 2^{122} chosen plaintexts, 2^{122} 11-round encryptions and 2^{102} bytes, respectively.

5 Conclusion

In this paper, we have presented new insight on impossible differential cryptanalysis of reduced-round Camellia with the FL/FL^{-1} and whitening layers. First, we propose impossible differential attacks on reduced-round Camellia for 75% of the keys, which are then extended to attacks that work for the whole key space. As a matter of fact, we attack 10-round Camellia-128, 11-round Camellia-192 and 12-round Camellia-256 which start from the first round and include the whitening layers. Meanwhile, we also attack 12-round Camellia-192 and 14-round Camellia-256 with two FL/FL^{-1} layers. Second, we construct a set of differentials including at least one 8-round impossible differential of Camellia with two layers FL/FL^{-1} . This impossible differential has the same length as the best known impossible differential of Camellia without the FL/FL^{-1} layer.

Therefore, our result shows that the keyed functions cannot thwart impossible differential attack effectively. On the basis of this set of differentials, we propose a new strategy to derive an effective attack on 11-round Camellia-128, 12-round Camellia-192 and 13-round Camellia-256, which do not start the first round but include the whitening and FL/FL^{-1} layers.

Acknowledgements. The authors are grateful to all anonymous reviewers for valuable suggestions and comments. The authors Ya Liu, Dawu Gu, Zhiqiang Liu and Wei Li are supported by the National Natural Science Foundation of China (No. 61073150 and No. 61003278), the Opening Project of Shanghai Key Laboratory of Integrate Administration Technologies for Information Security, the open research fund of State Key Laboratory of Information Security and the Fundamental Research Funds for the Central Universities. The authors Leibo Li, Xiaoyun Wang and Jiazhe Chen are supported by the National Natural Science Foundation of China (Grant No. 61133013 and No. 60931160442), and Tsinghua University Initiative Scientific Research Program (2009THZ01002).

References

1. Aoki, K., Ichikawa, T., Kanda, M., Matsui, M., Moriai, S., Nakajima, J., Tokita, T.: *Camellia*: A 128-Bit Block Cipher Suitable for Multiple Platforms - Design and Analysis. In: Stinson, D.R., Tavares, S. (eds.) SAC 2000. LNCS, vol. 2012, pp. 39–56. Springer, Heidelberg (2001)
2. 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)
3. Chen, J., Jia, K., Yu, H., Wang, X.: New Impossible Differential Attacks of Reduced-Round Camellia-192 and Camellia-256. In: Parampalli, U., Hawkes, P. (eds.) ACISP 2011. LNCS, vol. 6812, pp. 16–33. Springer, Heidelberg (2011)
4. CRYPTREC-Cryptography Research and Evaluation Committees: report. Archive (2002), <http://www.ipa.go.jp/security/enc/CRYPTREC/index-e.html>
5. Hatano, Y., Sekine, H., Kaneko, T.: Higher order differential attack of *Camellia* (II). In: Nyberg, K., Heys, H.M. (eds.) SAC 2002. LNCS, vol. 2595, pp. 129–146. Springer, Heidelberg (2003)
6. International Standardization of Organization (ISO): International standard - ISO/IEC 18033-3. Tech. rep., Information technology - Security techniques - Encryption algorithm - Part 3: Block Ciphers (July 2005)
7. Knudsen, L.R.: DEAL - a 128-bit block cipher. Tech. rep., Department of Informatics, University of Bergen, Norway. technical report (1998)
8. Kühn, U.: Improved Cryptanalysis of MISTY1. In: Daemen, J., Rijmen, V. (eds.) FSE 2002. LNCS, vol. 2365, pp. 61–75. Springer, Heidelberg (2002)
9. Lee, S., Hong, S., Lee, S., Lim, J., Yoon, S.: Truncated Differential Cryptanalysis of Camellia. In: Kim, K.-c. (ed.) ICISC 2001. LNCS, vol. 2288, pp. 32–38. Springer, Heidelberg (2002)
10. Duo, L., Chao, L., Feng, K.: New Observation on Camellia. In: Preneel, B., Tavares, S. (eds.) SAC 2005. LNCS, vol. 3897, pp. 51–64. Springer, Heidelberg (2006)

11. Duo, L., Li, C., Feng, K.: Square Like Attack on Camellia. In: Qing, S., Imai, H., Wang, G. (eds.) ICICS 2007. LNCS, vol. 4861, pp. 269–283. Springer, Heidelberg (2007)
12. Li, L., Chen, J., Jia, K.: New Impossible Differential Cryptanalysis of Reduced-Round Camellia. In: Lin, D., Tsudik, G., Wang, X. (eds.) CANS 2011. LNCS, vol. 7092, pp. 26–39. Springer, Heidelberg (2011)
13. Li, L., Chen, J., Wang, X.: Security of Reduced-Round Camellia against Impossible Differential Attack. IACR Cryptology ePrint Archive 2011, 524 (2011)
14. Liu, Y., Gu, D., Liu, Z., Li, W., Man, Y.: Improved Results on Impossible Differential Cryptanalysis of Reduced-Round Camellia-192/256. *Journal of Systems and Software* (accepted)
15. Lu, J., Dunkelman, O., Keller, N., Kim, J.-S.: New Impossible Differential Attacks on AES. In: Chowdhury, D.R., Rijmen, V., Das, A. (eds.) INDOCRYPT 2008. LNCS, vol. 5365, pp. 279–293. Springer, Heidelberg (2008)
16. Lu, J., Kim, J.-S., Keller, N., Dunkelman, O.: Improving the Efficiency of Impossible Differential Cryptanalysis of Reduced Camellia and MISTY1. In: Malkin, T. (ed.) CT-RSA 2008. LNCS, vol. 4964, pp. 370–386. Springer, Heidelberg (2008)
17. Lu, J., Wei, Y., Kim, J., Fouque, P.-A.: Cryptanalysis of Reduced Versions of the Camellia Block Cipher. In: Preproceeding of SAC (2011)
18. Lu, J., Wei, Y., Kim, J., Pasalic, E.: The Higher-Order Meet-in-the-Middle Attack and Its Application to the Camellia Block Cipher. In: Presented in Part at the First Asian Workshop on Symmetric Key Cryptography (ASK 2011) (August 2011), <https://sites.google.com/site/jiqiang/>
19. Mala, H., Shakiba, M., Dakhilalian, M., Bagherikaram, G.: New Results on Impossible Differential Cryptanalysis of Reduced-Round Camellia-128. In: Jacobson Jr., M.J., Rijmen, V., Safavi-Naini, R. (eds.) SAC 2009. LNCS, vol. 5867, pp. 281–294. Springer, Heidelberg (2009)
20. NESSIE: New European Schemes for Signatures, Integrity, and Encryption, final report of european project IST-1999-12324. Archive (1999), <http://www.cosic.esat.kuleuven.be/nessie/Bookv015.pdf>
21. Shirai, T.: Differential, Linear, Boomerange and Rectangle Cryptanalysis of Reduced-Round Camellia. In: Proceedings of 3rd NESSIE Workshop, Munich, Germany, November 6-7 (2002)
22. Sugita, M., Kobara, K., Imai, H.: Security of Reduced Version of the Block Cipher Camellia against Truncated and Impossible Differential Cryptanalysis. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 193–207. Springer, Heidelberg (2001)
23. Wu, W., Feng, D., Chen, H.: Collision Attack and Pseudorandomness of Reduced-Round Camellia. In: Handschuh, H., Hasan, M.A. (eds.) SAC 2004. LNCS, vol. 3357, pp. 252–266. Springer, Heidelberg (2004)
24. Wu, W., Zhang, W., Feng, D.: Impossible Differential Cryptanalysis of Reduced-Round ARIA and Camellia. *J. Comput. Sci. Technol.* 22(3), 449–456 (2007)

Improved Rebound Attack on the Finalist Grøstl

Jérémy Jean^{1,*,**}, María Naya-Plasencia^{2,*}, and Thomas Peyrin^{3,***}

¹ École Normale Supérieure, France

² University of Versailles, France

³ Nanyang Technological University, Singapore

Abstract. Grøstl is one of the five finalist hash functions of the SHA-3 competition. For entering this final phase, the designers have tweaked the submitted versions. This tweak renders inapplicable the best known distinguishers on the compression function presented by Peyrin [18] that exploited the internal permutation properties. Since the beginning of the final round, very few analysis have been published on Grøstl. Currently, the best known rebound-based results on the permutation and the compression function for the 256-bit version work up to 8 rounds, and up to 7 rounds for the 512-bit version. In this paper, we present new rebound distinguishers that work on a higher number of rounds for the permutations of both 256 and 512-bit versions of this finalist, that is 9 and 10 respectively. Our distinguishers make use of an algorithm that we propose for solving three fully active states in the middle of the differential characteristic, while the Super-Sbox technique only handles two.

Keywords: Hash Function, Cryptanalysis, SHA-3, Grøstl, Rebound Attack.

1 Introduction

Hash functions are one of the main families in symmetric cryptography. They are functions that, given an input of variable length, produce an output of a fixed size. They have many important applications, like integrity check of executables, authentication, digital signatures.

Since 2005, several new attacks on hash functions have appeared. In particular, the hash standards MD5 and SHA-1 were cryptanalysed by Wang et al. [21,22]. Due to the resemblance of the standard SHA-2 with SHA-1, the confidence in the former has also been somewhat undermined. This is why the American National Institute of Standards and Technology (NIST) decided to launch in 2008 a competition for finding a new hash standard, SHA-3. This competition received 64 hash function submissions and accepted 51 to enter the first round. Now, three

* Supported by the French Agence Nationale de la Recherche through the SAPHIR2 project under Contract ANR-08-VERS-014.

** Supported by the French *Délégation Générale pour l'Armement* (DGA).

*** Supported by the Lee Kuan Yew Postdoctoral Fellowship 2011 and the Singapore National Research Foundation Fellowship 2012.

years and two rounds later, only 5 hash functions remain in the final phase of the competition.

Amongst these finalists, there is only one AES-based function, though many were proposed. This hash function is `Grøstl` [2], and is at the origin of the introduction of a new cryptanalysis technique that has been widely deployed, improved and applied to a large number of SHA-3 candidates, hash functions and other types of constructions. This new technique, called rebound attack, was introduced by Mendel et al. [11] and has become one of the most important tools used to analyze the security margin of many SHA-3 candidates as well as their building blocks. As for `Grøstl` itself, it has been applied and improved in several occasions [3, 12, 13, 15, 18]. `Grøstl` is undoubtedly one of the SHA-3 candidates that have received the largest amount of cryptanalysis. When entering the final round, a tweak of the function was proposed, which prevents the application of the attacks from [18]; we denote `Grøstl-0` the original submission of the algorithm and `Grøstl` its tweaked version. Apart from the rebound results, the other main analysis communicated on `Grøstl` was at the presentation of [1] where a higher order property on 10 rounds of `Grøstl-256` permutation with a complexity of 2^{509} was shown. In Table 1, we report a summary of the best known results on both 256 and 512-bit tweaked versions of `Grøstl`, including the ones that we will present in the following.

In this paper, we propose new results regarding both versions of the finalist `Grøstl`. First, on `Grøstl-256`, we provide the best known rebound distinguishers on 9 rounds of the permutation. From these results, we show how to make some nontrivial observations on the the compression function, providing the best known analysis on the compression function exploiting the properties of the internal permutations. For `Grøstl-512`, we considerably increase the number of analyzed rounds, from 7 to 10, providing the best analysis known on the permutation. Both results are obtained using rebound-like attack techniques and an algorithm that we introduce that allows to solve three fully active rounds in the middle of the differential characteristic with a much lower cost than a generic algorithm. Additionally, we provide in Appendix A the direct application of our new techniques to the AES-based hash function `PHOTON`.

These results do not threaten the security of `Grøstl`, but we believe they will have an important role in better understanding `Grøstl`, and AES-based functions in general. In particular, we believe that our work will help determining the bounds and limits of rebound-like attacks in these types of constructions.

2 Generalities

2.1 Description of `Grøstl`

The hash function `Grøstl-0` has been submitted to the SHA-3 competition under two different versions: `Grøstl-0-256`, which outputs a 256-bit digest and

Table 1. Best known analysis on the finalist **Grøstl**. By best analysis, we mean the ones on the highest number of rounds

Target	Subtarget	Rounds	Time	Memory	Ideal	Reference
Grøstl-256	Permutation	8 (dist.)	2^{112}	2^{64}	2^{384}	[3]
		8 (dist.)	2^{48}	2^8	2^{96}	[19]
		9 (dist.)	2^{368}	2^{64}	2^{384}	Section 3
		10 (zero-sum)	2^{509}	—	2^{512}	[1]
Grøstl-512	Permutation	8 (dist.)	2^{280}	2^{64}	2^{448}	Section 4
		9 (dist.)	2^{328}	2^{64}	2^{384}	Section 4
		10 (dist.)	2^{392}	2^{64}	2^{448}	Section 4

Grøstl-0-512 with a 512-bit fingerprint. For the final round of the competition, the candidate have been tweaked to **Grøstl**, with corresponding versions **Grøstl-256** and **Grøstl-512**.

The **Grøstl** hash function handles arbitrary long messages by diving them into blocks after some padding and uses them to update iteratively an internal state (initialized to a predefined IV) with a compression function. This function is itself built upon two different permutations, namely P and Q . Each of those two permutations updates a large internal state using the well-understood wide-trail strategy of the AES. As an AES-like Substitution-Permutation Network, **Grøstl** enjoys a strong diffusion in each of the two permutations and by its wide-pipe design, the size of the internal states is ensured to be at least twice as large as the final digest.

The compression function f_{256} of **Grøstl-256** uses two permutations P_{256} and Q_{256} , which are similar to the two permutations P_{512} and Q_{512} used in the compression function f_{512} of **Grøstl-512**. More precisely, for a chaining value h and a message block m , the compression functions (Figure 1) produce the output (\oplus denotes the XOR operation):

$$\begin{aligned}
 f_{256}(h, m) &= P_{256}(h \oplus m) \oplus Q_{256}(m) \oplus h, \\
 \text{or: } f_{512}(h, m) &= P_{512}(h \oplus m) \oplus Q_{512}(m) \oplus h.
 \end{aligned}$$

The internal states are viewed as byte matrices of size 8×8 for the 256-bit version and 8×16 for the 512-bit one. The permutations strictly follow the design of the AES and are constructed as N_r iterations of the composition of four basic transformations:

$$R \stackrel{\text{def}}{=} \text{MixBytes} \circ \text{ShiftBytes} \circ \text{SubBytes} \circ \text{AddRoundConstant}.$$

All the linear operations are performed in the same finite field $GF(2^8)$ as in the AES, defined via the irreducible polynomial $x^8 + x^4 + x^3 + x + 1$ over $GF(2)$.




Fig. 1. The compression function of Grøst1 hash function using the two permutations P and Q

The **AddRoundConstant** (AC) operation adds a predefined round-dependent constant, which significantly differs between P and Q to prevent the internal differential attack [18] taking advantage of the similarities in P and Q . The **SubBytes** (SB) layer is the non-linear layer of the round function R and applies the same SBox as in the AES to all the bytes of the internal state. The **ShiftBytes** (Sh) transformation shifts bytes in row i by $\tau_P[i]$ positions to the left for permutation P and $\tau_Q[i]$ positions for permutation Q . We note that τ also differs from P to Q to emphasize the asymmetry between the two permutations. Finally, the **MixBytes** (Mb) operation applies a maximum-distance separable (MDS) circular constant matrix M independently to all the columns of the state. In Grøst1-256, $N_r = 10$, $\tau_P = [0, 1, 2, 3, 4, 5, 6, 7]$ and $\tau_Q = [1, 3, 5, 7, 0, 2, 4, 6]$, whereas for Grøst1-512, $N_r = 14$ and $\tau_P = [0, 1, 2, 3, 4, 5, 6, 11]$ and $\tau_Q = [1, 3, 5, 11, 0, 2, 4, 6]$.

Once all the message blocks of the padded input message have been processed by the compression function, a final output transformation is applied to the last chaining value h to produce the final n -bit hash value $h' = \text{trunc}_n(P(h) \oplus h)$, where trunc_n only keeps the last n bits.

2.2 Distinguishers

In this article, we will describe algorithms that find input pairs (X, X') for the permutation P (or the permutation Q), such that the input difference $\Delta_{IN} = X \oplus X'$ belongs to a subset of size IN and the output difference $\Delta_{OUT} = P(X) \oplus P(X')$ belongs to a subset of size OUT . The best known generic algorithm (this problem is different than the one studied in [8] where linear subspaces are considered) in order to solve this problem, known as limited-birthday problem, has been given in [3] and later a very close lower bound has been proven in [16]. For a randomly chosen n -bit permutation π , the generic algorithm can find such a pair with complexity $\max\{\min\{\sqrt{2^n/IN}, \sqrt{2^n/OUT}\}, 2^n/(IN \cdot OUT)\}$. If one is able to describe an algorithm requiring less computation power, then we consider that a distinguisher exists on the permutation π .

In the case of Grøst1, it is also interesting to look at not only the internal permutations P and Q , but also the compression function f itself. For that matter, we will generate compression function input values (h, m) such that

$\Delta_{IN} = m \oplus h$ belongs to a subset of size IN , and such that $\Delta_{IN} \oplus \Delta_{OUT} = f(h, m) \oplus f(m, h) \oplus h \oplus m$ belongs to a subset of size OUT . Then, one can remark that:

$$\begin{aligned} f(h, m) \oplus f(m, h) &= P_{256}(h \oplus m) \oplus Q_{256}(m) \oplus P_{256}(m \oplus h) \oplus Q_{256}(h) \oplus h \oplus m, \\ f(h, m) \oplus f(m, h) &= Q_{256}(m) \oplus Q_{256}(h) \oplus h \oplus m. \end{aligned}$$

Hence, it follows that:

$$f(h, m) \oplus f(m, h) \oplus h \oplus m = Q_{256}(m) \oplus Q_{256}(h).$$

Since the permutation Q is supposed to have no structural flaw, the best known generic algorithm requires $\max\{\min\{\sqrt{2^n/IN}, \sqrt{2^n/OUT}\}, 2^n/(IN \cdot OUT)\}$ operations (the situation is exactly the same as the permutation distinguisher with permutation Q) to find a pair (h, m) of inputs such that $h \oplus m \in IN$ and $f(h, m) \oplus f(m, h) \oplus h \oplus m \in OUT$. Note that both IN and OUT are specific to our attacks.

We emphasize that even if trivial distinguishers are already known for the **Grøst1** compression function (for example fixed-points), no distinguisher is known for the internal permutations. Moreover, our observations on the compression function use the differential properties of the internal permutations.

3 Distinguishers for Reduced Grøst1-256 Permutations

In this section, we describe a distinguisher for the permutation P_{256} of the **Grøst1-256** compression function reduced to 9 rounds. We emphasize that in the latest version of the **Grøst1** submission [20], the permutation Q_{256} has different coefficients in the **ShiftRows** transformation, but the technique we describe in the following applies to Q_{256} as well.

3.1 The Truncated Differential Characteristic

In the following, we will consider truncated differential characteristics, originally introduced by Knudsen [7] for block cipher analysis. With this technique, already proven to be efficient for AES-based hash functions cryptanalysis [5, 6, 10, 17], the attacker only checks if there is a difference in a byte (active byte, denoted by a black square in the Figures) or not (inactive byte, denoted by an empty square in the Figures) without caring about the actual value of the difference.

The truncated differential characteristic we use has the sequence of active bytes

$$8 \xrightarrow{R_1} 1 \xrightarrow{R_2} 8 \xrightarrow{R_3} 64 \xrightarrow{R_4} 64 \xrightarrow{R_5} 64 \xrightarrow{R_6} 8 \xrightarrow{R_7} 1 \xrightarrow{R_8} 8 \xrightarrow{R_9} 64,$$

where the size in the input and output differences subsets are both $IN = OUT = 2^{8 \times 8} = 2^{64}$, since there are eight active bytes in each extreme state

of the truncated characteristic. The actual truncated characteristic is reported in Appendix B.

Note that we have three fully active internal states in the middle of the differential characteristic, thus impossible to handle with the classical rebound or **SuperSBox** techniques.

3.2 Finding a Conforming Pair

The method to find a pair of inputs conforming to this truncated differential characteristic is similar to the rebound technique: we first find many solutions for the middle rounds (round 3 to round 6) and then we filter them out during the outwards probabilistic transitions through the **MixBytes** layers (round 2 and round 7). We denote $x \rightarrow y$ a non-null truncated differential transition mapping x active bytes to y active bytes in a column through a **MixBytes** (or **MixBytes**⁻¹) layer, and the MDS property ensures $x + y \geq 9$. Its differential probability is determined by the number $(8 - y)$ of inactive bytes on the output: $2^{-8(8-y)}$ if the MDS property is verified, 0 otherwise.

Therefore, since in our case we have two transitions $8 \rightarrow 1$ (see Figure 2), the outbound phase has a success probability of $(2^{-8 \times 7})^2 = 2^{-112}$ and is straightforward to handle once we found enough solutions for the inbound phase.

In order to find solutions for the middle rounds (see Figure 2), we propose an algorithm inspired by the ones in [14, 15]: As in [3, 8], instead of dealing with the classical 8-bit **SubBytes** SBoxes, one can consider 64-bit SBoxes (named **SuperSBoxes**) each composed of two **AES** SBox layers surrounding one **MixBytes** and one **AddRoundConstant** function¹. Indeed, the **ShiftBytes** can be taken out from the **SuperSBoxes** since it commutes with **SubBytes**.

We start by choosing the input difference δ_{IN} after the first **SubBytes** layer in state **S1** and the output difference δ_{OUT} after the last **MixBytes** layer in state **S12** in a way that the truncated characteristic holds in **S0** and **S12**. Note that since we have 8 active bytes in **S1** and **S12**, there are as many as $2^{2 \times 64} = 2^{128}$ different ways of choosing $(\delta_{IN}, \delta_{OUT})$. We continue by constructing the 8 forward **SuperSBox** independently by considering the 2^{64} possible input values for each of them in state **S3**: differences in **S1** can be directly propagated to **S3** since **MixBytes** is linear. This generates 8 independent lists, each of size 2^{64} and composed by paired values. Doing the same for the 8 backwards **SuperSBoxes** from state **S12**, we again get 8 independent lists of 2^{64} elements each, and we end up in state **S8** where the 8 forward and the 8 backward lists overlap. In the sequel, we denote L_i the i th forward **SuperSBox** list and L'_i the i th backward one, for $1 \leq i \leq 8$.

In terms of freedom degrees in state **S8**, we want to merge 16 lists of 2^{64} elements each for a merging condition on $2 \times 512 = 1024$ bits (512 for values and 512 for differences): we then expect $2^{16 \times 64} 2^{-1024} = 1$ solution as a result of the

¹ These **SuperSBoxes** are 64-bit large in the case of Grøst1, but only $4 \times 8 = 32$ bits for the **AES**.




Fig. 2. Inbound phase for the 9-round distinguisher attack on the *Grøst1* permutation P_{256} . The four rounds represented are the rounds 3 to 6 from the whole truncated differential characteristic. A gray byte indicates an active byte; hatched and coloured bytes emphasize one **SuperSBox**: there are seven similar others.

merging process. We detail a method in order to find this solution in time 2^{256} and memory 2^{64} (see Figure 3).

Step 1. We start by considering every possible combination of elements in each of the four lists L'_1, L'_2, L'_3 and L'_4 . There are 2^{256} possibilities.

Step 2. This fully constraints 2×4 bytes in each of the 8 lists $L_i, 1 \leq i \leq 8$ (i.e. the first 4 columns of the internal state). For each of them, we then expect $2^{64} 2^{-8 \times 8} = 1$ element to match the randomized bytes. These elements can be found with one operation by sorting the lists L_i beforehand. At this point, note that the second half of the state **S8** has been fully determined by the choice in L_1, \dots, L_8 .

Step 3. We now need to ensure that the 4 last lists L'_5, L'_6, L'_7 and L'_8 contain the elements imposed: those lists being of size 2^{64} each, this happens with probability $2^{64} 2^{-8 \times (2 \times 8)} = 2^{-64}$ independently on each list. Again, these elements can be found with one operation by sorting the lists L'_i beforehand.

All in all, trying all the 2^{256} elements in (L'_1, L'_2, L'_3, L'_4) , we expect to find $2^{256} 2^{-64 \times 4} = 1$ solution that will verify the 1024 bits of condition and we can find this solution with only a few operations.




Fig. 3. Steps to merge the 16 lists. Grey cells denote bytes fully constrained by a choice of elements in L'_1, \dots, L'_4 during the first step.

Hence, from random differences $(\delta_{IN}, \delta_{OUT})$, we find a pair of internal states of the permutation that conforms to the middle rounds in time 2^{256} and memory 2^{64} . To pass the probabilistic transitions of the outbound phase, we need to repeat the merging 2^{112} times by picking another couple of differences $(\delta_{IN}, \delta_{OUT})$. In total, we find a pair of inputs to the permutation that conforms to the truncated differential characteristic in time complexity 2^{368} and memory complexity 2^{64} .

3.3 Comparison with Ideal Case

In the ideal case, obtaining a pair whose input and output differences lie in a subset of size $IN = OUT = 2^{64}$ for a 512-bit permutation requires 2^{384} computations: we can directly conclude that this leads to a distinguishing attack on the 9-round reduced version of the **Grøst1-256** permutation with 2^{368} computations and 2^{64} memory. Similarly, as explained in Section 2.2, this result also induces a nontrivial observation on the 9-round reduced version of the **Grøst1-256** compression function with identical complexity.

Finally, one can also derive slightly cheaper distinguishers by aiming less rounds: instead of using the 9-round truncated characteristic from Appendix B, it is possible to remove either round 2 or 8 and spare one $8 \rightarrow 1$ truncated differential transition. Overall, the generic complexity remains the same and this gives a distinguishing attack on the 8-round reduced version of the **Grøst1-256** permutation with 2^{312} computations and 2^{64} memory. Unfortunately, this is worse than previously known results.

4 Distinguishers for Reduced **Grøst1-512** Permutations

The 512-bit version of the **Grøst1** hash function uses a non-square 8×16 matrix as 1024-bit internal state, which therefore presents a lack of optimal diffusion: a single difference generates a fully active state after three rounds where a square-state would need only two. This enables us to add an extra round to the generalization of the regular 9-round characteristic of AES-like permutation (Section 3) to reach 10 rounds.

4.1 The Truncated Differential Characteristic

To distinguish its permutation P_{512} ² reduced to 10 rounds, we use the truncated differential characteristic with the sequence of active bytes

$$64 \xrightarrow{R_1} 8 \xrightarrow{R_2} 1 \xrightarrow{R_3} 8 \xrightarrow{R_4} 64 \xrightarrow{R_5} 128 \xrightarrow{R_6} 64 \xrightarrow{R_7} 8 \xrightarrow{R_8} 1 \xrightarrow{R_9} 8 \xrightarrow{R_{10}} 64.$$

where the size of the input differences subset is $IN = 2^{512}$ and the size of the output differences subset is $OUT = 2^{64}$.

The actual truncated characteristic is appended in Appendix C. Again, we split the characteristic into two parts: the inbound phase involving a merging of lists in the four middle rounds (round 4 to round 7), and an outbound phase that behaves as a probabilistic filter ensuring both $8 \rightarrow 1$ transitions in the outward directions. Again, passing those two transitions with random values occurs with probability 2^{-112} .

4.2 Finding a Conforming Pair

In the following, we present an algorithm to solve the middle rounds in time 2^{280} and memory 2^{64} . In total, we will need to repeat this process 2^{112} times to get a pair of internal states that conforms to the whole truncated differential characteristic, which would then cost $2^{280+112} = 2^{392}$ in time and 2^{64} in memory. The strategy of this algorithm (see Figure 4) is similar to the ones presented in [14, 15] and the one from the previous section: we start by fixing the difference to a random value δ_{IN} in S1 and δ_{OUT} in S12 and linearly deduce the difference δ'_{IN} in S3 and δ'_{OUT} in S10. Then, we construct the 32 lists corresponding to the 32 **SuperSBoxes**: the 16 forward **SuperSBoxes** have an input difference fixed to δ'_{IN} and cover states S3 to S8, whereas the 16 backward **SuperSBoxes** spread over states S10 to S6 with an output difference fixed to δ'_{OUT} . In the sequel, we denote L_i the 16 forward **SuperSBoxes** and L'_i the backward ones, $1 \leq i \leq 16$.

The 32 lists overlap in S8, where we merge them on 2048 bits³ to find $2^{64 \times 32} 2^{-2048} = 1$ solution, since each list is of size 2^{64} . The naive way to find the solution would cost 2^{1024} in time by considering each element of the Cartesian product of the 16 lists L_i to check whether it satisfies the output 1024 bit difference condition. We describe now the algorithm that achieves the same goal in time 2^{280} .

First, we observe that due to the geometry of the non-square state, any list L_i intersects with only half of the L'_j . For instance, the first list L_1 associated to the first column of state S7 intersects with lists $L'_1, L'_6, L'_{11}, L'_{12}, L'_{13}, L'_{14}, L'_{15}$ and L'_{16} . We represent this property with a 16×16 array on Figure 5: the 16 columns correspond to the 16 lists L'_i and the lines to the L_i , $1 \leq i \leq 16$. The cell (i, j) is white if and only if L_i has a non-null intersection with the list L'_j , otherwise it is gray.

² It would work exactly the same way for the other permutation Q_{512} .

³ The 2048 bits come from 1024 bits of values and 1024 bits of differences.

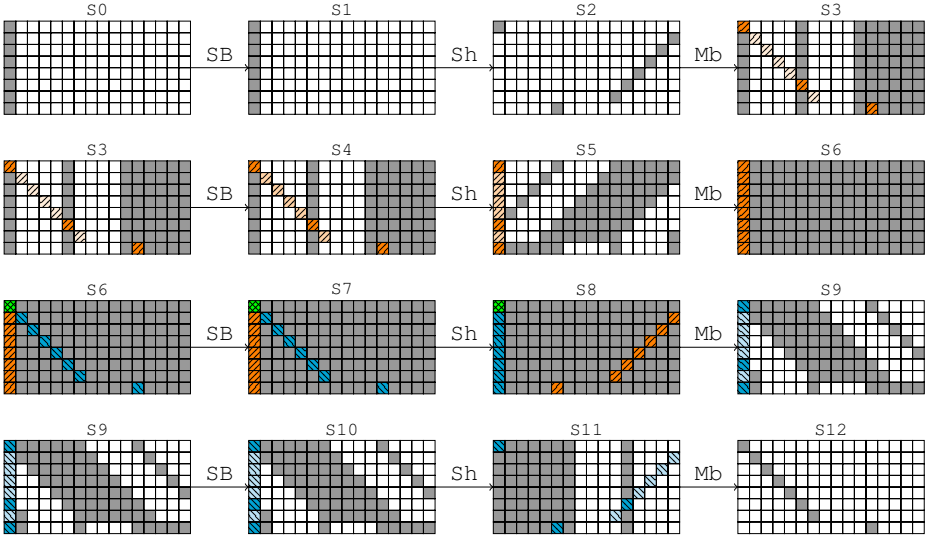


Fig. 4. Inbound phase for the 10-round distinguisher attack on the Grøstl-512 permutation P_{512} . The four rounds represented are the rounds 4 to 7 from the whole truncated differential characteristic C. A gray byte indicates an active byte; hatched and coloured bytes emphasize the **SuperSBoxes**.

Then, we note that the **MixBytes** transition between the states S8 and S9 constrains the differences in the lists L'_i : in the first column of S9 for example, only three bytes are active, so that the same column in S8 can only have $2^{3 \times 8}$ different differences, which means that knowing three out of the eight differences in an element of L'_1 is enough to deduce the other five. For a column-vector of differences lying in a n -dimensional subspace, we can divide the 2^{64} elements of the associated lists in 2^{8n} disjoint sets of 2^{64-8n} values each. So, whenever we know the n independent differences, the only freedom that remains lie in the values. The bottom line of Figure 5 reports the subspace dimensions for each L'_i .

Using a guess-and-determine approach, we derive a way to use the previous facts to find the solution to the merge problem in time 2^{280} . As stated before, we expect only one solution; that is, we want to find a single element in each of the 32 lists. We start by guessing the values and the differences of the elements associated to the lists L'_2 , L'_3 , L'_4 and L'_5 . For this, we will try all the possible combinations of their elements, there are $2^{4 \times 64} = 2^{256}$ in total. For each one of the 2^{256} tries, all the checked cells \checkmark now have known value and difference. From here, 8 bytes are known in each of the four lists L_5 , L_6 , L_7 and L_8 : this imposes a 64-bit constraint on those lists, which filter out a single element in each. Thereby, we determined the value and difference in the other 16 bytes marked by \checkmark in Figure 5. In lists L'_1 and L'_{16} , we have reached the maximum number of independent differences (three and two, respectively), so we can determine the differences for the other bytes of those columns: we mark them by \bullet . In L_4 , the 8 constraints (three \checkmark and two \bullet) filter out one element; then, we deduce the




Fig. 5. A ✓ means we know both value and difference for that byte, a ● means that we only determined the difference for that byte and white bytes are not constrained yet

correct element in L_4 and mark it by ✓. We can now determine the differences in L'_{15} since the corresponding subspace has a dimension equals to two.

At this point, no more byte can be determined based on the information propagated so far. We continue by guessing the elements remaining in L'_6 . Since there are already six byte-constraints on that list (three ✓), only 2^{16} elements conform to the conditions. The time complexity until now is thus $2^{256+16} = 2^{272}$.

Guessing the list L'_6 implies a 64-bit constraint of the list L_9 so that we get a single element out of it and determine four yet-unknown other bytes. This enables to learn the independent differences in L'_{14} and therefore, we filter an element from L_3 (two ✓ and four ●). At this stage, the list L'_1 is already fully constrained on its differences, so that we are left with a set of $2^{64-3 \times 8} = 2^{40}$ values constrained on five bytes (five ✓). Hence, we are able to determine all the unset values in L'_1 (Figure 6a).

Again, the lack of constraints prevent us to determine more bytes. We continue by guessing the 2^8 elements left in L_1 (two ✓ and three ●), which makes the time complexity increase to 2^{280} . The list L_1 being totally known, we derive the vector of differences in L'_{13} , which adds an extra byte-constraint on L_2 where only one element was left, and so fully determines it. From here, L'_7 becomes fully determined as well (four ✓) and so is L_{16} . In the latter, the differences being known, we were left with a set of $2^{64-2 \times 8} = 2^{48}$ values, which are now constrained on six bytes (six ✓).

We describe in Figure 6b the knowledge propagated so far, with time complexity 2^{280} and probability 1. We observe that L_{10} is overdetermined (four ✓ and one ●) by one byte. This means that we get the correct value with probability 2^{-8} , whereas L_{11} is filtered with probability 1. Similarly, the element of L'_8 happens to be correctly defined with probability 2^{-16} ; as for L'_9 and L'_{15} , with

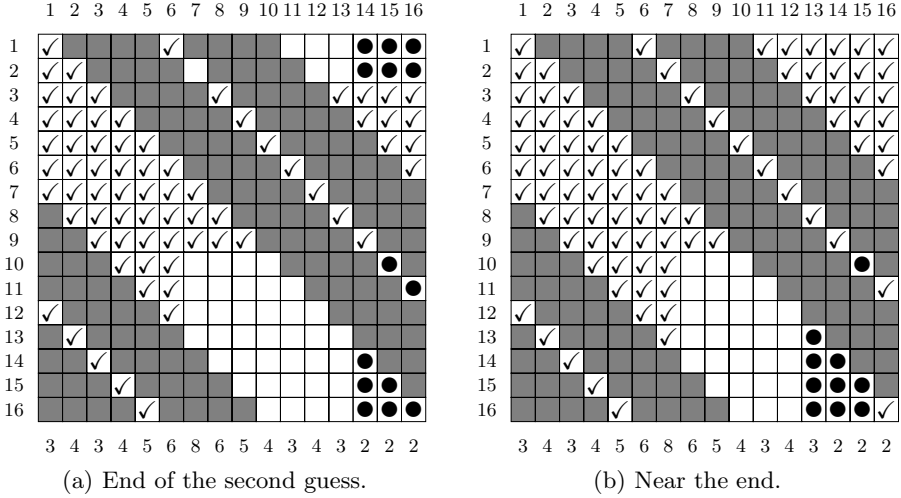


Fig. 6. A ✓ means we know both value and difference for that byte, a ● means that we only determined the difference for that byte and white bytes are not constrained yet

probability 1. We continue in L'_{11} by learning the full vector of differences, which constraints L_{12} on 11 bytes (five ✓ and one ●) so that we get a valid element with probability 2^{-24} . Finishing the guess and determine technique is done by filtering L'_{10} and L_{12} with probability 1, L_{16} with probability 2^{-40} and L_{13} , L_{14} and L_{15} with probability 2^{-64} each.

In total, for each guess, we successfully merge the 32 lists with probability

$$2^{-8-16-24-40-64-64-64} = 2^{-280},$$

but the whole procedure is repeated $2^{64 \times 4 + 16 + 8} = 2^{280}$ times, so we expect to find the one existing solution. All in all, we described a way to do the merge with time complexity 2^{280} and memory complexity 2^{64} . The final complexity to find a valid candidate for the whole characteristic is then 2^{392} computations and 2^{64} memory.

4.3 Comparison with Ideal Case

In the ideal case, obtaining a pair whose input difference lies in a subset of size $IN = 2^{512}$ and whose output difference lies in a subset of size $OUT = 2^{64}$ for a 1024-bit permutation requires 2^{448} computations. We can directly conclude that this leads to a distinguishing attack on the 10-round reduced version of the Grøst1-512 permutation with 2^{392} computations and 2^{64} memory. Similarly, as explained in Section 2.2, this results also induces a nontrivial observation on the 10-round reduced version of the Grøst1-512 compression function with identical complexity.

One can also derive slightly cheaper distinguishers by aiming less rounds while keeping the same generic complexity: instead of using the 10-round truncated characteristic from Appendix C, it is possible to remove either round 3 or 9 and spare one $8 \rightarrow 1$ truncated differential transition. Overall, this gives a distinguishing attack on the 9-round reduced version of the `Grøst1-512` permutation with 2^{336} computations and 2^{64} memory. By removing both rounds 3 and 9, we achieve 8 rounds with 2^{280} computations.

One can further gain another small factor for the 9-round case by using a $8 \rightarrow 2$ truncated differential transition instead of $8 \rightarrow 1$, for a final complexity of 2^{328} computations and 2^{64} memory. Indeed, the generic complexity drops to 2^{384} because we would now have $OUT = 2^{128}$.

5 Conclusion

In this paper, we have provided new and improved cryptanalysis results on the building blocks of both 256 and 512-bit versions of the finalist `Grøst1`. This is done by using a rebound-like approach as well as an algorithm that allows us to pass three fully active states in the middle of the differential characteristic with lower complexity than a general probabilistic approach. To the best of our knowledge, all previously known methods only manage to control two fully active states in the middle of the differential characteristic.

On `Grøst1-256`, we could provide the best known rebound distinguishers on 9 rounds of the permutation. For `Grøst1-512`, we have considerably increased the number of analyzed rounds, from 7 to 10, providing the best analysis known the permutation.

These results do not threaten the security of `Grøst1`, but we believe they will have an important role in better understanding AES-based functions in general. In particular, we believe that our work will help determining the bounds and limits of rebound-like attacks in these types of constructions. Future works could include the study of more AES-like functions in regards to this new cryptanalysis method.

References

1. Boura, C., Canteaut, A., De Cannière, C.: Higher-Order Differential Properties of KECCAK and *Luffa*. In: Joux, A. (ed.) FSE 2011. LNCS, vol. 6733, pp. 252–269. Springer, Heidelberg (2011)
2. Gauravaram, P., Knudsen, L.R., Matusiewicz, K., Mendel, F., Rechberger, C., Schl affer, M., Thomsen, S.S.: `Grøst1` -- a SHA-3 candidate
3. Gilbert, H., Peyrin, T.: Super-Sbox Cryptanalysis: Improved Attacks for AES-Like Permutations. In: Hong, S., Iwata, T. (eds.) FSE 2010. LNCS, vol. 6147, pp. 365–383. Springer, Heidelberg (2010)
4. Guo, J., Peyrin, T., Poschmann, A.: The `PHOTON` Family of Lightweight Hash Functions. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 222–239. Springer, Heidelberg (2011)

5. Jean, J., Fouque, P.-A.: Practical Near-Collisions and Collisions on Round-Reduced ECHO-256 Compression Function. In: Joux, A. (ed.) FSE 2011. LNCS, vol. 6733, pp. 107–127. Springer, Heidelberg (2011)
6. Jean, J., Naya-Plasencia, M., Schl affer, M.: Improved Analysis of ECHO-256. In: Miri, A., Vaudenay, S. (eds.) SAC 2011. LNCS, vol. 7118, pp. 19–36. Springer, Heidelberg (2012)
7. Knudsen, L.R.: Truncated and Higher Order Differentials. In: Preneel, B. (ed.) FSE 1994. LNCS, vol. 1008, pp. 196–211. Springer, Heidelberg (1995)
8. Lamberger, M., Mendel, F., Rechberger, C., Rijmen, V., Schl affer, M.: Rebound Distinguishers: Results on the Full Whirlpool Compression Function. In: [9], pp. 126–143
9. Matsui, M. (ed.): ASIACRYPT 2009. LNCS, vol. 5912. Springer, Heidelberg (2009)
10. Matusiewicz, K., Naya-Plasencia, M., Nikolic, I., Sasaki, Y., Schl affer, M.: Rebound Attack on the Full LANE Compression Function. In: [9], pp. 106–125
11. Mendel, F., Rechberger, C., Schl affer, M., Thomsen, S.S.: The Rebound Attack: Cryptanalysis of Reduced Whirlpool and Grøstl. In: Dunkelman, O. (ed.) FSE 2009. LNCS, vol. 5665, pp. 260–276. Springer, Heidelberg (2009)
12. Mendel, F., Peyrin, T., Rechberger, C., Schl affer, M.: Improved Cryptanalysis of the Reduced Grøstl Compression Function, ECHO Permutation and AES Block Cipher. In: Jacobson Jr., M.J., Rijmen, V., Safavi-Naini, R. (eds.) SAC 2009. LNCS, vol. 5867, pp. 16–35. Springer, Heidelberg (2009)
13. Mendel, F., Rechberger, C., Schl affer, M., Thomsen, S.S.: Rebound Attacks on the Reduced Grøstl Hash Function. In: Pieprzyk, J. (ed.) CT-RSA 2010. LNCS, vol. 5985, pp. 350–365. Springer, Heidelberg (2010)
14. Naya-Plasencia, M.: How to Improve Rebound Attacks. Cryptology ePrint Archive, Report 2010/607 (2010) (extended version), <http://eprint.iacr.org/>
15. Naya-Plasencia, M.: How to Improve Rebound Attacks. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 188–205. Springer, Heidelberg (2011)
16. Nikolić, I., Pieprzyk, J., Sokołowski, P., Steinfeld, R.: Known and Chosen Key Differential Distinguishers for Block Ciphers. In: Rhee, K.-H., Nyang, D. (eds.) ICISC 2010. LNCS, vol. 6829, pp. 29–48. Springer, Heidelberg (2011)
17. Peyrin, T.: Cryptanalysis of GRINDAHL. In: Kurosawa, K. (ed.) ASIACRYPT 2007. LNCS, vol. 4833, pp. 551–567. Springer, Heidelberg (2007)
18. Peyrin, T.: Improved Differential Attacks for ECHO and Grøstl. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 370–392. Springer, Heidelberg (2010)
19. Sasaki, Y., Li, Y., Wang, L., Sakiyama, K., Ohta, K.: Non-full-active Super-Sbox Analysis: Applications to ECHO and Grøstl. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 38–55. Springer, Heidelberg (2010)
20. Schl affer, M.: Updated Differential Analysis of Grøstl. Grøstl website (January 2011)
21. Wang, X., Yin, Y.L., Yu, H.: Finding Collisions in the Full SHA-1. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 17–36. Springer, Heidelberg (2005)
22. Wang, X., Yu, H.: How to Break MD5 and Other Hash Functions. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 19–35. Springer, Heidelberg (2005)

A Distinguishers for Other AES-Like Permutations

Using the same cryptanalysis technique, it is possible to study other AES-like schemes using permutations similar to the Grøst1 ones. For example, the recent lightweight hash function family PHOTON [4] is based on five different versions of AES-like permutations. We denote s the size of the cells ($s = 8$ for AES) and c the size of the square matrix representing the internal state ($c = 4$ for AES), the five versions (s, c) for PHOTON are then $(4, 5)$, $(4, 6)$, $(4, 7)$, $(4, 8)$ and $(8, 6)$ for increasing versions. All versions are defined to apply 12 rounds of an AES-like process, where the subkey additions are replaced by constant additions. Since the internal state is always square, by trivially adapting the method from Section 3 to the specific parameters of PHOTON, one can hope to obtain distinguishers for 9 rounds of the PHOTON internal permutations. However, we are able to do so only for the parameters $(4, 8)$ used in PHOTON-224/32/32 (see Table 2 with the comparison to previously known results). Indeed, the size c of the matrix plays an important role in the gap between the complexity of our algorithm and the generic one. The bigger is the matrix, the better will be the gap between the algorithm complexity and the generic one.

Table 2. Distinguishers on PHOTON internal permutation when applying the method from Section 3

Target	Subtarget	Rounds	Time	Memory	Ideal	Ref.
PHOTON-224/32/32	Permutation	8 (dist.)	2^8	2^4	2^{10}	[4]
		9 (dist.)	2^{184}	2^{32}	2^{192}	Section A

The same effect applies on AES in the known-key model, for which distinguishers on only 8 rounds are known as of today [3]. When attacking 9 rounds with the method from Section 3, the middle rounds will cost about 2^{64} operations per solution, while the two $4 \rightarrow 1$ truncated differential transitions during the outbound will be verified with probability $(2^{-24})^2 = 2^{-48}$. Overall, one solution for the whole characteristic is found with 2^{112} computation and 2^{32} memory, but the generic algorithm can find such a pair with only 2^{64} .

B 9-Round Grøst1-256 Permutation Truncated Characteristic




Fig. 7. The 9-round truncated differential characteristic used to distinguish the permutation P of Grøst1-256 from an ideal permutation

C 10-Round Grøstl-512 Permutation Truncated Characteristic




Fig. 8. The 10-round truncated differential characteristic used to distinguish the permutation P of Grøstl-512 from an ideal permutation

(Pseudo) Preimage Attack on Round-Reduced Grøstl Hash Function and Others

Shuang Wu¹, Dengguo Feng¹, Wenling Wu¹, Jian Guo², Le Dong¹,
and Jian Zou¹

¹ State Key Laboratory of Information Security, Institute of Software,
Chinese Academy of Sciences

² Institute for Infocomm Research, Singapore
wushuang@is.iscas.ac.cn

Abstract. The Grøstl hash function is one of the 5 final round candidates of the SHA-3 competition hosted by NIST. In this paper, we study the preimage resistance of the Grøstl hash function. We propose pseudo preimage attacks on Grøstl hash function for both 256-bit and 512-bit versions, *i.e.*, we need to choose the initial value in order to invert the hash function. Pseudo preimage attack on 5(out of 10)-round Grøstl-256 has a complexity of $(2^{244.85}, 2^{230.13})$ (in time and memory) and pseudo preimage attack on 8(out of 14)-round Grøstl-512 has a complexity of $(2^{507.32}, 2^{507.00})$. To the best of our knowledge, our attacks are the first (pseudo) preimage attacks on round-reduced Grøstl hash function, including its compression function and output transformation. These results are obtained by a variant of meet-in-the-middle preimage attack framework by Aoki and Sasaki. We also improve the time complexities of the preimage attacks against 5-round Whirlpool and 7-round AES hashes by Sasaki in FSE 2011.

Keywords: hash function, meet-in-the-middle, preimage attack, Grøstl, Whirlpool, AES.

1 Introduction

In FSE 2008, Gaëtan Leurent proposed the first preimage attack on the full MD4 hash function [12]. Based on this pioneering work, Aoki and Sasaki invented the technique of Meet-int-the-middle (MitM) preimage attack [2]. The basic idea of this technique is to divide the compression function into two concatenated sub-functions. The output values of two sub-functions can be independently calculated from the given input value in the forward direction and the backward direction. The steps of the forward and backward computation are called forward chunk and backward chunk. Then the MitM attack is applied to the output values of two sub-functions at the concatenating point of two chunks.

For hash functions based on block ciphers, the feedforward operations in the mode of operations like Davis-Meyer, Matyas-Meyer-Oseas and Miyaguchi-Preneel provide a chance for the applications of new technique called *splice-and-cut* [2]. The input and output of a compression function can be regarded as

concatenated through the feed-forward operation in these modes of operations. Then the compression function is in the form of a circle and any step can be selected as either the starting point or the matching point.

Improvements have been developed on both the starting point and the matching point. The *initial structure* technique [16] (also called *message stealing* [7]) and the *local collision*¹ [15] technique allows two sub-functions to share several steps without violating the independency in computing their own values, which provides more attackable rounds. The *partial matching* technique [2,16,7,1] takes advantage of the compression function’s diffusion properties at the matching point. Due to slow diffusion of the Feistel-like round function, part of the state value can remain independent of the other chunk while proceeding with more reversed rounds. The deterministic part of the state is used as the matching point. After finding a match of the partial values, the equality of the remaining part is calculated and checked. These techniques used in the MitM preimage attacks are illustrated in Fig. 1.




Fig. 1. Advanced techniques for MitM preimage attack

The MitM preimage attacks have been applied to full HAVAL-3/4 [15], MD4 [2,7], MD5 [16], Tiger [7], and round-reduced HAS-160 [8], RIPEMD [21], SHA-0/1 [3], SHA-2 [7,1]. The compression functions of these hash functions all use Feistel-like structures. In FSE 2011, Yu Sasaki proposed MitM preimage attack on AES hash mode for the first time [14]. He discussed how initial structure and partial matching can be used on AES-like structures and proposed direct applications to AES in different hash modes and round-reduced Whirlpool [4]. The development of the MitM attacks on hash functions has also inspired several attacks on block ciphers, such as KTANTAN [22] and XTEA [19].

Our Contributions. In this paper, we found a way to reduce the complexity of the MitM preimage attack on AES-like hash functions. By finding the optimal chunk separation with best balance between freedom degrees and the size of the matching point, the freedom degrees in the internal states are fully utilized.

Grøst1 [6] is one of the five finalists in the third round of SHA-3 [13] competition hosted by NIST. The Grøst1 hash function has been tweaked in the third

¹ The local collision technique was proposed by Joux et al. [5], which is originally used in the collision attacks. The similar idea can be used to construct the initial structure in the MitM preimage attack.

round. The original version is renamed to Grøstl-0 and the tweaked version is called Grøstl.

We found that Grøstl’s round-reduced output transformation can be inverted using the MitM techniques. Then we noticed that if we can control the initial value, preimage of the output transformation can be connected with a compression function. The Grøstl hash function uses wide-pipe chaining values, so we can actually match $2n$ -bit chaining value with a time complexity less than 2^n compression function calls. Since the initial value is chosen by us, this attack is a pseudo preimage attack.

The matching of double-sized states are based on a method of variant generalized birthday attack. The special property of Grøstl’s compression function makes this approach possible. We found that the matching can be regarded as a special three-sum problem. Since the elements in one of the three sets can be restricted in a subspace, we can reduce the complexity to less than 2^n .

The comparison of previous best attacks and our attacks on Grøstl are shown in Table 1. Note that the attacks on Grøstl-0 are not included in this table, since our attack is on the tweaked version.

We also improve the existing attacks against 5-round Whirlpool and 7-round AES hashing modes. While the previous result on 5-round Whirlpool applies to second preimage only, we improve the time complexity and also make the attack work for first preimages. We also improve the time complexity for the attacks against 7-round AES hashing modes. The details are presented in Appendix and the extended version [23] due to space limit.

Outline of This Paper. In Sect. 2, we describe the specification of the Grøstl hash function. In Sect. 3, we introduce the attack outline of the pseudo preimage attack on reduced round Grøstl. Attacks on Grøstl-256 and Grøstl-512 are illustrated in Sect. 4 and Sect. 5 respectively. Sect. 6 is the conclusion.

2 Specification of Grøstl

Grøstl is a double-pipe design, i. e., the size of the chaining value ($2n$ -bit) is twice as the hash size (n -bit). Message length should be less than $2^{73} - 577$ bits. The padding rule is not introduced here, since it’s not important in our attack.

The compression function of Grøstl is written as:

$$F(H, M) = P(H \oplus M) \oplus Q(M) \oplus H$$

Where H is the chaining value and M is the message block, both are of $2n$ bits. After all message blocks are processed, the last chaining value X is used as input of the output transformation, which is written as

$$\Omega(X) = Trunc_n(P(X) \oplus X)$$

The right half of $P(X) \oplus X$ is used as the hash value. The compression function and output transformation are illustrated in Fig. 2.

Table 1. Comparison of the attacks on Grøst1-256 and Grøst1-512

Algorithm	Target	Attack Type	Rounds	Time	Memory	Source
Grøst1-256	Hash Function	Collision	3	2^{64}	-	[18]
	Compression Function	Semi-Free-Start Collision	6	2^{112}	2^{64}	[18]
	Permutation	Distinguisher	8	2^{48}	2^8	[17]
	Output Transformation	Preimage	5	2^{206}	2^{48}	Sect. 4.1
	Hash Function	Pseudo Preimage	5	$2^{244.85}$	$2^{230.13}$	Sect. 4
Grøst1-512	Hash Function	Collision	3	2^{192}	-	[18]
	Compression Function	Semi-Free-Start Collision	7	2^{152}	2^{56}	[17]
	Output Transformation	Preimage	8	2^{495}	2^{16}	Sect. 5.1
	Hash Function	Pseudo Preimage	8	$2^{507.32}$	$2^{507.00}$	Sect. 5




Fig. 2. Compression function and output transformation of Grøst1

P and Q are AES-like permutations with 8×8 and 8×16 sized state for Grøst1-256 and Grøst1-512 separately. Grøst1-256 uses 10-round P , Q and Grøst1-512 uses 14-round P , Q . The round function of the permutations consists of the four operations:

- SubBytes(SB): applies the Substitution-Box to each byte.
- ShiftBytes(SR): cyclically shifts the i -th row leftwards for i positions.
- MixBytes(MC): multiplies each column of the state matrix by an MDS matrix:

$$C = circ(02, 02, 03, 04, 05, 03, 05, 07)$$

- AddRoundConstant(AC): XOR the round constant to the state.

The shift vectors used in P and Q are different. P in Grøstl-256 uses $(0,1,2,3,4,5,6,7)$ and P in Grøstl-512 uses $(0,1,2,3,4,5,6,11)$. In the description of our attack, we skip Q 's detail since it's not required.

An important property of the compression function has been pointed out in the submission document of Grøstl hash function [6]. Note that with $H' = H \oplus M$, the compression function can be written as

$$F(H, M) = P(H') \oplus H' \oplus Q(M) \oplus M.$$

So the generic preimage attack on the compression function with $2n$ -bit state costs 2^n computations, since solving the equation $F(H, M) = T$ can be regarded as a birthday problem. Then the collision attack on the compression function costs $2^{2n/3}$ computations, since $F(H_1, M_1) \oplus F(H_2, M_2) = 0$ is a (four-sum) generalized birthday problem [20].

3 Outline of the Attack on the Grøstl Hash Function

Suppose the hash size is n -bit and the state size is $2n$ -bit. In order to find a pseudo preimage (H, M) of the Grøstl hash function, let $X = F(H, M)$, then X is the preimage of the output transformation: $P(X) \oplus X = *||T$ where T is the target hash value and $*$ stands for arbitrary n -bit value. With $H' = H \oplus M$, we have

$$(P(H') \oplus H') \oplus (Q(M) \oplus M) \oplus X = 0 \tag{1}$$

If we have collected enough candidates for $P(H') \oplus H'$, $Q(M) \oplus M$ and X , the pseudo preimage attack turns into a three-sum problem. As we know, there is no generic solution for three-sum problem faster than birthday attack. But if we can restrict $P(H') \oplus H'$ in a subspace, it is possible to break the birthday bound. Here we restrict $P(H') \oplus H'$ in a subspace by finding its partial zero preimages.

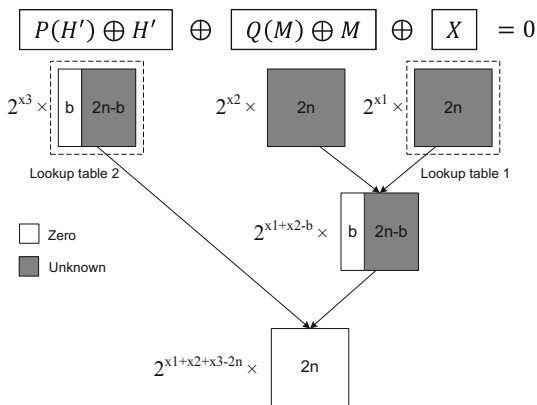


Fig. 3. Outline for pseudo preimage attack on the Grøstl hash function

As illustrated in Fig. 3, the attack process is similar to the generalized birthday attack [20]. With four parameters x_1, x_2, x_3 and b , this attack can be described in four steps:

1. Find 2^{x_1} preimages X of the output transformation and store them in lookup table L_1 .
2. Find 2^{x_3} H' such that leftmost b bits of $P(H') \oplus H'$ are all zero. Then store all $P(H') \oplus H'$ and H' in lookup table L_2 . This step can be regarded as finding partial zero preimages on $P(H') \oplus H'$.
3. Choose 2^{x_2} random M with correct padding and calculate $Q(M) \oplus M$. Then check if there is an X in L_1 with the same leftmost b bits as $Q(M) \oplus M$. We expect to find $2^{x_1+x_2-b}$ partial matches $Q(M) \oplus M \oplus X$ here, whose leftmost b bits are all zero.
4. For each of the $2^{x_1+x_2-b}$ $Q(M) \oplus M \oplus X$ found in step 3, check if its remaining $(2n - b)$ -bit value can be found in L_2 .

Once a final match is found, we have H', M and X which satisfies equation (1). So, $(H' \oplus M, M)$ is a pseudo preimage of **Grøst1**.

Note that to find an X is to find an n -bit partial preimage of $P(X) \oplus X$ and the truncation bits are fixed (the leftmost n -bits are truncated). But for $P(H') \oplus H'$, it's not necessary to find partial preimage for the leftmost b bits. In fact, we can choose any b bits as the zero bits. We will further discuss the differences between fixed position and chosen position partial preimage attacks later.

Suppose that for **Grøst1** with $2n$ -bit state, it takes $2^{C_1(2n,n)}$ computations to find a fixed position n -bit partial preimage and it takes $2^{C_2(2n,b)}$ computations to find a chosen position b -bit partial preimage of $P(X) \oplus X$. Now we calculate the complexity for each of the four attacking steps:

1. Step 1, building the look-up table 1 takes $2^{x_1+C_1(2n,n)}$ computations and 2^{x_1} memory.
2. Step 2, building the look-up table 2 takes $2^{x_3+C_2(2n,b)}$ computations and 2^{x_3} memory.
3. Step 3, calculating $Q(M) \oplus M$ for 2^{x_2} M and checking the partial match in table 1 takes 2^{x_2} Q calls, which is equivalent to 2^{x_2-1} compression function calls.
4. Step 4, checking the final match for $2^{x_1+x_2-b}$ candidates requires $2^{x_1+x_2-b}$ table look-ups, which can be equivalently regarded as $2^{x_1+x_2-b} C_{TL}$ compression function calls. C_{TL} is the complexity of one table lookup, where unit one is one compression function call. For 5-round **Grøst1-256** and 8-round **Grøst1-512**(the attacked versions), C_{TL} is chosen as $1/640$ and $1/2048$ respectively².

² The constant C_{TL} is chosen as the upper bound of the complexity that one table lookup takes, due to the fact that 5-round **Grøst1-256** software implementation composes of $(8 * 8) * 5 * 2 = 640$ s-box lookups, and other operations. In 8-round **Grøst1-512**, there are $(8 * 16) * 8 * 2 = 2048$ s-box lookups.

Then the overall complexity is:

$$2^{x_1+C_1(2n,n)} + 2^{x_3+C_2(2n,b)} + 2^{x_2-1} + 2^{x_1+x_2-b} \cdot C_{TL} \tag{2}$$

with memory requirement of $2^{x_1} + 2^{x_3}$.

In the following sections, we first show how to find partial preimages of the function $P(X) \oplus X$ and calculate the complexity $C_1(2n, n)$ and $C_2(2n, b)$. Then we need to choose optimal parameters x_1, x_2, x_3 and b to minimize the complexity with the restriction of $x_1 + x_2 + x_3 \geq 2n$ and $0 \leq b \leq 2n$. Since in order to find one final match, we need $2^{x_1+x_2+x_3-2n} \geq 1 \Rightarrow x_1 + x_2 + x_3 \geq 2n$.

4 Pseudo Preimage Attack on 5-Round Grøst1-256

In this section, first, we introduce the preimage attack on the output transformation, *i.e.*, the fixed position partial preimage attack on $P(X) \oplus X$ and calculate the complexity $C_1(512, 256)$. Then we introduce the chosen position partial preimage attack on $P(H') \oplus H'$ and give the expression of the function $f(b) = C_2(512, b)$. At last, we try to minimize the overall complexity by finding proper parameters for the generic attack introduced in Section 3.

4.1 Fixed Position Partial Preimage Attack on $P(X) \oplus X$

The chunk separation for this attack is shown in Fig. 4. Note that the yellow cells with a diagonal line are the truncated bytes, which can be regarded as free variables. In the last state of Fig. 4, the equations for the truncated byte can be directly removed since they are automatically fulfilled. The size of the full match is 256-bits for this MitM attack.




Fig. 4. Chunk separation of preimage attack on Grøst1-256’s output transformation

The Colors in the Chunk Separation. First, we explain what the colors stand for. Actually, we use the same colors as in [14] to illustrate the chunk separations. The blue bytes in the forward chunk can be determined by the blue bytes in the initial structure. The white color in the forward chunk stands for the bytes whose values are affected by both red bytes and blue bytes in the initial structure, and can't be pre-computed until the partial match is found. Similarly, in the backward chunk, red and white cells stand for the certain and uncertain bytes. The gray cells are constant bytes in the target value, the chaining value and the initial structure, which are known or can be chosen before the MitM attack.

Freedom Degrees and Size of the Matching Point. Before we apply the MitM attack, we need to know the freedom degrees in the forward and backward directions and the bit size of the matching point. The calculation method has been explained in [14]. More details can be found in the extended version of this paper [23].

We can find that, in Fig. 4, there are $D_2 = 2^{48}$ and $D_1 = 2^{64}$ freedom degrees in red and blue bytes respectively. In each of the four available columns, there are two bytes of matching point. So the size of the matching point is $m = 4 \times (2 \times 8) = 64$ bits.

The Attack Algorithm and Its Complexity. In this section, we consider a generic MitM attack algorithm with partial matching technique. Suppose there are 2^{D_1} and 2^{D_2} freedom degrees in the forward and backward chunks. The size of the matching point is m -bit and the full matching size is b -bit. Without loss of generality, assume that $D_1 \geq D_2$. Note that if $D_1 + D_2 \geq b$, we can't fully use all the freedom degrees. Here we use d_1 and d_2 to denote the actually used freedom degrees:

$$(d_1, d_2) = \begin{cases} (D_1, D_2), & \text{if } D_1 + D_2 \leq b; \\ (b/2, b/2), & \text{if } D_1 + D_2 > b \text{ and } D_2 \geq b/2; \\ (b - D_2, D_2), & \text{if } D_1 + D_2 > b \text{ and } D_2 < b/2. \end{cases} \quad (3)$$

This MitM preimage attack can be described in four steps.

1. Choose random constants in the initial structure.
2. With the chosen constants, for all 2^{d_2} values v_j^2 of the forward direction, calculate all the partial values p_j^2 and the full values f_j^2 at the matching point and store all the pairs (v_j^2, p_j^2) in a look up table L ;
3. For all 2^{d_1} values v_i^1 of the backward direction, calculate p_i^1 . Then check if p_i^1 is in table L . If we found one partial match that $p_i^1 = p_j^2$ for some j , calculate the full value f_i^1 using v_i^1 and check if $f_i^1 = f_j^2$;
4. If no full match has been found yet, go to step 1.

Then we calculate the complexity. Step 2 costs $2^{d_2} f_2$ calls and 2^{d_2} memory. Step 3 costs $2^{d_1} f_1$ calls. Consider two kinds of circumstances separately.

- If $d_1 + d_2 \geq m$. After step 3 is done, we expect $2^{d_1+d_2-m}$ good candidates that satisfy the m -bit matching point. Now check if the full value of all good candidates are matched. This step requires $2^{d_1+d_2-m}$ computations. The probability that a good candidate is a full match is 2^{m-b} . Then the probability that there exists one full match in $2^{d_1+d_2-m}$ good candidates is about $2^{(d_1+d_2-m)+(m-b)} = 2^{d_1+d_2-b}$. So, we need to repeat the attack $2^{b-d_1-d_2}$ times in order to find a full match. The complexity is:

$$2^{b-d_1-d_2} \cdot (2^{d_1} + 2^{d_2} + 2^{d_1+d_2-m}) = 2^b \cdot (2^{-d_1} + 2^{-d_2} + 2^{-m})$$

- If $d_1 + d_2 < m$. After step 3 is done, we can find one good candidate with probability of $2^{d_1+d_2-m}$. So, we need to repeat the attack $2^{m-d_1-d_2}$ times to find one good candidate, then we calculate the full value of the good candidate at the matching point to check if it is a full match, which cost one computation. So the complexity to find one good candidate and check its full value is $2^{m-d_1-d_2}(2^{d_1} + 2^{d_2}) + 1$. Then find and check 2^{b-m} good candidates to get a full match. The complexity is:

$$2^{b-m} \cdot (2^{m-d_1-d_2}(2^{d_1} + 2^{d_2}) + 1) = 2^b \cdot (2^{-d_1} + 2^{-d_2} + 2^{-m})$$

So, no matter in which case, the complexity to find one full match using this algorithm is always

$$2^b \cdot (2^{-d_1} + 2^{-d_2} + 2^{-m}) \quad (4)$$

computations and 2^{d_2} memory.

Application to Grøst1's Output Transformation. In Fig. 4, the freedom degrees are $D_1 = 48, D_2 = 64$, the partial and full matching size are $m = 64$ and $b = 256$ bits. Using the attack algorithm introduced in Section 4.1, we can calculate the complexity to invert 5-round Grøst1's output transformation. Here the complexity is measured by compression function calls. In the MitM attack it takes about half P calls, i. e. $1/4$ compression function calls to evaluate the matching point for one direction. Thus we can multiply 2^{-2} to the complexity: $2^{C_1(512,256)} = 2^{-2} \cdot 2^{256}(2^{-64} + 2^{-48} + 2^{-64}) \approx 2^{206}$ compression function calls with 2^{48} memory.

On the Choice of the Chunk Separation. We can prove that our chunk separation in Fig. 4 is optimal, which minimizes the complexity of inverting the output transformation.

Suppose there are b blue bytes and r red bytes in each column of the matching point. Then we show the relation between b, r , freedom degrees D_1, D_2 and the partial matching size m .

In the forward direction, r red bytes in one column of the matching point $\xrightarrow{AC, SB, SR, MC}$ r full red columns $\xrightarrow{AC, SB, SR}$ r red bytes in one column. Here we stops at the left end of the initial structure. In order to produce at least one byte of freedom degrees in the blue color, there are at least $r + 1$ blue columns

in the initial structure. Then there would be at most $8 - (r + 1) = 7 - r$ red columns in the initial structure.

In the backward direction, b blue bytes in one column of the matching point $\xrightarrow{SR^{-1}, SB^{-1}, AC^{-1}}$ $8 - b$ white columns $\xrightarrow{MC^{-1}, SR^{-1}, SB^{-1}, AC^{-1}}$ $8 - b$ white bytes in each columns.

Now we count the freedom degrees. There are $(7 - r)$ red columns in the initial structure and each column produces $8 - b$ free bytes. So, freedom degrees in red color is $D_2 = 8(7 - r)(8 - b)$ bits. The minimum freedom degrees in the blue color here is $D_1 = 2^{64}$. Size of the matching point in one column is $8(b + r - 8)$ bits, so there are $4 \times 8(b + r - 8)$ bits of matching point in total.

So the complexity is $2^{-2} \cdot 2^{256}(2^{-64} + 2^{-8(7-r)(8-b)} + 2^{-32(b+r-8)})$. The minimum complexity is 2^{206} when $b = 6, r = 4$ or $b = 5, r = 5$. Fig. 4 is the case of $b = 6, r = 4$.

4.2 Chosen Position Partial Preimage Attack on $P(H') \oplus H'$

Now, consider the attack model of chosen position partial preimage. In the partial preimage attack of $P(H') \oplus H'$, we can choose the positions of the target bits. In order to minimize the complexity, we choose this chunk separation to maximize the size of the matching point $m(b)$ within all possible b target bits.




Fig. 5. Chunk separation of chosen position partial preimage attack on $P(H') \oplus H'$ for Grøst1-256

First, we discuss the size of matching point and chosen positions in one column. If less than 8 bits of the red byte in one column are chosen, no matching point can be derived. if $b > 8$ bits of the red bytes are chosen, there are $b - 8$ bits of matching point. Since there are only two red bytes in one column in the last state of Fig. 5, even if $b > 16$, no more than 8 bits of matching point can be derived. In order to maximize $m(b)$, we choose at most 2 red bytes in one column and then chose the red bytes from another column. When $b > 128$, $m(b) = 64$, because there are 64 bits of matching point in total. The graph of $m(b)$ is shown in Fig. 6.

In this Figure, freedom degrees in the red and blue color are $D_2 = 40$ and $D_1 = 64$. Then we can calculate the complexity of chosen position partial preimage:

$$2^{C_2(512,b)} = 2^{-2} \cdot 2^b(2^{-d_1} + 2^{-d_2} + 2^{-m(b)})$$

where d_1 and d_2 are chosen according to equation (3), i.e.:

$$(d_1, d_2) = \begin{cases} (64, 40), & \text{if } b \geq 104; \\ (b - 40, 40), & \text{if } 80 \leq b < 104; \\ (b/2, b/2), & \text{if } b < 80. \end{cases}$$

The graph of $C_2(512, b)$ is shown in Fig. 7. When $b > 80$, $C_2(512, b) \approx b - 42$.




Fig. 6. Size of the matching point for chosen position truncations for Grøstl-256




Fig. 7. Complexity of chosen position partial preimage of $P(H') \oplus H'$ for Grøstl-256

4.3 Minimizing the Overall Complexity

By now, we have found $C_1(512, 256)$ and $C_2(512, b)$. So we can start to deal with the overall complexity in equation (2). In the expression of the complexity, b can




Fig. 8. Overall complexity of pseudo preimage attack on 5-round Grøstl-256

be integers from 0 to 512. For all $b \in [0, 512]$, optimal x_1, x_2 and x_3 are chosen to minimize the overall complexity. The graph of the minimum overall complexity for $b \in [0, 120]$ is shown in Fig. 8.

When $b = 31, x_1 \approx 36.93, x_2 \approx 244.93$ and $x_3 \approx 230.13$, the complexity is the lowest: $2^{244.85}$ compression function calls. Memory requirement is $2^{230.13}$. The chosen positions for the 31 bits ≈ 4 bytes are marked in Fig. 5.

5 Pseudo Preimage Attack on 8-Round Grøstl-512

The attack on Grøstl-512 uses the same method for the three-sum phase as in the attack on Grøstl-256. Here we skip the details of the attack algorithm and introduce the difference between the attacks on them only.

5.1 Fixed Position Partial Preimage Attack on $P(X) \oplus X$

The chunk separation for 8-round Grøstl-512 is shown in Fig. 9. Note that in this figure, we use a 2-round initial structure. Freedom degrees in the red and blue bytes are both 2^{16} . There are 4 bytes of matching point in total.

The parameters for the MitM preimage attack on the output transformation are $D_1 = D_2 = 16, m = 32$ and $n = b = 512$. So the complexity is $2^{C_1(1024, 512)} = 2^{-2} \cdot 2^{512}(2^{-16} + 2^{-16} + 2^{-32}) \approx 2^{495}$ compression function calls and 2^{16} memory.

On the Choice of the Chunk Separation. Actually, we searched for all the possible patterns of the chunk separation for 8-round Grøstl-512. The chunk separation in Fig. 9 is one of the best we found. The search algorithm is as follows:

Step 1. Search for the matching point.

We want to find good candidates in all the possible positions of the white columns in round 2 and round 6. Since there are 32 columns in two states, there are 2^{32} patterns in total.

For each of the pattern of white columns, we can calculate round 2 backward and round 6 forward and check if there are at least two byte of matching point. After the search for all the 2^{32} patterns, we found 1322 patterns with at least two bytes of matching point.

Step 2. Search for the initial structure.

Considering the mirror image and rotational similarity, there are only 120 distinct patterns in all the 1322 patterns of matching point. For each of the 120 patterns, we calculate forward from round 2 and backward from round 6.

If there is one white column in round 2, the number of possible patterns of the white bytes in the same column of round 3 is $2^8 - 1$, since there must be at least one white byte in this column. So size of the search space is $(2^8 - 1)^w$, where w is the number of white columns in both round 2 and round 6. In the 120 possible patterns, w is no more than 4, so the search space is at most $2^{32} \cdot 120 \approx 2^{39}$.

Using early-abort trick, we can directly skip some bad patterns in round 2 without knowing the pattern in round 6. Then the search space is reduced again and the search is practical.




Fig. 9. Chunk separation of preimage attack on Grøst1-512's output transformation

5.2 Chosen Position Partial Preimage Attack on $P(H') \oplus H'$

For chosen position partial preimage, we use another chunk separation in Fig. 10. The freedom degrees for the MitM preimage attack are $D_1 = 24, D_2 = 8$. Then we can calculate the complexity of chosen position partial preimage:

$$2^{C_2(1024,b)} = 2^{-2} \cdot 2^b(2^{-d_1} + 2^{-d_2} + 2^{-m(b)})$$

where d_1 and d_2 are chosen according to equation (3), *i.e.*,

$$(d_1, d_2) = \begin{cases} (24, 8), & \text{if } b \geq 32; \\ (b - 8, 8), & \text{if } 16 \leq b < 32; \\ (b/2, b/2), & \text{if } b < 16. \end{cases}$$

5.3 Minimizing the Overall Complexity

With the value and expression of $C_1(1024, 512)$ and $C_2(1024, b)$, we can deal with the overall complexity like we have done for Grøst1-256. When $b = 0, x_1 \approx 10.50, x_2 \approx 506.50$ and $x_3 \approx 507.00$, the overall complexity is the lowest: $2^{507.32}$. Memory requirement is $2^{507.00}$.




Fig. 10. Chunk separation of chosen position partial preimage attack on $P(H') \oplus H'$ for Grøst1-512

6 Conclusion

In this paper, we proposed pseudo preimage attacks on the hash functions of 5-round Grøst1-256 and 8-round Grøst1-512. This is the first pseudo preimage attack on round-reduced Grøst1 hash function, which is a wide-pipe design.

In order to invert the wide-pipe hash function, we have to match $2n$ -bit state value with less than 2^n computations. This is achieved by exploiting the special property of the Grøst1 compression function. After collecting enough partial preimages on the component $P(X) \oplus X$, the double-sized state values are matched using a variant of the generalized birthday attack.

There is an interesting observation that this attack works with any function Q . Thus our attack can be applied to the Grøst1 hash function with round-reduced permutation P and full-round permutation Q . However, our attacks do not threaten any security claims of Grøst1.

Acknowledgement. The authors would like to thank Kazumaro Aoki, Keting Jia, Mohammad Ali Orumiehchiha, Somitra Sanadhya, and Chunhua Su for their inspiring suggestions during the ASK 2011 workshop. The authors would also thank Lei Wang for useful discussions, Praveen Gauravaram for improving the editorial quality of this paper and reviewers of FSE 2012 for helpful comments.

This work is supported by the National Natural Science Foundation of China (No.60873259 and No.60903212), National Science and Technology Major Project of China(No.2011ZX03002-005-02) and the Knowledge Innovation Project of The Chinese Academy of Sciences.

References

1. Aoki, K., Guo, J., Matusiewicz, K., Sasaki, Y., Wang, L.: Preimages for Step-Reduced SHA-2. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 578–597. Springer, Heidelberg (2009)
2. Aoki, K., Sasaki, Y.: Preimage Attacks on One-Block MD4, 63-Step MD5 and More. In: Avanzi, R.M., Keliher, L., Sica, F. (eds.) SAC 2008. LNCS, vol. 5381, pp. 103–119. Springer, Heidelberg (2009)
3. Aoki, K., Sasaki, Y.: Meet-in-the-Middle Preimage Attacks Against Reduced SHA-0 and SHA-1. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 70–89. Springer, Heidelberg (2009)
4. Barreto, P.S.L.M., Rijmen, V.: The whirlpool hashing function. Submission to NESSIE (September 2000)
5. Chabaud, F., Joux, A.: Differential Collisions in SHA-0. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 56–71. Springer, Heidelberg (1998)
6. Gauravaram, P., Knudsen, L.R., Matusiewicz, K., Mendel, F., Rechberger, C., Schl affer, M., Thomsen, S.S.: Gr ostl – a SHA-3 candidate. Submission to NIST, Round 3 (2011)
7. Guo, J., Ling, S., Rechberger, C., Wang, H.: Advanced Meet-in-the-Middle Preimage Attacks: First Results on Full Tiger, and Improved Results on MD4 and SHA-2. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 56–75. Springer, Heidelberg (2010)
8. Hong, D., Koo, B., Sasaki, Y.: Improved Preimage Attack for 68-Step HAS-160. In: Lee, D., Hong, S. (eds.) ICISC 2009. LNCS, vol. 5984, pp. 332–348. Springer, Heidelberg (2010)
9. Joux, A.: Multicollisions in Iterated Hash Functions. Application to Cascaded Constructions. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 306–316. Springer, Heidelberg (2004)
10. Kelsey, J., Schneier, B.: Second Preimages on n -Bit Hash Functions for Much Less than 2^n Work. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 474–490. Springer, Heidelberg (2005)
11. Kiayias, A. (ed.): CT-RSA 2011. LNCS, vol. 6558. Springer, Heidelberg (2011)
12. Leurent, G.: MD4 is Not One-Way. In: Nyberg, K. (ed.) FSE 2008. LNCS, vol. 5086, pp. 412–428. Springer, Heidelberg (2008)
13. National Institute of Standards and Technology. Announcing Request for Candidate Algorithm Nominations for a New Cryptographic Hash Algorithm (SHA-3) Family. Federal Register 27(212), 62212–62220 (2007), http://csrc.nist.gov/groups/ST/hash/documents/FR_Notice_Nov07.pdf (October 17, 2008)
14. Sasaki, Y.: Meet-in-the-Middle Preimage Attacks on AES Hashing Modes and an Application to Whirlpool. In: Joux, A. (ed.) FSE 2011. LNCS, vol. 6733, pp. 378–396. Springer, Heidelberg (2011)
15. Sasaki, Y., Aoki, K.: Preimage Attacks on 3, 4, and 5-Pass HAVAL. In: Pieprzyk, J. (ed.) ASIACRYPT 2008. LNCS, vol. 5350, pp. 253–271. Springer, Heidelberg (2008)

16. Sasaki, Y., Aoki, K.: Finding Preimages in Full MD5 Faster Than Exhaustive Search. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 134–152. Springer, Heidelberg (2009)
17. Sasaki, Y., Li, Y., Wang, L., Sakiyama, K., Ohta, K.: Non-full-active Super-Sbox Analysis: Applications to ECHO and Grøstl. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 38–55. Springer, Heidelberg (2010)
18. Schl affer, M.: Updated Differential Analysis of Grøstl. Grøstl website (January 2011)
19. Sekar, G., Mouha, N., Velichkov, V., Preneel, B.: Meet-in-the-Middle Attacks on Reduced-Round XTEA. In: Kiayias [11], pp. 250–267
20. Wagner, D.: A Generalized Birthday Problem. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 288–303. Springer, Heidelberg (2002)
21. Wang, L., Sasaki, Y., Komatsubara, W., Ohta, K., Sakiyama, K.: (Second) Preimage Attacks on Step-Reduced RIPEMD/RIPEMD-128 with a New Local-Collision Approach. In: Kiayias [11], pp. 197–212
22. Wei, L., Rechberger, C., Guo, J., Wu, H., Wang, H., Ling, S.: Improved Meet-in-the-Middle Cryptanalysis of KTANTAN (Poster). In: Parampalli, U., Hawkes, P. (eds.) ACISP 2011. LNCS, vol. 6812, pp. 433–438. Springer, Heidelberg (2011)
23. Wu, S., Feng, D., Wu, W., Guo, J., Dong, L., Zou, J.: (Pseudo) Preimage Attack on Round-Reduced Grøstl Hash Function and Others (Extended Version). Cryptology ePrint Archive, Report 2012/206 (2012), <http://eprint.iacr.org/>

A Preimage Attack on Round-Reduced Whirlpool

A.1 Specification of Whirlpool

Whirlpool uses MD-strengthening structure, with narrow-pipe chaining value and no block counters. So it is vulnerable to generic attack, like the expandable messages [10] and multi-target pseudo preimage [12] attack. We will talk about the details later.

Whirlpool accepts any message with less than 2^{256} bits as input and the 256-bit binary expression of bit length is padded according to MD-strengthening, i. e. $M||1||0^*||length$. Size of the message block, the chaining value and the hash value is 512-bit.

Compression function of Whirlpool can be regarded as a block cipher called W in Miyaguchi-Preneel mode.

$$F(H, M) = W_H(M) \oplus M \oplus H$$

where block cipher W use AES-like iteration with 8×8 state of bytes and the $(8i + j)$ -th input byte of the message block is placed at the i -th row and j -th column of the state. Each round consists of four operations:

- SubBytes(SB): applies the Substitution-Box to each byte.
- ShiftColumns(SC): cyclically shift the i -th column downwards for i positions.
- MixRows(MR): multiply each row of the state matrix by an MDS matrix

$$C = circ(01, 01, 04, 01, 08, 05, 02, 09)$$

- AddRoundKey(AK): XOR the round key to the state.

Since the key schedule is not important in our attack, the description is omitted.

A.2 Improved Second Preimage Attack on Whirlpool

In [14], Yu Sasaki proposed a second preimage attack on 5-round **Whirlpool** using the MitM approach. In their attack, there are only 2^8 freedom degrees in both chunks, but the size of matching point is much larger (40 bytes=320 bits). The comparison of the preimage attacks on **Whirlpool** is shown in Table 2.

Table 2. Comparison of the preimage attacks on **Whirlpool**

Attack Type	Rounds	Time	Memory	Source
Second Preimage	5	2^{504}	2^8	[14]
Second Preimage	5	2^{448}	2^{64}	this section
Preimage	5	$2^{481.5}$	2^{64}	this section

In this section, we propose an improved chunk separation with more freedom degrees and a smaller matching point in Fig. 11.

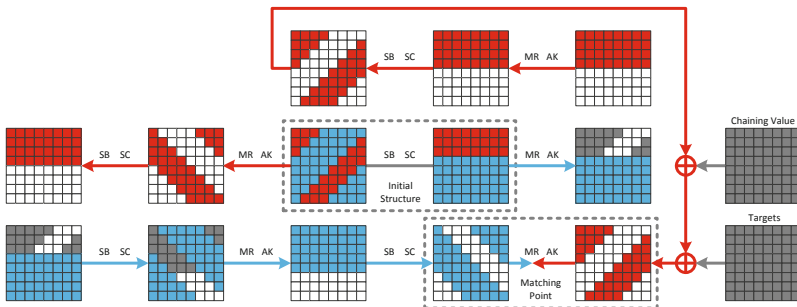


Fig. 11. Chunk separation for improved 2nd-preimage attack on 5-round **Whirlpool**

We use the same colors as in [14] to illustrate the chunk separations. The blue bytes in the forward chunk can be determined by the previously chosen blue bytes in the initial structure. The white bytes in the forward chunk stands for the bytes whose value are affected by the red bytes from initial structure and

can't be precomputed until the partial match is found. Similarly, in the backward chunk, red and white cells stand for the certain and uncertain bytes. The gray cells are constant bytes in the target value, the chaining value and the initial structure.

Since this is a second preimage attack, the second last chaining value and the last message block with proper padding are known. We choose random messages and get a random chaining value at the third last position. With this chaining value, apply MitM preimage attack of the compression function.

With chunk separation in Fig. 11, we have a MitM attack with $D_1 = 72$, $D_2 = 64$, $m = 64$ and $b = 512$. According to equation 4, the complexity can be computed as $2^{-1}2^{512}(2^{-72} + 2^{-64} + 2^{-64}) \approx 2^{448}$. Memory requirement is 2^{64} . Note that the complexity of computing the two chunks and checking the full match may be different. Here, we don't consider the difference between them and they are all regarded as the same cost of half compression function call.

A.3 First Preimage Attack on Whirlpool

This attack consists of three steps: First, find a preimage of the last block with proper padding. Second, construct an expandable message. At last, connect expandable message and the last block with MitM. The attack process is illustrated in Fig. 12.




Fig. 12. Outline of the first preimage attack on 5-round Whirlpool

Dealing with Message Padding. In order to apply the first preimage attack, the message padding must be dealt with properly. In our attack, the last message block consists of 255-bit message concatenated with one bit of “1” padding and 256-bit binary expression of the message length l . Since Whirlpool uses 512-bit message block, $l \equiv 255 \pmod{512}$. Then the last 9 bits of l are fixed to 011111111.

In Fig. 13, the initial structure is relocated at the beginning of the compression function for the convenience of the message padding, since in MP mode, the first state is the message block itself. Value of the black byte in the first state is fixed to $0xff$, because it is the last 8 bits of l . One red byte is marked with a “0”, which means the last bit of it is fixed to zero due to the message length l . There is another blue byte marked with a “1”, which comes from the “1-0” padding.

Parameters for this MitM attack are $D_1 = D_2 = 63$, $m = 64$ and $b = 512$. According to equation 4, the complexity is about 2^{449} computations and 2^{63} memory for the last block. When the attack on the last block is done, the remaining bits of the message length are fulfilled by expandable messages.

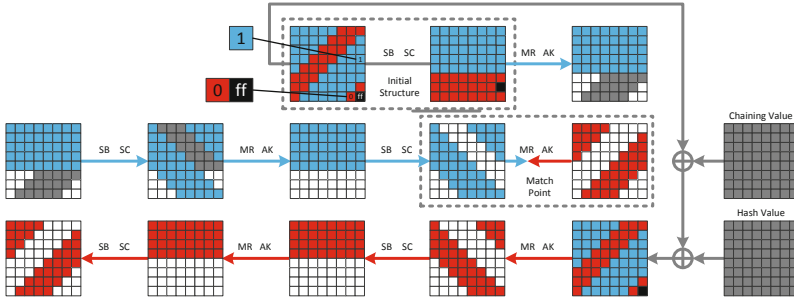


Fig. 13. Chunk separation for the last message block of Whirlpool

Expandable Messages. Expandable messages [10] can be constructed using either Joux’s multi-collision [9] or fix points of the compression function.

Expandable 2^k -collision can be constructed with $k \cdot 2^{n/2}$ computations and k memory. But its length can only be in the range of $[k, k + 2^k - 1]$ blocks. If the message length obtained from the last block is less than k (with a very small probability), we choose different random constants and repeat the attack.

Fix points of MP mode can be constructed by finding the zero preimages of the compression function in MMO mode, since

$$W_H(M) \oplus M \oplus H = H \Leftrightarrow W_H(M) \oplus M = 0.$$

This can be done using the same technique as in our 2nd-preimage attack, with complexity of $(2^{449}, 2^{64})$, which is an affordable cost for us. Note that for random H , the fix point exists with probability of $1 - e^{-1}$. If no fix point can be found for IV, we choose a random message block, compute the following chaining value and try to find fix point for this chaining value instead.

So, either way is fine to construct the expandable message here and has little influence on the overall complexity.

Turns Pseudo Preimage into Preimage. After preparing preimage for the last message block $F(H, M) = T$ and the expandable message $IV \xrightarrow{M^*} H'$. Now we can connect them to form a first preimage.

Suppose it takes 2^c to find a pseudo preimage. A traditional MitM approach can convert preimage attack on the compression function into preimage attack on the hash function works like this. First, find and store 2^k pseudo preimages with 2^{k+c} computations and 2^k memory. Then choose 2^{n-k} random message, calculate from IV to find a chaining value appearing in one of the pseudo preimages. The complexity is $2^{n-k} + 2^{k+c}$. Take the optimal $k = \frac{n-c}{2}$, the minimum complexity is $2^{\frac{n+c}{2}+1}$. Using the pseudo preimage attack described in Sect. A.2, the preimage attack has a complexity of $(2^{481.5}, 2^{64})$.

Practical Cryptanalysis of ARMADILLO2

María Naya-Plasencia^{1,*} and Thomas Peyrin^{2,**}

¹ University of Versailles, France

`maria.naya-plasencia@prism.uvsq.fr`

² Division of Mathematical Sciences, School of Physical and Mathematical Sciences,
Nanyang Technological University, Singapore

`thomas.peyrin@gmail.com`

Abstract. The ARMADILLO2 primitive is a very innovative hardware-oriented multi-purpose design published at CHES 2010 and based on data-dependent bit transpositions. In this paper, we first show a very unpleasant property of the internal permutation that allows for example to obtain a cheap distinguisher on ARMADILLO2 when instantiated as a stream-cipher. Then, we exploit the very weak diffusion properties of the internal permutation when the attacker can control the Hamming weight of the input values, leading to a practical free-start collision attack on the ARMADILLO2 compression function. Moreover, we describe a new attack so-called local-linearization that seems to be very efficient on data-dependent bit transpositions designs and we obtain a practical semi-free-start collision attack on the ARMADILLO2 hash function. Finally, we provide a related-key recovery attack when ARMADILLO2 is instantiated as a stream cipher. All collision attacks have been verified experimentally, they require negligible memory and a very small number of computations (less than one second on an average computer), even for the high security versions of the scheme.

Keywords: ARMADILLO2, hash function, stream-cipher, MAC, cryptanalysis, collision.

1 Introduction

Hash functions are among the most important and widely spread primitives in cryptography. Informally a hash function H is a function that takes an arbitrarily long message as input and outputs a fixed-length hash value of size n bits. The classical security requirements for such a function are collision resistance and (second)-preimage resistance. Namely, it should be impossible for an adversary to find a collision (two different messages that lead to the same hash value) in less than $2^{n/2}$ hash computations, or a (second)-preimage (a message hashing to a given challenge) in less than 2^n hash computations. In general, a hash function H

* Supported by the French Agence Nationale de la Recherche through the SAPHIR2 project under Contract ANR-08-VERS-014.

** Supported by the Lee Kuan Yew Postdoctoral Fellowship 2011 and the Singapore National Research Foundation Fellowship 2012.

is built from an iterative use of a n -bit output compression function h in a Merkle-Damgård-like operating mode [6,4]. The compression function takes a chaining variable CV (fixed to an initial value IV at the beginning) and a message block M as inputs and in order to allow security proofs on the operating mode, one requires the same security properties as a hash function, namely collision and (second)-preimage resistance. However, the compression function allows several flavors of security properties depending on how well the attacker can control the chaining variable:

- free-start collision: the attacker fully controls the chaining variable, i.e. both its value and difference
- semi-free-start collision: the attacker control partially the chaining variable, i.e. only its value, and the difference is null
- collision: the attacker does not control the chaining variable, the value is defined by the IV and the difference is null

For all three flavors, it should be impossible for an adversary to find a collision in less than $2^{n/2}$ compression function computations. Note that free-start collision is required as necessary assumption regarding the compression function in the Merkle-Damgård-like security proofs. Moreover, a semi-free-start collision means there exists initial values IV for which it is possible to find collisions for the hash function. Therefore, both these two notions are very important and should be verified for a secure compression function.

ARMADILLO2 [2] is a very novel primitive dedicated to hardware, defining a FIL-MAC, a stream cipher and a hash function. Originally, two versions were proposed, ARMADILLO and ARMADILLO2, the later being the recommended one. A key recovery attack on ARMADILLO was rapidly published by a subset of the designers [9]. ARMADILLO2 remained unbroken until Abdelraheem *et al.* [1] found a meet-in-the-middle technique that allows to invert the ARMADILLO2 main function. This cryptanalysis eventually led to a key recovery attack on the FIL-MAC and the stream cipher, and a (second)-preimage attack on the hash function. However, while being the first weakness published on ARMADILLO2, this work is an improved meet-in-the-middle technique, therefore requiring a lot of computations and memory, often close to the generic complexity. For example, the preimage attack on the 256-bit output hash function requires either 2^{208} computations and 2^{205} memory or 2^{249} computations and 2^{45} memory. With its data-dependent bit transpositions and original compression function construction, ARMADILLO2 is clearly not following the classical design trends for symmetric-key primitives (for example RC5 [7] and RC6 [8] use data-dependent rotations, while IDEA [5] use data-dependent multiplication). As a consequence, it would be interesting to look at this proposal without necessarily relying on known cryptanalysis techniques.

Our Contributions. In this paper, we first observe the very unpleasant property that the parity bit is preserved through all ARMADILLO2 internal permutations. This allows us for example to derive a very cheap distinguisher for the stream-cipher. Then, we analyze the differential diffusion of the permutations and we provide practical free-start collision attacks for all versions of the

compression function of ARMADILLO2. We extend our results by introducing a new technique, the *local linearization*, that seems very efficient against data-dependent bit transpositions. This method led us to practical semi-free-start collision attacks for all versions of ARMADILLO2. All attacks require very few computations (at most $2^{10.2}$ operations for 256-bit output version) and negligible memory. Moreover, our implementations validate our techniques and we provide collision examples. Finally, we provide a related-key recovery attack when ARMADILLO2 is instantiated as a stream cipher.

2 The ARMADILLO2 Function

We let $X[i]$ denote the i -th bit of a word X . Let C be an initial vector of size c and U be a message block of size m . The size of the register ($C||U$) is $k = c + m$, where $||$ denotes the concatenation operation. The internal ARMADILLO2 function transforms the vector (C, U) into (V_c, V_t) as described in Figure 1, $(V_c, V_t) = \text{ARMADILLO2}(C, U)$. The internal ARMADILLO2 function relies on a parameterized permutation on k bits Q , instantiated by Q_U and Q_X , where U is a m -bit parameter and X is a k -bit parameter.

Let σ_0 and σ_1 be two fixed bitwise permutations of size k . In [2], the permutations are not specifically defined but some criteria they should fulfill is given.

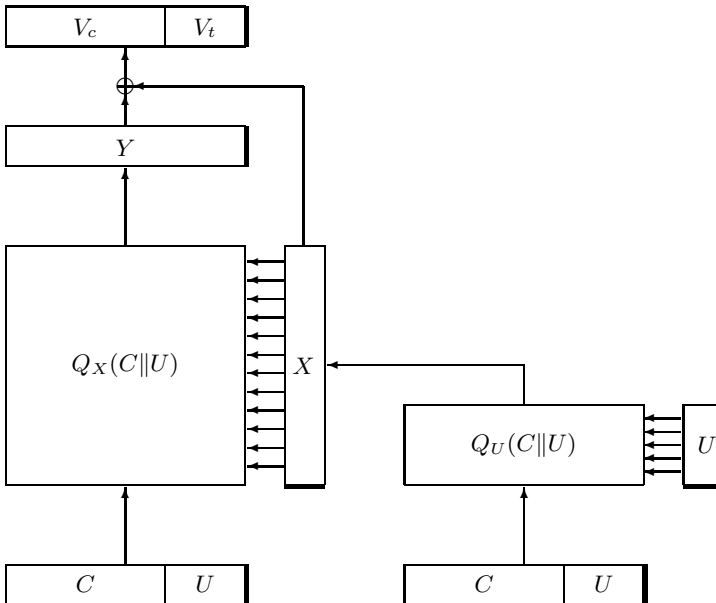


Fig. 1. The internal function of ARMADILLO2. The thick line at the side of a register represents the least significant bit.

We denote by cst a constant of size k defined by alternating 0's and 1's, i.e. : $cst = 1010 \cdots 10$. Using these notations, we can specify Q which is used twice in the internal ARMADILLO2 function. Let A be the a -bit parameter and B be the k -bit input of Q , the parameterized permutation Q_A can be divided into $a = |A|$ simple steps. The i -th step of Q_A (reading A from its least significant bit to its most significant one) is defined by:

- an elementary **bitwise permutation**: $B \leftarrow \sigma_{A[i]}(B)$, that is if the i -bit of A is 0 we apply σ_0 to B , otherwise we apply σ_1 .
- a **constant addition** (bitwise XOR) of cst : $B \leftarrow B \oplus cst$.

The internal ARMADILLO2 function first computes $X = Q_U(C||U)$, then $Y = Q_X(C||U)$, and finally outputs $(V_c, V_t) = Y \oplus X$.

Using this internal primitive, ARMADILLO2 builds a FIL-MAC, a stream-cipher and a hash function:

- **Stream-cipher**: the secret key is inserted in the C register and the output sequence is obtained by taking the k bits of the output (V_c, V_t) after one iteration. The keystream is composed of k -bit frames indexed by U (which is a public value).
- **Hash function**: it uses a strengthened Merkle-Damgård construction, where V_c represents the output of the compression function (i.e. the next chaining value or the hash digest), U is the incoming message block and C is the incoming chaining variable.
- **FIL-MAC**: the secret key is inserted in the C register and the challenge, considered known by the attacker, is inserted in the U register. The response to the challenge is the m -bit output V_t .

Five different sets of register sizes (k, c, m) are provided, namely $(128, 80, 48)$, $(192, 128, 64)$, $(240, 160, 80)$, $(288, 192, 96)$ and $(384, 256, 128)$.

3 First Tools

We denote $\text{HAM}(X)$ the Hamming weight of the word X . We recall from [1] that for two random k -bit words A and B of Hamming weight a and b respectively, the probability that $\text{HAM}(A \wedge B) = i$ (where \wedge stands for the bitwise AND function) is given by the formula

$$P_{\text{and}}(k, a, b, i) = \frac{\binom{a}{i} \binom{k-a}{b-i}}{\binom{k}{b}} = \frac{\binom{b}{i} \binom{k-b}{a-i}}{\binom{k}{a}}.$$

Moreover, we would like to deduce from it the probability that $\text{HAM}(A \oplus B) = i$ (where \oplus stands for the bitwise XOR function) for two randomly chosen k -bit words A and B of Hamming weight a and b respectively. We remark that $\text{HAM}(A \oplus B) = a + b - 2 \cdot \text{HAM}(A \wedge B)$ and therefore the probability that $\text{HAM}(A \oplus B) = i$ is given by the formula.

$$P_{\text{XOR}}(k, a, b, i) = \begin{cases} P_{\text{and}}(k, a, b, \frac{a+b-i}{2}) & \text{for } (a + b - i) \text{ even} \\ 0 & \text{for } (a + b - i) \text{ odd} \end{cases}$$

Since they have not been specified in the original ARMADILLO2 document, in the following we assume that σ_0 and σ_1 are randomly chosen bit permutations.

4 Parity Preservation

We call the parity bit of an a -bit word A the bit value $\bigoplus_{i=0}^{a-1} A[i]$. Regardless of the parameter A of the internal permutation Q_A , we have that **the parity of the input is always maintained through the permutation**. This can be easily verified by remarking that Q_A is composed of several identical rounds, all satisfying this property. Indeed, one round is composed of a bit permutation (which fully maintains the Hamming weight) and an XOR of the internal state with the constant $cst = 1010 \dots 10$. This constant being always the same during the whole ARMADILLO2 computation and its parity being even, the parity of the internal state remains the same after application of the XOR. Note that even if this constant was changed during the rounds, the attacker would only have to compute the parity of the XOR of all constants to be able to tell if the parity bit will be maintained or negated. This property is moreover maintained whatever number of rounds is applied in the permutations, thus the attack proposed in this section is independent of the number of rounds.

Distinguisher for the Stream Cipher Mode. We can exploit the previous property to build a cheap distinguisher on ARMADILLO2 when used as a stream-cipher. In the attack model, the whole output of the function is assumed to be known as it is a frame of the keystream. This output is generated by a XOR of internal states X and Y . Since permutations Q_U and Q_X will maintain the parity, their respective outputs X and Y will both have the same parity as $(C||U)$. As a consequence, the output of the function $X \oplus Y$ always has an even parity. For a random sequence, this will only happen with probability $1/2$, as for ARMADILLO2 this happens with probability 1. In other words, the entropy of the ARMADILLO2 function output is reduced by one bit.

5 Controlled Diffusion: Practical Free-Start Collision Attack

In this section, we show how an attacker can control the bit difference diffusion in ARMADILLO2 function by using the available inputs. This leads to a very cheap free-start collision attack against the compression function.

5.1 General Description

Assume that we insert a single bit difference in C , that is $\text{HAM}(\Delta C) = 1$, and no difference in U that is $\Delta U = 0$. We can use c distinct ΔC , one for each active bit position. The attack is depicted in Figure 2.




Fig. 2. A schematic view of the free-start collision attack on ARMADILLO2. The thick line at the side of a register represents the least significant bit and black circles stand for bit differences. The dashed box indicates the first round of Q_X , which contains a difference on its corresponding parameter input bit.

Difference Propagation in Q_U . Since we have no difference in U , the permutation Q_U always remains the same. We only have to study the propagation of the bit difference in C through Q_U . Note that one round of the internal permutation Q_U provides no difference diffusion since it is only composed of a bit permutation and a constant addition. Therefore, the single bit difference in C will be just transferred to some random bit position in X at the end of Q_U and we have $\text{HAM}(\Delta X) = 1$. We would like the single bit difference in X to be positioned in bit 0, i.e. $\Delta X = 00 \dots 01$ (this will later allow us to use the freedom degrees efficiently). For a randomly chosen value of U and C , this happens with probability

$$P_X = \frac{1}{k}.$$

Difference Propagation in Q_X . Since we have a single difference on the first bit of X (corresponding to the first step of Q_X), the permutation Q_X remains the same except for the first step where we switch from bit permutation σ_0 to σ_1 or from σ_1 to σ_0 . We denote by $P_{step}(in, out)$ the probability that in active bits are mapped to out active bits through a step of data-dependent permutation with a difference (i.e. σ_0 and σ_1 are swapped). Assume for the moment that after this first step, only b bits are active in the internal state. This happens with probability $P_{step}(1, b)$. Since the next rounds of the internal permutation Q_X provide no difference diffusion, we end up in Y with b active bits randomly distributed. We need to ensure that all the b active bits remaining in Y will go to the m -bit V_t part of the k -bit output, so that all differences will be truncated and we eventually obtain a collision on the output of the compression function. For $b \leq m$, this happens with probability

$$P_{out}(b) = P_{\text{and}}(k, m, b, b) = \frac{\binom{b}{b} \binom{k-b}{m-b}}{\binom{k}{m}} = \prod_{i=0}^{b-1} \frac{m-i}{k-i}.$$

During the feed-forward after Q_X the single active bit of X is already on the V_t part of the output. Overall the probability of obtaining a compression function collision for randomly chosen U and C values is:

$$P_{collision} = P_X \cdot \sum_{i=1}^{i=m} P_{step}(1, i) \cdot P_{out}(i).$$

the sum stopping at m because when $i > m$, we trivially have $P_{out}(i) = 0$. At this point our problem is that in order for the probability $P_{out}(i)$ to be high enough, we need the number i of active bits to be small. On the other side, if i is small, $P_{step}(1, i)$ will be very low (we do not explain how to compute $P_{step}(1, i)$ here as we will study a slightly more detailed problem in the next section). However, in this scenario we only considered an attacker that randomly chooses the value of U and C and the bit difference position in C , but we can do much better by using the available degrees of freedom efficiently.

5.2 Using the Freedom Degrees

First, note that the event related to the probability P_X only depends on the position of the bit difference in C and on the value of U . We can therefore attack Q_U in a first phase (by fixing the position of the bit difference in C and the value of U), and then independently attack Q_X by choosing the value of C .

Handling Q_U . We will see later that we would like C and U values to have an extremely low or extremely high Hamming weight. Therefore, we fix $\Delta X = 00\dots 01$ and test with the two values $U = 00\dots 00$ and $U = 11\dots 11$ how the bit difference will propagate through Q_U^{-1} (note that we are dealing with the inverse of Q_U , thus attacking backwards from ΔX). For each try, we have a probability

$P_{\text{and}}(k, c, 1, 1) = c/k$ that the single bit difference is mapped to the C part of the input. Since for all ARMADILLO2 versions we have $2c/k > 1$, we expect at least one of the two U candidates to satisfy $\Delta X = 00\dots 01$, $\text{HAM}(\Delta C) = 1$ and $\text{HAM}(\Delta U) = 0$. Overall, this phase costs us only 2 operations. We assume without loss of generality that the selected candidate has value $U = 00\dots 00$.

Handling Q_X . At the present time, everything is fixed except the value of C and we have $\Delta X = 00\dots 01$ and $U = 00\dots 00$. We now describe a simple criteria in order to choose the values of C such that the first round probability $P_{\text{step}}(1, i)$ in Q_X is high, even for small i . As an example, let's assume that $C = 0$, that is $\text{HAM}(C||U) = 0$. In that case, we trivially have that $P_{\text{step}}(1, 1) = 1$ (and $P_{\text{step}}(1, i) = 0$ for all other i) since changing the bit positions of the word $00\dots 00$ (switching from σ_0 to σ_1 or from σ_1 to σ_0) will not have any effect at all and the single bit difference in C will just be placed to some random bit position. Similarly, with a single one-bit in C , that is $\text{HAM}(C||U) = 1$, we have that $P_{\text{step}}(1, 1) = \frac{1}{128} + \frac{2 \cdot 127}{128^2}$ and $P_{\text{step}}(1, 3) = \frac{127 \cdot 126}{128^2}$ (and $P_{\text{step}}(i) = 0$ for all other i). More generally, we have to compute the probability $P_{\text{step}}(1, b, hw)$ which corresponds to the probability $P_{\text{step}}(1, b)$ knowing that the input word hamming weight is hw . This can be modeled as follows: choose two random k -bit words x and y both with Hamming weight hw (they represent $\sigma_0(C||U)$ and $\sigma_1(C||U)$) and compute $z = x \oplus y \oplus 1$ (the 1 represents the single bit difference in C). Then $P_{\text{step}}(1, b, hw)$ is the probability that $\text{HAM}(z) = b$ (note that $\text{HAM}(z)$ is always odd thus we have $P_{\text{step}}(1, 2i, hw) = 0$ for all i) and we have:

$$P_{\text{step}}(1, b, hw) = \frac{hw}{c} \cdot P_{\text{XOR}}(k, hw, hw - 1, b) + \frac{c - hw}{c} \cdot P_{\text{XOR}}(k, hw, hw + 1, b).$$

The complexity for handling Q_X is finally

$$\text{Comp} = \frac{1}{\sum_{i=1}^{i=m} P_{\text{step}}(1, i, hw) \cdot P_{\text{out}}(i)}.$$

5.3 Complexity Results

The number C of candidate values we can generate with Hamming weight hw is $\binom{c}{hw}$ and in order to have a good chance to find a collision after Q_X with this amount, we need to ensure that

$$\binom{c}{hw} \geq 1 / \sum_{i=1}^{i=m} P_{\text{step}}(1, i, hw) \cdot P_{\text{out}}(i).$$

One can check that in order to minimize the complexity Comp , the dominant factor of the sum is when i is small. Then, for i small, $P_{\text{step}}(1, i, hw)$ is higher when hw is close to 0 or close to k , in other words the input should have very low or very high Hamming weight. Since we previously chose $U = 00\dots 00$ our goal is to find for each ARMADILLO2 versions the smallest hw value hw_{min} that

ensures enough C candidate values to handle the collision probability in Q_X (but the same reasoning is possible with $U = 11..11$ and the biggest hw value hw_{max}). Overall, the full attack runs in $2 + Comp$ operations (i.e. compression function calls) and negligible memory in order to find a free-start collision for the ARMADILLO2 compression function. We depict in Table 1 our results relative to all proposed versions of ARMADILLO2. This attack has been implemented and verified in practice for $k = 128$ and we give free-start collision examples in the Appendix.

Table 1. Summary of results for free-start collision attack on the different size variants of the ARMADILLO2 compression function. The number of C candidates must always be enough so as to handle the collision probability in Q_X .

scheme parameters				attack parameters			
k	c	m	generic complexity	hw_{min}	nber of C candidates	collision prob. in Q_X	attack complexity
128	80	48	2^{40}	1	$2^{6.3}$	$2^{-4.1}$	$2^{7.5}$
192	128	64	2^{64}	1	2^7	$2^{-4.6}$	$2^{7.8}$
240	160	80	2^{80}	1	$2^{7.3}$	$2^{-4.7}$	$2^{8.1}$
288	192	96	2^{96}	1	$2^{7.6}$	$2^{-4.7}$	$2^{8.3}$
384	256	128	2^{128}	1	2^8	$2^{-4.8}$	$2^{8.7}$

6 Local Linearization: Practical Semi-free-Start Collision Attack

In this section, we show how one can obtain a semi-free-start collision attack (no difference on the input chaining variable) with a very low computational complexity for the ARMADILLO2 compression function.

6.1 General Description

The previous method only allows to add differences on the capacity part of the input, thus leading to free-start collision attacks. One can directly extend this technique to allow only differences in the message part of the input, but this only leads to semi-free-start collisions for randomly chosen bit permutations σ_0 and σ_1 with a not-so-high probability of success.

We would like to derive a semi-free-start collision attack that will output a result with very high probability. In order to achieve this goal we propose a new technique for data-dependent bit transposition ciphers, so-called **local linearization**: by guessing some part of the input we are able to render a few rounds of the internal permutation linear. Indeed, by knowing the g first bits of

U we completely determine the permutations applied during the first g rounds of Q_U . Therefore, for those g rounds the primitive Q_U only consists of known bit permutations and known constant additions. With this method we neutralize for the first g rounds the only non-linearity source: the fact that we don't know which bit permutation σ_0 or σ_1 is applied each round.

On a high-level view, our semi-free-start collision attack will force a collision on the X value at the output of Q_U thanks to the local linearization technique. This collision on X will ensure that the Q_X permutation will be the same for both inputs. Therefore, the difference Hamming weight on the input of Q_X will remain the same in the output. We then hope that those bit differences will be mapped in the truncated part of the output in order to eventually obtain the semi-free-start collision (no difference is feed-forwarded from X since we forced a collision on it). The attack is depicted in Figure 3.




Fig. 3. A schematic view of the semi-free-start collision attack on ARMADILLO2. The thick line at the side of a register represents the least significant bit and black circles stand for bit differences. The dashed box indicates the linearized part.

During a first phase, the input will be divided into two parts: the fixed and the unfixed part. The fixed part $z \in \{0, 1\}^g$ is composed of the g first bits of U and we choose random values for those g bits (so as to know the g first choices of σ_0 or σ_1). The unfixed part $w \in \{0, 1\}^{k-g}$ is composed of the rest of the input bits and

we will be set to a value later. We force the input difference to be contained in the fixed part and we denote it $\Delta z \in \{0, 1\}^g$ (since we are looking for semi-free-start collisions we obviously have $g \leq m$, otherwise we would have a difference in the input chaining variable C). Let $I_1 = (C_1||U_1)$ (resp. $I_2 = (C_2||U_2)$) be the k -bit value of the first input (resp. second output), we have:

$$I_1 = (x||z) \text{ and } I_2 = (x||z \oplus \Delta z).$$

and our goal is to have the collision $X = Q_{U_1}(I_1) = Q_{U_2}(I_2)$.

Assume for the moment that this collision on X happens. Then the same permutation Q_X will be used for both inputs I_1 and I_2 on the right side of Figure 1. As a consequence, no additional bit difference will be introduced during the computation of Q_X , but the bit difference positions will be randomly moved. In order to obtain a semi-free-start collision on the output of the function, we need the $b = \text{HAM}(\Delta z)$ active bits of the input to be mapped in the truncated part of the output through Q_X . As already explained in Section 5, this happens with probability

$$P_{out}(b) = P_{\text{and}}(k, m, b, b) = \prod_{i=0}^{i=b-1} \frac{m-i}{k-i}.$$

6.2 Colliding on X

We need now to evaluate the probability of getting a collision on X . Note that for any round, if there is no difference on the bit choosing the permutation to apply σ_0 or σ_1 , the bit differences at the input of this round will only have their position changed and cannot be erased. Therefore, if we want to obtain a collision on X , we need to obtain it at latest just after the last round of Q_U for which a difference is inserted on the side (in U). We consider from now on that the input difference Δz contains at least one active bit on its MSB, thus this last round is the g -th one.

We know the value of the g first bit of U , therefore we know exactly the permutation applied to I_1 and I_2 for the g first rounds of Q_U . For a collision after g rounds of Q_U , we want that

$$\begin{aligned} & \sigma_{U_1[g-1]}(\cdots(\sigma_{U_1[1]}(\sigma_{U_1[0]}(I_1) \oplus cst) \oplus cst) \cdots) \\ &= \sigma_{U_2[g-1]}(\cdots(\sigma_{U_2[1]}(\sigma_{U_2[0]}(I_2) \oplus cst) \oplus cst) \cdots) \end{aligned}$$

and since **all operations are linear**, this can be rewritten as

$$\rho(I_1) \oplus A = \rho'(I_2) \oplus B = \rho'(I_1 \oplus \Delta z) \oplus B = \rho'(I_1) \oplus \rho'(\Delta z) \oplus B$$

where

$$\begin{aligned} \rho &= \sigma_{U_1[g-1]} \circ \cdots \circ \sigma_{U_1[1]} \circ \sigma_{U_1[0]} & A &= \sigma_{U_1[g-1]}(\cdots(\sigma_{U_1[1]}(cst) \oplus cst) \cdots) \\ \rho' &= \sigma_{U_2[g-1]} \circ \cdots \circ \sigma_{U_2[1]} \circ \sigma_{U_2[0]} & B &= \sigma_{U_2[g-1]}(\cdots(\sigma_{U_2[1]}(cst) \oplus cst) \cdots). \end{aligned}$$

Finally, we end up with the equation

$$\rho(I_1) \oplus \rho'(I_1) = A \oplus B \oplus \rho'(\Delta z) \quad (1)$$

Since we know the value of the g first bit of U , we can compute the value of A and B . Moreover, assuming that we already chose a Δz , then the collision condition (1) can be rephrased as

$$I_1 \oplus \tau(I_1) = C$$

where $C = \rho^{-1}(A \oplus B \oplus \rho'(\Delta z))$ and $\tau = \rho^{-1} \circ \rho'$.

In order to study this system \mathcal{S} of k bit equations, we model τ as a random bit permutation and C as a random k -bit word. Note that since this equation system is linear finding the potential solutions requires only a few operations, but we would like to know how many such systems we need to generate before finding a solution, i.e. a collision on X . Thus, our goal is now to deduce the probability that this system has at least one solution and what is the average number of expected solutions.

The structure of this equation system is very particular and the number of independent groups of bit equations is exactly the number of cycles of the bit permutation τ . More precisely, let $\text{CYCLE}(\tau)$ represent the number of cycles of the permutation τ and let S_i denote the set of bits belonging to the i -th cycle of τ .

Theorem 1. *The equation system $\mathcal{S} : I_1 \oplus \tau(I_1) = C$ admits a solution if and only if for every cycle set S_i of τ the parity of the sum of the corresponding C bit is null, that is*

$$\bigoplus_{p \in S_i} C[p] = 0.$$

If this system is solvable, then the number of solutions that can be generated is exactly equal to $2^{\text{CYCLE}(\tau)}$.

The Idea of the Theorem is that when we want to find a solution for the system, we can start by fixing one bit a_0 to a random value. This bit is involved into two binary equations from \mathcal{S} . All equations having only two terms, one of the two equations directly links bit a_0 with say bit a_1 , and we can deduce the value of a_1 . The bit a_1 is in turn linked with bit a_2 through his second equation and we directly deduce the value of a_2 . This chain of dependency will eventually cycle (the new bit deduced will be a_0 again) and will be validated if and only if the sum of the C bits of the equations visited is null (otherwise we encounter a inconsistency). This check is then performed for all cycles.

Proof. Since τ is a bit permutation, the equation system \mathcal{S} can be represented as a collection of cycles, each cycle depicting the direct cyclical dependencies between some set of bits: if bit x and bit y are linked by one of the k equations,

then they belong to the same cycle. The vertex weight between two members x and y of the cycle is the value $C[x]$.

If we fix the bit value of a member of a cycle S_i , then this determines entirely all the other bits of that cycle (according to the vertices values). Then, if the XOR of all the vertex weights is different from zero, we have a direct contradiction. A solution can only exist if all cycles present no internal contradiction.

Each cycle can have either zero or two solutions (the two solutions being their mutual complement). If every cycle has no contradiction, then there exists exactly $2^{\text{CYCLE}(\tau)}$ distinct combinations of cycle solutions, each one leading to a distinct solution for the whole equation system \mathcal{S} . \square

From Theorem 1, we directly deduce that the probability that the system admits a solution is equal to $2^{-\text{CYCLE}(\tau)}$. The expected number of cycles for a randomly chosen permutation on k elements is $\log(k)$. Therefore, we have to try at least $2^{\log(k)}$ different equation systems before finding one admitting a solution. When one system admits a solution, we directly get $2^{\log(k)}$ solutions for free. Overall, the cost for finding one solution of the system is 1 on average (the average cost is the meaningful one here since we will have to find several inputs colliding on X during the whole attack).

6.3 Complexity Results

We now look for a solution such that the original guess of the g first bits of the input was right (with probability 2^{-g}) and such that the b bit differences in Q_X are mapped to the truncated part of the output (with probability $P_{out}(b)$). Overall, the total complexity of the semi-free-start collision attack is $2^g \cdot P_{out}^{-1}(b)$ with $b \leq g$. Minimizing g and b will minimize the overall complexity, but we need to ensure that we can go through enough equation systems in order to have a good chance to find a collision eventually. More precisely, we need

$$1/2 \cdot 2^g \cdot \binom{g}{b} \geq 2^g \cdot P_{out}^{-1}(b)$$

which can be rewritten as

$$\binom{g}{b} \geq 2 \cdot P_{out}^{-1}(b).$$

We depict in Table 2 our results relative to all proposed versions of ARMADILL02. This attack has been implemented and verified in practice for $k = 128$ and we give semi-free-start collision examples in the Appendix.

7 Related-Key Recovery in Stream Cipher Mode

In this section we will present a related key attack that will allow us to recover all key bits in practical time when using ARMADILL02 in the stream cipher mode. We will first present the main idea of this attack, and afterwards, we will give a more detailed analysis of the probabilities and complexities.

Table 2. Summary of results for semi-free-start collision attack on the different size variants of the ARMADILLO2 compression function

scheme parameters				attack parameters			
k	c	m	generic complexity	g	b	$P_{out}(b)$	time complexity
128	80	48	2^{40}	6	2	$2^{-2.9}$	$2^{8.9}$
192	128	64	2^{64}	7	2	$2^{-3.2}$	$2^{10.2}$
240	160	80	2^{80}	7	2	$2^{-3.2}$	$2^{10.2}$
288	192	96	2^{96}	7	2	$2^{-3.2}$	$2^{10.2}$
384	256	128	2^{128}	7	2	$2^{-3.2}$	$2^{10.2}$

7.1 Using Related-Keys for Recovering the Key

First of all, we consider a pair of related keys (K_1, K_2) that have one only bit of difference, that is $\text{HAM}(K_1 \oplus K_2) = \text{HAM}(\Delta_K) = 1$. Our analysis will work for any bit difference position d amongst all the bits of the key. Note that we expect a pair of keys valid for performing the related-key attack to appear after using about $(2^k/k)^{1/2}$ keys.

Let us consider a value of U for generating k bits of key-stream with each of both keys K_1 and K_2 . We use the index i for the intermediate states generated from the key K_i . We first make the following observations, important in order to understand the whole attack procedure:

- Since no difference is inserted in the U part (it is a public value) and since $\text{HAM}(\Delta_K) = 1$, we have $\text{HAM}(X_1 \oplus X_2) = 1$. Let e be the bit position of this difference in X .
- The first $(e - 1)$ intermediate states of Q_X will also have a difference of Hamming weight 1.

We assume that the attacker can choose the values of U . In this case, we can make the bit difference in the key to go from position d to any wanted position e in X through Q_U . We expect $2^m/k$ distinct values of U that make the bit difference go from position d to e for $e \in [0, k - 1]$. We denote by U_e each one of these k subgroups of U values.

The output of the function $(V_c, V_i) = X \oplus Y$ is known to the attacker, but concerning X he only knows the m bits of the U part (since U is known, he can deduce directly where the bits coming from U and C will be eventually located in X). Thus, he can recover m bits from the outputs of Q_X , Y_1 and Y_2 . If he could compute backward from Y_1 and Y_2 until the beginning of the e -th step of Q_X , the colliding positions of the bits known from Y_1 and from Y_2 will have the same values with maybe the exception of one, which would be the original single bit difference before the step e .

Our attack basically consists in choosing several values for U from U_e , for decreasing e values (starting from $e = k - 1$), that will gradually increase the number of key bits appearing in X after position e . Each time we will guess the value of the new key bits appearing and discard the guesses that will not lead to collisions on the bit values in the colliding positions just before step e when computing backward from Y_1 and Y_2 in Q_X . The complexity of this attack depends on the bit permutations σ_0 and σ_1 , but in the next subsection we give a complexity analysis assuming that these permutations are randomly chosen.

7.2 Generic Complexity Estimation

We start at $e = k - 1$. First, we choose the value of i (denoted i_{max}), that maximizes the probability $P_{\text{and}}(k, m, m, i)$ that we denote p_{max} . For instance, if we consider the smallest version of ARMADILLO2, where $k = 128$, $c = 80$ and $m = 48$, then we have $i_{max} = 18$ and the probability of obtaining 18 positions of known bits that collide is equal to $p_{max} = 2^{-2.72}$.

Amongst the values from U_{k-1} , we choose p_{max}^{-1} random ones. Each of them is introduced in the ARMADILLO2 function parametrized with the keys K_1 and K_2 . For each of the p_{max}^{-1} pairs of values, we guess the bit at position $k - 1$ of X_1 and of X_2 (for example 1 and 0 respectively since there is a difference on this bit position) and we end up with $2 \cdot p_{max}^{-1}$ pairs. Then, we can undo the last round of Q_X for the known bits from Y_1 and Y_2 . We consider that a guess passes the test if it verifies the conditions on the number of colliding values on the colliding bit positions. For one of these $2 \cdot p_{max}^{-1}$ pairs (in our example $(Q_1^{-1}(Y_1), Q_0^{-1}(Y_2))$), the number of colliding bit positions will be i_{max} . When this is the case, if the guess on the bit of X_1 and X_2 was incorrect, we have a probability of $2^{-i_{max}+1}$ to pass the test, while we will pass it with probability one if the guess was correct. Finally, we have determined one bit of each key K_1 and K_2 with a complexity of $2 \cdot p_{max}^{-1}$, which in our example would be $2^3.72$.

We can continue the process by considering $e = k - 2$ and p_{max}^{-1} values from U_{k-2} that have a key bit at position $k - 1$. Following the same method as before, we will recover one key bit, i.e. the one at position $k - 1$ in X when we have 18 colliding bits before the step $k - 1$ of Q_X . Let us remark here that in practice we do not have to wait for having a collision on 18 bits, but most of the time collisions on a different number of bits will also be enough for determining if a guess passes the test or not. We can repeat this step in order to obtain the biggest possible number of key bits and determining each bit will add at most a complexity of p_{max}^{-1} .

The next steps depend on the number of bits that we have already determined. All in all, we conjecture that when both bit permutations behave like random ones, the complexity will not exceed $2 \cdot c \cdot p_{max}^{-1}$.

8 Conclusion

We have presented some new and practical analysis of ARMADILLO2. Notably a free-start and semi-free-start collision attacks for the full ARMADILLO2 hash

functions. Extending this work to real collisions (i.e. with a predefined IV) might be possible but it is not very appealing because it is likely that several message blocks are required (all versions have $c > m$) and therefore the task of the cryptanalyst would be quite complex to handle. ARMADILLO2 should not be used in any security application since our attacks have a very low complexity. This work and the local-linearization method is a first step in order to evaluate the security of data-dependent bit transpositions cryptographic designs.

Acknowledgements. The authors would like to thank the anonymous referees and the ARMADILLO2 team for their helpful comments.

References

1. Abdelraheem, M.A., Blondeau, C., Naya-Plasencia, M., Videau, M., Zenner, E.: Cryptanalysis of ARMADILLO2. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 308–326. Springer, Heidelberg (2011)
2. Badel, S., Dağtekin, N., Nakahara Jr., J., Ouafi, K., Reffé, N., Sepehrdad, P., Sušil, P., Vaudenay, S.: ARMADILLO: A Multi-purpose Cryptographic Primitive Dedicated to Hardware. In: Mangard, S., Standaert, F.-X. (eds.) CHES 2010. LNCS, vol. 6225, pp. 398–412. Springer, Heidelberg (2010)
3. Brassard, G. (ed.): CRYPTO 1989. LNCS, vol. 435. Springer, Heidelberg (1990)
4. Damgård, I.: A Design Principle for Hash Functions. In: Brassard [3], pp. 416–427
5. Lai, X., Massey, J.L.: A Proposal for a New Block Encryption Standard. In: Damgård, I.B. (ed.) EUROCRYPT 1990. LNCS, vol. 473, pp. 389–404. Springer, Heidelberg (1991)
6. Merkle, R.C.: One Way Hash Functions and DES. In: Brassard [3], pp. 428–446
7. Rivest, R.L.: The RC5 Encryption Algorithm, pp. 86–96. Springer (1995)
8. Rivest, R.L., Robshaw, M.J.B., Yin, Y.L.: RC6 as the AES (2000)
9. Sepehrdad, P., Sušil, P., Vaudenay, S.: Fast Key Recovery Attack on ARMADILLO1 and Variants. In: Prouff, E. (ed.) CARDIS 2011. LNCS, vol. 7079, pp. 133–150. Springer, Heidelberg (2011)

A Implementation of the Collision Attacks for $k = 128$

We implemented all attacks for $k = 128$ and they require less than a second and negligible memory on an average computer (Intel Core2 Duo CPU @ 2.13 GHz) in order to find a collision. Since no specific σ_0 and σ_1 bit transpositions are defined for ARMADILLO2, we run the attack for many randomly chosen instances so as to ensure the soundness of our reasoning. We give here examples of (semi)-free-start collisions for ARMADILLO2 with a σ_0 and σ_1 bit transpositions instance that fulfill the criteria required in [2] for $k = 128$. Namely, we denote λ the second largest eigenvalue of the matrix $M = \frac{1}{4}(P_{\sigma_0} + P_{\sigma_0}^{128} + P_{\sigma_1} + P_{\sigma_1}^{128})$, then for the σ_0 and σ_1 instance found we have $\lambda = 0.87$. This means that there exists a distinguisher with advantage $\lambda^{256} = 2^{-51.4}$, while our attacks have much better advantage.

Free-Start Collision for ARMADILLO2 with $k = 128$, $c = 80$, $m = 48$:

```
ARMADILLO2(ffffffffffffffffbfff, ffffffff) =
ARMADILLO2(ffffdffffffffffbfff, ffffffff) =
dfb0d8f2b763ce97f785
```

Semi-Free-Start Collision for ARMADILLO2 with $k = 128$, $c = 80$, $m = 48$:

```
ARMADILLO2(6bc8c848de5ff533cd6f, 0850b04b82e2) =
ARMADILLO2(6bc8c848de5ff533cd6f, 0850b04b82f0) =
26827e3d614d2fc75d64
```

Bit Transpositions σ_0 and σ_1 Used:

$\sigma_0 = 62, 98, 14, 114, 36, 77, 55, 3, 28, 88, 29, 122, 57, 90, 66, 52, 44, 22, 95, 118, 69, 86,$
 $35, 56, 58, 82, 18, 97, 78, 21, 85, 101, 19, 65, 10, 6, 116, 121, 70, 99, 61, 102, 4, 91,$
 $39, 119, 79, 16, 84, 50, 113, 45, 93, 104, 73, 112, 8, 5, 51, 9, 105, 46, 64, 94, 41, 54,$
 $127, 67, 106, 23, 63, 49, 123, 15, 60, 81, 96, 72, 110, 37, 30, 89, 7, 92, 2, 68, 40, 32,$
 $53, 11, 71, 26, 103, 59, 109, 111, 38, 74, 20, 48, 24, 43, 126, 117, 13, 124, 31, 33,$
 $100, 125, 87, 27, 83, 128, 12, 42, 80, 107, 108, 17, 25, 120, 76, 75, 115, 47, 1, 34$

$\sigma_1 = 10, 60, 111, 78, 38, 57, 110, 75, 104, 56, 88, 79, 23, 99, 16, 22, 128, 94, 120, 24, 64,$
 $3, 6, 55, 42, 51, 43, 82, 114, 89, 26, 35, 61, 73, 77, 36, 28, 21, 105, 15, 67, 70, 113,$
 $65, 39, 80, 122, 31, 101, 100, 107, 124, 18, 46, 85, 19, 49, 14, 12, 71, 86, 68, 102, 91,$
 $58, 95, 1, 53, 83, 125, 66, 98, 81, 44, 48, 59, 27, 9, 119, 40, 45, 74, 92, 112, 93,$
 $69, 5, 108, 106, 115, 90, 13, 84, 126, 7, 109, 54, 127, 33, 121, 62, 87, 30, 29, 63, 2,$
 $97, 116, 4, 47, 11, 8, 34, 96, 118, 72, 52, 103, 37, 25, 123, 50, 76, 17, 20, 41, 117, 32$

On the (In)Security of IDEA in Various Hashing Modes^{*}

Lei Wei¹, Thomas Peyrin¹, Przemysław Sokołowski²,
San Ling¹, Josef Pieprzyk², and Huaxiong Wang¹

¹ Division of Mathematical Sciences, School of Physical and Mathematical Sciences,
Nanyang Technological University, Singapore
wl@pmail.ntu.edu.sg, thomas.peyrin@gmail.com

² Macquarie University, Australia

Abstract. In this article, we study the security of the IDEA block cipher when it is used in various simple-length or double-length hashing modes. Even though this cipher is still considered as secure, we show that one should avoid its use as internal primitive for block cipher based hashing. In particular, we are able to generate instantaneously free-start collisions for most modes, and even semi-free-start collisions, pseudo-preimages or hash collisions in practical complexity. This work shows a practical example of the gap that exists between secret-key and known or chosen-key security for block ciphers. Moreover, we also settle the 20-year-old standing open question concerning the security of the Abreast-DM and Tandem-DM double-length compression functions, originally invented to be instantiated with IDEA. Our attacks have been verified experimentally and work even for strengthened versions of IDEA with any number of rounds.

Keywords: IDEA, block cipher, hash function, cryptanalysis, collision, preimage.

1 Introduction

Hash functions are considered as a very important building block for many security and cryptography applications. Informally, a hash function H is a function that takes an arbitrarily long message as input and outputs a fixed-length hash value of size n bits. In cryptography, we want these functions to fulfill three security requirements, namely collision resistance and (second)-preimage resistance. It should be impossible for an adversary to find a collision (two different messages that lead to the same hash value) in less than $2^{n/2}$ hash computations, or a

^{*} The first, fourth and sixth authors are supported by the Singapore National Research Foundation under Research Grant NRF-CRP2-2007-03 and the first author is also supported by the Singapore Ministry of Education under Research Grant T206B2204 and by the NTU NAP Startup Grant M58110000. The second author is supported by the Lee Kuan Yew Postdoctoral Fellowship 2011 and the Singapore National Research Foundation Fellowship 2012.

(second)-preimage (a message hashing to a given challenge) in less than 2^n hash computations. Most of nowadays hash functions divide the whole input message into blocks after padding it, and then process the blocks in an iterative way. A very known and utilised example is the Merkle-Damgård algorithm [12,33], which uses an n -bit compression function h in order to process the m message blocks M_i : $CV_{i+1} = h(CV_i, M_i)$, where CV_i is the n -bit internal state (or chaining variable) that is initialized by a fixed public value $CV_0 = IV$ and the final hash value is H_m . This algorithm is very interesting because it allows to reduce the collision/preimage security of the hash function to the collision/preimage security of the compression function. However, in order to guarantee the soundness of the construction, a designer must ensure that an attacker can not break the collision/preimage resistance of the compression function. One can identify different security properties for a compression function:

- *free-start collision*: in less than $2^{n/2}$ computations, find two different pairs $(CV, M) \neq (CV', M')$ such that they lead to the same compression function output value: $h(CV, M) = h(CV', M')$,
- *semi-free-start collision*: in less than $2^{n/2}$ computations, find one chaining variable CV and two different message blocks $M \neq M'$ such that they lead to the same compression function output value: $h(CV, M) = h(CV, M')$,
- *preimage*: in less than 2^n computations, find one chaining variable CV and one message block M such that they lead to a given output challenge X : $h(CV, M) = X$.

Note that a semi-free-start collision for the compression function where the chaining variable CV is not chosen by the attacker directly leads to a collision for the whole hash function. In any case, a semi-free-start collision is very dangerous since it means that for some choices of IV , the attacker knows how to generate a collision. Even free-start collision are considered serious as they invalidate the collision resistance assumption on the compression function and we have seen many free-start collision attacks eventually turning into full hash collision attacks in the recent history (for example free-start collision attacks for MD5 were quickly identified [14], then upgraded to semi-free-start collision attacks [15] and eventually to full collision attacks [38]). As for preimage attacks on the compression function (also known as pseudo-preimages), they are very relevant since there exist a meet-in-the-middle algorithm that in most cases can turn them into a preimage attack for the full hash function.

The separation between a block cipher and a compression function has always been blurry. Constructions are known to turn the former into the latter [7,36] or the latter into the former [31]. For example, the Davies-Meyer mode [1] converts a secure block cipher E into a secure compression function and is incorporated in a large majority of the currently known hash functions. While very satisfying solutions exist to transform a secure n -bit block cipher into an n -bit compression function (Davies-Meyer, Miyaguchi-Preneel, Matyas-Meyer-Oseas modes [1] or see [7,36] for a systematic study of this problem), there is still a lot of research being actively conducted on double-block length compression functions (where

the block cipher size is n bits and the compression function output size is $2n$), from simple-key block ciphers such as AES-128 or double-key such as AES-256 [11].

A major difference between the cryptanalysis of block ciphers and compression functions is that the attacker can fully control the inner behavior of the compression function. In other words, the attacker can use more efficiently the freedom degrees available on the input (i.e. the number of independent binary variables he has to determine). A new security model for block ciphers, the so-called *known-key model* [24], was recently proposed in order to fill the gap between these two situations. In this model, the secret key is known to the adversary and its goal is to distinguish the behavior of a random instance of the block cipher from the one of a random permutation by constructing a set of (plaintext, ciphertext) pairs satisfying an *evasive* property. Such a property is easy to check but impossible to achieve with the same complexity and a non-negligible probability using oracle accesses to a random permutation and its inverse. In general, these known-key attacks are not regarded as problematic when the block cipher is used in a classical “secret key” setting. Moreover, it is rare that such threats are extended to attacks on the compression function.

A potential candidate for hashing is the 64-bit block cipher IDEA [26,39] that uses 128-bit keys. While a simple-length hashing mode would only provide a 64-bit hash output, insufficient for most of nowadays security applications, a double-block length construction (DBL) would allow 128-bit hash outputs which can be sufficient in some scenarios. As IDEA handles double-length keys, more freedom in the constructions is possible. In fact, the well known Abreast-DM and Tandem-DM modes were specifically created to perform hashing with IDEA (see page 2 and Section 6 of [39]). These modes were later studied in much details [16,17,28,30], but the security they provide when instantiated with IDEA remains a 20-year-old standing open question. In classical “secret key” setting, IDEA has already been studied a lot [2,3,4,5,6,9,10,13,18] and is still considered as a secure cipher despite its age and despite the current best attack [5] that requires 2^{63} data (half the codebook) and 2^{114} computations to recover the secret key for IDEA reduced to 7.5 rounds over a total of 8.5 (the attack on the full cipher from [5] is very marginal with $2^{126.8}$ computations and the one from [22] requires 2^{126} computations and 2^{52} chosen plaintexts). One can also cite the work of [6], that exposes a weak-key class of size 2^{64} . Note also that a first step towards analysis of IDEA in hashing mode was done in [21] where a 3-round chosen-key attack is described and in [9] where the authors show how to find a free-start near collision (only a subset of the output collides) when IDEA is plugged into the Hirose DBL mode [9] (and also a free-start collision if the internal constant c is controlled by the attacker).

Our Contribution. In this paper, we study the security of the IDEA block cipher [26,39] when plugged into various block cipher based compression function constructions, such as the classical Davies-Meyer mode [1], also DBL constructions such as Hirose [19,20], Abreast-DM and Tandem-DM [27,39], Peyrin *et al.* (II) [35] or MJH-Double [29]. Even if this cipher is still considered as secure in the classical “secret key” setting, its security remains an open problem in hashing

mode. Depending on the IDEA-based hash construction, we show that an attacker can find free-start collisions instantaneously, preimages or semi-free-start collisions practically. For some modes, we even describe a method to compute collisions for the whole hash function. These attacks are based on weak-keys utilisation, but in contrary to the “secret key” setting where the goal of the attacker is to exhibit the biggest weak-key class possible, in hashing mode the goal is to find and exploit the weakest of all keys. We use the fact that the key 0 in IDEA is extremely weak, actually rendering the whole encryption process a T-function [23], already known as dangerous for building a hash function [34]. While weak-keys are already known to be dangerous for block cipher-based hash functions, our method use a novel and non-trivial almost half-involution property for IDEA. Even strengthened versions of the cipher with any number of rounds can be attacked with about the same complexities. This work is one more example that one has to be very careful when hashing with a block cipher that presents any weakness when the key is known or controlled by the attacker. In particular, one should strictly avoid the use of a block cipher for which weak-keys exist, even if only a single weak-key is known.

2 The IDEA Block Cipher

The International Data Encryption Algorithm (IDEA) is a 64-bit block cipher handling 128-bit keys and designed by Lai and Massey [26,39] in 1990. While its use is reducing over the recent years, it remains deployed in practice and has not been broken yet despite its advanced age. It has a very simple design, performing 8.5 rounds composed of only 16-bit wide XOR, additions and multiplications. More precisely, one round is composed of three layers: first the key addition layer (denoted KA), a multiplication-addition layer (denoted MA) and a middle words switching layer (denoted S). For the eighth round, the switching is omitted.

Let X^i represent the 64-bit internal state of IDEA before application of the i -th round and we can view it as four 16-bit subwords $X^i = (X_1^i, X_2^i, X_3^i, X_4^i)$, with $1 \leq i \leq 9$. Also, $Y^i = (Y_1^i, Y_2^i, Y_3^i, Y_4^i)$ will stand for the intermediate internal state value of IDEA during the i -th round, right between the KA and the MA layers. We denote by \oplus the bitwise XOR operation, by \boxplus the addition modulo 2^{16} and by \odot the multiplication modulo $2^{16} + 1$, where the value 0 is considered as 2^{16} and vice-versa. Finally, $Z^i = (Z_1^i, Z_2^i, Z_3^i, Z_4^i, Z_5^i, Z_6^i)$ represents the six 16-bit subkeys used during the i -th round (only the first four subkeys for the last half round).

The KA layer simply incorporates four subkeys:

$$Y_1^i = X_1^i \odot Z_1^i, \quad Y_2^i = X_2^i \boxplus Z_2^i, \quad Y_3^i = X_3^i \boxplus Z_3^i, \quad Y_4^i = X_4^i \odot Z_4^i.$$

The MA layer first computes $B = Z_6^i \odot ((Y_2^i \oplus Y_4^i) \boxplus (Z_5^i \odot (Y_1^i \oplus Y_3^i)))$ and $A = B \boxplus (Z_5^i \odot (Y_1^i \oplus Y_3^i))$. Then, after application of the S layer we have:

$$X_1^{i+1} = Y_1^i \oplus B, \quad X_2^{i+1} = Y_3^i \oplus B, \quad X_3^{i+1} = Y_2^i \oplus A, \quad X_4^{i+1} = Y_4^i \oplus A.$$

All the subkeys are simply determined by choosing consecutive bits in the 128-bit master key according to some selection table (we refer to the IDEA specifications). Finally, ciphering the plaintext P with IDEA to obtain the ciphertext C is defined as: $C = \text{KA} \circ \text{S} \circ \{\text{S} \circ \text{MA} \circ \text{KA}\}^8(P)$.

Currently, the best cryptanalysis work published on IDEA [5] can reach 7.5 rounds with 2^{63} data (half the codebook) and 2^{114} computations. Concerning weak-keys, the current biggest weak-key class contains 2^{64} elements and has been published in [6].

3 Hashing with a Double-Length Key Block Cipher

We will study the security of the various block cipher-based constructions that can use IDEA as the internal primitive. Therefore, we only consider the ones that use a double-key block cipher. More precisely, we denote $C = E_K(P)$ the process of ciphering the 64-bit plaintext P with IDEA using the 128-bit key K .

3.1 Simple-Length Compression Function

A simple-length compression function construction with IDEA will provide a 64-bit output CV_{i+1} .

Davies-Meyer is the most usual simple-length mode [1] and it handles 128-bit message blocks: $CV_{i+1} = E_M(CV_i) \oplus CV_i$. Most standardized hash functions are actually implementing this mode, with an ad-hoc internal block cipher. While some weaknesses such as fixed-points are known, its security in terms of preimage and collision resistance have been studied and proved in the ideal cipher model [7]. Namely, we should expect at least 2^{32} and 2^{64} computations respectively to generate a (semi)-free-start collision or preimage for the compression function. Note that Miyaguchi-Preneel and Matyas-Meyer-Oseas simple-block length modes [1] are not considered in this article since they require the internal primitive to have the same block and key size, which is not the case for IDEA.

3.2 Double-Length Compression Function

A more interesting design strategy with IDEA would be to define double-block length constructions, in order to get 128-bit output, represented by two 64-bit words $CV1_i$ and $CV2_i$. This problem has already been studied a lot and remains a very active research domain, even when the internal primitive is a double-key block cipher.

Abreast-DM and Tandem-DM will of course be considered in this article since they both have been especially designed for IDEA [27,39]. Tandem-DM handles a 64-bit message block M .

We define $W = E_{CV1_i||M}(CV2_i)$ and then we have

$$\begin{aligned} CV1_{i+1} &= E_{M||W}(CV1_i) \oplus CV1_i, \\ CV2_{i+1} &= W \oplus CV2_i. \end{aligned}$$

Abreast-DM also handles a 64-bit message block M :

$$\begin{aligned} CV1_{i+1} &= E_{M||CV2_i}(\overline{CV1_i}) \oplus CV1_i, \\ CV2_{i+1} &= E_{CV1_i||M}(CV2_i) \oplus CV2_i, \end{aligned}$$

where \overline{X} stands for the bitwise complement of X .

Hirose proposed a construction that contains two independent block cipher instances [19], later improved to only a single instance [20] by using a constant c to simulate the two independent ciphers:

$$\begin{aligned} CV1_{i+1} &= E_{CV2_i||M}(CV1_i) \oplus CV1_i, \\ CV2_{i+1} &= E_{CV2_i||M}(CV1_i \oplus c) \oplus CV1_i \oplus c. \end{aligned}$$

Peyrin et al. described in [35] a compression function (denoted *Peyrin et al.*(II)) that utilizes 5 calls to independent $3n$ -to- n -bit compression functions, advising to be instantiated with double-key internal block ciphers such as AES-256 or IDEA. It handles two 64-bit message blocks $M1$ and $M2$:

$$\begin{aligned} CV1_{i+1} &= f_1(CV1_i, CV2_i, M1) \oplus f_2(CV1_i, CV2_i, M2) \oplus f_3(CV1_i, M1, M2), \\ CV2_{i+1} &= f_3(CV1_i, M1, M2) \oplus f_4(CV1_i, CV2_i, M1) \oplus f_5(CV2_i, M1, M2), \end{aligned}$$

where the functions f_i can be build for example by using the IDEA block cipher into a Davies-Meyer mode and we can simulate their independency by XORing distinct constants to the plaintext inputs, as it is done in [20]: $f_i(U, V, W) = E_{U||V}(W \oplus i) \oplus W$ (note that XORing the constants on the key input would be avoided in practice because it would lead to very frequent rekeying and therefore reduce the overall performance of the hash function). Since no real candidate was proposed by the authors, all possible position permutations of the three f_i inputs will be considered. Note that when cryptanalysing this scheme, we will attack the functions f_i independently. Thus, we will not use any weakness coming from potential dependencies between the functions f_i (apart of course that all 5 functions are based on IDEA).

MJH-Double is a rate 1 double-block length compression function recently published by Lee and Stam [29]. It uses a double-key block cipher and handles two 64-bit message blocks $M1$ and $M2$:

$$\begin{aligned} CV1_{i+1} &= E_{M2||CV2_i}(CV1_i \oplus M1) \oplus CV1_i \oplus M1, \\ CV2_{i+1} &= g \cdot (E_{M2||CV2_i}(f(CV1_i \oplus M1)) \oplus f(CV1_i \oplus M1)) \oplus CV1_i, \end{aligned}$$

where f is an involution with no fixed point and $g \neq 0, 1$ is a constant.

For all these double-block length proposals, the conjectured security is 2^{64} and 2^{128} computations respectively to generate a (semi)-free-start collision or preimage for the compression function or hash function.

4 Weak-Keys for IDEA

Weak-keys for IDEA has already been studied in details [6,10,18], but what we are looking for is slightly different. Indeed, for block cipher cryptanalysis, since the attacker can not control the key input he looks for the biggest possible class of weak-keys, so as to get the highest possible probability that a weak-key will indeed be chosen. In the case of compression function cryptanalysis, the key input is fully known or even controlled by the attacker. The goal is therefore not to find the biggest possible class of weak-keys, but to find the weakest possible key. As we will show for IDEA, even if only one weak-key is found, its weakness might directly lead to successful attacks on the whole compression or hash function.

4.1 Analysis of the Internal Functions

When looking at the internal round function of IDEA, one might wonder what would be a weak-key. In IDEA, the most annoying functions for the cryptanalyst are clearly the multiplications in $\mathbb{Z}_{2^{16}+1}$. Indeed, these operations are strongly non-linear and provide good diffusion between the different bit positions. On the contrary, XOR operations are linear and do not provide any diffusion between the bit positions, while the additions in $\mathbb{Z}_{2^{16}}$ can be easily approximated linearly and the diffusion between the bit positions only happens through the carry. Moreover, XOR and additions are even weaker in IDEA since no rotations are present, comparing with Addition-Rotation-XOR (ARX) designs. Here the rotation is done through the multiplications in $\mathbb{Z}_{2^{16}+1}$ and our goal is therefore to avoid them.

When adding $(a + b) \bmod 2^{16}$, we can avoid any diffusion by forcing one operand to 0. When multiplying $(a \odot b) = (a \cdot b) \bmod 2^{16} + 1$, the good diffusion will happen especially when $(a \cdot b) \geq 2^{16} + 1$. An easy way to avoid this is to fix one of the two operands to 1. In that case, we have $(a \odot 1) = (a \cdot 1) \bmod 2^{16} + 1 = a \bmod 2^{16}$. As already remarked in [10], a good choice is also 0, since

$$\begin{aligned} (a \odot 0) \bmod 2^{16} &= ((a \cdot 2^{16}) \bmod (2^{16} + 1)) \bmod 2^{16} \\ &= (((a \cdot 2^{16} + a) + (2^{16} + 1) - a) \bmod (2^{16} + 1)) \bmod 2^{16} \\ &= (0 + 2^{16} + 1 - a) \bmod 2^{16} = 1 - a \bmod 2^{16} \\ &= 2 + (2^{16} - 1 - a) \bmod 2^{16} = (2 + \bar{a}) \bmod 2^{16} \end{aligned}$$

and the multiplication is reduced to only a complement and an addition with a constant.

4.2 Weak-Keys Classes

Based on the remark that the operand 0 is very weak for both multiplications and additions, Daemen *et al.* [10] generated a class of weak-keys. A first obvious candidate is the null key (all bits set to zero), which will force all the subkeys to zero as well. As a consequence, all subkeys additions can be simply removed and all subkeys multiplications can be replaced by a complement (or XOR with 0xffff) and an addition with value 2. At this point, all the operations in IDEA with null key are either XOR or additions. Therefore, by inserting differences only on the Most Significant Bit (MSB) of the four 16-bit plaintext input words, the attacker is ensured that only the MSB of the four output words will contain a difference. Even better, the mapping from an MSB input difference pattern to an MSB output difference pattern is completely deterministic (is it linear since on the MSB no carry is propagated). Such a property is largely sufficient to consider the null key as weak. This reasoning can be generalized by observing that the attacker does not necessarily need all subkeys to be null, but only the ones that are multiplied to an internal word which contains a MSB difference. Since the MSB differential paths are quite sparse, many of the null constraints on the subkeys are relaxed and one finally gets 2^{35} weak-keys.

4.3 The Null Weak-Key

We show that the null key is particularly weak for hash function utilization. Even if other keys belong to a weak-key class, they do not present the same special properties as the null key.

Almost Half-Involution. When using the null key, we remark that all subkeys will be null as well. Then, all rounds layers will be the same and we write KA_0 and MA_0 the KA and MA layers with null subkeys. A nice practical feature of IDEA is that the decryption is done using the very same algorithm as encryption, but with different subkeys. The decryption subkeys for the MA layer are the same as the encryption ones since the MA layer is an involution (i.e. $MA=MA^{-1}$). The decryption subkeys for the KA layer are the respective multiplicative and additive inverses of the encryption subkeys. However, note that a null subkey is both its own multiplicative and additive inverse and the KA layer becomes an involution as well (i.e. $KA_0=KA_0^{-1}$). To summarize, using the null key, we are ensured that $KA_0=KA_0^{-1}$ and $MA_0=MA_0^{-1}$. Note that we trivially have $S=S^{-1}$.

Now, since the KA layer and S layer commute, IDEA with null key can be rewritten as

$$\begin{aligned}
 C &= KA_0 \circ S \circ \{S \circ MA_0 \circ KA_0\}^8(P) \\
 &= KA_0 \circ S \circ \{S \circ MA_0 \circ KA_0\}^3 \circ S \circ MA_0 \circ KA_0 \circ \{S \circ MA_0 \circ KA_0\}^4(P) \\
 &= \underbrace{KA_0 \circ MA_0 \circ \{S \circ KA_0 \circ MA_0\}^3}_{\sigma^{-1}} \circ \underbrace{KA_0 \circ S}_{\theta} \circ \underbrace{\{MA_0 \circ KA_0 \circ S\}^3 \circ MA_0 \circ KA_0}_{\sigma}(P)
 \end{aligned}$$

which eventually gives $C = \sigma^{-1} \circ \theta \circ \sigma(P)$. One can check that since KA_0 , MA_0 and S are involutions, the operation denoted by σ^{-1} is indeed the inverse of the one denoted by σ . Thus, using the notation

$$P \xrightarrow{\sigma^{-1}} U \xrightarrow{\theta} V \xrightarrow{\sigma} C$$

where U and V are internal state values, we have

$$P \xleftarrow{\sigma} U \xrightarrow{\theta} V \xrightarrow{\sigma} C.$$

We will use this almost half-involution property in Section 6 to find free-start collisions and even hash function collisions for some IDEA-based constructions.

T-function. When using the null key, we have already described that all operations remaining are either XOR or additions. These operations are triangular functions [23] (or T-functions) in the sense that any output bit at position i only depends on the input bits located at a position i or lower. A composition of T-functions is itself a T-function, therefore the whole permutation defined by IDEA with the null key is a T-function. As shown in [34], this property might be very dangerous in a hash function design. We will explain in Section 7 how to exploit this weakness and compute preimages by guessing the input words bit layer by bit layer.

5 Simple Collision Attacks

As shown by Daemen *et al.* [10], when using the null key for the encryption process of IDEA, differences inserted uniquely on the MSB of the four 16-bit input plaintext words will lead to differences on the MSB of the four 16-bit output ciphertext words. Moreover, since this difference mapping is linear (the difference on the carry is not propagated further than the MSB), all possible differential characteristics have a differential probability 1. For example, we denote by $\delta_{MSB} = 0x8000$ the 16-bit word with difference only on the MSB and by $\Delta_{MSB} = (\delta_{MSB}, \delta_{MSB}, \delta_{MSB}, \delta_{MSB})$ the 64-bit difference composed of 4 words with difference δ_{MSB} . Then, Δ_{MSB} propagates to itself with probability 1 through one round of IDEA, or through its last half-round. Therefore, we have with probability 1

$$\Delta_{MSB} \xrightarrow{\text{IDEA}_{K=0}} \Delta_{MSB}.$$

Note that instead of using δ_{MSB} only, one can generalize the input difference space and obtain other very good differential paths for the encryption of IDEA with the null key. However, we omit this generalization here since the methods described in later sections already provide much better attacks.

Davies-Meyer. Finding a free-start collision on Davies-Meyer mode instantiated with IDEA is very easy. Since the difference Δ_{MSB} is mapped to itself through the IDEA encryption process with the null key, the attacker only has to pick $M = 0$. Then, any value of CV with difference Δ_{MSB} applied to it will lead to a collision with probability 1. We give in the full version of the article examples of such a free-start collision.

Hirose. The same method as for Davies-Meyer mode can be applied to the Hirose mode in order to find free-start collisions. The attacker fixes $CV2 = 0$ and $M = 0$ so as to force the null key to both encryptions. Then, any value of $CV1$ with a difference Δ_{MSB} applied to it will lead to a collision with probability 1, since Δ_{MSB} will appear on the plaintext input of both encryptions with the null key. We give in the full version of the article examples of such a free-start collision.

Abreast-DM. This technique seems impossible to apply to the Abreast-DM mode since forcing a difference Δ_{MSB} on any of the two encryptions plaintext input will imply a difference inserted in the key input of the other encryption block. Therefore, one cannot use Δ_{MSB} difference on plaintext input with null key in both encryption blocks. Even if the attacker tries to attack only one encryption block with this method, the other block will not be controlled and he will have to deal with random differences on its output. These random differences cannot be dealt with some birthday technique because fixing all inputs of one encryption block will fix all inputs of the other one as well.

Tandem-DM. This technique seems impossible to apply to the Tandem-DM mode for the exact same reasons as for Abreast-DM.

Peyrin *et al.*(II). We have to separate in two groups the possible instances of this construction, obtained by permuting the position of the three inputs of each internal function f_i . If all compression function inputs $CV1$, $CV2$, $M1$ and $M2$ appear in at least one of the IDEA key inputs of any f_i internal function, then the attack will not apply. Indeed, since all inputs will be involved at least one time, the attacker will necessarily have to insert a difference in at least one IDEA key input and he will not be able to use the differential path with probability 1. Note that these instances would be avoided in practice because they would lead to more frequent re-keying and therefore reduce the overall performance of the hash function. If this condition is not met, then we can apply the following free-start collision attack. Let $X \in \{CV1, CV2, M1, M2\}$ denote the input that is missing in all the IDEA key inputs of the compression function. The attacker simply fixes the difference Δ_{MSB} on X (one can give any value to X) and all other inputs are set to 0 in order to get the null key in every internal IDEA. The attacker ends up with several Davies-Meyer in parallel, with either no difference at all or with null key and Δ_{MSB} as plaintext input difference. Thus, he obtains a collision with probability 1. If $X \notin \{CV1, CV2\}$, then this attack finds semi-free-start collisions.

MJH-Double. The MJH-Double mode prevents this simple attack since even if we fix $CV2 = 0$ and $M2 = 0$ in order to get the null key in both encryptions, it is hard to force the difference Δ_{MSB} on both their plaintext inputs. Indeed, the f operation will randomize the difference and in order for the attack to run, we would require $\Delta_{MSB} \xrightarrow{f} \Delta_{MSB}$ which is unlikely to happen.

6 Improved Collision Attacks

In this section, using the almost half-involution property with the null key, we will show how to get the same difference on the input and on the output of the IDEA ciphering process with good probability. Then, we will use this weakness to derive our collision attacks, for any number of rounds.

6.1 Exploiting the Almost Half-Involution

We have already shown in Section 4 that when the key is null, IDEA encryption process can be rewritten as

$$P \xleftarrow{\sigma} U \xrightarrow{\theta} V \xrightarrow{\sigma} C$$

where

$$\sigma = \{MA_0 \circ KA_0 \circ S\}^3 \circ MA_0 \circ KA_0 \quad \text{and} \quad \theta = KA_0 \circ S.$$

We denote ΔU the XOR difference between two 64-bit internal state values U and U' , i.e. $\Delta U = U \oplus U'$, and δU_i represents the 16-bit difference on the i -th word of ΔU , that is $\Delta U = (\delta U_1, \delta U_2, \delta U_3, \delta U_4)$. Let us consider two random 64-bit internal state values U and U' such that $\delta U_2 = \delta U_3$ and we denote this 16-bit difference δ_M . For truly random values U and U' , this condition happens with probability 2^{-16} . One can check that applying θ on U and U' to obtain V and V' respectively will lead to $\delta V_2 = \delta V_3 = \delta_M$ since layer S only switches the two middle words and layer KA_0 has no effect on them (addition of null subkeys).

Let δ_L and δ_R represent the difference on δU_1 and δU_4 respectively, i.e. $\Delta U = (\delta_L, \delta_M, \delta_M, \delta_R)$. Applying function θ to U and U' , we would like the same differences to appear on internal state V and V' : $\Delta V = (\delta_L, \delta_M, \delta_M, \delta_R)$. The previous condition with probability 2^{-16} already ensures the two middle differences being the same δ_M . Concerning differences δ_L and δ_R , they will both be unaffected by layer S, but they might be modified through layer KA_0 that applies a multiplication with a null subkey. Therefore, we need to study the probability that a random difference δ is mapped to itself through a multiplication by the null subkey. We show in the full version of the article that this probability is equal to $2^{-1.585}$ and finally we have $\Pr[(\delta_L, \delta_M, \delta_M, \delta_R) \xrightarrow{\theta} (\delta_L, \delta_M, \delta_M, \delta_R)] = 2^{-3.17}$.

At this point, we proved that for randomly chosen internal state values U and U' , we will observe with probability $2^{-19.17}$ the same difference on U and V , i.e. $\Delta U = \Delta V$.

One can see that computing backward from internal states U to P or forward from V to C , the function σ is applied. Our final goal is to have the same difference on P and C . However, this seems unlikely to happen since U and V have different values, the forward and backward computations of σ should be completely unrelated, even with the same input difference. Yet, this reasoning does not take into account the fact that while U and V have distinct values, they are far from being independent: $V = \theta(U)$ with θ being a very light function. Moreover, we remarked that almost each time that we got the same difference on P and C , the same differences were observed as well in all rounds of the forward and backward σ computations (the round success probability increasing with the number of rounds already processed). Because all the rounds are not independent and because U and V are strongly related, it is very difficult to compute theoretically the probability of observing the same difference on P and C and we leave this as an open problem. Therefore, we measured it by choosing random values of U , δ_L , δ_M , δ_R , computing $V = \theta(U)$, and checking for collisions on the difference of P and C . The probability obtained was $2^{-16.26}$ for about 2^{28} tests (note that this probability somehow contains the $2^{-3.17}$ probability computed previously, but we can not separate them because the two events are not independent).

To conclude, the probability that two randomly chosen internal state values U and U' give the same difference on P and C is equal to $2^{-16-16.26} = 2^{-32.26}$ (instead of 2^{-64} expected for a random function). In other words, using the birthday paradox, one can find such a pair with about $2^{16.13}$ computations.

Interestingly, we have observed that most of the pairs fulfilling the differential path for the full IDEA will also be valid for a strengthened version of the cipher with any number of additional rounds. Since the subkeys are always null, strengthening the cipher would mean that $\sigma = \{\text{MA}_0 \circ \text{KA}_0 \circ \text{S}\}^t \circ \text{MA}_0 \circ \text{KA}_0$ for any $t > 3$. We checked that the probability that two randomly chosen internal state values U and U' give the same difference on P and C tends to $2^{-32.54}$ when t tends to infinite. Thus, similarly to the method presented in the previous section, the attacks using this almost half-involution property will work for any number of rounds.

6.2 Improving Collision Attacks

Davies-Meyer. A first obvious application of having the same difference in P and C is collision search on Davies-Mayer mode, where the feed-forward will cancel the two differences in the output. The attack finds collisions for the whole hash function and the procedure is very simple: we start from the IV and add random differences in the first message block M_0 . This will cause random differences in the the first chaining variable CV_1 . For the second message block M_1 , we will set all its bits 0 ($M_1 = 0$), forcing the internal IDEA computation to use the null key. Since we estimated in the previous section that with the null key a random pair of inputs has a probability $2^{-32.26}$ to give the same input/output difference, one can use the birthday paradox to generate a collision on CV_2 with only $2^{16.13}$ distinct message blocks M_0 . We give in the full version of the article examples of hash collisions for the Davies-Meyer mode. Note that finding

semi-free-start collisions with this technique is impossible since we would have to insert differences in the message input, which forbids the use of the null key in the internal cipher.

Hirose. We already showed how to find free-start collisions for the Hirose mode. However, finding semi-free-start collisions with this technique is impossible since we would have to insert differences in the message input, which forbids the use of the null key in the internal cipher. Also, concerning hash collisions, it seems hard as well because forcing the null key during iteration i requires us to obtain a chaining variable $CV2_{i-1} = 0$ during the previous iteration. This half-preimage already costs the same complexity as a generic collision search on the entire compression function.

Abreast-DM. One can derive a free-start collision attack for the Abreast-DM compression function using this technique. The attacker first fixes $CV1 = 0$ and $M = 0$. Then, he builds a set of $2^{48.13}$ distinct values $CV2$ and checks if a pair of this set leads to a collision. The probability that a pair leads to a collision on the first (top) branch is $2^{-32.26}$ (since the internal cipher on this part has the null key), and 2^{-64} on the other half. Overall, using the birthday paradox on the set of $2^{48.13}$ values $CV2$ is sufficient to have a good chance to obtain a collision. Note that finding a semi-free-start collision for the compression function or a collision for the hash function seems impossible with this method, for the same reasons as the Hirose mode.

Tandem-DM. The situation of Tandem-DM is absolutely identical to the Abreast-DM one: one can find free-start collisions for compression function using this technique. The attacker first fixes $CV1 = 0$ and $M = 0$. Then, he builds a set of $2^{48.13}$ distinct values $CV2$ and checks if a pair of this set leads to a collision. The probability that a pair leads to a collision on the first (top) branch is $2^{-32.26}$ (since the internal cipher on this part has the null key), and 2^{-64} on the other half. Overall, using the birthday paradox on the set of $2^{48.13}$ values $CV2$ is sufficient to have a good chance to obtain a collision. Again, finding a semi-free-start collision for the compression function or a collision for the hash function seems impossible with this method, for the same reasons as the Hirose mode.

Peyrin *et al.*(II). We showed in previous section how to find (semi)-free-start collisions with probability 1 for a certain subset of Peyrin *et al.*(II) constructions, but here we provide attacks on a bigger subset. If all compression function inputs $CV1$, $CV2$, $M1$ and $M2$ appear in at least one of the IDEA key inputs of f_1 , f_2 , f_3 (left side) and in at least one of the IDEA key inputs of f_3 , f_4 , f_5 (right side), then the attack will not apply. Indeed, for both left side and right side of the compression function, the attacker will necessarily have to insert a difference in at least one key input (since all inputs will be involved) and he will not be able to use the null key completely. Note that these instances would be avoided in practice because they would lead to more frequent rekeying and therefore

reduce the overall performance of the hash function. However, if this condition is not met, then we can apply the following free-start collision attack. Let $X \in \{CV1, CV2, M1, M2\}$ denote the input that is missing in all the IDEA key inputs of f_1, f_2, f_3 (wlog the reasoning is the same with f_3, f_4, f_5). The attacker first fixes all inputs but X to 0 in order to get the null key in every internal IDEA on the left side. Then he chooses $2^{48.13}$ random values for X and checks among them if any pair collides on the whole compression function output. Since he has a probability $2^{-32.26}$ to get a collision on the left side and 2^{-64} on the right side, using a birthday search the attacker finds a solution with complexity $2^{48.13}$. Again, if $X \notin \{CV1, CV2\}$, then this attack finds semi-free-start collisions. However, finding a collision for the hash function seems impossible with this method, because at least one of the chaining variable inputs $CV1$ and $CV2$ will be present as key input for one of the IDEA internal encryption. Setting this word to 0 is equivalent to a half-preimage that already costs the same complexity as a generic collision search on the entire hash function.

MJH-Double. One can derive a semi-free-start collision attack on the MJH-Double compression function instantiated with IDEA. The attacker first fixes $CV2 = 0$ and $M2 = 0$ and this will force the null key in both encryptions. Now he chooses a random value for $CV1$ (note that actually this value could be fixed by the challenger) and builds a set of $2^{32.26}$ values $M1$. In this configuration, it is easy to see that one will have random differences on the plaintext inputs to both encryptions. Since the null key is used for both, we have a probability $2^{-64.52}$ that a pair of $M1$ leads to a collision after the feed-forward of both encryptions (on the output of the bottom block and just before the application of g on the top block). Therefore, with a birthday technique, one can find such a pair with only $2^{32.26}$ computations. Note that while this pair will directly lead to a collision on the bottom $CV1$ output, the difference on $M1$ is injected two times before computing the top $CV2$ output. Two times of the same difference will cancel themselves and we eventually get a full semi-free-start collision. Note that it seems hard to extend this attack to a hash collision since the attacker would require to force the incoming chaining variable $CV2$ to be equal to 0 and this half-preimage already costs the same complexity as a generic collision search on the entire hash.

7 Preimage Attacks

Due to space limitations, all results regarding preimage attacks are given in the full version of the article.

8 Results and Implementations

We depict in Table 1 our collision results for the block cipher to compression function modes considered in this article when instantiated with IDEA. We implemented all attacks of reasonable complexities and provide in the full version of the article the collision/preimage examples obtained.

Table 1. Summary of collision results for block cipher to compression function modes when instantiated with IDEA (we did not include MDC-2 as it does not provide ideal collision resistance). The results for Peyrin *et al.*(II) construction, marked with a *, depend on the instance considered (see relevant parts of Sections 5 and 6 for more details).

Mode	hash output size	compression function		hash function
		free-start collision attack	semi-free-start collision attack	collision attack
Davies-Meyer [1]	64	2^1		$2^{16.13}$
Hirose [19,20]	128	2^1		
Abreast-DM [27,39]	128	$2^{48.13}$		
Tandem-DM [27,39]	128	$2^{48.13}$		
Peyrin <i>et al.</i> (II) [35]	128	$2^1 / 2^{48.13*}$	$2^1 / 2^{48.13*}$	
MJH-Double [29]	128	$2^{32.26}$	$2^{32.26}$	

9 Conclusion

In this article, we showed collision and preimage attacks for several single and double-length block cipher based compression function constructions when instantiated with the block cipher IDEA. Namely, we analyzed all known double-key schemes such as Davies-Meyer, Hirose, Abreast-DM, Tandem-DM, Peyrin *et al.* (II) and MJH-Double. While most of these constructions are conjectured or proved to be secure in the ideal cipher model, we showed that their security is very weak when instantiated with the block cipher IDEA, which remains considered as secure in the secret key model. In particular, we answer in the negative for the 20-year-old standing open question concerning the security of the Abreast-DM and Tandem-DM instantiated with IDEA. All our practical attacks have been implemented and they can work even for any number of IDEA rounds. Our results indicate that one has to be very careful when hashing with a block cipher that presents any weakness when the key is known or controlled by the attacker. Also, since we extensively use the presence of weak-keys for IDEA, as a future work it would be interesting to look at the security of hash functions based on block ciphers for which some key sets are known to be weaker than others.

Acknowledgments. The authors would like to thank the anonymous referees for their helpful comments.

References

1. Menezes, A., van Oorschot, P., Vanstone, S.: CRC-Handbook of Applied Cryptography. CRC Press (1996)
2. Ayaz, E.S., Selçuk, A.A.: Improved DST Cryptanalysis of IDEA. In: Biham, E., Youssef, A.M. (eds.) SAC 2006. LNCS, vol. 4356, pp. 1–14. Springer, Heidelberg (2007)

3. Biham, E., Dunkelman, O., Keller, N.: New Cryptanalytic Results on IDEA. In: Lai and Chen [25], pp. 412–427
4. Biham, E., Dunkelman, O., Keller, N.: A New Attack on 6-Round IDEA. In: Biryukov, A. (ed.) FSE 2007. LNCS, vol. 4593, pp. 211–224. Springer, Heidelberg (2007)
5. Biham, E., Dunkelman, O., Keller, N., Shamir, A.: New Data-Efficient Attacks on Reduced-Round IDEA. Cryptology ePrint Archive, Report 2011/417 (2011)
6. Biryukov, A., Nakahara Jr., J., Preneel, B., Vandewalle, J.: New Weak-Key Classes of IDEA. In: Deng, R.H., Qing, S., Bao, F., Zhou, J. (eds.) ICICS 2002. LNCS, vol. 2513, pp. 315–326. Springer, Heidelberg (2002)
7. Black, J.A., Rogaway, P., Shrimpton, T.: Black-Box Analysis of the Block-Cipher-Based Hash-Function Constructions from PGV. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 320–335. Springer, Heidelberg (2002)
8. Brassard, G. (ed.): CRYPTO 1989. LNCS, vol. 435. Springer, Heidelberg (1990)
9. Chang, D.: Near-Collision Attack and Collision-Attack on Double Block Length Compression Functions based on the Block Cipher IDEA. Cryptology ePrint Archive, Report 2006/478 (2006), <http://eprint.iacr.org/>
10. Daemen, J., Govaerts, R., Vandewalle, J.: Weak Keys for IDEA. In: Stinson [37], pp. 224–231
11. Daemen, J., Rijmen, V.: The Design of Rijndael: AES - The Advanced Encryption Standard. Springer (2002)
12. Damgård, I.: A Design Principle for Hash Functions. In: Brassard [8], pp. 416–427
13. Demirci, H., Selçuk, A.A., Türe, E.: A New Meet-in-the-Middle Attack on the IDEA Block Cipher. In: Matsui and Zuccherato [32], pp. 117–129
14. den Boer, B., Bosselaers, A.: Collisions for the Compression Function of MD-5. In: Helleseht, T. (ed.) EUROCRYPT 1993. LNCS, vol. 765, pp. 293–304. Springer, Heidelberg (1994)
15. Dobbertin, H.: Cryptanalysis of MD5 compress. Presented at the Rump Session of EUROCRYPT 1996 (1996)
16. Fleischmann, E., Gorski, M., Lucks, S.: On the Security of TANDEM-DM. In: Dunkelman, O. (ed.) FSE 2009. LNCS, vol. 5665, pp. 84–103. Springer, Heidelberg (2009)
17. Fleischmann, E., Gorski, M., Lucks, S.: Security of Cyclic Double Block Length Hash Functions including Abreast-DM. Cryptology ePrint Archive, Report 2009/261 (2009), <http://eprint.iacr.org/>
18. Hawkes, P.: Differential-Linear Weak Key Classes of IDEA. In: Nyberg, K. (ed.) EUROCRYPT 1998. LNCS, vol. 1403, pp. 112–126. Springer, Heidelberg (1998)
19. Hirose, S.: Provably Secure Double-Block-Length Hash Functions in a Black-Box Model. In: Park, C., Chee, S. (eds.) ICISC 2004. LNCS, vol. 3506, pp. 330–342. Springer, Heidelberg (2005)
20. Hirose, S.: Some Plausible Constructions of Double-Block-Length Hash Functions. In: Robshaw, M. (ed.) FSE 2006. LNCS, vol. 4047, pp. 210–225. Springer, Heidelberg (2006)
21. Kelsey, J., Schneier, B., Wagner, D.: Key-Schedule Cryptanalysis of IDEA, G-DES, GOST, SAFER, and Triple-DES. In: Koblitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 237–251. Springer, Heidelberg (1996)
22. Khovratovich, D., Leurent, G., Rechberger, C.: Narrow-Bicliques: Cryptanalysis of Full IDEA. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 392–410. Springer, Heidelberg (2012)
23. Klimov, A., Shamir, A.: Cryptographic Applications of T-Functions. In: Matsui and Zuccherato [32], pp. 248–261

24. Knudsen, L.R., Rijmen, V.: Known-Key Distinguishers for Some Block Ciphers. In: Kurosawa, K. (ed.) ASIACRYPT 2007. LNCS, vol. 4833, pp. 315–324. Springer, Heidelberg (2007)
25. Lai, X., Chen, K. (eds.): ASIACRYPT 2006. LNCS, vol. 4284. Springer, Heidelberg (2006)
26. Lai, X., Massey, J.L.: A Proposal for a New Block Encryption Standard. In: Damgård, I.B. (ed.) EUROCRYPT 1990. LNCS, vol. 473, pp. 389–404. Springer, Heidelberg (1991)
27. Lai, X., Massey, J.L.: Hash Functions Based on Block Ciphers. In: Rueppel, R.A. (ed.) EUROCRYPT 1992. LNCS, vol. 658, pp. 55–70. Springer, Heidelberg (1993)
28. Lee, J., Kwon, D.: The Security of Abreast-DM in the Ideal Cipher Model. Cryptology ePrint Archive, Report 2009/225 (2009), <http://eprint.iacr.org/>
29. Lee, J., Stam, M.: MJH: A Faster Alternative to MDC-2. In: Kiayias, A. (ed.) CT-RSA 2011. LNCS, vol. 6558, pp. 213–236. Springer, Heidelberg (2011)
30. Lee, J., Stam, M., Steinberger, J.: The Collision Security of Tandem-DM in the Ideal Cipher Model. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 561–577. Springer, Heidelberg (2011)
31. Luby, M., Rackoff, C.: How to Construct Pseudorandom Permutations from Pseudorandom Functions. *SIAM J. Comput.* 17(2), 373–386 (1988)
32. Matsui, M., Zuccherato, R.J. (eds.): SAC 2003. LNCS, vol. 3006. Springer, Heidelberg (2004)
33. Merkle, R.C.: One Way Hash Functions and DES. In: Brassard [8], pp. 428–446
34. Muller, F., Peyrin, T.: Cryptanalysis of T-Function-Based Hash Functions. In: Rhee, M.S., Lee, B. (eds.) ICISC 2006. LNCS, vol. 4296, pp. 267–285. Springer, Heidelberg (2006)
35. Peyrin, T., Gilbert, H., Muller, F., Robshaw, M.J.B.: Combining Compression Functions and Block Cipher-Based Hash Functions. In: Lai and Chen [25], pp. 315–331
36. Preneel, B., Govaerts, R., Vandewalle, J.: Hash Functions Based on Block Ciphers: A Synthetic Approach. In: Stinson [37], pp. 368–378
37. Stinson, D.R. (ed.): CRYPTO 1993. LNCS, vol. 773. Springer, Heidelberg (1994)
38. Wang, X., Yu, H.: How to Break MD5 and Other Hash Functions. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 19–35. Springer, Heidelberg (2005)
39. Lai, X.: On the Design and Security of Block Ciphers. Hartung-Gorre Verlag, Konstanz (1992)

The Security of Ciphertext Stealing

Phillip Rogaway¹, Mark Wooding², and Haibin Zhang¹

¹ Dept. of Computer Science, University of California, Davis, USA

² Thales e-Security Ltd, UK

Abstract. We prove the security of CBC encryption with ciphertext stealing. Our results cover all versions of ciphertext stealing recently recommended by NIST. The complexity assumption is that the underlying blockcipher is a good PRP, and the security notion achieved is the strongest one commonly considered for chosen-plaintext attacks, indistinguishability from random bits (ind $\$$ -security). We go on to generalize these results to show that, when intermediate outputs are slightly delayed, one achieves ind $\$$ -security in the sense of an online encryption scheme, a notion we formalize that focuses on what is delivered across an online API, generalizing prior notions of blockwise-adaptive attacks. Finally, we pair our positive results with the observation that the version of ciphertext stealing described in Meyer and Matyas’s well-known book (1982) is not secure.

Keywords: blockwise-adaptive attacks, CBC, ciphertext stealing, cryptographic standards, modes of operation, provable security.

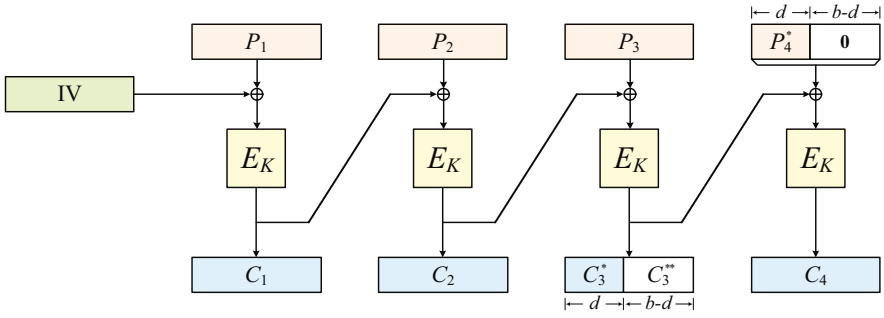
1 Introduction

CIPHERTEXT STEALING. Many blockcipher modes require the input be a sequence of complete blocks, each having a number of bits that is the blockcipher’s blocksize. One approach for dealing with inputs not of this form is *ciphertext stealing*. The classical combination is CBC encryption and ciphertext stealing, a mode going back to at least 1982 [14].

In 2010, NIST put out an addendum [8] to Special Publication 800-38A [7], the document that had defined blockcipher modes ECB, CBC, CFB, OFB, and CTR. The addendum defines three ways to enrich CBC with ciphertext stealing. The modes are named CBC-CS1, CBC-CS2, and CBC-CS3. See Fig. 1 for the definition of these modes, which differ only in the ordering of ciphertext bits.

Despite the classicism of ciphertext-stealing, its adoption in standards, and the strong preferences, these days, for proven-secure modes, there has, until now, been no proof offered for CBC with ciphertext stealing. This paper fills in this gap.

OUR CONTRIBUTIONS. We begin by looking at the NIST ciphertext-stealing modes, which we collectively call CBC-CS. Assuming a random IV, we show that the CBC-CS schemes achieve the strongest conventional form of chosen-plaintext-attack (CPA) security: what we call ind $\$$, indistinguishability from



```

10  algorithm CBC-CSIVK(P)
11  n ← ⌈|P|/b⌉
12  P1 ⋯ Pn-1 Pn* ← P where |P1| = ⋯ = |Pn-1| = b
13  Pn ← Pn* 0b-d where d ← |Pn*|
14  C0 ← IV;
15  C1 ⋯ Cn ← CBCIVK(P1 ⋯ Pn) where |C1| = ⋯ = |Cn| = b
16  Cn-1* ← MSBd(Cn-1)
17-1 return C1 ⋯ Cn-2 Cn-1* Cn                                     ⇐ for CS1
17-2 if d=b return C1 ⋯ Cn-2 Cn-1* Cn else return C1 ⋯ Cn-2 Cn Cn-1* ⇐ for CS2
17-3 return C1 ⋯ Cn-2 Cn Cn-1*                                     ⇐ for CS3

```

```

20  algorithm CBCIVK(P1 ⋯ Pn) where |P1| = ⋯ = |Pn| = b
21  C0 ← IV
22  for i ← 1 to n do Ci ← EK(Ci-1 ⊕ Pi)
23  return C1 ⋯ Cn

```

Fig. 1. Encryption under NIST modes CBC-CS1, CBC-CS2, and CBC-CS3. The schemes differ only in which version of line 17 is used. The schemes depend on a blockcipher $E: \mathcal{K} \times \{0, 1\}^b \rightarrow \{0, 1\}^b$ that determines the key space \mathcal{K} , the IV space \mathcal{IV} , and the message space $\mathcal{P} = \{0, 1\}^{\geq b}$. We insist that $K \in \mathcal{K}$, $IV \in \mathcal{IV}$, and $P \in \mathcal{P}$.

random bits under an adaptive chosen-plaintext attack. The definition, easily shown to imply all conventional formulations of CPA-style semantic security, formalizes that a ciphertext C is indistinguishable from as many random bits.

Next we show that *delayed* versions of CBC-CS achieve an analogous IND\$ notion that we define for *online* security. The idea of delayed CBC is from Fouque, Martinet, and Poupard [11]. Our formulation for online security generalizes their and subsequent work (further history and credits coming shortly). In particular, prior definitional approaches were specific to blockcipher-based schemes of a specified form—restrictions not in keeping with identifying a general notion of security. We levy no such restrictions, but do imagine that the encryption scheme is written to an incremental API (application programming interface). Each time a user presents a piece of plaintext to encrypt she will get back a corresponding chunk of ciphertext. The length of both is arbitrary. One can understand our definition of online security as establishing that a specified incremental API introduces no new security vulnerabilities. Technically, we reconceptualize an

encryption scheme *as* the incremental interface. We regard a general definition for online security—a definition motivated by cryptographic APIs and not the characteristics of any particular encryption mode—as an important and independent contribution of this paper.

The workings of delayed CBC—the naturalness of this scheme and how much one must delay—are clarified by freeing the definition of online security from a demand on a scheme being blockcipher-based. Now it is the security analysis, not the syntax, that surfaces by just how much one must delay—an amount that is, in fact, slightly different for the CS1/CS2 and the CS3 versions of the scheme. Absent a careful treatment of such matters the author of an incremental API could well get these things wrong, buffering more than what is necessary or less than what is needed.

Finally, we point out that a 30-year-old version of ciphertext stealing described in the book of Meyer and Matyas [14] is essentially wrong: it will not achieve any desirable security notion we know. The apparently unnoticed observation highlights the importance of having proofs in this domain, and underscores NIST’s wisdom in selecting the versions of ciphertext stealing that it did.

ADDITIONAL HISTORY. The provable-security treatment of CBC, and of other blockcipher-based encryption modes, begins with Bellare, Desai, Jokipii, and Rogaway [3]. The stronger ind $\$$ -definition that we adopt here is from Rogaway, Bellare, Black, and Krovetz [16]. For online security, the delayed-CBC scheme that we embellish with NIST’s versions of ciphertext stealing is due to Fouque, Martinet, and Poupard [11].

Our definition of online security springs from the line of work on blockwise-adaptive attacks that starts with Bellare, Kohno, and Namprempre [4] and Joux, Martinet, and Valette [13] and continues with Fouque, Martinet, and Poupard [11], Fouque, Joux, and Poupard [10], and Bard [2]. As explained, our own security definitions take a different turn by divorcing the notion of online security from its former association with blockcipher-based schemes. We instead assume an arbitrary symmetric encryption scheme that is presented to the user by way of an incremental API. The user provides the plaintext as a sequence of chunks and the encryption algorithm, buffering what it needs, returns corresponding ciphertext chunks. The approach echos Gennaro and Rohatgi [12], which likewise transplants a primitive (digital signatures) from a setting that sees messages as atomic to one that sees messages as something produced and consumed across an expanse of time.

DISCUSSION. A possible reaction to any discussion of ciphertext stealing is to say: forget it, use CTR mode instead. We are sympathetic to this point of view, knowing no convincing reason to favor CBC encryption over CTR mode, which natively handles plaintexts of arbitrary length. But the fact remains that CBC encryption is widely used, and that ciphertext stealing is a classical, standardized, and elegant way to extend it. This makes it worth attending to.

In justifying the use of ciphertext stealing in a mode that employed it, Matt Ball writes that “[d]espite lacking a formal security proof, ciphertext

stealing still has general approval in the cryptographic community” [1, p. 5]. Probably this statement is at some level true, but “general approval” is hard to gauge and far removed from being a proof.

We think that security notions that attend to the vulnerabilities introduced by the specifics of an envisioned API comprise an interesting direction in narrowing the gap between conventional abstractions of cryptographic primitives and what cryptographic practice actually exports. It is not just that protocols may segment conceptually atomic messages (the original motivation for dealing with blockwise adaptivity); rather, it is that the segmentation is actually surfaced to users, and therefore desirable to directly model.

We do not discuss the security of CBC-CS when the IV fails to be unpredictable; it would seem that no interesting or desirable security notion is achieved in this case. NIST SP800-38A appropriately demands an unpredictable IV for CBC [7, Appendix C].

The CBC-CS schemes predate NIST’s addendum [8]: CBC-CS2 goes back to at least 1996 [17], while older versions of ciphertext stealing go back to at least 1982 [14]. Looking at these schemes from a modern vantage is long overdue.

2 Preliminaries

NOTATION. Strings are assumed to be binary, elements of $\{0,1\}^*$. Both $A \parallel B$ and AB denote the concatenation of strings A and B . If X is a string then $|X|$ is its length. The empty string is denoted ε . Throughout this paper we fix an integer $b \geq 1$ called the *blocksize*. For a string X and a number $d \leq |X|$ let $\text{MSB}_d(X)$ and $\text{LSB}_d(X)$ be the leftmost and rightmost d bits of X .

BLOCKCIPHERS. A *blockcipher* is a map $E: \mathcal{K} \times \{0,1\}^b \rightarrow \{0,1\}^b$ where $\mathcal{K} \subseteq \{0,1\}^*$ is finite and $E_K(\cdot) = E(K, \cdot)$ is a permutation for each $K \in \mathcal{K}$. Let $\text{Perm}(b)$ be the set of all permutations on b bits. This may be regarded as a blockcipher with a $(2^b!)$ -size key space. Let $\mathbf{Adv}_E^{\text{prp}}(A) = \Pr[A^{E_K(\cdot)} \Rightarrow 1] - \Pr[A^{\pi(\cdot)} \Rightarrow 1]$ with $K \xleftarrow{\$} \mathcal{K}$ and $\pi \xleftarrow{\$} \text{Perm}(b)$. Similarly define $\mathbf{Adv}_E^{\text{prf}}(A) = \Pr[A^{E_K(\cdot)} \Rightarrow 1] - \Pr[A^{\rho(\cdot)} \Rightarrow 1]$ with $K \xleftarrow{\$} \mathcal{K}$ and $\rho \xleftarrow{\$} \text{Func}(b)$ the set of all functions from b bits to b bits. Here $E_K(\cdot)$ need not be a permutation.

ENCRYPTION SCHEMES. It has become traditional to regard blockciphers as fixed functions but encryption schemes as tuples, as in $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$. To simplify and unify matters we formalize an encryption scheme more like a blockcipher: an (IV-based, symmetric) *encryption scheme* is a function $\mathcal{E}: \mathcal{K} \times \mathcal{IV} \times \mathcal{P} \rightarrow \mathcal{P}$. We call \mathcal{K} , \mathcal{IV} , and \mathcal{P} the *key space*, *IV space*, and *message space*. For simplicity we assume that \mathcal{K} is finite and \mathcal{IV} is the set of all strings of some one particular length. We write $\mathcal{E}_K^{\text{IV}}(P)$ instead of $\mathcal{E}(K, \text{IV}, P)$. To keep things simple we require that $\mathcal{E}_K^{\text{IV}}(\cdot)$ be a length-preserving permutation for all $K \in \mathcal{K}$ and $\text{IV} \in \mathcal{IV}$. The condition implies that \mathcal{E} has a unique inverse, the map \mathcal{D} where $\mathcal{D}_K^{\text{IV}}(C) = P$ when $\mathcal{E}_K^{\text{IV}}(P) = C$. Because there is no formal need to specify the decryption direction \mathcal{D} of an encryption scheme \mathcal{E} , we never do so. Of course it is important in practice that \mathcal{E} and \mathcal{D} have efficient realizations, it is simply that this doesn’t show up in the statement of definitions or security results.

Let $\mathcal{E}: \mathcal{K} \times \mathcal{IV} \times \mathcal{P} \rightarrow \mathcal{P}$ be an IV-based encryption scheme and let A be an adversary (algorithm) with one of two types of oracles. A *real* encryption oracle $\text{Real}(\cdot)$ chooses a random $K \xleftarrow{\$} \mathcal{K}$ and then, on input $P \in \mathcal{P}$, returns $C \leftarrow IV \parallel \mathcal{E}_K^{IV}(P)$ for a random $IV \xleftarrow{\$} \mathcal{IV}$. A *fake* encryption oracle $\text{Fake}(\cdot)$ takes an input $P \in \mathcal{P}$ and returns $C \xleftarrow{\$} \{0, 1\}^c$ where $c = |IV| + |P|$ (for $IV \in \mathcal{IV}$). Define $\text{Adv}_{\mathcal{E}}^{\text{ind}\$}(A) = \Pr[A^{\text{Real}(\cdot)} \Rightarrow 1] - \Pr[A^{\text{Fake}(\cdot)} \Rightarrow 1]$. This “indistinguishability-from-random-bits” definition is easily shown to imply all conventional (CPA) formulations of indistinguishability and semantic security [3]; that we have selected a different syntax makes no difference in the proofs.

Note that even though the encryption function is formalized as taking, besides the key, an IV and a plaintext, the security definition does not allow the adversary to specify the IV; the adversary asks P and the IV is randomly generated, used, and returned. Our security notion thus formalizes security for random IVs, not, for example, security for nonce IVs.

3 Conventional Security of the CBC-CS Schemes

We begin with a simple proposition about the security of conventional CBC encryption (no ciphertext stealing) with a random IV. The result is needed insofar as we deduce the security of CBC-CS from it. Recall that the mode was defined in Fig. 1 and was proven secure by Bellare *et al.* [3]. That proof, however, is for a somewhat weaker definition than the one we use here. The proof below is a simple application of the game-playing technique [5,18].

Lemma 1. *Suppose A asks queries totaling at most σ blocks. Then we have $\text{Adv}_{\text{CBC}[\text{Perm}(b)]}^{\text{ind}\$}(A) \leq \sigma^2/2^b$.*

Proof. The difference between $\text{Adv}_{\text{CBC}[\text{Perm}(b)]}^{\text{ind}\$}(A)$ and $r = \text{Adv}_{\text{CBC}[\text{Func}(b)]}^{\text{ind}\$}(A)$ is at most $0.5\sigma^2/2^b$; this is a standard application of PRP/PRF switching [5]. It thus suffices to bound r by $r \leq 0.5\sigma^2/2^b$. To that end, consider the games of Fig. 2. Observe that, with $\mathcal{E} = \text{CBC}[\text{Func}(b)]$, $\Pr[A^{\text{Real}(\cdot)} \Rightarrow 1] = \Pr[A^{G_1(\cdot)} \Rightarrow 1]$, while $\Pr[A^{\text{Fake}(\cdot)} \Rightarrow 1] = \Pr[A^{G_0(\cdot)} \Rightarrow 1]$. As a consequence, we have that $r = \Pr[A^{G_1(\cdot)} \Rightarrow 1] - \Pr[A^{G_0(\cdot)} \Rightarrow 1]$ and, the two games being identical-until-*bad*, we know that $r \leq \Pr[A^{G_0}$ sets *bad*]. Because in game G_0 all of the C_i values are uniform and independent of P_i , so too all of the X_i values are uniform and independent of one another, so the probability that *bad* gets set—the probability some two of the X_i ’s collide—is at most $(1 + 2 + \dots + (\sigma - 1))/2^b \leq 0.5\sigma^2/2^b$. This completes the proof.

Turning now to the CBC-CS modes, we claim that these inherit CBC’s security with no quantitative degradation. The needed observation is that $\text{CBC-CS1}_K^{IV}(P)$ is just $\text{CBC}_K^{IV}(P0^*)$ (minimal padding to the next multiple of b bits) with some bits excised and some bits reordered. Which bits are excised and how bits are rearranged depends only on $|P|$. Thus if $\text{CBC}_K^{IV}(\cdot)$ looks random, so too will look $\text{CBC-CS1}_K^{IV}(\cdot)$. The same comments hold for CBC-CS2 and CBC-CS3; these are just different rearrangements of the bits of $\text{CBC}_K^{IV}(P0^*)$. The observation and proof are formalized by the proposition below.

100	algorithm Enc(P)	Game G_0
101	$P_1 \cdots P_n \leftarrow P$ where $ P_1 = \cdots = P_n = b$	Game G_1
102	$C_0, \dots, C_n \xleftarrow{\$} \{0, 1\}^b$	
103	for $i \leftarrow 1$ to n do	
104	$X_i \leftarrow P_i \oplus C_{i-1}$	
105	if $\rho(X_i)$ then $bad \leftarrow \text{true}$, $C_i \leftarrow \rho(X_i)$	
106	$\rho(X_i) \leftarrow C_i$	
107	return $C_0 C_1 \cdots C_n$	

Fig. 2. Proof of the ind $\$$ -security of CBC encryption with a random IV. This application of game-playing is probably simple and well-known enough to be considered folklore. Game G_1 includes the boxed statement following the setting of bad ; game G_0 omits it. Variable bad is initialized to **false** and ρ is initialized to everywhere undefined, a value treated as **false** if used as a boolean.

Theorem 1. *Let \mathcal{E} be any of CBC-CS1[Perm(b)], CBC-CS2[Perm(b)], or CBC-CS3[Perm(b)] and suppose adversary A asks queries totaling at most σ blocks. Then $\text{Adv}_{\mathcal{E}}^{\text{ind}\$}(A) \leq \sigma^2/2^b$.*

Proof. Suppose that A , asking σ total blocks of queries, gets advantage δ at distinguishing oracles $\mathcal{E} = \text{CBC-CS1}(\cdot)$ and $\$(\cdot)$. The first of these oracles chooses a random permutation $\pi \xleftarrow{\$} \text{Perm}(n)$ and then, when asked a query $P \in \{0, 1\}^{\geq b}$, returns $IV \parallel \text{CBC-CS1}_{\pi}^{IV}(P)$ for a random $IV \xleftarrow{\$} \{0, 1\}^b$; the second oracle, when asked a query P , returns a random string of length $b + |P|$. We construct from A an adversary B that, also asking σ blocks worth of queries, also gets advantage δ , but now at distinguishing between $\text{CBC}(\cdot)$ and $\$(\cdot)$. The first of these oracle chooses a random permutation $\pi \xleftarrow{\$} \text{Perm}(n)$ and then, when asked a query $P \in (\{0, 1\}^b)^+$, returns $IV \parallel \text{CBC}_{\pi}^{IV}(P)$ for a random $IV \xleftarrow{\$} \{0, 1\}^b$. Adversary B now works as follows: it runs A and when A generates a query of $P \in \{0, 1\}^{\geq b}$ adversary B queries its own oracle on $P' = P 0^*$, meaning P padded on the right with the minimal number of zero-bits so that P' is a multiple of b bits. Suppose this returns a ciphertext $C = C_0 C_1 \cdots C_n$ where $|C_i| = n$. Then B returns to A the string $C^* = C_0 C_1 \cdots C_{n-1}^* C_n$ where $C_{n-1}^* = \text{MSB}_d(C_{n-1})$ and $d = b - (|P| \bmod b)$. We observe that $\Pr[B^{\text{CBC-CS1}(\cdot)} \Rightarrow 1] = \Pr[A^{\text{CBC}(\cdot)} \Rightarrow 1]$ (we have reordered bits exactly as required by CBC-CS1) and that $\Pr[B^{\$(\cdot)} \Rightarrow 1] = \Pr[A^{\$(\cdot)} \Rightarrow 1]$ (reordered and pruned uniform random bits are still uniform), and so $\delta = \text{Adv}_{\text{CBC-CS1}[\text{Perm}(b)]}^{\text{ind}\$}(A) = \text{Adv}_{\text{CBC}[\text{Perm}(b)]}^{\text{ind}\$}(B)$. By Proposition 1 we thus have $\delta \leq 0.5 \sigma^2/2^b$. This establishes the first of the three results. The analogous results for CBC-CS2 and CBC-CS3 are obtained simply by modifying the string C^* returned to A : for CBC-CS2 return $C^* = C_0 C_1 \cdots C_{n-2} C_{n-1}^* C_n$ when $|P|$ is a multiple of b and $C^* = C_0 C_1 \cdots C_{n-2} C_n C_{n-1}^*$ otherwise; for CBC-CS3 always return $C^* = C_0 C_1 \cdots C_{n-2} C_n C_{n-1}^*$. This completes the theorem.

The proof’s simplicity stems from having unidentified a clean abstraction boundary: directly modifying the proof of Lemma 1 to attend to the ciphertext stealing would be much more complex.

Finally, one can pass from the information-theoretic result to its complexity-theoretic analog in the standard way, trading the family of random permutations for a conventional blockcipher. Stating the result for completeness, we have the following.

Corollary 1. *Let $E: \mathcal{K} \times \{0, 1\}^b \rightarrow \{0, 1\}^b$ be a blockcipher and let \mathcal{E} be any of the encryption schemes CBC-CS1[E], CBC-CS2[E], or CBC-CS3[E]. Suppose A asks queries that total σ blocks, runs in time t , and achieves advantage $\delta = \text{Adv}_{\mathcal{E}}^{\text{ind}^S}(A)$. Then there is an adversary B , explicitly known and constructed from A in a blackbox manner, that asks at most σ queries, runs in time at most $t + \lambda\sigma$, and achieves advantage $\text{Adv}_E^{\text{prp}}(B) \geq \delta - \sigma^2/2^b$. Here λ is an absolute constant depending only on details of the model of computation.*

4 Defining Online Security

SYNTAX. We adjust the syntax of an encryption scheme to accommodate the staged presentation of plaintexts and ciphertexts. Rather than messages being atomic objects that get encrypted all at once, messages may be arbitrarily partitioned into chunks, each of which gets fed into a stateful encryption engine. Breaking with former treatments, we do not assume that chunks are single blocks, nor multiples of blocks, where the length of a block is the blocksize of some underlying blockcipher. Instead, we provide a general definition where one assumes nothing about the structure of the underlying encryption scheme (in particular, there is no assumption that it is blockcipher-based). As each installment of plaintext is provided to the encryption interface, it is up to the algorithm to decide how much ciphertext to spit out. The algorithm will thus return not only a ciphertext chunk, but also an updated state.

Realizing the idea above, we choose to define an *online encryption scheme* as a function $\mathcal{E}: \mathcal{K} \times \mathcal{V} \times \{0, 1\} \times \{0, 1\}^* \rightarrow \{0, 1\}^* \times \mathcal{V}$. We write $\mathcal{E}_K^{V, \delta}(P)$ for $\mathcal{E}(K, V, \delta, P)$. We call \mathcal{K} and \mathcal{V} the *key space* and *state space*, respectively. The key space is finite and the state space is a finite set of strings. The third argument to \mathcal{E} , a bit, is the *end-of-message indicator*. The final argument to \mathcal{E} is the next chunk of *message*. An online encryption scheme \mathcal{E} must have an associated *IV space* $\mathcal{IV} \subseteq \mathcal{V}$ and *message space* $\mathcal{P} \subseteq \{0, 1\}^*$. The former contains strings of some one fixed length. Formally, an online encryption scheme is the tuple $(\mathcal{E}, \mathcal{IV}, \mathcal{P})$, but we will usually use the first component as shorthand for the whole.

We also impose a number of “syntactic” requirements on an online encryption scheme $(\mathcal{E}, \mathcal{IV}, \mathcal{P})$. First we define some additional notation. We write $(C_1, \dots, C_n) \leftarrow \mathcal{E}_K^{IV}(P_1, \dots, P_n)$ for the sequence:

$$\begin{aligned} &V_0 \leftarrow IV \\ &\mathbf{for } i \leftarrow 1 \text{ to } n - 1 \mathbf{ do } (C_i, V_i) \leftarrow \mathcal{E}_K^{V_{i-1}, 0}(P_i) \\ &(C_n, V_n) \leftarrow \mathcal{E}_K^{V_{n-1}, 1}(P_n) \\ &\mathbf{return } (C_1, \dots, C_n). \end{aligned}$$

Alternatively, we can think of $\mathcal{E}_K^{IV}(P_1, \dots, P_n)$ as returning a single string, setting $\mathcal{E}_K^{IV}(P_1, \dots, P_n)$ to $C = C_1 \cdots C_n$ where $(C_1, \dots, C_n) \leftarrow \mathcal{E}_K^{IV}(P_1, \dots, P_n)$.

Now fix an online encryption scheme $(\mathcal{E}, \mathcal{IV}, \mathcal{P})$.

- The *consistency requirement* says that you get the same ciphertext regardless of how you split up the plaintext. More formally, if $P_1 \parallel \dots \parallel P_n = P'_1 \parallel \dots \parallel P'_{n'}$ and $P \in \mathcal{P}$ then $\mathcal{E}_K^{IV}(P_1, \dots, P_n) = \mathcal{E}_K^{IV}(P'_1, \dots, P'_{n'})$. We can therefore write this as $\mathcal{E}_K^{IV}(P)$ without ambiguity.
- The *invertibility requirement* is that $\mathcal{E}_K^{IV}(\cdot)$ is injective on \mathcal{P} (for all $K \in \mathcal{K}$ and $IV \in \mathcal{IV}$).
- The *length requirement* is that the length of the first and second components of $\mathcal{E}_K^{V, \delta}(P)$ depend only on $|V|$, $|P|$, and δ . This ensures that, when $(C_1, \dots, C_m) \leftarrow \mathcal{E}_K^{IV}(P_1, \dots, P_m)$, the lengths of C_1, C_2, \dots, C_m reveal nothing about $P = P_1 \dots P_m$ beyond how it was partitioned up.

INDISTINGUISHABILITY. We define a very strong form of indistinguishability for an online encryption scheme: indistinguishability from random bits. Fix an online encryption scheme $(\mathcal{E}, \mathcal{IV}, \mathcal{P})$ and consider the following two \mathcal{E} -dependent oracles.

- $\text{Real}(i, M, \delta)$: At the beginning, set $K \xleftarrow{\$} \mathcal{K}$ and $V_i \xleftarrow{\$} \mathcal{IV}$ for all $i \in \mathbb{N}$. Then, on query $(i, P, \delta) \in \mathbb{N} \times \{0, 1\}^* \times \{0, 1\}$, compute $(C, V_i) \xleftarrow{\$} \mathcal{E}_K^{V_i, \delta}(P)$ and return C .
- $\text{Fake}(i, P, \delta)$: At the beginning, set $K \xleftarrow{\$} \mathcal{K}$ and $V_i \xleftarrow{\$} \mathcal{IV}$ for all $i \in \mathbb{N}$. Then, on query $(i, P, \delta) \in \mathbb{N} \times \{0, 1\}^* \times \{0, 1\}$, compute $(C, V_i) \xleftarrow{\$} \mathcal{E}_K^{V_i, \delta}(P)$ and return $|C|$ random bits.

We define $\text{Adv}_{\mathcal{E}}^{\text{IND}^{\$}}(A) = \Pr[A^{\text{Real}} \Rightarrow 1] - \Pr[A^{\text{Fake}} \Rightarrow 1]$. Informally, an online encryption scheme is $\text{IND}^{\$}$ -secure if an adversary can't distinguish the ciphertexts it is receiving from random bits.

DISCUSSION. Some of our high-level definitional choices differ for conventional and online encryption schemes. A conventional encryption scheme does not spit out its IV, while an online scheme does. The former is needed to match NIST's definitions for the CBC-CS schemes, but it works less well in the online setting, as here it is important that the algorithm can decide if and when to release the IV. Typically, the IV does get discharged, and as the first part of the ciphertext, so we say that an online encryption scheme $(\mathcal{E}, \mathcal{IV}, \mathcal{P})$ is *IV-prefixed* if $C = \mathcal{E}_K^{IV}(P)$ is always IV followed by some $|P|$ additional bits (assuming $K \in \mathcal{K}$ and $IV \in \mathcal{IV}$). An IV-prefixed online encryption scheme $(\mathcal{E}, \mathcal{IV}, \mathcal{P})$ determines a conventional encryption scheme $\hat{\mathcal{E}}$ in the natural way, setting $\hat{\mathcal{E}}_K^{IV}(P)$ to be $\mathcal{E}_K^{IV}(P)$ stripped of its initial $|IV|$ bits. Conversely, a conventional encryption scheme $\hat{\mathcal{E}}: \mathcal{K} \times \mathcal{IV} \times \mathcal{P} \rightarrow \mathcal{P}$ is realized by an IV-prefixed online encryption scheme $(\mathcal{E}, \mathcal{IV}, \mathcal{P})$ if the latter determines the former in the manner just defined. In this way one can speak of an encryption scheme $\mathcal{E}: \mathcal{K} \times \mathcal{IV} \times \mathcal{P} \rightarrow \mathcal{P}$ as being online; the statement means that it has a secure online realization (the notion of security soon to be defined).

While our notions make sense regardless of whether or not \mathcal{V} is finite, its being finite is the essence of what it means to be online: that one can encrypt (and decrypt) streaming messages without having to buffer more than a constant

number of bits. Equivalently, that one can implement an incremental API with a fixed-size context. Our notions allow one to consider things in a more quantitative manner, using $|\mathcal{V}|$ as a measure of worth. We say that $\mathcal{E}: \mathcal{K} \times \mathcal{V} \times \{0, 1\} \times \{0, 1\}^* \rightarrow \{0, 1\}^* \times \mathcal{V}$ uses v -bits of state if v is the smallest number such that $\mathcal{V} \subseteq \{0, 1\}^{\leq v}$.

Since we concern ourselves only with chosen-plaintext security, we do not formalize syntax or security for the decryption direction of an online encryption scheme. Still, we comment that if an incremental encryption scheme is online then it has an online (that is, finite state-space) decryption.

We regard the initialization vector IV as the initial value of the saved state V . The embedding of the IV space into the state space doesn't prevent a scheme from performing "special" initialization; one can always distinguish the first chunk of a message from subsequent chunks of message by arranging that point in $\mathcal{T}\mathcal{V}$ are never returned as a modified state.

An online encryption function has control over if and when the IV is revealed. This can be essential for security: in particular, the Delayed CBC scheme we will soon describe is insecure if the IV is revealed too soon.

Note that the IND $\$$ -definition allows interleaved querying of multiple streams; this is the purpose of the index i . Fouque, Martinet, and Poupard earlier observed that, with respect to their definitions for online indistinguishability, this made for a stronger security notion [11]. The same is true for us; it is easy to see that if the adversary were restricted to asking a sequence of messages with nondecreasing indexes, a restriction that amounts to forbidding the interleaving of encryptions, the resulting security notion would be properly weaker.

We do not find it necessary to demand that, once an oracle query $(i, \cdot, 1)$ is made, there are no subsequent queries (i, \cdot, \cdot) . Nonetheless, this is the expected behavior, as the setting of $\delta = 1$ is meant to indicate that the message is complete.

5 Online Security of the CBC-CS Schemes

DELAYED CBC. We now present an online version of CBC mode. For the moment, assume all messages have a multiple of b bits. The most obvious approach for defining an online version of CBC is to just spit out ciphertext blocks as they are formed. But this does not work: if an adversary knows C_{i-1} it can choose P_i such that $C_{i-1} \oplus P_i = P_j \oplus C_{j-1}$ for some $j < i$, whence C_i will be C_j if the adversary has a "real" encryption oracle, while this is unlikely if the adversary has a "fake" encryption oracle. We can defend against this attack and, more broadly, get online-secure scheme, simply by *delaying* the last ciphertext block from each plaintext chunk, holding onto it until the relevant blockcipher has already been made. The idea is due to Fouque, Martinet, and Poupard [11]. The contents of this section are a strengthening and extension of that work, adding ciphertext stealing, employing less restrictive syntax, and establishing a stronger notion of security.

The algorithm, detailed in Fig. 3, is called *delayed CBC*, or DCBC. The state consists of two parts: a *pending ciphertext block*, which initially contains a randomly generated IV , and *unprocessed plaintext*, a partial block, possibly empty,

```

30  algorithm DCBC $_{K}^{V, \delta}(P)$ 
31  if  $|V| < b$  then return error
32   $C_0 P_0 \leftarrow V$  where  $|C_0| = b$ 
33   $P \leftarrow P_0 P$ ;  $n \leftarrow \lfloor |P|/b \rfloor$ 
34   $P_1 \dots P_n P^* \leftarrow P$  where  $|P_1| = \dots = |P_n| = b$ 
35  if  $\delta = 1$  and  $P^* \neq \varepsilon$  then return error
36  for  $i \leftarrow 1$  to  $n$  do  $C_i \leftarrow E_K(P_i \oplus C_{i-1})$ 
37  if  $\delta = 0$  then  $(C, V') \leftarrow (C_0 \dots C_{n-1}, C_n P^*)$ 
38  if  $\delta = 1$  then  $(C, V') \leftarrow (C_0 \dots C_n, \varepsilon)$ 
39  return  $(C, V')$ 

```

Fig. 3. Mode DCBC. An online encryption scheme, encryption now depends on the saved state $V \in \{0, 1\}^*$. The first b bits of V comprise the *pending ciphertext*, C_0 , while the remaining 0 to $b-1$ bits are *unprocessed plaintext*, P_0 . Bit δ signals if the plaintext is over.

carried over from the previous message chunk. If the blockcipher acts on b bits then the state will be at most $v = 2b-1$ bits. In the pseudocode of Fig. 3, regard $C_i \dots C_j$ as the empty string if $i > j$.

Informally, the algorithm of Fig. 3 proceeds as follows. The algorithm receives a key K , a state V , an end-of-message indicator δ , and a plaintext chunk P . It parses the state into a b -bit delayed ciphertext block C_0 , and a partial plaintext block P_0 with $0 \leq |P_0| < b$. The algorithm then adjusts the incoming plaintext chunk P by prefixing it with P_0 . Next it splits P into b -bit blocks P_1, \dots, P_n , leaving a leftover and possibly empty partial block P^* . Since DCBC can only cope with messages that are an integral number of blocks long, the algorithm fails (it reports an error) if $P^* \neq \varepsilon$ when $\delta = 1$. The algorithm next performs the CBC encryption: for $1 \leq i \leq n$, set $C_i = E_K(P_i \oplus C_{i-1})$. Finally, if $\delta = 1$ the algorithm outputs all of $C = C_0 \dots C_n$, and clears the state. If $\delta = 0$ then it outputs $C = C_0 \dots C_{n-1}$, holding $V' = C_n \parallel P^*$ in the revised state.

ONLINE SECURITY OF DCBC. We now show that Delayed CBC achieves IND\$ security.

Theorem 2. *Suppose that adversary A asks queries totaling at most σ blocks (each query P contributes $\lceil |P|/b \rceil$ blocks). Then $\mathbf{Adv}_{\text{DCBC}[\text{Perm}(b)]}^{\text{IND\$}}(A) \leq \sigma^2/2^b$.*

Proof. Let $r = \mathbf{Adv}_{\text{DCBC}[\text{Func}(b)]}^{\text{IND\$}}(A)$. As in the proof of Lemma 1, the PRF/PRP switching gives us that $|\mathbf{Adv}_{\text{DCBC}[\text{Perm}(b)]}^{\text{IND\$}}(A) - r| \leq \sigma^2/2^{b+1}$. It remains to show that $r \leq \sigma^2/2^{b+1}$, for which we use the games in Fig. 4.

The games are constructed so that, with $\mathcal{E} = \text{DCBC}[\text{Func}(b)]$, we have $\Pr[A^{\text{Real}(\cdot, \cdot)} \Rightarrow 1] = \Pr[A^{G_1(\cdot, \cdot)} \Rightarrow 1]$ and $\Pr[A^{\text{Fake}(\cdot, \cdot)} \Rightarrow 1] = \Pr[A^{G_0(\cdot, \cdot)} \Rightarrow 1]$. Furthermore, games G_0 and G_1 are identical-until-*bad*, and hence we get $r = |\mathbf{Adv}[A^{G_1(\cdot, \cdot)} \Rightarrow 1] - \mathbf{Adv}[A^{G_0(\cdot, \cdot)} \Rightarrow 1]| = \Pr[A^{G_0}$ sets *bad*].

In game G_0 , all of the C_i values are uniform and independent: all except C_0 in the initial call are generated explicitly by the oracle—and that C_0 is the initial state V_j , chosen uniformly as part of the initialization.

<pre> 300 algorithm Enc(j, P, δ) 301 if $V_j < b$ then return error 302 $C_0 P_0 \leftarrow V_j$ where $C_0 = b$ 303 $P \leftarrow P_0 P; n \leftarrow \lfloor P /b \rfloor$ 304 $P_1 \dots P_n P^* \leftarrow P$ where $P_1 = \dots = P_n = b$ 305 if $\delta = 1$ and $P^* \neq \varepsilon$ then return error 306 for $i \leftarrow 1$ to n do 307 $X_i \leftarrow P_i \oplus C_{i-1}$ 308 $C_i \xleftarrow{\\$} \{0, 1\}^b$ 309 if $\rho(X_i) \neq \text{undefined}$ then $bad \leftarrow \text{true}$, $C_i \leftarrow \rho(X_i)$ 310 else $\rho(X_i) \leftarrow C_i$ 311 if $\delta = 0$ then $(C, V'_j) \leftarrow (C_0 \dots C_{n-1}, C_n P^*)$ 312 if $\delta = 1$ then $(C, V'_j) \leftarrow (C_0 \dots C_n, \varepsilon)$ 313 return C </pre>	Game G_0 Game G_1
--	---

Fig. 4. Proof of the IND $\$$ -security of Delayed CBC. Game G_1 includes the boxed statement following the setting of bad ; game G_0 omits it. The variable bad is initialized to **false**, V_j is initialized to a random b -bit string chosen uniformly at random for each $j \in \mathbb{N}$, and ρ is initialized to everywhere **undefined**.

Since the P_i are determined solely by the adversary’s inputs, we can think of them as being selected directly by the adversary. We claim that the adversary must choose each P_i before receiving any information about C_{i-1} . For $i > 1$ this is clear, since C_{i-1} is chosen uniformly at random after P_i has been determined. It remains to show that C_0 is uniformly distributed and independent of the adversary’s view until P_1 is determined. (Ensuring this property is the reason for delaying the ciphertext block.) We do this inductively, and separately for each index $j \in \mathbb{N}$. The base case is the first encryption query with index j : then $C_0 = V_j$ is the randomly selected initialization vector. Here the adversary can’t know anything about its value at this stage since it hasn’t been used in any computations at all. The state is empty and we return an immediate error if the previous call’s end-of-message indicator was set, so there is no C_0 to concern ourselves with. In the remaining case, the value of C_0 is equal to the value of C_n from the previous encryption query with the same index; the inductive step, therefore, is to show that C_n is uniform and independent of the adversary’s view if C_0 is also and $\delta = 0$. But nothing dependent on C_n is part of the oracle’s output if $\delta = 0$, and C_n is either freshly generated (if $n > 0$), or equal to C_0 and therefore uniform and independent of the adversary by the induction hypothesis (if $n = 0$).

It immediately follows that each P_i is independent of C_{i-1} , and therefore all of the X_i values are uniform and independent of one another. Hence the probability that two X_i collide—and bad is set—is at most $\sigma^2/2^{b+1}$, completing the proof.

As usual, it is easy to pass from the information-theoretic setting to complexity-theoretic one.

DELAYED CBC WITH CIPHERTEXT STEALING. The algorithms DCBC-CS1, DCBC-CS2, and DCBC-CS3 are defined in Fig. 5. Implicitly, the modes are

```

40  algorithm DCBC-CSKV, δ(P)
41  if |V| < b then return error
42  C-1C0 P0 ← V where |C-1| ∈ {0, b}, |C0| = b, |P0| < b
43  P ← P0 P
44  if δ = 0 then P1 ⋯ Pn P* ← P where n ← ⌊|P|/b⌋, |P1| = ⋯ = |Pn| = b
45  else P1 ⋯ Pn ← P 0b-d where n ← ⌈|P|/b⌋, d ← b + |P| - nb, |P1| = ⋯ = |Pn| = b
46  for i ← 1 to n do Ci ← EK(Pi ⊕ Ci-1)
47  if δ = 0 then
48-1   (C, V') ← (C0 ⋯ Cn-1, Cn P*)           ⇐ for CS1
48-2   (C, V') ← (C0 ⋯ Cn-1, Cn P*)           ⇐ for CS2
48-3   (C, V') ← (P* = ε)? (C-1C0 ⋯ Cn-2, Cn-1Cn) :
                                     (C-1C0 ⋯ Cn-1, Cn P*)           ⇐ for CS3
49  if δ = 1 then
50    if n > 0 then Cn-1 ← MSBd(Cn-1)
51-1   (C, V') ← (C0 ⋯ Cn-2Cn-1Cn, ε)           ⇐ for CS1
51-2   (C, V') ← (d = b)? (C0 ⋯ Cn-2Cn-1Cn, ε) :
                                     (C0 ⋯ Cn-2CnCn-1, ε)           ⇐ for CS2
51-3   (C, V') ← (C-1C0 ⋯ Cn-2CnCn-1, ε)           ⇐ for CS3
52  return (C, V')

```

Fig. 5. Delayed CBC with ciphertext stealing: DCBC-CS. Each online scheme depends on $E: \mathcal{K} \times \{0, 1\}^b \rightarrow \{0, 1\}^b$. String $C_{-1}C_0$ is pending ciphertext (with C_{-1} used only for DCBC-CS3). String P_0 is unprocessed plaintext from the prior call.

all parameterized by a blockcipher $E: \mathcal{K} \times \{0, 1\}^b \rightarrow \{0, 1\}^b$. The state V once again maintains two portions: the *pending ciphertext* and the *unprocessed plaintext*. The pending ciphertext is a single block—or possibly two blocks in the cases of DCBC-CS3—that the algorithm retains until it is “safe” to spit this out. This is followed by 0 to $b - 1$ bits of unprocessed plaintext. The dividing line between the two portions is always clear from the length of the string V . Note that for DCBC-CS3, the state has grown from $2b - 1$ bits to $2b$ bits while, for DCBC-CS1 and DCBC-CS2 the state remains at $2b - 1$ bits.

The IND\$ security of the DCBC-CS schemes can be inferred from the IND\$ security of the DCBC schemes. This is done in the proof below.

Theorem 3. *Let \mathcal{E} be any of DCBC-CS1[Perm(b)], DCBC-CS2[Perm(b)], or DCBC-CS3[Perm(b)], and suppose A asks queries totaling at most σ blocks. Then $\text{Adv}_{\mathcal{E}}^{\text{IND}\$}(A) \leq \sigma^2/2^b$.*

Proof. We use a different description of DCBC-CS, shown in Fig. 6, now writing the algorithm in terms of DCBC. The state vector consists of three components: a state W for DCBC, which is not interpreted; an additional delayed ciphertext block C_{-1} , which corresponds to C_{-1} in Fig. 5; and a length $0 \leq \ell < b$, which keeps track of the amount of unprocessed plaintext maintained in W , so that $\ell = |P_0|$.

The theorem will follow from three observations about this new description of DCBC. First, *DCBC-CS* is functionally identical to DCBC-CS. Second, if the


```

400 algorithm  $DCBC-CS_K^{V, \delta}(P)$ 
401 if  $|V| \leq b$  then  $(W, C_{-1}, \ell) \leftarrow (V, \varepsilon, 0)$  else  $[W, C_{-1}, \ell] \leftarrow V$ 
402  $m \leftarrow \ell + |P|$ ,  $d \leftarrow m - b \lfloor (m-1)/b \rfloor$ 
403 if  $\delta = 1$  then  $P \leftarrow P 0^{b-d}$ 
404  $(C, W') \leftarrow DCBC_K^{W, \delta}(P)$ 
405  $C_0 \cdots C_n \leftarrow C$  where  $n \leftarrow |C|/b - 1$  and  $|C_0| = \cdots = |C_n| = b$ 
406 if  $\delta = 0$  then
407-1  $(C', C'_{-1}) \leftarrow (C_0 \cdots C_n, \varepsilon)$   $\Leftarrow$  for CS1
407-2  $(C', C'_{-1}) \leftarrow (C_0 \cdots C_n, \varepsilon)$   $\Leftarrow$  for CS2
407-3  $(C', C'_{-1}) \leftarrow (d=b)? (C_0 \cdots C_{n-1}, C_n) : (C_0 \cdots C_n, \varepsilon)$   $\Leftarrow$  for CS3
408  $\ell' \leftarrow (d=b)? 0 : d$ 
409 if  $\delta = 1$  then
410 if  $n > 0$  then  $C_{n-1} \leftarrow MSB_d(C_{n-1})$ 
411-1  $C' \leftarrow C_0 \cdots C_{n-2} C_{n-1} C_n$   $\Leftarrow$  for CS1
411-2  $C' \leftarrow (d=b)? C_0 \cdots C_{n-2} C_{n-1} C_n : C_0 \cdots C_{n-2} C_n C_{n-1}$   $\Leftarrow$  for CS2
411-3  $C' \leftarrow C_0 \cdots C_{n-2} C_n C_{n-1}$   $\Leftarrow$  for CS3
412  $\ell' \leftarrow 0$ ,  $C'_{-1} \leftarrow \varepsilon$ 
413 return  $(C, [W', C'_{-1}, \ell'])$ 

```

Fig. 6. Defining DCBC-CS in terms of DCBC. The notation $[x_1, \dots, x_n]$ denotes an unambiguous non-compressing encoding of the items x_1, \dots, x_n ; used on the left-hand side of an assignment, it implies a decoding operation.

call to function DCBC at line 404 were to instead call a function that returned a random strings of the appropriate length, then so too would $DCBC-CS$. This observation is immediate, since the strings returned $DCBC-CS'$ are derived from those returned by $DCBC_K^{V, \delta}$ by discarding and reordering particular fixed bits. Third, $DCBC-CS$ can be implemented using only oracle access to the DCBC function: it doesn't need to inspect or interpret the DCBC state vector W , nor examine the key K , and it uses the state only in the “single-threaded” way permitted by the online IND $\$$ oracle.

Consequently, for any adversary A attacking DCBC-CS, we can construct an adversary B attacking DCBC: B will run A against a simulated oracle built from B 's (real or fake) DCBC oracle using $DCBC-CS$ and, in the end, output A 's guess as its own. We have

$$\begin{aligned}
\mathbf{Adv}_{DCBC-CS}^{\text{IND}\$}(A) &= \Pr[A^{\text{Real}} \Rightarrow 1] - \Pr[A^{\text{Fake}} \Rightarrow 1] \\
&= \Pr[A^{DCBC-CS[\text{Real}]} \Rightarrow 1] - \Pr[A^{DCBC-CS[\text{Fake}]} \Rightarrow 1] \\
&= \Pr[B^{\text{Real}} \Rightarrow 1] - \Pr[B^{\text{Fake}} \Rightarrow 1] \\
&= \mathbf{Adv}_{\varepsilon}^{\text{IND}\$}(B) \leq \sigma^2/2^b
\end{aligned}$$

appealing to Theorem 2 for the final inequality.

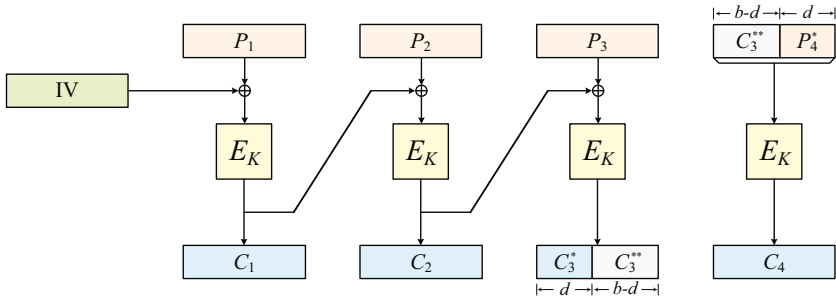
As before, one can immediately conclude the corresponding complexity-theoretic statement, which would read as follows.

Corollary 2. Let $E: \mathcal{K} \times \{0, 1\}^b \rightarrow \{0, 1\}^b$ be a blockcipher and let \mathcal{E} be any of the encryption schemes DCBC-CS1[E], DCBC-CS2[E], or DCBC-CS3[E]. Suppose A asks queries that total σ blocks, runs in time t , and achieves advantage $\delta = \text{Adv}_{\mathcal{E}}^{\text{IND}_S}(A)$. Then there is an adversary B , explicitly known and constructed from A in a blackbox manner, that asks at most σ queries, runs in time $t + \lambda\sigma$, and achieves advantage $\text{Adv}_E^{\text{PRP}}(B) \geq \delta - \sigma^2/2^b$. Here λ is an absolute constant depending only on details of the model of computation.

6 Insecurity of the Meyer-Matyas CBC-CS

The CBC ciphertext-stealing construction by Meyer and Matyas, what we will call CBC-CSX, is defined in Fig. 7. This well-known scheme—it has been used since the early 1980’s under the IBM CUSP architecture—is susceptible to a simple chosen-plaintext attack, a fact that appears not to have been pointed out before. Thus NIST did well in choosing not to standardize this form of ciphertext stealing, but the alternative, “correct” variant.

Here is an attack on the ind $\$$ -security of CBC-CSX. The adversary makes two encryption queries: $M = 1^b 0^{b-1}$ and $M' = 1^b 0^{b-1}$. As the IV is randomized, asking the same plaintext twice is not without purpose. The oracle returns $C = C_0 C_1 C_2$ and $C' = C'_0 C'_1 C'_2$ where C_0 and C'_0 are randomly chosen IVs and $|C_1| = |C'_1| = b - 1$. If $C_2 = C'_2$ the adversary returns 1; otherwise, it returns 0.



```

90 algorithm CBC-CSXIVK(P)
91   n ← ⌈|P|/b⌉
92   P1 ⋯ Pn-1 Pn* ← P where |P1| = ⋯ = |Pn-1| = b and |Pn*| = d
93   C0 ← IV
94   for i ← 1 to n - 1 do Ci ← EK(Pi ⊕ Ci-1)
95   Cn ← EK((LSBb-d(Cn-1) || Pn*)
96   Cn-1* ← MSBd(Cn-1)
97   return C0 C1 ⋯ Cn-2 Cn-1* Cn

```

Fig. 7. Mode CBC-CSX. The mode is insecure and should not be used. This version of ciphertext stealing is from Meyer and Matyas [14]. The mode depends on a blockcipher $E: \mathcal{K} \times \{0, 1\}^b \rightarrow \{0, 1\}^b$. That can be ideal, and the IV random, and still the mode will fail to achieve standard (CPA) privacy definitions.

Now if the adversary is given a CBC-CSX oracle, the probability that $C_2 = C'_2$ is at least $1/2$; otherwise, it's about $1/2^b$. Thus we have a trivial but effective ind $\$$ -attack.

We remark that, not surprisingly, CBC-CSX is not secure under conventional, weaker notions of security, like left-or-right indistinguishability [3]; a similar attack can easily be described. It is not that the definition is too strong; from a modern point of view, the scheme is simply wrong.

Acknowledgments. Authors Rogaway and Zhang received support for this project under NSF grant CNS 0904380. Many thanks to the NSF for their support.

References

1. Ball, M.: Follow-up to NIST's consideration of XTS-AES as standardized by IEEE Std 1619-2007. Public comments to NIST (2008), <http://tinyurl.com/nist-ball-xts>
2. Bard, G.V.: Blockwise-Adaptive Chosen-Plaintext Attack and Online Modes of Encryption. In: Galbraith, S.D. (ed.) *Cryptography and Coding 2007*. LNCS, vol. 4887, pp. 129–151. Springer, Heidelberg (2007)
3. Bellare, M., Desai, A., Jokipii, E., Rogaway, P.: A concrete security treatment of symmetric encryption: analysis of the DES modes of operation. In: *FOCS 1997*, pp. 394–403. IEEE Press (1997)
4. Bellare, M., Kohno, T., Namprempre, C.: Breaking and provably repairing the SSH authenticated encryption scheme: a case study of the encode-then-encrypt-and-MAC paradigm. *ACM Transactions on Information and System Security (TISSEC)* 7(2), 206–241 (2004); Earlier version from *CCS 2002*
5. Bellare, M., Rogaway, P.: The Security of Triple Encryption and a Framework for Code-Based Game-Playing Proofs. In: Vaudenay, S. (ed.) *EUROCRYPT 2006*. LNCS, vol. 4004, pp. 409–426. Springer, Heidelberg (2006)
6. Boldyreva, A., Taesombut, N.: Online encryption schemes: New security notions and constructions. In: Okamoto, T. (ed.) *CT-RSA 2004*. LNCS, vol. 2964, pp. 1–14. Springer, Heidelberg (2004)
7. Dworkin, M.: Recommendation for block cipher modes of operation: method and techniques. NIST Special Publication 800-38A, 2001 Edition (December 2001)
8. Dworkin, M.: Recommendation for block cipher modes of operation: three variants of ciphertext stealing for CBC mode. Addendum to NIST Special Publication 800-38A (October 2010)
9. Fouque, P., Joux, A., Martinet, G., Valette, F.: Authenticated On-line Encryption. In: Matsui, M., Zuccherato, R.J. (eds.) *SAC 2003*. LNCS, vol. 3006, pp. 145–159. Springer, Heidelberg (2004)
10. Fouque, P., Joux, A., Poupard, G.: Blockwise Adversarial Model for On-line Ciphers and Symmetric Encryption Schemes. In: Handschuh, H., Hasan, M.A. (eds.) *SAC 2004*. LNCS, vol. 3357, pp. 212–226. Springer, Heidelberg (2004)
11. Fouque, P., Martinet, G., Poupard, G.: Practical Symmetric On-Line Encryption. In: Johansson, T. (ed.) *FSE 2003*. LNCS, vol. 2887, pp. 362–375. Springer, Heidelberg (2003)
12. Gennaro, R., Rohatgi, P.: How to Sign Digital Streams. In: Kaliski Jr., B.S. (ed.) *CRYPTO 1997*. LNCS, vol. 1294, pp. 180–197. Springer, Heidelberg (1997)

13. Joux, A., Martinet, G., Valette, F.: Blockwise-Adaptive Attackers Revisiting the (In)Security of Some Provably Secure Encryption Modes: CBC, GEM, IACBC. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 17–30. Springer, Heidelberg (2002)
14. Meyer, C., Matyas, M.: Cryptography: a new dimension in data security. John Wiley & Sons, New York (1982)
15. NIST. Proposal to extend CBC mode by “ciphertext stealing.” Anonymous draft (May 6, 2007), Available from NIST’s website
16. Rogaway, P., Bellare, M., Black, J.: OCB: A block-cipher mode of operation for efficient authenticated encryption. *ACM Transactions on Information and System Security* 6(3), 365–403 (2003); Earlier version, with Krovetz, T.: ACM CCS 2001
17. Schneier, B.: *Applied Cryptography: Protocols, Algorithms, and Source Code in C*, 2nd edn. Wiley, New York (1996)
18. Shoup, V.: Sequences of games: a tool for taming complexity. ePrint archive 2004/332 Revised (2006)
19. Vaudenay, S.: Security Flaws Induced by CBC Padding – Applications to SSL, IPSEC, WTLS.. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 534–545. Springer, Heidelberg (2002)

McOE: A Family of Almost Foolproof On-Line Authenticated Encryption Schemes

Ewan Fleischmann, Christian Forler, and Stefan Lucks

Bauhaus-University Weimar, Germany

{ewan.fleischmann,christian.forler,stefan.lucks}@uni-weimar.de

Abstract. On-Line Authenticated Encryption (OAE) combines privacy with data integrity and is on-line computable. Most block cipher-based schemes for Authenticated Encryption can be run on-line and are provably secure against *nonce-respecting* adversaries. But they fail badly for more general adversaries. This is not a theoretical observation only – in practice, the reuse of nonces is a frequent issue¹.

In recent years, cryptographers developed *misuse-resistant* schemes for Authenticated Encryption. These guarantee excellent security even against general adversaries which are allowed to reuse nonces. Their disadvantage is that encryption can be performed in an off-line way, only.

This paper considers OAE schemes dealing both with nonce-respecting and with general adversaries. It introduces MCOE, an efficient design for OAE schemes. For this we present in detail one of the family members, MCOE-X, which is a design solely based on a standard block cipher. As all the other member of the MCOE family, it provably guarantees reasonable security against general adversaries as well as standard security against nonce-respecting adversaries.

Keywords: authenticated encryption, on-line encryption, provable security, misuse resistant.

1 Introduction

On-Line Authenticated Encryption (OAE). Application software often requires a network channel that guarantees the privacy and authenticity of data being communicated between two parties. Cryptographic schemes able to meet both of these goals are commonly referred to as Authenticated Encryption (AE) schemes. The ISO/IEC 19772:2009 standard for AE [21] defines generic composition (Encrypt-then-MAC [4]) and five dedicated AE schemes: OCB2 [38], SIV [41] (denoted as “Key Wrap” in [21]), CCM [13], EAX [6], and GCM [34]. To integrate an AE-secure channel most seamlessly into a typical software architecture, application developers expect it to encrypt in an *on-line* manner meaning that the i -th ciphertext block can be written before the $(i + 1)$ -th plaintext block

¹ A prominent example is the PlayStation 3 ‘jailbreak’ [20], where application developers used a constant that was actually supposed to be a nonce for a digital signature scheme.

has to be read. A restriction to off-line encryption, where usually the entire plaintext must be known in advance (or read more than once) is an encumbrance to software architects.

Nonces and their reuse. Goldwasser and Micali [18] formalized encryption schemes as stateful or probabilistic, because otherwise important security properties are lost. Rogaway [37,39,40] proposed an unified point of view, by always defining a cryptographic scheme as a deterministic algorithm that takes an user supplied nonce (a *number used once*). So the application programmer – and not the encryption scheme – is responsible for flipping coins or maintaining state. This reflects cryptographic practice since the algorithm itself is often implemented by a multi-purpose cryptographic library which is more or less application-agnostic.

In theory, the concept of a nonce is simple. In practice, it is challenging to ensure that a nonce is *never* reused. Flawed implementations of nonces are ubiquitous [9,20,28,44,45]. Apart from implementation failures, there are fundamental reasons why software developers can't always prevent nonce reuse. A persistently stored counter, which is increased and written back each time a new nonce is needed, may be reseted by a backup – usually after some previous data loss. Similarly, the internal and persistent state of an application may be duplicated when a virtual machine is cloned, etc.

Related Work and Our Contribution. We aim to achieve *both simultaneously*: security against nonce-reusing adversaries (sometimes also called nonce-misusing adversaries) *and* support for on-line-encryption in terms of an AE scheme. Apart from generic composition (Encrypt-then-Mac, EtM), none of the ISO/IEC 19772:2009 schemes – in fact, no previously published AE scheme at all – achieves both of these goals, cf. Table 1. In this table, we classify a vast variety of provably secure block cipher-based AE scheme with respect to their on-line-ability and against which adversaries (nonce-respecting versus -reusing) they are proven secure.

Since EtM is not a concrete scheme but merely a generic construction technique, there are some challenges left in order to make it full on-line secure: First, an appropriate on-line cipher has to be chosen. Second, a suitable, on-line computable, secure deterministic MAC must be selected. And, third, the EtM scheme requires at least two *independent* keys to be secure. Since two schemes are used in parallel, is likely to squander resources in terms of run time and – important for hardware designers – in terms of space. Since EtM first has to be turned into an OAE scheme by making the appropriate choices, we don't include it in our analysis.

As it turned out, we actually found nonce-reuse attacks for *all* of those schemes, cf. Table 2, Appendix A, and, especially, Appendix 1 in the full version of this paper [14]. We present a new construction method for efficient AE schemes, called MCOE-X, that is actually able to fill the apparent gap in the upper-right. It belongs to the family of MCOE schemes [14]. We argue that closing this gap is both practically relevant and theoretically interesting.

Table 1. Classification of provably secure block cipher-based AE Schemes. CCM and SSH-CTR are considered off-line because encryption requires prior knowledge of the message length. Note that the family of McOE schemes, because of being on-line, satisfies a slightly weaker security definition against nonce-reusing adversaries than SIV, HBS, and BTM.

secure ...	against nonce-respecting adversaries	ag. nonce-reusing adversaries
on-line	CCFB[33] CHM[22] CIP[23] CWC[29] EAX[6] GCM[34] IACBC[26] IAPM[26] McOE OCB1-3[40,38,30] RPC[10] TAE[31] XCBC[17]	McOE (this paper)
off-line	BTM[24] CCM[13] HBS[25] SIV[41] SSH-CTR[36]	BTM[24] HBS[25] SIV[41]

Table 2. Overview of our **nonce-reuse** attacks on published AE schemes, excluding SIV, HBS and BTM, which have been explicitly designed to resist nonce-reuse. Almost all attacks achieve an advantage close to 1. An “attack workload” of X means that the adversary is restricted to at most X units of time and at most X chosen texts. Details are given in Appendix A and in the full version of this paper [14].

	privacy attack workload	authenticity attack workload		privacy attack workload	authenticity attack workload
CCFB [33]	$O(1)$	$O(1)$	IAPM [26]	$O(1)$	$O(1)$
CCM [13]	$O(1)$	$\ll 2^{(n/2)}$ [15]	OCB1 [40]	$O(1)$	$O(1)$
CHM [22]	$O(1)$	$O(1)$	OCB2 [38]	$O(1)$	$O(1)$
CIP [23]	$O(1)$	$O(1)$	OCB3 [30]	$O(1)$	$O(1)$
CWC [29]	$O(1)$	$O(1)$	RPC [10]	$O(1)$	$O(1)$
EAX [6]	$O(1)$	$O(1)$	TAE [31]	$O(1)$	$O(1)$
GCM [34]	$O(1)$	$O(1)$	XCBC [17]	$O(2^{n/4})$?
IACBC [26]	$O(1)$	$O(1)$			

Initial Value (IV) based AE schemes maximally forgiving of repeated IV’s have been addressed in [41], coining the notion of “misuse resistance” and proposing SIV as a solution. SIV and related schemes (HBS [25] and BTM [24]) actually provide excellent security against nonce-reusing adversaries, though there are other potential misuse cases, cf. the Appendix of the full version of this paper [14]. Their main disadvantage is that they are inherently off-line: For encryption, one must either keep the entire plaintext in memory, or read the plaintext twice.

Ideally, an adversary seeing the encryptions of two (equal-length) plaintexts P_1 and P_2 can’t even decide if $P_1 = P_2$ or not. When using a nonce more than once, deciding about $P_1 = P_2$ is easy. SIV and its relatives ensure that nothing else is feasible for nonce-reusing adversaries. In the case of on-line encryption, where the first few bits of the encryption of a lengthy message must not depend on the last few bits of that message, there is unavoidably something beyond $P_1 = P_2$. The adversary can compare any two ciphertexts for their longest common prefix, and then conclude about common prefixes of the secret plaintexts. Our notion

of *misuse resistance* means that this is all the adversary can gain. Even in the case of a nonce-reuse, the adversary

1. can't do anything beyond determining the length of common plaintext prefixes and
2. the scheme still provides the usual level of authenticity for AE (INT-CTXT).

The first property is common for on-line ciphers/permutations (OPRP) [1]. Recently, [43] studied the design of on-line ciphers from tweakable block ciphers bearing some similarities to our approach, especially to TC3. In contrast to the MCOE family, the constructions from [43] provide no authentication. The MCOE schemes are, *e.g.*, based on a normal block cipher *or* a tweakable block cipher.

Design Principles for AE Schemes. The question how to provide authenticated encryption (without stating that name) when given a secure on-line cipher is studied in [3], the revised and full version of [1]. The first idea in [3] only provides security if all messages are of the same length. The second idea repairs that by prepending the message's length to the message, at the cost of being off-line, since the message length must be known at the beginning of the encryption process. The third idea is to prepend and append a random W to a message M and then to perform the on-line encryption of $(W||M||W)$. This looks promising, but the same W is used for two different purposes, putting different constraints on the generation of W . For privacy, it suffices that W behaves like a nonce, not requiring secrecy or unpredictability. Even if W is not a nonce, but the same W is used for the encryption of several messages, all the adversary can determine are the lengths of common plaintexts prefixes, as we required for nonce-reuse. On the other hand, authenticity actually assumes a *secret or unpredictable* W , rather than a nonce. If the adversary can guess W before choosing a message, she asks for the authenticated encryption of $(M||W)$. Then she can predict the authenticated encryption of M without actually asking for it.

The MCOE family replaces the "random" W by a proper nonce and a value τ which is *key-dependent*, performing a nonce-dependent on-line encryption of $(M||\tau)$. The encryption can also depend on some associated data, which turns MCOE into a family of schemes for OAEAD (*On-Line Authenticated Encryption with Associated Data*).

Roadmap. In this paper we focus on one member of the MCOE [14] family of schemes called MCOE-X. In Section 2 we describe a concrete block cipher based OAE scheme – called MCOE-X– and provide performance data when MCOE-X is instantiated with either AES-128 or Threefish-512 as the underlying block cipher. Section 3 deals with general notions and definitions, and Section 4 defines the security of OAE. The main result of the paper, the full MCOE-X scheme and its analysis, is presented in Section 5. The discussion in Section 6 concludes the paper. The appendix deals with misuse attacks against published AE schemes.

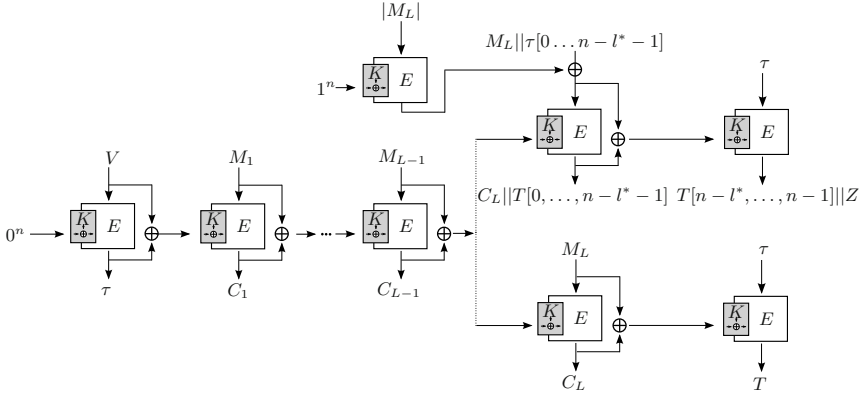


Fig. 1. The MCOE-X-AES/MCOE-X-Threefish encryption process. If, after the last complete message block has been encrypted, there is some incomplete block left, MCOE-X performs tag-splitting (upper variant), Else, the tag can be computed without splitting (lower variant). The key used for the block cipher E is computed by the injective function $K \oplus W$ which is given the secret key K and the chaining value input W . The tag returned is the n -bit value T . The $n - l$ -bit value Z is discarded. The decryption process works in a similar way from 'left to right' only the block cipher component E is replaced by its counterpart E^{-1} apart from one exception: the first call computing τ .

2 Practical On-Line Authenticated Encryption Using AES and Threefish

We start with the fruits of our analysis by giving two concrete instances of OAE schemes (illustrated in Figure 1) including performance data and reference source code². One instance, MCOE-X-AES uses AES-128 as the core component while MCOE-X-Threefish uses the block cipher Threefish-512, a cipher with 512-bit block size and key size, which is the core working component inside the SHA-3 finalist Skein[35]. We also introduce the *tag-splitting* (TS) method for processing messages whose length is not a multiple of the block length. Without TS, we would have to pad such messages and then encrypt the padded messages – resulting in an expanded ciphertext. The effect of TS is similar to the well-known length preserving method called *ciphertext stealing* (CTS), *e.g.* [12]. But the technique itself is quite different since CTS requires to process the last block before the last but one, which is not possible for MCOE-X.

Let E_K be a block cipher taking a k -bit key K and a plaintext/ciphertext of size n -bit. Note that for our chosen instances, AES-128 and Threefish-512, we have $n = k$. The pseudo code for these two MCOE-X instances is given in Table 4 – on the upper side without TS, on the lower side with TS.

² The reference source code is available on request; it will be published as open source.

Table 3. Performance values (cycles-per-byte, single core), measured on an Core i5 540M for AES-128 and Threefish-512. MCOE-X is the main contribution in the current paper, MCOE-D invokes the underlying block cipher twice and MCOE-G uses Galois field arithmetic. For a comparison, we also provide the performance of unauthenticated AES-CBC. The AES software implementation is based on Gladman [16], whereas the hardware implementation is based on the Intel AES-NI Sample Library[11]. The Threefish implementation is based on the NIST/SHA-3 reference source as provided by the Skein authors [35]. Finally, the implementation of Galois field NI multiplication (GF-NI) is based on the example-code from [19].

Block cipher	Impl.	Message length in Bytes							
		64	256	512	1024	2048	8192	32768	
MCOE-X-AES	software	31.2	23.9	22.7	22	21.7	21.5	21.5	
MCOE-X-AES	AES-NI	14.2	11.2	10.7	10.5	10.4	10.3	10.3	
MCOE-X-Threefish	software	19.5	9.9	8.3	7.5	7.1	6.8	6.7	
MCOE-D-AES	software	40.1	29.4	27.6	26.7	26.3	25.9	25.9	
MCOE-D-AES	AES-NI	11.6	8.3	7.2	6.7	6.4	6.3	6.2	
MCOE-G-AES	software	33	25.4	24.1	23.5	23.2	22.9	22.8	
MCOE-G-AES	GF-NI/AES-NI	12.5	9.7	9.3	9	8.9	8.8	8.8	
AES-CBC encryption	software	38.3	13.5	13.3	13.2	13.2	13.1	13.1	
AES-CBC encryption	AES-NI	4	3.6	3.5	3.5	3.5	3.5	3.5	

The algorithms without TS, **EncryptAuthenticate** and **DecryptAuthenticate**, are simplified algorithms for messages that are aligned on n -bit boundaries, *i.e.* $M = (M_1, \dots, M_L) \in (\{0, 1\}^n)^L$ for some integer L . The TS algorithms are **EncryptAuthenticateSplitTag** and **DecryptAuthenticateSplitTag**. they can handle arbitrarily sized messages, *i.e.*, $M = (M_1, \dots, M_L) \in (\{0, 1\}^n)^{L-1} || \{0, 1\}^{l^*}$ where L and l^* are integers with $0 < l^* < n$ and $||$ denotes the string concatenation operator. See Figure 1 and Table 4.

In addition to MCOE-X, we introduce two further authenticated encryption schemes following the MCOE design principles. The first one is called MCOE-D and is based on the THC-CBC construction [7]. The ratio of this scheme is 2-1, *i.e.* the block cipher is invoked twice to encipher resp. decipher one message block. The second one is called MCOE-G and is based on the HCBC-2 construction [2]. This scheme updates the chaining value by invoking a universal hash function, *i.e.*, a n -bit Galois-Field multiplication.

Remarks. For MCOE-X we actually do need related key resistance for the block cipher E since the adversary can 'partially control' some relations among keys used in the computation. This is not true for the other mentioned constructions.

All MCOE schemes are easily extended to smoothly handle associated data, *i.e.* data that is not encrypted but only authenticated. This is discussed in more detail in Section 5.

Table 4. Instances of MCOE-X: upper side is for messages whose size is evenly divisible by the block size n ; Lower side is for arbitrarily sized messages (TS-variant); see text for details

<p>EncryptAuthenticate(V, M)</p> <ol style="list-style-type: none"> 1. $\tau \leftarrow E_K(V)$ 2. $U \leftarrow V \oplus \tau \oplus K$ 3. for $i = 1, \dots, L$ loop $C_i \leftarrow E_U(M_i)$ $U \leftarrow M_i \oplus C_i \oplus K$ 4. $T \leftarrow E_U(\tau)$ 5. return (C_1, \dots, C_L, T) 	<p>DecryptAuthenticate(V, C, T)</p> <ol style="list-style-type: none"> 1. $\tau \leftarrow E_K(V)$ 2. $U \leftarrow V \oplus \tau \oplus K$ 3. for $i = 1, \dots, L$ loop $M_i \leftarrow E_U^{-1}(C_i)$ $U \leftarrow M_i \oplus C_i \oplus K$ 4. if $T = E_U(\tau)$ then return (M_1, \dots, M_L) else return \perp
<p>EncryptAuthenticateSplitTag(V, M)</p> <ol style="list-style-type: none"> 1. $\tau \leftarrow E_K(V)$ 2. $U \leftarrow V \oplus \tau \oplus K$ 3. for $i = 1, \dots, L - 1$ loop $C_i \leftarrow E_U(M_i)$ $U \leftarrow M_i \oplus C_i \oplus K$ 4. $M^* \leftarrow (M_L \tau[0 \dots n - l^* - 1])$ 5. $M^* \leftarrow M^* \oplus E_{K \oplus 1^n}(M_L)$ 6. $C^* \leftarrow E_U(M^*)$ 7. Parse $C_L T[0 \dots n - l^* - 1] \leftarrow C^*$ 8. $U \leftarrow M^* \oplus C^* \oplus K$ 9. $C^{**} \leftarrow E_U(\tau)$ 10. $T[n - l^* \dots n - 1] \leftarrow C^{**}[0 \dots l^* - 1]$ 11. return $(C_1, \dots, C_{L-1}, C_L^*, T)$ 	<p>DecryptAuthenticateSplitTag(V, C, T)</p> <ol style="list-style-type: none"> 1. $\tau \leftarrow E_K(V)$ 2. $U \leftarrow V \oplus \tau \oplus K$ 3. for $i = 1, \dots, L - 1$ loop $M_i \leftarrow E_U^{-1}(C_i)$ $U \leftarrow M_i \oplus C_i \oplus K$ 4. $C^* \leftarrow C_L T[0 \dots n - l^* - 1]$ 5. $M^* \leftarrow E_U^{-1}(C^*)$ 6. $U \leftarrow M^* \oplus C^* \oplus K$ 7. $M^* \leftarrow M^* \oplus E_{K \oplus 1^n}(C_L)$ 8. Parse $M_L \tau'[0 \dots n - l^* - 1] \leftarrow M^*$ 9. $T' \leftarrow E_U(\tau)$ 10. if $\tau'[0 \dots n - l^* - 1] = \tau[0 \dots n - l^* - 1]$ and $T'[0 \dots l^* - 1] = T[n - l^* \dots n - 1]$ then return (M_1, \dots, M_L) else return \perp

3 On-Line Authenticated Encryption and Related Notions

Length of Longest Common Prefix (LLCP $_n$). The length of a string $x \in \{0, 1\}^n$ is denoted by $|x| := n$. For integers $n, \ell, d \geq 1$, set $D_n^d = (\{0, 1\}^n)^d$, and $D_n^* := \bigcup_{d \geq 0} D_n^d$, and $D_{\ell, n} = \bigcup_{0 \leq d \leq \ell} D_n^d$. Note that D_n^0 only contains the empty string. For $M \in D_n^d$, we write $M = (M_1, \dots, M_d)$ with $M_1, \dots, M_d \in D_n$. For $P, R \in D_n^*$, say, $P \in D_n^p$ and $R \in D_n^r$, we define the *length of the longest common n -prefix* of P and R as

$$\text{LLCP}_n(P, R) = \max_i \{P_i = R_i, \dots, P_i = R_i\}.$$

Let \mathcal{Q} a non-empty set of strings in D_n^* . Then we define $\text{LLCP}_n(\mathcal{Q}, P)$ as $\max_{q \in \mathcal{Q}} \{\text{LLCP}_n(q, P)\}$, e.g., if $P \in \mathcal{Q}$, then $\text{LLCP}_n(\mathcal{Q}, P) = |P|/n$.

For convenience, we introduce a notation for a *restriction on a set*. If $\mathcal{Q} = \{0, 1\}^a \times \{0, 1\}^b \times \{0, 1\}^c$, we write $\mathcal{Q}_{|b,c} = \{(B, C) \mid \exists A : (A, B, C) \in \mathcal{Q}\}$. This generalizes in the obvious way.

3.1 Block Ciphers and On-Line Permutations

Block Ciphers. An (k, n) block cipher is a keyed family of permutations consisting of two paired algorithms $E : \{0, 1\}^k \times D_n \rightarrow D_n$ and $E^{-1} : \{0, 1\}^k \times D_n \rightarrow D_n$, accepting a k -bit key and an input from D_n for some $k, n > 0$. For $n > 0$, $Block(k, n)$ is the set of all (k, n) block ciphers. For any $E \in Block(k, n)$ and a fixed key $K \in \{0, 1\}^k$, the decryption $E_K^{-1}(Y) := E^{-1}(K, Y)$ is the inverse function of encryption $E_K(X) := E(K, X)$, so that $E_K^{-1}(E_K(X)) = X$ holds for any $X \in D_n$. We follow the usual convention to write oracles, that are provided to an algorithm, as superscripts. We define the related key PRP-security of a block cipher E by the success probability of an adversary trying to differentiate between the block cipher and a random permutation.

Definition 1. Let $E \in Block(k, n)$ and denote by E^{-1} the corresponding inverse. Let $\varphi : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^k$. A fixed related key adversary A has access to an E oracle with two parameters such that she can query either $E_{\varphi(K, \cdot)}(\cdot)$ or its inverse. Let $PERM(n, n)$ be the set of n -bit permutations such that the first parameter models the permutation and the second parameter the value that is to be permuted, i.e. for $\pi \in PERM(n, n)$ it holds that $\pi(Z, \cdot)$ is a random permutation for any given value of Z . The related-key (RK) advantage [32] of A in breaking E is then defined as

$$\begin{aligned} Adv_E^{RK-CPA-PRP}(A) &= |\Pr[K \xleftarrow{\$} \{0, 1\}^k : A^{E_{\varphi(K, \cdot)}(\cdot)} \Rightarrow 1] \\ &\quad - \Pr[\pi \xleftarrow{\$} Perm(n, n) : A^{\pi(\cdot, \cdot)} \Rightarrow 1]| \\ Adv_{E, E^{-1}}^{RK-CCA-PRP}(A) &= |\Pr[K \xleftarrow{\$} \{0, 1\}^k : A^{E_{\varphi(K, \cdot)}(\cdot), E_{\varphi(K, \cdot)}^{-1}(\cdot)} \Rightarrow 1] \\ &\quad - \Pr[\pi \xleftarrow{\$} Perm(n, n) : A^{\pi(\cdot, \cdot), \pi^{-1}(\cdot, \cdot)} \Rightarrow 1]|. \end{aligned}$$

On-Line Permutations. We aim for larger permutations that not only permute single blocks but can handle multiple/variable block messages. Such a permutation, from D_n^* to D_n^* , is (n) -on-line if the i -th block of the output is determined completely by the first i blocks of the input.

Definition 2. Let $n, k \geq 0$, $K \in \{0, 1\}^k$, $V \in D_n$. A function $\Pi : \{0, 1\}^k \times D_n^* \rightarrow D_n^*$ is an (n) -on-line permutation if for any fixed K, V the function $\Pi(K, V, \cdot)$ is a permutation and there exists for any message $M = (M_1, \dots, M_m)$ a family of functions $\tilde{\pi}^i : \{0, 1\}^k \times \{0, 1\}^n \times D_n^i \rightarrow D_n$, $i = 1, \dots, m$ such that

$$\begin{aligned} \Pi(K, V, M) &= \tilde{\pi}_K^1(V, M_1) || \tilde{\pi}_K^2(V, M[1..2]) \\ &\quad || \dots || \\ &\quad \tilde{\pi}_K^{m-1}(V, M[1..m-1]) || \tilde{\pi}_K^m(V, M[1..m]), \end{aligned}$$

where $M[a \dots b] := M_a || M_{a+1} || \dots || M_b$ with “||” being the concatenation of strings, holds.

An encryption scheme is (n) -on-line if the encryption function is (n) -on-line. A thorough discussion of on-line encryption and its properties can be found in [1].

3.2 Authenticated Encryption (With Associated Data)

An authenticated encryption scheme is a tuple $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$. Its aim is to provide privacy and data integrity. The key generation function \mathcal{K} takes no input and returns a randomly chosen key K from the key space, *e.g.* from $\{0, 1\}^k$. The encryption algorithm \mathcal{E} and the decryption algorithm \mathcal{D} are deterministic algorithms that map values from $\{0, 1\}^k \times \mathcal{H} \times D_n^*$ to a string or – if the input is invalid – the value \perp . The header \mathcal{H} consists either only of the initial value/nonce $V \in D_n$ (if no data is to be authenticated/checked in the encryption/decryption process) or is a combination of V and a value from D_n^* . So $\mathcal{H} \subset D_n^+$ in either case. For sake of convenience, we usually write $\mathcal{E}_K^H(M)$ for $\mathcal{E}(K, H, M)$ and $\mathcal{D}_K^H(M)$ for $\mathcal{D}(K, H, M)$, where the message M is chosen from D_n^* , $H \in \mathcal{H}$ and a key from the key space. We require $\mathcal{D}_K^H(\mathcal{E}_K^H(M)) = M$ for any possible K, M, H , and define the tag size for a message $M \in D_n^*$ and header $H \in \mathcal{H}$ as $\text{TAG}(H, M) := |\mathcal{E}_K^H(M)| - |M|$. We denote an authenticated encryption scheme with the requirement that the initial vector V is only used once in a *nonce based* scheme. Otherwise, we call such a scheme *deterministic*. Similarly, we call an adversary *nonce-respecting* (NR) if no nonce is used twice for any query. Otherwise, the adversary is called *nonce-ignoring* (NI).

4 Security Notions for On-Line Authenticated Encryption

Authenticated (On-Line) Encryption tries to achieve privacy and authenticity at the same time. Therefore we need security notions to handle this twofold goal. For AE, there have been notions and their relations introduced for deterministic [42] and nonce based [4,5,27,37,40] AE schemes. In order to have one convenient toolset of notions, we adopt the notion of CCA3 security suggested in [42] as a *natural strengthening* of CCA2 security.

We parameterize our definition in order to define different – but closely related – notions by explicitly stating whether we mean an on-line or off-line scheme, $\omega \in \{\text{AE}, \text{OAE}\}$ and stating the adversary behavior as either nonce-respecting or nonce-ignoring, $\nu \in \{\text{NR}, \text{NI}\}$.

Definition 3 ($\text{CCA3}(\omega, \nu)$). *Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an authenticated encryption scheme with header space \mathcal{H} and message space D_n^* , and fix an adversary A . The advantage of A breaking Π is defined as*

$$\text{Adv}_{\Pi}^{\text{CCA3}(\omega, \nu)}(A) = \left| \Pr \left[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\cdot, \cdot), \mathcal{D}_K(\cdot, \cdot)} \Rightarrow 1 \right] - \Pr \left[A^{\mathcal{S}^{\omega}(\cdot, \cdot), \perp(\cdot, \cdot)} \Rightarrow 1 \right] \right|$$

The adversary's random-bits oracle, $\mathcal{S}^{\text{AE}}(\cdot, \cdot)$ or $\mathcal{S}^{\text{OAE}}(\cdot, \cdot)$, returns on a query with header $H \in \mathcal{H}$ and plaintext $X \in D_n^*$ a random string of length $|\mathcal{E}_K(M)|$ which is either on-line or not, depending on the variable ω . The $\perp(\cdot, \cdot)$ oracle returns \perp on every input. We assume *wlog.* that the adversary A never ask a query which

Game G_{CPA} , G_{CCA3}	<pre> 10 Encrypt(H, M) 11 if ($\nu = \text{NR}$ and $V \in B$) then 12 return \perp; 13 if ($b=1$) then 14 $C \leftarrow \mathcal{E}_K(H, M)$; 15 else 16 $C \leftarrow \mathcal{S}^\omega(H, M)$; 17 $B \leftarrow B \cup \{V\}$; 18 $Q \leftarrow Q \cup \{(H, C)\}$; 19 return C; </pre>	<pre> 20 Decrypt(H, C) 21 if ($(H, C) \in Q$) then 22 return \perp; 23 if ($b=1$) then 24 $M \leftarrow \mathcal{D}_K(H, C)$; 25 else 26 $M \leftarrow \perp(H, C)$; 27 return M; </pre>
--	---	---

Fig. 2. $G_{\text{CPA}}(\omega, \nu)$ is the $\text{CPA}_{\Pi}^{(\omega, \nu)}$ -Game and $G_{\text{CCA3}}(\omega, \nu)$ the $\text{CCA3}_{\Pi}^{(\omega, \nu)}$ -Game where $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$. Game G_{CCA3} contains the code in the box while G_{CPA} does not. The oracle $\mathcal{S}^{\text{AE}}(H, M)$ returns a string of length $|M| + \text{TAG}(H, M)$, this string is on-line compatible if $\omega = \text{OAE}$. V denotes the last block of the header representing the nonce/initial value.

answer is already known. It is easy to see that we can rewrite the term given in Definition 3 as

$$\left| \Pr \left[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\cdot, \cdot), \mathcal{D}_K(\cdot, \cdot)} \Rightarrow 1 \right] - \Pr \left[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\cdot, \cdot), \perp(\cdot, \cdot)} \Rightarrow 1 \right] \right| \quad (1)$$

$$+ \Pr \left[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\cdot, \cdot), \perp(\cdot, \cdot)} \Rightarrow 1 \right] - \Pr \left[A^{\mathcal{S}^\omega(\cdot, \cdot), \perp(\cdot, \cdot)} \Rightarrow 1 \right] \Big|. \quad (2)$$

One can interpret (1) as the advantage that an adversary has on the integrity of the ciphertext and (2) as the advantage that an CPA adversary has on the privacy. Using this decomposition as a motivational starting point, we now define ciphertext integrity and what we mean by a CPA adversary on authenticated encryption schemes. From now on, our definitions are based on the game playing methodology. For example, we can restate Definition 3 using the game G_{CCA3} given in Figure 2 as

$$\text{Adv}_{\Pi}^{\text{CCA3}(\omega, \nu)}(A) = 2 \left| \Pr[A^{G_{\text{CCA3}}(\omega, \nu)} \Rightarrow 1] - 0.5 \right|.$$

We denote $\text{Adv}_{\Pi}^{\text{CCA3}(\omega, \nu)}(q, t, \ell)$ as the maximum advantage over all $\text{CCA3}(\omega, \nu)$ adversaries run in time at most t , ask a total maximum of q queries to \mathcal{E} and \mathcal{D} , and whose total query length is not more than ℓ blocks.

4.1 Privacy and Integrity Notions for Authenticated Encryption Schemes.

Similarly, we define the privacy and integrity of an authenticated (on-line) encryption scheme $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ with header space D_n^+ , message space D_n^* and tag-size function $\text{TAG}(H, M)$ as follows.

Definition 4. Let $G_{\text{CPA}}(\omega, \nu)$ be the $\text{CPA}_{\Pi}^{\omega, \nu}$ game given in Figure 2. Fix an adversary A . The advantage of A breaking Π is defined as

$$\text{Adv}_{\Pi}^{\text{CPA}(\omega, \nu)}(A) \leq 2 \left| \Pr[A^{G_{\text{CPA}}(\omega, \nu)} \Rightarrow 1] - 0.5 \right|.$$

Game $G_{INT-CTXT}$ 1 Initialize (ν) $K \leftarrow \mathcal{K}()$; 3 Finalize () 4 return win;	10 Encrypt (H, M) 11 if ($\nu = \text{NR}$ and 12 $V \in B$) then 13 return \perp ; 14 $C \leftarrow \mathcal{E}_K(H, M)$; 15 $B \leftarrow B \cup \{V\}$; 16 $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(H, C)\}$; 17 return C ;	20 Verify (H, C) 21 $M \leftarrow \mathcal{D}_K(H, C)$; 22 if ($(H, C) \notin \mathcal{Q}$ 23 and $M \neq \perp$) then 24 $\text{win} \leftarrow \text{true}$; 25 return ($M \neq \perp$);
--	---	--

Fig. 3. Game $G_{INT-CTXT}(\nu)$ is the $\text{INT-CTXT}_{\Pi}^{\omega, \nu}$ game where $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$. V denotes the last block of the header representing the nonce/initial value.

Definition 5. Let $G_{INT-CTXT}(\nu)$ be the $\text{INT-CTXT}_{\Pi}^{\nu}$ game given in Figure 3. Fix an adversary A . The advantage of A breaking Π is defined as

$$\text{Adv}_{\Pi}^{\text{INT-CTXT}(\nu)}(A) \leq \Pr[A^{G_{INT-CTXT}(\nu)} \Rightarrow 1].$$

We denote $\text{Adv}_{\Pi}^{\text{CPA}(\omega, \nu)}(q, t, \ell)$ and $\text{Adv}_{\Pi}^{\text{INT-CTXT}(\nu)}(q, t, \ell)$ as the maximum advantage over all $\text{CPA}(\omega, \nu)$ resp. $\text{INT-CTXT}(\nu)$ adversaries run in time at most t , ask a total maximum of q queries to \mathcal{E} and \mathcal{D} , and whose total query length is not more than ℓ blocks.

4.2 CCA3 Is Equal to INT-CTXT Plus CPA

We now give a generalization of Theorem 3.2 from Bellare and Namprempre [4]. It simply states the equivalence of a scheme being CCA3 secure and both INT-CTXT and CPA secure. These statements hold in the on-line and offline case.

Theorem 1. Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an authenticated encryption scheme. Fix $\omega \in \{\text{AE}, \text{OAE}\}$ and $\nu \in \{\text{NR}, \text{NI}\}$. Let A be an $\text{CCA3}(\omega, \nu)_{\Pi}$ -adversary running in time t , making q queries with a total length of at most ℓ blocks. Then there are a $\text{CPA}(\omega, \nu)$ -adversary A_p and an $\text{INT-CTXT}(\omega, \nu)$ -adversary A_c such that

$$\text{Adv}_{\Pi}^{\text{CCA3}(\omega, \nu)}(A) \leq \text{Adv}_{\Pi}^{\text{CPA}(\omega, \nu)}(A_p) + \text{Adv}_{\Pi}^{\text{INT-CTXT}(\omega, \nu)}(A_c).$$

Furthermore, A_c and A_p run in time $O(t)$ and both make at most q queries in each case.

The proof is given in the full version of this paper [14].

5 The On-Line Authenticated Encryption Scheme McOE-X

In this section, we present McOE-X, a construction for an OAE scheme. We prove that McOE-X achieves our two-fold goal. First, it guarantees a certain minimum, well defined, security against a nonce-ignoring adversary. And, second,

we show – in the full version of the paper [14] – that the complete MCOE family of OAE schemes (including MCOE-X) is fully secure against a nonce-respecting adversary.

Since we already have presented two MCOE-X instances in Section 2, we proceed by formally defining MCOE-X and giving its pseudocode. Indeed this is very similar to the results presented in Section 2, but here our definitions are slightly more general. Instead of fixing the key computation function to $K \oplus V$, where R is the chaining value and K the secret key, we here use a key derivation function $\varphi(K, R)$. By this we make sure that our proof also works for tweakable block ciphers - with K as key and R as tweak - leading to more efficient design.

Definition 6 (McOE-X). *Let $k, n \in \mathbb{N}$ with $k \geq n$, $E \in \text{Block}(k, n)$, and $\varphi : \{0, 1\}^k \times \{0, 1\}^v \rightarrow \{0, 1\}^k$ such that $\varphi(K, \cdot)$ is injective. The encryption function takes a header $H \in D_n^{L_H}$, a message M and returns a ciphertext C and a tag $T \in D_n$. The decryption function takes a header $H \in D_n^{L_H}$, a ciphertext C and a tag $T \in D_n$ and returns either a plaintext M or the fail symbol \perp .*

- (i) 'Non-TS'. *Let $M, C \in D_N^L$ for some integer L , then MCOE-X is defined by the algorithms **EncryptAuthenticate** and **DecryptAuthenticate** given in Table 5.*
- (ii) 'TS'. *Let $M, C \in D_N^L || \{0, 1\}^{l^*}$ for some integers L and l^* , $0 < l^* < n$, then MCOE-X/TS is defined by the algorithms **EncryptAuthenticateSplitTag** and **DecryptAuthenticateSplitTag** given in Table 5.*

We now proceed to show the security of MCOE-X. For this we use the results of Theorem 1 and show the INT-CTXT and RK-CPA-PRP security separately.

Theorem 2

- (i) *Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be a MCOE-X scheme as in Definition 6 (i). We further assume that the block cipher E is secure against related key attacks. Then*

$$\text{Adv}_{\Pi}^{\text{CCA3(OAE,NI)}}(q, \ell, t) \leq \frac{2(q + \ell)(q + \ell + 1) + 3q + 2\ell}{2^n - (q + \ell)} + 3\text{Adv}_{E, E^{-1}}^{\text{RK-CCA-PRP}}(q + \ell).$$

- (ii) *Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be a MCOE-X scheme as in Definition 6 (ii). We further assume that the block cipher E is secure against related key attacks. Then*

$$\text{Adv}_{\Pi}^{\text{CCA3(OAE,NI)}}(q, \ell, t) \leq \frac{4(q + \ell + 2)(q + \ell + 3) + 6(2q + \ell)}{2^n - (q + \ell)} + \frac{3q(q + 1)}{2^n - q} + \frac{q}{2^{n/2} - q} + 3\text{Adv}_{E, E^{-1}}^{\text{RK-CCA-PRP}}(2q + \ell).$$

Proof. The proof of (i) follows from Theorem 1 together with Lemmas 1 and 2. Due to the lack of space the proof of (ii) it is skipped here and is available in the full version of the paper [14].

Table 5. Instances of MCOE-X: Left side is for messages whose size is evenly divisible by the block size n ; Right side is for arbitrarily sized messages (TS-variant); see text for details

<p>EncryptAuthenticate(H, M)</p> <ol style="list-style-type: none"> 1. $U \leftarrow \varphi(K, 0^n)$ 2. for $i = 1, \dots, L_H - 1$ do $U \leftarrow \varphi(K, H_i \oplus E_U(H_i))$ 3. $\tau \leftarrow E_U(H_{L_H})$ 4. $U \leftarrow \varphi(K, H_{L_H} \oplus \tau)$ 5. for $i = 1, \dots, L$ do $C_i \leftarrow E_U(M_i)$ $U \leftarrow \varphi(K, M_i \oplus C_i)$ 6. $T \leftarrow E_U(\tau)$ 7. return (C_1, \dots, C_L, T) 	<p>DecryptAuthenticate(H, C, T)</p> <ol style="list-style-type: none"> 1. $U \leftarrow \varphi(K, 0^n)$ 2. for $i = 1, \dots, L_H - 1$ do $U \leftarrow \varphi(K, H_i \oplus E_U(H_i))$ 3. $\tau \leftarrow E_U(H_{L_H})$ 4. $U \leftarrow \varphi(K, H_{L_H} \oplus \tau)$ 5. for $i = 1, \dots, L$ do $M_i \leftarrow E_U^{-1}(C_i)$ $U \leftarrow \varphi(K, M_i \oplus C_i)$ 6. if $T = E_U(\tau)$ then return (M_1, \dots, M_L) else return \perp
<p>EncryptAuthenticate(H, C, T)</p> <ol style="list-style-type: none"> 1. $U \leftarrow \varphi(K, 0^n)$ 2. for $i = 1, \dots, L_H - 1$ do $U \leftarrow \varphi(K, H_i \oplus E_U(H_i))$ 3. $\tau \leftarrow E_U(H_{L_H})$ 4. $U \leftarrow \varphi(K, H_{L_H} \oplus \tau)$ 5. for $i = 1, \dots, L - 1$ do $C_i \leftarrow E_U(M_i)$ $U \leftarrow \varphi(K, M_i \oplus C_i)$ 6. $M^* \leftarrow M_L \tau[0 \dots n - l^* - 1]$ 7. $M^* \leftarrow M^* \oplus E_{K \oplus 1^n}(M_L)$ 8. $C^* \leftarrow E_U(M^*)$ 9. Parse $C_L T[0 \dots n - l^* - 1] \leftarrow$ 10. C^* 11. $U \leftarrow \varphi(K, M^* \oplus C^*)$ 12. $C^{**} \leftarrow E_U(\tau)$ 13. $T[n - l^* \dots n - 1] \leftarrow$ $C^{**}[0 \dots l^* - 1]$ 14. return (C_1, \dots, C_L, T) 	<p>DecryptAuthenticateSplitTag(H, C, T)</p> <ol style="list-style-type: none"> 1. $U \leftarrow \varphi(K, 0^n)$ 2. for $i = 1, \dots, L_H - 1$ do $U \leftarrow \varphi(K, H_i \oplus E_U(H_i))$ 3. $\tau \leftarrow E_U(H_{L_H})$ 4. $U \leftarrow \varphi(K, H_{L_H} \oplus \tau)$ 5. for $i = 1, \dots, L$ do $M_i \leftarrow E_U^{-1}(C_i)$ $U \leftarrow \varphi(K, M_i \oplus C_i)$ 6. $C^* \leftarrow C_{L+1} T[0 \dots n - l^* - 1]$ 7. $M^* \leftarrow E_U^{-1}(C^*)$ 8. $U \leftarrow \varphi(K, M^* \oplus C^*)$ 9. $M^* \leftarrow M^* \oplus E_{K \oplus 1^n}(C_L)$ 10. Parse $M_L \tau'[0 \dots n - l^* - 1] \leftarrow M^*$ 11. $T' \leftarrow E_U(\tau)$ 12. if $\tau'[0 \dots n - l^* - 1] = \tau[0 \dots n -$ $l^* - 1]$ and $T'[0 \dots l^* - 1] = T[n -$ $l^* \dots n - 1]$ then return (M_1, \dots, M_L) else return \perp

Lemma 1. Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be a MCOE-X scheme as in Definition 6 (i). Let q be the number of total queries an adversary A is allowed to ask and ℓ be an integer representing the total length in blocks of the queries to \mathcal{E} and \mathcal{D} . Then,

$$\mathbf{Adv}_{\Pi}^{\text{INT-CTXT(NI)}}(q, \ell, t) \leq \frac{(q + \ell)(q + \ell + 1)}{2^n - (q + \ell)} + \frac{2q + \ell}{2^n - (q + \ell)} + \mathbf{Adv}_{E, E^{-1}}^{\text{RK-CCA-PRP}}(q + \ell).$$

Proof (Lemma 1). Our bound is derived by game playing arguments. Consider games G_1 - G_3 of Figure 4 and a fixed adversary A asking at most q queries with a total length of at most ℓ blocks. The functions **Initialize** and **Finalize** are identical for all games in this proof. Lets denote G_0 as the Game INT-CTXT(NI) as defined in Figure 3. Definition 5 states that

$$\mathbf{Adv}_{\Pi}^{\text{INT-CTXT(NI)}}(A) \leq \Pr[A^{G_0} \Rightarrow 1].$$

In G_1 , the encryption and verify placeholders are replaced by their specific McOE-X counterparts as of Definition 6. Clearly, $\Pr[A^{G_0} \Rightarrow 1] = \Pr[A^{G_1} \Rightarrow 1]$. We now discuss the differences between G_1 and G_2 . The set B is initialized to $\{\varphi(K, 0^n)\}$ and then collects new key-input values U which are computed during the encryption or verification process (in lines 204, 207, 213, 223, 226, 232 and 237). We note that, since φ is injective, a collision for the chaining values follows if there is a collision in the U values.

In lines 203 and 222, the LLCP_n oracle is inquired. Finally, the variable **bad** is set to **true** if one of the if-conditions in lines 208, 214, 227, 233, or 238 is **true**. *None* of these modifications affect the values returned to the adversary and therefore

$$\Pr[A^{G_1} \Rightarrow 1] = \Pr[A^{G_2} \Rightarrow 1].$$

For our further discussion we require another game G_4 which is explained in more detail later in this proof³. It follows that

$$\begin{aligned} \Pr[A^{G_2} \Rightarrow 1] &= \Pr[A^{G_3} \Rightarrow 1] + |\Pr[A^{G_2} \Rightarrow 1] - \Pr[A^{G_3} \Rightarrow 1]| \\ &\leq \Pr[A^{G_3} \Rightarrow 1] + \Pr[A^{G_3} \text{ sets bad}] \\ &\leq \Pr[A^{G_4} \Rightarrow 1] + |\Pr[A^{G_3} \Rightarrow 1] - \Pr[A^{G_4} \Rightarrow 1]| + \Pr[A^{G_3} \text{ sets bad}]. \end{aligned} \tag{3}$$

We now proceed to upper bound any of the three terms contained in (3) – in right to left order. The success probability of game G_3 does not differ from the success probability of G_2 unless a chaining value U occurs twice. In this case, the adversary must (i) either have ‘found’ a collision for $E_{\varphi(K,X)}(Y) \oplus Y$, *i.e.* she stumbles over (X, Y) and (X', Y') such that $E_{\varphi(K,X)}(Y) \oplus Y = E_{\varphi(K,X')}(Y') \oplus Y'$ or, (ii), must have found a preimage of $\varphi(K, 0^n)$, which is always the starting point of our chain. Note that that value $\varphi(K, 0^n)$ is initially stored in the set B . In both cases, the variable **bad** would have been set to **true**, and it follows [8] that

$$\Pr[A^{G_3} \text{ sets bad}] \leq \frac{(q + \ell)(q + \ell + 1)}{2^n - (q + \ell)} + \frac{q + \ell}{2^n - (q + \ell)}.$$

³ Since the difference is very minor, we do not provide an extra figure.

```

1  Initialize()
2   $K \xleftarrow{\$} \mathcal{K}()$ ;
3   $B \leftarrow \{\varphi(K, 0^n)\}$ ;

100 Encrypt( $H, M$ ) Game  $G_1$ 
101  $L_H \leftarrow |H|/n$ ;  $L \leftarrow |M|/n$ ;
102  $U \leftarrow \varphi(K, 0^n)$ ;
103 for  $i = 1, \dots, L_H$  do
104    $\tau \leftarrow E_U(H_i)$ ;
105    $U \leftarrow \varphi(K, H_i \oplus \tau)$ ;
106 for  $i = 1, \dots, L$  do
107    $C_i \leftarrow E_U(M_i)$ ;
108    $U \leftarrow \varphi(K, C_i \oplus M_i)$ ;
109  $T \leftarrow E_U(\tau)$ ;
110  $\mathcal{Q} \leftarrow (H, M, C, T)$ ;
111 return  $(C_1, \dots, C_L, T)$ ;

200 Encrypt( $H, M$ ) Game  $G_2, \boxed{G_3}$ 
201  $L_H \leftarrow |H|/n$ ;  $L \leftarrow |M|/n$ ;
202  $A \leftarrow A \cup H$ ;
203  $p \leftarrow \text{LLCP}_n(\mathcal{Q}_{H,M}, (H, M))$ ;
204  $U \leftarrow \varphi(K, 0^n)$ ;
205 for  $i = 1, \dots, L_H$  do
206    $\tau \leftarrow E_U(H_i)$ ;
207    $U \leftarrow \varphi(K, H_i \oplus \tau)$ ;
208   if  $(U \in B \text{ and } i > p)$  then
209     bad  $\leftarrow$  true;  $U \xleftarrow{\$} \{0,1\}^n \setminus B$ ;
210    $B \leftarrow B \cup U$ ;
211 for  $i = 1, \dots, L$  do
212    $C_i \leftarrow E_U(M_i)$ ;
213    $U \leftarrow \varphi(K, C_i \oplus M_i)$ ;
214   if  $(U \in B \text{ and } i + L_H > p)$  then
215     bad  $\leftarrow$  true;  $U \xleftarrow{\$} \{0,1\}^n \setminus B$ ;
216    $B \leftarrow B \cup U$ ;
217  $T \leftarrow E_U(\tau)$ ;
218  $\mathcal{Q} \leftarrow (H, M, C, T)$ ;
219 return  $(C_1, \dots, C_L, T)$ ;

4  Finalize()
5  return win;

112 Verify( $H, C, T$ ) Game  $G_1$ 
113  $L_H \leftarrow |H|/n$ ;  $L \leftarrow |C|/n$ ;
114  $U \leftarrow \varphi(K, 0^n)$ ;
115 for  $i = 1, \dots, L_H$  do
116    $\tau \leftarrow E_U(H_i)$ ;
117    $U \leftarrow \varphi(K, H_i \oplus \tau)$ ;
118 for  $i = 1, \dots, L$  do
119    $M_i \leftarrow E_U^{-1}(C_i)$ ;
120    $U \leftarrow \varphi(K, C_i \oplus M_i)$ ;
121 if  $(T = E_U(\tau) \text{ and } (H, C) \notin \mathcal{Q}_{H,C})$ 
122 then win  $\leftarrow$  true;
123  $\mathcal{Q} \leftarrow (H, \perp, C, \perp)$ ;
124 return  $(T = E_U(\tau))$ ;

220 Verify( $H, C, T$ ) Game  $G_2, \boxed{G_3}$ 
221  $L_H \leftarrow |H|/n$ ;  $L \leftarrow |C|/n$ ;
222  $p \leftarrow \text{LLCP}_n(\mathcal{Q}_{H,M}, (H, M))$ ;
223  $U \leftarrow \varphi(K, 0^n)$ ;
224 for  $i = 1, \dots, L_H$  do
225    $\tau \leftarrow E_U(H_i)$ ;
226    $U \leftarrow \varphi(K, H_i \oplus \tau)$ ;
227   if  $(U \in B \text{ and } i > p)$  then
228     bad  $\leftarrow$  true;  $U \xleftarrow{\$} \{0,1\}^n \setminus B$ ;
229    $B \leftarrow B \cup U$ ;
230 for  $i = 1, \dots, L-1$  do
231    $M_i \leftarrow E_U^{-1}(C_i)$ ;
232    $U \leftarrow \varphi(K, C_i \oplus M_i)$ ;
233   if  $(U \in B \text{ and } i + L_H > p)$  then
234     bad  $\leftarrow$  true;  $U \xleftarrow{\$} \{0,1\}^n \setminus B$ ;
235    $B \leftarrow B \cup U$ ;
236  $M_L \leftarrow E_U^{-1}(C_L)$ ;
237  $U \leftarrow \varphi(K, C_L \oplus M_L)$ ;
238 if  $(U \in B \text{ and } H \notin A)$  then
239   bad  $\leftarrow$  true;  $U \xleftarrow{\$} \{0,1\}^n \setminus B$ ;
240 if  $(T = E_U(\tau) \text{ and } (H, C, T) \notin \mathcal{Q}_{H,C,T})$ 
241 then win  $\leftarrow$  true;
242  $\mathcal{Q} \leftarrow (H, \perp, C, \perp)$ ;
243  $B \leftarrow B \cup U$ ;
244 return  $(T = E_U(\tau))$ ;

```

Fig. 4. Games G_1 - G_3 for the proof of Lemma 1. Game G_3 contains the code in the box while G_2 does not.

We now describe the new game G_4 . It is equal to G_3 *except* that the block cipher E and its inverse E^{-1} are replaced by randomly chosen functions **EncryptBlock** and **DecryptBlock**, which are modeled as pseudo random permutations. We assume that they are implemented via lazy sampling. More precisely, the call $E_K(A)$ is replaced by an invocation of **EncryptBlock** $_K(A)$ and the call $E_K^{-1}(A)$ is replaced by an invocation of **DecryptBlock** $_K(A)$. We now upper bound the difference between G_3 and G_4 .

So, by definition of G_4 , we have

$$|\Pr[A^{G_3} \Rightarrow 1] - \Pr[A^{G_4} \Rightarrow 1]| \leq \mathbf{Adv}_{E,E^{-1}}^{\text{RK-CCA-PRP}}(q + \ell).$$

Finally, we have to upper bound the advantage for the adversary A to win the game G_4 . A can only win this game if the condition in line 238 (resp. 438 for game G_4) is **true**. As usual, we assume *wlog.* that A doesn't ask a question if the answer is already known which implies that $(H, C, T) \notin \mathcal{Q}_{H,C,T}$. For our analysis we distinguish between three cases. So we formally adjust line 240 (*i.e.* choose as the tag computation operation either E or E^{-1}) such that we always have enough randomness left for our result.

Case 1: H has already been used in an *Encrypt* or *Verify* query before and $U \in B$. Since we already have computed τ in the past, the chance of success is upper bounded by the probability $\Pr[E_U^{-1}(T) = \tau]$ which can be upper bounded by $1/(2^n - (q + \ell))$.

Case 2: H has never been used before, also U has never been used as a chaining value. Then the tagging operation uses a 'new key' – essentially due since φ is injective – and therefore the output of $E_U(\tau)$ is uniformly distributed and the success probability is $\leq 1/2^n$.

Case 3: $H \in A$ but U has never been used as a chaining value. The chance of success is upper bounded by $\Pr[E_U^{-1}(T) = \tau]$ which can be upper bounded by $1/2^n$.

Note that the 'missing' fourth case has been explicitly excluded by line 240 (resp. 440). Since these three cases are mutually exclusive, we can upper bound the success probability for q queries as

$$\Pr[A^{G_4} \Rightarrow 1] \leq \frac{q}{2^n - (q + \ell)}.$$

Our claim follows by adding up the individual bounds. □

Lemma 2. *Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be a MCOE-X scheme as in Definition 6 (i). Let q be the number of total queries an adversary A is allowed to ask and ℓ be an integer representing the total length of the queries to \mathcal{E} and \mathcal{D} . Then,*

$$\mathbf{Adv}_{\Pi}^{\text{CPA(AOE,NI)}}(q, \ell, t) \leq 2 \left(\frac{(q + \ell)(q + \ell + 1)}{2^n - (q + \ell)} + \frac{q + \ell}{2^n - (q + \ell)} + \mathbf{Adv}_E^{\text{RK-CPA-PRP}}(q + \ell) \right).$$

The proof is given in the full version of this paper [14].

6 Discussion

New Challenges for Research. At the this point of time, cryptographic research has developed an impressive number of good schemes for encryption, authentication, and authenticated encryption. Many of these schemes have been proven secure under standard assumptions on the underlying primitives. In practice, however, such schemes are often used in a way that undermines security. Trying to design cryptosystems as “misuse resistant” as possible still stands as a challenge for cryptographers.

Furthermore, our research seems to pose new challenges for the design of symmetric primitives. Ideally, we would like to implement MCOE using a tweakable n -bit block cipher with n -bit tweaks, supporting fast random tweak changes. Due to the current lack of such a primitive, we designed MCOE-X, which requires an ordinary n -bit block cipher being secure against XOR-related key attacks, and supporting fast random key changes. Much beyond MCOE, cryptosystem designers could benefit from new tweak-agile tweakable block ciphers and new key-agile ordinary block ciphers.

It is mentionable that MCOE-X, when using Threefish-512 in software, performs considerably better as when using software or even hardware AES-128. (Note that Threefish-512 actually is a tweakable block cipher, but the 128-bit tweak is too short for MCOE.) As an alternative, we developed further variants of MCOE using double encryption and Galois field arithmetic. These two variants also don’t expose the underlying block cipher to related-key attacks.

Conclusion. Originally, this research has been inspired by the search for a default authenticated encryption mode of operation for a general-purpose cryptographic library. It should offer, by default, a huge failure tolerance for practical software developers and still allow being used in an on-line manner.

Since the well-known schemes as, such as OCB and SIV, did not fit our requirements, we searched for other ways to achieve the security and functionality we were looking for. Apart from MCOE, generic composition (Encrypt-then-Mac) of a secure on-line cipher for encryption and a secure deterministic MAC for authentication, using two independent keys might be another solution. As it turned out, using MCOE, one can save the additional key and the time to generate the MAC by using a slightly tweaked on-line cipher for both encryption and authentication.

Acknowledgments. We like to thank Jakob Wenzel for very helpful comments, Phil Rogaway for making us aware of the Galois field native instructions, and the participants of the Dagstuhl Seminar on Symmetric Cryptography 2012 for inspiring discussions.

References

1. Bellare, M., Boldyreva, A., Knudsen, L.R., Namprempre, C.: Online Ciphers and the Hash-CBC Construction. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 292–309. Springer, Heidelberg (2001)
2. Bellare, M., Boldyreva, A., Knudsen, L.R., Namprempre, C.: On-Line Ciphers and the Hash-CBC Constructions. IACR Cryptology ePrint Archive, 2007:197 (2007)

3. Bellare, M., Boldyreva, A., Knudsen, L.R., Namprempre, C.: Online Ciphers and the Hash-CBC Construction. Cryptology ePrint Archive, Report 2007/197; full version of [1] (2007), <http://eprint.iacr.org/>
4. Bellare, M., Namprempre, C.: Authenticated Encryption: Relations among Notions and Analysis of the Generic Composition Paradigm. *J. Cryptology* 21(4), 469–491 (2008)
5. Bellare, M., Rogaway, P.: Encode-Then-Encipher Encryption: How to Exploit Nonces or Redundancy in Plaintexts for Efficient Cryptography. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 317–330. Springer, Heidelberg (2000)
6. Bellare, M., Rogaway, P., Wagner, D.: The EAX Mode of Operation. In: Roy, B., Meier, W. (eds.) FSE 2004. LNCS, vol. 3017, pp. 389–407. Springer, Heidelberg (2004)
7. Black, J.A., Cochran, M., Shrimpton, T.: On the Impossibility of Highly-Efficient Blockcipher-Based Hash Functions. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 526–541. Springer, Heidelberg (2005)
8. Black, J.A., Rogaway, P., Shrimpton, T.: Black-Box Analysis of the Block-Cipher-Based Hash-Function Constructions from PGV. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, p. 320. Springer, Heidelberg (2002)
9. Borisov, N., Goldberg, I., Wagner, D.: Intercepting Mobile Communications: The Insecurity of 802.11. In: MOBICOM, pp. 180–189 (2001)
10. Buonanno, E., Katz, J., Yung, M.: Incremental Unforgeable Encryption. In: Matsui, M. (ed.) FSE 2001. LNCS, vol. 2355, pp. 109–124. Springer, Heidelberg (2002)
11. Intel Corporation. AES-NI Sample Library v1.2 (2010), <http://software.intel.com/en-us/articles/download-the-intel-aesni-sample-library/>
12. Daemen, J.: Hash Function and Cipher Design: Strategies Based on Linear and Differential Cryptanalysis. Ph.D. thesis, Katholieke Universiteit Leuven, Leuven, Belgium (March 1995)
13. Dworkin, M.: Special Publication 800-38C: Recommendation for block cipher modes of operation: the CCM mode for authentication and confidentiality. National Institute of Standards and Technology, U.S. Department of Commerce (May 2005)
14. Fleischmann, E., Forler, C., Lucks, S.: McOE: A Foolproof On-Line Authenticated Encryption Scheme. IACR Cryptology ePrint Archive, 2011:644 (2011)
15. Fouque, P.-A., Martinet, G., Valette, F., Zimmer, S.: On the Security of the CCM Encryption Mode and of a Slight Variant. In: Bellare, S.M., Gennaro, R., Keromytis, A.D., Yung, M. (eds.) ACNS 2008. LNCS, vol. 5037, pp. 411–428. Springer, Heidelberg (2008)
16. Gladman, B.: Brian Gladman's AES Implementation (June 19, 2006), <http://gladman.plushost.co.uk/oldsite/AES/index.php>
17. Gligor, V.D., Donescu, P.: Fast Encryption and Authentication: XCBC Encryption and XECB Authentication Modes. In: Matsui, M. (ed.) FSE 2001. LNCS, vol. 2355, pp. 92–108. Springer, Heidelberg (2002)
18. Goldwasser, S., Micali, S.: Probabilistic Encryption. *J. Comput. Syst. Sci.* 28(2), 270–299 (1984)
19. Gueron, S., Kounavis, M.E.: Efficient implementation of the Galois Counter Mode using a carry-less multiplier and a fast reduction algorithm. *Inf. Process. Lett.* 110(14–15), 549–553 (2010)
20. Hotz, G.: Console Hacking 2010 - PS3 Epic Fail. 27th Chaos Communications Congress (2010), http://events.ccc.de/congress/2010/Fahrplan/attachments/1780_27c3_console_hacking_2010.pdf

21. ISO/IEC. 19772:2009, Information technology – Security techniques – Authenticated Encryption (2009)
22. Iwata, T.: New Blockcipher Modes of Operation with Beyond the Birthday Bound Security. In: Robshaw, M. (ed.) FSE 2006. LNCS, vol. 4047, pp. 310–327. Springer, Heidelberg (2006)
23. Iwata, T.: Authenticated Encryption Mode for Beyond the Birthday Bound Security. In: Vaudenay, S. (ed.) AFRICACRYPT 2008. LNCS, vol. 5023, pp. 125–142. Springer, Heidelberg (2008)
24. Iwata, T., Yasuda, K.: BTM: A Single-Key, Inverse-Cipher-Free Mode for Deterministic Authenticated Encryption. In: Jacobson Jr., M.J., Rijmen, V., Safavi-Naini, R. (eds.) SAC 2009. LNCS, vol. 5867, pp. 313–330. Springer, Heidelberg (2009)
25. Iwata, T., Yasuda, K.: HBS: A Single-Key Mode of Operation for Deterministic Authenticated Encryption. In: Dunkelman, O. (ed.) FSE 2009. LNCS, vol. 5665, pp. 394–415. Springer, Heidelberg (2009)
26. Jutla, C.S.: Encryption Modes with Almost Free Message Integrity. *J. Cryptology* 21(4), 547–578 (2008)
27. Katz, J., Yung, M.: Unforgeable Encryption and Chosen Ciphertext Secure Modes of Operation. In: Schneier, B. (ed.) FSE 2000. LNCS, vol. 1978, pp. 284–299. Springer, Heidelberg (2001)
28. Kohno, T.: Attacking and Repairing the WinZip Encryption Scheme. In: ACM Conference on Computer and Communications Security, pp. 72–81 (2004)
29. Kohno, T., Viega, J., Whiting, D.: CWC: A High-Performance Conventional Authenticated Encryption Mode. In: Roy, B., Meier, W. (eds.) FSE 2004. LNCS, vol. 3017, pp. 408–426. Springer, Heidelberg (2004)
30. Iwata, T.: New Blockcipher Modes of Operation with Beyond the Birthday Bound Security. In: Robshaw, M. (ed.) FSE 2006. LNCS, vol. 4047, pp. 310–327. Springer, Heidelberg (2006)
31. Liskov, M., Rivest, R.L., Wagner, D.: Tweakable Block Ciphers. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 31–46. Springer, Heidelberg (2002)
32. Lucks, S.: Ciphers Secure against Related-Key Attacks. In: Roy, B., Meier, W. (eds.) FSE 2004. LNCS, vol. 3017, pp. 359–370. Springer, Heidelberg (2004)
33. Lucks, S.: Two-Pass Authenticated Encryption Faster Than Generic Composition. In: Gilbert, H., Handschuh, H. (eds.) FSE 2005. LNCS, vol. 3557, pp. 284–298. Springer, Heidelberg (2005)
34. McGrew, D.A., Viega, J.: The Security and Performance of the Galois/Counter Mode (GCM) of Operation. In: Canteaut, A., Viswanathan, K. (eds.) INDOCRYPT 2004. LNCS, vol. 3348, pp. 343–355. Springer, Heidelberg (2004)
35. Ferguson, N., Lucks, S., Schneier, B., Whiting, D., Bellare, M., Kohno, T., Callas, J., Walker, J.: Skein source code and test vectors,
<http://www.skein-hash.info/downloads>
36. Paterson, K.G., Watson, G.J.: Plaintext-Dependent Decryption: A Formal Security Treatment of SSH-CTR. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 345–361. Springer, Heidelberg (2010)
37. Rogaway, P.: Authenticated-Encryption with Associated-Data. In: ACM Conference on Computer and Communications Security, pp. 98–107 (2002)
38. Rogaway, P.: Efficient Instantiations of Tweakable Blockciphers and Refinements to Modes OCB and PMAC. In: Lee, P.J. (ed.) ASIACRYPT 2004. LNCS, vol. 3329, pp. 16–31. Springer, Heidelberg (2004)
39. Rogaway, P.: Nonce-Based Symmetric Encryption. In: Roy, B., Meier, W. (eds.) FSE 2004. LNCS, vol. 3017, pp. 348–359. Springer, Heidelberg (2004)

40. Rogaway, P., Bellare, M., Black, J., Krovetz, T.: OCB: a block-cipher mode of operation for efficient authenticated encryption. In: ACM Conference on Computer and Communications Security, pp. 196–205 (2001)
41. Rogaway, P., Shrimpton, T.: A Provable-Security Treatment of the Key-Wrap Problem. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 373–390. Springer, Heidelberg (2006)
42. Rogaway, P., Shrimpton, T.: Deterministic Authenticated-Encryption: A Provable-Security Treatment of the Key-Wrap Problem. Cryptology ePrint Archive, Report 2006/221; full version of [41] (2006), <http://eprint.iacr.org/>
43. Rogaway, P., Zhang, H.: Online Ciphers from Tweakable Blockciphers. In: Kiayias, A. (ed.) CT-RSA 2011. LNCS, vol. 6558, pp. 237–249. Springer, Heidelberg (2011)
44. Sabin, T.: Vulnerability in Windows NT’s SYSKEY encryption. BindView Security Advisory (1999), <http://marc.info/?l=ntbugtraq&m=94537191024690&w=4>
45. Wu, H.: The Misuse of RC4 in Microsoft Word and Excel. Cryptology ePrint Archive, Report 2005/007 (2005), <http://eprint.iacr.org/>

A Misuse-Attacks: The Weak Point of Current Authenticated Encryption (AE) Schemes

We now give a short overview on one of the attack patterns we have successfully used (cf. Table 2). A more detail led analysis (including more attack patterns) can be found in the full version of this paper [14].

Cipher-block-chaining (CBC) is an unauthenticated encryption mode which is sometimes used as the encryption component of an AE scheme. It is well known that, for constant nonces, the ciphertext of two different plaintexts do reveal the full keystream. It was to be expected that a scheme using counter mode or CBC inherits the nonce reuse issue from that mode. But, as it turned out, common AE schemes also fail at the authenticity frontier, as was already indicated in Table 2 using the following ‘linear tag’ attack pattern. Schemes susceptible to this attack are CWC [29], GCM [34], EAX [6], and CHM [22].

Linear Tag Attack. Assume an AE scheme which generate a keystream $S = F_K(V)$ depending on a secret key K and a nonce V encryption a message M by computing a ciphertext $C = S \oplus M$. For AE schemes using the encrypt-then-authenticate paradigm, we rewrite the authentication tag T as

$$T = f(V) \oplus g(C),$$

where V is the nonce, C is the ciphertext, and f and g are some key-dependent functions. This enables the adversary to mount the following attack:

- Encrypt the plaintext M under the nonce V to (C, T) with $T = f(V) \oplus g(C)$.
- Encrypt the plaintext $M' \neq M$ with $|M'| = |M|$ under the nonce $V' \neq V$ to (C', T') with the tag $T' = f(V') \oplus g(C')$.
- Set $M'' := M' \oplus C' \oplus C$. Encrypt M'' under the nonce V' to (C'', T'') . Observe $C'' = C$, thus $T'' = f(V') \oplus g(C)$.
- Set $T^* = T \oplus T' \oplus T'' = f(V) \oplus g(C')$, The adversary accepts (C', T^*) under V .

Cycling Attacks on GCM, GHASH and Other Polynomial MACs and Hashes

Markku-Juhani Olavi Saarinen

Revere Security

4500 Westgrove Drive, Suite 335, Addison, TX 75001, USA

mjos@reveresecurity.com

Abstract. The Galois/Counter Mode (GCM) of operation has been standardized by NIST to provide single-pass authenticated encryption. The GHASH authentication component of GCM belongs to a class of Wegman-Carter polynomial hashes that operate in the field $\text{GF}(2^{128})$. We present message forgery attacks that are made possible by its extremely smooth-order multiplicative group which splits into 512 subgroups. GCM uses the same block cipher key K to both encrypt data and to derive the generator H of the authentication polynomial for GHASH. In present literature, only the trivial weak key $H = 0$ has been considered. We show that GHASH has much wider classes of weak keys in its 512 multiplicative subgroups, analyze some of their properties, and give experimental results on AES-GCM weak key search. Our attacks can be used not only to bypass message authentication with garbage but also to target specific plaintext bits if a polynomial MAC is used in conjunction with a stream cipher. These attacks can also be applied with varying efficiency to other polynomial hashes and MACs, depending on their field properties. Our findings show that especially the use of short polynomial-evaluation MACs should be avoided if the underlying field has a smooth multiplicative order.

Keywords: Cryptanalysis, Galois/Counter Mode, AES-GCM, Cycling Attacks, Weak Keys.

1 Introduction

Authenticated encryption modes and algorithms provide confidentiality and integrity protection in a single processing step. This results in performance and cost advantages as data paths can be shared.

The Galois/Counter Mode (GCM) has been standardized by NIST [1] to be used in conjunction with a 128-bit block cipher for providing authenticated encryption functionality. When paired with the AES [2] algorithm, the resulting AES-GCM combination has been used as a replacement to dedicated hash-based HMAC [3] in popular cryptographic protocols such as SSH [4], IPSec [5] and TLS [6].

In AES-GCM, data is encrypted using the Counter Mode (CTR). A single AES key K is used to both encrypt data and to derive authentication secrets. The component that is used by GCM to produce a message authentication code is called GHASH. GCM also supports Additional Authenticated Data (AAD) which is authenticated using GHASH but transmitted as plaintext.

The GHASH algorithm belongs to a widely studied class of Wegman-Carter [7,8] polynomial MACs. These were originally proposed in context of polynomial evaluation independently by three authors [9,10,11]. A good overview of their genealogy and evolution is by Bernstein [12,13]. The security bounds known for these algorithms indicate that a n -bit tag will give $2^{-\frac{n}{2}}$ security against forgery [12,14].

In this paper we give further evidence that this is not only the security lower bound but an upper bound as well. It can be argued that universal hashes sacrifice communication bandwidth for convenience as traditional hash-based MACs are designed to reach the information theoretic 2^{-n} bound against message forgery and are therefore technically somewhat inferior, especially for short MACs. The security against cycling attacks depends very sharply on the properties of the underlying field.

This paper is structured as follows. We give a description of GHASH in Section 2, followed by a key observation regarding collisions derived from cycles in Section 3. Section 4 contains an analysis of cycle lengths and group orders. In Section 5 we discuss the probability of successful forgery. Section 6 briefly considers targeted attacks against underlying protocols. Section 7 contains a test and experimental results related to cycle lengths. We discuss the security of other polynomial mac constructions in Section 8 and conclude in Section 9.

2 Description of GHASH

Let X be a concatenation of unencrypted authenticated data, CTR-encrypted ciphertext, and padding. This data is split into m 128-bit blocks X_i :

$$X = X_1 \parallel X_2 \parallel \cdots \parallel X_m.$$

AES is used to derive the root authentication key $H = E_K(0)$. The same AES key K is also used as the data encryption key. In the present work we assume that H is unknown to the attacker as the scheme would be otherwise trivially breakable.

GHASH is based on operations in the finite field $\text{GF}(2^{128})$. Horner's rule is used in this field to evaluate the polynomial Y .

$$Y_m = \sum_{i=1}^m X_i \times H^{m-i+1}. \quad (1)$$

Figure 1 illustrates how this value is usually computed (together with the CTR mode). The authentication tag is finalized with $T = Y_m + E_K(\text{IV} \parallel 0^{31} \parallel 1)$, assuming that a 96-bit Initialization Vector (IV) is used. The IV value must never be reused as that would lead to the ‘‘forbidden attack’’ discussed by Joux in [15].

3 Collisions from Weak Keys

It has been observed that if $E_K(0) = H = 0$, the polynomial Y evaluates to zero and the security of GHASH breaks down. In fact, some sources assume that this pathological case is the only weak key [16]. AES keys K that produce this fixed point are not

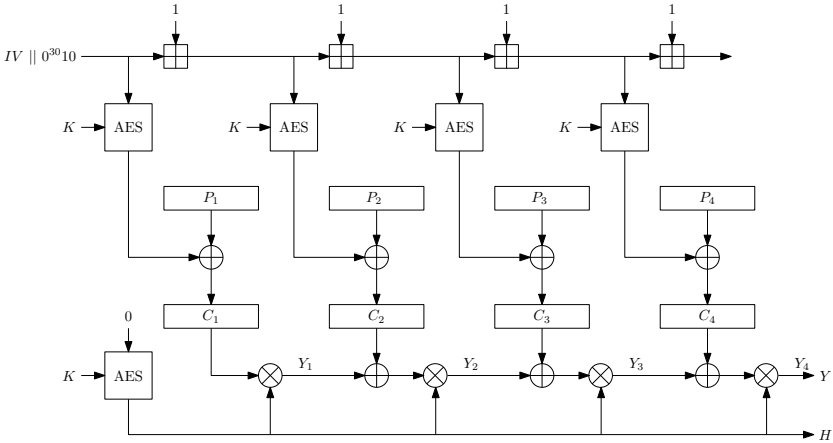


Fig. 1. Basic operation of first four rounds of GCM-CTR (without unencrypted authenticated data or padding). Here \boxplus denotes regular modular addition, \oplus bitwise XOR operation, and \otimes multiplication in $\text{GF}(2^{128})$. The counter is initialized with IV and incremented by 1 for each block. This is used to to produce a keystream that is XORed over plaintext blocks P_i to produce ciphertext blocks C_i (or vice versa). The lower half of the diagram shows how the authentication tag is processed; each authenticated block is XORed over the state Y and multiplied with $H = E_K(0)$. The final processing of the authentication tag Y is omitted from this picture.

known.¹ However, It is easy to see why such keys should exist for AES, especially when the size of K is more than 128 bits.

Our main observation is that sometimes the powers of H will repeat in a relatively short cycle. A trivial example occurs when H is equal to the identity element 1, which will lead to all powers being equal. Due to the commutativity of addition in Equation 1, a GHASH collision can be achieved by swapping any two ciphertext blocks X_i and X_j . This amounts to message forgery.

More generally, if we know that $H^{m-i+1} = H^{m-j+1}$ with $i \neq j$, we may simply swap ciphertext blocks X_i and X_j and the resulting authentication tag stays unmodified which amounts to message forgery. This can be easily observed from Equation 1. Elementary group theory tells us that the powers of H will repeat in cycles which are determined by $n = \text{ord}(H)$, the multiplicative order of H . Hence we may produce collisions by swapping X_i and X_{i+nm} for arbitrary i and m .

4 Cycle Lengths and Group Orders

From Lagrange’s theorem in group theory we know that all subgroups divide the group of order $2^{128} - 1$. Numbers of this type factor into Fermat numbers

$$2^{2^n} - 1 = \prod_{i=1}^n 2^{2^{i-1}} + 1. \tag{2}$$

¹ Some block ciphers such as GOST allow such fixed-point keys to be very easily found.

We can easily obtain the full factorization of $2^{128} - 1$:

$$3 * 5 * 17 * 257 * 641 * 65537 * 274177 * 6700417 * 67280421310721. \quad (3)$$

As this is a “smooth number”, we can see that there are classes of H and therefore K values that produce cycles of length $n = 1, 3, 5, 15, 17, 51, \dots$; any one of the $2^9 = 512$ subset products of the primes in Equation 3 is a valid group order.²

4.1 Illustrating Multiplicative Subgroup Cycles

Due to the peculiar way finite field arithmetic is defined in the GCM standard [1], the identity element with $\text{ord}(H) = 1$ is:

$$H = 80 \ 00 \ 00 \ 00 \ 00 \ 00 \ 00 \ 00 \ 00 \ 00 \ 00 \ 00 \ 00 \ 00 \ 00 \ 00 \ 00 \ 00$$

Apparently this was considered as the “first bit” by those who originally implemented GCM. Otherwise standard polynomial arithmetic is used with the field representation defined by the reducing polynomial $x^{128} + x^7 + x^2 + x + 1$.

The following two elements will produce a cycle of length $\text{ord}(H) = 3$ (the cycle obviously goes through the identity as well):

$$\begin{aligned} H &= 10 \ D0 \ 4D \ 25 \ F9 \ 35 \ 56 \ E6 \ 9F \ 58 \ CE \ 2F \ 8D \ 03 \ 5A \ 94 \\ H &= 90 \ D0 \ 4D \ 25 \ F9 \ 35 \ 56 \ E6 \ 9F \ 58 \ CE \ 2F \ 8D \ 03 \ 5A \ 94 \end{aligned}$$

These four elements have $\text{ord}(H) = 5$:

$$\begin{aligned} H &= 46 \ 36 \ BD \ BD \ 1C \ 76 \ 43 \ D3 \ 4E \ E4 \ BB \ 1B \ F9 \ CA \ 08 \ 4F \\ H &= 92 \ 17 \ 8D \ 40 \ 26 \ DA \ 1D \ CA \ 42 \ 96 \ 77 \ 87 \ 30 \ EB \ 9A \ 9E \\ H &= 82 \ C7 \ C0 \ 65 \ DF \ EF \ 4B \ 2C \ DD \ CE \ B9 \ A8 \ BD \ E8 \ C0 \ 0A \\ H &= D6 \ E6 \ F0 \ 98 \ E5 \ 43 \ 15 \ 35 \ D1 \ BC \ 75 \ 34 \ 74 \ C9 \ 52 \ DB \end{aligned}$$

We do not know which actual AES keys produce these H values, nor do we recommend testing against these particular values as the probability of hitting them is exceedingly small.

Note that a cycle of length such as $15 = 3 * 5$ also contains the beforementioned component groups of order 1, 3 and 5, in addition to the 8 unique elements that can act as a generator of the cycle of order 15. This is entirely analogous to arithmetic in the addition group of integers modulo 15; 0 will generate a “cycle” of one element when repeatedly added to itself, 5 and 10 will generate a cycles of order 3, the four elements { 3, 6, 9, 12 } cycles of order 5 and the rest of the numbers will have order 15. This is illustrated in Figure 2.

5 Message Forgery

We know that the field $\text{GF}(2^{128})$ offers a generous serving of $2^9 = 512$ different multiplicative subgroups. Figure 3 shows that these are quite evenly distributed in the range due to the nearly log-uniform progression of the factors.

² The term *smooth number* comes from factorization theory and indicates that a number factors into a large number of small primes.

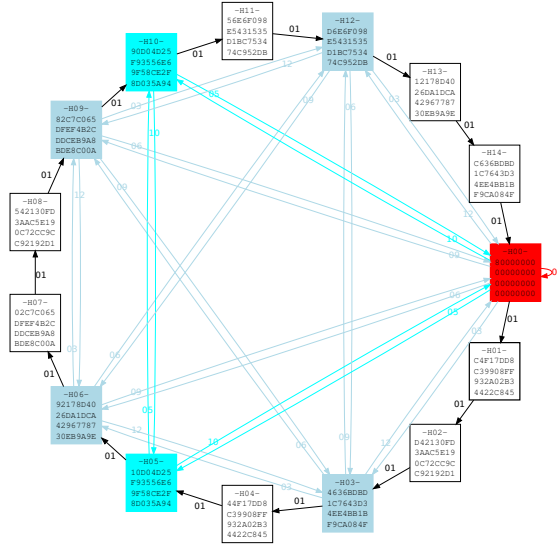


Fig. 2. The case of $H = C4 F1 7D D8 C3 99 08 FF 93 2A 02 B3 44 22 C8 45$ (clockwise). This is one of eight elements that generate a multiplicative subgroup in $GF(2^{128})$ which is isomorphic to the additive group \mathbb{Z}_{15} . The identity element and subgroups of sizes 3 and 5 are also demonstrated. There are 512 multiplicative subgroups of different sizes in this particular field.

In our attack the adversary does not know H but will simply attempt a blind forgery by swapping two (or more) message blocks in transit as discussed in Section 3.

It is easy to show that it is sufficient that the group order divides the distance between swapped elements. Since each subgroup of size n has exactly n elements, we arrive at the following observation:

Theorem 1. *Let n be a number satisfying $\gcd(2^{128} - 1, n) = n$. Blindly swapping blocks X_i and X_j , where $i \equiv j \pmod n$ will result in a successful forgery with probability of at least $\frac{n+1}{2^{128}}$ for some random H .*

Proof. The distance congruence implies that the distance between X_i and X_j is a multiple of n . The $\gcd(2^{128} - 1, n) = n$ condition implies that n is one of the $2^9 = 512$ possible multiplicative subgroup sizes in $GF(2^{128})$. If indeed $\text{ord}(H) \mid n$ then $H^i = H^j$ and the forgery is successful due to commutativity of equation 1. We observe that the cycles are unique; there are n members in a subgroup of size n and the set of n elements is unique to each subgroup size. Hence the probability of hitting one of these cycle elements is $\frac{n}{2^{128}}$. In addition there is the pathological case $H = 0$ which completes the proof. \square

If the \gcd condition given in Theorem 1 does not hold, we have no reason to expect that the forgery is successful with a probability higher than $\frac{1}{2^{128}}$.




Fig. 3. GCM / GHASH: probability of hitting a multiplicative subgroup (cycle) of given (or smaller) size with a random authentication generator H in $GF(2^{128})$. For comparison we also graph the security for $GF(2^{127})$, which is entirely contained in the lower and right borders of the graph due to the fact that its multiplicative group order $2^{127} - 1$ is a prime.

Assuming that an oracle has indicated a successful message forgery, any number of consecutive forgeries can be produced with probability 1 if the key remains unchanged (IV may change).

6 Targeted Multiple Bit Forgeries

Our attacks enable elaborate message forgeries against authenticated encryption hybrids such as GCM due to the fact that the CTR encryption mode behaves like a stream cipher; flipping a ciphertext bit will result the corresponding plaintext bit to be flipped. This is especially true for lightweight protocols that combine a short binary polynomial MAC with a stream cipher.

If $\text{ord}(H)|(i - j)$ the authentication tag will remain valid as long as the equation

$$X_i \times H^{m-i+1} + X_j \times H^{m-j+1} = c \tag{4}$$

holds for some (unknown) constant c related to the authentication tag. If we write $H^{m-i+1} = H^{m-j+1} = H_c$, this can be simplified to

$$X_i + X_j = c \times H_c^{-1}. \tag{5}$$

We see that the authentication tag will be valid if the sum of ciphertext blocks on the left side of Equation 5 remains constant. One may therefore flip *individual bits* in block

X_i if the corresponding bit in X_j is also flipped. Any number of such modifications can be done to a message without affecting the probability of success (assuming that the same distance is used) indicated by Theorem 1.

7 Testing for AES-GCM Weak Keys

We know that finding weak H values is easy, so a natural question arises on how to determine weak AES keys K that produce these weak H roots.

To determine group order, we use a simple algorithm which is related to the Silver-Pohlig-Hellman algorithm for discrete logarithms [17]. Our algorithm is based on the following elementary observation:

Theorem 2. *Let p be one of the prime divisors given in Equation 3. If and only if p divides $\text{ord}(H)$ we have*

$$H^{\frac{2^{128}-1}{p}} \neq 1. \tag{6}$$

Proof. Let g be a generator of the full multiplicative group; $\text{ord}(g) = 2^{128} - 1$. Then each element $H \neq 0$ can be expressed as a power $H = g^h$ for some $h, 0 \leq h < 2^{128} - 1$. Raising an element to power q , where $q \mid 2^{128} - 1$, sets the index modulo q to zero: $(g^h)^q = g^{qh}$. Since $\frac{2^{128}-1}{p}$ is divisible with all prime divisors q_i of the group order except p , we see that the condition of Equation 6 only holds if $h \neq 0 \pmod{p}$, which is equivalent to the condition $p \mid \text{ord}(H)$. \square

By performing the exponentiation test of Theorem 2 for each one of the nine prime divisors of $2^{128} - 1$ in Equation 3, we may completely determine the multiplicative order of H .

7.1 An Efficient Algorithm for Subgroup Size

Raising a finite field element to a Fermat $F_n = 2^{2^n} + 1$ power can be done efficiently. It is well known that squaring operation is “linear” in $\text{GF}(2^n)$ [18]. For $\text{GF}(2^{128})$, a unique 128×128 bit matrix \mathbf{M}_0 exists that satisfies

$$X^2 = \mathbf{M}_0 X \tag{7}$$

for all X . In the following $\mathbf{M}_0 X$ denotes a matrix multiplication where X is interpreted as a vector of 128 bits and $X \times X = X^2$ is a multiplication where X is interpreted as a (polynomial) member of $\text{GF}(2^{128})$.

By squaring \mathbf{M}_0 , we obtain $\mathbf{M}_1 = \mathbf{M}_0^2$ which satisfies $X^4 = \mathbf{M}_1 X$ for all X . By repeating this process we can rapidly compute $\mathbf{M}_0, \mathbf{M}_1, \dots, \mathbf{M}_6$ that satisfy

$$X^{2^{2^i}} = \mathbf{M}_i X. \tag{8}$$

Once the matrices (table lookups) \mathbf{M}_i have been initialized, raising the authentication key H to a Fermat number power can be achieved with:

$$H^{F_n} = \mathbf{M}_n H \times H. \tag{9}$$

Therefore this operation can be made with a table lookup (multiplication with M_n) and a single Galois Field multiplication. The matrices need to be computed only once as they are independent from particular H .

Since $2^{128} - 1 = \prod_{i=0}^6 F_i$, checking whether the group order is of H is divisible with Fermat number F_i involves raising H to all Fermat powers F_j *except* F_i . For example, to check whether or not group order is divisible with $F_3 = 257$, we may see if this equation holds:

$$M_6(M_5(M_4(M_2(M_1(M_0H \times H) \times H) \times H) \times H) \times H) \times H = 1. \quad (10)$$

The Fermat numbers F_5 and F_6 are not primes (unlike F_0, F_1, F_2, F_3 and F_4 which are indeed the only known Fermat primes). Here the technique involves first powering H to all Fermat powers except $F_5 = 641 * 6700417$ or $F_6 = 274177 * 67280421310721$. Then then we use a conventional square-multiply exponentiation method to individually check these two subfactors.

In practice the matrix M_i multiplication is implemented as byte-based table lookups with seven $16 \times 256 \times 128$ - bit tables. The initialization of these tables is very fast as M_{i+1} can be developed from M_i with a loop of $16 * 256$ table lookups. Significant speedups are achieved by reusing partial results.

7.2 Experimental Results

Using the techniques outlined in the previous subsection, we have developed a reasonably efficient cycle determination code specifically for GCM's $GF(2^{128})$, together with an AES-128 key setup and encryption function for deriving H values from K values.

Our implementation is currently able to fully determine the order of 25000 AES keys per second on a low-end Linux laptop that has a single 1.7 GHz AMD V140 processor.

Over couple of days we tested 2^{32} AES-128 keys and found progressively smaller subgroups:

$n \approx 2^{126.4}$	$K = 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 02$
$n \approx 2^{125.6}$	$K = 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 03$
...	
$n \approx 2^{96.52}$	$K = 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 24\ 3E\ 8B\ 40$
$n \approx 2^{96.00}$	$K = 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 37\ 48\ CF\ CE$
$n \approx 2^{93.93}$	$K = 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 42\ 87\ 3C\ C8$
$n \approx 2^{93.41}$	$K = 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ EC\ 69\ 7A\ A8$

As indicated by Figure 3, a significantly smaller group than $2^{128-32} = 2^{96}$ was found with 2^{32} effort, due to the large number of multiplicative subgroup sizes available in $GF(2^{128})$.

There is clearly room for improvement. The search is fully parallelizable, and hence a massively parallel FPGA or GPU-based search could be performed to find subgroups of magnitude $n \approx 2^{64}$ or less.

8 Other Polynomial-Evaluation MACs

The security of Polynomial-evaluation MACs against attacks of this type can be determined from the factorization of the group size in straightforward fashion. Trivial changes can introduce radical differences.

One may consider this difference by comparing the binary field $GF(2^{127})$ and the prime field $GF(2^{127} - 1)$. Here the binary field is perfectly secure due to the fact that $2^{127} - 1$ is indeed a prime (if the message is processed in 127-bit blocks). However, the latter prime field has a multiplicative order $2^{127} - 2$ which factors spectacularly into 15 pieces and is exceptionally weak against a cycling attack! We note that the HASH127 MAC is based on the latter [19]. This is illustrated in Figure 3.

If a prime field is to be used, we recommend Sophie Germain primes where $q = (p - 1)/2$ is also a prime. Such a field has well-understood cycle properties which may be easily determined using the Legendre symbol from elementary number theory. A practical alternative to GCM would use a Sophie Germain prime such as $GF(2^{128} + 12451)$, which is slightly larger than the 2^{128} to deter trivial collisions.

It is clear that risks rise quadratically when GCM is used with a 64-bit block cipher as suggested in Appendix A of [20]. There is a substantial risk of hitting a bad long-term key and therefore we recommend against using the 64-bit GCM.

9 Conclusions and Future Work

We have shown that the GHASH algorithm has other weak key classes besides the trivial $H = 0$ case considered in current literature [16]. This is a result of the multiplicative group of $GF(2^{128})$ having a particularly smooth order.

Our attacks allow specific plaintext bits to be targeted by modifying ciphertext bits, which can have a devastating effect when a short polynomial MAC over a binary field is combined with a stream cipher in a (lightweight) communication protocol. The probability of randomly hitting an exploitable weak key with a AES-GCM cryptographic protocol such as SSH [4], IPsec [5] or TLS [6] is very small.

However, malicious players may exploit subtle weaknesses in cryptographic protocols in surprising ways. One feature of cycle attacks is that an attacker may first test for short cycles and then force a re-keying event if the test fails; once a long-term key with a short cycle is found, she may exploit it any number of times.

We have also described a straightforward method of detecting GHASH weak keys. We performed an exhaustive experiment that found many AES-128 keys that produce H with order below $n \approx 2^{96}$.

We suggest that binary fields $GF(2^n)$ with prime $2^n - 1$ or Sophie Germain prime fields are used in constructions of this type as this minimizes the total number of weak keys. This was illustrated with the surprising observation that $GF(2^{127})$ is perfectly secure against this type of attack while GCM's $GF(2^{128})$ is not.

One interesting future research direction and open question is the feasibility of mapping the weak H values to K symmetric keys with various block ciphers other than AES.

References

1. NIST: Recommendation for block cipher modes of operation: Galois/counter mode (GCM) and GMAC. NIST Special Publication 800-38D (2007)
2. NIST: The advanced encryption standard (AES). FIPS Publication 197 (2001)
3. Bellare, M., Canetti, R., Krawczyk, H.: Keying Hash Functions for Message Authentication. In: Koblitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 1–15. Springer, Heidelberg (1996)
4. Igoe, K., Solinas, J.: AES Galois counter mode for the secure shell transport layer protocol. IETF Request for Comments 5647 (2009)
5. Law, L., Solinas, J.: Suite B cryptographic suites for IPsec. IETF Request for Comments 4869 (2007)
6. Salter, M., Rescorla, E., Housley, R.: Suite B profile for transport layer security (TLS). IETF Request for Comments 5430 (2009)
7. Wegman, M.N., Carter, J.L.: New classes and applications of hash functions. In: 20th Annual Symposium on Foundations of Computer Science. IEEE Computer Society Press, New York (1979)
8. Wegman, M.N., Carter, J.L.: New hash functions and their use in authentication and set equality. *Journal of Computer and System Sciences* 22, 265–279 (1981)
9. den Boer, B.: A simple and key-economical unconditional authentication scheme. *Journal of Computer Security* 2, 65–71 (1993)
10. Taylor, R.: An Integrity Check Value Algorithm for Stream Ciphers. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 40–48. Springer, Heidelberg (1994)
11. Bierbrauer, J., Johansson, T., Kabatianskii, G., Smeets, B.: On Families of Hash Functions via Geometric Codes and Concatenation. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 331–342. Springer, Heidelberg (1994)
12. Bernstein, D.J.: Stronger Security Bounds for Wegman-Carter-Shoup Authenticators. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 164–180. Springer, Heidelberg (2005)
13. Bernstein, D.J.: The Poly1305-AES Message-Authentication Code. In: Gilbert, H., Handschuh, H. (eds.) FSE 2005. LNCS, vol. 3557, pp. 32–49. Springer, Heidelberg (2005)
14. Sarkar, P.: A trade-off between collision probability and key size in universal hashing using polynomials. *Designs, Codes and Cryptography* 58(3), 271–278 (2011)
15. Joux, A.: Authentication failures in NIST version of GCM. NIST Comment (2006)
16. Handschuh, H., Preneel, B.: Key-Recovery Attacks on Universal Hash Function Based MAC Algorithms. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 144–161. Springer, Heidelberg (2008)
17. Pohlig, S., Hellman, M.: An improved algorithm for computing logarithms over $GF(p)$ and its cryptographic significance. *IEEE Transactions on Information Theory* 24(1), 106–110 (1978)
18. Ferguson, N.: Authentication weaknesses in GCM. NIST Comment (May 2005)
19. Bernstein, D.J.: Floating-point arithmetic and message authentication (1999), <http://cr.ypt.to/papers.html#hash127>
20. McGrew, D.A., Viega, J.: The Galois/counter mode of operation (GCM). Submission to NIST (2005)

Collision Attacks on the Reduced Dual-Stream Hash Function RIPEMD-128

Florian Mendel¹, Tomislav Nad², and Martin Schl affer²

¹ Katholieke Universiteit Leuven, ESAT/COSIC and IBBT, Belgium

² Graz University of Technology, IAIK, Austria

Abstract. In this paper, we analyze the security of RIPEMD-128 against collision attacks. The ISO/IEC standard RIPEMD-128 was proposed 15 years ago and may be used as a drop-in replacement for 128-bit hash functions like MD5. Only few results have been published for RIPEMD-128, the best being a preimage attack for the first 33 steps of the hash function with complexity $2^{124.5}$. In this work, we provide a new assessment of the security margin of RIPEMD-128 by showing attacks on up to 48 (out of 64) steps of the hash function. We present a collision attack reduced to 38 steps and a near-collisions attack for 44 steps, both with practical complexity. Furthermore, we show non-random properties for 48 steps of the RIPEMD-128 hash function, and provide an example for a collision on the compression function for 48 steps.

For all attacks we use complex nonlinear differential characteristics. Due to the more complicated dual-stream structure of RIPEMD-128 compared to its predecessor, finding high-probability characteristics as well as conforming message pairs is nontrivial. Doing any of these steps by hand is almost impossible or at least, very time consuming. We present a general strategy to analyze dual-stream hash functions and use an automatic search tool for the two main steps of the attack. Our tool is able to find differential characteristics and perform advanced message modification simultaneously in the two streams.

Keywords: hash functions, RIPEMD-128, collisions, near-collisions, differential characteristic, message modification, automatic tool.

1 Introduction

In the last few years, the cryptanalysis of hash functions has become an important topic within the cryptographic community. Especially the collision attacks on the MD4 family of hash functions have weakened the security assumptions of many commonly used hash functions. Still, most of the existing cryptanalytic work has been published for this particular family of hash functions [17,19,20]. In fact, practical collisions have been shown for MD4, MD5, RIPEMD and SHA-0. For SHA-1, a collision attack has been proposed with a complexity of about 2^{63} [18]. However, some members of this family including the ISO/IEC standard RIPEMD-128 (the successor of RIPEMD) seems to be more resistant against these attacks. In this paper, we analyze the security of RIPEMD-128 against collision attacks and show that the security margin is less than expected.

Related Work. Since its proposal 15 years ago only a few results have been published for RIPEMD-128. Most published results are concerning the preimage resistance of the hash function [13, 16]. The best currently known attack is a preimage attack for 33 steps and 36 intermediate steps of the hash function with a complexity only slightly faster than the generic complexity of 2^{128} [16]. The only work regarding the collision resistance of RIPEMD-128 has been published by Mendel et al. [11], where the application of the differential attacks on RIPEMD by Dobbertin [5] and Wang et al. [17] is studied. However, due to the increased number of steps and the fact that the two streams are more different than in RIPEMD, they concluded that RIPEMD-128 is secure against this type of attacks.

Our Contribution. In this paper, we first provide a general strategy to analyze dual-stream hash functions in Sect. 2. We analyze different methods to find high-probability differential characteristics which work for both streams. Similar as in the attack on RIPEMD [17], characteristics in two streams are impossible with a high probability. Therefore, in our attacks an automatic search tool is essential for finding valid differential characteristics [4, 10]. This is especially important in the first round of a hash function where characteristics are usually quite dense. In this first round, one usually assumes that conditions imposed by the characteristic can be fulfilled efficiently using message modification techniques. However, message modification is much more difficult in the dual-stream case since two state words are updated using a single message word. This reduced freedom could in general be compensated with hand-tuned advanced message modification techniques [8, 9, 15, 20]. However, another contribution of our work is to provide a fully automatic tool which can be used to find conforming message pairs in the first round of a dual-stream hash function.

Table 1. Summary of our new and previous results on RIPEMD-128

component	attack	steps	complexity	generic	reference
hash	collision	38	example, 2^{14}	2^{64}	Sect. 4
hash	near-collision	44	example, 2^{32}	$2^{47.8}$	Sect. 5.1
hash	non-randomness	48	2^{70}	2^{76}	Sect. 5.2
compression	collision	48	example, 2^{40}	2^{64}	Sect. 5.3
hash	preimage	33	$2^{124.5}$	2^{128}	[13]
hash	preimage	interm. 35	2^{121}	2^{128}	[13]
hash	preimage	interm. 36	$2^{126.5}$	2^{128}	[16]

We apply our attack strategy and tools to the ISO/IEC standard RIPEMD-128 which we describe in Sect. 3. Using our automatic tools, we are able to construct the first practical collisions for up to 38 steps of RIPEMD-128 with a complexity of 2^{14} . We describe the collision attack in details in Sect. 4. The attack can be extended (Sect. 5) to practical near-collisions on 44 steps with complexity 2^{32} . Furthermore, we provide a theoretical distinguisher of the hash function for 48

steps (3 out of 4 rounds) and show that 3 rounds of the RIPEMD-128 compression function are not collision free. Our results are summarized in Table 1, together with all known previous results. Finally, we conclude in Sect. 6 and discuss directions of future work on hash functions with parallel state update transformation.

2 Cryptanalysis of Dual-Stream Hash Functions

In this section, we describe our attack strategy for the cryptanalysis of dual-stream hash functions. The general attack strategy is based on the recent results in cryptanalysis of the MD4-family of hash functions [17, 20]. However, the application of this strategy is nontrivial in the case of dual stream hash functions. Since in each step, one message word is used to update two state words, the freedom of an attacker in finding valid differential characteristics and performing message modification is limited. Hence, a more careful analysis is required to overcome this problem.

2.1 Collision Attacks on Hash Functions

In the following, we first give a brief overview of the attack strategy used in the recent collision attacks on the MD4-family of hash functions [17, 20]. All attacks basically use the same strategy which we adopt for dual-stream hash functions. The high-level strategy can be summarized as follows:

1. Find a characteristic for the hash function that holds with high probability after the first round of the hash function.
2. Find a characteristic (not necessary with high probability) for the first round of the hash function.
3. Use message modification techniques to fulfill conditions imposed by the characteristic in the first round. This increases the probability of the characteristic.
4. Use random trials to find values for the remaining free message bits such that the message follows the characteristic.

The most difficult and important part of the attack is to find a good differential characteristic for both the first round and the remaining rounds of the hash function, since this defines the final attack complexity. There are several methods to find good differential characteristics. The second important part of the attack is to find conforming inputs for the complex nonlinear differential characteristic in the first round of the hash function using message modification techniques.

2.2 Collision Attacks on Dual-Stream Hash Functions

In the following, we will describe our approach to construct good differential characteristics and find colliding message pairs for dual-stream hash functions. We focus on hash functions like RIPEMD-128, but the general idea is applicable to any hash function with two or more streams.

Finding Suitable Differential Characteristics. If the two streams of the hash function are the same except for constant additions, the same differential characteristic can be used in both streams. For instance, in the case of RIPEMD, the permutation and rotation values are indeed equal for both streams. Hence, it is sufficient to find a collision-producing characteristic for only one stream (3 rounds) and apply it simultaneously to both streams [17]. Nevertheless, the number of necessary conditions increases for two streams. Hence, it is more likely to have contradicting conditions. In fact, Wang et al. reported that among 30 selected collision-producing characteristics only one can produce a real collision.

If the two streams are more different, we first need to find a differential characteristic for the hash function after round 1, which holds with a high probability in both streams. One approach is to find such characteristics is to use a linearized model of the hash function and algorithms from coding theory [2, 7, 14]. This works quite well for hash functions with a regular message expansion and step update transformation (like SHA-1), and can be applied to dual-stream hash functions in a straight-forward way.

However, the linearization approach does not work well for hash functions with a permutation of words in the message expansion and different rotation values in the state update transformation (RIPEMD-128 and RIPEMD-160). One usually gets linear differential characteristics with high Hamming weight and hence, a high complexity. However, for such hash functions, we can still make use of the approach of Wang et al. in the attacks on MD4, RIPEMD and MD5 [17, 20]. The idea is to use differences in one or more message words to find local (or inner) collisions within a few steps in the last round(s) of the hash function. Then a suitable characteristic for the remaining steps, preferably also using short local collisions, has to be constructed. Although this is obviously more difficult for dual-stream hash functions, we were able to construct such high-probability differential characteristics for reduced RIPEMD-128 (see Sect. 4.1).

Once, the characteristic after round 1 is fixed we need to find a characteristic (not necessary with high probability) for the first round of the hash function for both streams. Note that in the previous part of the attack it might still be possible to construct inner collisions with hand by choosing the differences carefully. However, to construct a valid nonlinear differential characteristic for both streams in the first round, an automatic search tool is needed. While one can use complex differential characteristics in both streams, we aim for differential characteristics that are sparse in at least one of the two streams, since such sparse characteristics will then also reduce the complexity of the message modification step.

Using Message Modification Techniques. Once we have fixed the differential characteristic for both streams we start with the message search. In the first round, the freedom of the whole message block can be used to get a conforming message pair for the first 16 steps. For single-stream hash function, basic message modification techniques simply choose conforming state words and invert each step update transformation to get the message word [20]. However, as already noted by Wang et al. [17], message modification is more complicated for

two streams since the conditions on two state words need to be fulfilled using a single message word. While in RIPEMD the same message word is used in the same step of the left and right stream, this is not the case in RIPEMD-128, which significantly increases the complexity of message modification.

In the attack on RIPEMD, two techniques have been proposed exploiting the freedom of other message bits using carry effects, the Boolean function and previous message words. The same rotation values in RIPEMD allow an easier application of this idea since it is still possible to fulfill conditions from LSB to MSB. However, for streams with different rotation values, previously corrected conditions may become invalid again. In general, conditions on two state words using a single message word can be fulfilled using advanced message modification techniques. Many dedicated techniques have been proposed in recent years [8, 9, 15, 20], which could also be used to fulfill conditions in the first round of dual-stream hash functions.

To simplify the message modification we use a more general approach. Instead of complicated, dedicated techniques, we use an automated tool for the message modification in the first round. To be more precise, we use the same tool as for the differential path search in the first round. Instead of searching for valid differential characteristics in both streams, we search for valid bit-wise assignments of 0's and 1's to the message and state bits in the first round. Since we solve for conforming message words bit-wise, a different message word permutation, different rotation values and carry effects are handled automatically, similar as in the search for differential characteristics. Moreover, this approach can be generalized to any ARX based design.

The disadvantage of our automated bit-wise approach is a slightly higher complexity, compared to a hand-tuned word-wise approach. However, this increased costs can be amortized by randomizing message words at the end of round 1 to find solutions efficiently for the high-probability characteristic of the remaining rounds.

2.3 Automatic Search Tool

The application of the above strategies is far from being trivial and requires an advanced set of techniques and tools to be successful. Due to the increased complexity of dual-stream hash functions with different streams, finding good differential characteristics by hand is almost impossible. Therefore, we have developed an automatic tool which can be used for finding complex nonlinear differential characteristics as well as for solving nonlinear equations involving conditions on state words and free message bits, i.e. to find confirming message pairs. Our tool is based on the approach of Mendel et al. [10] to find both complex nonlinear differential characteristics and conforming message pairs for SHA-2.

The basic idea is to consider differential characteristics which impose arbitrary conditions on pairs of bits using generalized conditions [4]. Generalized conditions are inspired by signed-bit differences and take all 16 possible conditions on a pair of bits into account. Table 2 lists all these possible conditions and introduces the notation for the various cases.

Table 2. Notation for possible generalized conditions on a pair of bits [4]

(X_i, X_i^*)	(0, 0)	(1, 0)	(0, 1)	(1, 1)	(X_i, X_i^*)	(0, 0)	(1, 0)	(0, 1)	(1, 1)
?	✓	✓	✓	✓	3	✓	✓	-	-
-	✓	-	-	✓	5	✓	-	✓	-
x	-	✓	✓	-	7	✓	✓	✓	-
0	✓	-	-	-	A	-	✓	-	✓
u	-	✓	-	-	B	✓	✓	-	✓
n	-	-	✓	-	C	-	-	✓	✓
1	-	-	-	✓	D	✓	-	✓	✓
#	-	-	-	-	E	-	✓	✓	✓

Using these generalized conditions and propagating them in a bitsliced manner, we can construct complex differential characteristics in an efficient way. The basic idea of the search algorithm is to randomly pick a bit from a set of bit positions with predefined conditions, impose a more restricted condition and compute how this new condition propagates. This is repeated until an inconsistency is found or all unrestricted bits from the set are eliminated. Note that this general approach can be used for both, finding differential characteristics and conforming message pairs.

For example, the search strategy for finding nonlinear characteristics works as follows (for a more detailed description of the search algorithm or how the conditions are propagated we refer to [4, 10]):

1. Define a set of unrestricted bits (?) and unsigned differences (**x**).
2. Pick a random bit from the set.
3. Impose a zero-difference (-) on unrestricted bits (?), or randomly choose a sign (**u** or **n**) for unsigned differences (**x**).
4. Check how the new conditions propagate.
5. If an inconsistency occurs, remember the last bit and jump back until this bit can be restricted without leading to a contradiction.
6. Repeat from step 2 until all bits from the set have been restricted.

We use the same strategy to find conforming input pairs for a given differential characteristic. Instead of picking an unrestricted bit (?) we pick an undetermined bit without difference (-) and assign randomly a value (0 or 1) until a solution is found:

1. Define a set of undetermined bits without difference (-).
2. Pick a random bit from the set.
3. Randomly choose the value of the bit (0 or 1).
4. Check how the new conditions propagate.
5. If an inconsistency occurs, remember the last bit and jump back until this bit can be restricted without leading to a contradiction.
6. Repeat from step 2 until all bits from the set have been restricted.

Note that the efficiency of finding a conforming message pair can be increased if the undetermined bits without difference (–) are picked in a specific order. The order strongly depends on the specific hash function. In general, fully determining word after word turns out to be a good approach for word-wise defined ARX-based hash functions. Using this approach, we can instantly (milliseconds) find solutions for the first round of dual-stream hash functions without the need for hand-tuned advanced message modification techniques.

3 Description of RIPEMD-128

RIPEMD-128 was designed by Dobbertin, Bosselaers and Preneel in [6] as a replacement for RIPEMD. It is an iterative hash functions based on the Merkle-Damg ard design principle [3,12] and processes 512-bit input message blocks and produces a 128-bit hash value. To guarantee that the message length is a multiple of 512 bits, an unambiguous padding method is applied. For the description of the padding method we refer to [6].




Fig. 1. Structure of the RIPEMD-128 compression function

Like its predecessor, the function of RIPEMD-128 consists of two parallel streams. In each stream the state variables are updated corresponding to the message block and combined with the previous chaining value after the last step, depicted in Figure 1. While RIPEMD consists of two parallel streams of MD4, the two streams are designed differently in the case of RIPEMD-128. In the following, we describe the compression function in detail.

Each stream of the compression function of RIPEMD-128 basically consists of two parts: the state update transformation and the message expansion. Furthermore, RIPEMD-128 consists of a feed-forward where the input and output state words are added in a different order. For a detailed description we refer to [6].

State Update Transformation. The state update transformation of each stream starts from a (fixed) initial value IV of four 32-bit words $B_{-4}, B_{-3}, B_{-2}, B_{-1}$. and updates them in 4 rounds of 16 steps each. In each step one message word is used to update the four state variables. Figure 2 shows one step of the state update transformation of each stream of RIPEMD-128.




Fig. 2. The step update transformation of RIPEMD-128

The function f is different in each round. f_r is used for the r -th round in the left stream, and f_{5-r} is used for the r -th round in the right stream ($r = 1, \dots, 4$):

$$\begin{aligned}
 f_1(x, y, z) &= x \oplus y \oplus z, \\
 f_2(x, y, z) &= (x \wedge y) \vee (\neg x \wedge z), \\
 f_3(x, y, z) &= (x \vee \neg y) \oplus z, \\
 f_4(x, y, z) &= (x \wedge z) \vee (y \wedge \neg z).
 \end{aligned}$$

A step constant K_r is added in every step; the constant is different for each round and for each stream. For the actual values of the constants we refer to [6], since we do not need them in the analysis. For both streams the following rotation values s given in Table 3 are used.

Table 3. The rotation values s for each step and each stream of RIPEMD-128

	Step	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
left stream	Round 1	11	14	15	12	5	8	7	9	11	13	14	15	6	7	9	8
	Round 2	7	6	8	13	11	9	7	15	7	12	15	9	11	7	13	12
	Round 3	11	13	6	7	14	9	13	15	14	8	13	6	5	12	7	5
	Round 4	11	12	14	15	14	15	9	8	9	14	5	6	8	6	5	12
right stream	Round 1	8	9	9	11	13	15	15	5	7	7	8	11	14	14	12	6
	Round 2	9	13	15	7	12	8	9	11	7	7	12	7	6	15	13	11
	Round 3	9	7	15	11	8	6	6	14	12	13	5	14	13	13	7	5
	Round 4	15	5	8	11	14	14	6	14	6	9	12	9	12	5	15	8

Message Expansion. The message expansion of RIPEMD-128 is a permutation of the 16 message words in each round. Different permutations are used for the left and the right stream. For both streams the message words are permuted according to Table 4.

Table 4. The index of the message words m_i which are used as the expanded message words W_i in each step and each stream of RIPEMD-128

	Step	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
left stream	Round 1	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
	Round 2	7	4	13	1	10	6	15	3	12	0	9	5	2	14	11	8
	Round 3	3	10	14	4	9	15	8	1	2	7	0	6	13	11	5	12
	Round 4	1	9	11	10	0	8	12	4	13	3	7	15	14	5	6	2
right stream	Round 1	5	14	7	0	9	2	11	4	13	6	15	8	1	10	3	12
	Round 2	6	11	3	7	0	13	5	10	14	15	8	12	4	9	1	2
	Round 3	15	5	1	3	7	14	6	9	11	8	12	2	10	0	4	13
	Round 4	8	6	4	1	3	11	15	0	5	12	2	13	9	7	10	14

Feed-Forward. After the last step of the state update transformation, the initial values B_{-4}, \dots, B_{-1} and the output values of the last step of the left stream B_{63}, \dots, B_{60} and the last step of the right stream B'_{63}, \dots, B'_{60} are combined, resulting in the final value of one iteration (feed-forward). The result is the final hash value or the initial value for the next message block:

$$\begin{aligned}
 & B_{-1} \boxplus B_{62} \boxplus B'_{61} \\
 & B_{-4} \boxplus B_{63} \boxplus B'_{62} \\
 & B_{-3} \boxplus B_{60} \boxplus B'_{63} \\
 & B_{-2} \boxplus B_{61} \boxplus B'_{60}
 \end{aligned}$$

4 Collision Attacks on RIPEMD-128

To find collisions in reduced RIPEMD-128 we use the strategy proposed in Sect. 2.2. The attack consists of 3 major parts given as follows:

1. **Starting Point:** Find a good start setting, i.e. differences in only a few specific message words that may lead in a differential characteristic with high probability after step 15.
2. **Differential Characteristic:** Search for a high-probability differential characteristic for the whole hash function where at most one stream has a low probability in step 0-15.
3. **Message Pair:** Find a colliding message pair using automated message modification and random trials.

4.1 Finding a Starting Point

In MD4-like hash functions, differences are introduced and canceled using differences in the expanded message words. Since RIPEMD-128 has two streams with different permutation of message words, the first step in the attack is to determine those message words which may contain differences. We have several constraints such that the whole attack can be carried out efficiently.

First of all, we aim for a high probability differential characteristics after step 15 in both streams. Such high probability differential characteristics can be constructed if the differences introduced by the message words are canceled immediately using local collisions spanning over only a few steps. The shortest local collision in the MD4 step update goes over 4 steps. However, due to the different message permutation used in each stream, it is difficult to achieve short local collisions in both streams simultaneously.

Another possibility is to cancel all differences in each stream as early as possible in round 2 and find message words, such that new differences are introduced late in round 3. A further constraint is to have a short local collision and hence sparse differential characteristic in one stream between step 0-15 such that the message modification part can be carried out more efficiently (see Sect. 2.2).

A single message word which seems to be a good choice is m_{13} . In this case, we get one short local collision between round 1 and round 2 in the left stream and one slightly longer local collision between round 1 and round 2 in the right stream. Both local collisions end in the first few steps of round 2. Furthermore, the message word m_{13} introduces differences very late in the last few steps of round 3 (see Fig. 3). Note that a similar approach was used by Dobbertin in the attack on RIPEMD [5]. Unfortunately, no local collision spanning over 5 steps in the left stream between round 1 and 2 can be constructed which renders the attack impossible.




Fig. 3. Using only message word m_{13}

A better choice is to use differences in two message words, like it was done by Wang et al. in the attack on RIPEMD [17]. If we choose differences in m_0 and m_6 then we get for the left stream one local collision over 6 steps in round 1, and another local collision over 4 steps in round 2. Note that in the right stream a short local collision over 4 steps (step 16-20) is actually impossible. This is

due to the fact that for f_3 (ONX-function), a local collision over 4 steps with differences in only two message does not exist. Hence, we combine in the right stream the two local collisions resulting in one long local collision between step 3 and 20. In round 3, the first difference is added in step 38. Hence, using this starting point we can get a collision for 38 steps of RIPEMD-128.




Fig. 4. Using message words m_0 and m_6

4.2 Finding a Differential Characteristic

Once we have fixed the starting point, i.e. the message words which may contain differences, we use an automated tool to find high-probability differential characteristics. Note that we do not fix the message difference prior to the search to allow the tool to find an optimal solution.

In order to get a differential characteristics resulting in a low attack complexity, we aim for a low Hamming weight difference in state word B_{21} . The best we could find is a differential characteristic with 2 differences in B_{21} (see Table 8). Furthermore, the Boolean function XOR in the first round of the left stream provides less freedom in constructing local collisions than the non-linear functions. Hence, we first search for a differential characteristic in the left stream.

Once the characteristic in the left stream is fixed, we use an arbitrary first message block to fulfill the conditions on the chaining value. Since we have 14 conditions on the chaining value (see Table 8), finding the 1st block has a complexity of about 2^{14} .

Next, we search for a differential characteristic in the right stream. To get a low complexity for the message search in round 2, we search for characteristics with only a few differences in state words B'_{14} and B'_{15} . Using our search tool, we can find many differential characteristic for the left and right stream within only a few minutes on an ordinary PC. A colliding differential characteristic for 38 steps of RIPEMD-128 is given in Table 8.

4.3 Finding a Confirming Message Pair

To fulfill all conditions imposed by the differential characteristic in the first round, we need to apply message modification techniques. Since we have many conditions in the first 6 steps of the left stream and the first 15 steps of the right stream this may not be an easy task. However, using our tool and generalized conditions, we can do message modification for the first 16 steps efficiently

and immediately within milliseconds on a PC. Of course, by hand-tuning basic message modification the complexity might be improved, but using our tool this phase of the message search can be fully automated. Furthermore, the cost of message modification is fully amortized by randomizing e.g. message word m_{12} to find a solution also for the high-probability characteristic in round 2 (and 3). Using the approximately 2^{30} possible value for m_{12} , we can find a solution for the differential characteristic (complexity 2^{14} after round 1) including message modification in less than a second on our PC. The resulting message pair for a collision on 38 steps of RIPEMD-128 is given in Table 5.

Table 5. Collision for 38 steps of RIPEMD-128

M_1	9431bddf 7b9827d6 f54a64a9 df41a58a fd707a50 dad10eb6 48b0cc76 be66cb8c ab3b7afa 084ba98e ab0a4798 2a4b0d06 a79bf8b7 3fd6008a 4da2112d 849c5b9c
M_2	952bc70f d0840848 eaafffa57 0ca3c38a 45383ffb ddc6a9a1 796f1e20 0b9ff55f ddb80113 f0ffe1b5 b7d75dc0 82c7298f f2c442f4 96cbf293 c441d662 06e9eec2
M_2^*	952bc50e d0840848 eaafffa57 0ca3c38a 45383ffb ddc6a9a1 79ef1e21 0b9ff55f ddb80113 f0ffe1b5 b7d75dc0 82c7298f f2c442f4 96cbf293 c441d662 06e9eec2
ΔM_2	00000201 00000000 00000000 00000000 00000000 00000000 00800001 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
H_2	a0a00507 fd4c7274 ba230d53 87a0d10a
H_2^*	a0a00507 fd4c7274 ba230d53 87a0d10a
ΔH_2	00000000 00000000 00000000 00000000

5 Extending the Attack to More Steps

In this section, we will show how the collision attack on 38 steps can be extended to more steps of the hash function by using a weaker attack setting, i.e. near-collisions and subspace distinguisher. Furthermore, we present a free-start collision for 48 steps of RIPEMD-128 compression function.

5.1 Near-Collisions for the Hash Function

It is easy to see that by appending 6 steps to the characteristic for 38 steps one gets a near-collision for 44 steps of the hash function with only 6 differences in the hash value. However, note that while in the collision attack one can always append a message block with the correct padding this can not be done for a near-collision. Hence, in order to construct a near-collision for the hash function the padding has to be fixed on beforehand. Luckily, we have such a high amount of freedom in our attack the we can easily fix m_{15} , m_{14} and parts of m_{13} in the attack to guarantee that the padding is correct. The result is a practical near-collision (see Table 6) for 44 steps of RIPEMD-128 with complexity of 2^{32} . Note that the generic attack to find a near-collision with only 6 differences in the hash value has a complexity of about $2^{47.8}$.

Table 6. Near-collision for 44 steps of RIPEMD-128

M_1	2ca95052 425a8f73 08be4537 c790e019 0dcc7d4e 29075123 75327262 8d0d4803 1e57a6a4 73550688 59263eb1 98c6f6ce f03b8b4b 62d3fdf7 638db196 68c0b7b3
M_2	aa1437ef f3646663 c339343a 52c43a1a 779995d5 7b6bd784 e927bb74 5e7cb217 7af2ac15 93392ccf 07e847cf 86318b70 d9d33105 809693dd 000003b8 00000000
M_2^*	aa1435ee f3646663 c339343a 52c43a1a 779995d5 7b6bd784 e9a7bb75 5e7cb217 7af2ac15 93392ccf 07e847cf 86318b70 d9d33105 809693dd 000003b8 00000000
ΔM_2	00000201 00000000 00000000 00000000 00000000 00000000 00800001 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
H_2	92dd7ef7 b1f15ee4 b3e6a250 9db2131b
H_2^*	929d5ef7 b1f15ee4 b3e6a250 bdb21b5f
ΔH_2	00402000 00000000 00000000 20000844

5.2 Non-randomness for the Hash Function

In this section, we show non-random properties for 48 steps (3 rounds) of the hash function. It is based on the differential q -multicollision distinguisher and the differential characteristic for 44 steps which is extended to 48 steps.

Differential q -multicollisions were introduced by Biryukov et al. in the cryptanalysis of the block cipher AES-256 [1]. Note that in [1] the attack is described for a block cipher. However, it can be easily adapted for a hash function. Below we repeat the basic definition and lemma, we need for the attack on RIPEMD-128.

Definition 1. *A set of one difference and q inputs*

$$\{\Delta M; (M^1), (M^2), \dots, (M^q)\}$$

is called a differential q -multicollision for $h(\cdot)$ if

$$\begin{aligned} h(M^1) \oplus h(M^1 \oplus \Delta M) &= h(M^2) \oplus h(M^2 \oplus \Delta M) \\ &= \dots = h(M^q) \oplus h(M^q \oplus \Delta M). \end{aligned}$$

The complexity of the generic attack is measured in the number of queries.

Lemma 1. *To construct a differential q -multicollision for an ideal has function with an n -bit output an adversary needs at least*

$$O(q \cdot 2^{\frac{q-1}{q+1} \cdot n})$$

queries on the average for small q .

The proof for Lemma 1 works similar as in [1] for an ideal cipher. Finally, we construct a differential q -multicollision to show non-random properties for RIPEMD-128 reduced to 48 steps. The attack has a complexity of about $4 \cdot 2^{68}$ while the generic attack has a complexity of about 2^{76} .

5.3 Collisions for the Compression Function

When attacking the compression function an adversary has additional the possibility to inject difference in the chaining input. Using this additional freedom

and the same techniques as for the collision attack on the RIPEMD-128 hash function (see Section 4), we can construct a collision for the compression function of RIPEMD-128 reduced to 48 steps. In Table 9 the differential characteristic is shown, resulting in a practical collision for 48 steps of the compression function with a complexity of 2^{40} . The example is given in Table 7.

Table 7. Free-start collision for 48 steps of RIPEMD-128

H_0	5a1d2fbd cd6d40c7 128dd546 900e0e65
H_0^*	5a1927bd edad5cc7 128dd542 900e0e65
ΔH_0	00040800 20c01c00 00000004 00000000
M_1	06083719 9ae0b19b 7ffae1ec 637041ad 28d722d7 fa0082c3 5e78f84e 416ee5e7 faf2b4fc 56738a9f 363c6155 cc7d7ae3 0cb5fc95 b362a16f 6cac81a9 cc11fedd
M_1^*	06083719 9ae0b19b 7ffae1ec 637041ad 28d722d7 fa0082c3 5e78f84e 416ee5e7 faf2b4fc 56738a9f 363c6155 cc7d7ae3 0cb5fc95 b362a16f 6cac81a9 cc11fedd
ΔM_1	00000200 00000000 00000000 00000000 00000000 00000000 00000001 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
H_1	e6428c57 a9f1f589 fc045baf a9cdbc1f
H_1^*	e6428c57 a9f1f589 fc045baf a9cdbc1f
ΔH_1	00000000 00000000 00000000 00000000

6 Conclusions and Future Work

In this work, we have presented new results on the ISO/IEC standard RIPEMD-128, a dual-stream hash function where the message permutation and rotation values are different in the two streams. More specifically, we have presented a collision attack on reduced RIPEMD-128 and get practical collisions for 38 steps of the hash function with a complexity of about 2^{14} . Furthermore, our attack can be extended to near-collisions on 44 steps with complexity 2^{32} and a theoretical distinguisher on the hash function for 48 steps (3 out of 4 rounds) with complexity 2^{70} . Furthermore, we present practical collisions for the RIPEMD-128 compression function, also reduced to 48 steps with complexity 2^{40} .

Apart from these new results, we have outlined a strategy to analyze ARX-based dual-stream hash functions more efficiently. More precisely, we have shown how to automate the most difficult parts of an attack involving more than one stream: finding a differential characteristic and performing message modification in the first round. In particular, message modification had to be hand-tuned or was omitted in previous attacks on ARX-based hash functions. What remains for an attacker is to determine a good starting point (possibly using tools from coding theory) and to assist the tools in the order of guessing words or parts of the state, to improve the overall complexity.

Ideally, these tools can immediately be applied to more complicated hash functions. However, the obtained results depend mainly on the choice of the starting point for the nonlinear tool. If no good starting point can be found or the search space is too large, no attack can be obtained. Future work is to analyze also other, stronger dual-stream hash functions like RIPEMD-160. Furthermore,

the tools and techniques used in this paper can also be applied to other ARX-based hash functions, where more than one state word is updated using a single message word. Examples are SHA-2 or the SHA-3 candidates Blake and Skein.

Acknowledgments. This work was supported in part by the Research Council KU Leuven: GOA TENSE (GOA/11/007), by the IAP Programme P6/26 BCRYPT of the Belgian State (Belgian Science Policy) and by the European Commission through the ICT programme under contract ICT-2007-216676 ECRYPT II. In addition, this work was supported by the Research Fund KU Leuven, OT/08/027 and by the Austrian Science Fund (FWF, project P21936).

References

1. Biryukov, A., Khovratovich, D., Nikolić, I.: Distinguisher and Related-Key Attack on the Full AES-256. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 231–249. Springer, Heidelberg (2009)
2. Brier, E., Khazaei, S., Meier, W., Peyrin, T.: Linearization Framework for Collision Attacks: Application to CubeHash and MD6. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 560–577. Springer, Heidelberg (2009)
3. Damg ard, I.B.: A Design Principle for Hash Functions. In: Brassard, G. (ed.) CRYPTO 1989. LNCS, vol. 435, pp. 416–427. Springer, Heidelberg (1990)
4. De Canni ere, C., Rechberger, C.: Finding SHA-1 Characteristics: General Results and Applications. In: Lai, X., Chen, K. (eds.) ASIACRYPT 2006. LNCS, vol. 4284, pp. 1–20. Springer, Heidelberg (2006)
5. Dobbertin, H.: RIPEMD with Two-Round Compress Function is Not Collision-Free. *J. Cryptology* 10(1), 51–70 (1997)
6. Dobbertin, H., Bosselaers, A., Preneel, B.: RIPEMD-160: A Strengthened Version of RIPEMD. In: Gollmann, D. (ed.) FSE 1996. LNCS, vol. 1039, pp. 71–82. Springer, Heidelberg (1996)
7. Indestege, S., Preneel, B.: Practical Collisions for EnRUPT. In: Dunkelman, O. (ed.) FSE 2009. LNCS, vol. 5665, pp. 246–259. Springer, Heidelberg (2009)
8. Joux, A., Peyrin, T.: Hash Functions and the (Amplified) Boomerang Attack. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 244–263. Springer, Heidelberg (2007)
9. Kl ima, V.: Tunnels in Hash Functions: MD5 Collisions Within a Minute. *IACR Cryptology ePrint Archive* 2006, 105 (2006)
10. Mendel, F., Nad, T., Schl affer, M.: Finding SHA-2 Characteristics: Searching through a Minefield of Contradictions. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 288–307. Springer, Heidelberg (2011)
11. Mendel, F., Pramstaller, N., Rechberger, C., Rijmen, V.: On the Collision Resistance of RIPEMD-160. In: Katsikas, S.K., L opez, J., Backes, M., Gritzalis, S., Preneel, B. (eds.) ISC 2006. LNCS, vol. 4176, pp. 101–116. Springer, Heidelberg (2006)
12. Merkle, R.C.: One Way Hash Functions and DES. In: Brassard, G. (ed.) CRYPTO 1989. LNCS, vol. 435, pp. 428–446. Springer, Heidelberg (1990)
13. Ohtahara, C., Sasaki, Y., Shimoyama, T.: Preimage Attacks on Step-Reduced RIPEMD-128 and RIPEMD-160. In: Lai, X., Yung, M., Lin, D. (eds.) *Inscrypt* 2010. LNCS, vol. 6584, pp. 169–186. Springer, Heidelberg (2011)

14. Pramstaller, N., Rechberger, C., Rijmen, V.: Exploiting Coding Theory for Collision Attacks on SHA-1. In: Smart, N.P. (ed.) *Cryptography and Coding 2005*. LNCS, vol. 3796, pp. 78–95. Springer, Heidelberg (2005)
15. Sugita, M., Kawazoe, M., Imai, H.: Gröbner Basis Based Cryptanalysis of SHA-1. *IACR Cryptology ePrint Archive* 2006, 98 (2006)
16. Wang, L., Sasaki, Y., Komatsubara, W., Ohta, K., Sakiyama, K.: (Second) Preimage Attacks on Step-Reduced RIPEMD/RIPEMD-128 with a New Local-Collision Approach. In: Kiayias, A. (ed.) *CT-RSA 2011*. LNCS, vol. 6558, pp. 197–212. Springer, Heidelberg (2011)
17. Wang, X., Lai, X., Feng, D., Chen, H., Yu, X.: Cryptanalysis of the hash functions MD4 and RIPEMD. In: Cramer, R. (ed.) *EUROCRYPT 2005*. LNCS, vol. 3494, pp. 1–18. Springer, Heidelberg (2005)
18. Wang, X., Yao, A., Yao, F.: New Collision Search for SHA-1. Presented at rump session of *CRYPTO* (2005)
19. Wang, X., Yin, Y.L., Yu, H.: Finding Collisions in the Full SHA-1. In: Shoup, V. (ed.) *CRYPTO 2005*. LNCS, vol. 3621, pp. 17–36. Springer, Heidelberg (2005)
20. Wang, X., Yu, H.: How to Break MD5 and Other Hash Functions. In: Cramer, R. (ed.) *EUROCRYPT 2005*. LNCS, vol. 3494, pp. 19–35. Springer, Heidelberg (2005)

A Differential Characteristics and Conditions

Table 8. Characteristic for a collision on 38 steps of RIPEMD-128. Bits with gray background have one additional conditions.

i	∇B_i	$\nabla B'_i$	∇m_i
-4	-----		
-3	-----		
-2	-----		
-1	-----		
0	unnnnnnnnnnnnnnnnnnnnn		u-----u
1	n-----n-----nnnnnnnnnn	-0-----0-----	
2	unnnnnnnnnnnnnnnnnnnnn	-0-----0-----	
3	-----	--0100--u-----u--0110--	
4	-----	--1101--1-1-----1--1111--	--0-----
5	-----	--unnn00--1-1-----1--unnn-00	
6	-----	--000010--n-u--00--n--0111-10	--n-----n
7	-----	001nnuuu--0--11111--1001-nu	
8	-----	110100--1--un11n--u	
9	-----	un1n00--1--unn--1--1--	
10	-----	-n0u1-----0-10000--1--	
11	-----	--0nuu-----01n11--n--	
12	-----	--110-----nuuu--	
13	-----	--01-----11-1--	
14	-----	-----00-1--0	
15	-----	-----n--n	
16	-----	-----n--n	
17	-----	-----0--0	
18	-----		
19	-----		
20	-----		
21	-----		
22	-----	-----0--0	
23	-----	-----1--1	
24	-----		
25	-----		
26	-----		
27	-----		
28	-----		
29	-----		
30	-----		
31	-----		
32	-----		
33	-----		
34	-----		
35	-----		
36	-----		
37	-----		

Table 9. Characteristic for a free-start collision for 48 steps of RIPEMD-128 compression function. Bits with gray background have one additional conditions.

i	∇B_i	$\nabla B'_i$	∇m_i
-4	-----u-----u-----		
-3	-0-----00-----011-----1-		
-2	-00-00--10-011--101--10--u--		
-1	-1n-11--nu-011--nnn--10--1-1		
0	un1nnnnn00nu01--un101nuunnnn0n0	010-1u----nuu--1-101101-010-1-1	-----0-----u-----1
1	nnnnnnnnnnnnnnnnnn--010--01--n-u	1nn-n1----n00-01-110110-n11-00u	-----100-----
2	-0-----10--unnnnnnnnnnnnnnnnnnn	u1n--1000-1n10-0n-un-nnn-unu--1	-----11-----0--1-----
3	-1-----00-----110--10--0	0n0--n1111-01u-u--01uu--1--nu1	-----1-----1
4	-----110--10--1	u11--unn0u11111--001--10--0nu	-----110--111
5		u1--011uuun010-0-1nu--0-01u1	
6		0u--10u11uu0n-1-n--10u--0u	-----n
7		00--n1n0101-n-0--111--11	-----0-----
8		-1--0unnnnn0000u0000un1--0	-----0-----
9		110--n11u10111-10n00	0-----1-----0-----
10		-01--0unnnnnnnn1u011	-----
11		1011--nuuuuu	-----0-----
12		100--010	-----
13		101--	-----11
14		0-	-----
15		-n	-----1-----
16		-n	-----
17		-0-	-----
18			-----
19			-----
20			-----
21			-----
22			-----
23			-----
24			-----
25			-----
26			-----
27			-----
28			-----
29			-----
30			-----
31			-----
32			-----
33			-----
34			-----
35			-----
36			-----
37			-----
38			-----
39			-----
40			-----
41			-----
42			-----
43			-----
44			-----
45			-----
46			-----
47			-----

Bicliques for Preimages: Attacks on Skein-512 and the SHA-2 Family^{*}

Dmitry Khovratovich¹, Christian Rechberger², and Alexandra Savelieva³

¹ Microsoft Research Redmond, USA

² DTU MAT, Denmark

³ National Research University Higher School of Economics, Russia

Abstract. We present a new concept of biclique as a tool for preimage attacks, which employs many powerful techniques from differential cryptanalysis of block ciphers and hash functions.

The new tool has proved to be widely applicable by inspiring many authors to publish new results of the full versions of AES, KASUMI, IDEA, and Square. In this paper, we show how our concept leads to the first cryptanalysis of the round-reduced Skein hash function, and describe an attack on the SHA-2 hash function with more rounds than before.

Keywords: SHA-2, SHA-256, SHA-512, Skein, SHA-3, hash function, meet-in-the-middle attack, splice-and-cut, preimage attack, initial structure, biclique.

1 Introduction

Major breakthrough in preimage attacks on hash functions happened in 2008 when the so-called *splice-and-cut* framework was introduced. Its applications to MD4 and MD5 [2,25], and later to Tiger [11] brought amazing results. Internal properties of message schedule appeared to be limiting application of this framework to SHA-x family [1,3]. However, the concept of *biclique* introduced in this paper allows to mitigate such obstacles, as demonstrated by our results.

This is the first work on the concept of biclique cryptanalysis. We focus on new definitions and algorithms and concentrate on the hash function setting. As applications, we present an attack on the Skein hash function (the only one existing so far) and the SHA-2 family. Our findings are summarized in Table 1.

Splice-and-cut framework and its progress. Both splice-and-cut and meet-in-the-middle attacks exploit the property that a part of a primitive does not make use of particular key/message bits (called *neutral bits*). If the property holds, the computation of this part remains the same when we flip those bits in the other

^{*} This work was supported by the European Commission under contract ICT-2007-216646 (ECRYPT II) and the Federal Target Program “Scientific and scientific-pedagogical personnel of innovative Russia“ in 2009-2013 under contract No. P965 from 27 May, 2010.

part of a primitive. Assume that neutral bits can be found for both parts (also called *chunks*). Then a cryptanalyst prepares a set of independent computations for all possible values of those bits and subsequently checks for a match in the middle. Efficiency of the attack depends on the number of neutral bits.

Sasaki and Aoki observed [2,25] that compression functions with permutation-based message schedule are vulnerable to this kind of attack as chunks can be long. They proposed several ways to improve splice-and-cut framework, including a very interesting trick referred to as *initial structure* [3,26]. It can be informally defined as an overlapping of chunks, where neutral bits, although formally belonging to both chunks, are involved in the computation of the proper chunk only. In our work we aimed to explore the potential of initial structure.

Our Contributions. We replace the idea of the initial structure with a more formal and general concept of *biclique* that provides us with a new level of understanding. In terms of graph theory, bicliques are structures represented by two sets of states with each state having a relation with all states in another set. We derive a system of functional equations linking internal states several rounds apart, and show that it is equivalent to a system of differentials, so the full structure of states can be built out of a set of trails. The differential view allows us to apply numerous tools from collision search and differential cryptanalysis. We propose three generic and flexible algorithms for constructing the bicliques.

Our simple example of biclique application is an attack on round-reduced Skein-512 hash function, the SHA-3 finalist which currently lacks any other attacks in the hash setting. We penetrate 22 rounds of Skein-512, which is comparable to the best attacks on the compression function that survived the last tweak. Our attack on the compression function of Skein-512 covers 37 rounds.

Our second group of applications is the SHA-2 family. We heavily use differential trails in SHA-2, message modification techniques from SHA-1 and SHA-0, trail backtracking techniques from RadioGatun, Grindahl, SHA-1, etc., to build attacks on 45-round SHA-256 and 50-round SHA-512 (both the best attacks in the hash mode). For the compression functions, we penetrate up to 7 more rounds, reaching 52 rounds and violating the security of about 80% of SHA-256.

Table 1. New (second) preimage attacks on Skein-512 and the SHA-2 family

Reference	Target	Steps	Complexity			Memory (words)
			Pseudo-preimage	Second Preimage	Preimage	
Section 4	Skein-512	22	2^{508}	2^{511}	-	2^6
Section 6	Skein-512	37	$2^{511.2}$	-	-	2^{64}
[1,11]	SHA-256	43	$2^{251.9}$	$2^{254.9}$	$2^{254.9}$	2^6
Section 5	SHA-256	45	2^{253}	$2^{255.5}$	$2^{255.5}$	2^6
Section 6	SHA-256	52	2^{255}	-	-	2^6
[1,11]	SHA-512	46	2^{509}	$2^{511.5}$	$2^{511.5}$	2^6
Section 5	SHA-512	50	2^{509}	$2^{511.5}$	$2^{511.5}$	2^4
Section 6	SHA-512	57	2^{511}	-	-	2^6

2 Biclques

In this section we introduce preimage attacks with bicliques. We consider hash functions with block cipher based compression functions $H = E_N(X) \oplus X$, where E is the block cipher keyed with parameter N (notation \xrightarrow{N} and \xleftarrow{N} will be used for E computed in forward and backward direction respectively). Depending on the design, parameters (N, X) will be either (M, CV) for the most popular Davies-Meyer mode, or (CV, M) for Matyas-Meyer-Oseas mode, where CV is the chaining variable and M is the message.

Let f be a sub-cipher of E , and $\mathcal{N} = \{N[i, j]\}$ be a group of parameters for f . Then a *biclique of dimension d* over f for \mathcal{N} is a pair of sets $\{Q_i\}$ and $\{P_j\}$ of 2^d states each such that

$$Q_i \xrightarrow[f]{N[i, j]} P_j. \quad (1)$$

A biclique is used in the preimage search as follows (Figure 1). First, we note that if $N[i, j]$ yeilds a preimage, then

$$E : X \xrightarrow{N[i, j]} Q_i \xrightarrow[f]{N[i, j]} P_j \xrightarrow{N[i, j]} H.$$

An adversary selects a variable v outside of f (w.l.o.g. between P_j and H) and checks, for appropriate choices of sub-ciphers g_1 and g_2 , if

$$\exists i, j : P_j \xrightarrow[g_1]{N[i, j]} v \stackrel{?}{=} v \xleftarrow[g_2]{N[i, j]} Q_i.$$

A positive answer yields a candidate preimage. Here, to compute v from Q_i , the adversary first computes X and then derives the output of E as $X \oplus H$.

To benefit from the meet-in-the-middle framework, the variable v is chosen so that g_1 and g_2 are independent of i and j respectively:

$$P_j \xrightarrow[g_1]{N[*, j]} v \stackrel{?}{=} v \xleftarrow[g_2]{N[i, *]} Q_i.$$

Then the complexity of testing 2^{2d} messages for preimages is computed as follows:

$$C = 2^d(C_{g_1} + C_{g_2}) + C_{bicl} + C_{recheck},$$

where C_{bicl} is the biclique construction cost, and $C_{recheck}$ is the complexity of rechecking the remaining candidates on the full state. We explain how to amortize the biclique construction in the next section. Clearly, one needs 2^{n-2d} bicliques of dimension d to test 2^n parameters.

3 Biclique Construction Algorithms

Here we introduce several algorithms for the biclique construction. They differ in complexity and requirements to the dimension of a biclique and properties of the mapping f .




Fig. 1. Biclique of dimension 2 in the meet-in-the-middle attack on a Davies-Meyer compression function

Consider a single mapping in Equation (1)

$$Q_0 \xrightarrow[f]{N[0,0]} P_0. \tag{2}$$

We call this a *basic computation*. Consider the other mappings as differentials to the basic computation:

$$\nabla_i \xrightarrow[f]{\Delta_{i,j}^N} \Delta_j, \tag{3}$$

so that

$$Q_i = Q_0 \oplus \nabla_i, \quad P_j = P_0 \oplus \Delta_j, \quad N[i,j] = N[0,0] \oplus \Delta_{i,j}^N.$$

Vice versa, if a computation (2) is a solution to 2^{2d} differentials in (3), then it is a basic computation for a biclique.

The algorithms presented below allow us to reduce the number of differentials needed for a biclique, and hence construct a biclique efficiently.

Algorithm 1. Let the differences in the set \mathcal{N} be defined as the following linear function:

$$\Delta_{i,j}^N = \Delta_j^N \oplus \nabla_i^N \tag{4}$$

Let us fix Q_0 and construct P_j as follows:

$$Q_0 \xrightarrow[f]{N[0,j]} P_j. \tag{5}$$

As a result, we get a set of trails:

$$0 \xrightarrow[f]{\Delta_j^N} \Delta_j. \tag{6}$$

Let us also construct Q_i out of P_0 :

$$Q_i \xleftarrow[f]{N[i,0]} P_0, \tag{7}$$

and get another set of trails:

$$\nabla_i \xleftarrow[f]{\nabla_i^N} 0. \tag{8}$$

Suppose that the trails (8) do not affect active non-linear elements in the trails (6). Then Q_i are solutions to the trails (6), so we get the biclique equation:

$$Q_i \xrightarrow[f]{N[i,j]} P_j. \tag{9}$$

Assume that the computation (7) does not affect active non-linear elements in the trails (6) with probability 2^{-t} . The probability that 2^d computations affect no condition is 2^{-t2^d} . Equation (9) is satisfied with probability 2^{-t2^d} , so we need 2^{t2^d} solutions to Equation (6) to build a biclique (feasible for small d).

This algorithm is used in the preimage attack on the hash function Skein-512. For non-ARX primitives with predictable diffusion it can even be made deterministic. For example, for AES [8] and Square it is easy to build truncated differential trails that do not share active non-linear components with probability 1. As a result, the biclique construction can be simply explained using a picture of trails (Figure 2). ■




Fig. 2. Biclique construction out of non-interleaving trails

Algorithm 2. (Modification of Algorithm 1 for the case when the hash function operates in DM mode, and we can control internal state and injections of message M within the biclique). Assume that the mapping f uses several independent parts (blocks) of message M via message injections (like in SHA-2). Consider a message group with property (4) but do not define the messages yet. Choose a state Q_0 satisfying sufficient conditions to build sets of trails (6) and (8) that do not share active non-linear components. Then find $N[0,0]$ such that $Q_0 \xrightarrow[f]{N[0,0]} P_0$ conforms to both sets of trails. Since the sets do not share active non-linear components, we get

$$Q_i \xrightarrow[f]{N[i,j]} P_j,$$

where $Q_i = Q_0 \oplus \nabla_i$, $P_j = P_0 \oplus \Delta_j$.

We can control message injections in f , and therefore, are able to define $N[0, 0]$ block by block similarly to the trail backtracking [5]. The procedure that ensures that the message $N[0, 0]$ is well-defined, and the trails (6) and (8) do not contradict, was first proposed in [1] and referred to as *message compensation*. ■

Algorithm 3. (for bicliques of dimension 1) We apply this rebound-style [20] algorithm if the mapping f is too long for differential trails with reasonable number of sufficient conditions. Then we split it into two parts f_1 and f_2 and consider two differential trails with probabilities p and q , respectively:

$$0 \xrightarrow[f_1]{\Delta^N} \Delta, \quad \nabla \xrightarrow[f_2]{\nabla^N} 0. \tag{10}$$

We fix the state S between f_1 and f_2 , and consider a quartet of states:

$$S, S \oplus \Delta, S \oplus \nabla, S \oplus \Delta \oplus \nabla.$$

Suppose that a quartet of states is a quartet in the middle of the boomerang attack, which happens with probability p^2q^2 for a random N under an appropriate independency assumption. Then we derive input states Q_0, Q_1 and output states P_0, P_1 , which are linked as follows (see also Figure 3):

$$\begin{aligned} Q_0 &\xrightarrow[f_1]{N} S \xrightarrow[f_2]{N} P_0; & Q_0 &\xrightarrow[f_1]{N \oplus \Delta^N} S \oplus \Delta \xrightarrow[f_2]{N \oplus \Delta^N} P_1; \\ Q_1 &\xrightarrow[f_1]{N \oplus \nabla^N} S \oplus \nabla \xrightarrow[f_2]{N \oplus \nabla^N} P_0; & Q_1 &\xrightarrow[f_1]{N \oplus \Delta^N \oplus \nabla^N} S \oplus \Delta \oplus \nabla \xrightarrow[f_2]{N \oplus \Delta^N \oplus \nabla^N} P_1. \end{aligned}$$

Therefore, we get a biclique, where the set of parameters \mathcal{N} is defined as follows:

$$N[0, 0] = N; N[0, 1] = N \oplus \Delta^N; N[1, 0] = N \oplus \nabla^N; N[1, 1] = N \oplus \Delta^N \oplus \nabla^N. \blacksquare$$




Fig. 3. Rebound-style algorithm for biclique construction

We use Algorithm 2 in the attacks on SHA-256 and SHA-512, and Algorithm 3 is applied in the preimage attack on the Skein compression function. In practice, we use freedom in the internal state and in the message injection fulfill conditions in both trails with tools like message modification and auxiliary paths.

4 Simple Case: Second Preimage Attack on Skein-512

Skein [10] is a SHA-3 finalist, and gets a lot of cryptanalytic attention. Differential [4] and rotational cryptanalysis [17] led the authors of Skein to tweak the design twice. As a result, the rotational property, which allowed cryptanalysts to penetrate the highest number of rounds, does not exist anymore in the new version of Skein. Hence the best known attack are near-collisions on up to 24 rounds (rounds 20-43) of the compression function of Skein [4,27]. Very recently near-collisions attacks on up to 32 rounds of Skein-256 were demonstrated [29].

The cryptanalysis of Skein in the hash setting is very limited. Rotational attacks did not extend to the Skein hash function, and the differential attacks were not applied in this model. In our paper, we demonstrate the first attack in this arguably much more relevant setting. At the time of publication this is the only cryptanalytic attack on round-reduced version of Skein hash function.

We chose to give the simplest example in the strongest model rather than to attack the highest number of rounds. The attack we present is on a 22-round version of Skein-512 hash function. In addition to the biclique concept, an interesting feature of our attack is the application of statistical hypothesis test at the matching phase. This technique is applied for the first time and allows to cover more rounds than direct or symbolic (indirect) matching with the same computational complexity.

4.1 Second Preimage Attack on the Reduced Skein-512 Hash

Details of Skein specification are provided in Appendix A. We consider Skein-512 reduced to rounds 3–24. In the hash function setting we are given the message M and the tweak value T , and have to find a second preimage. We produce several pseudo-preimages (CV, M') to a call of the compression function that uses 512 bits of M , and then find a valid prefix that maps the original IV to one of the chaining values that we generated. Let f map the state after round 11 to the state before round 16. We construct a biclique of dimension 3 for f following Algorithm 1 (Section 3):

1. Define $\Delta_j^N = (0, j \ll 58, j \ll 58, 0, 0, 0, 0, 0)$ and $\nabla_i^N = (0, 0, 0, i \ll 55, i \ll 55, 0, 0, 0)$.
2. Generate Q_0 and compute P_0, P_1, \dots, P_7 . If the trails $0 \xrightarrow[f]{\Delta_j^N} \Delta_j$ are not based on the linear difference propagation, repeat the step.
3. Compute Q_i and check if the condition on active non-linear elements is fulfilled. If so, output a biclique.

We use a differential trail that follows a linear approximation that is a variant of the 4-round differential trail, which can be obtained in a similar way to the one presented in the paper [4]. The number of active bits is given in Table 2. Further details of the trail are provided Appendix A in Table 5. For the trails based on

the 3-bit difference Δ_j^N we have 206 sufficient conditions in total. Computations of Q_i out of P_0 do not affect those conditions with probability $2^{-0.3}$ (verified experimentally). Therefore, for the eight states P_j the probability is $2^{-0.3 \cdot 8} \approx 2^{-3}$. We construct a 4-round biclique with complexity at most $2^{206+3} = 2^{209}$. Note that we have $1024 - 209 = 815$ degrees of freedom left.

Table 2. Number of active bits in the most dense Δ -trail in 4 rounds of Skein-512

	I^0	I^1	I^2	I^3	I^4	I^5	I^6	I^7	Conditions in the round
S^{12-A}								3	3
S^{13-A}				6	3				9
S^{14-A}	6		3				3	12	24
S^{15-A}	3	6	3	24	12	6	6	3	63
S^{15-P}	21	9	12	4	3	18	3	37	107 (message addition)

Probabilistic matching. The matching variable v consists of bits 30, 31, 53 of the word 1 after round 24. Due to carry effects, there is a small probability that those bits require the knowledge of the full message to be computed in both directions. This probability was experimentally estimated as 0.09. The matching bits can be computed from both chunks independently with probability 0.91, so with probability $\approx 2^{-0.1}$ we have a type-I error [22], i.e. a false positive, and the candidate is discarded (insisting on probability 1, as in earlier work, would have resulted in an attack on a smaller number of rounds).

Layout of the attack is as follows:

1. Build a biclique of dimension 3 in rounds 12-15 with key additions (key addition + 4 rounds + key addition).
2. Compute forward chunk in rounds 16-19, backward chunks in rounds 8-11, and bits $I_{30,31,53}^1$ of the the state S^{24-P} in both directions in the partial matching procedure.
3. Check for the match in these bits, produce 2^3 key candidates, which get reduced to $2^{2.9}$ due to the type-I error. Check them for the match on the full state.
4. Generate a new biclique out of the first one by change of key bits.
5. Repeat steps 2-5 $2^{507.5}$ times and generate $2^{507.5-509+2.9} = 2^{1.6}$ full pseudo-preimages.
6. Match one of the pseudo-preimages with the real CV_0 .

Complexity. The biclique construction cost can be made negligible, since many bicliques can be produced out of one. Indeed, we are able to flip most of the bits in the message so that the biclique computation between the message injections remain unaffected, and only output states are changed. Every new biclique needs half of rounds 8-11 and 16-19 recomputing, and half of rounds 3-5 and 21-24 computing to derive the value of the matching variable. Hence each biclique tests 2^6 preimage candidates at cost of $(2 + 2 + 1.5) \cdot 8 + (2 + 2 + 2) \cdot 8 = 92$ rounds of

22-round Skein, or $2^{2.3}$ calls of the compression function, taking a recheck into account. As a result, a full pseudo-preimage is found with complexity $2^{508.4}$. We need $2^{1.6} \approx 3$ pseudo-preimages to match one of $2^{510.4}$ prefixes, so the total complexity is $2^{511.2}$.

5 Preimage Attacks on the SHA-2 Hash Functions

The SHA-2 family is the object of very intensive cryptanalysis in the world of hash functions. We briefly review parts of the specification [23] needed for the cryptanalysis in Appendix B. In contrast to its predecessors, collision attacks are no longer the major threat with the best attack on 24 rounds of the hash function [13,24]. So far the best attacks on the SHA-2 family are preimage attacks on the hash function in the splice-and-cut framework [1] and a boomerang distinguisher that is only applicable for the compression function [18].

We demonstrate that our concept of biclique adds two rounds to the attack on SHA-256, four rounds to the attack on SHA-512, and many more when attacking the compression functions. The full layout of our attacks is provided in Table 3. The biclique is based on a 6-round trail with few conditions, easy to use as a ∇ -differential. The number of attacked rounds depends significantly on its position, because:

- message injections in rounds 14-15 are partially determined by the padding rules;
- chunks do not bypass the feedforward operation due to high nonlinearity of the message schedule;
- chunks do not have maximal length, otherwise the biclique trail becomes too dense.

SHA-256. We construct a 6-round biclique with Algorithm 2, Section 3 and place it in rounds 17-22 (see Appendix C for more details of the attack).

SHA-512. The biclique is similar to the one we build for SHA-256. However, our attack on SHA-512 does not fix all the 129 padding bits of the last block. This approach still allows to generate short second preimages by using the first preimage to invest the last block that includes the padding and perform the preimage attack in the last chaining input as the target.

For a preimage attack without a first preimage, expandable messages such as described in [16] can be used. This adds no noticeable cost as the effort is only slightly above the birthday bound.

In addition, the compression function attack needs to fulfill the following two properties. Firstly, the end of the message (before the length encoding, i.e., the LSB of W^{13}) has to be '1'. Secondly, the length needs to be an exact multiple of the block length, i.e., fix the last nine bits of W^{15} to "110111111" (895). In total 11 bits need to be fixed.

Details of the attack are presented in Appendix D.

Table 3. Parameters of the preimage attack on reduced SHA-2 hash functions

Target	Attack layout					
SHA-256 hash function (45-round version)	Biclique					
	Rounds	Dimension	Δ^M bits:	∇^M bits:	Complexity	Freedom used
	17-22	3	$W_{25,26,27}^{17}$	$W_{22,23,31}^{22}$	2^{32}	416
	Message compensation					
	Equations			Constants used in the biclique		
	9			2		
Chunks		Matching				
Forward	Backward	Partial matching	Matching bits	Complexity per match		
2-16	23-36	$37 \rightarrow 38 \leftarrow 1$	$A_{0,1,2,3}^{38}$	2^3		
SHA-512 hash function (50-round version)	Biclique					
	Rounds	Dimension	Δ^M bits:	∇^M bits:	Complexity	Freedom used
	21-26	3	$W_{60,61,62}^{21}$	$W_{53,54,55}^{26}$	2^{32}	832
	Message compensation					
	Equations			Constants used in the biclique		
	9			2		
Chunks		Matching				
Forward	Backward	Partial matching	Matching bits	Complexity per match		
6-20	27-40	$41 \rightarrow 43 \leftarrow 5$	$A_{0,1,2}^{43}$	2^3		

6 Attacks on the Compression Functions: SHA-2 and Skein

6.1 Preimage Attacks on the Skein Compression Functions

In this section we provide an attack on the 37-round Skein-512 compression function. Control over the tweak value gives us additional freedom both in chunks and the biclique construction.

The attack parameters are listed in Table 4. We build a biclique in rounds 16-23, and apply the attack to rounds 2-38. Bicliques are constructed by Algorithm 3 (Section 3). We use two differential trails: based on Δ^M (Δ -trail) for rounds 16-19 (including key addition in round 19) and based on ∇^M (∇ -trail) for rounds 20-23. The Δ - and ∇ - trails are based on the evolution of a single difference in the linearized Skein and have probability 2^{-52} and 2^{-29} respectively.

The biclique is constructed as follows. First, we restrict to rounds 19-20, where the compression function can be split into two independent 256-bit transformations. A simple approach with table lookups gives a solution to restricted trails with amortized cost 1 (more efficient methods certainly exist). Then we extend this solution to an 8-round biclique by the bits of K^5 . We use K^5 in the messagemodification-like process and adjust the sufficient conditions in rounds 16-23. We have 221 degrees of freedom for that (computed experimentally). As many as 96 bits of freedom do not affect the biclique at all and are used to reduce the amortized cost to only a single round.

Table 4. Parameters of the preimage attack on the 37-round Skein-512 compression function

Biclique					
Rounds	Dimension	Δ^M bits	∇^M bits	Complexity	Freedom used
16-23	1	$K[0]$	$K[4]_{63}$	2^{256}	162
Chunks		Matching			
Forward	Backward	Partial matching	Matching bit	Matching pairs	Complexity
8-15	24-31	$32 \rightarrow 38 = 2 \leftarrow 7$	I_{25}^3	2^2	$2^{1.1}$

In the matching part we recompute 29 rounds per biclique. However, a single key bit flip affects only half of rounds 12-15 and 24-27, and also we need to compute only a half of rounds 2-5 and 35-38. In total, we recompute 42 rounds, or $2^{1.2}$ calls of the compression function per structure, and get 2 candidates matching on one bit. The full preimage is found with complexity $2^{511.2}$.

6.2 Preimage Attacks on the SHA-2 Compression Functions

In this section we provide short description of attacks on the SHA-2 compression functions. The number of rounds we obtain for the compression function setting is in both cases comparable to [18], the latter however does not allow extension to the hash function nor does it violate any “traditional” security requirement. The preimage attack on the compression function is relevant if it is faster than 2^n , though not all these attacks are convertible to the hash function attacks. As a result, we can apply the splice-and-cut attack with the minimum gain to squeeze out the maximum number of rounds. This implies that we consider bicliques of dimension 1. In differential terms, we consider single bit differences Δ_1^M and ∇_1^M . As a result, we get sparse trails with few conditions, and can extend them to more rounds.

- Build 11-round biclique out of a 11-round ∇ -trail in rounds 17-27 (SHA-256) and 21-31 (SHA-512). The trail is a variant of the trail in Table 7 that starts with one-bit difference.
- Construct message words in the biclique as follows. In SHA-256 fix all the message words to constants, then apply the difference Δ_1^M to W^{17} , and assume the linear evolution of Δ_1^M when calculating ΔW^{17+i} from W^2, \dots, W^{17} . Assume also the linear evolution of ∇M when calculating ∇W^{27-i} from W^{28}, \dots, W^{42} . Analogously for SHA-512.
- Build the biclique using internal message words as freedom, then spend the remaining 5 message words to ensure the Δ and ∇ -trails in the message schedule. As a result, we get the longest possible chunks (2-16 and 28-42 in SHA-256).

Therefore, we gain 5 more rounds in the biclique, and two more rounds in the forward chunk. This results in a 52-round attack on the SHA-256 compression function, and a 57-round attack on the SHA-512 compression function.

7 Discussion and Conclusions

We introduced a new concept of bicliques for meet-in-the-middle attacks. We presented several applications of biclique cryptanalysis, including the best preimage attacks so far on SHA-256, SHA-512, and the SHA-3 finalist Skein. In line with most cryptanalytic work, we focused on obtaining results on as many rounds as possible. Though all the functions in this paper are ARX-based, our technique can also be applied to other narrow-pipe designs.

Overall, the differential view gives a cryptanalyst much more freedom and flexibility compared to previous attacks. We can outline the following benefits of applying the biclique concept:

- Use of differential trails in a biclique with a small number of sufficient conditions;
- Deterministic algorithms to build a biclique, which can be adapted for a particular primitive;
- Use of various tools from differential cryptanalysis like trail backtracking [5], message modification and neutral bits [6,15,21,28], condition propagation [9], and rebound techniques [20].

Status of SHA-2 and Skein-512. For SHA-256, SHA-512, and Skein-512, we considered both the hash function and the compression function setting. In all settings we obtained cryptanalytic results on more rounds than any other known method. Based on these results we conclude that Skein-512 is more resistant against splice-and-cut cryptanalysis than SHA-512.

Other Applications of Biclique Cryptanalysis. Soon after the initial circulation of this work, the idea of biclique cryptanalysis found other applications. Bicliques have large potential in attacks on block ciphers, as has been demonstrated by recent attacks on the full versions of popular block ciphers. Among them we mention key recovery faster than brute force for AES-128, AES-192, and AES-256 by Bogdanov et al. [8]. Cryptanalysis of AES employed algorithms for biclique construction which are partly covered in Section 3. In this context we also mention new and improved results on Kasumi by Jia et al. [14] and IDEA by Biham et al. [7] as well as more results announced both publicly [12,19,30] and privately.

Acknowledgements. Part of this work was done while Christian Rechberger was with KU Leuven and visiting MSR Redmond, and while Alexandra Savelieva was visiting MSR Redmond. The authors would like to thank Eik List and anonymous reviewers for useful comments on earlier versions of the paper.

References

1. Aoki, K., Guo, J., Matusiewicz, K., Sasaki, Y., Wang, L.: Preimages for Step-Reduced SHA-2. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 578–597. Springer, Heidelberg (2009)

2. Aoki, K., Sasaki, Y.: Preimage Attacks on One-Block MD4, 63-Step MD5 and More. In: Avanzi, R.M., Keliher, L., Sica, F. (eds.) SAC 2008. LNCS, vol. 5381, pp. 103–119. Springer, Heidelberg (2009)
3. Aoki, K., Sasaki, Y.: Meet-in-the-Middle Preimage Attacks Against Reduced SHA-0 and SHA-1. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 70–89. Springer, Heidelberg (2009)
4. Aumasson, J.-P., Çalik, Ç., Meier, W., Özen, O., Phan, R.C.-W., Varıcı, K.: Improved Cryptanalysis of Skein. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 542–559. Springer, Heidelberg (2009)
5. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: RadioGatun, a belt-and-mill hash function. In: NIST Cryptographic Hash Workshop (2006), <http://radiogatun.noekeon.org/>
6. Biham, E., Chen, R.: Near-Collisions of SHA-0. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 290–305. Springer, Heidelberg (2004)
7. Biham, E., Dunkelman, O., Keller, N., Shamir, A.: New Data-Efficient Attacks on Reduced-Round IDEA. Cryptology ePrint Archive, Report 2011/417 (2011), <http://eprint.iacr.org/>
8. Bogdanov, A., Khovratovich, D., Rechberger, C.: Biclique cryptanalysis of the full AES. In: Lee, D.H. (ed.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 344–371. Springer, Heidelberg (2011), <http://eprint.iacr.org/2011/449>
9. De Cannière, C., Rechberger, C.: Finding SHA-1 Characteristics: General Results and Applications. In: Lai, X., Chen, K. (eds.) ASIACRYPT 2006. LNCS, vol. 4284, pp. 1–20. Springer, Heidelberg (2006)
10. Ferguson, N., Lucks, S., Schneier, B., Whiting, D., Bellare, M., Kohno, T., Callas, J., Walker, J.: The Skein hash function family, version 1.3 (October 1, 2010)
11. Guo, J., Ling, S., Rechberger, C., Wang, H.: Advanced Meet-in-the-Middle Preimage Attacks: First Results on Full Tiger, and Improved Results on MD4 and SHA-2. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 56–75. Springer, Heidelberg (2010)
12. Hong, D.: Biclique attack on the full HIGHT. To appear in ICISC 2011 (2011)
13. Indestege, S., Mendel, F., Preneel, B., Rechberger, C.: Collisions and Other Non-random Properties for Step-Reduced SHA-256. In: Avanzi, R.M., Keliher, L., Sica, F. (eds.) SAC 2008. LNCS, vol. 5381, pp. 276–293. Springer, Heidelberg (2009)
14. Jia, K., Yu, H., Wang, X.: A meet-in-the-middle attack on the full KASUMI. Cryptology ePrint Archive, Report 2011/466 (2011), <http://eprint.iacr.org/>
15. Joux, A., Peyrin, T.: Hash Functions and the (Amplified) Boomerang Attack. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 244–263. Springer, Heidelberg (2007)
16. Kelsey, J., Schneier, B.: Second Preimages on n -Bit Hash Functions for Much Less than 2^n Work. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 474–490. Springer, Heidelberg (2005)
17. Khovratovich, D., Nikolić, I., Rechberger, C.: Rotational Rebound Attacks on Reduced Skein. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 1–19. Springer, Heidelberg (2010)
18. Lamberger, M., Mendel, F.: Higher-order differential attack on reduced SHA-256 (2011), <http://eprint.iacr.org/2011/037.pdf>
19. Mala, H.: Biclique cryptanalysis of the block cipher SQUARE. Cryptology ePrint Archive, Report 2011/500 (2011), <http://eprint.iacr.org/>

20. Mendel, F., Rechberger, C., Schläffer, M., Thomsen, S.S.: The Rebound Attack: Cryptanalysis of Reduced Whirlpool and Grøstl. In: Dunkelman, O. (ed.) FSE 2009. LNCS, vol. 5665, pp. 260–276. Springer, Heidelberg (2009)
21. Naito, Y., Sasaki, Y., Shimoyama, T., Yajima, J., Kunihiko, N., Ohta, K.: Improved Collision Search for SHA-0. In: Lai, X., Chen, K. (eds.) ASIACRYPT 2006. LNCS, vol. 4284, pp. 21–36. Springer, Heidelberg (2006)
22. Neyman, J., Pearson, E.S.: The testing of statistical hypotheses in relation to probabilities a priori. In: Proc. Camb. Phil. Soc. (1933)
23. NIST. FIPS-180-2: Secure Hash Standard (August 2002), <http://www.itl.nist.gov/fipspubs/>
24. Sanadhya, S.K., Sarkar, P.: New Collision Attacks against Up to 24-Step SHA-2. In: Chowdhury, D.R., Rijmen, V., Das, A. (eds.) INDOCRYPT 2008. LNCS, vol. 5365, pp. 91–103. Springer, Heidelberg (2008)
25. Sasaki, Y., Aoki, K.: Preimage Attacks on Step-Reduced MD5. In: Mu, Y., Susilo, W., Seberry, J. (eds.) ACISP 2008. LNCS, vol. 5107, pp. 282–296. Springer, Heidelberg (2008)
26. Sasaki, Y., Aoki, K.: Finding Preimages in Full MD5 Faster Than Exhaustive Search. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 134–152. Springer, Heidelberg (2009)
27. Su, B., Wu, W., Wu, S., Dong, L.: Near-Collisions on the Reduced-Round Compression Functions of Skein and BLAKE. Cryptology ePrint Archive, Report 2010/355 (2010), <http://eprint.iacr.org/>
28. Wang, X., Yin, Y.L., Yu, H.: Finding Collisions in the Full SHA-1. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 17–36. Springer, Heidelberg (2005)
29. Yu, H., Chen, J., Jia, K., Wang, X.: Near-Collision Attack on the Step-Reduced Compression Function of Skein-256. Cryptology ePrint Archive, Report 2011/148 (2011), <http://eprint.iacr.org/>
30. Chen, S.Z., Xu, T.M.: Biclique Attack of the Full ARIA-256. Cryptology ePrint Archive, Report 2012/011 (2012), <http://eprint.iacr.org/>

A Skein Specification and Details of Differential Trail Design

Skein-512 is based on the block cipher Threefish-512 — a 512-bit block cipher with a 512-bit key parametrized by a 128-bit tweak. Both the internal state I and the key K consist of eight 64-bit words, and the tweak T is two 64-bit words. The compression function $F(CV, T, M)$ of Skein is defined as:

$$F(CV, T, M) = E_{CV, T}(M) \oplus M,$$

where $E_{K, T}(P)$ is the Threefish cipher, CV is the previous chaining value, T is the tweak, and M is the message block. The tweak value is a function of parameters of message block M .

Threefish-512 transforms the plaintext P in 72 rounds as follows:

$$P \rightarrow \text{Add } K^0 \rightarrow 4 \text{ rounds} \rightarrow \text{Add } K^1 \rightarrow \dots \rightarrow \text{Add } K^{18} \rightarrow C.$$

The subkey $K^s = (K_0^s, K_1^s, \dots, K_7^s)$ is produced out of the key $K = (K[0], K[1], \dots, K[7])$ as follows:

$$K_j^s = K[(s + j) \bmod 9], \quad 0 \leq j \leq 4; \quad K_5^s = K[(s + 5) \bmod 9] + T[s \bmod 3];$$

$$K_6^s = K[(s + 6) \bmod 9] + T[(s + 1) \bmod 3]; \quad K_7^s = K[(s + 7) \bmod 9] + s,$$

where the additions are all modulo 2^{64} , s is a round counter, $T[0]$ and $T[1]$ are tweak words, $T[2] = T[0] + T[1]$, and $K[8] = C_{240} \oplus \bigoplus_{j=0}^7 K[j]$ with constant C_{240} optimized against rotation attacks.

One round transforms the internal state as follows. Eight words I^0, I^1, \dots, I^7 are grouped into pairs and each pair is processed by a simple 128-bit function MIX. Then all the words are permuted by the operation PERM. Details of these operations are irrelevant for the high-level description and can be found in [10]. We use the following notation for the internal states in round r :

$$S^{r-A} \xrightarrow{\text{MIX}} S^{r-M} \xrightarrow{\text{PERM}} S^{r-P}$$

Local collision in Skein-512. If an attacker controls both the IV and the tweak he is able to introduce difference in these inputs so that one of subkeys has zero difference. As a result, he gets a differential which has no difference in internal state for 8 rounds. The lowest weight of input and output differences is achieved with the combination $\Delta K[6] = \Delta K[7] = \Delta T[1] = \delta$, which gives difference $(0, 0, \dots, 0, \delta)$ in the subkey K^0 and $(\delta, 0, 0, \dots, 0)$ in K^8 , and zero difference in the subkey K^4 . The local collisions for further rounds are constructed analogously. We use the following differences in the compression function attack to make a local collision in rounds 8-15 and 24-31:

$$\Delta K[0] = \Delta T[0] = \Delta T[1] = 1 \lll 63; \quad \Delta K[3] = \Delta K[4] = \Delta T[1] = 1 \lll 63.$$

B Specification of the SHA-2 Family of Hash Functions

The SHA-2 hash functions are based on a compression function that updates the state of eight 32-bit state variables A, \dots, H according to the values of 16 32-bit words M_0, \dots, M_{15} of the message. SHA-384 and SHA-512 operate on 64-bit words. For SHA-224 and SHA-256, the compression function consists of 64 rounds, and for SHA-384 and SHA-512 — of 80 rounds. The full state in round r is denoted by S^r .

The i -th step uses the i -th word W^i of the expanded message. The message expansion works as follows. An input message is split into 512-bit or 1024-bit message blocks (after padding). The message expansion takes as input a vector M with 16 words and outputs a vector W with n words. The words W^i of the expanded vector are generated from the initial message M according to the following equations (n is the number of steps of the compression function):

$$W^i = \begin{cases} M^i & \text{for } 0 \leq i < 15 \\ \sigma_1(W^{i-2}) + W^{i-7} + \sigma_0(W^{i-15}) + W^{i-16} & \text{for } 15 \leq i < n \end{cases}. \quad (11)$$

Table 6. Details of SHA-2 hash family internal operation

Function	SHA-2 Family	
	SHA-224 and SHA-256	SHA-384 and SHA-512
Ch(x, y, z)	$x \wedge y \oplus \bar{x} \wedge z$	
Maj(x, y, z)	$x \wedge y \oplus x \wedge z \oplus y \wedge z$	
$\Sigma_0(x)$	$(x \ggg 2) \oplus (x \ggg 13) \oplus (x \ggg 22)$	$(x \ggg 28) \oplus (x \ggg 34) \oplus (x \ggg 39)$
$\Sigma_1(x)$	$(x \ggg 6) \oplus (x \ggg 11) \oplus (x \ggg 25)$	$(x \ggg 14) \oplus (x \ggg 18) \oplus (x \ggg 41)$
$\sigma_0(x)$	$(x \ggg 7) \oplus (x \ggg 18) \oplus (x \gg 3)$	$(x \ggg 1) \oplus (x \ggg 8) \oplus (x \gg 7)$
$\sigma_1(x)$	$(x \ggg 17) \oplus (x \ggg 19) \oplus (x \gg 10)$	$(x \ggg 19) \oplus (x \ggg 61) \oplus (x \gg 6)$

1. Fix a group of 6-round differential trails $\nabla_i \xrightarrow{\nabla_i^M} 0$ (the one based on 3-bit difference is listed in Table 7). Derive a set of sufficient conditions on the internal states (Table 8).
2. Fix the message compensation equations with constants c_1, c_2, \dots, c_9 (Section C.2).
3. Fix an arbitrary Q_0 and modify it so that most of conditions in the computation $Q_0 \rightarrow P_0$ are fulfilled. Derive Q_i out of Q_0 by applying ∇_i .
4. Fix a group of 2-round trails (the one based on 3-bit difference is given in Table 7) ($\Delta W^{17} \rightarrow \Delta S^{19}$) as a Δ -trail (Equation (6)) in rounds 17-19.
5. Choose $W^{17}, W^{18}, \dots, W^{22}$ and constants c_8, c_9 so that the conditions in the computations $Q_0 \rightarrow P_j, j = 0, \dots, 7$ are fulfilled. Produce all P_j .

Finally, we produce Q_0, \dots, Q_7 and P_0, \dots, P_7 that conform to the biclique equations.

The complexity of building a single biclique is estimated as 2^{32} . 7 message words left undefined in the message compensation equations give us enough freedom to reuse a single biclique up to 2^{256} times. The complexity to recalculate the chunks is upper bounded by 2^2 calls of the compression function. The total amortized complexity of running a single biclique and producing 2^2 matches on 4 bits is 2^3 calls of the compression function. Since we need 2^{252} matches, the complexity of the pseudo-preimage search is 2^{253} . A full preimage can be found with complexity approximately $2^{1+(253+256)/2} \approx 2^{255.5}$ by restarting the attack procedure $2^{\frac{256-253}{2}} = 2^{1.5}$ times. Memory requirements are $\approx 2^{1.5} \times 24$ words.

C.2 Message Compensation

Since any consecutive 16 message words in SHA-2 bijectively determine the rest of the message block at an iteration of compression function, we need to place the initial structure within a 16-round block and define such restrictions on message dependencies that maximize the length of chunks.

We discovered that with W^{17} and W^{22} selected as the words with for a 6-step initial structure, it is possible to expand 16-round message block $\{W^{12}, \dots, W^{27}\}$ by 10 steps backwards and 9 steps forwards, so that $\{W^2, \dots, W^{16}\}$ are calculated independently of W^{17} , and $\{W^{23}, \dots, W^{36}\}$ are calculated independently

of W^{22} . Below we define the message compensation conditions that make such chunk separation possible (neutral bit words are outlined in frames):

$$\begin{aligned}
 -\sigma_1(W^{25}) + W^{27} &= c_1; & -W^{19} - \sigma_1(W^{24}) + W^{26} &= c_2 \\
 -\sigma_1(W^{23}) + W^{25} &= c_3 & -\boxed{W^{17}} + W^{24} &= c_4 \\
 -\sigma_1(W^{21}) + W^{23} &= c_5; & -\sigma_1(W^{19}) + W^{21} &= c_6 \\
 -\sigma_1(\boxed{W^{17}}) + W^{19} &= c_7; & W^{12} + \sigma_0(W^{13}) &= c_8; \\
 W^{13} + \boxed{W^{22}} &= c_9
 \end{aligned} \tag{12}$$

$W^{14}, \dots, W^{16}, W^{18}$, and W^{20} can be chosen independently of both W^{17} and W^{22} , so we can assign W^{14} and W^{15} with 64-bit length of the message to satisfy padding rules (additionally, 1 bit of W^{13} needs to be fixed). W^{18} and W^{20} are additional freedom for constructing the biclique. We use bits 25, 26, 27 as neutral in W_{17} . To prevent this difference from interleaving with the backward trail difference in round 19, we restrict the behavior of the forward trail as specified in Table 7 (aggregated conditions are given in Table 8).

C.3 Trails

The basic differential trail for the biclique is a 6-round trail in the backward direction ($\Delta_Q \leftarrow \nabla M$) that starts with the difference in bits 22, 23, and/or 31 in W_{22} . The trail is briefly depicted in Table 7 with references to the sufficient conditions (which work out for all the 7 possible differences) in Table 8.

Table 7. Details for biclique in SHA-256. Differential ∇ - and Δ - trails (active bits).

$$A' = \{6, 11, 12, 16, 17, 20, 23, 24, 29, 30\}, \Phi = \Sigma_1\{25, 26, 27\} = \{0, 1, 2, 14, 15, 16, 19, 20, 21\},$$

* refers to an arbitrary difference.

Trail	R-nd	A	B	C	D	E	F	G	H	W	Cond-s
∇	17	-	-	22,23,31	-	-	A'	-	*	-	1
∇	18	-	-	-	22,23,31	-	-	A'	-	-	3,4
∇	19	-	-	-	-	22,23,31	-	-	A'	-	7-11
∇	20	-	-	-	-	-	22,23,31	-	-	-	12
∇	21	-	-	-	-	-	-	22,23,31	-	-	13
∇	22	-	-	-	-	-	-	-	22,23,31	-	
∇	23	-	-	-	-	-	-	-	-	22,23,31	
Δ	18	*	-	-	-	25,26,27	-	-	-	-	2
Δ	19	*	*	-	-	Φ	25,26,27	-	-	-	5,6

With three neutral bits we construct a biclique with 8 starting points for chunks in each direction. First, we choose the initial state A_{17}, \dots, H_{17} so that the conditions 1 and 5 are fulfilled. Then we proceed with a standard trail backtracking procedure modifying the starting state if needed. Next, in round 18 we check whether the value of E stops carries in the forward trail. If not, we

Table 8. Sufficient conditions for the Δ - and ∇ -trails in SHA-256

F – how the conditions are fulfilled (IC – initial configuration, SM – state modification).
 C – total number of independent conditions; D_W – conditions fulfilled by message words.
 A^i – i -th bit of A ; $A = \Sigma_1\{22, 23, 31\} = \{6, 11, 12, 16, 17, 20, 25, 29, 30\}$

Round	Conditions	Purpose	F	C	D_W
17	1: $A^{22,23,31} = B^{22,23,31}$	Absorption (MAJ)	IC	3	0
	2: $(W^{\oplus} E_{18})^{25,26,27} = 0$	Stop forw. carry	SM	6	0
18	3: $E^{A'} = 1$,	Absorption (IFF)	SM	9	0
	4: $(D \oplus E_{19})^{22,23,31} = 0$	Stop carry	SM	3	0
19	5: $F^{25,26,27} = G^{25,26,27}$,	Absorption (IFF)	IC	9	0
	6: $(S1 \oplus E_{19})^{\Phi} = 0$	Stop forw. carry	SM	2	0
	7: $F^{22,31} = G^{22,31}$	Absorption (IFF)	SM	2	0
	8: $F^{23} \neq G^{23}$	Pass (IFF)	SM	1	0
	9: $CH^{25} \neq S1^{25}$	Force carry (H)	SM	1	0
	10: $(S1 \oplus H)^A = 1$	Stop carry (H)	SM	9	0
	11: $(CH \oplus H)^{24} = 0$	Force carry (H)	SM	1	0
20	11': $(CH \oplus H)^{23} = 0$	Force carry (H)	SM	1	0
	12: $E^{22,23,31} = 0$	Absorption (IFF)	W^{19}	21	21
21	13: $E^{22,23,31} = 1$	Absorption (IFF)	W^{20}	21	21

change the value of D in the starting state accordingly. Then we sequentially modify the initial state to fulfill the conditions 2-11.

The last two conditions are affected by the message words W_{19} and W_{20} . We need to fulfill three bit conditions for every W_{17} , used in the attack. Therefore, we spend $3 \cdot 8 \cdot 2 = 48$ degrees of freedom in message words $W_{17}, W_{18}, W_{19}, W_{20}, W_{21}$. Note that there is a difference in W_{19} determined by the difference in W_{17} due to the message compensation. We have fixed the constants c_6 and c_7 from Eq. 12 while defining W_{19} and W_{21} . In total, we construct the biclique in about 2^{32} time required to find proper W_{19} and W_{20} .

Amount of freedom used. In total, we have 512 degrees of freedom in the message and 256 degrees of freedom in the state. The biclique is determined by the state in round 17 and message words W_{17} – W_{21} . The choice of W_{19} and W_{21} is equivalent to the choice of constants c_6, c_7 in Eq. 12. We spend $256 + 5 \cdot 32 = 416$ degrees of freedom for the biclique fulfilling as few as $47 + 42$ (Table 8) conditions. After the biclique is fixed, there are $768 - 416 = 352$ degrees of freedom left. We spend $32 + 32 + 2 = 66$ for the padding, leaving 286 degrees of freedom. Therefore, one biclique is enough for the full attack.

D Details on the 50-Round SHA-512 Attack

D.1 Biclique Construction

Steps of the algorithm are similar to those in Section C.1, except the trails are described in Table 9. By applying similar reasoning, we estimated that a full preimage can be found with complexity $\approx 2^{511.5}$ and memory $\approx 2^{1.5} \times 24$ words.

D.2 Message Compensation

The system of compensation equations is defined as follows:

$$\begin{aligned}
 -\sigma_1(W^{29}) + W^{31} &= c_1; & -W^{23} - \sigma_1(W^{28}) + W^{30} &= c_2; & W^{17} + \boxed{W^{26}} &= c_9 \\
 -\sigma_1(W^{27}) + W^{29} &= c_3; & -\boxed{W^{21}} + W^{28} &= c_4; & -\sigma_1(W^{25}) + W^{27} &= c_5; \\
 -\sigma_1(W^{23}) + W^{25} &= c_6; & -\sigma_1(\boxed{W^{21}}) + W^{23} &= c_7; & W^{16} + \sigma_0(W^{17}) &= c_8;
 \end{aligned}$$

We use 1 LSB of W^{13} and 10 LSB of W^{15} for padding. The choice of constants c_8, c_9 and fixed lower 53 bits of W^{26} provide us with sufficient freedom. By choosing c_9 we define lower 53 bits of W^{17} . Having c_8 chosen, we derive 45 lower bits of W^{16} fixed due to σ_0 . We get lower 37 bits of W^{15} , 29 bits of W^{14} and 21 bit of W^{13} fixed. As we need only one LSB of W^{13} and 10 LSB of W^{15} to be fixed, we use lower 33 bits of W^{26} and c_9 , and lower 25 bits of c_8 .

D.3 Trails

The basic differential trail for the biclique is a 6-round trail in the backward direction ($\Delta_Q \leftarrow \nabla M$) that starts with the difference in bits 53, 54, and/or 55 in W^{26} . We also use bits 60, 61, 62 as neutral in W^{21} . To prevent this difference from interleaving with the backward trail difference in round 19, we restrict the behavior of the forward trail. The trails are depicted in Table 9.

Table 9. Details for biclique in SHA-512. Differential ∇ - and Δ -trails (active bits).

$\Lambda = \Sigma_1\{53, 54, 55\} = \{12, 13, 14, 35, 36, 37, 39, 40, 41\}$, $\Phi = \Sigma_1\{60, 61, 62\} = \{17, 20, 21, 42, 43, 44, 46, 47, 48\}$,
 * refers to an arbitrary difference.

Trail	Round	A	B	C	D	E	F	G	H	Cond-s
∇	21	-	-	53,54,55	-	-	Λ	-	*	3
∇	22	-	-	-	53,54,55	-	-	Λ	-	12
∇	23	-	-	-	-	53,54,55	-	-	Λ	12
∇	24	-	-	-	-	-	53,54,55	-	-	24
∇	25	-	-	-	-	-	-	53,54,55	-	24
∇	26	-	-	-	-	-	-	-	53,54,55	
Δ	22	*	-	-	-	60,61,62	-	-	-	3
Δ	23	*	*	-	-	Φ	60,61,62	-	-	18

Converting Meet-In-The-Middle Preimage Attack into Pseudo Collision Attack: Application to SHA-2

Ji Li¹, Takanori Isobe², and Kyoji Shibutani²

¹ Sony China Research Laboratory, China
Ji.Li@sony.com.cn

² Sony Corporation, Japan
{Takanori.Isobe,Kyoji.Shibutani}@jp.sony.com

Abstract. In this paper, we present a new technique to construct a collision attack from a particular preimage attack which is called a partial target preimage attack. Since most of the recent meet-in-the-middle preimage attacks can be regarded as the partial target preimage attack, a collision attack is derived from the meet-in-the-middle preimage attack. By using our technique, pseudo collisions of the 43-step reduced SHA-256 and the 46-step reduced SHA-512 can be obtained with complexities of 2^{126} and $2^{254.5}$, respectively. As far as we know, our results are the best pseudo collision attacks on both SHA-256 and SHA-512 in literature. Moreover, we show that our pseudo collision attacks can be extended to 52 and 57 steps of SHA-256 and SHA-512, respectively, by combined with the recent preimage attacks on SHA-2 by bicliques. Furthermore, since the proposed technique is quite simple, it can be directly applied to other hash functions. We apply our algorithm to several hash functions including Skein and BLAKE, which are the SHA-3 finalists. We present not only the best pseudo collision attacks on SHA-2 family, but also a new insight of relation between a meet-in-the-middle preimage attack and a pseudo collision attack.

Keywords: hash function, narrow-pipe, SHA-2, Skein, BLAKE, meet-in-the-middle attack, preimage attack, pseudo collision attack.

1 Introduction

Cryptographic hash functions play a central role in the modern cryptography. A secure hash function, which produces a fixed length hash value from an arbitrary length message, is required to satisfy at least three security properties: preimage resistance, second preimage resistance and collision resistance.

While there has not been a generic method to convert a collision attack into a preimage attack, it has been known that the preimage attack that can find at least two distinct preimages from the same target can be directly converted into a collision attack. However, the converted collision attack is often not efficient due to that the birthday bound of a collision attack ($2^{n/2}$) is far lower than

the generic bound of the preimage attack (2^n), where n is the bit size of the hash value. Thus, it is left as open question that how to convert an efficient preimage attack into an efficient collision attack. In the case of the reduced SHA-256 regarding the number of attacked rounds, a preimage attack, covering 43 steps [4], is much better than the best known collision attack, with only 27 steps [17]. Moreover, basically, a collision attack and a preimage attack require quite different techniques. In other words, in general, the techniques used for the collision attack do not work well for a preimage attack, and vice versa. In fact, most of the recent collision attacks are based on a differential attack [32,31], in contrast to that most of the recent preimage attacks are based on a meet-in-the-middle (MITM) attack [2]. Though converting the differential collision attack to a (pseudo) preimage attack was discussed in [8], there is no generic way to construct a collision attack from a MITM preimage attack.

In this paper, we give a generic method to convert a particular preimage attack into a collision attack. By using our technique, an efficient collision attack which works faster than a generic collision attack can be constructed from a partial target preimage attack even if the complexity of the preimage attack is more than the birthday bound ($2^{n/2}$). Our method is especially fit for converting a MITM preimage attack into a pseudo collision attack, since most of the recent MITM preimage attacks can be considered as the partial target preimage attack as long as its matching point is located in the end of the compression function. We first apply our algorithm to SHA-256 and SHA-512 and show the best pseudo collision attacks on them in literature. Specifically, pseudo collisions of the 43-step (out of 64-step) reduced SHA-256 and the 46-step (out of 80-step) reduced SHA-512 can be derived faster than a generic attack. Combined with the recent preimage attacks on SHA-2 [14], these attacks are extended to the 52-step and 57-step reduced SHA-256 and SHA-512, respectively. Then we show some other applications of our conversion techniques including a pseudo collision attack on the 37-round reduced Skein-512 and pseudo collision attacks on the 4-round reduced BLAKE-256/512 without the initialization function. While it seems hard to extend our pseudo collision attacks to collision attacks, the proposed conversion technique is a generic, and thus it is expected to be widely used for security evaluations of hash functions.

This paper is organized as follows. Some security notions and a meet-in-the-middle preimage attack are introduced in Section 2. Section 3 introduces our approach for constructing a pseudo collision attack. Then, applications of our technique to SHA-256 and SHA-512 are presented in Section 4. The result on Skein is described in Section 5. Finally, we conclude in Section 6.

2 Preliminaries

In this section, we first give security notions used throughout this paper, then briefly refer a meet-in-the-middle (MITM) preimage attack.

2.1 Security Notions

Let f be a compression function which outputs an n -bit chaining variable h_i from an n -bit input chaining variable h_{i-1} and a k -bit input message m_i , i.e., $h_i = f(h_{i-1}, m_i)$. Similarly, let H be an iterated hash function consisting of f , which produces an n -bit hash value d from an initial value $IV (= h_0)$ and an arbitrary length message M , i.e., $d = H(IV, M) = f(\cdots f(f(IV, m_1), m_2), \cdots, m_t)$, where $pad(M) = (m_1|m_2|\cdots|m_t)$ and pad denotes a padding function. This type of hash function, in which the size of an intermediate chaining variable is the same as that of a hash value, is called a *narrow-pipe* hash function. On the other hand, a hash function having a larger internal state size is called a *wide-pipe* hash function, i.e., the size of a final hash value is smaller than that of a chaining variable. We use the terminology introduced in [15] for a collision attack and a pseudo (or free-start) collision attack on hash functions as follows.

Definition 1 (Collision attack). *Given IV , find (M, M') such that $M \neq M'$ and $H(IV, M) = H(IV, M')$.*

Definition 2 (Free-start or pseudo collision attack). *Find (IV, IV', M, M') such that $H(IV, M) = H(IV', M')$ and $(IV, M) \neq (IV', M')$.*

Additionally, we give several definitions for (pseudo) preimage attacks on hash functions and (pseudo) preimage attacks on compression functions.

Definition 3 (Preimage attack). *Given IV and $d (= H(IV, M))$, find M' such that $H(IV, M') = d$.*

Definition 4 (Pseudo preimage attack). *Given $d (= H(IV, M))$, find (IV', M') such that $H(IV', M') = d$.*

Definition 5 ((t -bit) partial target preimage attack). *Given IV and t -bit partial target of $d (= H(IV, M))$, find M' such that t -bit of $d' (= H(IV, M'))$ is the same as the t -bit of d at the same position, and the other part of d' is randomly obtained.*

Definition 6 (Preimage attack on compression function). *Given h_{i-1} and $h_i (= f(h_{i-1}, m_i))$, find m'_i such that $f(h_{i-1}, m'_i) = h_i$.*

Definition 7 (Pseudo preimage attack on compression function). *Given $h_i (= f(h_{i-1}, m_i))$, find (h'_{i-1}, m'_i) such that $f(h'_{i-1}, m'_i) = h_i$.*

Definition 8 ((t -bit) partial target preimage attack on compression function). *Given h_{i-1} and t -bit partial target of $h_i (= f(h_{i-1}, m_i))$, find m'_i such that t -bit of $h'_i (= f(h_{i-1}, m'_i))$ is the same as the t -bit of h_i at the same position, and the other part of h'_i is randomly obtained.*

Definition 9 ((t -bit) pseudo partial target preimage attack on compression function). *Given t -bit partial target of $h_i (= f(h_{i-1}, m_i))$, find (h'_{i-1}, m'_i) such that t -bit of $h'_i (= f(h'_{i-1}, m'_i))$ is the same as the t -bit of h_i at the same position, and the other part of h'_i is randomly obtained.*




Fig. 1. Meet-in-the-middle preimage attack

2.2 Meet-In-The-Middle Preimage Attack

The basic concept of the MITM preimage attack was introduced in [22,16]. Since then, the MITM preimage attacks have been drastically improved and applied to several hash functions [2,28,27,3,13,4,10]. Also, the techniques for the MITM preimage attacks on hash functions have been extended to the attacks on several block ciphers [7,12].

As shown in Fig. 1,¹ in the MITM preimage attack on a compression function, the compression function f is assumed to be divided into two sub-functions: f_1 (forward process) and f_2 (backward process) so that the w -bit matching point z calculated by f_1 does not depend on m_2 which is some message bits of m , and z calculated by f_2 does not depend on m_1 which is other message bits of m . Such m_1 and m_2 are called neutral bits of f_2 and f_1 , respectively. Then, the MITM preimage attack finds a preimage m' such that $f(x, m') = y$ from a given x and $y (= f(x, m))$ as follows.

- Step 1.** Choose a random m except for m_1 and m_2 .
- Step 2.** For all possible m_1 , calculate w -bit $z_1 (= f_1(x, m_1))$, and add a pair of $(z_1^{(i)}, m_1^{(i)})$ to a list, where $(1 \leq i \leq 2^{|m_1|})$, and $|*|$ denotes the bit size of $*$.
- Step 3.** For all possible m_2 , calculate w -bit $z_2 (= f_2^{-1}(x \oplus y, m_2))$, and add a pair of $(z_2^{(j)}, m_2^{(j)})$ to a list, where $(1 \leq j \leq 2^{|m_2|})$.
- Step 4.** Compare two lists to find pairs satisfying $z_1^{(p)} = z_2^{(q)}$. If such pair is found, then check if the other bits of the matching point derived from $m_1^{(p)}$ and $m_2^{(q)}$ are the same value.
- Step 5.** If the other parts are also the same, then outputs such m including $m_1^{(p)}$ and $m_2^{(q)}$. Otherwise, go back to Step 1 and repeat the computation.

From Steps 2 and 3, we have $2^{|m_1|}$ and $2^{|m_2|}$ values of w -bit z_1 and z_2 , i.e., we have $2^{|m_1|+|m_2|}$ values of $(z_1 \oplus z_2)$. Since the probability of $(z_1 \oplus z_2 = 0)$ is 2^{-w} , we have $2^{|m_1|+|m_2|} \cdot 2^{-w}$ pairs such that $z_1 = z_2$ in Step 4. Thus, by repeating this algorithm about $2^{n-w} \cdot 2^{-(|m_1|+|m_2|)} \cdot 2^w$ times, we expect to obtain a desired preimage. The required computation for the one process from

¹ Here, we show the MITM preimage attack on Davies-Meyer mode as an example. MITM preimage attacks on other modes like Matyas-Meyer-Oseas mode can be performed in a similar way.

Step 1 to 4 is at most $\max(2^{|m_1|}, 2^{|m_2|})$ calls of the compression function. Thus, the total computation to find a preimage of the compression function is about $2^n \cdot 2^{-(|m_1|+|m_2|)} \cdot \max(2^{|m_1|}, 2^{|m_2|})$.²

For a narrow-pipe hash function, by replacing x and y by IV and d , this MITM preimage attack on a compression function can be directly converted into a preimage attack on a hash function. However, for an attack on a hash function, some of the message bits related to the padding bits are required to be controlled by the attacker to set appropriate padding data.

3 Method to Convert Preimage Attack into Collision Attack

In this section, we present how to efficiently convert a particular preimage attack into a pseudo collision attack. First, we introduce a generic technique to construct a pseudo collision attack from a partial target preimage attack. Then, we introduce the MITM preimage attack whose matching point is located at the end of the compression function. We show that such class of the MITM preimage attack is regarded as the partial target preimage attack. Finally, we show that a pseudo collision attack can be efficiently constructed from the MITM preimage attack whose matching point is at the end by showing how to efficiently obtain many partial target preimages.

3.1 Generic Conversion of Partial Target Preimage Attack into Collision Attack

We consider the oracle \mathcal{A} that can find a t -bit partial target preimage with a complexity of 2^s . Also, \mathcal{A} is assumed to return different M' for each call. Obviously, we can construct a collision attack with a complexity of $2^s \cdot 2^{(n-t)/2}$ by iteratively calling \mathcal{A} as follows.

- Set t -bit random data as d'
- Call \mathcal{A} with the parameter IV and d' in $2^{(n-t)/2}$ times

After this procedure, we have $2^{(n-t)/2}$ of $(n-t)$ -bit random data, and thus there exists a colliding data with a high probability. Once the colliding data are found, we have a collision of the hash function since the rest of the hash value d' is fixed. The total complexity is $2^{(n-t)/2} \cdot 2^s$. The memory requirement can be reduced to the memory requirement of finding a partial target preimage by using memory free birthday attack [29,21]. This conversion itself can be applied to not only a narrow-pipe hash function but also a wide-pipe hash function, since the required complexity depends only on the size of the digest. The basic

² The estimated complexity does not depend on the size of the matching point w . However, as discussed in [10], if w is extremely small like $w = 1$, the total complexity is dominated by the recomputations in Step 4 which is ignored in our estimation. Thus, in our evaluation, we assume that w is sufficiently large.




Fig. 2. MITM preimage attack with the matching point in the last step

concept of this attack that fixes t -bit of the target with the complexity of 2^s has been used to find a collision of (new) FORK-256 in [22] and a collision and a second preimage of LUX in [33]. However, the method does not work if the partial target preimage attack is not efficient, i.e., ($s \geq t/2$). In this case, the required complexity in total will be higher than $2^{n/2}$.

3.2 Meet-In-The-Middle Attack with Matching Point in Last Step

We consider a similar model explained in Section 2.2. The difference from the model shown in Fig. 1 is that the matching point is restricted to be in the last step as shown in Fig. 2. In this scenario, the MITM pseudo preimage attack on a compression function finds a preimage m' and a random x' such that $f(x', m') = y$ from a given $y (= f(x, m))$ as follows.

- Step 1.** Choose a random m except for m_1 and m_2 , and a random starting state S .
- Step 2.** For all possible m_1 , calculate w -bit $z_1 (= f_1(S, m_1))$, and add a pair of $(z_1^{(i)}, m_1^{(i)})$ to a list, where $(1 \leq i \leq 2^{|m_1|})$.
- Step 3.** For all possible m_2 , calculate w -bit $z_2 (= f_2^{-1}(S, m_2))$, and add a pair of $(z_2^{(j)}, m_2^{(j)})$ to a list, where $(1 \leq j \leq 2^{|m_2|})$.
- Step 4.** Compare two lists to find pairs satisfying that $z_1^{(p)} \oplus z_2^{(q)}$ equals the t -bit of y . If such pair is found, then check if the XORed other bits of the matching point derived from $m_1^{(p)}$ and $m_2^{(q)}$ is the same as the rest of y .
- Step 5.** If the XORed other bits are also the same as y , then output such m including $m_1^{(p)}$ and $m_2^{(q)}$, and x' calculated from the data of the matching point. Otherwise, go back to Step 1 and repeat the computation.

Note that, this attack basically cannot obtain a preimage from the given x unlike the attack described in Section 2.2, since x' will be randomly derived. Thus, this attack is considered as a pseudo preimage attack on a compression function. However, for a narrow-pipe hash, it has been known that a pseudo preimage attack on a compression function can be converted into a preimage attack on a hash function assuming that the attacker can set valid padding bits [19,10]. The estimated complexity to find a desired pseudo preimage is the same as that presented in Section 2.2, i.e., $2^n \cdot 2^{-(|m_1|+|m_2|)} \cdot \max(2^{|m_1|}, 2^{|m_2|})$.

3.3 Conversion of MITM Preimage Attack into Pseudo Collision Attack

If we can construct the MITM pseudo preimage attack whose matching point is located at the end of the compression function, we can control part of the output variables as explained in the previous subsection. In other words, the MITM pseudo preimage attack described in the previous subsection can be regarded as the pseudo partial target preimage attack on a compression function. For the MITM preimage attack, at least $2^{t/2}$ computations are required to derive a preimage of an t -bit partial target. Thus, the directly converted pseudo collision attack will at least have the complexity of $2^{(n-t)/2+t/2} = 2^{n/2}$, that is not an efficient pseudo collision attack.

In order to overcome this problem, we exploit extra freedom of a neutral word after finding a partial target preimage. For example, in the case of $t = 10$ and $|m_1| = |m_2| = 8 (> t/2)$, we can find $2^6 (= 2^{8+8}/2^{10})$ 10-bit partial target preimages with the complexity of 2^8 . It essentially means that a 10-bit partial target preimage is found with the complexity of $2^2 (= 2^8/2^6) < 2^5 (= 2^{10/2})$. When $t \leq w$, the required complexity to find a partial target preimage from a given t -bit partial target is estimated as

$$2^{t-(|m_1|+|m_2|)} \cdot \max(2^{|m_1|}, 2^{|m_2|}),$$

where recall that w denotes the bit size of the matching point. In particular, $s < t/2$, which is the condition for a successful attack as mentioned in Section 3.1, holds when $\min(|m_1|, |m_2|) > t/2$, where recall that 2^s represents the required complexity to find a t -bit partial target preimage. Therefore, if we can move the matching point of the MITM attack to the end of the compression function and there is enough freedom in neutral words, we can construct an efficient pseudo collision attack on a compression function.

Moreover, for a narrow-pipe hash function, it has been known that a (pseudo) collision attack on a compression function can be directly converted to a (pseudo) collision attack on a hash function by appending another message block illustrated in Fig. 3, which is called multi-block message technique. By using the multi-block message technique, an attacker can append arbitrary messages. Thus, unlike the conversion to a (pseudo) preimage attack on a hash function, for the conversion to a pseudo collision attack on a hash function, there is no restriction on controllability of message bits for a MITM pseudo preimage attack on a compression function. This will relax conditions on the position of the matching point for the MITM pseudo preimage attack on a compression function, and thus may allow us to attack larger number of steps. Note that, for a wide-pipe hash function, even though a (pseudo) collision attack on a compression function can not be directly converted to a (pseudo) collision attack on a hash function by using multi-block message, we still can convert a MITM pseudo preimage attack on a hash function to a pseudo collision attack on a hash function since the conversion of a partial target preimage attack into a collision attack is generic. Furthermore, in our attack, t -bit of the colliding digest can be determined by the attacker unlike the usual collision attack that derives a completely random

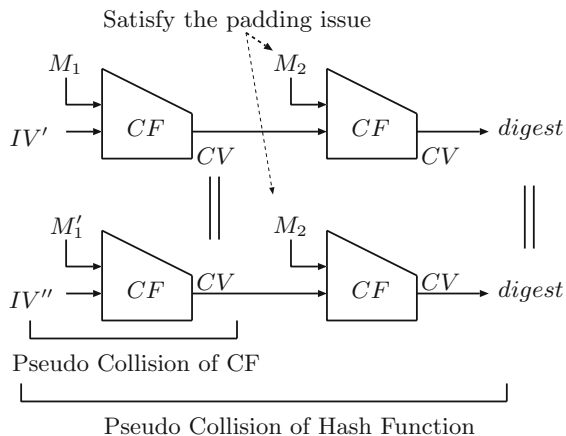


Fig. 3. Multi-block pseudo collision

digest. This is another feature of our approach. On the other hand, the required complexity of our converted (pseudo) collision attack is likely to be high due to a few gains from the MITM procedure, though it is still more efficient than the generic attack. This is considered as one of the limitations of our approach.

4 Pseudo Collision Attacks on SHA-2

In this section, we apply our conversion technique to SHA-2. At first, we briefly describe the algorithm of SHA-2. Then, we review the previous collision attacks on SHA-2. After that, we introduce the known MITM preimage attack on the 43-step SHA-256 presented in [4]. After we modify these results in order to fit our conversion technique, i.e., moving the matching point to the end of the compression function, we show the pseudo collision attack on the 43-step SHA-256. Moreover, we present the pseudo collision attack on the 46-step SHA-512 based on the MITM preimage attack on the 46-step SHA-512 [4]. Furthermore, pseudo collision attacks on the 40-step reduced SHA-224 and SHA-384 are demonstrated as well. Finally, we discuss pseudo collision attacks based on the recent MITM preimage attacks [14], which significantly improve the results of [4] in terms of the number of attacked steps by using *bicliques*. These results on SHA-2 are summarized in Table 1.

4.1 Description of SHA-2

While our target is both SHA-256 and SHA-512, we only explain the structure of SHA-256, since SHA-512 is structurally equivalent to SHA-256 except for the number of steps, the amount of rotations and the word size. The compression function of SHA-256 consists of a message expansion function and a state update function. The message expansion function expands a 512-bit message block into

Table 1. Summary of collision attacks on the reduced SHA-2

algorithm	type of attack	steps	complexity	based attack	paper
SHA-256	collision	24	$2^{28.5}$	-	[11]
	collision	27	(practical)	-	[17]
	semi-free-start-collision*1	24	2^{17}	-	[11]
	semi-free-start-collision*1	32	(practical)	-	[17]
	pseudo-near-collision	31	2^{32}	-	[11]
	pseudo collision	42	2^{123}	[4]	Our (Section 4.7)
	pseudo collision	43	2^{126}	[4]	Our (Section 4.4)
	pseudo collision	45	$2^{126.5}$	[14]	Our (Section 4.9)
SHA-224	pseudo collision	52	$2^{127.5}$	[14]	Our (Section 4.9)
	pseudo collision	40	2^{110}	[4]	Our (Section 4.8)
SHA-512	collision	24	$2^{28.5}$	-	[11]
	pseudo collision	42	2^{244}	[4]	Our (Section 4.7)
	pseudo collision	46	$2^{254.5}$	[4]	Our (Section 4.6)
	pseudo collision	50	$2^{254.5}$	[14]	Our (Section 4.9)
	pseudo collision	57	$2^{255.5}$	[14]	Our (Section 4.9)
SHA-384	pseudo collision	40	2^{183}	[4]	Our (Section 4.8)

*1: semi-free-start-collision attack finds (IV', M, M') such that $H(IV', M) = H(IV', M')$ and $M \neq M'$.

64 32-bit message words (W_0, \dots, W_{63}) as follows:

$$W_i = \begin{cases} M_i & (0 \leq i < 16), \\ \sigma_1(W_{i-2}) + W_{i-7} + \sigma_0(W_{i-15}) + W_{i-16} & (16 \leq i < 64), \end{cases}$$

where the functions $\sigma_0(X)$ and $\sigma_1(X)$ are defined by

$$\begin{aligned} \sigma_0(X) &= (X \ggg 7) \oplus (X \ggg 18) \oplus (X \gg 3), \\ \sigma_1(X) &= (X \ggg 17) \oplus (X \ggg 19) \oplus (X \gg 10). \end{aligned}$$

The state update function updates eight 32-bit chaining variables, A, B, \dots, G, H in 64 steps as follows:

$$\begin{aligned} T_1 &= H_i + \Sigma_1(E_i) + Ch(E_i, F_i, G_i) + K_i + W_i, \\ T_2 &= \Sigma_0(A_i) + Maj(A_i, B_i, C_i), \\ A_{i+1} &= T_1 + T_2, \quad B_{i+1} = A_i, \quad C_{i+1} = B_i, \quad D_{i+1} = C_i, \\ E_{i+1} &= D_i + T_1, \quad F_{i+1} = E_i, \quad G_{i+1} = F_i, \quad H_{i+1} = G_i, \end{aligned}$$

where K_i is the i -th step constant and the functions Ch, Maj, Σ_0 and Σ_1 are given as follows:

$$\begin{aligned} Ch(X, Y, Z) &= XY \oplus \overline{X}Z, \\ Maj(X, Y, Z) &= XY \oplus YZ \oplus XZ, \\ \Sigma_0(X) &= (X \ggg 2) \oplus (X \ggg 13) \oplus (X \ggg 22), \\ \Sigma_1(X) &= (X \ggg 6) \oplus (X \ggg 11) \oplus (X \ggg 25). \end{aligned}$$

After 64 steps, a feed-forward process is executed with initial state variables by using word-wise addition modulo 2^{32} .

4.2 Known Collision Attacks on SHA-2

The first collision attack on reduced SHA-256 was presented in [18] which is a 19-step near collision attack. Since then, the collision attacks on SHA-2 have been improved [20,23,25,24,26,11,17]. The previously published best collision attacks in terms of the number of attacked steps are the 27 steps on SHA-256 [17] and the 24 steps on SHA-512 [11,25]. A non-random property, which is a second-order differential collision, of the 47-step reduced SHA-256 compression function was reported in [6].

4.3 Known MITM Preimage Attack on 43-Step SHA-256 [4]

The MITM preimage attack on the 43-step SHA-256 presented in [4] uses the 33-step two chunks W_j, \dots, W_{j+32} including the 4-step initial structure (IS), the 2-step partial fixing (PF), the 7-step partial matching (PM) and the 1-step indirect partial matching (IPM). In the following, we review the details of these techniques.

33-step Two Chunks with the 4-Step IS. The message words of length 33 is divided into two chunks as $\{W_j, \dots, W_{j+14}, W_{j+18}\}$ and $\{W_{j+15}, W_{j+16}, W_{j+17}, W_{j+19}, \dots, W_{j+32}\}$. Using message compensation technique [4], the first chunk and the second chunk are independent from W_{j+15} and W_{j+18} , respectively. In particular, the following constraints ensure the above message words to be neutral words with respect to each chunk;

$$\begin{aligned} W_{j+17} &= \sigma_1(W_{j+15}), & W_{j+19} &= \sigma_1^2(W_{j+15}), & W_{j+21} &= \sigma_1^3(W_{j+15}), \\ W_{j+22} &= W_{z+5}, & W_{j+23} &= \sigma_1^4(W_{j+15}), & W_{j+24} &= 2\sigma_1(W_{j+15}), \\ W_{j+25} &= \sigma_1^5(W_{j+15}), \end{aligned} \tag{1}$$

where $\sigma_1^2(X)$ means $\sigma_1 \circ \sigma_1(X)$.

These two chunks include the 4-step IS, which essentially exchanges the order of the words W_i and W_{i+3} by exploiting the absorption property of the function Ch . After the swapping, the final output after the step $(i + 3)$ still keeps unchanged. Here, W_{j+18} is moved to the first chunk and W_{j+15}, W_{j+16} and W_{j+17} are moved to the second chunk.

In the forward direction, a state value of $p_{j+33} = A_{j+33} || \dots || H_{j+33}$ can be computed independently of the first chunk. In the backward direction, a state value of $p_j = A_j || \dots || H_j$ can be computed independently of the second chunk. Note that the 33-step two-chunk is valid regardless of the choice of j for $j > 0$.

7-step PM. In the backward computation, A_j can be computed from p_{j+7} without knowing $\{W_j, \dots, W_{j+6}\}$ for any j as used in [13].

2-step PF. PF is a technique to enhance PM by fixing a part of a neutral word. The equation for H_{j-1} is as follows:

$$\begin{cases} H_{j-1} = A_j - \Sigma_0(B_j) - Maj(B_j, C_j, D_j) - \Sigma_1(F_j) \\ \quad - Ch(F_j, G_j, H_j) - K_{j-1} - W_{j-1}, \\ W_{j-1} = W_{j+15} - \sigma_1(W_{j+13}) - W_{j+8} + \sigma_0(W_j). \end{cases}$$

If we fix the lower ℓ bits of W_{j+15} , which is assumed to be a neutral word for the other chunk, the lower ℓ bits of H_{j-1} can be computed without using the value of the higher $(32 - \ell)$ bits of W_{j+15} . Furthermore, the equation for H_{j-2} is expressed as follows:

$$\begin{cases} H_{j-2} = A_{j-1} - \Sigma_0(B_{j-1}) - Maj(B_{j-1}, C_{j-1}, D_{j-1}) - \Sigma_1(F_{j-1}) \\ \quad - Ch(F_{j-1}, G_{j-1}, H_{j-1}) - K_{j-2} - W_{j-2}, \\ W_{j-2} = W_{j+14} - \sigma_1(W_{j+12}) - W_{j+7} + \sigma_0(W_{j-1}). \end{cases}$$

The lower $(\ell - 18)$ bits of H_{j-2} can be computed if we can obtain the lower ℓ bits of $Ch(F_{j-1}, G_{j-1}, H_{j-1})$ and the lower $(\ell - 18)$ bits of $\sigma_0(W_{j-1})$. Note that these values can be computed by using only the lower ℓ bits of W_{j+15} . Thus, when we fix the lower ℓ bits of W_{j+15} , the lower $(\ell - 18)$ bits of H_{j-2} can be computed without knowing the higher $(32 - \ell)$ bits of W_{j+15} . Therefore, by combining the 7-step PM with the 2-step PF, 9 steps can be skipped in the backward computation.

1-step IPM. For the forward computation, A_{j+34} can be expressed as a sum of two independent functions ψ_F, ξ_F of each neutral word as follows;

$$\begin{cases} A_{j+34} = \Sigma_0(A_{j+33}) + Maj(A_{j+33}, B_{j+33}, C_{j+33}) + H_{j+33} + \Sigma_1(A_{j+33}) \\ \quad + Ch(A_{j+33}, B_{j+33}, C_{j+33}) + K_{j+33} + W_{j+33}, \\ W_{j+33} = \sigma_1(W_{j+31}) + W_{j+26} + \sigma_0(W_{j+18}) + W_{j+17}, \end{cases} \\ \Rightarrow A_{j+34} = \psi_F(W_{j+15}) + \xi_F(W_{j+18}).$$

Then, we can compute $\psi_F(W_{j+15})$ and $\xi_F(W_{j+18})$ independently. It is equivalent to move the computation of $\xi_F(W_{j+18})$ to the backward chunk. In this case, $\xi_F(W_{j+18}) = \sigma_0(W_{j+18})$.

Attack Overview. These techniques enable us to construct the 43 ($= 33 + 7 + 2 + 1$)-step attack on SHA-256. Here, we have the freedom of choice of j as long as 36 steps (W_{j-2} to W_{j+34}) is located sequentially.

For the actual attack in [4], j is chosen as $j = 3$, because W_{13}, W_{14} and W_{15} can be freely chosen to satisfy the message padding rule. The matching state is the lower 4 bits of A_{37} . In addition, the number of fixed bits ℓ for PF is chosen as $\ell = 23$. Then, neutral words of W_{18} and W_{21} have 5- and 4-bit freedom degrees, respectively. As a result, a pseudo preimage is found with the complexity of $2^{251.9}$. After that, pseudo preimages are converted into a preimage with the complexity of $2^{254.9}$. See [4] for more details about this attack.

4.4 Pseudo Collision Attack on 43-Step SHA-256

As discussed in Section 3.3, to convert a MITM preimage attack into a pseudo collision attack, the matching point is located into the end of the compression function, i.e., the addition of the feed-forward. As mentioned in section 4.3, the




Fig. 4. 43-step pseudo collision attack on SHA-256

matching point of the 43-step MITM preimage attack is selected at the state after the step 37 ($j = 3$) due to the padding bits.

However, for a (pseudo) collision attack, we do not need to control message words for satisfying the padding rules, since we can generate correct padding by simply adding another message block as discussed in Section 3.3. It means that the last block of a compression function is used only for satisfying the padding condition in the collision attack when pseudo collision can be found before the last compression function as shown in Fig. 3. As a result, for a (pseudo) collision attack, we can move the matching point to the state after the step 43 ($j = 9$) that is the end of the compression function.³

Let a 256-bit output of the compression function be $CV = \{Z_A || \dots || Z_H\}$, where each word is 32 bits. For $j = 9$, W_{24} and W_{27} are neutral words, and the matching point is the lower 4 bits of $A_{43}(= A_0 \oplus Z_A)$.

In order to construct the pseudo collision attack, we give the efficient method to obtain 4-bit partial target preimages by using the MITM technique [4]. Figure 4 shows the overview of the 43-step pseudo collision attack.

Attack Procedure

1. Choose the lower 4 bits of Z_A , which are target values.
2. Randomly choose the value of p_{25} and message W_{25} . Randomly fix the lower 23 bits of W_{24} . Then we can find 2^5 values of W_{24} on average from 9 free bits that correctly construct the 4-step initial structure and store them in the table T_W .
3. Randomly choose message words not related to the initial structure and the neutral words, i.e., W_{19} , W_{20} , W_{21} , W_{22} , W_{23} and W_{29} (called an initial configuration).

³ It is also pointed out in [10] as the matching point can be rotated to the end of the compression function.

4. For all 2^5 possible W_{24} in T_W , compute $W_{26}, W_{28}, W_{30}, W_{31}, W_{32}, W_{33}$ and W_{34} following Eq. (1). Compute forward and find $\psi_F(W_{24})$. Then, store the pairs $(W_{24}, \psi_F(W_{24}))$ in a list L_F .
5. For all 2^4 possible values (the lower 4 bits) of W_{27} , compute backward and find $\xi_F(W_{27})$ and the lower 4 bits of A_0 . Then, store the pairs $(W_{27}, Z_A \oplus A_0 - \sigma_0(W_{27}))$ in a list L_B .
6. If a match is found, i.e., $\psi_F(W_{24}) = Z_A \oplus A_0 - \sigma_0(W_{27})$, then compute two group of states $A_{43}, B_{43}, \dots, H_{43}$ and A_0, B_0, \dots, H_0 with corresponding W_{24} and W_{27} , respectively. Then obtain $2^5 (= 2^9/2^4)$ CV whose 4-bit are fixed, i.e., the lower 4 bits of Z_A , and store these in a List L_1 .
7. Repeat (3)-(6) 2^{121} times with different values of the initial configuration.

After the above procedures, we obtain $2^{126} (= 2^5 \times 2^{121})$ pairs whose 4 bits are fixed.⁴ Thus, there exists a colliding pair with a high probability, because of the equation of $(2^{126} = 2^{(256-4)/2})$.

Evaluation. We assume that the complexity for the 1-step function and the 1-step message expansion is 1/43 compression function operation of the 43-step SHA-256. As estimated in [10], the complexity of Step 2 in the presented attack is 2^9 , and that of Steps 3-6 is $2^{4.878}$, which is the complexity for finding 2^5 4-bit partial target preimages. Thus, whole complexity of the pseudo collision attack on the 43-step SHA-256 is estimated as $2^{126} \approx 2^9 + (2^{121} \times 2^{4.878})$.

4.5 Known MITM Preimage Attack on 46-Step SHA-512 [4]

The MITM preimage attack on the 46-step SHA-512 presented in [4] uses the 31-step two chunks W_j, \dots, W_{j+30} including the 2-step IS, the 8-step PF for W_{j-1}, \dots, W_{j-6} and W_{j+31}, W_{j+32} and the 7-step PM. In this attack, we can choose j as long as 39 step (W_{j-6} to W_{j+32}) are located sequentially. For the actual attack in [10], j is chosen as $j = 6$ to satisfy the padding rule. Then, the neutral words W_{21} and W_{22} have 4 and 3-bit freedom degrees, respectively, and the bit size of the matching point is 3. Thus, a preimage of the 46-step SHA-512 is found with the complexity of $2^{511.5}$. See [4] for more details about this attack.

4.6 Pseudo Collision Attack on 46-Step SHA-512

Similarly to the attack on the reduced SHA-256, we can move the matching point to the end of the compression function, because the padding issue can be avoided by using multi-block message technique in the pseudo collision attack. In the case of SHA-512, since the bit size of the matching point is 3, we utilize the 3-bit partial target preimages for the attack. Then, the complexity of the attack is estimated as $2^{254.5} = (2^{(512-3)/2})$.

⁴ It is noted that we need a slightly more than 2^{121} times repeated experiments to get 2^{126} pairs that will achieve a probability higher than 2^{-1} . However the difference is so small that we ignore it here.

4.7 Pseudo Collision Attacks on 42-step SHA-256 and 42-step SHA-512

We consider pseudo collision attacks on smaller number of rounds of SHA-2 in order to save the time complexity. For the 42-step reduced SHA-256, we can use 10 bits of freedom in both directions to find a 10-bit partial target preimage as discussed in Section 5.4 of [4]. This implies that a 10-bit partial target preimage is obtained with the complexity $1 (< 2^5)$. Thus, a pseudo collision is found with the complexity of $2^{123} (= 2^{(256-10)/2} \times 2^{10}/2^{10})$. Similarly to this, for the 42-step reduced SHA-512, we can use 24 bits of freedom in both directions to find a 24-bit partial target preimage as discussed in Section 6.5 of [4]. Therefore, a pseudo collision of the 42-step reduced SHA-512 is found with the complexity of $2^{244} (= 2^{(512-24)/2} \times 2^{24}/2^{24})$.

4.8 Pseudo Collision Attacks on Reduced SHA-224 and SHA-384

The pseudo collision attack on the 43-step SHA-256 described in Section 4.4 is applicable to the 43-step SHA-224 in the similar manner. However, we can not use the multi-block message technique straightforwardly, because the pseudo collision attack on SHA-224 needs to be done in the last compression function whose output Z_H is disregarded. Thus, due to the padding issue, we can mount only pseudo collision attack on a compression function of 43-step, not a hash function. The estimated complexity is 2^{110} for this attack.

However, the smaller number of rounds of SHA-224 hash function can be attacked by using another MITM attack. The 40-step SHA-224 hash function can be attacked by using the same two chunks for the 43-step preimage attack on SHA-256 in [4], i.e., the case of $j = 3$. The 7-step partial matching for backward computation are replaced by the 4-step one. Then the message words W_{13} , W_{14} and W_{15} are left as free message words to satisfy the padding rule. Instead of the lower 4 bits of Z_A , we use the lower 4 bits of Z_D as the target value. Here, we need additional one step: when finding matches at the lower 4 bits of A_{37} , we compute forward from the matching point to the end of the compression function (40-th step) by using these values that are computed forward from the starting point. Since $A_{37} = D_{40} = D_0 \oplus Z_D$ for the 40-step SHA-224, the lower 4 bits of Z_D will keep unaffected by the additional step. Thus, we can still get a partial target preimage. It can be converted into a pseudo collision attack on a hash function, because we can set W_{13} , W_{14} and W_{15} to follow the padding rule.

The detail of the attack procedure is as follows.

1. Choose the lower 4 bits of Z_D , which are target values.
2. Randomly choose the value of p_{19} and message W_{19} . Randomly fix the lower 23 bits of W_{18} . Then we can find 2^5 values of W_{18} on average from 9 free bits that correctly construct the 4-step initial structure and store them in the table T_W .
3. Randomly choose message words not related to the initial structure and the neutral words, i.e., W_{13} , W_{14} , W_{15} , W_{16} , W_{17} , W_{23} (called an initial configuration [4]).

4. For all 2^5 possible W_{18} in T_W , compute $W_{20}, W_{22}, W_{24}, W_{25}, W_{26}, W_{27}, W_{28}$ following Eq. (1). Compute forward and find $\psi_F(W_{18})$. Store the pairs $(W_{18}, \psi_F(W_{18}))$ in a list L_F .
5. For all 2^4 possible values (the lower 4 bits) of W_{21} , compute backward and find $\xi_F(W_{21})$ and the lower 4 bits of A_{37} ($= D_{40} = Z_D \oplus D_0$). Store the pairs $(W_{21}, Z_D \oplus D_0 - \sigma_0(W_{27}))$ in a list L_B .
6. If a match is found, i.e., $\psi_F(W_{24}) = Z_D \oplus D_0 - \sigma_0(W_{27})$, then compute forward to get the states $A_{40}, B_{40}, \dots, H_{40}$ with corresponding W_{24} and W_{27} , respectively. D_{40} will keep unaffected in this step. Then obtain 2^5 ($= 2^9/2^4$) CV whose 4 bits are fixed, i.e., the lower 4 bits of Z_D , and store these in a List.
7. Repeat (3)-(6) 2^{105} times with different values of the initial configuration.

The complexity of the attack is estimated as 2^{110} .

Similarly, the pseudo collision attack on the 46-step SHA-512 hash function described in 4.6 can also be applied to the 46-step SHA-384 compression function with the complexity of $2^{190.5} = (2^{(384-3)/2})$. For a pseudo collision attack on the reduced SHA-384 hash function, we use the 43-step preimage attack on SHA-384 [4]. Combining the result in [4] with our conversion technique, a pseudo collision attack on the 40-step SHA-384 hash function can be constructed. The matching bit is 18 when chosen parameter of partial matching as $\ell = 27$. The complexity of the pseudo collision attack on the 40-step SHA-384 is estimated as $2^{(384-18)/2} = 2^{183}$. These 40-step pseudo collision attacks give examples that the matching point is not at but near the end of compression function. That is compatible to solve padding problem.

4.9 Application to Other Results of SHA-2

Recently, the MITM preimage attacks on the reduced SHA-2 are improved by using ‘‘biclques’’ technique which is considered as generalized initial structure [14]. This technique enables us to construct longer initial structures than those of the attacks [4]. In the following, let us consider pseudo collision attacks based on [14].

For SHA-256, the 36-step two independent chunks including the 6-step IS based on bicliques are constructed. Combining the 2-step PM with the 7-step PM and the 1-step IPM, the MITM preimage attack on the 45-step SHA-2 is derived. In this attack, both neutral words have 3-bit freedom degrees, and the matching point is 4-bit. Since our conversion technique does not need to consider the padding issue, the matching point can be moved to the end of the compression function similar to the 43-step attack. Then, we can convert it into the 45-step pseudo collision attack on SHA-256 with the complexity of $2^{126.5} (= 2^{(256-3)/2})^5$. Similarly, we can construct the 50-step pseudo collision attack on SHA-512 based on the 50-step MITM preimage attack [14]. In this attack, both neutral words have 3-bit freedom degrees, and the bit size of the matching point is 3. Thus, the complexity of the attack is estimated as $2^{254.5} (= 2^{(512-3)/2})$.

⁵ Our attack uses only 3 bits for the matching and find 3-bit partial target preimages, because this setting is optimal with respect to the time complexity.

In addition, [14] showed pseudo preimage attacks on the 52-step SHA-256 and the 57-step SHA-512. For the setting of a pseudo preimage attack, the cost of converting a pseudo preimage to a preimage is omitted. Thus, larger number of rounds can be attacked. Note that in these attacks, the amount of freedom degrees for both neutral words are only 1-bit, and the bit size of the matching point is 1. In order to construct a pseudo collision attack by using our conversion technique, it is sufficient to obtain a pseudo preimage on a compression function, i.e., a preimage on a hash function is not needed. Therefore, the above explained pseudo preimage attacks can also be converted into pseudo collision attacks in a similar way. The complexities of the pseudo collision attacks on the 52-step SHA-256 and the 57-step SHA-512 are estimated as $2^{127.5}$ ($= 2^{(256-1)/2}$) and $2^{255.5}$ ($= 2^{(512-1)/2}$), respectively.

5 Application to Skein

In this section, we show pseudo collision attacks on the reduced Skein-512 [9] based on the preimage attacks presented in [14].

5.1 Description of Skein

Skein is built from the tweakable block cipher Threefish $E_{K,T}(P)$, where K , T and P denote a key, a tweak and a plaintext message, respectively. The compression function $F(CV, T, M)$ of Skein outputs the next chaining variable as $F(CV, T, M) = E_{CV,T}(M) \oplus M$, where CV is the previous chaining variable and M is an input message block.

Threefish-512 supports a 512-bit block and a 512-bit key, and operates on 64-bit words. The subkey $K^s = (K_0^s, K_1^s, \dots, K_7^s)$ injected every four rounds is generated from the secret key $K = K[0], K[1], \dots, K[7]$ as follows:

$$\begin{aligned} K_j^s &= K[(s + j) \bmod 9], (0 \leq j \leq 4); & K_5^s &= K[(s + 5) \bmod 9] + T[s \bmod 3]; \\ K_6^s &= K[(s + 6) \bmod 9] + T[(s + 1) \bmod 3]; & K_7^s &= K[(s + 7) \bmod 9] + s, \end{aligned}$$

where s denotes a round counter, $T[0]$ and $T[1]$ denote tweak words, $T[2] = T[0] + T[1]$, and $K[8] = C_{240} \oplus \bigoplus_{j=0}^7 K[j]$ with a constant C_{240} . Each Threefish-512 round consists of four *MIX* functions followed by a permutation of the eight 64-bit words. The 128-bit function *MIX* processes the pairs of eight words of internal state I^0, I^1, \dots, I^7 after key addition.

5.2 Known Pseudo Preimage Attacks on Skein [14].

We briefly review two MITM preimage attacks on Skein-512 presented in [14]: one is a preimage attack on the 22-round reduced Skein-512 hash function starting from the 3rd round, and the other is a preimage attack on the 37-round reduced Skein-512 compression function starting from the 2nd round.

For the 22-round attack, the 3-dimension biclique at rounds 12-15 is obtained with the complexity of 2^{200} . Since many bicliques can be produced out of one,

the cost of constructing the bicliques is negligible in the total complexity of the attack. In this attack, we can obtain 2^3 pairs matched in 3 bits by $2^{2.3}$ calls of the 22-round Skein-512 compression function. As a result, a preimage of the 22-round reduced Skein is found with the complexity of $2^{511.2}$.

Table 2. Parameters of the (pseudo) preimage attacks on the reduced Skein-512 [14]

Parameters of the preimage attack on the 22-round Skein-512 hash function						
Chunks			Matching			
Forward	Backward	Biclique	Partial matching	Matching bits	Total matching pairs	Complexity
8-11	16-19	12-15	$20 \rightarrow 24 = 3 \leftarrow 7$	$I_{30,31,53}^1$	2^3	$2^{2.3}$

Parameters of the pseudo preimage attack on the 37-round Skein-512 compression function						
Chunks			Matching			
Forward	Backward	Biclique	Partial matching	Matching bits	Total matching pairs	Complexity
8-15	24-31	16-23	$32 \rightarrow 38 = 2 \leftarrow 7$	I_{25}^3	2	$2^{1.2}$

Considering a pseudo preimage attack on the compression function, it is natural to assume that tweak bits T can also be controlled by the attacker. Due to additional freedom, the pseudo preimage attack on the 37-round reduced Skein-512 is feasible by using the 1-dimension biclique at rounds 16-23. In this attack, we can obtain 2 pairs matched in 1 bit by $2^{1.2}$ calls of the 37-round Skein-512 compression function. Consequently, a pseudo preimage of the 37-round reduced Skein is found with the complexity of $2^{511.2}$.

The parameters for the preimage attacks on the 22-round and the 37-round reduced Skein-512 hash function and compression function are summarized in Table 2. See [14] for more details about this attack.

5.3 Pseudo Collision Attacks on Skein

Since the matching point used in the MITM preimage attack on the 22-round reduced Skein-512 hash function [14] is located in the end of the compression function, our conversion technique can directly convert it to the pseudo collision attack on the 22-round reduced Skein-512. In this attack, the neutral words have 3-bit freedom degrees, and the bit size of the matching point is 3. As reported in [14], a 3-bit matching candidate can be found with the complexity of $2^{2.3}/2^3$. Thus, the complexity of the pseudo collision attack on the 22-round reduced Skein-512 hash function is estimated as $2^{253.8}$ ($= 2^{(512-3)/2} \times 2^{2.3}/2^3$).

The pseudo preimage attack on the 37-round reduced Skein compression function can be converted into a pseudo collision attack on a hash function in a similar way. The required complexity for the pseudo collision attack on the 37-round reduced Skein hash function is estimated as $2^{255.7}$ ($= 2^{(512-1)/2} \times 2^{1.2}/2$).

6 Conclusion

In this paper, we gave a generic method to convert preimage attacks to pseudo collision attacks. It provides a new insight to evaluate the security of hash

functions. The essence of the method is converting a partial target preimage attack to a pseudo collision attack. That is especially compatible to meet-in-the-middle preimage attacks since it can be converted into a partial target preimage attack if the matching point can be moved to the end of a hash function or a compression function and enough freedom on neutral bits are left.

Using the proposed approach, we presented the best pseudo collision attacks on SHA-2 based on the known preimage attacks, which has been left as open question. We showed pseudo collision attacks on the 43- and 46-step reduced SHA-256 and SHA-512 based on the MITM preimage attacks presented in [4]. Also, pseudo collision attacks on the 52- and 57-step reduced SHA-256 and SHA-512 based on the more advanced MITM preimage attacks in [14] were demonstrated. We also applied the conversion technique to other hash functions including Skein and BLAKE with the meet-in-the-middle preimage attacks, which showed the widely usage of this method. The pseudo collision attacks on the 22- and 37-round reduced Skein-512 were presented. The 4-round reduced BLAKE-256/512 without the initialization function can be attacked by the converted pseudo collision attack (see Appendix A). Our technique may also apply to other hash functions, such as Tiger [1]. Based on the MITM preimage attack on the full Tiger [10], we might construct the pseudo collision attack on the full Tiger. We believe that the technique can be used for more hash algorithms once their preimage or pseudo preimage attacks are found.

By this method, now we only can get pseudo collision attacks. It is left as future works that how to construct collision attacks from known preimage attacks.

Acknowledgments. The author would like to thank the anonymous reviewers for their helpful comments.

References

1. Anderson, R.J., Biham, E.: Tiger: A Fast New Hash Function. In: Gollmann, D. (ed.) FSE 1996. LNCS, vol. 1039, pp. 89–97. Springer, Heidelberg (1996)
2. Aoki, K., Sasaki, Y.: Preimage Attacks on One-Block MD4, 63-Step MD5 and More. In: Avanzi, R.M., Keliher, L., Sica, F. (eds.) SAC 2008. LNCS, vol. 5381, pp. 103–119. Springer, Heidelberg (2009)
3. Aoki, K., Sasaki, Y.: Meet-in-the-Middle Preimage Attacks Against Reduced SHA-0 and SHA-1. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 70–89. Springer, Heidelberg (2009)
4. Aoki, K., Guo, J., Matusiewicz, K., Sasaki, Y., Wang, L.: Preimages for Step-Reduced SHA-2. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 578–597. Springer, Heidelberg (2009)
5. Aumasson, J.-P., Henzen, L., Meier, W., Phan, R.C.-W.: SHA-3 proposal BLAKE (version 1.3). Submission to NIST (December 2010), <http://131002.net/blake/blake.pdf>
6. Biryukov, A., Lamberger, M., Mendel, F., Nikolić, I.: Second-Order Differential Collisions for Reduced SHA-256. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 270–287. Springer, Heidelberg (2011)

7. 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)
8. De Cannière, C., Rechberger, C.: Preimages for Reduced SHA-0 and SHA-1. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 179–202. Springer, Heidelberg (2008)
9. Ferguson, N., Lucks, S., Schneier, B., Whiting, D., Bellare, M., Kohno, T., Callas, J., Walker, J.: The Skein hash function family (version 1.3, October 1, 2010), <http://www.schneier.com/skein1.3.pdf>
10. Guo, J., Ling, S., Rechberger, C., Wang, H.: Advanced Meet-in-the-Middle Preimage Attacks: First Results on Full Tiger, and Improved Results on MD4 and SHA-2. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 56–75. Springer, Heidelberg (2010)
11. Indestege, S., Mendel, F., Preneel, B., Rechberger, C.: Collisions and Other Non-random Properties for Step-Reduced SHA-256. In: Avanzi, R.M., Keliher, L., Sica, F. (eds.) SAC 2008. LNCS, vol. 5381, pp. 276–293. Springer, Heidelberg (2009)
12. Isobe, T.: A Single-Key Attack on the Full GOST Block Cipher. In: Joux, A. (ed.) FSE 2011. LNCS, vol. 6733, pp. 290–305. Springer, Heidelberg (2011)
13. Isobe, T., Shibutani, K.: Preimage Attacks on Reduced Tiger and SHA-2. In: Dunkelman, O. (ed.) FSE 2009. LNCS, vol. 5665, pp. 139–155. Springer, Heidelberg (2009)
14. Khovratovich, D., Rechberger, C., Savelieva, A.: Bicliques for Preimages: Attacks on Skein-512 and the SHA-2 Family. In: Canteaut, A. (ed.) FSE 2012. LNCS, vol. 7549, pp. 244–263. Springer, Heidelberg (2012)
15. Lai, X., Massey, J.L.: Hash Functions Based on Block Ciphers. In: Rueppel, R.A. (ed.) EUROCRYPT 1992. LNCS, vol. 658, pp. 55–70. Springer, Heidelberg (1993)
16. Leurent, G.: MD4 is Not One-Way. In: Nyberg, K. (ed.) FSE 2008. LNCS, vol. 5086, pp. 412–428. Springer, Heidelberg (2008)
17. Mendel, F., Nad, T., Schläffer, M.: Finding SHA-2 characteristics: Searching through a minefield of contradictions. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 288–307. Springer, Heidelberg (2011)
18. Mendel, F., Pramstaller, N., Rechberger, C., Rijmen, V.: Analysis of Step-Reduced SHA-256. In: Robshaw, M.J.B. (ed.) FSE 2006. LNCS, vol. 4047, pp. 126–143. Springer, Heidelberg (2006)
19. Menezes, A.J., van Oorschot, P.C., Vanstone, S.A.: Handbook of Applied Cryptography. CRC Press (1997)
20. Nikolić, I., Biryukov, A.: Collisions for Step-Reduced SHA-256. In: Nyberg, K. (ed.) FSE 2008. LNCS, vol. 5086, pp. 1–15. Springer, Heidelberg (2008)
21. Quisquater, J.-J., Delescaille, J.-P.: How Easy Is Collision Search? Application to DES (Extended Summary). In: Quisquater, J.-J., Vandewalle, J. (eds.) EUROCRYPT 1989. LNCS, vol. 434, pp. 429–434. Springer, Heidelberg (1990)
22. Saarinen, M.-J.O.: A Meet-in-the-Middle Collision Attack Against the New FORK-256. In: Srinathan, K., Rangan, C.P., Yung, M. (eds.) INDOCRYPT 2007. LNCS, vol. 4859, pp. 10–17. Springer, Heidelberg (2007)
23. Sanadhya, S.K., Sarkar, P.: 22-step collisions for SHA-2. CoRR, abs/0803.1220 (2008)
24. Sanadhya, S.K., Sarkar, P.: Attacking Reduced Round SHA-256. In: Bellare, S.M., Gennaro, R., Keromytis, A.D., Yung, M. (eds.) ACNS 2008. LNCS, vol. 5037, pp. 130–143. Springer, Heidelberg (2008)

25. Sanadhya, S.K., Sarkar, P.: New Collision Attacks against Up to 24-Step SHA-2. In: Chowdhury, D.R., Rijmen, V., Das, A. (eds.) INDOCRYPT 2008. LNCS, vol. 5365, pp. 91–103. Springer, Heidelberg (2008)
26. Sanadhya, S.K., Sarkar, P.: Non-linear Reduced Round Attacks against SHA-2 Hash Family. In: Mu, Y., Susilo, W., Seberry, J. (eds.) ACISP 2008. LNCS, vol. 5107, pp. 254–266. Springer, Heidelberg (2008)
27. Sasaki, Y., Aoki, K.: Finding Preimages in Full MD5 Faster Than Exhaustive Search. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 134–152. Springer, Heidelberg (2009)
28. Sasaki, Y., Aoki, K.: Preimage Attacks on 3, 4, and 5-Pass HAVAL. In: Pieprzyk, J. (ed.) ASIACRYPT 2008. LNCS, vol. 5350, pp. 253–271. Springer, Heidelberg (2008)
29. Sedgewick, R., Szymanski, T.G., Yao, A.C.-C.: The complexity of finding cycles in periodic functions. *SIAM J. Comput.* 11(2), 376–390 (1982)
30. Wang, L., Ohta, K., Sakiyama, K.: Free-start preimages of round-reduced Blake compression function. Rump session at ASIACRYPT 2009 (2009)
31. Wang, X., Yin, Y.L., Yu, H.: Finding Collisions in the Full SHA-1. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 17–36. Springer, Heidelberg (2005)
32. Wang, X., Yu, H.: How to Break MD5 and Other Hash Functions. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 19–35. Springer, Heidelberg (2005)
33. Watanabe, D.: OFFICIAL COMMENT: LUX. NIST mailing list (2009), http://csrc.nist.gov/groups/ST/hash/sha-3/Round1/documents/LUX_Comments.pdf

Appendix

A Application to BLAKE

We apply our technique to BLAKE hash function family consisting of BLAKE-224, BLAKE-256, BLAKE-384 and BLAKE-512 [5]. We utilize the result presented in [30] which showed a pseudo preimage attack on the 4-round reduced BLAKE compression function without the initialization function. While the practical impact on the attack for this reduced BLAKE compression function is debatable, a pseudo collision on the reduced BLAKE can be directly derived by using our conversion technique. As a result, we can find a pseudo collision of the 4-round reduced BLAKE-256 compression function without the initialization with the complexity of 2^{112} . Similarly, a pseudo collision of the 4-round reduced BLAKE-512 compression function without the initialization can be found with the complexity of 2^{224} .

A.1 Description of BLAKE

The compression function of BLAKE-256 consists of *initialization*, *round function* and *finalization*.

Table 3. Message and Constants Permutation

i	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
σ_0	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
σ_1	14	10	4	8	9	15	13	6	1	12	0	2	11	7	5	3
σ_2	11	8	12	0	5	2	15	13	10	14	3	6	7	1	9	4
σ_3	7	9	3	1	13	12	11	14	2	6	5	10	4	0	15	8

Initialization. 8 words of chaining variables h_0, \dots, h_7 are transformed into 16 words of an initial state v_0, \dots, v_{15} as $v_i = h_i$ for $0 \leq i < 8$, where $h_i, v_j \in \{0, 1\}^{32}$. The other 8 words of the initial state v_i ($8 \leq i < 16$) are determined from the given salts s_0, \dots, s_3 and counter t_0, t_1 , where $s_i, t_j \in \{0, 1\}^{32}$.

Round Function. An initial state v is updated by 14 round functions with message words m_0, \dots, m_{15} and constants c_0, \dots, c_7 , where $m_i, c_j \in \{0, 1\}^{32}$. Each round function includes the following steps, $G_0(v_0, v_4, v_8, v_{12}), G_1(v_1, v_5, v_9, v_{13}), G_2(v_2, v_6, v_{10}, v_{14}), G_3(v_3, v_7, v_{11}, v_{15}), G_4(v_0, v_5, v_{10}, v_{15}), G_5(v_1, v_6, v_{11}, v_{12}), G_6(v_2, v_7, v_8, v_{13}), G_7(v_3, v_4, v_7, v_{14})$. The function $G_i(a, b, c, d)$ is defined as:

$$\begin{aligned}
 a &\leftarrow a + b + (m_{\sigma_r(2i)} \oplus c_{\sigma_r(2i+1)}), & d &\leftarrow (d \oplus a) \ggg 16, \\
 c &\leftarrow c + d, & b &\leftarrow (b \oplus c) \ggg 12, \\
 a &\leftarrow a + b + (m_{\sigma_r(2i+1)} \oplus c_{\sigma_r(2i)}), & d &\leftarrow (d \oplus a) \ggg 8, \\
 c &\leftarrow c + d, & b &\leftarrow (b \oplus c) \ggg 7,
 \end{aligned}$$

where permutations $\sigma_r(j)$ ($0 \leq j < 16$) of the first 4 rounds refer to Table 3. The functions G_0 to G_3 and G_4 to G_7 denote the column transforms and the diagonal transforms, respectively.

Finalization. After the round functions, the new chaining variables are extracted with the updated state, the salts and the feed-forward of the initial chaining variables as follows.

$$\begin{aligned}
 h'_0 &\leftarrow h_0 \oplus s_0 \oplus v_0 \oplus v_8 & h'_1 &\leftarrow h_1 \oplus s_1 \oplus v_1 \oplus v_9 \\
 h'_2 &\leftarrow h_2 \oplus s_2 \oplus v_2 \oplus v_{10} & h'_3 &\leftarrow h_3 \oplus s_3 \oplus v_3 \oplus v_{11} \\
 h'_4 &\leftarrow h_4 \oplus s_0 \oplus v_4 \oplus v_{12} & h'_5 &\leftarrow h_5 \oplus s_1 \oplus v_5 \oplus v_{13} \\
 h'_6 &\leftarrow h_6 \oplus s_2 \oplus v_6 \oplus v_{14} & h'_7 &\leftarrow h_7 \oplus s_3 \oplus v_7 \oplus v_{15}
 \end{aligned}$$

BLAKE-512 operates on 64-bit words and outputs 512 bits. The compression function of BLAKE-512 is similar to that of BLAKE-256 except for the number of rounds (16 instead of 14), and the constants and the amount of rotation used in G functions.

A.2 Known MITM Preimage Attacks on 4-Round Compression Function of BLAKE [30]

In the setting of the pseudo preimage attack on the reduced BLAKE compression function presented in [30], the initialization step is disregarded, and an attacker

can select a random start value from the start of round functions (the end of initialization step) as shown in Fig. 5.




Fig. 5. MITM preimage attack for finalization

Figure 6 shows the overview of the pseudo preimage attack on the 4-round reduced BLAKE compression function without the initialization. Let an input state of the round i be $v^{i-1} = \{v_0^{i-1}, \dots, v_{15}^{i-1}\}$, where $v_j^i \in \{0, 1\}^{32}$. In this attack, message words m_4 and m_6 are used as the neutral words, and the starting point of the attack is the state after the column transformation of the round 3. In the forward computation from the starting point, v_6^4, v_{14}^4 can be computed without using m_6 . Similarly, in the backward computation, v_6^0 can be computed without using m_4 . Therefore, storing m_4, v_6^4, v_{14}^4 in a list L_F , and m_6, v_6^0 in a list L_B , we expect to find matching pairs satisfying $h'_6 = v_6^0 \oplus v_6^4 \oplus v_{14}^4$. As a result, a pseudo preimage of the 4-round reduced BLAKE without the initialization is found with the complexity of 2^{224} .

A.3 Pseudo Collision Attacks on BLAKE Compression Function

Since the matching point of the known pseudo preimage attack is at the end of the compression function, a pseudo collision attack can be directly constructed from it.

Attack Procedure

1. Randomly choose the 7-th word of the output value h'_6 , which is the target value.
2. Randomly choose the values of state words and message words except for m_4 and m_6 .
3. For all 2^{32} possible m_4 , compute forward and find v_6^4 and v_{14}^4 . Store the pairs $(m_4, v_6^4 \oplus v_{14}^4)$ in a list L_F
4. For all 2^{32} possible m_6 , compute forward and find v_6^0 . Store the pairs $(m_4, h'_6 \oplus v_6^0)$ in a list L_B .
5. Compare the value $v_6^4 \oplus v_{14}^4$ and $h'_6 \oplus v_6^0$ in two lists L_F and L_B .
6. Once matching, compute states $v_0^0, v_1^0, \dots, v_{15}^0$ and $v_0^4, v_1^4, \dots, v_{15}^4$. Compute output values $h'_0, h'_1, \dots, h'_{15}$ according to finalization steps and store with message words together. Then obtain 2^{32} items in which the value of h'_6 are fixed.

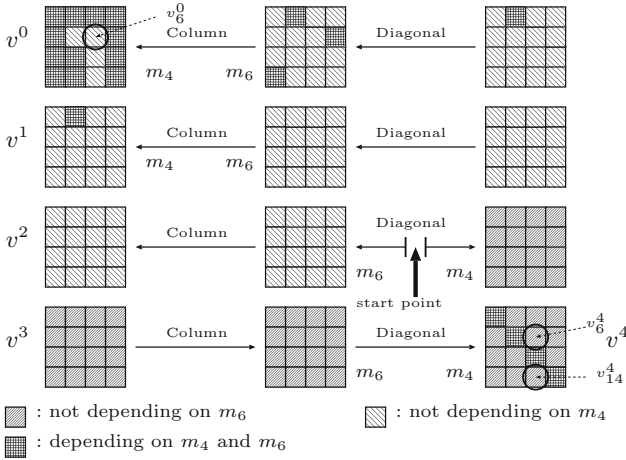


Fig. 6. Pseudo preimage attacks on reduced BLAKE compression function

7. Repeat steps (2) - (6) 2^{80} times.

We can obtain 2^{112} items in which the value of h'_6 are fixed. A colliding pair exists with a high probability that the other 224 bits of output values are also same. Finally, we can find a pseudo collision of the 4-round reduced BLAKE-256 compression function with the complexity of $2^{112} = 2^{80} \cdot 2^{32}$.

The attack is applicable to the reduced BLAKE-512 in a similar way, since the components of BLAKE-512 are similar to those of BLAKE-256. In BLAKE-224, the variable h'_7 is truncated and discarded. However, the truncation does not affect our conversion, since we use h'_6 as a partial target preimage. Thus, a pseudo collision attack on the 4-round reduced BLAKE-224 without the initialization can be constructed with the complexity of $2^{96} (= 2^{(224-32)/2})$. For BLAKE-384, in contrast to the other variants, the variable h'_6 is discarded by the truncation as well. Therefore, it is hard to straightforwardly apply our conversion to the reduced BLAKE-384, since h'_6 cannot be used as a partial target preimage.

UNAF: A Special Set of Additive Differences with Application to the Differential Analysis of ARX^{*}

Vesselin Velichkov^{1,2,**}, Nicky Mouha^{1,2,***},
Christophe De Cannière^{1,2,†}, and Bart Preneel^{1,2}

¹ Department of Electrical Engineering ESAT/SCD-COSIC,
Katholieke Universiteit Leuven. Kasteelpark Arenberg 10, B-3001 Heverlee, Belgium

² Interdisciplinary Institute for BroadBand Technology (IBBT), Belgium

Vesselin.Velichkov@gmail.com, Nicky.Mouha@esat.kuleuven.be

Abstract. Due to their fast performance in software, an increasing number of cryptographic primitives are constructed using the operations addition modulo 2^n , bit rotation and XOR (ARX). However, the resistance of ARX-based ciphers against differential cryptanalysis is not well understood. In this paper, we propose a new tool for evaluating more accurately the probabilities of additive differentials over multiple rounds of a cryptographic primitive. First, we introduce a special set of additive differences, called UNAF (unsigned non-adjacent form) differences. Then, we show how to apply them to find good differential trails using an algorithm for the automatic search for differentials. Finally, we describe a key-recovery attack on stream cipher Salsa20 reduced to five rounds, based on UNAF differences.

Keywords: UNAF, ARX, Salsa20, additive differential probability, differential cryptanalysis.

1 Introduction

Differential cryptanalysis [4] and linear cryptanalysis [14] have shown to be two of the most powerful techniques in the cryptanalysis of symmetric-key cryptographic primitives. Security against linear and differential cryptanalysis is therefore typically a major design criterion for modern ciphers. An example of this is the wide-trail design strategy, used to provide provable resistance against linear and differential cryptanalysis for the AES block cipher [6].

* This work was supported in part by the Research Council K.U.Leuven: GOA TENSE, and by the IAP Program P6/26 BCRYPT of the Belgian State (Belgian Science Policy), and in part by the European Commission through the ICT program under contract ICT-2007-216676 ECRYPT II.

** DBOF Doctoral Fellow, K.U.Leuven, Belgium.

*** This author is funded by a research grant of the Institute for the Promotion of Innovation through Science and Technology in Flanders (IWT-Vlaanderen).

† Postdoctoral Fellow of the Research Foundation – Flanders (FWO).

In order to achieve a fast performance in software, an increasing number of cryptographic primitives are built using the operations addition modulo 2^n , rotation and XOR (ARX). Examples include the block cipher FEAL [17], the Salsa20 stream cipher family [3], as well as the SHA-3 finalists BLAKE [2] and Skein [9]. Although ARX-based algorithms are very popular, their resistance to differential cryptanalysis [4] is not well understood.

The probability with which differences propagate through a sequence of operations must be calculated efficiently and accurately, in order to correctly assess the security of a cipher against differential cryptanalysis. Lipmaa et al. studied the xor-differential probability of addition (xdp^+) in [12], and the additive differential probability of XOR (adp^\oplus) in [13]. These results were generalized using the S-functions framework, introduced by Mouha et al. [15].

As shown by Velichkov et al. [18], the additive differential probability of ARX (adp^{ARX}) can differ significantly from the multiplication of the differential probability of the separate components – addition, rotation and XOR. Although an algorithm was proposed in [18] for the exact calculation of adp^{ARX} , unfortunately their method does not scale to analyze larger components. The accurate calculation of the probability of a differential characteristic therefore still remains an open problem for ARX constructions.

In this paper we take a different approach. Namely, we do not calculate the *exact* differential probability of a component consisting of more than one ARX operations. Instead, we multiply the differential probabilities of several ARX operations in order to estimate the total probability. As we want to avoid that this calculation differs significantly from the actual probability (e.g. due to dependencies between the inputs as noted in [18]), we propose to use a new type of difference: the UNAF difference, which represents a set of specially chosen additive differences.

We apply UNAF differences to the cryptanalysis of the ARX-based stream cipher Salsa20. A general algorithm for automatic search of differentials is briefly discussed. We apply it to find several differentials for three rounds of Salsa20. By multiplying the probabilities adp^{ARX} of separate ARX components, we estimate that the best differential has a probability of 2^{-10} . Using UNAF differences, the same probability is evaluated as 2^{-4} . Experimentally, we estimate the probability of this differential to be $2^{-3.39}$. We observe that the probability obtained using UNAF differences is much closer to the experimental value.

Finally, we apply UNAF differences to mount key-recovery attack on a version of Salsa20 reduced to 5 rounds. Note that this is not the best known attack on Salsa20. It is therefore provided only as a demonstration of a practical application of UNAF differences. Furthermore, we expect that our attack can be extended to more rounds.

The outline of the paper is as follows. In Sect. 2, we describe the UNAF framework. It is applied to the differential analysis of stream cipher Salsa20 in Sect. 3. Sect. 4 concludes the paper. Notation is defined in Table 1.

Table 1. Notation

Symbol	Meaning
n	Number of bits in a word
x	n -bit word
$x[i]$	Select the $(i \bmod n)$ -th bit (or element) of the n -bit word x , $x[0]$ is the least-significant bit (or element)
$ x $	The absolute value of x
\bar{x}	The negation of x i.e. $\bar{x} = -x$ (e.g. $\bar{1} = -1$)
$\#A$	Number of elements in the set A
$+$, $-$	Addition modulo 2^n , subtraction modulo 2^n
\oplus	Exclusive-OR (XOR)
$\lll t$	Left bit rotation by t positions
$\alpha \rightarrow \beta$	Input difference α propagates to output difference β
w_i^r	32-bit word i from the input state to round $r + 1$ of Salsa20
Δ_i^r	Additive difference in word i of the input to round $r + 1$ of Salsa20
0_i^r	Zero difference in word i of the input to round $r + 1$ of Salsa20
$\{\Delta_i^U\}_i^r$	UNAF difference in word i of the input to round $r + 1$ of Salsa20
ARX	The sequence of the operations: $+$, \lll , \oplus as a single operation
HW (x)	Hamming weight of x (number of non-zero bits in x)

2 The UNAF Framework

In this section, we describe the UNAF framework. We define UNAF differences and state the main UNAF theorem. The UNAF differential probability of **ARX** (udp^{ARX}) is defined and a general algorithm for the automatic search for high-probability differentials is briefly discussed.

2.1 Preliminaries

Before we give the formal definition of UNAF differences, we first recall a few related concepts: the binary-signed digit (BSD) difference and the non-adjacent form (NAF) difference.

Definition 1. (BSD difference) *A BSD difference is a difference whose bits are signed and take values in the set $\{\bar{1}, 0, 1\}$:*

$$\Delta^{\pm}a : \Delta^{\pm}a[i] = (a_2[i] - a_1[i]) \in \{\bar{1}, 0, 1\}, \quad 0 \leq i < n. \quad (1)$$

An additive difference Δ^+a can be composed of more than one BSD difference $\Delta^{\pm}a$. From any BSD difference, the corresponding additive difference can be computed as: $\Delta^+a = \sum_{i=0}^{n-1} \Delta^{\pm}a[i] \cdot 2^i$.

All BSD differences corresponding to Δ^+a can be obtained by replacing 01 with $1\bar{1}$ and vice versa and by replacing $0\bar{1}$ with $\bar{1}1$ and vice versa [7,16]. Note also that the number of pairs (a_1, a_2) that satisfy the n -bit difference Δ^+a is

2^n , while the number of pairs that satisfy any of its BSD differences $\Delta^\pm a$ is 2^k , where k is the number of zeros in the word $\Delta^\pm a$. Therefore, the following inequality holds: $2^k \leq 2^n$, $k = n - \text{HW}(\Delta^\pm a)$.

The non-adjacent form (NAF) difference is a special BSD difference and is defined as follows:

Definition 2. (NAF) *A NAF (non-adjacent form) difference is a BSD difference in which no two adjacent bits are non-zero:*

$$\Delta^N a : \nexists i : (\Delta^N a[i] \neq 0) \wedge (\Delta^N a[i+1] \neq 0), \quad 0 \leq i < n - 1 . \quad (2)$$

For every additive difference $\Delta^+ a$, there is exactly one NAF difference $\Delta^N a$ (ignoring the sign of the MSB). No other BSD difference has a lower Hamming weight than $\Delta^N a$ [16]. We illustrate this with the following example:

Example 1. Let $n = 4$ and $\Delta^+ a = 3$. Then all possible BSD differences corresponding to $\Delta^+ a$ are $0011, 010\bar{1}, 01\bar{1}1, 1\bar{1}\bar{1}1, \bar{1}\bar{1}\bar{1}1, \bar{1}10\bar{1}$ and $\bar{1}\bar{1}0\bar{1}$. Of them, only $010\bar{1}$ is in non-adjacent form (NAF). It also has the lowest Hamming weight among all BSD differences, namely 2.

By enumerating all possible combinations of signs of the non-zero bits of $\Delta^N a$, we can construct a special set of additive differences. What is special about this set, is that all of its elements correspond to the same unsigned NAF difference. This set is a UNAF difference and is denoted by $\Delta^U a$. More formally:

Definition 3. (UNAF) *A UNAF difference is a set of additive differences that correspond to the same unsigned NAF difference (i.e. a NAF difference with the signs ignored):*

$$\Delta^U a = \{ \Delta^+ x : |\Delta^N x| = |\Delta^N a| \} . \quad (3)$$

It is easy to see that the size of the UNAF set $\Delta^U a$ is 2^k , where k is the Hamming weight of the n -bit word $\Delta^N a$, excluding the MSB. We further clarify the concept of a UNAF difference with the following example:

Example 2. Consider again an example where $n = 4$. Let $\Delta^+ a = 3$, thus $\Delta^N a = 010\bar{1}$. Then, $\Delta^U a = \{ \Delta^+ x_1 = 3, \Delta^+ x_2 = -3, \Delta^+ x_3 = 5, \Delta^+ x_4 = -5 \}$. This follows from $|\Delta^N x_1| = |\Delta^N x_2| = |\Delta^N x_3| = |\Delta^N x_4| = |\Delta^N a|$, because $|010\bar{1}| = |0\bar{1}01| = |0101| = |0\bar{1}0\bar{1}| = 0101$.

2.2 Main UNAF Theorem

The main UNAF theorem provides the motivation for applying UNAF differences to the differential analysis of ARX. Before we state it, we define the additive differential probability of XOR (adp^\oplus).

The differential probability of the operation XOR, when differences are expressed using addition modulo 2^n , is denoted by adp^\oplus . For fixed additive differences α, β and γ , adp^\oplus is equal to the number of pairs (a_1, b_1) for which the equality $((a_1 + \alpha) \oplus (b_1 + \beta)) - (a_1 \oplus b_1) = \gamma$ holds, divided by the total number of such pairs. More formally, $\text{adp}^\oplus(\alpha, \beta \rightarrow \gamma)$ is defined as:

Definition 4. (adp^\oplus)

$$\begin{aligned} \text{adp}^\oplus(\alpha, \beta \rightarrow \gamma) &= \frac{\#\{(a_1, b_1) : c_2 - c_1 = \gamma\}}{\#\{(a_1, b_1)\}} \\ &= 2^{-2n} \cdot \#\{(a_1, b_1) : c_2 - c_1 = \gamma\} , \end{aligned} \quad (4)$$

where $c_1 = a_1 \oplus b_1$, $c_2 = (a_1 + \alpha) \oplus (b_1 + \beta)$ and 2^{2n} is the total number of pairs (a_1, b_1) .

Efficient algorithms for the computation of adp^\oplus were studied in [13,15]. Next we state the main UNAF theorem. Its proof is given in Appendix A.

Theorem 1. (Main UNAF theorem) *If the probability with which input additive differences Δ^+a and Δ^+b propagate to output difference Δ^+c through XOR is non-zero, then the probability with which any of the input additive differences belonging to the corresponding UNAF sets resp. Δ^Ua and Δ^Ub propagate to any of the output additive differences belonging to the UNAF set Δ^Uc is also non-zero:*

$$\begin{aligned} \text{adp}^\oplus(\Delta^+a, \Delta^+b \rightarrow \Delta^+c) > 0 &\implies \text{adp}^\oplus(\Delta^+a_i, \Delta^+b_j \rightarrow \Delta^+c_k) > 0 , \\ \forall i, j, k : \Delta^+a_i \in \Delta^Ua, \Delta^+b_j \in \Delta^Ub, \Delta^+c_k \in \Delta^Uc . \end{aligned} \quad (5)$$

Theorem 1 states that if a given additive differential is possible w.r.t. the XOR operation, then all additive differentials whose inputs and outputs belong to the same UNAF sets, are also possible. This is illustrated with the following example.

Example 3. Let $n = 4$ and $\Delta^+a = 5$, $\Delta^+b = 1$, $\Delta^+c = 6$. Because $\text{adp}^\oplus(5, 1 \rightarrow 6) = 0.15625 > 0$, we can use Theorem 1 to show that $\text{adp}^\oplus(\Delta^+a_i, \Delta^+b_j \rightarrow \Delta^+c_k) > 0$ for any $\Delta^+a_i \in \Delta^Ua = \{3, -3, 5, -5\}$, $\Delta^+b_j \in \Delta^Ub = \{1, -1\}$ and $\Delta^+c_k \in \Delta^Uc = \{6, -6\}$.

In the next section we investigate the probability with which UNAF differences propagate through the ARX operation.

2.3 The UNAF Differential Probability of ARX

The UNAF differential probability of ARX represents the probability with which the sets of input additive differences Δ^Ua , Δ^Ub and Δ^Ud propagate to the set of output additive differences Δ^Ue . It is defined as:

Definition 5. (udp^{ARX})

$$\begin{aligned} \text{udp}^{\text{ARX}}(\Delta^Ua, \Delta^Ub, \Delta^Ud \xrightarrow{t} \Delta^Ue) &= \\ \frac{\#\{(a_1, b_1, d_1) : \Delta^+a \in \Delta^Ua, \Delta^+b \in \Delta^Ub, \Delta^+d \in \Delta^Ud, \Delta^+e \in \Delta^Ue\}}{\#\{(a_1, b_1, d_1) : \Delta^+a \in \Delta^Ua, \Delta^+b \in \Delta^Ub, \Delta^+d \in \Delta^Ud\}} , \end{aligned} \quad (6)$$

where

$$\Delta^+e = e_2 - e_1 = \text{ARX}(a_1 + \Delta^+a, b_1 + \Delta^+b, d_1 + \Delta^+d, t) - \text{ARX}(a_1, b_1, d_1, t),$$

and $\text{ARX}(x, y, z, t) = ((x + y) \lll t) \oplus z$.

The probability udp^{ARX} is computed using a method conceptually similar to the one proposed for the computation of adp^{ARX} in [18]. The main difference is that in this case we are dealing with *sets* of input and output additive differences. Details on this computation are provided in Appendix B.

2.4 An Algorithm for Finding the Best Output Difference

To demonstrate how the UNAF framework can be used to construct high-probability differential characteristics, we have developed a general algorithm for the automatic search of differentials. It is capable of computing the highest probability output difference from a given operation. The proposed algorithm is applicable to any type of difference and any operation. The only condition is that the propagation of the difference through the operation can be represented as an S-function. The method to find the best output difference is based on the A* search algorithm [11].

Space constraints do not allow us to present the algorithm here in detail. However, a full description of the algorithm accompanied by pseudo-code can be found in Appendix C. Furthermore, a software toolkit that implements this algorithm is available.¹

In the following sections we describe an application of the algorithm and of UNAF differences to the differential analysis of stream cipher Salsa20.

3 Applications

We describe several applications of the UNAF framework to the differential analysis of stream cipher Salsa20. UNAF differences can be used to obtain more accurate estimations of the probabilities of differentials through multiple rounds of ARX operations. We describe a key-recovery attack using UNAF differentials on a version of Salsa20, reduced to 5 rounds.

3.1 Description of Salsa20

Salsa20 is a stream cipher proposed by Bernstein in [3]. It is one of the finalists of the eSTREAM competition [8]. Salsa20 operates on 32-bit words. The inputs are a 256-bit key (k_0, k_1, \dots, k_7) , a 64-bit nonce (v_0, v_1) , a 64-bit counter (t_0, t_1) and four predefined 32-bit constants c_0, c_1, c_2, c_3 . These inputs are mapped to a two-dimensional square matrix as follows:

$$\begin{bmatrix} c_0 & k_0 & k_1 & k_2 \\ k_3 & c_1 & v_0 & v_1 \\ t_0 & t_1 & c_2 & k_4 \\ k_5 & k_6 & k_7 & c_3 \end{bmatrix} \rightarrow \begin{bmatrix} w_0^0 & w_1^0 & w_2^0 & w_3^0 \\ w_4^0 & w_5^0 & w_6^0 & w_7^0 \\ w_8^0 & w_9^0 & w_{10}^0 & w_{11}^0 \\ w_{12}^0 & w_{13}^0 & w_{14}^0 & w_{15}^0 \end{bmatrix}. \quad (7)$$

¹ <http://www.ecrypt.eu.org/tools/s-function-toolkit>




Fig. 1. Round $s = r + 1$ of Salsa20

The basic operation of Salsa20 is the *quarterround*. One *quarterround* transforms four of the input words to round $r + 1$: $w_0^r, w_1^r, w_2^r, w_3^r$ into four output words: $w_0^{r+1}, w_1^{r+1}, w_2^{r+1}, w_3^{r+1}$ by the means of four consecutive ARX operations:

$$w_1^{r+1} = w_1^r \oplus ((w_0^r + w_3^r) \lll 7) = \text{ARX}(w_0^r, w_3^r, w_1^r, 7) , \tag{8}$$

$$w_2^{r+1} = w_2^r \oplus ((w_1^{r+1} + w_0^r) \lll 9) = \text{ARX}(w_1^{r+1}, w_0^r, w_2^r, 9) , \tag{9}$$

$$w_3^{r+1} = w_3^r \oplus ((w_2^{r+1} + w_1^{r+1}) \lll 13) = \text{ARX}(w_2^{r+1}, w_1^{r+1}, w_3^r, 13) , \tag{10}$$

$$w_0^{r+1} = w_0^r \oplus ((w_3^{r+1} + w_2^{r+1}) \lll 18) = \text{ARX}(w_3^{r+1}, w_2^{r+1}, w_0^r, 18) . \tag{11}$$

One *round* of Salsa20 consists of four parallel applications of the *quarterround* transformation. Each transformation is applied to the elements (in permuted order) of one of the four columns of the input state matrix, followed by a permutation of the words, as shown on Fig. 1.

Salsa20 has a total of 20 rounds, although versions with eight and twelve rounds have been proposed, resp. Salsa20/8 and Salsa20/12. The output state after the last round is added to the initial input state by means of a feed-forward operation. This produces sixteen 32-bit words (512 bits) of key stream.

3.2 Estimating the Probability of Differentials Using UNAF Differentials

We apply the algorithm of Sect. 2.4 to search for high probability differential characteristics in Salsa20. We use a greedy strategy in which at every ARX operation we select the output UNAF difference with the highest probability, before proceeding with the next ARX operation. In this way we find the following truncated differential for three rounds:

$$\Delta_8^0 = 0x80000000 \rightarrow \Delta_9^3 = 0x80000000 . \tag{12}$$

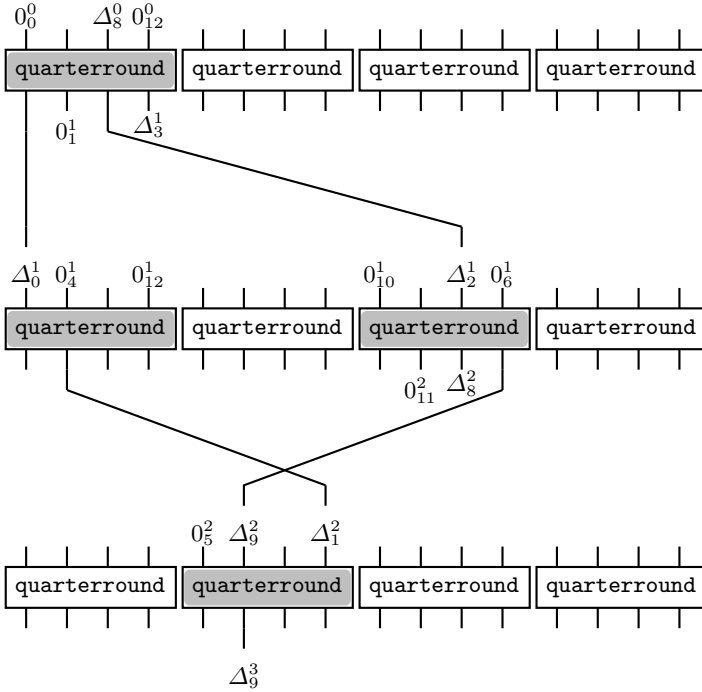


Fig. 2. Three round differential characteristic satisfying the differential $\Delta_8^0 \rightarrow \Delta_9^3$

The expression (12) implies that all words of the input state have zero difference, except for the word at position 8, which has difference $0x80000000$. A three round differential characteristic that satisfies (12) is shown on Fig. 2. The probability with which the differential (12) holds, obtained experimentally over 2^{20} chosen plaintexts, is $p_{\text{exper}} = 2^{-3.39}$.

We compute two theoretical estimations of p_{exper} . The first estimation is based on single additive differences and is denoted \hat{p}_{add} . It is computed as a multiplication of adp^{ARX} probabilities:

$$\hat{p}_{\text{add}} = \prod \text{adp}^{\text{ARX}} = 2^{-10} . \tag{13}$$

The second estimation of p_{exper} is based on UNAF differences and is denoted \hat{p}_{unaf} . It is computed as a multiplication of udp^{ARX} probabilities:

$$\hat{p}_{\text{unaf}} = \prod \text{udp}^{\text{ARX}} = 2^{-4} . \tag{14}$$

The computations (13) and (14) are shown in Table 2 and Table 3 respectively.

Clearly \hat{p}_{unaf} is a better estimation of p_{exper} than \hat{p}_{add} . The reason is that multiple differential characteristics connect the input and output differences of

Table 2. The estimated probability \hat{p}_{add} (13) of the differential (12); adp^{ARX} refers to $\text{adp}^{\text{ARX}}((\Delta^+a + \Delta^+b), \Delta^+d \xrightarrow{t} \Delta^+e)$

Δ	Δ^+a	Δ^+b	Δ^+d	t	$\Delta^+e = \Delta$	adp^{ARX}
Δ_2^1	0	0	80000000	9	80000000	1
Δ_3^1	80000000	0	0	13	fffff000	2^{-1}
Δ_0^1	fffff000	80000000	0	18	40020000	$2^{-2.41}$
Δ_1^2	40020000	0	0	7	01000020	$2^{-2.99}$
Δ_8^2	0	0	80000000	9	80000000	1
Δ_9^2	80000000	0	0	13	fffff000	2^{-1}
Δ_9^3	0	01000020	fffff000	7	80000000	$2^{-2.58}$

$\hat{p}_{\text{add}} = 2^{-10}$

Table 3. The estimated probability \hat{p}_{unaf} (14) of the differential (12); udp^{ARX} refers to $\text{udp}^{\text{ARX}}(\Delta^Ua, \Delta^Ub, \Delta^Ud \xrightarrow{t} \Delta^Ue)$

Δ^U	Δ^Ua	Δ^Ub	Δ^Ud	t	$\Delta^Ue = \Delta^U$	udp^{ARX}
$\{\Delta^U\}_2^1$	0	0	80000000	9	80000000	1
$\{\Delta^U\}_3^1$	80000000	0	0	13	00001000	1
$\{\Delta^U\}_0^1$	00001000	80000000	0	18	40020000	$2^{-0.41}$
$\{\Delta^U\}_1^2$	40020000	0	0	7	01000020	$2^{-0.99}$
$\{\Delta^U\}_8^2$	0	0	80000000	9	80000000	1
$\{\Delta^U\}_9^2$	80000000	0	0	13	00001000	1
$\{\Delta^U\}_9^3$	0	01000020	00001000	7	80000000	$2^{-2.58}$

$\hat{p}_{\text{unaf}} = 2^{-4}$

the differential (12). The estimation \hat{p}_{add} is based upon a single one among all possible characteristics, while the estimation \hat{p}_{unaf} takes into account several characteristics at once. This effect is illustrated in Fig. 3. Note that the input $\{\Delta^U\}_8^0$ and output $\{\Delta^U\}_9^3$ UNAF sets contain a single element – the additive difference 80000000. Because of that $\{\Delta^U\}_8^0 = \Delta_8^0$ and $\{\Delta^U\}_9^3 = \Delta_9^3$ and therefore the estimations (13) and (14) can be compared to each other.

In the case where the output UNAF set contains more than one element (i.e. $\{\Delta^U\}_9^3 \neq \Delta_9^3$), we propose to divide the resulting probability by the size of the output UNAF set $\#\Delta^U$:

$$\hat{p}_{\text{unaf}} = \frac{\prod \text{udp}^{\text{ARX}}}{\#\Delta^U} . \tag{15}$$

The estimation (15) is based on the assumption that all additive differences from the output UNAF set Δ^U hold with approximately the same (or very close) probabilities. For the case of Salsa20, our experiments confirm this assumption.

We use (15) to estimate the probabilities with which several differences from the output state after Salsa20/3 hold, given input UNAF difference $\{\Delta^U\}_8^0 = 0x80000000$. The results are shown in Table 4 and in Fig. 4.




Fig. 3. A single UNAF characteristic, satisfying the differential $\Delta_8^0 \rightarrow \Delta_9^3$. It is composed of multiple additive characteristics.

The results presented in Table 4 and Fig. 4 show that although the probability estimations $\hat{p}_{\text{unaf}}/\#\Delta^U$ computed using UNAF differences with (15) deviate from the values obtained experimentally p_{exper} , they are still more accurate than the estimations \hat{p}_{add} based on single additive differences and computed with (13).

3.3 Key-Recovery Attack on Salsa20/5

In this section, we apply UNAF differences to mount a key-recovery attack on a version of stream cipher Salsa20 reduced to 5 rounds, denoted as Salsa20/5. Although its complexity is lower than exhaustive key search, the attack does not improve the best known attack on the cipher. Therefore it is described only as a demonstration of a practical application of UNAF differences.

Using the best-first search algorithm from Sect. 2.4 we find the following UNAF differential for 3 rounds of Salsa20:

$$\{\Delta^U\}_8^0 = 0x80000000 \rightarrow \{\Delta^U\}_{11}^3 = 0x01000024 . \quad (16)$$

Table 4. Estimating the probabilities of differentials for three rounds of Salsa20 using UNAF differences

i	Δ_i^3	$\{\Delta^U\}_i^3$	\hat{p}_{add}	$\hat{p}_{\text{unaf}}/\#\Delta^U$	p_{exper}
9	80000000	80000000	$2^{-10.00}$	$2^{-4.00}$	$2^{-3.38}$
13	ffe00100	00200100	$2^{-15.75}$	$2^{-7.75}$	$2^{-4.93}$
14	ff00001c	01000024	$2^{-16.29}$	$2^{-8.31}$	$2^{-6.35}$
1	00e00fe4	01201024	$2^{-23.01}$	$2^{-13.04}$	$2^{-10.18}$
2	00000800	00000800	$2^{-35.59}$	$2^{-16.62}$	$2^{-11.08}$
3	fff000a0	001000a0	$2^{-41.48}$	$2^{-20.04}$	$2^{-14.68}$
6	01038020	01048020	$2^{-41.76}$	$2^{-21.91}$	$2^{-15.68}$
7	ffefc000	00104000	$2^{-44.65}$	$2^{-22.15}$	$2^{-17.42}$

The input UNAF set $\{\Delta^U\}_8^0 = 0x80000000$ consists of one element: the additive difference $0x80000000$. The output UNAF set $\{\Delta^U\}_{11}^3 = 0x01000024$ contains the following 2^3 additive differences: $0x01000024$, $0x0100001c$, $0x00ffffe4$, $0x00ffffdc$, $0xff000024$, $0xff00001c$, $0xfeffffe4$, $0xfeffffdc$. The probability that an additive difference Δ_{11}^3 falls into the set $\{\Delta^U\}_{11}^3$ was determined experimentally to be $p_{\text{exper}} = 2^{-3.38}$.

In our attack, we first invert the feed-forward operation to compute the differences $\Delta_5^5, \Delta_6^5, \dots, \Delta_{10}^5$ of the state after round 5. Next, we guess 5 of the 8 words of the secret key, in order to compute the differences $\Delta_1^5, \Delta_2^5, \Delta_3^5, \Delta_4^5, \Delta_{11}^5$. Therefore, we do not only know the differences $\Delta_1^5, \Delta_2^5, \dots, \Delta_{11}^5$, but also the corresponding values of the word pairs. This allows us to compute the differences $\Delta_{12}^4, \Delta_{13}^4, \Delta_{14}^4$ from the state after round 4. Using the latter, we can finally compute the UNAF difference $\{\Delta^U\}_{11}^3$. If it is equal to $0x01000024$, then our guess of the key words was correct with some probability. This process is illustrated in Appendix D.

Since the probability of the differential (16) is $2^{-3.38} \geq 2^{-4}$, from $M = 2^6$ chosen plaintext pairs we expect that $2^{-4} \cdot 2^6 = 2^2 = 4$ pairs will follow the differential (i.e. will satisfy the output difference $\{\Delta^U\}_{11}^3$).

We assume that a pair encrypted under a wrong key results in a uniformly random difference. The probability that this difference falls into the set $\{\Delta^U\}_{11}^3$ is $P_{\text{rand}} = 2^3/2^{32} = 2^{-29}$. Therefore the probability that at least 4 plaintext pairs turn out to be all false positives (i.e. they satisfy the differential, but are encrypted under a wrong key) can be calculated using the binomial distribution:

$$\sum_{i=4}^{64} \binom{64}{i} (2^{-29})^i (1 - 2^{-29})^{64-i} \approx 2^{-96.72} . \quad (17)$$

As explained, because we guess 160 bits (5 words) of the secret key, in the attack we have to make 2^{160} guesses. For each guess, we encrypt 2^6 chosen plaintext pairs and we partially decrypt the resulting ciphertext pairs for 2 rounds in order to compute the output difference. From 2^{160} guesses, the expected number of wrong keys that result in at least 4 pairs with the right difference is $2^{-96.72} \cdot 2^{160} \approx$




Fig. 4. Three estimates of the probabilities of eight differentials for three rounds of Salsa20, based on the data from Table 4: (1) estimation obtained experimentally, (2) based on UNAF differences and (3) based on single additive differences.

2^{63} . For each of those keys, we guess the remaining 96 bits (3 words) i.e. we make 2^{96} guesses per candidate key. For each guess we encrypt one plaintext pair (i.e. two encryptions are performed) under the full key and check if the encryption matches the corresponding ciphertext pair. This results in $2 \cdot 2^{63} \cdot 2^{96} = 2^{160}$ additional operations. Thus we estimate the total number of encryptions of our attack to be:

$$2 \cdot 2^6 \cdot 2^{160} + 2 \cdot 2^{63} \cdot 2^{96} = 2^{167} + 2^{160} \approx 2^{167} . \tag{18}$$

Therefore the presented attack on Salsa20/5 has data complexity 2^7 chosen plaintexts and time complexity 2^{167} encryptions. As shown in Table 5, it is comparable to the attack proposed by Crowley [5].

Table 5. Overview of key-recovery attacks on Salsa20

Rounds	Reference	Time	Data	Type of Differences
Salsa20/5	Our result	2^{167}	2^7	Additive
Salsa20/5	Crowley [5]	2^{165}	2^6	XOR
Salsa20/6	Fischer et al. [10]	2^{177}	2^{16}	XOR
Salsa20/7	Aumasson et al. [1]	2^{151}	2^{26}	XOR
Salsa20/8	Aumasson et al. [1]	2^{251}	2^{31}	XOR

4 Conclusion

In this paper, we introduced UNAF differences. These are sets of specially chosen additive differences used to estimate the probabilities of differentials through sequences of ARX operations more accurately.

We presented the main UNAF theorem, which shows how a UNAF difference groups several possible additive differences together. Further, we investigated the propagation of UNAF differences through the ARX operation. We defined the UNAF differential probability of ARX and noted that it can be computed efficiently using the S-functions framework proposed by Mouha et al.

UNAF differences were applied to the cryptanalysis of the stream cipher Salsa20. We found that for three rounds of Salsa20, the probability of the best differential based on additive differences is estimated as 2^{-10} . Evaluating the same probability using UNAF differences leads to the value 2^{-4} . The latter is closer to the the probability of the differential $2^{-3.39}$ that was determined experimentally.

A general algorithm for the automatic search for differentials was briefly discussed. It was used to find high-probability UNAF differentials for three rounds of Salsa20. One of them was used to mount a key-recovery attack on Salsa20 reduced to five rounds. The attack has a time complexity of 2^7 and a data complexity of 2^{167} . It therefore does not improve the best-known attack on the cipher. Nevertheless, to the best of our knowledge, this is the first cryptanalysis result on Salsa20 that is based on additive differences. Furthermore, we expect that the attack can be extended to more rounds. One possibility in this direction is to group two or more ARX operations and consider them as a single operation. Another is to improve the method for finding differential characteristics for multiple rounds.

The results in this paper were obtained for the Salsa20 stream cipher. We see the application of UNAF differences to other ARX-based ciphers as another interesting topic for future research.

Acknowledgments. The authors would like to thank Florian Mendel for his detailed comments and for his suggestions on the overall structure of the paper. His feedback was critical for improving the quality of the final text. We thank prof. Vincent Rijmen for reviewing the draft version of the paper. We also thank our colleagues at COSIC for the useful discussions, as well as the anonymous reviewers for their detailed comments.

References

1. Aumasson, J.-P., Fischer, S., Khazaei, S., Meier, W., Rechberger, C.: New Features of Latin Dances: Analysis of Salsa, ChaCha, and Rumba. In: Nyberg, K. (ed.) FSE 2008. LNCS, vol. 5086, pp. 470–488. Springer, Heidelberg (2008)
2. Aumasson, J.-P., Henzen, L., Meier, W., Phan, R.C.-W.: SHA-3 proposal BLAKE. Submission to the NIST SHA-3 Competition (Round 2) (2008)

3. Bernstein, D.J.: The Salsa20 Family of Stream Ciphers. In: Robshaw, M., Billet, O. (eds.) *New Stream Cipher Designs*. LNCS, vol. 4986, pp. 84–97. Springer, Heidelberg (2008)
4. Biham, E., Shamir, A.: Differential Cryptanalysis of DES-like Cryptosystems. *J. Cryptology* 4(1), 3–72 (1991)
5. Crowley, P.: Truncated differential cryptanalysis of five rounds of Salsa20. In: *SASC 2006 Workshop: Stream Ciphers Revisited*. eSTREAM, ECRYPT Stream Cipher Project, Report 2005/073 (2005), <http://www.ecrypt.eu.org/stream>
6. Daemen, J., Rijmen, V.: *The Design of Rijndael: AES - The Advanced Encryption Standard*. Springer (2002)
7. Ebeid, N.M., Hasan, M.A.: On binary signed digit representations of integers. In: *Des. Codes Cryptography*, vol. 42(1), pp. 43–65 (2007)
8. eSTREAM. ECRYPT stream cipher project, <http://www.ecrypt.eu.org/stream>
9. Ferguson, N., Lucks, S., Schneier, B., Whiting, D., Bellare, M., Kohno, T., Callas, J., Walker, J.: *The Skein Hash Function Family*. Submission to the NIST SHA-3 Competition (Round 2) (2009)
10. Fischer, S., Meier, W., Berbain, C., Biase, J.-F., Robshaw, M.J.B.: Non-randomness in eSTREAM Candidates Salsa20 and TSC-4. In: Barua, R., Lange, T. (eds.) *INDOCRYPT 2006*. LNCS, vol. 4329, pp. 2–16. Springer, Heidelberg (2006)
11. Hart, P.E., Nilsson, N.J., Raphael, B.: A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics* 4(2), 100–107 (1968)
12. Lipmaa, H., Moriai, S.: Efficient Algorithms for Computing Differential Properties of Addition. In: Matsui, M. (ed.) *FSE 2001*. LNCS, vol. 2355, pp. 336–350. Springer, Heidelberg (2002)
13. Lipmaa, H., Wallén, J., Dumas, P.: On the Additive Differential Probability of Exclusive-Or. In: Roy, B., Meier, W. (eds.) *FSE 2004*. LNCS, vol. 3017, pp. 317–331. Springer, Heidelberg (2004)
14. Matsui, M., Yamagishi, A.: A New Method for Known Plaintext Attack of FEAL Cipher. In: Rueppel, R.A. (ed.) *EUROCRYPT 1992*. LNCS, vol. 658, pp. 81–91. Springer, Heidelberg (1993)
15. Mouha, N., Velichkov, V., De Cannière, C., Preneel, B.: The Differential Analysis of S-Functions. In: Biryukov, A., Gong, G., Stinson, D.R. (eds.) *SAC 2010*. LNCS, vol. 6544, pp. 36–56. Springer, Heidelberg (2011)
16. Reitwiesner, G.W.: Binary arithmetic. *Advances in Computers* 1, 231–308 (1960)
17. Shimizu, A., Miyaguchi, S.: Fast Data Encipherment Algorithm FEAL. In: Price, W.L., Chaum, D. (eds.) *EUROCRYPT 1987*. LNCS, vol. 304, pp. 267–278. Springer, Heidelberg (1988)
18. Velichkov, V., Mouha, N., De Cannière, C., Preneel, B.: The Additive Differential Probability of ARX. In: Joux, A. (ed.) *FSE 2011*. LNCS, vol. 6733, pp. 342–358. Springer, Heidelberg (2011)

A Proof of Theorem 1

The following Lemma provides the condition under which the probability adp^{\oplus} is non-zero.

Lemma 1 (Theorem 2 of [13]). *All differences Δ^+a , Δ^+b and Δ^+c for which $\text{adp}^\oplus(\Delta^+a, \Delta^+b \rightarrow \Delta^+c) > 0$, are $\Delta^+a = \Delta^+b = \Delta^+c = 0$, and*

$$\Delta^+a = \Delta^+a[n - 1 \dots q + 1] \parallel \Delta^+a[q] \parallel 0^* , \tag{19}$$

$$\Delta^+b = \Delta^+b[n - 1 \dots q + 1] \parallel \Delta^+b[q] \parallel 0^* , \tag{20}$$

$$\Delta^+c = \Delta^+c[n - 1 \dots q + 1] \parallel \Delta^+c[q] \parallel 0^* , \tag{21}$$

where $\neg(\Delta^+a[q] = \Delta^+b[q] = \Delta^+c[q] = 0)$ and $\Delta^+a[q] \oplus \Delta^+b[q] = \Delta^+c[q]$. Each of the sub-word differences $\Delta^+a[n - 1 \dots q + 1]$, $\Delta^+b[n - 1 \dots q + 1]$ and $\Delta^+c[n - 1 \dots q + 1]$ can take any arbitrary value. The symbol $*$ represents the Kleene star.

We proceed next with the proof of Theorem 1.

Proof. From Reitwiesner’s algorithm for the construction of the NAF [16], it follows that if the first non-zero bit (starting from the LSB) of Δ^+a_i is at position q , then the first non-zero bit of its NAF representation $\Delta^N a_i$ is also at position q . Since all Δ^+a_i in (5) belong to the same UNAF set $\Delta^U a$, the first non-zero bit for all of them is in the same position q . The same observation holds for Δ^+b_j and Δ^+c_k . From $\text{adp}^\oplus(\Delta^+a, \Delta^+b \rightarrow \Delta^+c) > 0$ and Lemma 1, it follows that $\Delta^+a[q] \oplus \Delta^+b[q] = \Delta^+c[q]$. Therefore $\Delta^+a_i[q] \oplus \Delta^+b_j[q] = \Delta^+c_k[q], \forall i, j, k$. Again by Lemma 1, it follows that if Δ^+a is replaced by any Δ^+a_i belonging to the same UNAF set $\Delta^U a$, the resulting probability adp^\oplus is still non-zero. The same observation can be made for Δ^+b and Δ^+c , which completes the proof. \square

B Computation of udp^{ARX}

The probability udp^{ARX} can be efficiently computed using the S-function framework [15,18]. We briefly describe this computation below. It is also a part of a toolkit that will be made publicly available.

The propagation of input UNAF differences $\Delta^U a$, $\Delta^U b$ and $\Delta^U d$ to output UNAF difference $\Delta^U e$ is represented as an S-function. The latter is used to compute 16 adjacency matrices. Each of them corresponds to a given value of the i -th bit of each of the four UNAF differences and connects a set of possible input states to a set of possible output states.

The differential $(\Delta^U a[i], \Delta^U b[i], \Delta^U d[i + t] \xrightarrow{t} \Delta^U e[i + t])$ at bit position i is written as the bit string $w[i] \leftarrow (\Delta^U a[i] \parallel \Delta^U b[i] \parallel \Delta^U d[i + t] \parallel \Delta^U e[i + t])$. At each bit position $0 \leq i < n$, the index $w[i] \in \{0, \dots, 15\}$ selects one of the 16 adjacency matrices $A_{w[i]}$. The probability udp^{ARX} is computed as follows:

$$\begin{aligned} \text{udp}^{\text{ARX}}(\Delta^U a, \Delta^U b, \Delta^U d \xrightarrow{t} \Delta^U e) = \\ \sum_{j=0}^{14} L_j \left(\prod_{i=n-t}^{n-1} A_{w[i]} \right) R \left(\prod_{i=0}^{n-t-1} A_{w[i]} \right) C_j . \end{aligned} \tag{22}$$

In (22), the summation is performed over each of the 14 possible initial states. The reason for having multiple initial states is the bit rotation by t positions, as

explained in [18]. The multiplication by the projection matrix R at bit position t is necessary because of the rotation operation. The column vectors C_j , $0 \leq j < 15$ represent the 15 possible initial states. The row vectors L_j , $0 < j < 15$ represent their corresponding final states. For further details, we refer to [18].

Note that the matrices $A_{w[i]}$ are of dimension 540×540 , but these can be minimized to 60×60 by combining equivalent states using the algorithm of [15, §3.5].

C An Algorithm for Finding the Best Output Difference

Let \square be an operation that takes a finite number of n -bit input words a_1, b_1, d_1, \dots and computes an n -bit output word $c_1 = \square(a_1, b_1, d_1, \dots)$. Let \bullet be a type of difference. Let $\alpha, \beta, \zeta, \dots$ and γ be differences of type \bullet such that $a_1 \bullet a_2 = \alpha$, $b_1 \bullet b_2 = \beta$, $d_1 \bullet d_2 = \zeta$, \dots and $c_1 \bullet c_2 = \gamma$ for some a_2, b_2, d_2, \dots and some c_2 . The differential probability with which input differences $\alpha, \beta, \zeta, \dots$ propagate to output difference γ with respect to the operation \square is denoted as $\bullet \text{dp}^\square(\alpha, \beta, \zeta, \dots \rightarrow \gamma)$. Finally, let the difference \bullet be such that it is possible to express its propagation through the operation \square as an S-function consisting of N states. Therefore, there exist adjacency matrices $A_{w[i]}$ such that the probability $\bullet \text{dp}^\square$ can be efficiently computed as $LA_{w[n-1]} \dots A_{w[1]}A_{w[0]}C$, where $L = [1 \ 1 \ \dots \ 1]$ is a $1 \times N$ matrix and $C = [1 \ 0 \ \dots \ 0]^T$ is an $N \times 1$ matrix (as in [15]). The problem is to find an output difference γ such that its probability p_γ over all possible output differences is maximal:

$$p_\gamma = \bullet \text{dp}^\square(\alpha, \beta, \zeta, \dots \rightarrow \gamma) = \max_j \bullet \text{dp}^\square(\alpha, \beta, \zeta, \dots \rightarrow \gamma_j) . \quad (23)$$

We represent (23) as a problem of finding the shortest path in an *node-weighted binary tree*. We define the binary tree $T = (N, E)$, where N is the set of nodes and E is the set of edges. The height of T is $n + 1$ with a dummy start node positioned at level -1 and the leaves positioned at level $n - 1$. Each node at level i : $0 \leq i < n$ contains a value of $\gamma[i]$, where $i = 0$ is the LSB and $i = n - 1$ is the MSB. Every node on level i has two children at level $i + 1$. Since the input differences $\alpha, \beta, \zeta, \dots$ are fixed, at every bit position i we can choose between two matrices $A_{w[i]}$, corresponding to the two possibilities for the output difference $\gamma[i]$.

To find the output difference with the highest probability, we use the A* search algorithm [11]. In this algorithm, an evaluation function f can be computed for every node in the search tree. The f -function represents the weight of a node, and is based on the cost of the path from the start node, and a heuristic that estimates the distance to the goal node. The algorithm always expands the node with the highest f -value (corresponding to the highest probability). The A* search algorithm guarantees that the optimal solution will be found, provided that the evaluation function f never underestimates the probability of the best output difference. After introducing some definitions, we will define an evaluation function f and prove in Theorem 2 that this f satisfies the required condition.

Let vector $X_i = [x_{i,0} \ x_{i,1} \ \cdots \ x_{i,N-1}]$ be a transition probability vector, i.e. $x_{i,r} \geq 0$ for $0 \leq r < N$ and $\sum_{r=0}^{N-1} x_{i,r} \leq 1$. We define H_r as a column vector of length N , of which the r -th element (counting from 0) is 1 and all other elements are 0. The cost of a node at level i is then denoted by $\|X_i\|$ (the 1-norm of X_i) and is calculated as $\|A_{w[i]}A_{w[i-1]} \cdots A_{w[0]}C\|$. Let us define a sequence of row vectors $\hat{G}_{i,r}$, $0 \leq r < N$ and $0 \leq i < n$. Each $\hat{G}_{i,r}$ is a product of matrices $LA_{w[n-1]}A_{w[n-2]} \cdots A_{w[i+1]}$, where each of the A -matrices are chosen such that $\hat{G}_{i,r}H_r$ is maximized. The choice of the A -matrices may differ for different values of r . We define row vector G_i as the product of matrices $LA_{w[n-1]}A_{w[n-2]} \cdots A_{w[i+1]}$, where the A -matrices are chosen such that G_iX_i is maximized. For a node at level i with cost $\|X_i\|$, the evaluation function f is defined as $\sum_{r=0}^{N-1} \hat{G}_{i,r}H_r x_{i,r}$.

Theorem 2. *The evaluation function $f = \sum_{r=0}^{N-1} \hat{G}_{i,r}H_r x_{i,r}$ never underestimates the probability of the best output difference.*

Proof. The following inequality holds: $\hat{G}_{i,r}H_r \geq G_iH_r$ for $0 \leq r < N$. The latter can be proven by contradiction: if $\hat{G}_{i,r}H_r < G_iH_r$ for some r , then $\hat{G}_{i,r}$ is not the product of A -matrices that maximizes $\hat{G}_{i,r}H_r$, which contradicts its definition. Because probabilities are non-negative, we can multiply both sides of the inequality by the state probability $x_{i,r}$, to obtain $\hat{G}_{i,r}H_r x_{i,r} \geq G_iH_r x_{i,r}$, $0 \leq r < N$. By summing the left and the right sides of the N inequalities, we obtain $\sum_{r=0}^{N-1} \hat{G}_{i,r}H_r x_{i,r} \geq \sum_{r=0}^{N-1} G_iH_r x_{i,r} = G_iX_i$. By definition, G_iX_i is the best choice of A -matrices, starting from transition probability X_i . This proves that the left-hand side of the inequality never underestimates the probability, which proves the theorem. \square

Before we can apply the A^* algorithm to compute the best output difference, we must determine the values of $\hat{G}_{i,r}H_r$ for $0 \leq i < n$ and $0 \leq r < N$. This is done by again running the A^* algorithm for the most significant bit, then for the two most significant bits, and so on until we process the entire word. For the MSB, we define $\hat{G}_{n-1,r} = L$ for $0 \leq r < N$. For the two MSBs, we run the A^* algorithm for every $0 \leq r < N$, setting the transition probability vector X_{n-2} to H_r . This allows us to compute $\hat{G}_{n-2,r}H_r$. This process is continued until $\hat{G}_{0,r}H_r$ for $0 \leq r < N$ is calculated. Having calculated all values of $\hat{G}_{i,r}H_r$, we then use the A^* algorithm to search for the best output difference by setting the state transition probability vector $X_{-1} = C$. Pseudo-code of the entire A^* search algorithm is provided in Algorithm 1.

D Attack on Salsa20/5 Using UNAF Differences

Fig. 5 illustrates the attack presented in Sect. 3.3. Gray boxes denote guessed words and white boxes denote words that are either known or can be computed.

Algorithm 1. Find the Best Output Diff. of Type \bullet w.r.t. Operation \square .

Input: Matrices $A_{w[i]}$ for $\bullet\text{dp}^\square$; input diffs. $\alpha, \beta, \zeta, \dots$; num. states N .

Output: Output difference γ and probability p_γ such that

$$p_\gamma = \bullet\text{dp}^\square(\alpha, \beta, \zeta, \dots \rightarrow \gamma) = \max_j \bullet\text{dp}^\square(\alpha, \beta, \zeta, \dots \rightarrow \gamma_j) .$$

```

1: Define struct node = {index,  $\gamma$ ,  $f_{\text{index}-1}$ ,  $\hat{H}_{\text{index}-1}$ }
2: Init priority queue of nodes ordered by  $f$ :  $Q = \emptyset$ 
3: Init output difference:  $\gamma \leftarrow \emptyset$ 
4: for  $i = n - 1$  downto 0 do
5:   if  $i = n - 1$  then
6:      $\hat{G}_i \leftarrow L = [1 \ 1 \ \dots \ 1]$ 
7:   else
8:      $\hat{G}_i \leftarrow [\hat{G}_{i,0} \ \hat{G}_{i,1} \ \dots \ \hat{G}_{i,N-1}]$ 
9:   end if
10:  if  $i = 0$  then
11:     $N = 1$ 
12:  end if
13:  for  $r = 0$  to  $N - 1$  do
14:    Reset priority queue:  $Q = \emptyset$ 
15:    Init the total probability of node  $v_{i-1}$ :  $f_{i-1} \leftarrow 1$ 
16:    Init the transition probability vector  $v_i$ :  $\hat{H}_{i-1} \leftarrow \hat{H}_{i-1,r}$ 
17:    Init node  $v_i \leftarrow \{i, \gamma, f_{i-1}, \hat{H}_{i-1}\}$ 
18:    Add new node to the queue:  $Q.\text{push}(v_i)$ 
19:     $v_{\text{best}} \leftarrow Q.\text{top}()$ ;  $\{j, \gamma, f_{j-1}, \hat{H}_{j-1}\} \leftarrow v_{\text{best}}$ 
20:    while  $j \neq n$  do
21:      Remove  $v_{\text{best}}$  from the queue:  $Q.\text{pop}()$ 
22:      for  $q = 0$  to 1 do
23:        Set the  $j$ -th bit of  $\gamma$ :  $\gamma[j] \leftarrow q$ 
24:        Estimate the total probability:  $f_j \leftarrow \hat{G}_j A_{w[j]}^q \hat{H}_{j-1}$ 
25:        Compute the transition probability vector:  $\hat{H}_j \leftarrow A_{w[j]}^q \hat{H}_{j-1}$ 
26:        Init child of  $v_{\text{best}}$ : node  $v_{j+1}^q \leftarrow \{j + 1, \gamma, f_j, \hat{H}_j\}$ 
27:        Add the child to the queue:  $Q.\text{push}(v_{j+1}^q)$ 
28:      end for
29:      Extract the node with the lowest total cost:  $v_{\text{best}} \leftarrow Q.\text{top}()$ 
30:       $\{j, \gamma, f_{j-1}, \hat{H}_{j-1}\} \leftarrow v_{\text{best}}$ 
31:    end while
32:     $v_{\text{best}} \leftarrow Q.\text{top}()$ ;  $f_{\text{best}} \leftarrow \text{get\_cost}(v_{\text{best}})$ 
33:    Set the  $r$ -th element of  $\hat{G}_i$ :  $\hat{G}_{i,r} \leftarrow f_{\text{best}}$ 
34:  end for
35: end for
36: Extract the node with highest total probability:  $v_{\text{best}} \leftarrow Q.\text{top}()$ 
37: Get the output difference associated to  $v_{\text{best}}$ :  $\gamma, p_\gamma \leftarrow \text{get\_gamma}(v_{\text{best}})$ 
38: return  $\gamma, p_\gamma$ 

```




Fig. 5. Key-recovery attack on Salsa20/5 using the 3-round UNAF differential $\{\Delta^U\}_8^0 \rightarrow \{\Delta^U\}_{11}^3$. Gray boxes denote guessed words; white boxes denote words that are either known or can be computed.

ElimLin Algorithm Revisited

Nicolas T. Courtois¹, Pouyan Sepehrdad^{2,*},
Petr Sušil^{2,**}, and Serge Vaudenay²

¹ University College London, UK

`n.courtois@ucl.ac.uk`

² EPFL, Lausanne, Switzerland

`{pouyan.sepehrdad, petr.susil, serge.vaudenay}@epfl.ch`

Abstract. ElimLin is a simple algorithm for solving polynomial systems of multivariate equations over small finite fields. It was initially proposed as a single tool by Courtois to attack DES. It can reveal some hidden linear equations existing in the ideal generated by the system. We report a number of key theorems on ElimLin. Our main result is to characterize ElimLin in terms of a sequence of intersections of vector spaces. It implies that the linear space generated by ElimLin is invariant with respect to any variable ordering during elimination and substitution. This can be seen as surprising given the fact that it eliminates variables. On the contrary, monomial ordering is a crucial factor in Gröbner basis algorithms such as F4. Moreover, we prove that the result of ElimLin is invariant with respect to any affine bijective variable change. Analyzing an overdefined dense system of equations, we argue that to obtain more linear equations in the succeeding iteration in ElimLin some restrictions should be satisfied. Finally, we compare the security of LBlock and MIBS block ciphers with respect to algebraic attacks and propose several attacks on Courtois Toy Cipher version 2 (CTC2) with distinct parameters using ElimLin.

Keywords: block ciphers, algebraic cryptanalysis, systems of sparse polynomial equations of low degree.

*[Breaking a good cipher should require]
“as much work as solving a system of simultaneous equations
in a large number of unknowns of a complex type.”*

Claude Elwood Shannon [45]

1 Introduction

Various techniques exist in cryptanalysis of symmetric ciphers. Some involve statistical analysis and some are purely deterministic. One of the latter methods is *algebraic attack* recognized as early as 1949 by Shannon [45]. Any algebraic attack consists of two distinct stages:

* This work has been supported in part by the European Commission through the ICT program under contract ICT-2007-216646 ECRYPT II.

** Supported by a grant of the Swiss National Science Foundation, 200021_134860/1.

- Writing the cipher as a system of polynomial equations of low degree often over $\text{GF}(2)$ or $\text{GF}(2^k)$, which is feasible for any cipher [48,20,42].
- Recovering the secret key by solving such a large system of polynomial equations.

Algebraic attacks have been successful in breaking several stream ciphers (see [1,18,11,24,19,14,23,10] for instance) and a few block ciphers such as Keeloq [37] and GOST [15], but they are not often as successful as statistical attacks. On the other hand, they often require low data complexity. This is not the case for statistical attacks.

General purpose algebraic attack techniques were developed in the last few years by Courtois, Bard, Meier, Faugère, Raddum, Semaev, Vielhaber, Dinur and Shamir to solve these systems [16,21,20,18,11,30,31,44,47,23,24]. The problem of solving such polynomial systems of multivariate equations is called MQ problem and is known to be NP hard for a random system. Currently, for a random system in which the number of equations is equal to the number of unknowns, there exists no technique faster than an exhaustive key search which can solve such systems. On the other hand, the equations derived from symmetric ciphers turn out to be overdefined and sparse for most ciphers. So, they might be easier to solve. This sparsity is coming from the fact that due to the limitations in hardware and the need for lightweight algorithms, simple operations arise in the definition of cryptosystems. They are also overdefined due to the non-linear operations.

The traditional method for solving overdefined polynomial systems of equations are known to be various Gröbner basis algorithms such as Buchberger algorithm [9], F4 and F5 [30,31] and XL [21]. The most critical drawback of the Gröbner basis approach is the elimination step where the degree of the system increases. This leads to an explosion in memory space and even the most current efficient implementations of Faugère algorithm [30,31] under PolyBoRi framework [8] or Magma [40] are not capable of handling *large* systems of equations efficiently. On the other hand, they are faster than other methods for overdefined dense systems or when the equations are over $\text{GF}(q)$ where $q > 2$. In fact, together with SAT solvers, they are currently the most successful methods for solving polynomial systems.

Nevertheless, due to the technical reasons mentioned above, the system of equations extracted from symmetric ciphers turns out to be sparse. Unfortunately, the Gröbner basis algorithms can not exploit this property. In such cases, algorithms such as XSL [20], SAT solving techniques [4,27,3], Raddum-Semaev algorithm [44] and ElimLin [16] are of interest.

In this paper, we study the elimination algorithm ElimLin that falls within the remit of Gröbner basis algorithms, though it is conceptually much simpler and is based on a mix of simple linear algebra and substitution. It maintains the degree of the equations and it does not require any fixed ordering on the set of all monomials. On the contrary, we need to work with ad-hoc monomial orderings to preserve the sparsity and make it run faster. This simple algorithm reveals

some hidden linear equations existing in the ideal generated by the system. We show in Sec. 7 that ElimLin does not find all such linear equations.

As far as the authors are aware, no clue has been found yet which demonstrates that ElimLin at some stage stops working. This does not mean that ElimLin can break any system. As mentioned earlier, for a random system this problem is NP hard and Gröbner basis algorithms behave much better for such dense random systems. But, the equations derived from cryptosystems are often not random (see [32] for the huge difference between a random system and the algebraic representation of cryptographic protocols). What we mean here is that if for some small number of rounds ElimLin performs well but then it stops working for more rounds, we can increase the number of samples and it will become effective again. The bottleneck is having an efficient data structure for implementing ElimLin together with a rigorous theory behind it to anticipate its behaviour. These two factors are currently missing in the literature.

Except two simple theorems by Bard (see Chapter 12, Section 5 of [4]), almost nothing has been done regarding the theory behind ElimLin. As ElimLin can also be used as a pre-processing step in any algebraic attack, building a proper theory is vital for improving the state of the art algebraic attacks. We are going to shed some lights on the way this ad-hoc algorithm works and the theory behind it.

In this paper, we show that the output of ElimLin is invariant with respect to any variable ordering. This is a surprising result, i.e., while the spaces generated are different depending on how substitution is performed, we prove that their intersection is exactly the same. Furthermore, we prove that no affine bijective variable change can modify the output of ElimLin. Then, we prove a theorem on how the number of linear equations evolves in each iteration of ElimLin.

An unannounced competition is currently running for designing lightweight cryptographic primitives. This includes several designs which have appeared in the last few years (see [7,22,39,34,29,36,46,2,35,6]). These designs mainly compete over the gate equivalent (GE) and throughput. This might not be a fair comparison of efficiency, since they do not provide the same level of security with respect to distinct types of attacks. In this paper, we compare the two lightweight Feistel-based block ciphers MIBS [38] and LBlock [49] and show that with the same number of rounds, LBlock provides a much lower level of security compared to MIBS with respect to algebraic attacks. In fact, we attack both ciphers with ElimLin and F4 algorithm. Finally, we provide several algebraic attacks against Courtois Toy Cipher version 2 (CTC2) with distinct parameters using ElimLin.

In Sec. 2, we elaborate the ElimLin algorithm. Then, we remind some basic theorems on ElimLin in Sec. 3. As our main contribution (Theorem 7), we prove in Sec. 4 that ElimLin can be formulated as an intersection of vector spaces. We also discuss its consequences in Sec. 4.2 and prove a theorem regarding the evolution of linear equations in Sec. 4.3. We perform some attacks simulations on CTC2, LBlock and MIBS block ciphers in Sec. 5.2, 5.3 and 5.4 respectively. In Sec. 6, we compare ElimLin and F4. We mention some open problems and a conjecture in Sec. 7 and we conclude.

2 ElimLin Algorithm

ElimLin stands for **E**liminate **L**inear and it is a technique for solving polynomial systems of multivariate equations of low degree d mostly: 2, 3, or 4 over a finite field specifically GF(2). It is also known as “inter-reduction” step in all major algebra systems. As a single tool, it was proposed in [16] to attack DES. It broke 5-round DES. Later, it was applied to break 5-round PRESENT block cipher [43] and to analyze the resistance of Snow 2.0 stream cipher against algebraic attacks [17]. It is a simple but a powerful algorithm which can be applied to any symmetric cipher and is capable of breaking their reduced versions. There is no specific requirement for the system except that there should exist at least one linear term, otherwise ElimLin trivially fails. The key question for such an algorithm is to predict its behavior. Currently, very similar to most other types of algebraic attacks such as [47,23,24], multiple parts of the algorithm are heuristic, so it is worthwhile to prove which factors can improve its results, make it run faster or does not have any influence on its ultimate result. This will yield a better understanding of how ElimLin works.

ElimLin is composed of two sequential distinct stages, namely:

- *Gaussian Elimination*: All the linear equations in the linear span of initial equations are found. They are the intersection between two vector spaces: The vector space spanned by all monomials of degree 1 and the vector space spanned by all equations.
- *Substitution*: Variables are iteratively eliminated in the whole system based on linear equations until there is no linear equation left. Consequently, the remaining system has fewer variables.

This routine is iterated until no linear equation is obtained in the linear span of the system. See Fig. 1 for a more precise definition of the algorithm. Clearly, the algorithm shall depend on ordering strategies to apply in step 5, 11, and 12 of Fig. 1. We will see that it is not, i.e., the span of the resulting \mathcal{S}_L is invariant.

We observe that new linear equations are derived in each iteration of the algorithm that did not exist in the former spans. This phenomenon is called *avalanche effect* in ElimLin and is the consequence of Theorem 7. At the end, the system is solved linearly (when \mathcal{S}_L is large enough) or ElimLin fails. If the latter occurs, we can increase the data complexity¹ and re-run the attack.

3 State of the Art Theorems

The only theoretical analysis of ElimLin was done by Bard in [4]. He proved the following theorem and corollary for **one** iteration of ElimLin:

Theorem 1 ([4]). *All linear equations in the linear span of a polynomial equation system \mathcal{S}^0 are found in the linear span of linear equations derived by performing the first iteration of ElimLin algorithm on the system.*

¹ For instance, the number of plaintext-ciphertext pairs.

```

1: Input : A system of polynomial equations  $\mathcal{S}^0 = \{\text{Eq}_1^0, \dots, \text{Eq}_{m_0}^0\}$  over  $\text{GF}(2)$ .
2: Output : An updated system of equations  $\mathcal{S}^T$  and a system of linear equations  $\mathcal{S}_L$ .
3: Set  $\mathcal{S}_L \leftarrow \emptyset$  and  $\mathcal{S}^T \leftarrow \mathcal{S}^0$  and  $k \leftarrow 1$ .
4: repeat
5:   Perform Gaussian elimination  $\text{Gauss}(\cdot)$  on  $\mathcal{S}^T$  with an arbitrary ordering of equations and monomials to eliminate non-linear monomials.
6:   Set  $\mathcal{S}_{L'} \leftarrow$  Linear equations from  $\text{Gauss}(\mathcal{S}^T)$ .
7:   Set  $\mathcal{S}^T \leftarrow \text{Gauss}(\mathcal{S}^T) \setminus \mathcal{S}_{L'}$ .
8:   Set flag.
9:   for all  $\ell \in \mathcal{S}_{L'}$  in an arbitrary order do
10:    if  $\ell$  is a trivial equation then
11:      if  $\ell$  is unsolvable then
12:        Terminate and output “No Solution”.
13:      end if
14:    else
15:      Unset flag.
16:      Let  $x_{t_k}$  be a monomial from  $\ell$ .
17:      Substitute  $x_{t_k}$  in  $\mathcal{S}^T$  and  $\mathcal{S}'_L$  using  $\ell$ .
18:      Insert  $\ell$  in  $\mathcal{S}_L$ .
19:       $k \leftarrow k + 1$ 
20:    end if
21:  end for
22: until flag is set.
23: Output  $\mathcal{S}_T$  and  $\mathcal{S}_L$ .

```

Fig. 1. ElimLin algorithm

The following corollary (also from [4]) is the direct consequence of the above theorem.

Corollary 2. *The linear equations generated after performing the first Gaussian elimination in ElimLin algorithm form a basis for all possible linear equations in the linear span of the system.*

This shows that any method to perform Gaussian elimination does not affect the linear space obtained at an arbitrary iteration of ElimLin. All linear equations derived from one method exist in the linear span of the equations cumulated from another method. This is trivial to see.

4 Algebraic Representation of ElimLin

4.1 ElimLin as an Intersection of Vector Spaces

We also formalize ElimLin in an algebraic way. This representation is used in proving Theorem 7. First, we define some notations.

We call an *iteration* a Gaussian elimination preceding a substitution; The system of equations for ElimLin can be stored as a matrix \mathcal{M}_α of dimension

$m_\alpha \times T_\alpha$, where each m_α rows represents an equation and each T_α columns represents a monomial at iteration α . Also, r_α denotes the rank of M_α . Let n_α be the number of variables at iteration α . We use a reverse lexicographical ordering of columns during Gaussian elimination to accumulate linear equations in the last rows of the matrix. Any arbitrary ordering can be used instead. In fact, we use the same matrix representation as described in [4].

Let $K = \text{GF}(2)$ and $x = (x_1, \dots, x_n)$ be a set of free variables. We denote by $K[x]$ the ring of multivariate polynomials over K . For $\mathcal{S} \subset K[x]$, we denote $\text{Span}(\mathcal{S})$ the K -vector subspace of $K[x]$ spanned by \mathcal{S} . Let $\gamma = (\gamma_1, \gamma_2, \dots, \gamma_n)$ be a power vector in \mathbf{N}^n . The term x^γ is defined as the product $x^\gamma = x_1^{\gamma_1} \times x_2^{\gamma_2} \times \dots \times x_n^{\gamma_n}$. The total degree of x^γ is defined as $\text{deg}(x^\gamma) \stackrel{\text{def}}{=} \gamma_1 + \gamma_2 + \dots + \gamma_n$. Let $\text{Ideal}(\mathcal{S})$ be the ideal spanned by \mathcal{S} and $\text{Root}(\mathcal{S})$ be the set of all tuples $m \in K^n$ such that $f(m) = 0$ for all $f \in \mathcal{S}$. Let

$$R_d = \text{Span}(\text{monomials of degree } \leq d) / \text{Ideal}(x_1^2 - x_1, x_2^2 - x_2, \dots, x_n^2 - x_n)$$

Let \mathcal{S}^α be \mathcal{S}^T after the α -th iteration of ElimLin and \mathcal{S}^0 be the initial system. Moreover, n_L^α is the number of non-trivial linear equations in \mathcal{S}_L^α at the α -th iteration. We denote \mathcal{S}_L^α the \mathcal{S}_L after the α -th iteration. Also, $C^\alpha \stackrel{\text{def}}{=} \#\mathcal{S}_L^\alpha$.

Let assume that \mathcal{S}^0 has degree bounded by d . We denote by $\text{Var}(f)$ the set of variables x_i expressed in f . Let x_{t_1}, \dots, x_{t_k} be the sequence of eliminated variables. We define $\mathbf{V}_k = \{x_1, \dots, x_n\} \setminus \{x_{t_1}, \dots, x_{t_k}\}$. Also, let $\ell_1, \ell_2, \dots, \ell_k$ be the sequence of linear equations as they are used during elimination (step 11 of Fig. 1). Hence, we have $x_{t_k} \in \text{Var}(\ell_k) \subseteq \mathbf{V}_{k-1}$.

We prove the following crucial lemma which we use later to prove Theorem 7.

Lemma 3. *After the α -th iteration of ElimLin, an arbitrary equation Eq_i^α in the system $(\mathcal{S}^\alpha \cup \mathcal{S}_L^\alpha)$ for an arbitrary i can be represented as*

$$\text{Eq}_i^\alpha = \sum_{t=1}^{m_0} \beta_{ti}^\alpha \cdot \text{Eq}_t^0 + \sum_{t=1}^{C^\alpha} \ell_t(x) \cdot g_{ti}^\alpha(x) \tag{1}$$

where $\beta_{ti}^\alpha \in K$ and $g_{ti}^\alpha(x)$ is a polynomial in R_{d-1} and $\text{Var}(g_{ti}^\alpha) \subseteq \mathbf{V}_t$.

Proof. Let x_{t_1} be one of the monomials existing in the first linear equation $\ell_1(x)$ and this specific variable is going to be eliminated. Substituting x_{t_1} in an equation $x_{t_1} \cdot h(x) + z(x)$, where $h(x)$ has degree at most $d - 1$, $x_{t_1} \notin \text{Var}(h)$ and $x_{t_1} \notin \text{Var}(z)$ is identical to subtracting $h(x) \cdot \ell_1(x)$. Consequently, the proof follows by induction on α . □

Now, we prove the inverse of the above lemma.

Lemma 4. *For each i and each α , there exists $\beta'_{ti}^\alpha \in K$ and $g'_{ti}^\alpha(x)$ such that*

$$\text{Eq}_i^0 = \sum_{t=1}^{m_\alpha} \beta'_{ti}^\alpha \cdot \text{Eq}_t^\alpha + \sum_{t=1}^{C^\alpha} \ell_t(x) \cdot g'_{ti}^\alpha(x) \tag{2}$$

where $g'_{ti}^\alpha(x)$ is a polynomial in R_{d-1} and $\text{Var}(g'_{ti}^\alpha) \subseteq \mathbf{V}_t$.

Proof. Gaussian elimination and substitution are invertible operations. We can use a similar induction as the previous lemma to prove the above equation. \square

In the next lemma, we prove that \mathcal{S}_L^α contains all linear equations which can be written in the form of Eq. (1).

Lemma 5. *If there exists $\ell \in R_1$ and some β_t and $g_t''(x)$ such that*

$$\ell(x) = \sum_{t=1}^{m_0} \beta_t \cdot \text{Eq}_t^0 + \sum_{t=1}^{C^\alpha} \ell_t(x) \cdot g_t''(x) \tag{3}$$

at iteration α , where $g_t''(x)$ is a polynomial in R_{d-1} , then there exists $u_t \in K$ and $v_t \in K$ such that

$$\ell(x) + \sum_{t=1}^{C^\alpha} u_t \cdot \ell_t(x) = \sum_{t=1}^{m_\alpha} v_t \cdot \text{Eq}_t^\alpha$$

So, $\ell(x) \in \text{Span}(\mathcal{S}_L^\alpha)$.

Proof. We define u_k iteratively: u_k is the coefficient of x_{t_k} in

$$\ell(x) + \sum_{t=1}^{k-1} u_t \cdot \ell_t(x)$$

for $k = 1, \dots, C^\alpha$. So, $\text{Var}(\ell(x) + \sum_{t=1}^k u_t \cdot \ell_t(x)) \subseteq \mathbf{V}_k$. By substituting Eq_i^0 from Eq. (2) in Eq. (3) and integrating u_t and g_t'' in g_{ti}^α , we obtain

$$\underbrace{\ell(x) + \sum_{t=1}^{C^\alpha} u_t \cdot \ell_t(x)}_{\subseteq \mathbf{V}_1} = \underbrace{\sum_{t=1}^{m_\alpha} v_t \cdot \text{Eq}_t^\alpha}_{\subseteq \mathbf{V}_1} + \underbrace{\sum_{t=1}^{C^\alpha} \ell_t(x) \cdot g_t'(x)}_{\Rightarrow \subseteq \mathbf{V}_1} \tag{4}$$

with $g_t'(x) \in R_{d-1}$. All $g_t'(x)$ where $t > 1$ can be written as $\bar{g}_t(x) + x_{t_1} \cdot \bar{\bar{g}}_t(x)$ with $\text{Var}(\bar{g}_t) \subseteq \mathbf{V}_1$, $\text{Var}(\bar{\bar{g}}_t) \subseteq \mathbf{V}_1$ and $\bar{\bar{g}}_t(x) \in R_{d-2}$. Since,

$$\begin{aligned} \ell_1(x) \cdot g_1'(x) + \ell_t(x) \cdot g_t'(x) &= \ell_1(x) \cdot \underbrace{(g_1'(x) + \ell_t(x) \cdot \bar{\bar{g}}_t(x))}_{\text{new } g_1'(x)} \\ &\quad + \underbrace{\ell_t(x)}_{\subseteq \mathbf{V}_1} \cdot \underbrace{(\bar{g}_t(x) + \bar{\bar{g}}_t(x) \cdot (x_{t_1} - \ell_1(x)))}_{(\text{new } g_t'(x)) \subseteq \mathbf{V}_1} \end{aligned}$$

we can re-arrange the sum in Eq. (4) using the above representation and obtain $\text{Var}(g_t') \subseteq \mathbf{V}_1$ for all $t > 1$. Also, x_{t_1} only appears in $\ell_1(x)$ and $g_1'(x)$. So, the coefficient of x_{t_1} in the expansion of $\ell_1(x) \cdot g_1'(x)$ must be zero. In fact, we have

$$\begin{aligned} \ell_1(x) \cdot g_1'(x) &= (x_{t_1} + (\ell_1(x) - x_{t_1})) \cdot (\bar{g}_1(x) + x_{t_1} \cdot \bar{\bar{g}}_1(x)) \\ &= x_{t_1} \cdot (\bar{\bar{g}}_1(x) \cdot (1 + \ell_1(x) - x_{t_1}) + \bar{g}_1(x)) + \bar{g}_1(x) \cdot (\ell_1(x) - x_{t_1}) \end{aligned}$$

So, $\bar{g}_1(x) = \bar{g}_1(x) \cdot (x_{t_1} - \ell_1(x) - 1)$ and we deduce,

$$g'_1(x) = \bar{g}_1(x) \cdot (\ell_1(x) + 1)$$

over $\text{GF}(2)$. But, then $\ell_1(x) \cdot g'_1(x) = 0$ over R , since $\ell_1(x) \cdot (\ell_1(x) + 1) = 0$. Finally, we iterate and obtain

$$\ell(x) + \sum_{t=1}^{C^\alpha} u_t \cdot \ell_t(x) = \sum_{t=1}^{m_\alpha} v_t \cdot \text{Eq}_t^\alpha$$

□

From another perspective, ElimLin algorithm can be represented as in Fig. 2. In fact, as a consequence of Lemma 3 and Lemma 5, Fig. 2 presents a unique characterization of $\text{Span}(\mathcal{S}_L)$ in terms of a fixed point:

- 1: **Input** : A set \mathcal{S}^0 of polynomial equations in R_d .
- 2: **Output** : A system of linear equations \mathcal{S}_L .
- 3: Set $\bar{\mathcal{S}}_L := \emptyset$.
- 4: **repeat**
- 5: $\bar{\mathcal{S}}_L \leftarrow \text{Span}(\mathcal{S}^0 \cup (R_{d-1} \times \bar{\mathcal{S}}_L)) \cap R_1$
- 6: **until** $\bar{\mathcal{S}}_L$ unchanged
- 7: Output \mathcal{S}_L : a basis of $\bar{\mathcal{S}}_L$.

Fig. 2. ElimLin algorithm from another perspective

Lemma 6. *At the end of ElimLin, $\text{Span}(\mathcal{S}_L)$ is the smallest subset $\bar{\mathcal{S}}_L$ of R_1 , such that*

$$\bar{\mathcal{S}}_L = \text{Span}(\mathcal{S}^0 \cup (R_{d-1} \times \bar{\mathcal{S}}_L)) \cap R_1$$

Proof. By induction, at step α we have $\bar{\mathcal{S}}_L \subseteq \text{Span}(\mathcal{S}_L^\alpha)$, using Lemma 5. Also, $\mathcal{S}_L^\alpha \subseteq \bar{\mathcal{S}}_L$ using Lemma 3. So, $\bar{\mathcal{S}}_L = \text{Span}(\mathcal{S}_L^\alpha)$ at step α . Since

$$\bar{\mathcal{S}}_L \mapsto \text{Span}(\mathcal{S}^0 \cup (R_{d-1} \times \bar{\mathcal{S}}_L)) \cap R_1$$

is increasing, we obtain the above equation. □

ElimLin eliminates variables, thus it looks very unexpected that the number of linear equations in each step of the algorithm is invariant with respect to any variable ordering in the substitution step and the Gaussian elimination. We finally prove this important invariant property. Concretely, we formalize ElimLin as a sequence of intersection of vector spaces. Such intersection in each iteration is between the vector space spanned by the equations and the vector space generated by all monomials of degree 1 in the system. This implies that if ElimLin runs for α iterations (finally succeeds or fails), it can be formalized as a sequence of intersections of α pairs of vector spaces. These intersections of vector spaces only depend on the vector space of the initial system.

Theorem 7. *The following relations exist after running ElimLin on a polynomial system of equations Q :*

1. $\text{Root}(\mathcal{S}^0) = \text{Root}(\mathcal{S}^T \cup \mathcal{S}_L)$
2. *There is no linear equation in $\text{Span}(\mathcal{S}^T)$.*
3. *$\text{Span}(\mathcal{S}_L)$ is uniquely defined by \mathcal{S}^0 .*
4. *\mathcal{S}_L consists of linearly independent linear equations.*
5. *The complexity is $O(n_0^{d+1}m_0^2)$, where d is the degree of the system and n_0 and m_0 are the initial number of variables and equations, respectively.*

Proof (1). Due to Lemma 3 and Lemma 4, \mathcal{S}^0 and $(\mathcal{S}^T \cup \mathcal{S}_L)$ are equivalent. So, a solution of \mathcal{S}^0 is also a solution of $(\mathcal{S}^T \cup \mathcal{S}_L)$ and vice versa.

Proof (2). Since ElimLin stops on \mathcal{S}^T , the Gaussian reduction did not find any linear polynomial.

Proof (3). Due to Lemma 6.

Proof (4). \mathcal{S}_L includes a basis for $\bar{\mathcal{S}}_L$. So, it consists of linearly independent equations.

Proof (5). n_0 is an upper bound on $\#\mathcal{S}_L$ due to the fact that \mathcal{S}_L consists of linearly independent linear equations. So, the number of iterations is bounded by n_0 . The total number of monomials is bounded by

$$T_0 \leq \sum_{i=0}^d \binom{n_0}{i} = O(n_0^d)$$

The complexity of Gaussian elimination is $O(m_0^2 T_0)$, since we have T_0 columns and m_0 equations. Therefore, overall, the complexity of ElimLin is $O(n_0^{d+1}m_0^2)$. □

4.2 Affine Bijective Variable Change


In the next theorem, we prove that the result of ElimLin algorithm does not change for any affine bijective variable change. It is an open problem to find an appropriate non-linear variable change which improves the result of the ElimLin algorithm.

Theorem 8. *Any affine bijective variable change $A : \text{GF}(2)^{n_0} \rightarrow \text{GF}(2)^{n_0}$ on a n_0 -variable system of equations \mathcal{S}^0 does not affect the result of ElimLin algorithm, implying that the number of linear equations generated at each iteration is invariant with respect to an affine bijective variable change.*

Proof. In Lemma 6, we showed that $\text{Span}(\mathcal{S}_L)$ is the output of the algorithm in Fig. 2, iterating

$$\bar{\mathcal{S}}_L \leftarrow \text{Span}(\mathcal{S}^0 \cup (R_{d-1} \times \bar{\mathcal{S}}_L)) \cap R_1$$

We represent the composition of a polynomial f_1 with respect to A by $\text{Com}(f_1)$. We then show that there is a commutative diagram.



We consider two parallel executions of the algorithm in Fig. 2, one with \mathcal{S}^0 and the other with $\text{Com}(\mathcal{S}^0)$.

If we compose the polynomials in \mathcal{S}^0 with respect to A , in the above relation R_{d-1} remains the same. Since the transformation A is affine,

$$\text{Com}(\text{Span}(\mathcal{S}^0 \cup (R_{d-1} \times \bar{\mathcal{S}}_L)) \cap R_1) = \text{Span}(\text{Com}(\mathcal{S}^0) \cup (R_{d-1} \times \text{Com}(\bar{\mathcal{S}}_L))) \cap R_1$$

So, at each iteration, the second execution has the result of applying Com to the result of the first one. \square

4.3 Linear Equations Evolution

An open problem regarding ElimLin is to predict how the number of linear equations evolves in the preceding iterations. In the following theorem, we give a necessary (but not sufficient) condition for a dense overdefined system of equations to have an additional linear equation in the next iteration of ElimLin. Proving a similar result for a sparse system is not straightforward.

Theorem 9. *If we apply ElimLin to an overdefined dense system of quadratic equations over GF(2), for $n_L^{\alpha+1} > n_L^\alpha$ to hold, it is necessary to have*

$$\frac{b_\alpha}{2} - a_\alpha < n_L^\alpha < \frac{b_\alpha}{2} + a_\alpha$$

where $b_\alpha = 2n_\alpha - 1$ and $a_\alpha = \frac{\sqrt{b_\alpha^2 - 8n_L^\alpha}}{2}$.

Proof. For the system to generate linear equations, it is necessary that the *sufficient rank condition* [4] is satisfied. More clearly, we must have $r_\alpha > T_\alpha - 1 - n_\alpha$, otherwise no linear equations will be generated. This is true if the system of equations is overdefined. Hence, we obtain,

$$n_L^\alpha = r_\alpha + n_\alpha + 1 - T_\alpha \tag{5}$$

If some columns of the matrix \mathcal{M}_α are pivotless, it will shift the diagonal strand of ones to the right. Therefore, n_L^α will be more than what the above equation expresses. Assuming the system of equations is dense, this phenomenon happens with a very low probability. Suppose the above equation is true with high probability, then we get

$$n_L^{\alpha+1} = r_{\alpha+1} + n_{\alpha+1} + 1 - T_{\alpha+1} \tag{6}$$

In the $(\alpha + 1)$ -th iteration, the number of variables is reduced by n_L^α . Thus, $n_{\alpha+1} = n_\alpha - n_L^\alpha$. If the system of equations is dense, in a quadratic system,

$$T_\alpha = \binom{n_\alpha}{2} + n_\alpha + 1$$

and so,

$$T_{\alpha+1} = \binom{n_\alpha - n_L^\alpha}{2} + n_\alpha - n_L^\alpha + 1$$

Consequently, we have

$$T_\alpha - T_{\alpha+1} = n_L^\alpha \left(n_\alpha - \frac{1}{2}(n_L^\alpha - 1) \right) \quad (7)$$

Therefore, using Eq. (5), Eq. (6) and Eq. (7), we obtain,

$$\begin{aligned} n_L^{\alpha+1} &= (r_{\alpha+1} - r_\alpha) + (r_\alpha + n_\alpha - T_\alpha + 1) + n_L^\alpha \left(-\frac{1}{2}n_L^\alpha + n_\alpha - \frac{1}{2} \right) \\ &= n_L^\alpha \left(-\frac{1}{2}n_L^\alpha + n_\alpha + \frac{1}{2} \right) - (r_\alpha - r_{\alpha+1}) \end{aligned}$$

If $n_L^{\alpha+1} > n_L^\alpha$, then $n_L^\alpha \left(-\frac{1}{2}n_L^\alpha + n_\alpha + \frac{1}{2} \right) - (r_\alpha - r_{\alpha+1}) > n_L^\alpha$ and this leads to

$$n_L^{\alpha 2} + (1 - 2n_\alpha)n_L^\alpha + 2(r_\alpha - r_{\alpha+1}) < 0$$

$\Delta = (1 - 2n_\alpha)^2 - 8(r_\alpha - r_{\alpha+1})$, and if the above inequality holds, Δ should be positive and assuming $b_\alpha = 2n_\alpha - 1$, then, $b_\alpha - \sqrt{\Delta} < 2n_L^\alpha < b_\alpha + \sqrt{\Delta}$.

Considering Δ is positive, $n_\alpha > \sqrt{2(r_\alpha - r_{\alpha+1})} + \frac{1}{2}$. We also know that $r_{\alpha+1} \leq r_\alpha - n_L^\alpha$, which together lead to $n_\alpha > \frac{1}{2} + \sqrt{2n_L^\alpha}$. Therefore, for $n_L^{\alpha+1} > n_L^\alpha$, it is necessary to have $n_\alpha > \frac{1}{2} + \sqrt{2n_L^\alpha}$, but not visa versa. Simplifying $b_\alpha - \sqrt{\Delta} < 2n_L^\alpha < b_\alpha + \sqrt{\Delta}$ and deploying $r_\alpha - r_{\alpha+1} \geq n_L^\alpha$ results in

$$b_\alpha - 2a_\alpha < 2n_L^\alpha < b_\alpha + 2a_\alpha$$

where $b_\alpha = 2n_\alpha - 1$ and $2a_\alpha = \sqrt{b_\alpha^2 - 8n_L^\alpha}$.

Notice that $n_\alpha > \frac{1}{2} + \sqrt{2n_L^\alpha}$, which was obtained in the first stage of the proof, has been originated from the fact that $b_\alpha^2 - 8n_L^\alpha$ should be non-negative. \square

5 Attacks Simulations

In this section, we present our experimental results against CTC2, LBlock and MIBS block ciphers. The simulations for CTC2 were run on an ordinary PC with a 1.8 Ghz CPU and 2 GB RAM. All the other simulations were run on an ordinary PC with a 2.8 Ghz CPU and 4 GB RAM. The amount of RAM required by our implementation is negligible.

In our attacks, we build a system of quadratic equations with variables representing plaintext, ciphertext, key and state bits, which allows to express the system of equations of high degree as quadratic equations. Afterwards, for each sample we set the plaintext and ciphertext according to the result of the input/output of the cipher. In order to test the efficiency of the algebraic attack, we guess some bits of the key and set the key variables corresponding to the guess. Then, we run the solver (ElimLin, F4 or SAT solver) to recover the remaining key bits and test whether the guess was correct. Therefore, the complexity of our algebraic attack can be bounded by $2^g \cdot \mathcal{C}(\text{solver})$, where $\mathcal{C}(\text{solver})$ represents the running time of the solver and g is the number of bits we guess. $\mathcal{C}(\text{solver})$ is represented as the the “*Running Time*” in all the following tables.

For a comparison with a brute force attack, we consider a fair implementation of the cipher, which requires 10 CPU cycles per round. This implies that the algebraic attack against t rounds of the cipher is faster than an exhaustive search for the 1.8 Ghz and 2.8 Ghz CPU iff recovering c bits of the key is faster than $5.55 \cdot t \cdot 2^{c-31}$ and $3.57 \cdot t \cdot 2^{c-31}$ seconds respectively. This is already twice faster than the complexity of exhaustive search. All the attacks reported in the following tables are faster than exhaustive search with the former argument. In fact, we consider the cipher to be broken for some number of rounds if the algebraic attack that recovers $(\#key - g)$ key bits is faster than an exhaustive key search over $(\#key - g)$ bits of the key.

5.1 Simulations Using F4 Algorithm under PolyBoRi Framework

The most efficient implementation of the F4 algorithm is available under PolyBoRi framework [8] running alone or under SAGE algebra system. PolyBoRi is a C++ library designed to compute Gröbner basis of an ideal applied to Boolean polynomials. A Python interface is used, surrounding the C++ core. It uses zero-suppressed binary decision diagrams (ZDDs) [33] as a high level data structure for storing Boolean polynomials. This representation stores the monomials more efficiently in memory and it makes the Gröbner basis computation faster compared to other algebra systems.

We use polybori-0.8.0 for our attacks. Together with ElimLin, we also attack LBlock and MIBS with F4 algorithm and then compare PolyBoRi’s efficiency with our implementation of ElimLin.

5.2 Simulations on CTC2

Courtois Toy Cipher (CTC) is an SPN-based block cipher devised by Courtois [13] as a toy cipher to evaluate algebraic attacks on smaller variants of cryptosystems. It was designed to show that it is possible to break a cipher using an ordinary PC deploying a small number of known or chosen plaintext-ciphertext pairs.

Since the system of equations of well-known ciphers such as AES is often large, it is not feasible by the current algorithms and computer capacities to solve them in a reasonable time, therefore smaller but similar versions such as CTC can be exploited to evaluate the resistance of ciphers against algebraic cryptanalysis. This turns out to yield a benchmark on understanding the algebraic structure of ciphers. Ultimately, this might lead to break of a larger system later.

CTC was not designed to be resistant against all known types of attacks like linear and differential cryptanalysis. Nevertheless, in [25], it was attacked by linear cryptanalysis. Subsequently, CTC Version 2 or CTC2 was proposed [12] to resolve the flaw exists in CTC structure. CTC2 is very similar to CTC with a few changes. It is an SPN-based network with scalable number of rounds, block and key size. For the full specification, refer to [12]. In CT-RSA 2009, differential and differential-linear attacks could reach up to 8 rounds of CTC2 [26], but as stated before, the objective of the CTC designer was not applying statistical attacks to his design. Finally, there is a cube attack on 4 rounds of one variant of this cipher in [41].

Since block size and key size are flexible in CTC2 cipher, we break various versions with distinct parameters (see Table 1) using ElimLin. The block size is specified by a parameter B , which specifies the number of parallel S-boxes per round. CTC2 S-box is 3×3 , hence the block size is computed as $3B$. We guess some LSB bits of the key and we show that recovering the remaining is faster than exhaustive search.

It might be possible that during the intermediate steps of ElimLin, a quadratic equation in only key bits (possibly linear) appears. In such cases, approximately $\mathcal{O}(\#key^2)$ samples are enough to break the system. This is due to the fact that we can simply change the plaintext-ciphertext pair and generate a new linearly independent equation in the key. Finally, when we have enough such equations, we solve a system of quadratic equations in only key bits using the linearization technique. When such phenomenon occurs, intuitively the cipher is close to be broken but not yet. We can increase the number of samples and most often it makes the cipher thoroughly collapse.

5.3 Simulations on LBlock

LBlock is a new lightweight Feistel-based block cipher, aimed at constrained environments, such as RFID tags and sensor networks [49] proposed at ACNS 2011. It operates on 64-bit blocks, uses a key of 80 bits and iterates 32 rounds. For a detailed specification of the cipher, refer to [49]. As far as the authors are aware, there is currently no cryptanalysis results published on this cipher.

We break 8 rounds of LBlock using 6 samples deploying an ordinary PC by ElimLin. Our results are summarized in Table 2. In the same scenario, PolyBoRi crashes due to running out of memory.

Table 1. CTC2 simulations using ElimLin up to 6 rounds with distinct parameters

B	N_r	$\#key$	g	Running Time ¹ (in hours)	Running Time ² (in hours)	Data	Attack notes	
16	3	48	0	0.03		5 KP	ElimLin	
16	3	48	0		0.12	14 KP	ElimLin	
64	3	192	155		0.03	1 KP	ElimLin	
85	3	255	210		0.04	1 KP	ElimLin	
16	4	48	0	0.01	0.05	2 CP	ElimLin	
16	4	48	0			4 CP	ElimLin	
40	4	120	85	0.00		1 KP	ElimLin	
40	4	120	85		0.84	16 KP	ElimLin	
48	4	144	100		0.12	4 KP	ElimLin	
64	4	192	148	0.05		1 KP	ElimLin	
64	4	192	155		2.21	5 KP	ElimLin	
85	4	255	220	0.29		1 KP	ElimLin	
85	4	255	215	0.64		1 KP	ElimLin	
85	4	255	220		0.26	2 KP	ElimLin	
85	4	255	215		0.90	3 KP	ElimLin	
85	4	255	210		1.33	4 KP	ElimLin	
16	5	48	0	3	0.03	8 CP	ElimLin	
40	5	120	85			2 CP	ElimLin	
32	6	96	60	2.5	3	16 CP	ElimLin	
40	6	120	80	1		8 CP	ElimLin	
64	6	192	155	2.4		4 CP	ElimLin	
85	6	255	210	3		2 CP	ElimLin	
85	6	255	220			16 CP	ElimLin	
85	6	255	210			180.5	64 CP	ElimLin
128	6	384	344			4.5	2 CP	ElimLin

B : Number of S-boxes per round. To obtain the block size, B is multiplied by 3.

N_r : Number of rounds

g : Number of guessed LSB of the key

Running Time¹: Running time until we achieve equations only in key variables (no other internal variables). When this is achieved, the cipher is close to be broken, but not yet (see Sec. 5.2).

Running Time²: Attack running time for recovering $(\#key - g)$ bits of the key.

KP: Known plaintext

CP: Chosen plaintext

5.4 Simulations on MIBS

Similar to the LBlock block cipher, MIBS is also a lightweight Feistel-based block cipher, aimed at constrained environments, such as RFID tags and sensor networks [38]. It operates on 64-bit blocks, uses keys of 64 or 80 bits and iterates 32 rounds. For a detailed specification of the cipher, see [38].

Currently, the best cryptanalysis results is a linear attack reaching 18-round MIBS with data complexity 2^{61} and time complexity of 2^{76} [5]. In fact, statistical attacks often require very large number of samples. This is not always achievable in practice.

Table 2. Algebraic attack complexities on reduced-round LBlock using ElimLin and PolyBoRi

N_r	$\#key$	g	Running Time (in hours)	Data	Attack notes
8	80	32	0.252	6 KP	ElimLin
8	80	32	crashed	6 KP	PolyBoRi

N_r : Number of rounds

g : Number of guessed LSB of the key

KP: Known plaintext

CP: Chosen plaintext

We break 4 and 3 rounds of MIBS80 and MIBS64 using 32 and 2 samples deploying an ordinary PC by ElimLin. Our results are summarized in Table 3. In 2 out of 3 experiments, PolyBoRi crashes due to running out of memory. This is the first algebraic analysis of the cipher.

The designers in [38] have evaluated the security of their cipher with respect to algebraic attacks. They used the complexity of XSL algorithm for this evaluation, which is not a precise measurement for evaluating resistance of a cipher against algebraic attacks, since effectiveness of XSL is still controversial and under speculation. There are better methods such as SAT solvers [3] which solve MQ problem faster than expected due to the system being overdefined and sparse.

Let assume XSL can be precise enough to evaluate the security of a cipher with respect to algebraic attacks. According to [20,38], the complexity of XSL can be evaluated with the work factor. For MIBS, work factor is computed as:

$$WF = \Gamma^\omega \left((\text{Block Size}) \cdot N_r^2 \right)^{\omega \lceil \frac{T}{r} \rceil}$$

where Γ is a parameter which depends only on the S-box. For MIBS, $\Gamma = 85.56$. The value $r = 21$ is the number of equations the S-box can be represented with. $T = 37$ is the number of monomials in that representation. $\omega = 2.37$ is the exponent of the Gaussian elimination complexity. The work factor for attacking 5-round MIBS is $WF = 2^{65.65}$ which is worse than an exhaustive key search for MIBS64. Deploying SAT solving techniques using MiniSAT 2.0 [28], we can break 5 rounds of MIBS64 (see Table 3). Our strategy is exactly the same as [3]. Table 3 already shows that we can do better than $2^{65.65}$ for MIBS64. We can perform a very similar attack on MIBS80. This already shows that considering the complexity of XSL is not a precise measure to evaluate the security of a cipher against algebraic cryptanalysis. Complexity of attacking such system with XL is extremely high.

We believe that due to the similarity between the structure of MIBS and LBlock, we can compare them with respect to algebraic attacks. As can be seen from the table of attacks, LBlock is much weaker. This is not surprising though, since the linear layer of LBlock is much weaker than MIBS, since it is nibble-wise instead of bit-wise. So, we could attack twice more rounds of LBlock. Thus, although LBlock is lighter with respect to the number of gates, but it provides a lower level of security with respect to algebraic attacks.

Table 3. Algebraic attack complexities on reduced-round MIBS using ElimLin, PolyBoRi and MiniSAT 2.0

N_r	#key	g	Running Time (in hours)	Data	Attack notes
4	80	20	0.137	32 KP	ElimLin
4	80	20	crashed	32 KP	PolyBoRi
5	64	16	0.395	6 KP	MiniSAT 2.0
5	64	16	crashed	6 KP	PolyBoRi
3	64	0	0.006	2 KP	ElimLin
3	64	0	0.002	2 KP	PolyBoRi

N_r : Number of rounds

g : Number of guessed LSB of the key

KP: Known plaintext

CP: Chosen plaintext

6 A Comparison between ElimLin and PolyBoRi

Gröbner basis is currently one of the most successful methods for solving polynomial systems of equations. However, it has its own restrictions. The main bottleneck of the Gröbner basis techniques is the memory requirement and therefore most of the Gröbner basis attacks use relatively small number of samples. It is worthwhile to mention that ElimLin is a subroutine in Gröbner basis computations. But, ElimLin algorithm as a single tool requires a large number of samples to work.

The Gröbner basis solve the system by reductions according to a pre-selected ordering, which can lead to high degree dense polynomials. ElimLin uses the fact that multiple samples provide an additional information to the solver, and therefore the key might be found even if when we restrict the reduction to degree 2.

Next, we compare the current state of the art implementation of F4 algorithm PolyBoRi and our implementation of ElimLin. In the cases where ElimLin behaves better than PolyBoRi, it does not mean that ElimLin is superior to F4 algorithm. In fact, it just means that there exists a better implementation for ElimLin than for F4 for some particular systems of equations. F4 uses a fixed ordering for monomials and therefore it does not preserve the sparsity in its intermediate steps. On the other hand, our implementation of ElimLin performs several sparsity preserving techniques by changing the ordering. This drops the total number of monomials and makes it memory efficient.

Table 2 and Table 3 show that PolyBoRi requires too much memory and crashes for a large number of samples. At the same time, our implementation of ElimLin is slightly slower than PolyBoRi implementation attacking 2 samples of 3-round MIBS64 as in Table 3. This demonstrates that our implementation of ElimLin can be more effective than PolyBoRi and vice versa, depending on memory requirements of PolyBoRi. However, whenever the system is solvable by our implementation of ElimLin, our experiments revealed that PolyBoRi does not give a significant advantage over ElimLin because the memory requirements are too high.

While PolyBoRi may yield a solution for a few samples, the success of ElimLin is determined by the number of samples provided to the algorithm. The evaluation of the number of sufficient samples in ElimLin is still an open problem.

We see that often preserving the degree by simple linear algebra techniques can outperform the more sophisticated Gröbner basis algorithms, mainly due to the structural properties of the system of equations of a cryptographic primitive (such as sparsity). ElimLin takes advantage of such structural properties and uncovers some hidden linear equations using multiple samples. According to our experiments, PolyBoRi does not seem to be able to take advantage of these structural properties as would be expected which results in higher memory requirements than would be necessary and ultimately their failure for large systems, even though it is clearly possible for the algorithm to find the solution in reasonable time. Finally, we need more efficient implementations and data structures for both ElimLin and Gröbner basis algorithms.

7 Further Work and Some Conjectures

An interesting area of research is to estimate the number of linear equations in ElimLin or anticipate how this number evolves in the succeeding iterations or evaluate after how many iterations ElimLin finishes. Also, to anticipate how many samples is enough to make the system collapse by ElimLin. Last but not least, it is prominent to find a very efficient method for implementing ElimLin and to find the most appropriate data structure to choose.

There are some evidence which illustrate that ElimLin does not reveal all hidden linear equations in the structure of the cipher up to a specific degree. We give an example, demonstrating such an evidence:

Assume there exists an equation in the system which can be represented as $\ell(x)g(x) + 1 = 0$ over $\text{GF}(2)$, where $\ell(x)$ is a polynomial of degree one and $g(x)$ is a polynomial of degree at most $d - 1$. Running ElimLin on this single equation trivially fails. But, if we multiply both sides of the equation by $\ell(x)$, we obtain $\ell(x)g(x) + \ell(x) = 0$. Summing these two equations, we derive $\ell(x) = 1$. This hidden linear equation can be simply captured by the XL algorithm, but can not be captured by ElimLin. There exist multiple other examples which demonstrate that ElimLin does not generate all the hidden linear equations. To generate all such linear equations, the degree-bounded Gröbner basis can be used.

For big ciphers, for example the full AES, it is also plausible that:

Conjecture 1. *For each number of rounds X , there exists Y such that AES is broken by ElimLin given Y Chosen or Known Plaintext-Ciphertext pairs.*

Disproving the above conjecture leads to the statement that “AES can not be broken by algebraic attack at degree 2”. But maybe this conjecture is true, then the capacities of the ElimLin attack are considerable and it works for any number of rounds X . As a consequence, if for $X = 14$ this Y is not too large, say less than 2^{64} , the AES-256 will be broken faster than brute force by ElimLin at degree

2, which is much simpler than Gröbner basis objective of breaking it at degree 3 or 4 with 1 KP.

ElimLin is a polynomial time algorithm. If it can be shown that a polynomial number of samples is enough to gain a high success rate for ElimLin, this can already be considered a breakthrough in cryptography. Unfortunately, the correctness of this statement is not clear.

8 Conclusion

In this paper, we proved that ElimLin can be formulated in terms of a sequence of intersections of vector spaces. We showed that different monomial orderings and any affine bijective variable change do not influence the result of the algorithm. We did some predictions on the evolution of linear equations in the succeeding iterations in ElimLin. We presented multiple attacks deploying ElimLin against CTC2, LBlock and MIBS block ciphers.

References

1. Armknecht, F., Ars, G.: Algebraic Attacks on Stream Ciphers with Gröbner Bases. In: Gröbner Bases, Coding, and Cryptography, pp. 329–348 (2009)
2. Aumasson, J.-P., Henzen, L., Meier, W., Naya-Plasencia, M.: QUARK: A Lightweight Hash. In: Mangard, S., Standaert, F.-X. (eds.) CHES 2010. LNCS, vol. 6225, pp. 1–15. Springer, Heidelberg (2010)
3. Bard, G., Courtois, N., Jefferson, C.: Efficient Methods for Conversion and Solution of Sparse Systems of Low-Degree Multivariate Polynomials over $GF(2)$ via SAT-Solvers. Presented at ECRYPT Workshop Tools for Cryptanalysis (2007), <http://eprint.iacr.org/2007/024.pdf>
4. Bard, G.V.: Algebraic Cryptanalysis. Springer (2009)
5. Bay, A., Nakahara Jr., J., Vaudenay, S.: Cryptanalysis of Reduced-Round MIBS Block Cipher. In: Heng, S.-H., Wright, R.N., Goi, B.-M. (eds.) CANS 2010. LNCS, vol. 6467, pp. 1–19. Springer, Heidelberg (2010)
6. Bogdanov, A., Knežević, M., Leander, G., Toz, D., Varıcı, K., Verbauwhede, I.: SPONGENT: A Lightweight Hash Function. In: Preneel, B., Takagi, T. (eds.) CHES 2011. LNCS, vol. 6917, pp. 312–325. Springer, Heidelberg (2011)
7. Bogdanov, 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)
8. Brickenstein, M., Dreyer, A.: PolyBoRi: A framework for Gröbner basis computations with Boolean polynomials. In: Electronic Proceedings of MEGA 2007 (2007), <http://www.ricam.oeaw.ac.at/mega2007/electronic/26.pdf>
9. Buchberger, B.: Bruno Buchberger's PhD thesis 1965: An algorithm for finding the basis elements of the residue class ring of a zero dimensional polynomial ideal. *Journal of Symbolic Computation* 41(3-4), 475–511 (2006)
10. Courtois, N.T.: Higher Order Correlation Attacks, XL Algorithm and Cryptanalysis of Toyocrypt. In: Lee, P.J., Lim, C.H. (eds.) ICISC 2002. LNCS, vol. 2587, pp. 182–199. Springer, Heidelberg (2003)

11. Courtois, N.T.: Fast Algebraic Attacks on Stream Ciphers with Linear Feedback. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 176–194. Springer, Heidelberg (2003)
12. Courtois, N.: CTC2 and Fast Algebraic Attacks on Block Ciphers Revisited. In: Cryptology ePrint Archive (2007), <http://eprint.iacr.org/2007/152.pdf>
13. Courtois, N.: How Fast can be Algebraic Attacks on Block Ciphers? In: Symmetric Cryptography. Dagstuhl Seminar Proceedings, vol. 07021 (2007)
14. Courtois, N.: The Dark Side of Security by Obscurity - and Cloning MiFare Classic Rail and Building Passes, Anywhere, Anytime. In: SECRYPT, pp. 331–338 (2009)
15. Courtois, N.: Algebraic Complexity Reduction and Cryptanalysis of GOST. In: Cryptology ePrint Archive (2011), <http://eprint.iacr.org/2011/626>
16. Courtois, N.T., Bard, G.V.: Algebraic Cryptanalysis of the Data Encryption Standard. In: Galbraith, S.D. (ed.) Cryptography and Coding 2007. LNCS, vol. 4887, pp. 152–169. Springer, Heidelberg (2007)
17. Courtois, N.T., Debraize, B.: Algebraic Description and Simultaneous Linear Approximations of Addition in Snow 2.0. In: Chen, L., Ryan, M.D., Wang, G. (eds.) ICICS 2008. LNCS, vol. 5308, pp. 328–344. Springer, Heidelberg (2008)
18. Courtois, N.T., Meier, W.: Algebraic Attacks on Stream Ciphers with Linear Feedback. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 345–359. Springer, Heidelberg (2003)
19. Courtois, N.T., O’Neil, S., Quisquater, J.-J.: Practical Algebraic Attacks on the Hitag2 Stream Cipher. In: Samarati, P., Yung, M., Martinelli, F., Ardagna, C.A. (eds.) ISC 2009. LNCS, vol. 5735, pp. 167–176. Springer, Heidelberg (2009)
20. Courtois, N.T., Pieprzyk, J.: Cryptanalysis of Block Ciphers with Overdefined Systems of Equations. In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 267–287. Springer, Heidelberg (2002)
21. Courtois, N.T., Klimov, A.B., Patarin, J., Shamir, A.: Efficient Algorithms for Solving Overdefined Systems of Multivariate Polynomial Equations. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 392–407. Springer, Heidelberg (2000)
22. 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)
23. Dinur, I., Shamir, A.: Cube Attacks on Tweakable Black Box Polynomials. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 278–299. Springer, Heidelberg (2009)
24. Dinur, I., Shamir, A.: Breaking Grain-128 with Dynamic Cube Attacks. In: Joux, A. (ed.) FSE 2011. LNCS, vol. 6733, pp. 167–187. Springer, Heidelberg (2011)
25. Dunkelman, O., Keller, N.: Linear Cryptanalysis of CTC. In: Cryptology ePrint Archive (2006), <http://eprint.iacr.org/2006/250.pdf>
26. Dunkelman, O., Keller, N.: Cryptanalysis of CTC2. In: Fischlin, M. (ed.) CT-RSA 2009. LNCS, vol. 5473, pp. 226–239. Springer, Heidelberg (2009)
27. Eén, N., Sörensson, N.: MiniSat 2.0. An open-source SAT solver package, <http://www.cs.chalmers.se/Cs/Research/FormalMethods/MiniSat/>
28. Een, N., Sorensson, N.: Minisat - A SAT Solver with Conflict-Clause Minimization. In: Theory and Applications of Satisfiability Testing (2005)
29. Engels, D., Saarinen, M.-J.O., Schweitzer, P., Smith, E.M.: The Hummingbird-2 Lightweight Authenticated Encryption Algorithm. In: Juels, A., Paar, C. (eds.) RFIDSec 2011. LNCS, vol. 7055, pp. 19–31. Springer, Heidelberg (2012)

30. Faugère, J.: A new efficient algorithm for computing Gröbner bases (F4). *Journal of Pure and Applied Algebra* 139(1-3), 61–88 (1999)
31. Faugère, J.: A new efficient algorithm for computing Gröbner bases without reduction to zero (F5). In: *Symbolic and Algebraic Computation - ISSAC*, pp. 75–83 (2002)
32. Fusco, G., Bach, E.: Phase transition of multivariate polynomial systems. *Journal of Mathematical Structures in Computer Science* 19(1) (2009)
33. Ghasemzadeh, M.: A New Algorithm for the Quantified Satisfiability Problem, Based on Zero-suppressed Binary Decision Diagrams and Memoization. PhD thesis, University of Potsdam, Germany (2005)
34. Gong, Z., Nikova, S., Law, Y.W.: KLEIN: A New Family of Lightweight Block Ciphers. In: Juels, A., Paar, C. (eds.) *RFIDSec 2011*. LNCS, vol. 7055, pp. 1–18. Springer, Heidelberg (2012)
35. Guo, J., Peyrin, T., Poschmann, A.: The PHOTON Family of Lightweight Hash Functions. In: Rogaway, P. (ed.) *CRYPTO 2011*. LNCS, vol. 6841, pp. 222–239. Springer, Heidelberg (2011)
36. Guo, J., Peyrin, T., Poschmann, A., Robshaw, M.J.B.: The LED Block Cipher. In: Preneel, B., Takagi, T. (eds.) *CHES 2011*. LNCS, vol. 6917, pp. 326–341. Springer, Heidelberg (2011)
37. Indestege, S., Keller, N., Dunkelman, O., Biham, E., Preneel, B.: A Practical Attack on KeeLoq. In: Smart, N.P. (ed.) *EUROCRYPT 2008*. LNCS, vol. 4965, pp. 1–18. Springer, Heidelberg (2008)
38. Izadi, M., Sadeghiyan, B., Sadeghian, S., Arabnezhad, H.: MIBS: A New Lightweight Block Cipher. In: Garay, J.A., Miyaji, A., Otsuka, A. (eds.) *CANS 2009*. LNCS, vol. 5888, pp. 334–348. Springer, Heidelberg (2009)
39. Knudsen, L., Leander, G., Poschmann, A., Robshaw, M.J.B.: PRINTCIPHER: A Block Cipher for IC-Printing. In: Mangard, S., Standaert, F.-X. (eds.) *CHES 2010*. LNCS, vol. 6225, pp. 16–32. Springer, Heidelberg (2010)
40. Magma, software package, <http://magma.maths.usyd.edu.au/magma/>
41. Mroczkowski, P., Szmids, J.: The Cube Attack on Courtois Toy Cipher. In: *Cryptology ePrint Archive* (2009), <http://eprint.iacr.org/2009/497.pdf>
42. Murphy, S., Robshaw, M.J.B.: Essential Algebraic Structure within the AES. In: Yung, M. (ed.) *CRYPTO 2002*. LNCS, vol. 2442, pp. 1–16. Springer, Heidelberg (2002)
43. Nakahara Jr., J., Sepehrdad, P., Zhang, B., Wang, M.: Linear (Hull) and Algebraic Cryptanalysis of the Block Cipher PRESENT. In: Garay, J.A., Miyaji, A., Otsuka, A. (eds.) *CANS 2009*. LNCS, vol. 5888, pp. 58–75. Springer, Heidelberg (2009)
44. Raddum, H., Semaev, I.: Solving Multiple Right Hand Sides linear equations. *Journal of Designs, Codes and Cryptography* 49(1-3), 147–160 (2008)
45. Shannon, C.E.: Communication theory of secrecy systems. *Bell System Technical Journal* 28 (1949)
46. Shibutani, K., Isobe, T., Hiwatari, H., Mitsuda, A., Akishita, T., Shirai, T.: *Piccolo*: An Ultra-Lightweight Blockcipher. In: Preneel, B., Takagi, T. (eds.) *CHES 2011*. LNCS, vol. 6917, pp. 342–357. Springer, Heidelberg (2011)
47. Vielhaber, M.: Breaking ONE.FIVIUM by AIDA an Algebraic IV Differential Attack. In: *Cryptology ePrint Archive* (2007), <http://eprint.iacr.org/2007/413>
48. Weinmann, R.: Evaluating Algebraic Attacks on the AES. Master’s thesis, Technische Universität Darmstadt (2003)
49. Wu, W., Zhang, L.: LBlock: A Lightweight Block Cipher. In: Lopez, J., Tsudik, G. (eds.) *ACNS 2011*. LNCS, vol. 6715, pp. 327–344. Springer, Heidelberg (2011)

Short-Output Universal Hash Functions and Their Use in Fast and Secure Data Authentication

Long Hoang Nguyen and A.W. Roscoe

Oxford University, Department of Computer Science
hn2503@gmail.com, Bill.Roscoe@cs.ox.ac.uk

Abstract. Message authentication codes usually require the underlining universal hash functions to have a long output so that the probability of successfully forging messages is low enough for cryptographic purposes. To take advantage of fast operation on word-size parameters in modern processors, long-output universal hashing schemes can be securely constructed by concatenating several different instances of a short-output primitive. In this paper, we describe a new method for short-output universal hash function termed *digest()* suitable for very fast software implementation and applicable to secure message authentication. The method possesses a higher level of security relative to other well-studied and computationally efficient short-output universal hashing schemes. Suppose that the universal hash output is fixed at one word of b bits, then the collision probability of ours is 2^{1-b} compared to 6×2^{-b} of MMH, whereas $2^{-b/2}$ of NH within UMAC is far away from optimality. In addition to message authentication codes, we show how short-output universal hashing is applicable to manual authentication protocols where universal hash keys are used in a very different and interesting way.

1 Introduction

Universal hash functions (or UHF) first introduced by Carter and Wegman [4] have many applications in computer science, including randomised algorithms, database, cryptography and many others. A UHF takes two inputs which are a key k and a message m : $h(k, m)$, and produces a fixed-length output. Normally what we require of a UHF is that for any pair of distinct messages m and m' the collision probability $h(k, m) = h(k, m')$ is small when key k is randomly chosen from its domain. In the majority of cryptographic uses, UHF's usually have long outputs so that combinatorial search is made infeasible. For example, UHF's can be used to build secure message authentication codes or MAC schemes where the intruder's ability to forge messages is bounded by the collision probability of the UHF. In a MAC, parties share a secret universal hash key and an encryption key, a message is authenticated by hashing it with the shared universal hash key and then encrypting the resulting hash. The encrypted hash value together with the message is transmitted as an authentication tag that can be validated by the verifier. We note however that our new construction presented here is applicable

to other cryptographic uses of universal hashing, e.g., manual authentication protocols as seen later as well as non-cryptographic applications.

Since operating on short-length values of 16, 32 or 64 bits is fast and convenient in ordinary computers, long-output UHF's can be securely constructed by concatenating the results of multiple instances of short-output UHF's to increase computational efficiency. To our knowledge, a number of short-output UHF schemes have been proposed, notably MMH (Multilinear-Modular-Hashing) of Halevi and Krawczyk [8] and NH within UMAC of Black et al. [3]. We note that widely studied polynomial universal hashing schemes GHASH, PolyP and PolyQ [12] can also be designed to produce a short output. While polynomial based UHF's only require short and fixed length keys, they suffer from an unpleasant property relating to security as will be discussed later in the paper.

Our main contribution presented in Section 3 is the introduction of a new short-output UHF algorithm termed $digest(k, m)$ that can be efficiently computed on any modern microprocessors. The main advantage of ours is that it provides a higher level of security regarding both collision and distribution probabilities relative to MMH and NH described in Section 4. Our $digest()$ algorithm operates on word-size parameters via word multiplication and word addition instructions, i.e. finite fields or non-trivial reductions are excluded, because the emphasis is on high speed implementation using software.

Let us suppose that the universal hash output is fixed at one word of b bits then the collision probability of ours is 2^{1-b} compared to 6×2^{-b} of MMH, whereas $2^{-b/2}$ of NH is much weaker in security. For clarity, the security bounds of our constructions as well as MMH and NH are independent of the length of message being hashed, which is the opposite of polynomial universal hashing schemes mentioned earlier. For multiple-word output universal hashing constructions as required in MACs, the advantage in security of ours becomes more apparent. When the universal hash output is extended to n words or $n \times b$ bits for any $n \in \mathbb{N}^*$, then the collision probability of ours is 2^{n-nb} as opposed to $6^n \times 2^{-nb}$ of MMH and $2^{-nb/2}$ of NH. There is however a trade-off between security and computational cost as illustrated by our estimated operation counts and software implementations of these constructions. On a 1GHz AMD Athlon processor, one version of $digest()$ (where the collision probability ϵ_c is 2^{-31}) achieves peak performance of 0.53 cycles/byte (or cpb) relative to 0.31 cpb of MMH (for $\epsilon_c = 2^{-29.5}$) and 0.23 cpb of NH (for $\epsilon_c = 2^{-32}$). Another version of $digest(k, m)$ for $\epsilon_c = 2^{-93}$ achieves peak performance of 1.54 cpb. For comparison purpose, 12.35 cpb is the speed of SHA-256 recorded on our computer. A number of files that provide the software implementations in C programming language of NH, MMH and our proposed constructions can be downloaded from [1] so that the reader can run them and adapt them for other uses of the schemes.

We will briefly discuss the motivation of designing as well as the elegant graphical structure of our $digest()$ scheme which, we have only recently discovered, relates to the multiplicative universal hashing schemes of Dietzfelbinger et al. [5], Krawczyk [11] and Mansour et al. [15]. The latter algorithms are however not efficient when the input message is of a significant size.

Although researchers from cryptographic community have mainly studied UHF's to construct message authentication codes, we would like to point out that short-output UHF on its own has found applications in manual authentication protocols [7,13,14,16,10,17,18,19,20,26]. In the new family of authentication protocols, data authentication can be achieved without the need of passwords, shared private keys as required in MACs, or any pre-existing security infrastructures such as a PKI. Instead human owners of electronic devices who seek to exchange their data authentically would need to manually compare a short string of bits that is often outputted from a UHF. Since humans can only compare short strings, the UHF ideally needs to have a short output of say 16 or 32 bits. There is however a fundamental difference in the use of universal hash keys between manual authentication protocols and message authentication codes, it will be clear in Section 5 that none of the short-output UHF schemes including ours should be used directly in the former. Thus we will propose a general framework where any short-output UHF's can be used efficiently and securely to digest a large amount of data in manual authentication protocols.

While existing universal hashing methods are already as fast as the rate information is generated, authenticated and transmitted in high-speed network traffic, one may ask whether we need another universal hashing algorithm. Besides keeping up with network traffic, as excellently explained by Black et al. [3] — *the goal is to use the smallest possible fraction of the CPU's cycles (so most of the machine's cycles are available for other work), by the simplest possible hash mechanism, and having the best proven bounds.* This is relevant to MACs as well as manual authentication protocols where large data are hashed into a short string, and hence efficient short-output UHF constructions possessing a higher (or optimal) level of security are needed.

2 Notation and Definitions

We define M , K and b the bit length of the message, the key and the output of a universal hash function. We denote $R = \{0, 1\}^K$, $X = \{0, 1\}^M$ and $Y = \{0, 1\}^b$.

Definition 1. [11] A ϵ -balanced universal hash function, $h : R \times X \rightarrow Y$, must satisfy that for every $m \in X \setminus \{0\}$ and $y \in Y$: $\Pr_{\{k \in R\}}[h(k, m) = y] \leq \epsilon$

Many existing UHF constructions [3,8,11] as well as our newly proposed scheme rely on (integer or matrix) multiplications of message and key, and hence non-zero input message is required; for otherwise $h(k, 0) = 0$ for any key $k \in R$.

Definition 2. [11,24] A ϵ -almost universal hash function, $h : R \times X \rightarrow Y$, must satisfy that for every $m, m' \in X$ ($m \neq m'$): $\Pr_{\{k \in R\}}[h(k, m) = h(k, m')] \leq \epsilon$

Since it is useful particularly in manual authentication protocols discussed later to have both the collision and distribution probabilities bounded, we combine Definitions 1 and 2 as follows.

Definition 3. An ϵ_d -balanced and ϵ_c -almost universal hash function, $h : R \times X \rightarrow Y$, satisfies

- for every $m \in X \setminus \{0\}$ and $y \in Y$: $\Pr_{\{k \in R\}}[h(k, m) = y] \leq \epsilon_d$
- for every $m, m' \in X$ ($m \neq m'$): $\Pr_{\{k \in R\}}[h(k, m) = h(k, m')] \leq \epsilon_c$

3 Integer Multiplication Construction

We first discuss the multiplicative universal hashing algorithm of Dietzfelbinger et al. [5] which obtains a very high level of security. Although this scheme is not efficient with long input data, it strongly relates to our *digest()* method that makes use of word multiplication instructions.

We note that there are two other universal hashing schemes which use arithmetic that computer likes to do to increase computational efficiency, namely MMH of Halevi and Krawczyk [8] and NH of Black et al. [3]. Both of which will be compared against our construction in Section 4.

3.1 Multiplicative Universal Hashing

Suppose that we want to compute a b -bit universal hash of a M -bit message, then the universal hash key k is drawn randomly from $R = \{1, 3, \dots, 2^M - 1\}$, i.e. k must be odd. Dietzfelbinger et al. [5] define:

$$h(k, m) = (k * m \bmod 2^M) \operatorname{div} 2^{M-b}$$

It was proved that the collision probability of this construction is $\epsilon_c = 2^{1-b}$ on equal length inputs [5]. While this has a simple description, for long input messages of several KB or MB, such as documents and images, it will become very time consuming to compute the integer multiplication involved in this algorithm.

3.2 Word Multiplicative Construction

In this section, we will define and prove the security of a new short-output universal hashing scheme termed *digest(k, m)* that can be calculated using word multiplications instead of an arbitrarily long integer multiplication as seen in Equation 1 or an example from Figure 1.

Let us divide message m into b -bit blocks $\langle m_1, \dots, m_{t=M/b} \rangle$. An $(M + b)$ -bit key $k = \langle k_1, \dots, k_{t+1} \rangle$ is selected randomly from $R = \{0, 1\}^{M+b}$. A b -bit *digest(k, m)* is defined as

$$\operatorname{digest}(k, m) = \sum_{i=1}^t [m_i * k_i + (m_i * k_{i+1} \operatorname{div} 2^b)] \bmod 2^b \quad (1)$$

Here, $*$ refers to a word multiplication of two b -bit blocks which produces a $2b$ -bit output, whereas both $+$ and \sum are additions modulo 2^b . We note that $(\operatorname{div} 2^b)$ is equivalent to a right shift ($\gg b$) and hence is efficient to compute.




Fig. 1. A b -bit output $digest(k, m)$: each parallelogram represents the expansion of a word multiplication between a b -bit key block and a b -bit message block

To see why this scheme is related to the multiplicative method of Dietzfelbinger et al. [5], one can study Figure 1 where all word multiplications involved in Equation 1 are elegantly arranged into the same shape as the overlap of the expanded multiplication between m and k .¹

Essentially what we are doing here is to obtain a short b -bit window in the middle of the product without computing the whole product.² Such an idea is very similar to the SQUASH hash function of Shamir [23] that produces an excellent numeric approximation of the b middle bits by computing a longer window of $b + u$ bits with u additional lower order bits so that the full effect of the carry bits is significantly restored. There are however three crucial differences between ours and SQUASH: (1) we do not need to compute the extra u lower order words or bits, (2) we partially ignore the carry between words, and (3) as opposed to SQUASH, the security of $digest()$ does not rely on the Rabin public key encryption scheme [23]. The first two of these make ours much faster in computation.

Operation Count. To give an estimated operation count for an implementation of $digest()$, which will be subsequently compared against universal hashing schemes MMH and NH, we consider a machine with the same properties as one used by Halevi and Krawczyk [8].³

¹ If we further ignore the effect of the carry in (word) multiplications of both $digest()$ and the scheme of Dietzfelbinger et al. then they become very similar to the Toeplitz matrix based construction of Krawczyk [11] and Mansour et al. [15]. Such a carry-less multiplication instruction is available in a new Intel processor [2].

² This idea was first reported in our patent application [21] dated back to 2006.

³ The same operation count given here is applicable to a ($2b = 64$)-bit machine. In the latter, a multiplication of two 32-bit unsigned integer is stored in a single 64-bit register, and *High* and *Low* are the upper and lower 32-bit halves of the register.

- ($b = 32$)-bit machine and arithmetic operations are done in registers.
- A multiplication of two 32-bit integers yields a 64-bit result that is stored in 2 registers.

A pseudo-code for $digest()$ on such machine may be as follows. For a 'C' implementation, please see [1].

```

digest(key, msg)
1.   Sum = 0
2.   load key[1]
3.   for i = 1 to t
4.     load msg[i]
5.     load key[i + 1]
6.     ⟨High1, Low1⟩ = msg[i] * key[i]
7.     ⟨High2, Low2⟩ = msg[i] * key[i + 1]
8.     Sum = Sum + Low1 + High2
9.   return Sum

```

This consists of $2t = 2M/b$ word multiplications (MULT) and $2t = 2M/b$ addition modulo 2^b (ADD). That is each message-word requires 1 MULT and 2 ADD operations. As in [8], a MULT/ADD operation should include not only the actual arithmetic instruction but also loading the message- and key-words to registers and/or loop handling.

The following theorem shows that the switch from a single (arbitrarily long) multiplication of Dietfelbinger et al. into word multiplications of $digest()$ does not weaken the security of the construction. Namely the same collision probability of 2^{1-b} is retained while optimality in distribution is achieved. Moreover this change not only greatly increases computational efficiency but also removes the restriction of odd universal hash key as required in Dietfelbinger et al.

Theorem 1. For any $t, b \geq 1$, $digest()$ of Equation 1 satisfies Definition 3 with the distribution probability $\epsilon_d = 2^{-b}$ and the collision probability $\epsilon_c = 2^{1-b}$ on equal length inputs.

Proof. We first consider the collision property. For any pair of distinct messages of equal length: $m = m_1 \cdots m_t$ and $m' = m'_1 \cdots m'_t$, without loss of generality we assume that $m_1 > m'_1$.⁴ A digest collision is equivalent to:

$$\sum_{i=1}^t [m_i * k_i + (m_i * k_{i+1} \operatorname{div} 2^b)] = \sum_{i=1}^t [m'_i * k_i + (m'_i * k_{i+1} \operatorname{div} 2^b)] \pmod{2^b}$$

There are two possibilities as follows.

⁴ Please note that when $m_i = m'_i$ for all $i \in \{1, \dots, j\}$ then in the following calculation we will assume that $m_{j+1} > m'_{j+1}$.

WHEN $m_1 - m'_1$ is odd. The above equality can be rewritten as

$$(m_1 - m'_1)k_1 = y \pmod{2^b} \tag{2}$$

where

$$y = (m'_1 k_2 \operatorname{div} 2^b) - (m_1 k_2 \operatorname{div} 2^b) + \sum_{i=2}^t [(m'_i - m_i) * k_i + (m'_i * k_{i+1} \operatorname{div} 2^b) - (m_i * k_{i+1} \operatorname{div} 2^b)]$$

We note that y depends only on keys k_2, \dots, k_{t+1} , and hence we fix k_2 through k_{t+1} in our analysis. Since $m_1 - m'_1$ is odd, i.e. $m_1 - m'_1$ and 2^b are co-prime, there is at most one value of k_1 satisfying Equation 2. The collision probability in this case is therefore $\epsilon_c = 2^{-b} < 2^{1-b}$.

WHEN $m_1 - m'_1$ is even. A digest collision can be rewritten as

$$(m_1 - m'_1)k_1 + (m_1 k_2 \operatorname{div} 2^b) - (m'_1 k_2 \operatorname{div} 2^b) + (m_2 - m'_2)k_2 = y \pmod{2^b} \tag{3}$$

where

$$y = (m'_2 k_3 \operatorname{div} 2^b) - (m_2 k_3 \operatorname{div} 2^b) + \sum_{i=3}^t [(m'_i - m_i) * k_i + (m'_i * k_{i+1} \operatorname{div} 2^b) - (m_i * k_{i+1} \operatorname{div} 2^b)] \pmod{2^b}$$

We note that y depends only on keys k_3, \dots, k_{t+1} . If we fix k_3 through k_{t+1} in our analysis, we need to find the number of pairs (k_1, k_2) such that Equation 3 is satisfied. We arrive at

$$\epsilon_c = \operatorname{Prob}_{\{k_1, k_2\}} [(m_1 - m'_1)k_1 + (m_1 k_2 \operatorname{div} 2^b) - (m'_1 k_2 \operatorname{div} 2^b) + (m_2 - m'_2)k_2 = y]$$

Let us define

$$\begin{aligned} m_1 k_2 &= u 2^b + v \\ m'_1 k_2 &= u' 2^b + v' \end{aligned}$$

Since we assumed $m_1 > m'_1$, we have $u \geq u'$ and $(m_1 - m'_1)k_2 = (u - u')2^b + v - v'$.

- When $v \geq v'$: $(m_1 k_2 \operatorname{div} 2^b) - (m'_1 k_2 \operatorname{div} 2^b) = (m_1 - m'_1)k_2 \operatorname{div} 2^b$
- When $v < v'$: $(m_1 k_2 \operatorname{div} 2^b) - (m'_1 k_2 \operatorname{div} 2^b) = [(m_1 - m'_1)k_2 \operatorname{div} 2^b] + 1$

Let $c = m_1 - m'_1$ and $d = m_2 - m'_2 \pmod{2^b}$, we then have $1 \leq c < 2^b$ and:

$$\epsilon_c \leq p_1 + p_2$$

where

$$p_1 = \operatorname{Prob}_{\substack{\{0 \leq k_1 < 2^b\} \\ \{0 \leq k_2 < 2^b\}}} [ck_1 + (ck_2 \operatorname{div} 2^b) + dk_2 = y \pmod{2^b}]$$

and

$$p_2 = \text{Prob}_{\left\{ \begin{smallmatrix} 0 \leq k_1 < 2^b \\ 0 \leq k_2 < 2^b \end{smallmatrix} \right\}} [ck_1 + (ck_2 \text{ div } 2^b) + dk_2 = y - 1 \pmod{2^b}]$$

Using Lemma 1, we have $p_1, p_2 \leq 2^{-b}$, and thus $\epsilon_c \leq 2^{1-b}$.

As regards distribution, since $m = m_1 \cdots m_t > 0$ as specified in Definition 3, without loss of generality we assume that $m_1 \geq 1$. If we fix k_3 through k_{t+1} and for any $y \in \{0, \dots, 2^b - 1\}$, then the distribution probability ϵ_d is equivalent to:

$$\epsilon_d = \text{Prob}_{\left\{ \begin{smallmatrix} 0 \leq k_1 < 2^b \\ 0 \leq k_2 < 2^b \end{smallmatrix} \right\}} [m_1 k_1 + (m_1 k_2 \text{ div } 2^b) + m_2 k_2 = y \pmod{2^b}]$$

Since $1 \leq m_1 < 2^b$, we can use Lemma 1 to deduce that $\epsilon_d = 2^{-b}$. □

Lemma 1. Let $1 \leq c < 2^b$ and $0 \leq d < 2^b$, then for any $y \in \{0, \dots, 2^b - 1\}$ we have

$$\text{Prob}_{\left\{ \begin{smallmatrix} 0 \leq k_1 < 2^b \\ 0 \leq k_2 < 2^b \end{smallmatrix} \right\}} [ck_1 + (ck_2 \text{ div } 2^b) + dk_2 = y \pmod{2^b}] = 2^{-b}$$

Proof. We write $c = s2^l$ with s odd and $0 \leq l < b$. Since s and 2^b are co-prime, there exist a unique inverse modulo 2^b of s , we call it s^{-1} . Our equation now becomes:

$$2^l sk_1 + (2^l sk_2 \text{ div } 2^b) + ds^{-1} sk_2 = y \pmod{2^b}$$

Let $sk_1 = \gamma \pmod{2^{b-l}}$ and $sk_2 = \alpha 2^{b-l} + \beta \pmod{2^b}$, we then have $0 \leq \gamma < 2^{b-l}$ and $0 \leq \alpha < 2^l$. The above equation becomes:

$$\begin{aligned} 2^l \gamma + \alpha + ds^{-1}(\alpha 2^{b-l} + \beta) &= y \pmod{2^b} \\ 2^l \gamma + \alpha(1 + ds^{-1}2^{b-l}) + \beta ds^{-1} &= y \pmod{2^b} \\ 2^l \gamma + \alpha x &= z \pmod{2^b} \end{aligned}$$

where $x = 1 + ds^{-1}2^{b-l} \pmod{2^b}$ which is always odd because $l < b$, and $z = y - \beta ds^{-1} \pmod{2^b}$. Since z is independent of γ and α , we fix z in our analysis. We can then use Lemma 2 to derive that there is a unique pair (γ, α) satisfying the above equation.

Since $0 \leq \gamma < 2^{b-l}$ and $0 \leq \alpha < 2^l$, γ and α together determine b bits of the combination of k_1 and k_2 . Consequently there are at most 2^b different pairs (k_1, k_2) satisfying the condition that we require in this lemma. □

Lemma 2. Let $0 \leq l < b$ and $x \in \{1, 3, \dots, 2^b - 1\}$ then for any $z \in \{0, \dots, 2^b - 1\}$ there is a unique pair (γ, α) such that $0 \leq \gamma < 2^{b-l}$, $0 \leq \alpha < 2^l$, and $2^l \gamma + \alpha x = z \pmod{2^b}$.

Proof. If there exist two distinct pairs (γ, α) and (γ', α') satisfying this condition, then

$$2^l \gamma + \alpha x = 2^l \gamma' + \alpha' x = z \pmod{2^b}$$

which implies that

$$2^l(\gamma - \gamma') = (\alpha' - \alpha)x \pmod{2^b}$$

This leads to two possibilities.

- When $\alpha' = \alpha$ then $2^l(\gamma - \gamma') = 0$, which means that $2^{b-l} | (\gamma - \gamma')$. The latter is impossible because $0 \leq \gamma, \gamma' < 2^{b-l}$ and $\gamma \neq \gamma'$.
- When $\alpha' \neq \alpha$ and since x is odd, we must have $2^l | (\alpha' - \alpha)$. This is also impossible because $0 \leq \alpha, \alpha' < 2^l$.

□

REMARKS. The bound given by Theorem 1 for the distribution probability ($\epsilon_d = 2^{-b}$) is tight: let $m = 0^{b-1}1$ and any y and note that any key $k = k_1k_2$ with $k_1 = y$ satisfying this equation $\text{digest}(k, m) = y$. The bound given by Theorem 1 for the collision probability $\epsilon_c = 2^{1-b}$ also appears to be tight, i.e. it cannot be reduced to 2^{-b} . To verify this bound, we have implemented exhaustive tests on single-word messages with small value of b . For example, when $b = 7$, we look at all possible pairs of two different ($b = 7$)-bit messages in combination with all ($2b = 14$)-bit keys, the obtained collision probability is $2^{-7} \times 1.875$.

We end this section by pointing out that truncation is secure in this digest construction. For any $b' \in \{1, \dots, b-1\}$, we define

$$\text{trunc}_{b'}(\text{digest}(k, m)) = \sum_{i=1}^t [m_i * k_i + (m_i * k_{i+1} \text{ div } 2^b)] \text{ mod } 2^{b'} \quad (4)$$

where $\text{trunc}_{b'}()$ takes the first b' least significant bits of the input. We then have the following theorem whose proof is very similar to the proof of Theorem 1, and hence it is not given here.

Theorem 2. For any $n, t \geq 1$, $b \geq 1$ and any integer $b' \in \{1, \dots, b-1\}$, $\text{trunc}_{b'}(\text{digest}())$ of Equation 4 satisfies Definition 3 with the distribution probability $\epsilon_d = 2^{-b'}$ and the collision probability $\epsilon_c = 2^{1-b'}$ on equal length inputs.

3.3 Extending $\text{digest}()$

If we want to use digest functions as the main ingredient of a message authentication code, we need to reduce the collision probability without increasing the word bitlength b that is dictated by architecture characteristics. One possibility is to hash our message with several random and independent keys, and concatenate the results. If we concatenate the results from n independent instances of the digest function, the collision probability drops from 2^{1-b} to 2^{n-nb} . This solution however requires n times as much key material.

A much better and well-studied approach is to use the Toeplitz-extension: given one key we left shift the key by one word to get the next key and digest again. The resulting construction is called $\text{digest}_{MW}()$, where MW stands




Fig. 2. A $3b$ -bit (or three-word) output $\text{digest}_{MW}(k, m)$: each parallelogram represents the expansion of a word multiplication between a b -bit key block and a b -bit message block

for multiple-word output. The structure of $\text{digest}_{MW}()$ is again graphically illustrated by an example in Figure 2 that shows a similar connection between $\text{digest}_{MW}()$ and the multiplicative hashing scheme of Dietfelbinger et al.

We define a n -blocks or $(n \times b)$ -bit output $\text{digest}_{MW}(k, m)$ as follows. We still divide m into b -bit blocks $\langle m_1, \dots, m_{t=M/b} \rangle$. However, an $(M + bn)$ -bit key $k = \langle k_1, \dots, k_{t+n} \rangle$ will be chosen randomly from $R = \{0, 1\}^{M+bn}$ to compute a nb -bit digest.

For all $i \in \{1, \dots, n\}$, we then define:

$$d_i = \text{digest}(k_{i \dots t+i}, m) = \sum_{j=1}^t [m_j k_{i+j-1} + (m_j k_{i+j} \text{ div } 2^b)] \pmod{2^b}$$

And

$$\text{digest}_{MW}(k, m) = \langle d_1 \cdots d_n \rangle$$

The following theorem and its proof show that $\text{digest}_{MW}()$ enjoys the best bound for both collision and distribution probabilities that one could hope for.

Theorem 3. For any $n, t \geq 1$ and $b \geq 1$, $\text{digest}_{MW}()$ satisfies Definition 3 with the distribution probability $\epsilon_d = 2^{-nb}$ and the collision probability $\epsilon_c = 2^{n-nb}$ on equal length inputs.

Proof. We first consider the collision property of a digest function. For any pair of distinct messages of equal length: $m = m_1 \cdots m_t$ and $m' = m'_1 \cdots m'_t$, without loss of generality we assume that $m_1 > m'_1$. Please note that when $t = 1$ or

$m_i = m'_i$ for all $i \in \{1, \dots, t - 1\}$ then in the following calculation we will assume that $m_{t+1} = m'_{t+1} = 0$.

For $i \in \{1, \dots, n\}$, we define Equality E_i as

$$E_i : \sum_{j=1}^t [m_j k_{i+j-1} + (m_j k_{i+j} \operatorname{div} 2^b)] = \sum_{j=1}^t [m'_j k_{i+j-1} + (m'_j k_{i+j} \operatorname{div} 2^b)] \pmod{2^b}$$

and thus the collision probability is: $\epsilon_c = \operatorname{Prob}_{\{k \in R\}}[E_1 \wedge \dots \wedge E_n]$.

Since all arithmetic operations are done over modulo 2^b , for simplicity we ignore $(\operatorname{mod} 2^b)$ in our notation.

WHEN $m_1 - m'_1$ is odd. We proceed by proving that for all $i \in \{1, \dots, n\}$

$$\operatorname{Prob}[E_i \text{ is true} \mid E_{i+1}, \dots, E_n \text{ are true}] \leq 2^{-b}$$

For Equality E_n , the claim is satisfied due to Theorem 1. We notice that Equalities E_{i+1} through E_n depend only on keys k_{i+1}, \dots, k_{n+t} , whereas Equality E_i depends also on key k_i . Fix k_{i+1} through k_{n+t} such that Equalities E_{i+1} through E_n are satisfied. We prove that there is at most one value of k_i satisfying E_i . To achieve this we let

$$z = (m'_1 k_{i+1} \operatorname{div} 2^b) - (m_1 k_{i+1} \operatorname{div} 2^b) + \sum_{j=2}^t [(m'_j - m_j) k_{i+j-1} + (m'_j k_{i+j} \operatorname{div} 2^b) - (m_j k_{i+j} \operatorname{div} 2^b)]$$

we then rewrite Equality E_i as

$$(m_1 - m'_1) k_i = z$$

Since we assumed $m_1 - m'_1$ is odd, there is at most one value of k_i satisfying this equation.

WHEN $m_1 - m'_1$ is even. We write $m_1 - m'_1 = 2^l s$ with s odd and $0 < l < b$, and $s' = (m'_2 - m_2) s^{-1}$. We further denote $sk_i = x_i 2^{b-l} + y_i$ for $i \in \{1, \dots, n + t\}$, where $0 \leq x_i < 2^l$ and $0 \leq y_i < 2^{b-l}$.

For $i \in \{1, \dots, n\}$, if we define $b_i \in \{0, 1\}$ and

$$f(y_i, x_{i+1}) = 2^l y_i + x_{i+1} [(m_2 - m'_2) s^{-1} 2^{b-l} + 1]$$

$$g(k_{i+2}, \dots, k_{i+t}) = (m'_2 k_{i+2} \operatorname{div} 2^b) + \sum_{j=3}^t [m'_j k_{i+j-1} + (m'_j k_{i+j} \operatorname{div} 2^b)] - (m_2 k_{i+2} \operatorname{div} 2^b) - \sum_{j=3}^t [m_j k_{i+j-1} + (m_j k_{i+j} \operatorname{div} 2^b)]$$

then, using similar trick as in the proof of Lemma 1, Equality E_i can be rewritten as

$$\begin{aligned} (m_1 - m'_1)k_i + ((m_1 - m'_1)k_{i+1} \operatorname{div} 2^b) + (m_2 - m'_2)k_{i+1} &= g(k_{i+2}, \dots, k_{i+t}) - b_i \\ 2^l sk_i + (2^l sk_{i+1} \operatorname{div} 2^b) + (m_2 - m'_2)s^{-1}sk_{i+1} &= g(k_{i+2}, \dots, k_{i+t}) - b_i \\ 2^l y_i + x_{i+1} + (m_2 - m'_2)s^{-1}(x_{i+1}2^{b-l} + y_{i+1}) &= g(k_{i+2}, \dots, k_{i+t}) - b_i \end{aligned}$$

Rearranging gives

$$\begin{aligned} 2^l y_i + x_{i+1}[(m_2 - m'_2)s^{-1}2^{b-l} + 1] &= s'y_{i+1} - b_i + g(k_{i+2}, \dots, k_{i+t}) \\ f(y_i, x_{i+1}) &= s'y_{i+1} - b_i + g(k_{i+2}, \dots, k_{i+t}) \end{aligned}$$

Putting Equalities E_1 through E_n together, we have

$$\begin{aligned} E_1 : f(y_1, x_2) &= s'y_2 - b_1 + g(k_3, \dots, k_{1+t}) \\ E_2 : f(y_2, x_3) &= s'y_3 - b_2 + g(k_4, \dots, k_{2+t}) \\ E_3 : f(y_3, x_4) &= s'y_4 - b_3 + g(k_5, \dots, k_{3+t}) \\ &\vdots \\ E_{n-1} : f(y_{n-1}, x_n) &= s'y_n - b_{n-1} + g(k_{n+1}, \dots, k_{n+t-1}) \\ E_n : f(y_n, x_{n+1}) &= s'y_{n+1} - b_n + g(k_{n+2}, \dots, k_{n+t}) \end{aligned}$$

We fix k_{n+2} through k_{t+n} . We note that there are 2^{b-t} values for y_{n+1} and two values for b_n . For each pair (y_{n+1}, b_n) there is a unique pair (y_n, x_{n+1}) satisfying Equality E_n due to Lemma 2. Similarly, for each tuple $\langle y_n, k_{n+1}, b_{n-1}, b_n \rangle$ there is also a unique pair (y_{n-1}, x_n) satisfying Equality E_{n-1} . We will continue this process until we reach the pair (y_1, x_2) in Equality E_1 . Since Equalities E_1 through E_n do not depend on x_1 and there are 2^l values for x_1 , there will be at most $2^l 2^n 2^{b-l} = 2^{n+b}$ different tuples $\langle k_1 \cdots k_{n+1} \rangle$ satisfying Equalities E_1 through E_n . And thus the collision probability $\epsilon_c = 2^{n+b}/2^{(n+1)b} = 2^{n-nb}$.

Similar argument also leads to our bound on the distribution probability $\epsilon_d = 2^{-nb}$. □

REMARKS. Even though Theorems 1 and 3 address the collision property, their proofs can be easily adapted to show that our constructions are also ϵ_c -almost- Δ -universal [8] as in the case of the MMH scheme considered in the next section. The latter property requires that for every $m, m' \in X$ where $m \neq m'$ and $a \in Y$: $\Pr_{\{k \in R\}}[\operatorname{digest}(k, m) - \operatorname{digest}(k, m') = a] \leq \epsilon_c$.

Operation Count. The advantage of this scheme is the ability to reuse the result of each word multiplication in the computation of two adjacent digest output words as seen in Figure 2 and the following pseudo-code, e.g. the multiplication $m_1 k_2$ is instrumental in the computation of both d_1 and d_2 . Using the same machine as specified in subsection 3.2, each message-word therefore requires $(n + 1)$ MULT and $2n$ ADD operations.

A pseudo-code for $digest_{MW}()$ on such machine may be as follows

$digest_{MW}(key, msg)$

1. For $i = 1$ to n
2. $d[i] = 0$
3. load $key[i]$
4. For $j = 1$ to t
5. load $msg[j]$
6. load $key[j + n]$
7. $\langle High[0], Low[0] \rangle = msg[j] * key[j]$
8. For $i = 1$ to n
9. $\langle High[i], Low[i] \rangle = msg[j] * key[j + i]$
10. $d[i] = d[i] + Low[i - 1] + High[i]$
11. return $\langle d[1] \dots d[n] \rangle$

4 Comparative Analysis

In this section, we mainly compare our new digest scheme against well-studied universal hashing algorithms MMH of Halevi and Krawczyk [8] and NH of Black et al. [3] described in Subsections 4.1 and 4.2 respectively. Since $digest()$ can be extended to produce multiple-word output as in the case of MMH and NH to build MACs, our analysis consider both single- and multiple-word output schemes. We note that NH is the building block of not only UMAC but also UHASH16 and UHASH32 [3]. For completeness, we will discuss another widely studied UHF family based on polynomial over finite field, e.g. GHASH, PolyP, PolyQ and PolyR [12]. While the polynomial universal hashing schemes only require short keys, they suffer from two unpleasant properties: (1) the collision probability decreases linearly with the message length, and (2) they are less efficient, especially in software implementation, than our digest functions as well as MMH and NH due to the involved modular arithmetic operations.

The properties of the three main schemes – MMH, NH and $digest()$ – are summarised in Table 1 where the upper and lower halves correspond to single-word (b bits) and respectively multiple-word (nb bits) output schemes for any $n \geq 1$. This table indicates that the security level obtained in our digest algorithm is higher than both MMH and NH with respect to the same output length. In particular, the collision probability of $digest()$ is a third of MMH, while NH must double the output length to achieve the same order of security. For multiple-word output schemes, this advantage in security of our proposed digest algorithm becomes even more significant as seen in the lower half of Table 1.

We end this section by providing implementation results in Table 2 of Section 4.3. As described earlier, C files which contain the implementations of NH, MMH and $digest()$ as well as their multiple-word output versions can be downloaded from [1] which allows readers to test the speed of the constructions.

Table 1. A summary on the main properties of *digest()*, MMH and NH. MULT operates on b -bit inputs, whereas ADD operates on inputs of either b or $2b$ bits.

Scheme	Key length	MULTs/word	ADDs/word	ϵ_c	ϵ_d	Output bitlength
<i>digest</i>	$M + b$	2	2	2^{1-b}	2^{-b}	b
MMH	M	1	1	6×2^{-b}	2^{2-b}	b
NH	M	1/2	3/2	2^{-b}	2^{-b}	$2b$
<i>digest_{MW}</i>	$M + nb$	$n + 1$	$2n$	2^{n-nb}	2^{-nb}	nb
MMH _{MW}	$M + (n - 1)b$	n	n	$6^n \times 2^{-nb}$	2^{2n-nb}	nb
NH _{MW}	$M + 2(n - 1)b$	$n/2$	$3n/2$	2^{-nb}	2^{-nb}	$2nb$

4.1 MMH

Fix a prime number $p \in [2^b, 2^b + 2^{b/2}]$. The b -bit output MMH universal hash function is defined for any $k = k_1, \dots, k_t$ and $m = m_1, \dots, m_t$ as follows

$$\text{MMH}(k, m) = \left[\left[\left[\sum_{i=1}^t m_i * k_i \right] \bmod 2^{2b} \right] \bmod p \right] \bmod 2^b$$

It was proved in [8] that the collision probability of MMH is $\epsilon_c = 6 \times 2^{-b}$ as opposed to only 2^{1-b} of *digest()*. By using the same proof technique presented in [8], it is also not hard to show that the distribution probability of MMH is $\epsilon_d = 2^{2-b}$, as opposed to 2^{-b} of *digest()*.⁵

For single-word output, each message word in MMH requires 1 ($b \times b$) MULT and 1 ADD modulo 2^{2b} . We note however that this does not include the cost of the final reduction modulo p . For n -word output MMH, using “the Toeplitz matrix approach”, the scheme is defined as

$$\text{MMH}_{MW}(k, m) = \text{MMH}(k_{1\dots t}, m) \parallel \text{MMH}(k_{2\dots t+1}, m) \parallel \dots \parallel \text{MMH}(k_{n\dots t+n-1}, m)$$

MMH_{MW} obtains $\epsilon_c = 6^n 2^{-nb}$ and $\epsilon_d = 2^{2n-nb}$, which are considerably weaker than *digest_{MW}*() ($\epsilon_c = 2^{n-nb}, \epsilon_d = 2^{-nb}$).

4.2 NH

The $2b$ -bit output NH universal hash function is defined for any $k = k_1, \dots, k_t$ and $m = m_1, \dots, m_t$, where t is even, as follows

$$\text{NH}(k, m) = \sum_{i=1}^{t/2} (k_{2i-1} + m_{2i-1})(k_{2i} + m_{2i}) \bmod 2^{2b}$$

⁵ There is a Square Hash variant of MMH introduced by Etzel et al. [6] that is defined as follows: $\text{SQH}(k, m) = [\sum_{i=1}^t ((m_i + k_i) \bmod 2^b)^2] \bmod p$. The collision probability ϵ_c of SQH is 2^{1-b} . While squaring an integer is more efficient than multiplication, SQH is not really faster than MMH because the summation of the squares does not fit into $2b$ bits and hence extra words are required to store and add when long data are hashed. Etzel et al. then optimise the implementation by ignoring the carry bits in the computation, but this makes the collision probability bound bigger than 2^{1-b} .

The downside of NH relative to MMH and our digest method is the level of security obtained, namely with a $2b$ -bit output, which is twice the length of both $digest()$ and MMH, NH was shown to have the collision probability $\epsilon_c = 2^{-b}$ and the distribution probability $\epsilon_d = 2^{-b}$, which are far from optimality. Its computational cost is however lower than the other twos, i.e. each message-word requires only $1/2 (b \times b)$ MULT, 1 ADD modulo 2^b , and $1/2$ ADD modulo 2^{2b} .

For $2n$ -word output, also using “the Toeplitz matrix approach”, we have $\epsilon_c = 2^{-nb}$ and $\epsilon_d = 2^{-nb}$. Each message-word requires $n/2$ MULT and $3n/2$ ADD operations as seen below.

$$NH_{MW}(k, m) = NH(k_{1\dots t}, m) \parallel NH(k_{3\dots t+2}, m) \parallel \dots \parallel NH(k_{2n-1\dots t+2(n-1)}, m)$$

Table 2. Performance (cycles/byte) of $digest$, MMH and NH constructions. In each row, the length of NH is always twice the length of MMH and $digest$.

<i>digest</i>			MMH			NH		
Output bitlength	ϵ_c	Speed (cpb)	Output bitlength	ϵ_c	Speed (cpb)	Output bitlength	ϵ_c	Speed (cpb)
32	2×2^{-32}	0.53	32	6×2^{-32}	0.31	64	2^{-32}	0.23
64	$2^2 \times 2^{-64}$	1.05	64	$6^2 \times 2^{-64}$	0.57	128	2^{-64}	0.39
96	$2^3 \times 2^{-96}$	1.54	96	$6^3 \times 2^{-96}$	0.76	192	2^{-96}	0.62
160	$2^5 \times 2^{-160}$	2.13	160	$6^5 \times 2^{-160}$	1.37	320	2^{-160}	1.15
256	$2^8 \times 2^{-256}$	3.44	256	$6^8 \times 2^{-256}$	2.31	512	2^{-256}	1.90

4.3 Implementations of MMH, NH and Digest Constructions

We have tested the implementations of $digest()$, MMH, NH as well as their multiple-word output versions on a workstation with a 1GHz AMD Athlon(tm) 64 X2 Dual Core Processor (4600+ or 512 KB caches) running the 2.6.30 Linux kernel. All source codes were written in C making use of GCC 4.4.1 compiler. The number of cycles elapsed during execution was measured by the $clock()$ instruction in the normal way (as in UMAC [25]) in our C implementations [1].

For comparison, we recompiled publicly available source codes for SHA-256 and SHA-512 [22] whose reported speeds on our workstation are 12.35 cpb and 8.54 cpb respectively.

For application of these primitives in MACs, normally a tree hash structure is used to significantly reduce the length of universal hash key. The downside of this method is that the resulting collision probability is not constant but proportional to the depth of the tree. For this reason, previously reported speeds for MMH and NH [3,8] and our results do not include the cost of key generation.

Table 2 shows the results of the experiments, which were averaged over a large number of random data inputs of at least 8 kilobytes. The speeds are in cycles/byte or cpb. Our digest constructions, at the cost of higher security, are slightly slower than MMH and NH due to extra multiplications, but still considerably faster than standard cryptographic hash functions SHA-256/512.

4.4 Polynomial Universal Hashing Schemes

In this section, we will study another well-studied class of UHF based on polynomial over finite fields, including GHASH within Galois Counter Mode, PolyP, PolyQ and PolyR [12].

For simplicity, we will give a simple version of polynomial universal hashing that is the core of PolyP, PolyQ, PolyR and GHASH. Let the set of all messages be $\{m = \langle m_1, \dots, m_t \rangle; m_i \in \mathbb{F}_p\}$, here p is the largest prime number less than 2^b and the message length is $M = tb$ bits. For any key $k \in \mathbb{F}_p$, we define:

$$\text{Poly}(k, m) = m_1 + m_2k + m_3k^2 + \dots + m_tk^{t-1} \pmod{p}$$

Such a scheme does have two nice properties as follows

- The key length of the b -bit output $\text{Poly}()$ scheme is fixed at b bits regardless of the message length.
- $\text{Poly}()$ provides collision resistance for both equal and unequal length messages. Suppose that the bit lengths of two different messages m and m' are bt and bt' , then the collision probability is $\max\{t - 1, t' - 1\}/p$. On the other hand, MMH, NH and $\text{digest}()$ only ensure collision resistance for equal length data, but not unequal length messages. The latter is intuitively because unequal length messages in $\text{digest}()$, MMH and NH require unequal length keys, which make them incomparable for collision analysis.

The main disadvantage of a polynomial universal hashing scheme is that its security bound depends on the message length, which is the opposite of MMH, NH and $\text{digest}()$. Namely, the collision probability of $\text{Poly}()$ is $\epsilon = (t - 1)2^{-b}$ that is no where near the level of security obtained by our digest function when message is of a significant size. The security downside of polynomial universal hash functions does have a negative impact on their use in manual authentication where short-output but highly secure universal hash functions are required.

5 Short-Output Universal Hash Functions in Manual Authentication Protocols

In addition to MAC schemes, short-output universal hash functions have found use in manual authentication protocols where parties A and B want to authenticate their public data $m_{A/B}$ to each other without the need for passwords, shared private keys as in MACs, or pre-established security infrastructures such as a PKI. Instead authentication is bootstrapped from human trust and interactions.

Using notation taken from authors' work [17,18,19] the N -indexed arrow (\rightarrow_N) indicates an unreliable and high-bandwidth (or normal) link where messages can be maliciously altered, whereas the E -indexed arrow (\rightarrow_E) represents an authentic and unspoofable (or empirical) channel. The latter is not a private channel (anyone can overhear it) and it is usually very low-bandwidth since it is implemented by humans, e.g., human conversations or manual data transfers between devices. $\text{hash}()$ is a cryptographic hash function. Long random keys

$k_{A/B}$ are generated by A/B , and k_A is kept secret until after k_B is revealed in Message 2. Operators \parallel and \oplus denote bitwise concatenation and exclusive-or.

A pairwise manual authentication protocol [13,14,17]	
1.	$A \rightarrow_N B : m_A, \text{hash}(A \parallel k_A)$
2.	$B \rightarrow_N A : m_B, k_B$
3.	$A \rightarrow_N B : k_A$
4.	$A \leftarrow_E B : h(k^*, m_A \parallel m_B)$ where $k^* = k_A \oplus k_B$

To ensure devices agree on the same data $m_A \parallel m_B$, their human owners manually compare the universal hash in Message 4. As human interactions are expensive, the universal hash function needs to have a short output of $b \in [16, 32]$ bits.

As seen from the above protocol, the universal hash key k^* always varies randomly and uniformly from one to another protocol run. In other words, no value of k^* is used to hash more than one message because $k_{A/B}$ instrumental in the computation of k^* are randomly chosen in each protocol run. This is fundamentally different from MACs which often use the same private key to hash multiple messages for a period of time,⁶ and hence attacks which rely on the reuse of a single private key in multiple sessions are irrelevant in manual authentication protocols. What we then want to understand is the collision and distribution properties of the universal hash function. We stress that this analysis is also applicable to group manual authentication protocols [17,18,19,26].

Should *digest()*, MMH or NH be used directly in Message 4 of the above protocol, random and fresh keys $k_{A/B}$ of similar size as $m_A \parallel m_B$ must be generated whenever the protocol is run. Obviously one can generate a long random key stream from a short seed via a pseudo-random number generator, but it can be computationally expensive especially when the authenticated data $m_{A/B}$ are of a significant size. Of course we can use one of the polynomial universal hashing functions (e.g. PolyP32, PolyQ32 or PolyR16_32 all defined in [12]) which require a short key. But since humans only can compare short value over the empirical channel, it is intolerable that the security bound of the universal hash function degrades linearly along with the length of data being authenticated.

One possibility suggested in [7,20] is to truncate the output of a cryptographic hash function to the b least significant bits:

$$h(k, m) = \text{trunc}_b(\text{hash}(k \parallel m))$$

Although it can be computationally infeasible to search for a full cryptographic hash collision, it is not clear whether the truncated solution is sufficiently secure because the definition of a hash function does not normally specify the distribution of individual groups of bits.

What we therefore propose is a combination of cryptographic hashing and short-output universal hash functions. Without loss of generality, we use our

⁶ Even when a new universal hash key is generated every time a message is hashed in MAC as recommended by Handschuh and Preneel [9], the long key is still derived from the same private and much shorter seed or key shared between the parties.

digest method in the following construction which is also applicable to MMH and NH. Let $hash()$ be a B -bit cryptographic hash function, e.g. SHA-2/3. First the input key is split into two parts of unequal lengths $k = k_1 \parallel k_2$, where k_1 is $B+b$ bits and k_2 is at least 80 bits. Then our modified construction $digest'()$ which takes an arbitrarily length message m is computed as follows⁷

$$digest'(k, m) = digest(k_1, hash(m \parallel k_2))$$

We denote θ_c the hash collision probability of $hash()$, and it should be clear that $\theta_c \gg 2^{-b}$ given that $b \in [16, 32]$. The following theorem shows that this construction essentially preserves both the collision and distribution security bounds, and at the same time removes the restriction on equal length input messages because the hash function $hash()$ always produces a fixed length value.

Theorem 4. $digest'()$ satisfies Definition 3 with the distribution probability $\epsilon_d = 2^{-b}$ and the collision probability $\epsilon_c = 2^{1-b} + \theta_c$.

Proof. Let l_1 and l_2 denote the bitlengths of keys k_1 and k_2 respectively.

We first consider collision property of $digest'()$. For any pair of distinct messages m and m' , as key k_2 varies uniformly and randomly the probability that $hash(m \parallel k_2) = hash(m' \parallel k_2)$ is θ_c . So there are two possibilities:

- When $hash(m \parallel k_2) = hash(m' \parallel k_2)$ then $digest(k_1, hash(m \parallel k_2)) = digest(k_1, hash(m' \parallel k_2))$ for any key $k_1 \in \{0, 1\}^{l_1}$.
- When $hash(m \parallel k_2) \neq hash(m' \parallel k_2)$ then $digest(k_1, hash(m \parallel k_2)) = digest(k_1, hash(m' \parallel k_2))$ with probability 2^{1-b} .

Consequently the collision probability of $digest'()$ is

$$\theta_c + (1 - \theta_c)2^{1-b} < \theta_c + 2^{1-b}$$

As regards distribution probability of $digest'()$, we fix message m of arbitrarily length and a b -bit value y in our analysis. For each value of k_2 , there will be at most 2^{l_1-b} different keys k_1 such that $digest(k_1, hash(m \parallel k_2)) = y$. Since there are 2^{l_2} different keys k_2 , there will be at most $2^{l_1-b}2^{l_2} = 2^{l_1+l_2-b}$ different pairs (k_1, k_2) or different keys k such that $digest(k_1, hash(m \parallel k_2)) = y$. The distribution probability of $digest'()$ is therefore 2^{-b} \square

There is another short-output universal hash function called UHASH16 (and also UHASH32) of Krovetz [3] with 16-bit output and $\epsilon_c \approx 2^{-15}$. The universal hash key length is fixed at around 2^{14} bits or 2KB that is much longer than $B+b+80$ of $digest'()$, but UHASH16 has the advantage of being more efficient than a cryptographic hash function required in $digest'()$.

UHASH16 is the result of a rather non-trivial combination (and tree hash structure) of NH and polynomial universal hashing. First, NH is used to compress

⁷ The concatenation of m and k_2 is hashed to make it much harder for the intruder to search for collision because a large number of bits of the hash input will not be controlled by the intruder.

input messages into strings which are typically many times (e.g. 512 times in UHASH16) shorter than the original input message. Second, the compressed message is hashed with a polynomial universal hash function into a 128-bit string. Finally, the 128-bit string is hashed using an inner-product hash into a 16-bit string. We should point out that there is a trade-off between key length and computational efficiency, i.e. the shorter is the NH key the more computation is required for polynomial hashing in the second stage.

Acknowledgements. Nguyen's work on this paper was supported by a research grant from the US Office of Naval Research. Roscoe's was partially supported by funding from the US Office of Naval Research.

The authors would like to thank Dr. Andrew Ker at Oxford University for his help with statistical analysis of the digest constructions.

Progresses on the security proof of our digest functions were first made when Nguyen visited Professor Bart Preneel and Dr. Frederik Vercauteren at the Computer Security and Industrial Cryptography (COSIC) research group at the Katholieke Universiteit Leuven in September and October 2010. The authors would like to thank them for their time and support as well as Drs Nicky Mouha and Antoon Bosselaers at COSIC for pointing out the relevance of the multiplicative universal hashing scheme of Dietzfelbinger et al.[5] and the literature on hash function implementation and speed measurement for benchmarking.

We also received helpful comments from many anonymous referees as well as had fruitful discussions and technical feedbacks from Professor Serge Vaudenay and Dr. Atefeh Mashatan when Nguyen visited the Security and Cryptography Laboratory (LASEC) at the Swiss Federal Institute of Technologies (EPFL) in February and March 2011. The feedbacks significantly improve the technical quality and presentation of the paper.

References

1. <http://www.cs.ox.ac.uk/publications/publication5935-abstract.html>
2. <http://software.intel.com/en-us/articles/carry-less-multiplication-and-its-usage-for-computing-the-gcm-mode/>
3. Black, J., Halevi, S., Krawczyk, H., Krovetz, T., Rogaway, P.: UMAC: Fast and Secure Message Authentication. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 216–233. Springer, Heidelberg (1999)
4. Carter, J.L., Wegman, M.N.: Universal Classes of Hash Functions. *Journal of Computer and System Sciences* 18, 143–154 (1979)
5. Dietzfelbinger, M., Hagerup, T., Katajainen, J., Penttonen, M.: A reliable randomized algorithm for the closest-pair problem. *Journal Algorithms* 25, 19–51 (1997)
6. Etzel, M., Patel, S., Ramzan, Z.: SQUARE HASH: Fast Message Authentication via Optimized Universal Hash Functions. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 234–251. Springer, Heidelberg (1999)
7. Gehrman, C., Mitchell, C., Nyberg, K.: Manual Authentication for Wireless Devices. *RSA Cryptobytes* 7(1), 29–37 (2004)
8. Halevi, S., Krawczyk, H.: MMH: Software Message Authentication in the Gbit/Second Rates. In: Biham, E. (ed.) FSE 1997. LNCS, vol. 1267, pp. 172–189. Springer, Heidelberg (1997)

9. Handschuh, H., Preneel, B.: Key-Recovery Attacks on Universal Hash Function Based MAC Algorithms. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 144–161. Springer, Heidelberg (2008)
10. Nguyen, L.H. (ed.): Information Technology – Security Techniques – Entity authentication – Part 6: Mechanisms using manual data transfer, ISO/IEC 9798-6 (2010)
11. Krawczyk, H.: New Hash Functions for Message Authentication. In: Guillou, L.C., Quisquater, J.-J. (eds.) EUROCRYPT 1995. LNCS, vol. 921, pp. 301–310. Springer, Heidelberg (1995)
12. Krovetz, T., Rogaway, P.: Fast Universal Hashing with Small Keys and No Preprocessing: The PolyR Construction. In: Won, D. (ed.) ICISC 2000. LNCS, vol. 2015, pp. 73–89. Springer, Heidelberg (2001)
13. Laur, S., Nyberg, K.: Efficient Mutual Data Authentication Using Manually Authenticated Strings. In: Pointcheval, D., Mu, Y., Chen, K. (eds.) CANS 2006. LNCS, vol. 4301, pp. 90–107. Springer, Heidelberg (2006)
14. Lindell, A.Y.: Comparison-Based Key Exchange and the Security of the Numeric Comparison Mode in Bluetooth v2.1. In: Fischlin, M. (ed.) CT-RSA 2009. LNCS, vol. 5473, pp. 66–83. Springer, Heidelberg (2009)
15. Mansour, Y., Nisan, N., Tiwari, P.: The Computational Complexity of Universal Hashing. In: ACM STOC, pp. 235–243 (1990)
16. Mashatan, A., Stinson, D.: Practical Unconditionally Secure Two-channel Message Authentication. *Designs, Codes and Cryptography* 55, 169–188 (2010)
17. Nguyen, L.H., Roscoe, A.W.: Authentication protocols based on low-bandwidth unspoofable channels: A comparative survey. *Journal of Computer Security* 19(1), 139–201 (2011)
18. Nguyen, L.H., Roscoe, A.W.: Efficient group authentication protocol based on human interaction. In: FCS-ARSPA, pp. 9–31 (2006)
19. Nguyen, L.H., Roscoe, A.W.: Authenticating ad-hoc networks by comparison of short digests. *Information and Computation* 206(2-4), 250–271 (2008)
20. Pasini, S., Vaudenay, S.: SAS-Based Authenticated Key Agreement. In: Yung, M., Dodis, Y., Kiayias, A., Malkin, T. (eds.) PKC 2006. LNCS, vol. 3958, pp. 395–409. Springer, Heidelberg (2006)
21. Roscoe, A.W., Nguyen, L.H.: Security in computing networks. World Intellectual Property Organization. Application number: PCT/GB2006/004113. Publication number: WO/2007/052045. Filed on November 03, 2006. Publication date, May 10 (2007)
22. <http://www.aarongifford.com/computers/sha.html>
23. Shamir, A.: SQUASH – A New MAC with Provable Security Properties for Highly Constrained Devices Such as RFID Tags. In: Nyberg, K. (ed.) FSE 2008. LNCS, vol. 5086, pp. 144–157. Springer, Heidelberg (2008)
24. Stinson, D.R.: Universal Hashing and Authentication Codes. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 74–85. Springer, Heidelberg (1992)
25. The performance of UMAC can be found at, <http://fastcrypto.org/umac/>
26. Valkonen, J., Asokan, N., Nyberg, K.: Ad Hoc Security Associations for Groups. In: Buttyán, L., Gligor, V.D., Westhoff, D. (eds.) ESAS 2006. LNCS, vol. 4357, pp. 150–164. Springer, Heidelberg (2006)

Lapin: An Efficient Authentication Protocol Based on Ring-LPN

Stefan Heyse¹, Eike Kiltz¹, Vadim Lyubashevsky^{2,*},
Christof Paar¹, and Krzysztof Pietrzak^{3,**}

¹ Ruhr-Universität Bochum

² INRIA / ENS, Paris

³ IST Austria

Abstract. We propose a new authentication protocol that is provably secure based on a *ring* variant of the learning parity with noise (LPN) problem. The protocol follows the design principle of the LPN-based protocol from Eurocrypt'11 (Kiltz et al.), and like it, is a two round protocol secure against *active* attacks. Moreover, our protocol has small communication complexity and a very small footprint which makes it applicable in scenarios that involve low-cost, resource-constrained devices.

Performance-wise, our protocol is more efficient than previous LPN-based schemes, such as the many variants of the Hopper-Blum (HB) protocol and the aforementioned protocol from Eurocrypt'11. Our implementation results show that it is even comparable to the standard challenge-and-response protocols based on the AES block-cipher. Our basic protocol is roughly 20 times slower than AES, but with the advantage of having 10 times smaller code size. Furthermore, if a few hundred bytes of non-volatile memory are available to allow the storage of some off-line pre-computations, then the online phase of our protocols is only twice as slow as AES.

Keywords: HB protocols, RFID authentication, LPN problem, Ring-LPN problem.

1 Introduction

Lightweight shared-key authentication protocols, in which a tag authenticates itself to a reader, are extensively used in resource-constrained devices such as radio-frequency identification (RFID) tags or smart cards. The straight-forward approach for constructing secure authentications schemes is to use low-level symmetric primitives such as block-ciphers, e.g. AES [DR02]. In their most basic form, the protocols consist of the reader sending a short challenge c and the tag responding with $\text{AES}_K(c)$, where K is the shared secret key. The protocol is secure if AES fulfils a strong, *interactive* security assumption, namely that it behaves like a strong pseudo-random function.

Authentication schemes based on AES have some very appealing features: they are extremely fast, consist of only 2 rounds, and have very small communication complexities. In certain scenarios, however, such as when low-cost and resource-constrained

* Supported in part by the European Research Council.

** Supported by the European Research Council / ERC Starting Grant (259668-PSPC).

devices are involved, the relatively large gate-count and code size used to implement AES may pose a problem. One approach to overcome the restrictions presented by low-weight devices is to construct a low-weight block cipher (e.g. PRESENT [BKL⁺07]), while another approach has been to deviate entirely from block-cipher based constructions and build a *provably-secure* authentication scheme based on the hardness of some mathematical problem. In this work, we concentrate on this second approach.

Ideally, one would like to construct a scheme that incorporates all the beneficial properties of AES-type protocols, while also acquiring the additional provable security and smaller code description characteristics. In the past decade, there have been proposals that achieved some, but not all, of these criteria. Most of these proposals are extensions and variants of the Hopper-Blum (HB) protocol, recently a protocol following a different blueprint has been proposed by Kiltz et al. [KPC⁺11]. Our proposal can be seen as a continuation of this line of research that contains all the advantages enjoyed by LPN-based protocols, while at the same time, getting even closer to enjoying the benefits of AES-type schemes.

OVERVIEW OF OUR RESULTS. In this work we present a new symmetric authentication protocol which (i) is provably-secure against active attacks (as defined in [JW05]) based on the Ring-LPN assumption, a natural variant of the standard LPN (learning parity with noise) assumption; (ii) consists of 2 rounds; (iii) has small communication complexity (approximately 1300 bits); (iv) has efficiency comparable to AES-based challenge-response protocols (depending on the scenario), but with a much smaller code size. To demonstrate the latter we implemented the tag part of our new protocol in a setting of high practical relevance – a low-cost 8-bit microcontroller which is a typical representative of a CPU to be found on lightweight authentication tokens, and compared its performance (code size and running time) with an AES implementation on the same platform.

PREVIOUS WORKS. Hopper and Blum [HB00, HB01] proposed a 2-round authentication protocol that is secure against *passive* adversaries based on the hardness of the LPN problem (we remind the reader of the definition of the LPN problem in Section 1.2). The characteristic feature of this protocol is that it requires very little workload on the part of the tag and the reader. Indeed, both parties only need to compute vector inner products and additions over F_2 , which makes this protocol (thereafter named HB) a good candidate for lightweight applications.

Following this initial work, Juels and Weis constructed a protocol called HB⁺ [JW05] which they proved to be secure against more realistic, so called *active* attacks. Subsequently, Katz et al. [KS06a], [KS06b, KSS10] provided a simpler security proof for HB⁺ as well as showed that it remains secure when executed in parallel. Unlike the HB protocol, however, HB⁺ requires three rounds of communication between tag and reader. From a practical aspect, 2 round authentication protocols are often advantageous over 3 round protocols. They often show a lower latency which is especially pronounced on platforms where the establishment of a communication in every directions is accompanied by a fixed initial delay. An additional drawback of both HB and HB⁺ is that their communication complexity is on the order of hundreds of thousands of bits, which makes them almost entirely impractical for lightweight authentication tokens because of

timing and energy constraints. (The contactless transmission of data on RFIDs or smart cards typically requires considerably more energy than the processing of the same data.)

To remedy the overwhelming communication requirement of HB^+ , Gilbert et al. proposed the three-round $\text{HB}^\#$ protocol [GRS08a]. A particularly practical instantiation of this protocol requires fewer than two thousand bits of communication, but is no longer based on the hardness of the LPN problem. Rather than using independent randomness, the $\text{HB}^\#$ protocol utilized a Toeplitz matrix, and is thus based on a plausible assumption that the LPN problem is still hard in this particular scenario.

A feature that the HB , HB^+ , and $\text{HB}^\#$ protocols have in common is that at some point the reader sends a random string r to the tag, which then must reply with $\langle r, s \rangle + e$, the inner product of r with the secret s plus some small noise e . The recent work of Kiltz et al. [KPC⁺11] broke with this approach, and they were able to construct the first 2-round LPN-based authentication protocol (thereafter named EC11) that is secure against active attacks. In their challenge-response protocol, the reader sends some challenge bit-string c to the tag, who then answers with a noisy inner product of a random r (which the tag chooses itself) and a session-key $K(c)$, where $K(c)$ selects (depending on c) half of the bits from the secret s . Unfortunately, the EC11 protocol still inherits the large communication requirement of HB and HB^+ . Furthermore, since the session key $K(c)$ is computed using bit operations, it does not seem to be possible to securely instantiate EC11 over structured (and hence more compact) objects such as Toeplitz matrices (as used in $\text{HB}^\#$ [GRS08a]).

1.1 Our Contributions

PROTOCOL. In this paper we propose a variant of the EC11 protocol from [KPC⁺11] which uses an “algebraic” derivation of the session key $K(c)$, thereby allowing to be instantiated over a carefully chosen ring $\mathbb{R} = \mathbb{F}_2[X]/(f)$. Our scheme is no longer based on the hardness of LPN, but rather on the hardness of a natural generalization of the problem to rings, which we call Ring-LPN (see Section 3 for the definition of the problem.) The general overview of our protocol is quite simple. Given a challenge c from the reader, the tag answers with $(r, z = r \cdot K(c) + e) \in \mathbb{R} \times \mathbb{R}$, where r is a random ring element, e is a low-weight ring element, and $K(c) = sc + s'$ is the session key that depends on the shared secret key $K = (s, s') \in \mathbb{R}^2$ and the challenge c . The reader accepts if $e' = r \cdot K(c) - z$ is a polynomial of low weight, cf. Figure 1 in Section 4. Compared to the HB and HB^+ protocols, ours has one less round and a dramatically lower communication complexity. Our protocol has essentially the same communication complexity as $\text{HB}^\#$, but still retains the advantage of one fewer round. And compared to the two-round EC11 protocol, ours again has the large savings in the communication complexity. Furthermore, it inherits from EC11 the simple and tight security proof that, unlike three-round protocols, does not use rewinding.

We remark that while our protocol is provably secure against active attacks, we do not have a proof of security against man-in-the-middle ones. Still, as argued in [KSS10], security against active attacks is sufficient for many use scenarios (see also [JW05, KW05, KW06]). We would like to mention that despite man-in-the-middle attacks being outside our “security model”, we think that it is still worthwhile investigating whether such attacks do in fact exist, because it presently seems that all previous

Table 1. Summary of implementation results

Protocol	Time (cycles)		Code size (bytes)
	online	offline	
Ours: reducible f (§5.1)	30,000	82,500	1,356
Ours: irreducible f (§5.2)	21,000	174,000	459
AES-based [LLS09, Tik]	10,121	0	4,644

man-in-the middle attacks against HB-type schemes along the lines of Gilbert et al. [GRS05] and of Ouafi et al. [OOV08] do not apply to our scheme. In Appendix A, however, we do present a man-in-the-middle attack that works in time approximately $n^{1.5} \cdot 2^{\lambda/2}$ (where n is the dimension of the secret and λ is the security parameter) when the adversary can influence on the order of $n^{1.5} \cdot 2^{\lambda/2}$ interactions between the reader and the tag. To resist this attack, one could simply double the security parameter, but we believe that even for $\lambda = 80$ (and $n > 512$, as it is currently set in our scheme) this attack is already impractical because of the extremely large number of interactions that the adversary will have to observe and modify.

IMPLEMENTATION. We demonstrate that our protocol is indeed practical by providing a lightweight implementation of the tag part of the protocol. (The reader is typically not run on a constrained device and therefore we do not consider its performance.) The target platform was an AVR ATmega163 [Atm] based smart card. The ATmega163 is a small 8-bit microcontroller which is a typical representative of a CPU to be found on lightweight authentication tokens. The main metrics we consider are run time and code size. We compare our results with a challenge-response protocol using an AES implementation optimized for the target platform. A major advantage of our protocol is its very small code size. The most compact implementation requires only about 460 bytes of code, which is an improvement by factor of about 10 over AES-based authentication. Given that EEPROM or FLASH memory is often one of the most precious resources on constrained devices, our protocol can be attractive in certain situations. The drawback of our protocol over AES on the target platform is an increase in clock cycles for one round of authentication. However, if we have access to a few hundred bytes of non-volatile data memory, our protocol allows precomputations which make the on-line phase only a factor two or three slower than AES. But even without precomputations, the protocol can still be executed in a few 100 msec, which will be sufficient for many real-world applications, e.g. remote keyless entry systems or authentication for financial transactions. Table 1 gives a summary of the results, see Section 5 for details.

We would like to stress at this point that our protocol is targeting lightweight tags that are equipped with (small) CPUs. For ultra constrained tokens (such as RFIDs in the price range of a few cents targeting the EPC market) which consist nowadays of a small integrated circuit, even compact AES implementations are often considered too costly. (We note that virtually all current commercially available low-end RFIDs do not have any crypto implemented.) However, tokens which use small microcontrollers are far more common, e.g., low-cost smart cards, and they do often require strong authentication. Also, it can be speculated that computational RFIDs such as the WISP [Wik]

will become more common in the future, and hence software-friendly authentication methods that are highly efficient such as the protocol provided here will be needed.

1.2 LPN, Ring-LPN, and Related Problems

The security of our protocols relies on the new Ring Learning Parity with Noise (Ring-LPN) problem which is a natural extension of the standard Learning Parity with Noise (LPN) problem to rings. It can also be seen as a particular instantiation of the Ring-LWE (Learning with Errors over Rings) problem that was recently shown to have a strong connection to lattices [LPR10]. We will now briefly describe and compare these hardness assumptions, and we direct the reader to Section 3 for a formal definition of the Ring-LPN problem.

The decision versions of these problems require us to distinguish between two possible oracles to which we have black-box access. The first oracle has a randomly generated secret vector $s \in \mathbb{F}_2^n$ which it uses to produce its responses. In the LPN problem, each query to the oracle produces a uniformly random matrix¹ $A \in \mathbb{F}_2^{n \times n}$ and a vector $As + e = t \in \mathbb{F}_2^n$ where e is a vector in \mathbb{F}_2^n each of whose entries is an independently generated Bernoulli random variable with probability of 1 being some public parameter τ between 0 and 1/2. The second oracle in the LPN problem outputs a uniformly-random matrix $A \in \mathbb{F}_2^{n \times n}$ and a uniformly random vector $t \in \mathbb{F}_2^n$.

The only difference between LPN and Ring-LPN is in the way the matrix A is generated (both by the first and second oracle). While in the LPN problem, all its entries are uniform and independent, in the Ring-LPN problem, only its first column is generated uniformly at random in \mathbb{F}_2^n . The remaining n columns of A depend on the first column and the underlying ring $R = \mathbb{F}_2[X]/(f(X))$. If we view the first column of A as a polynomial $r \in R$, then the i^{th} column (for $0 \leq i \leq n - 1$) of A is just the vector representation of rX^i in the ring R . Thus when the oracle returns $As + e$, this corresponds to it returning the polynomial $r \cdot s + e$ where the multiplication of polynomials r and s (and the addition of e) is done in the ring R . The Ring-LPN^R assumption states that it is hard to distinguish between the outputs of the first and the second oracle described above. In Section 3, we discuss how the choice of the ring R affects the security of the problem.

While the standard Learning Parity with Noise (LPN) problem has found extensive use as a cryptographic hardness assumption (e.g., [HB01, JW05, GRS08b, GRS08a, ACPS09, KSS10]), we are not aware of any constructions that employed the Ring-LPN problem. There have been some previous works that considered some relatively similar “structured” versions of LPN. The HB[♯] authentication protocol of Gilbert et al. [GRS08a] made the assumption that for a random Toeplitz matrix $S \in \mathbb{F}_2^{m \times n}$, a uniformly random vector $a \in \mathbb{F}_2^n$, and a vector $e \in \mathbb{F}_2^m$ whose coefficients are distributed as Ber_τ , the output $(a, Sa + e)$ is computationally indistinguishable from (a, t) where t is uniform over \mathbb{F}_2^m .

Another related work, as mentioned above, is the recent result of Lyubashevsky et al. [LPR10], where it is shown that solving the decisional Ring-LWE (Learning with

¹ In the more common description of the LPN problem, each query to the oracle produces one random sample in \mathbb{F}_2^n . For comparing LPN to Ring-LPN, however, it is helpful to consider the oracle as returning a matrix of n random independent samples on each query.

Errors over Rings) problem is as hard as quantumly solving the worst case instances of the shortest vector problem in *ideal* lattices. The Ring-LWE problem is quite similar to Ring-LPN, with the main difference being that the ring R is defined as $F_q[X]/(f(X))$ where $f(X)$ is a cyclotomic polynomial and q is a prime such that $f(X)$ splits completely into $\deg(f(X))$ distinct factors over F_q .

Unfortunately, the security proof of our authentication scheme does not allow us to use a polynomial $f(X)$ that splits into low-degree factors, and so we cannot base our scheme on lattice problems. For a similar reason (see the proof of our scheme in Section 4 for more details), we cannot use samples that come from a Toeplitz matrix as in [GRS08a]. Nevertheless, we believe that the Ring-LPN assumption is very natural and will find further cryptographic applications, especially for constructions of schemes for low-cost devices.

2 Definitions

2.1 Rings and Polynomials

For a polynomial $f(X)$ over F_2 , we will often omit the indeterminate X and simply write f . The degree of f is denoted by $\deg(f)$. For two polynomials a, f in $F_2[X]$, $a \bmod f$ is defined to be the unique polynomial r of degree less than $\deg(f)$ such that $a = fg + r$ for some polynomial $g \in F_2[X]$. The elements of the ring $F_2[X]/(f)$ will be represented by polynomials in $F_2[X]$ of maximum degree $\deg(f) - 1$. In this paper, we will only be considering rings $R = F_2[X]/(f)$ where the polynomial f factors into *distinct* irreducible factors over F_2 . For an element a in the ring $F_2[X]/(f)$, we will denote by \hat{a} , the CRT (Chinese Remainder Theorem) representation of a with respect to the factors of f . In other words, if $f = f_1 \dots f_m$ where all f_i are irreducible, then

$$\hat{a} \doteq (a \bmod f_1, \dots, a \bmod f_m).$$

If f is itself an irreducible polynomial, then $\hat{a} = a$. Note that an element $\hat{a} \in R$ has a multiplicative inverse iff, for all $1 \leq i \leq m$, $a \not\equiv 0 \pmod{f_i}$. We denote by R^* the set of elements in R that have a multiplicative inverse.

2.2 Distributions

For a distribution D over some domain, we write $r \stackrel{\$}{\leftarrow} D$ to denote that r is chosen according to the distribution D . For a domain Y , we write $U(Y)$ to denote the uniform distribution over Y . Let Ber_τ be the Bernoulli distribution over F_2 with parameter (bias) $\tau \in]0, 1/2[$ (i.e., $\Pr[x = 1] = \tau$ if $x \leftarrow \text{Ber}_\tau$). For a polynomial ring $R = F_2[X]/(f)$, the distribution Ber_τ^R denotes the distribution over the polynomials of R , where each of the coefficients of the polynomial is drawn independently from Ber_τ . For a ring R and a polynomial $s \in R$, we write $\Lambda_\tau^{R,s}$ to be the distribution over $R \times R$ whose samples are obtained by choosing a polynomial $r \stackrel{\$}{\leftarrow} U(R)$ and another polynomial $e \stackrel{\$}{\leftarrow} \text{Ber}_\tau^R$, and outputting $(r, rs + e)$.

2.3 Authentication Protocols

An authentication protocol Π is an interactive protocol executed between a Tag \mathcal{T} and a reader \mathcal{R} , both PPT algorithms. Both hold a secret x (generated using a key-generation algorithm KG executed on the security parameter λ in unary) that has been shared in an initial phase. After the execution of the authentication protocol, \mathcal{R} outputs either accept or reject. We say that the protocol has completeness error ε_c if for all $\lambda \in \mathbb{N}$, all secret keys x generated by $\text{KG}(1^\lambda)$, the honestly executed protocol returns reject with probability at most ε_c . We now define different security notions of an authentication protocol.

PASSIVE ATTACKS. An authentication protocol is secure against *passive* attacks, if there exists no PPT adversary \mathcal{A} that can make the reader \mathcal{R} return accept with non-negligible probability after (passively) observing any number of interactions between reader and tag.

ACTIVE ATTACKS. A stronger notion for authentication protocols is security against *active* attacks. Here the adversary \mathcal{A} runs in two stages. First, she can interact with the honest tag a polynomial number of times (with concurrent executions allowed). In the second phase \mathcal{A} interacts with the reader only, and wins if the reader returns accept. Here we only give the adversary one shot to convince the verifier.² An authentication protocol is (t, q, ε) -secure against active adversaries if every PPT \mathcal{A} , running in time at most t and making q queries to the honest reader, has probability at most ε to win the above game.

3 Ring-LPN and Its Hardness

The decisional Ring-LPN^R (Ring Learning Parity with Noise in ring R) assumption, formally defined below, states that it is hard to distinguish uniformly random samples in $\mathbb{R} \times \mathbb{R}$ from those sampled from $A_{\tau}^{\mathbb{R},s}$ for a uniformly chosen $s \in \mathbb{R}$.

Definition 1 (Ring-LPN^R). *The (decisional) Ring-LPN _{τ} ^R problem is (t, q, ε) -hard if for every distinguisher \mathcal{D} running in time t and making q queries,*

$$\left| \Pr \left[s \stackrel{s}{\leftarrow} \mathbb{R} : \mathcal{D}^{A_{\tau}^{\mathbb{R},s}} = 1 \right] - \Pr \left[\mathcal{D}^{U(\mathbb{R} \times \mathbb{R})} = 1 \right] \right| \leq \varepsilon.$$

3.1 Hardness of LPN and Ring-LPN

One can attempt to solve Ring-LPN using standard algorithms for LPN, or by specialized algorithms that possibly take advantage of Ring-LPN's additional structure. Some work towards constructing the latter type of algorithm has recently been done by Hanrot et al. [HLPS11], who show that in certain cases, the algebraic structure of the Ring-LPN and Ring-LWE problems makes them vulnerable to certain attacks. These attacks essentially utilize a particular relationship between the factorization of the polynomial $f(X)$ and the distribution of the noise.

² By using a hybrid argument one can show that this implies security even if the adversary can interact in $k \geq 1$ independent instances concurrently (and wins if the verifier accepts in at least one instance). The use of the hybrid argument loses a factor of k in the security reduction.

Ring-LPN with an Irreducible $f(X)$ When $f(X)$ is irreducible over F_2 , the ring $F_2[X]/(f)$ is a field. For such rings, the algorithm of Hanrot et al. does not apply, and we do not know of any other algorithm that takes advantage of the added algebraic structure of this particular Ring-LPN instance. Thus to the best of our knowledge, the most efficient algorithms for solving this problem are the same ones that are used to solve LPN, which we will now very briefly recount.

The computational complexity of the LPN problem depends on the length of the secret n and the noise distribution Ber_τ . Intuitively, the larger the n and the closer τ is to $1/2$, the harder the problem becomes. Usually the LPN problem is considered for constant values of τ somewhere between 0.05 and 0.25. For such constant τ , the fastest asymptotic algorithm for the LPN problem, due to Blum et al. [BKW03], takes time $2^{\Omega(n/\log n)}$ and requires approximately $2^{\Omega(n/\log n)}$ samples from the LPN oracle. If one has access to fewer samples, then the algorithm will perform somewhat worse. For example, if one limits the number of samples to only polynomially-many, then the algorithm has an asymptotic complexity of $2^{\Omega(n/\log \log n)}$ [Lyu05]. In our scenario, the number of samples available to the adversary is limited to n times the number of executions of the authentication protocol, and so it is reasonable to assume that the adversary will be somewhat limited in the number of samples he is able to obtain (perhaps at most 2^{40} samples), which should make our protocols harder to break than solving the Ring-LPN problem. Leveil and Fouque [LF06] made some optimizations to the algorithm of Blum et al. and analyzed its precise complexity. To the best of our knowledge, their algorithm is currently the most efficient one and we will refer to their results when analyzing the security of our instantiations.

In Section 5, we base our scheme on the hardness of the Ring-LPN^R problem where the ring is $R = F_2[X]/(X^{532} + X + 1)$ and $\tau = 1/8$. According to the analysis of [LF06], the fastest algorithm to solve an LPN problem of dimension 512 with $\tau = 1/8$ would require 2^{77} memory (and thus at least that much time) to solve when given access to approximately as many samples (see [LF06, Section 5.1]). Since our dimension is somewhat larger and the number of samples will be limited in practice, it is reasonable to assume that this instantiation has 80-bit security.

Note: After the appearance of this paper in the pre-proceedings of FSE, Tanja Lange pointed us to an unpublished paper of Paul Kirchner [Kir11]. In section 4.3.2 of that paper, the author uses the fact that the secret can come from the same distribution as the error³, rather than being completely uniform, to achieve a slightly better running time for solving certain instances of the LPN problem using generalized birthday attacks. We have not yet studied the paper in detail to see how this improvement can be used in conjunction with the work of [LF06], but it is conceivable that this improved algorithm for LPN along with some additional techniques [Lan12], would require a slight increase in the parameters of our scheme.

Ring-LPN with a Reducible $f(X)$. For efficiency purposes, it is sometimes useful to consider using a polynomial $f(X)$ that is not irreducible over F_2 . This will allow us to use the CRT representation of the elements of $F_2[X]/(f)$ to perform multiplications,

³ For readers familiar with the lattice literature, this is analogous to the result of Applebaum et al. [ACPS09, Lemma 2]

which in practice turns out to be more efficient. Ideally, we would like the polynomial f to split into as many small-degree polynomials f_i as possible, but there are some constraints that are placed on the factorization of f both by the security proof, and the possible weaknesses that a splittable polynomial introduces into the Ring-LPN problem.

If the polynomial f splits into $f = \prod_{i=1}^m f_i$, then it may be possible to try and solve the Ring-LPN problem modulo some f_i rather than modulo f . Since the degree of f_i is smaller than the degree of f , the resulting Ring-LPN problem may end up being easier. In particular, when we receive a sample $(r, rs + e)$ from the distribution $\Lambda_r^{\mathbb{R},s}$, we can rewrite it in CRT form as

$$(\widehat{r}, \widehat{rs + e}) = ((r \bmod f_1, rs + e \bmod f_1), \dots, (r \bmod f_m, rs + e \bmod f_m)),$$

and thus for every f_i , we have a sample

$$(r \bmod f_i, (r \bmod f_i)(s \bmod f_i) + e \bmod f_i),$$

where all the operations are in the ring (or field) $\mathbb{F}_2[X]/(f_i)$. Thus solving the (decision) Ring-LPN problem in $\mathbb{F}_2[X]/(f)$ reduces to solving the problem in $\mathbb{F}_2[X]/(f_i)$. The latter problem is in a smaller dimension, since $\deg(s) > \deg(s \bmod f_i)$, but the error distribution of $(e \bmod f_i)$ is quite different than that of e . While each coefficient of e is distributed independently as Ber_τ , each coefficient of $(e \bmod f_i)$ is distributed as the distribution of a sum of certain coefficients of e , and therefore the new error is larger.⁴ Exactly which coefficients of e , and more importantly, how many of them, combine to form every particular coefficient of e' depends on the polynomial f_i . For example, if

$$f(X) = (X^3 + X + 1)(X^3 + X^2 + 1)$$

and $e = \sum_{i=0}^5 e_i X^i$, then,

$$e' = e \bmod (X^3 + X + 1) = (e_0 + e_3 + e_5) + (e_1 + e_3 + e_4 + e_5)X + (e_2 + e_4 + e_5)X^2,$$

and thus every coefficient of the error e' is comprised of at least 3 coefficients of the error vector e , and thus $\tau' > \frac{1}{2} - \frac{(1-2\tau)^3}{2}$.

In our instantiation of the scheme with a reducible $f(X)$ in Section 5, we used the $f(X)$ such that it factors into f_i 's that make the operations in CRT form relatively fast, while making sure that the resulting Ring-LPN problem modulo each f_i is still around 2^{80} -hard.

4 Authentication Protocol

In this section we describe our new 2-round authentication protocol and prove its active security under the hardness of the Ring-LPN problem. Detailed implementation details will be given in Section 5.

⁴ If we have k elements $e_1, \dots, e_k \stackrel{\$}{\leftarrow} \text{Ber}_\tau$, then a simple calculation shows that the element $e' = e_1 + \dots + e_k$ is distributed as $\text{Ber}_{\tau'}$ where $\tau' = \frac{1}{2} - \frac{(1-2\tau)^k}{2}$.




Fig. 1. Two-round authentication protocol with active security from the Ring-LPN^R assumption

4.1 The Protocol

Our authentication protocol is defined over the ring $R = \mathbb{F}_2[X]/(f)$ and involves a “suitable” mapping $\pi : \{0, 1\}^\lambda \rightarrow R$. We call π *suitable* for ring R if for all $c, c' \in \{0, 1\}^\lambda$, $\pi(c) - \pi(c') \in R \setminus R^*$ iff $c = c'$. We will discuss the necessity and existence of such mappings after the proof of Theorem 1

- **Public parameters.** The authentication protocol has the following public parameters, where τ, τ' are constants and n depend on the security parameter λ .
 R, n ring $R = \mathbb{F}_2[X]/(f)$, $\deg(f) = n$
 $\pi : \{0, 1\}^\lambda \rightarrow R$ mapping
 $\tau \in (0, 1/2)$ parameter of Bernoulli distribution
 $\tau' \in (\tau, 1/2)$ acceptance threshold
- **Key Generation.** Algorithm $\text{KG}(1^\lambda)$ samples $s, s' \xleftarrow{\$} R$ and returns s, s' as the secret key.
- **Authentication Protocol.** The Reader \mathcal{R} and the Tag \mathcal{T} share secret value $s, s' \in R$. To be authenticated by a Reader, the Tag and the Reader execute the authentication protocol from Figure 1.

4.2 Analysis

For our analysis we define for $x, y \in]0, 1[$ the following constant:

$$c(x, y) := \left(\frac{x}{y}\right)^x \left(\frac{1-x}{1-y}\right)^{1-x}.$$

We now state that our protocol is secure against active adversaries. Recall that active adversaries can arbitrarily interact with a Tag oracle in the first phase and tries to impersonate the Reader in the 2nd phase.

Theorem 1. *If ring mapping π is suitable for ring R and the Ring-LPN $_R$ problem is (t, q, ε) -hard then the authentication protocol from Figure 1 is (t', q, ε') -secure against active adversaries, where*

$$t' = t - q \cdot \exp(R) \quad \varepsilon' = \varepsilon + q \cdot 2^{-\lambda} + c(\tau', 1/2)^{-n} \quad (4.1)$$

and $\exp(R)$ is the time to perform $O(1)$ exponentiations in R . Furthermore, the protocol has completeness error $\varepsilon_c(\tau, \tau', n) \approx c(\tau', \tau)^{-n}$.

Proof. The completeness error $\varepsilon_c(\tau, \tau', n)$ is (an upper bound on) the probability that an honestly generated Tag gets rejected. In our protocol this is exactly the case when the error e' has weight $\geq n \cdot \tau'$, i.e.

$$\varepsilon_c(\tau, \tau', n) = \Pr[\mathbf{wt}(e') > n \cdot \tau' : e \stackrel{\$}{\leftarrow} \text{Ber}_\tau^R]$$

Levieil and Fouque [LF06] show that one can approximate this probability as $\varepsilon_c \approx c(\tau', \tau)^{-n}$.

To prove the security of the protocol against active attacks we proceed in sequences of games. Game $_0$ is the security experiment describing an active attack on our scheme by an adversary \mathcal{A} making q queries and running in time t' , i.e.

- Sample the secret key $s, s' \stackrel{\$}{\leftarrow} R$.
- (1st phase of active attack) \mathcal{A} queries the tag \mathcal{T} on $c \in \{0, 1\}^\lambda$ and receives (r, z) computed as illustrated in Figure 1.
- (2nd phase of active attack) \mathcal{A} gets a random challenge $c^* \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda$ and outputs (r, z) . \mathcal{A} wins if the reader \mathcal{R} accepts, i.e. $\mathbf{wt}(z - r \cdot (s \cdot \pi(c^*) + s')) \leq n \cdot \tau'$.

By definition we have $\Pr[\mathcal{A} \text{ wins in Game}_0] \leq \varepsilon'$.

Game $_1$ is as Game $_0$, except that all the values (r, z) returned by the Tag oracle in the first phase (in return to a query $c \in \{0, 1\}^\lambda$) are uniform random elements $(r, z) \in R^2$. We now show that if \mathcal{A} is successful against Game $_0$, then it will also be successful against Game $_1$.

Claim. $|\Pr[\mathcal{A} \text{ wins in Game}_1] - \Pr[\mathcal{A} \text{ wins in Game}_0]| \leq \varepsilon + q \cdot 2^{-\lambda}$

To prove this claim, we construct an adversary \mathcal{D} (distinguisher) against the Ring-LPN problem which runs in time $t = t' + \exp(R)$ and has advantage

$$\varepsilon \geq |\Pr[\mathcal{A} \text{ wins in Game}_1] - \Pr[\mathcal{A} \text{ wins in Game}_0]| - q \cdot 2^{-\lambda}$$

\mathcal{D} has access to a Ring-LPN oracle \mathcal{O} and has to distinguish between $\mathcal{O} = A_\tau^{R,s}$ for some secret $s \in R$ and $\mathcal{O} = U(R \times R)$.

- \mathcal{D} picks a random challenge $c^* \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda$ and $a \stackrel{\$}{\leftarrow} R$. Next, it runs \mathcal{A} and simulates its view with the unknown secret s, s' , where $s \in R$ comes from the oracle \mathcal{O} and s' is implicitly defined as $s' := -\pi(c^*) \cdot s + a \in R$.
- In the 1st phase, \mathcal{A} can make q (polynomial many) queries to the Tag oracle. On query $c \in \{0, 1\}^\lambda$ to the Tag oracle, \mathcal{D} proceeds as follows. If $\pi(c) - \pi(c^*) \notin R^*$, then abort. Otherwise, \mathcal{D} queries its oracle $\mathcal{O}()$ to obtain $(r', z') \in R^2$. Finally, \mathcal{D} returns (r, z) to \mathcal{A} , where

$$r := r' \cdot (\pi(c) - \pi(c^*))^{-1}, \quad z := z' + ra. \quad (4.2)$$

- In the 2nd phase, \mathcal{D} uses $c^* \in \{0, 1\}^\lambda$ to challenge \mathcal{A} . On answer (r, z) , \mathcal{D} returns 0 to the Ring-LPN game if $\mathbf{wt}(z - r \cdot a) > n \cdot \tau'$ or $r \notin \mathbb{R}^*$, and 1 otherwise. Note that $s\pi(c^*) + s' = (\pi(c^*) - \pi(c^*))s + a = a$ and hence the above check correctly simulates the output of a reader with the simulated secret s, s' .

Note that the running time of \mathcal{D} is that of \mathcal{A} plus $O(q)$ exponentiations in \mathbb{R} .

Let bad be the event that for at least one query c made by \mathcal{A} to the Tag oracle, we have that $\pi(c) - \pi(c^*) \notin \mathbb{R}^*$. Since c^* is uniform random in \mathbb{R} and hidden from \mathcal{A} 's view in the first phase we have by the union bound over the q queries

$$\begin{aligned} \Pr[\text{bad}] &\leq q \cdot \Pr_{c^* \in \{0,1\}^\lambda} [\pi(c) - \pi(c^*) \in \mathbb{R} \setminus \mathbb{R}^*] \\ &= q \cdot 2^{-\lambda}. \end{aligned} \quad (4.3)$$

The latter inequality holds because π is suitable for \mathbb{R} .

Let us now assume bad does not happen. If $\mathcal{O} = \Lambda_r^{\mathbb{R}, s}$ is the real oracle (i.e., it returns (r', z') with $z' = r's + e$) then by the definition of (r, z) from (4.2),

$$z = (r's + e) + ra = r(\pi(c)s - \pi(c^*)s + a) + e = r(s\pi(c) + s') + e.$$

Hence the simulation perfectly simulates \mathcal{A} 's view in Game_0 . If $\mathcal{O} = U(\mathbb{R} \times \mathbb{R})$ is the random oracle then (r, z) are uniformly distributed, as in Game_1 . That concludes the proof of Claim 4.2.

We next upper bound the probability that \mathcal{A} can be successful in Game_1 . This bound will be information theoretic and even holds if \mathcal{A} is computationally unbounded and can make an unbounded number of queries in the 1st phase. To this end we introduce the minimal soundness error, ε_{ms} , which is an upper bound on the probability that a tag (r, z) chosen independently of the secret key is valid, i.e.

$$\varepsilon_{\text{ms}}(\tau', n) := \max_{(z, r) \in \mathbb{R} \times \mathbb{R}^*} \Pr_{\substack{s, s' \xleftarrow{\$} \mathbb{R} \\ e' \xleftarrow{\$} \mathbb{R}}} [\mathbf{wt}(z - r \cdot (s \cdot \pi(c^*) + s')) \leq n\tau']$$

As $r \in \mathbb{R}^*$ and $s' \in \mathbb{R}$ is uniform, also $e' = z - r \cdot (s \cdot \pi(c^*) + s')$ is uniform, thus ε_{ms} is simply

$$\varepsilon_{\text{ms}}(\tau', n) := \Pr_{e' \xleftarrow{\$} \mathbb{R}} [\mathbf{wt}(e') \leq n\tau']$$

Again, it was shown in [LF06] that this probability can be approximated as

$$\varepsilon_{\text{ms}}(\tau', n) \approx c(\tau', 1/2)^{-n}. \quad (4.4)$$

Clearly, ε_{ms} is a trivial lower bound on the advantage of \mathcal{A} in forging a valid tag, by the following claim in Game_1 one cannot do any better than this.

Claim. $\Pr[\mathcal{A} \text{ wins in } \text{Game}_1] = \varepsilon_{\text{ms}}(\tau', n)$

To see that this claim holds one must just observe that the answers \mathcal{A} gets in the first phase of the active attack in Game_1 are independent of the secret s, s' . Hence \mathcal{A} 's advantage is $\varepsilon_{\text{ms}}(\tau', n)$ by definition.

Claims 4.2 and 4.2 imply (4.1) and conclude the proof of Theorem 1.

We require the mapping $\pi : \{0, 1\}^\lambda \rightarrow R$ used in the protocol to be *suitable* for R , i.e. for all $c, c' \in \{0, 1\}^\lambda$, $\pi(c) - \pi(c') \in R \setminus R^*$ iff $c = c'$. In Section 5 we describe efficient suitable maps for any $R = F_2[X]/(f)$ where f has no factor of degree $\leq \lambda$. This condition is necessary, as no suitable mapping exists if f has a factor f_i of degree $\leq \lambda$: in this case, by the pigeonhole principle, there exist distinct $c, c' \in \{0, 1\}^\lambda$ such that $\pi(c) = \pi(c') \pmod{f_i}$, and thus $\pi(c) - \pi(c') \in R \setminus R^*$.

We stress that for our security proof we need π to be suitable for R , since otherwise (4.3) is no longer guaranteed to hold. It is an interesting question if this is inherent, or if the security of our protocol can be reduced to the Ring-LPN^R problem for arbitrary rings $R = F_2[X]/(f)$, or even $R = F_q[X]/(f)$ (This is interesting since, if f has factors of degree $\ll \lambda$, the protocol could be implemented more efficiently and even become based on the worst-case hardness of lattice problems). Similarly, it is unclear how to prove security of our protocol instantiated with Toeplitz matrices.

5 Implementation

There are two objectives that we pursue with the implementation of our protocol. First, we will show that the protocol is in fact practical with concrete parameters, even on extremely constrained CPUs. Second, we investigate possible application scenarios where the protocol might have additional advantages. From a practical point of view, we are particularly interested in comparing our protocol to classical symmetric challenge-response schemes employing AES. Possible advantages of the protocol at hand are (i) the security properties and (ii) improved implementation properties. With respect to the former aspect, our protocol has the obvious advantage of being provably secure under a reasonable and static hardness assumption. Even though AES is arguably the most trusted symmetric cipher, it is “merely” computationally secure with respect to known attacks.

In order to investigate implementation properties, constrained microprocessors are particularly relevant. We chose an 8-bit AVR ATmega163 [Atm] based smartcard, which is widely used in myriads of embedded applications. It can be viewed as a typical representative of a CPU used in tokens that are in need for an authentication protocol, e.g., computational RFID tags or (contactless) smart cards. The main metrics we consider for the implementation are run-time and code size. We note at this point that in many lightweight crypto applications, code size is the most precious resource once the run-time constraints are fulfilled. This is due to the fact that EEPROM or flash memory is often heavily constrained. For instance, the WISP, a computational RFID tag, has only 8 kBytes of program memory [Wik, MSP].

We implemented two variants of the protocol described in Section 4. The first variant uses a ring $R = F_2[X]/(f)$, where f splits into five irreducible polynomials; the second variant uses a field, i.e., f is irreducible. For both implementations, we chose parameters which provide a security level of $\lambda = 80$ bits, i.e., the parameters are chosen such that ε' in (4.1) is bounded by 2^{-80} and the completeness ε_c is bounded by 2^{-40} . This security level is appropriate for the lightweight applications which we are targeting.

5.1 Implementation with a Reducible Polynomial

From an implementation standpoint, the case of reducible polynomial is interesting since one can take advantage of arithmetic based on the Chinese Remainder Theorem.

PARAMETERS. To define the ring $R = \mathbb{F}_2[X]/(f)$, we chose the reducible polynomial f to be the product of the $m = 5$ irreducible pentanomials specified by the following powers with non-zero coefficients: $(127, 8, 7, 3, 0)$, $(126, 9, 6, 5, 0)$, $(125, 9, 7, 4, 0)$, $(122, 7, 4, 3, 0)$, $(121, 8, 5, 1, 0)$ ⁵. Hence f is a polynomial of degree $n = 621$. We chose $\tau = 1/6$ and $\tau' = .29$ to obtain minimal soundness error $\varepsilon_{\text{ms}} \approx c(\tau', 1/2)^{-n} \leq 2^{-82}$ and completeness error $\varepsilon_c \leq 2^{-42}$. From the discussion of Section 3 the best known attack on Ring-LPN $_{\tau}^R$ with the above parameters has complexity $> 2^{80}$. The mapping $\pi : \{0, 1\}^{80} \rightarrow R$ is defined as follows. On input $c \in \{0, 1\}^{80}$, for each $1 \leq i \leq 5$, pad $c \in \{0, 1\}^{80}$ with $\deg(f_i) - 80$ zeros and view the result as coefficients of an element $v_i \in \mathbb{F}_2[X]/(f_i)$. This defines $\pi(c) = (v_1, \dots, v_5)$ in CRT representation. Note that, for fixed $c, c^* \in \{0, 1\}^{80}$, we have that $\pi(c) - \pi(c^*) \in R \setminus R^*$ iff $c = c^*$ and hence π is *suitable* for R .

IMPLEMENTATION DETAILS. The main operations are multiplications and additions of polynomials that are represented by 16 bytes. We view the CRT-based multiplication in three stages. In the first stage, the operands are reduced modulo each of the five irreducible polynomials. This part has a low computational complexity. Note that only the error e has to be chosen in the ring and afterwards transformed to CRT representation. It is possible to save the secret key (s, s') and to generate r directly in the CRT representation. This is not possible for e because e has to come from Ber_{τ}^R . In the second stage, one multiplication in each of the finite fields defined by the five pentanomials has to be performed. We used the right-to-left comb multiplication algorithm from [HMV03]. For the multiplication with $\pi(c)$ we exploit the fact that only the first 80 coefficients can be non-zero. Hence we wrote one function for *normal* multiplication and one for *sparse* multiplication. The latter is more than twice as fast as the former. The subsequent reduction takes care of the special properties of the pentanomials, thus code reuse is not possible for the different fields. The third stage, constructing the product polynomial in the ring, is shifted to the prover (RFID reader) which normally has more computational power than the tag \mathcal{T} . Hence the response (r, z) is sent in CRT form to the reader. If non-volatile storage — in our case we need $2 \cdot 5 \cdot 16 = 160$ bytes — is available we can heavily reduce the response time of the tag. At an arbitrary point in time, choose e and r according to their distribution and precompute $tmp_1 = r \cdot s$ and $tmp_2 = r \cdot s' + e$. When a challenge c is received afterwards, tag \mathcal{T} only has to compute $z = tmp_1 \cdot \pi(c) + tmp_2$. Because $\pi(c)$ is sparse, the tag can use the *sparse* multiplication and response very quickly. The results of the implementation are shown in Table 2 in Section 5.3. Note that all multiplication timings given already include the necessary reductions and addition of a value according to Figure 1.

5.2 Implementation with an Irreducible Polynomial

PARAMETERS. To define the field $F = \mathbb{F}_2[X]/(f)$, we chose the irreducible trinomial $f(X) = X^{532} + X + 1$ of degree $n = 532$. We chose $\tau = 1/8$ and $\tau' = .27$ to

⁵ $(127, 8, 7, 3, 0)$ refers to the polynomial $X^{127} + X^8 + X^7 + X^3 + 1$.

obtain minimal soundness error $\varepsilon_{\text{ms}} \approx c(\tau', 1/2)^{-n} \leq 2^{-80}$ and completeness error $\varepsilon_c \approx 2^{-55}$. From the discussion in Section 3 the best known attack on Ring-LPN $_{\tau}^F$ with the above parameters has complexity $> 2^{80}$. The mapping $\pi : \{0, 1\}^{80} \rightarrow F$ is defined as follows. View $c \in \{0, 1\}^{80}$ as $c = (c_1, \dots, c_{16})$ where c_i is a number between 1 and 32. Define the coefficients of the polynomial $v = \pi(c) \in F$ as zero except all positions i of the form $i = 16 \cdot (j - 1) + c_j$, for some $j = 1, \dots, 16$. Hence $\pi(c)$ is sparse, i.e., it has exactly 16 non-zero coefficients. Since π is injective and F is a field, the mapping π is suitable for F .

IMPLEMENTATION DETAILS. The main operation for the protocol is now a 67-byte multiplication. Again we used the right-to-left comb multiplication algorithm from [HMOV03] and an optimized reduction algorithm. Like in the reducible case, the tag can do similar precomputations if $2 \cdot 67 = 134$ bytes non-volatile storage are available. Because of the special type of the mapping $v = \pi(c)$, the gain of the *sparse* multiplication is even larger than in the reducible case. Here we are a factor of 7 faster, making the response time with precomputations faster, although the field is larger. The results are shown in Table 3 in Section 5.3.

5.3 Implementation Results

All results presented in this section consider only the clock cycles of the actual arithmetic functions. The communication overhead and the generation of random bytes is excluded because they occur in every authentication scheme, independent of the underlying cryptographic functions. The time for building e from Ber_{τ}^R out of the random bytes and converting it to CRT form is included in *Overhead*. Table 2 and Table 3 shows the results for the ring based and field based variant, respectively.

Table 2. Results for the ring based variant w/o precomputation

Aspect	time code size	
	in cycles	in bytes
Overhead	17,500	264
Mul	$5 \times 13,000$	164
sparse Mul	$5 \times 6,000$	170
total	112,500	1356

The overall code size is not the sum of the other values because, as mentioned before, the same multiplication code is used for all *normal* and *sparse* multiplications, respectively, while the reduction code is different for every field (≈ 134 byte each). The same code for reduction is used independently of the type of the multiplication for the same field. If precomputation is acceptable, the tag can answer the challenge after approximately 30,000 clock cycles, which corresponds to a 15 msec if the CPU is clocked at 2 MHz.

For the field-based protocol, the overall performance is slower due to the large operands used in the multiplication routine. But due to the special mapping $v = \pi(c)$,

Table 3. Results for the field based variant w/o precomputation

Aspect	time	code size
	in cycles	in bytes
Overhead	3,000	150
Mul	150,000	161
sparse Mul	21,000	148
total	174,000	459

here the tag can do a sparse multiplications in only 21,000 clocks cycles. This allows the tag to respond in 10.5 msec at 2 MHz clock rate if non-volatile storage is available.

As mentioned in the introduction, we want to compare our scheme with a conventional challenge-response authentication protocol based on AES. The tag's main operation in this case is one AES encryption. The implementation in [LLS09] states 8,980 clock cycles for one encryption on a similar platform, but unfortunately no code size is given; [Tik] reports 10121 cycles per encryption and a code size of 4644 bytes.⁶ In comparison with these highly optimized AES implementations, our scheme is around eleven times slower when using the ring based variant without precomputations. If non-volatile storage allows precomputations, the ring based variant is only three times slower than AES. But the code size is by a factor of two to three smaller, making it attractive for Flash constrained devices. The field based variant without precomputations is 17 to 19 times slower than AES, but with precomputations it is only twice as slow as AES, while only consuming one tenths of the code size. From a practical point of view, it is important to note that even our slowest implementation is executed in less than 100 msec if the CPU is clocked at 2 MHz. This response time is sufficient in many application scenarios. (For authentications involving humans, a delay of 1 sec is often considered acceptable.)

The performance drawback compared to AES is not surprising, but it is considerably less dramatic compared to asymmetric schemes like RSA or ECC [GPW⁺04]. But exploiting the special structure of the multiplications in our scheme and using only a small amount of non-volatile data memory provides a response time in the same order of magnitude as AES, while keeping the code size much smaller.

6 Conclusions and Open Problems

We proposed a new [KPC⁺11]-like authentication protocol with provable security against active attacks based on the Ring-LPN assumption, consisting of only two rounds, and having small communication complexity. Furthermore, our implementations on an 8-bit AVR ATmega163 based smartcard demonstrated that it has very small code size and its efficiency can be of the same order as traditional AES-based authentication protocols. Overall, we think that its features make it very applicable in scenarios that involve low-cost, resource-constrained devices.

⁶ An internet source [Poe] claims to encrypt in 3126 cycles with code size of 3098 bytes but since this is unpublished material we do not consider it in our comparison.

Our protocol cannot be proved secure against man-in-the-middle (MIM) attacks, but using a recent transformation from [DKPW12] we can get a MIM secure scheme with small extra cost (one application of a universal hash function.) Still, finding a more direct construction which achieves MIM security (or proving that the current protocol already has this property) but doesn't require any hashing remains an interesting open problem.

We believe that the Ring-LPN assumption is very natural and will find further cryptographic applications, especially for constructions of schemes for low-cost devices. In particular, we think that if the LPN-based line of research is to lead to a practical protocol in the future, then the security of this protocol will be based on a hardness assumption with some "extra algebraic structure", such as Ring-LPN in this work, or LPN with Toeplitz matrices in the work of Gilbert et al. [GRS08a]. More research, however, needs to be done on understanding these problems and their computational complexity. In terms of Ring-LPN, it would be particularly interesting to find out whether there exists an equivalence between the decision and the search versions of the problem similar to the reductions that exist for LPN [BFKL93, Reg09, KS06a] and Ring-LWE [LPR10].

Acknowledgements. We would like to thank the anonymous referees of this conference and those of the *ECRYPT Workshop on Lightweight Cryptography* for very useful comments, and in particular for the suggestion that the scheme is somewhat vulnerable to a man-in-the-middle attack whenever an adversary observes two reader challenges that are the same. We hope that the attack we described in Appendix A corresponds to what the reviewer had in mind. We also thank Tanja Lange for pointing us to the paper of [Kir11] and for discussions of some of her recent work. This work was partially supported by the European Research Council.

References

- [ACPS09] Applebaum, B., Cash, D., Peikert, C., Sahai, A.: Fast Cryptographic Primitives and Circular-Secure Encryption Based on Hard Learning Problems. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 595–618. Springer, Heidelberg (2009)
- [Atm] Atmel, ATmega163 datasheet, www.atmel.com/atmel/acrobat/doc1142.pdf
- [BFKL93] Blum, A., Furst, M.L., Kearns, M., Lipton, R.J.: Cryptographic Primitives Based on Hard Learning Problems. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 278–291. Springer, Heidelberg (1994)
- [BKL⁺07] Bogdanov, A., Knudsen, L.R., Leander, G., Paar, C., Poschmann, A., Robshaw, M.J.B., 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)
- [BKW03] Blum, A., Kalai, A., Wasserman, H.: Noise-tolerant learning, the parity problem, and the statistical query model. *J. ACM* 50(4), 506–519 (2003)
- [DKPW12] Dodis, Y., Kiltz, E., Pietrzak, K., Wichs, D.: Message Authentication, Revisited. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 355–374. Springer, Heidelberg (2012)
- [DR02] Daemen, J., Rijmen, V.: The design of rijndael: AES - the advanced encryption standard. Springer (2002)

- [GPW⁺04] Gura, N., Patel, A., Wander, A., Eberle, H., Shantz, S.C.: Comparing Elliptic Curve Cryptography and RSA on 8-bit CPUs. In: Joye, M., Quisquater, J.-J. (eds.) CHES 2004. LNCS, vol. 3156, pp. 119–132. Springer, Heidelberg (2004)
- [GRS05] Gilbert, H., Robshaw, M., Sibert, H.: An active attack against HB+ – a provably secure lightweight authentication protocol, Cryptology ePrint Archive, Report 2005/237 (2005), <http://eprint.iacr.org/>
- [GRS08a] Gilbert, H., Robshaw, M.J.B., Seurin, Y.: $HB^\#$: Increasing the Security and Efficiency of HB^+ . In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 361–378. Springer, Heidelberg (2008)
- [GRS08b] Gilbert, H., Robshaw, M., Seurin, Y.: How to Encrypt with the LPN Problem. In: Aceto, L., Damgard, I., Goldberg, L.A., Halldórsson, M.M., Ingólfssdóttir, A., Walukiewicz, I. (eds.) ICALP 2008, Part II. LNCS, vol. 5126, pp. 679–690. Springer, Heidelberg (2008)
- [HB00] Hopper, N., Blum, M.: A secure human-computer authentication scheme. Tech. Report CMU-CS-00-139, Carnegie Mellon University (2000)
- [HB01] Hopper, N.J., Blum, M.: Secure Human Identification Protocols. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 52–66. Springer, Heidelberg (2001)
- [HLPS11] Hanrot, G., Lyubashevsky, V., Peikert, C., Stehlé, D.: Personal communication (2011)
- [HMOV3] Hankerson, D., Menezes, A.J., Vanstone, S.: Guide to elliptic curve cryptography. Springer-Verlag New York, Inc., Secaucus (2003)
- [JW05] Juels, A., Weis, S.A.: Authenticating Pervasive Devices with Human Protocols. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 293–308. Springer, Heidelberg (2005)
- [Kir11] Kirchner, P.: Improved generalized birthday attack. Cryptology ePrint Archive, Report 2011/377 (2011), <http://eprint.iacr.org/>
- [KPC⁺11] Kiltz, E., Pietrzak, K., Cash, D., Jain, A., Venturi, D.: Efficient authentication from hard learning problems. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 7–26. Springer, Heidelberg (2011)
- [KS06a] Katz, J., Shin, J.S.: Parallel and Concurrent Security of the HB and HB^+ Protocols. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 73–87. Springer, Heidelberg (2006)
- [KS06b] Katz, J., Smith, A.: Analyzing the HB and HB^+ protocols in the “large error” case. Cryptology ePrint Archive, Report 2006/326 (2006), <http://eprint.iacr.org/>
- [KSS10] Katz, J., Shin, J.S., Smith, A.: Parallel and concurrent security of the HB and HB^+ protocols. Journal of Cryptology 23(3), 402–421 (2010)
- [KW05] Kfir, Z., Wool, A.: Picking virtual pockets using relay attacks on contactless smart-card. In: International Conference on Security and Privacy for Emerging Areas in Communications Networks, pp. 47–58 (2005)
- [KW06] Kirschenbaum, I., Wool, A.: How to build a low-cost, extended-range RFID skimmer. In: Proceedings of the 15th USENIX Security Symposium (SECURITY 2006), pp. 43–57. USENIX Association (August 2006)
- [Lan12] Lange, T.: Personal communication (2012)
- [LF06] Leveil, É., Fouque, P.-A.: An Improved LPN Algorithm. In: De Prisco, R., Yung, M. (eds.) SCN 2006. LNCS, vol. 4116, pp. 348–359. Springer, Heidelberg (2006)
- [LLS09] Lee, H., Lee, K., Shin, Y.: AES implementation and performance evaluation on 8-bit microcontrollers. CoRR abs/0911.0482 (2009)
- [LPR10] Lyubashevsky, V., Peikert, C., Regev, O.: On Ideal Lattices and Learning with Errors over Rings. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 1–23. Springer, Heidelberg (2010)

- [Lyu05] Lyubashevsky, V.: The Parity Problem in the Presence of Noise, Decoding Random Linear Codes, and the Subset Sum Problem. In: Chekuri, C., Jansen, K., Rolim, J.D.P., Trevisan, L. (eds.) APPROX 2005 and RANDOM 2005. LNCS, vol. 3624, pp. 378–389. Springer, Heidelberg (2005)
- [MSP] MSP430 datasheet
- [OOV08] Ouafi, K., Overbeck, R., Vaudenay, S.: On the Security of HB[#] against a Man-in-the-Middle Attack. In: Pieprzyk, J. (ed.) ASIACRYPT 2008. LNCS, vol. 5350, pp. 108–124. Springer, Heidelberg (2008)
- [Poe] Poettering, B.: AVRAES: The AES block cipher on AVR controllers, <http://point-at-infinity.org/avraes/>
- [Reg09] Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. J. ACM 56(6) (2009)
- [Tik] Tikkanen, J.: AES implementation on AVR ATmega, 328 p., http://cs.ucsb.edu/~koc/cs178/projects/JT/avr_aes.html
- [Wik] WISP Wiki, WISP 4.0 DL hardware, <http://wisp.wikispaces.com/WISP+4.0+DL>

A Man-In-The-Middle Attack

In this section, we sketch a man-in-the-middle attack against the protocol in Figure 1 that recovers the secret key in time approximately $O(n^{1.5} \cdot 2^{\lambda/2})$ when the adversary is able to insert himself into that many valid interactions between the reader and the tag. For a ring $R = \mathbb{F}_2[X]/(f)$ and a polynomial $g \in R$, define the vector \mathbf{g} to be a vector of dimension $\deg(f)$ whose i^{th} coordinate is the X^i coefficient of g . Similarly, for a polynomial $h \in R$, let $\text{Rot}(h)$ be a $\deg(f) \times \deg(f)$ matrix whose i^{th} column (for $0 \leq i < \deg(f)$) is $h \cdot X^i$, or in other words, the coefficients of the polynomial $h \cdot X^i$ in the ring R . From this description, one can check that for two polynomials $g, h \in R$, the product $\overrightarrow{g \cdot h} = \text{Rot}(g) \cdot \mathbf{h} \bmod 2 = \text{Rot}(h) \cdot \mathbf{g} \bmod 2$.

We now move on to describing the attack. The i^{th} (successful) interaction between a reader \mathcal{R} and a tag \mathcal{T} consists of the reader sending the challenge c_i , and the tag replying with the pair (r_i, z_i) where $z_i - r_i \cdot (s \cdot \pi(c_i) + s')$ is a low-weight polynomial of weight at most $n \cdot \tau'$. The adversary who is observing this interaction will forward the challenge c_i untouched to the tag, but reply to the reader with the ordered pair $(r_i, z'_i = z_i + e_i)$ where e_i is a vector that is strategically chosen with the hope that the vector $z'_i - r_i \cdot (s \cdot \pi(c_i) + s')$ is *exactly* of weight $n \cdot \tau'$. It's not hard to see that it's possible to choose such a vector e_i so that the probability of $z'_i - r_i \cdot (s \cdot \pi(c_i) + s')$ being of weight $n \cdot \tau'$ is approximately $1/\sqrt{n}$. The response (r_i, z'_i) will still be valid, and so the reader will accept. By the birthday bound, after approximately $2^{\lambda/2}$ interactions, there will be a challenge c_j that is equal to some previous challenge c_i . In this case, the adversary replies to the reader with (r_i, z''_i) , where the polynomial z''_i is just the polynomial z'_i whose first bit (i.e. the constant coefficient) is flipped. What the adversary is hoping for is that the reader accepted the response (r_i, z'_i) but rejects (r_i, z''_i) . Notice that the only way this can happen is if the first bit of z'_i is equal to the first bit of $r_i \cdot (s \cdot \pi(c_i) + s')$, and thus flipping it, increases the error by 1 and makes the reader reject. We now explain how finding such a pair of responses can be used to recover the secret key.

Since the polynomial expression $z'_i - r_i \cdot (s \cdot \pi(c_i) + s') = z'_i - r_i \cdot \pi(c_i) \cdot s - r_i \cdot s'$ can be written as matrix-vector multiplications as

$$z'_i - \text{Rot}(r_i \cdot \pi(c_i)) \cdot \mathbf{s} - \text{Rot}(r_i) \cdot \mathbf{s}' \pmod 2,$$

if we let the first bit of z'_i be β_i , the first row of $\text{Rot}(r_i \cdot \pi(c_i))$ be \mathbf{a}_i and the first row of $\text{Rot}(r_i)$ be \mathbf{b}_i , then we obtain the linear equation

$$\langle \mathbf{a}_i, \mathbf{s} \rangle + \langle \mathbf{b}_i, \mathbf{s}' \rangle = \beta_i.$$

To recover the entire secret s, s' , the adversary needs to repeat the above attack until he obtains $2n$ linearly-independent equations (which can be done with $O(n)$ successful attacks), and then use Gaussian elimination to recover the full secret.

Higher-Order Masking Schemes for S-Boxes

Claude Carlet¹, Louis Goubin², Emmanuel Prouff^{3,*},
Michael Quisquater², and Matthieu Rivain⁴

¹ LAGA, Université de Paris 8

`claude.carlet@univ-paris8.fr`

² Université de Versailles St-Quentin-en-Yvelines

`louis.goubin@prism.uvsq.fr`

`michael.quisquater@prism.uvsq.fr`

³ Agence Nationale de la Sécurité des Systèmes d'Information

`e.prouff@gmail.com`

⁴ CryptoExperts

`matthieu.rivain@cryptoexperts.com`

Abstract. Masking is a common countermeasure against side-channel attacks. The principle is to randomly split every sensitive intermediate variable occurring in the computation into $d + 1$ shares, where d is called the *masking order* and plays the role of a security parameter. The main issue while applying masking to protect a block cipher implementation is to design an efficient scheme for the s-box computations. Actually, masking schemes with arbitrary order only exist for Boolean circuits and for the AES s-box. Although any s-box can be represented as a Boolean circuit, applying such a strategy leads to inefficient implementation in software. The design of an efficient and generic higher-order masking scheme was hence until now an open problem. In this paper, we introduce the first masking schemes which can be applied in software to efficiently protect any s-box at any order. We first describe a general masking method and we introduce a new criterion for an s-box that relates to the best efficiency achievable with this method. Then we propose concrete schemes that aim to approach the criterion. Specifically, we give optimal methods for the set of *power functions*, and we give efficient heuristics for the general case. As an illustration we apply the new schemes to the DES and PRESENT s-boxes and we provide implementation results.

1 Introduction

Side-channel analysis is a class of cryptanalytic attacks that exploit the physical environment of a cryptosystem to recover some *leakage* about its secrets. It is often more efficient than a cryptanalysis mounted in the so-called *black-box model* where no leakage occurs. In particular, *continuous side-channel attacks* in which the adversary gets information at each invocation of the cryptosystem are especially threatening. Common attacks as those exploiting the running-time,

* Part of this work has been done while the author was at Oberthur Technologies.

the power consumption or the electromagnetic radiations of a cryptographic computation fall into this class.

Many implementations of block ciphers have been practically broken by continuous side-channel analysis — see for instance [6, 18, 20, 22] — and securing them has been a longstanding issue for the embedded systems industry. A sound approach is to use *secret sharing* [3, 30], often called *masking* in the context of side-channel attacks. This approach consists in splitting each sensitive variable of the implementation (*i.e.* variables depending on the secret key) into $d+1$ shares, where d is called the *masking order*. It has been shown that the complexity of mounting a successful side-channel attack against a masked implementation increases exponentially with the masking order [7]. Starting from this observation, the design of efficient masking schemes for different ciphers has become a foreground issue.

The DES cipher has been the focus of first designs, with the notable work of Goubin and Patarin in [13]. Further schemes have been subsequently published, in particular for the AES cipher, applying masking in hardware or software with different area-time-memory trade-offs [2, 4, 21, 23, 26, 29]. All these schemes deal with *first-order masking*, namely the intermediate variables are split in two shares (a mask and a masked variable). As a result, they only thwart *first order* side-channel attacks in which the adversary exploits the leakage of a single intermediate computation. During the last years, several works have demonstrated that this defense strategy was not sufficient for long term security purpose and that *higher-order attacks* could be successfully performed against cryptographic implementations (see *e.g.* [22]). This has raised the need for secure and efficient higher-order masking schemes.

Higher-Order Masking. The principle of higher-order masking is to split every sensitive variable x occurring during the computation into $d+1$ shares x_0, \dots, x_d in such a way that the following relation is satisfied for a group operation \perp :

$$x_0 \perp x_1 \perp \dots \perp x_d = x . \quad (1)$$

In the rest of the paper, we shall consider that \perp is the addition over some field of characteristic 2. Usually, the d shares x_1, \dots, x_d (called *the masks*) are randomly picked up and the last one x_0 (called *the masked variable*) is processed such that it satisfies (1). When d random masks are involved per sensitive variable the masking is said to be *of order d* . The tuple $(x_i)_i$ is further called a *d th-order encoding of x* .

When higher-order masking is involved to protect a block cipher implementation, a so-called *masking scheme* must be designed to enable the computation on masked data. Such a scheme must ensure that the final shares correspond to the expected ciphertext on the one hand, and it must ensure the d th-order security property for the chosen order d on the other hand. The latter property states that every tuple of d or less intermediate variables is independent of any sensitive variable. When satisfied, it guarantees that no attack of order lower than or equal to d is possible.

Most block cipher structures (*e.g.* AES or DES) are iterative, meaning that they apply several times a same transformation, called *round*, to an internal state initially filled with the plaintext. The round itself is composed of a key addition, one or several linear transformation(s) and one or several non-linear s-box(es). Key addition and linear transformations are easily handled as linearity enables to process each share independently. The main difficulty in designing masking schemes for block ciphers hence lies in masking the s-box(es).

Masking and S-Boxes. Whereas many solutions have been proposed to deal with the case of first-order masking (see *e.g.* [2, 4, 21, 25]), only a few solutions exist for the higher-order case. A scheme has been proposed by Schramm and Paar in [29] which generalizes the (first-order) table recomputation method described in [2, 21]. Although the authors apply their method in the particular case of an AES implementation, it is generic and can be applied to protect any s-box. Unfortunately, this scheme has been shown to be vulnerable to a 3rd-order attack whatever the chosen masking order [8]. In other words, it only provides 2nd-order security. Further schemes were proposed by Rivain, Dottax and Prouff in [26] with formal security proofs but still limited to 2nd-order security.

The first scheme achieving d th-order security for an arbitrary chosen d has been designed by Ishai, Sahai and Wagner in [14]. The here-called *ISW scheme* consists in masking the Boolean representation of an algorithm which is composed of logical operations NOT and AND. Securing a NOT for any order d is straightforward since $x = \bigoplus_i x_i$ implies $\text{NOT}(x) = \text{NOT}(x_0) \oplus x_1 \cdots \oplus x_d$. The main contribution of [14] is a method to secure the AND operation for any arbitrary order d (the description of this scheme is recalled in Section 2.1). Although the ISW scheme is an important theoretical result, its practical application faces some issues. At the hardware level, the obtained circuits may have prohibitive area requirements, especially for being used in embedded systems (privileged targets of side-channel attacks). Moreover, Mangard *et al.* have shown in [19, 20] that masking at the hardware level is sensitive to *glitches* which induce unpredicted flaws in masked circuits. Preventing glitches can be done thanks to synchronization elements (*e.g.* registers or latches) [24] or by performing additional sharing [23] but in both cases, the circuit size is still significantly increased. On the other hand, a direct application of the ISW scheme to secure an s-box computation in software would consist in taking the Boolean representation of the s-box and in processing every logical operation successively in a masked way. Since the Boolean representation of common s-boxes involves a huge number of logical operations, the resulting implementation would likely be inefficient.

In the particular case of AES, a solution has been proposed by Rivain and Prouff in [27] to efficiently mask the s-box processing at any order. Specifically, the authors use the algebraic structure of the AES s-box, which is the composition of an affine function over \mathbb{F}_2^8 with the power function $x \mapsto x^{254}$ over \mathbb{F}_{256} , and they show that it can be expressed as a sequence of operations involving a few linear functions over \mathbb{F}_2^8 (easy to mask) and four multiplications over \mathbb{F}_{256} . The latter are secured by applying the ISW scheme (generalized to \mathbb{F}_{256}).

Subsequently, Kim, Hong and Lim have presented in [15] an extension of Rivain and Prouff's scheme, which is based on the tower-field approach from [28]. On the other hand, Genelle, Prouff and Quisquater have proposed in [12] a higher-order scheme based on the alternate use of Boolean masking and multiplicative masking. Although schemes in [15] and [12] achieve better performances than [27], they are still restricted to the AES s-box and their generalization to any s-box (or subclasses) is an open issue.

Our Contribution. The present paper introduces the first higher-order masking scheme which can be applied to efficiently protect any s-box processing in software. We first give a general method that extends the Rivain and Prouff approach to mask any s-box and we introduce a new criterion for an s-box that relates to the best efficiency achievable with our method. Then we give concrete schemes that aim to approach the so-called *masking complexity*. Specifically, we give optimal methods for the set of *power functions*, and we give efficient heuristics for the general case. As an illustration we apply our scheme to the DES and PRESENT s-boxes and we provide implementation results.

2 Higher-Order Masking of any S-Box

In this section, we describe a general method to mask any s-box and we introduce a related *masking complexity* criterion.

2.1 General Method

An s-box is a function from $\{0, 1\}^n$ to $\{0, 1\}^m$ with $m \leq n$ and n small (typically $n \in \{4, 6, 8\}$). We shall use the terminology of (n, m) s-box when the dimensions need to be specified. To design a higher-order masking scheme for such a function, our approach is to express it as a sequence of affine functions over \mathbb{F}_2^n , and multiplications over \mathbb{F}_{2^n} . Such a strategy is always possible since any (n, m) s-box can be represented by a polynomial function $x \mapsto \sum_{i=0}^{2^n-1} a_i x^i$ over \mathbb{F}_{2^n} where the a_i are constant coefficients in \mathbb{F}_{2^n} . The a_i can be obtained from the s-box look-up table by applying Lagrange's Interpolation Theorem. When m is strictly lower than n , the m -bit outputs can be embedded into \mathbb{F}_{2^n} by padding them to n -bit outputs (*e.g.* by setting most significant bits to 0). The padding is then removed after the polynomial evaluation. We recall hereafter the Lagrange Interpolation Theorem applied to our context.

Theorem 1 (Lagrange Interpolation). *Let S be a function $\mathbb{F}_{2^n} \rightarrow \mathbb{F}_{2^n}$. Then, for every $x \in \mathbb{F}_{2^n}$, we have:*

$$S(x) = \sum_{\alpha \in \mathbb{F}_{2^n}} S(\alpha) \ell_\alpha(x), \quad (2)$$

where, for every $\alpha \in \mathbb{F}_{2^n}$, ℓ_α is defined as:

$$\ell_\alpha(x) = \prod_{\substack{\beta \in \mathbb{F}_{2^n} \\ \beta \neq \alpha}} \frac{x - \beta}{\alpha - \beta}. \quad (3)$$

Remark 1. The ℓ_α are called the *Lagrange basis polynomials* and satisfy $\ell_\alpha(x) = 1$ if $x = \alpha$ and $\ell_\alpha(x) = 0$ otherwise. In particular, every ℓ_α is a monic polynomial of degree $2^n - 1$, and we have $\ell_\alpha(x) = (x + \alpha)^{2^n - 1} + 1$. Moreover, the coefficients of $S(x)$ can be directly computed from the Mattson-Solomon polynomial by:

$$a_i = \begin{cases} S(0) & \text{if } i = 0 \\ \sum_{k=0}^{2^n-2} S(\alpha^k) \alpha^{-ki} & \text{if } 1 \leq i \leq 2^n - 2 \\ S(1) + \sum_{i=0}^{2^n-2} a_i & \text{if } i = 2^n - 1 \end{cases}$$

for every primitive element α of \mathbb{F}_{2^n} .

The polynomial representation of an s-box is based on four kinds of operations over \mathbb{F}_{2^n} : additions, scalar multiplications (*i.e.* multiplications by constants), squares, and regular multiplications (*i.e.* of two different variables). Except for the latter, all these operations are \mathbb{F}_2^n -linear (or \mathbb{F}_2^n -affine), that is the corresponding function over \mathbb{F}_2^n are linear (resp. affine). The processing of any s-box can then be performed as a sequence of \mathbb{F}_2^n -affine functions (themselves composed of additions, squares and scalar multiplications over \mathbb{F}_{2^n}) and of regular multiplications over \mathbb{F}_{2^n} , called *nonlinear multiplications* in the following. Masking an s-box processing can hence be done by masking every affine function and every nonlinear multiplication independently. We recall hereafter how this can be done for each category.

Masking of \mathbb{F}_2^n -affine functions. Let $x = \sum_i x_i$ be a shared variable. Every affine function g with additive part c_g satisfies:

$$g(x) = \begin{cases} \sum_{i=0}^d g(x_i) & \text{if } d \text{ is even,} \\ c_g + \sum_{i=0}^d g(x_i) & \text{if } d \text{ is odd.} \end{cases}$$

The masked processing of g then simply consists in evaluating g for every share x_i , and possibly correcting one of them by addition of c_g . Such a processing clearly achieves d th-order security as the shares are all processed independently.

Masking of nonlinear multiplications. Every nonlinear multiplication can be processed by using the ISW scheme. Let $a, b \in \mathbb{F}_{2^n}$ and let $(a_i)_{0 \leq i \leq d}$ and $(b_i)_{0 \leq i \leq d}$ be d th-order encoding of a and b . To securely compute a d th-order encoding $(c_i)_{0 \leq i \leq d}$ of $c = ab$, the ISW method over \mathbb{F}_{2^n} performs as follows:¹

1. For every $0 \leq i < j \leq d$, pick up a random value $r_{i,j}$ in \mathbb{F}_{2^n} .
2. For every $0 \leq i < j \leq d$, compute $r_{j,i} = (r_{i,j} + a_i b_j) + a_j b_i$.
3. For every $0 \leq i \leq d$, compute $c_i = a_i b_i + \sum_{j \neq i} r_{i,j}$.

It can be checked that the obtained shares are a sound encoding of c . Namely, we have:

$$\sum_{i=0}^d c_i = \left(\sum_{i=0}^d a_i \right) \left(\sum_{i=0}^d b_i \right) = ab = c.$$

¹ The use of brackets indicates the order in which the operations are performed, which is mandatory for the security of the scheme.

In [14] it is shown that the above computation achieves $(d/2)$ th-order security. A tighter security proof is given in [27] which shows that d th-order security is actually achieved as long as the masks of the two inputs are independent.

Remark 2. Another method to process a masked multiplication at an arbitrary order is used in [10] to achieve provable security under specific leakage assumptions. However this method requires more operations and more random bits than the ISW scheme does. For this reason, the ISW scheme should be preferred in a usual d th-order security model.

2.2 Masking Complexity

The scheme described in the previous section secures the computation of any (n, m) s-box S by masking its polynomial representation over \mathbb{F}_{2^n} . The evaluation of such a polynomial is composed of \mathbb{F}_2^n -affine functions g and of nonlinear multiplications. The masked processing of each \mathbb{F}_2^n -affine function g merely involves $d + 1$ evaluations of g itself, while it involves $(d + 1)^2$ field multiplications, $2d(d + 1)$ field additions and the generation of $nd(d + 1)/2$ random bits for each nonlinear multiplication. The masked processing of \mathbb{F}_2^n -affine functions hence quickly becomes negligible compared to the masked processing of nonlinear multiplications as d grows. This observation motivates the following definition of the *masking complexity* for an s-box.

Definition 1 (Masking Complexity). *Let m and n be two integers such that $m \leq n$. The masking complexity of a (n, m) s-box is the minimal number of nonlinear multiplications required to evaluate its polynomial representation over \mathbb{F}_{2^n} .*

The following proposition directly results from this definition.

Proposition 1. *The masking complexity of an s-box is invariant when composed with \mathbb{F}_2^n -affine bijections in input and/or in output.*

Remark 3. Since field isomorphisms are \mathbb{F}_2 -linear bijections, the choice of the irreducible polynomial to represent field elements does not impact the masking complexity of an s-box.

In the next sections, we address the issue of finding polynomial evaluations of an s-box that aim at minimizing the number of nonlinear multiplications. Those constructions will enable us to deduce upper bounds on the masking complexity of an s-box. We first study the case of power functions whose polynomial representation has a single monomial (*e.g.* the AES s-box). For these functions, we exhibit the exact masking complexity by deriving addition chains with minimal number of nonlinear multiplications. We then address the general case and provide efficient heuristics to evaluate any s-box with a low number of nonlinear multiplications.

3 Optimal Masking of Power Functions

In this section, we consider s-boxes for which the polynomial representation over \mathbb{F}_{2^n} is a single monomial. These s-boxes are usually called *power functions* in the literature. We describe a generic method to compute the masking complexity of such s-boxes. Our method involves the notion of *cyclotomic class*.

Definition 2. Let $\alpha \in [0; 2^n - 2]$. The cyclotomic class of α is the set C_α defined by:

$$C_\alpha = \{\alpha \cdot 2^i \bmod 2^n - 1; i \in [0; n - 1]\}.$$

We have the following proposition.

Proposition 2. Let $\mu(m)$ denote the multiplicative order of 2 modulo m and let φ denote the Euler's totient function. For every divisor δ of $2^n - 1$, the number of distinct cyclotomic classes $C_\alpha \subseteq [0; 2^n - 2]$ with $\gcd(\alpha, 2^n - 1) = \delta$ is $\varphi(\frac{2^n - 1}{\delta}) / \mu(\frac{2^n - 1}{\delta})$. It follows that the total number of distinct cyclotomic classes of $[0; 2^n - 2]$ equals:

$$\sum_{\delta | (2^n - 1)} \frac{\varphi(\delta)}{\mu(\delta)}.$$

Proof. Proposition 2 can be deduced from the following facts:

- An integer $\alpha \in [0; 2^n - 2]$ satisfies $\gcd(\alpha, 2^n - 1) = \delta$ if and only if $\alpha = \delta\beta$, with $\gcd(\beta, \frac{2^n - 1}{\delta}) = 1$. There are thus $\varphi(\frac{2^n - 1}{\delta})$ integers $\alpha \in [0; 2^n - 2]$ such that $\gcd(\alpha, 2^n - 1) = \delta$.
- For any α such that $\gcd(\alpha, 2^n - 1) = \delta$ (hence of the form $\alpha = \delta\beta$ with $\gcd(\beta, \frac{2^n - 1}{\delta}) = 1$), we have $\alpha \cdot 2^i \equiv \alpha \cdot 2^j \bmod 2^n - 1$ if and only if $\beta \cdot 2^i \equiv \beta \cdot 2^j \bmod \frac{2^n - 1}{\delta}$, that is, if and only if $2^i \equiv 2^j \bmod \frac{2^n - 1}{\delta}$. Hence C_α has cardinality $\#C_\alpha = \mu(\frac{2^n - 1}{\delta})$.

The set of integers $\alpha \in [0; 2^n - 2]$ such that $\gcd(\alpha, 2^n - 1) = \delta$ is partitioned into cyclotomic classes, each of them having cardinality $\mu(\frac{2^n - 1}{\delta})$. Hence the number of such cyclotomic classes is $\varphi(\frac{2^n - 1}{\delta}) / \mu(\frac{2^n - 1}{\delta})$. It follows that the total number of distinct cyclotomic classes of $[0; 2^n - 2]$ equals $\sum_{\delta | (2^n - 1)} \varphi(\frac{2^n - 1}{\delta}) / \mu(\frac{2^n - 1}{\delta}) = \sum_{\delta | (2^n - 1)} \varphi(\delta) / \mu(\delta)$. □

The study of cyclotomic classes is interesting in our context since a power x^α can be computed from a power x^β without any nonlinear multiplication if and only if α and β lie in the same cyclotomic class. Hence, all the power functions with exponents within a given cyclotomic class have the same masking complexity and computing the masking complexity for all the power functions over \mathbb{F}_{2^n} thus amounts to compute this complexity for each cyclotomic class over \mathbb{F}_{2^n} . In what follows, we perform such a computation for fields \mathbb{F}_{2^n} of small dimensions n .

To compute the masking complexity for an element in a cyclotomic class, we use the following observation: determining the masking complexity of a power function $x \mapsto x^\alpha$ amounts to find the addition chain for α with the least number of additions which are not doublings (see [16] for an introduction to addition chains). This kind of addition chain is usually called a *2-addition chain*.² Let $(\alpha_i)_i$ denote some addition chain. At step i , it is possible to obtain any element within the cyclotomic classes $(C_{\alpha_j})_{j \leq i}$ using doublings only. As we are interested in finding the addition chain with the least number of additions which are not doublings, the problem we need to solve is the following: given some $\alpha \in C_\alpha$, find the shortest chain $C_{\alpha_0} \rightarrow C_{\alpha_1} \rightarrow \dots \rightarrow C_{\alpha_k}$ where $C_{\alpha_0} = C_1$, $C_{\alpha_k} = C_\alpha$ and for every $i \in [1; k]$, there exists $j, \ell < i$ such that $\alpha_i = \alpha'_j + \alpha'_\ell$ where $\alpha'_j \in C_{\alpha_j}$ and $\alpha'_\ell \in C_{\alpha_\ell}$.

We shall denote by \mathcal{M}_k^n the class of exponents α such that $x \mapsto x^\alpha$ has a masking complexity equal to k . The family of classes $(\mathcal{M}_k^n)_k$ is a partition of $[0; 2^n - 2]$ and each \mathcal{M}_k^n is the union of one or several cyclotomic classes. For a small dimension n , we can proceed by exhaustive search to determine the shortest 2-addition chain(s) for each cyclotomic class. We implemented such an exhaustive search from which we obtained the masking complexity classes \mathcal{M}_k^n for $n \leq 11$ (note that in practice most s-boxes have dimension $n \leq 8$). Table 1 summarizes the obtained results for $n \in \{4, 6, 8\}$ (usual dimensions). Results for other dimensions are summarized in appendix. Additionally, Table 2 gives the optimal 2-addition chains (in exponential notation) corresponding to every cyclotomic class for $n = 8$.

It is interesting to note that for every n , the *inverse function* $x \mapsto x^{2^n-2}$ related to the cyclotomic class $C_{2^{n-1}-1}$ always has the highest masking complexity. In particular, the inverse function $x \mapsto x^{254}$ (for $n = 8$) used in the AES has a masking complexity of 4 as it was conjectured in [27].

4 Efficient Heuristics for General S-Boxes

We now address the general case of an s-box having a polynomial representation $\sum_{j=0}^{2^n-1} a_j x^j$ over \mathbb{F}_{2^n} . A straightforward solution is to successively compute every power x^j using $x^j = (x^{j/2})^2$ if j is even and $x^j = x^{j-1}x$ if j is odd, while updating the polynomial value by adding the monomial $a_j x^j$ at every step. Such a method requires $2^{n-1} - 1$ nonlinear multiplications. As we show hereafter, less naive methods exist that substantially lower the number of nonlinear multiplications. We propose two different methods and then compare their efficiency.

² This problem has been studied in the general setting where the multiplication by q (and not specifically by 2) is considered *free* and the obtained addition chains are called *q-addition chains* [31]. The purpose is to find efficient exponentiation methods in \mathbb{F}_q (as in such field the Frobenius map $x \mapsto x^q$ is efficient). To the best of our knowledge, apart from a specific application to the SFLASH signature algorithm in [1], the case of 2-addition chains has not been particularly investigated.

Table 1. Cyclotomic classes for $n \in \{4, 6, 8\}$ w.r.t. the masking complexity k

k	Cyclotomic classes in \mathcal{M}_k^n
$n = 4$	
0	$C_0 = \{0\}, C_1 = \{1, 2, 4, 8\}$
1	$C_3 = \{3, 6, 12, 9\}, C_5 = \{5, 10\}$
2	$C_7 = \{7, 14, 13, 11\}$
$n = 6$	
0	$C_0 = \{0\}, C_1 = \{1, 2, 4, 8, 16, 32\}$
1	$C_3 = \{3, 6, 12, 24, 48, 33\}, C_5 = \{5, 10, 20, 40, 17, 34\}, C_9 = \{9, 18, 36\}$
2	$C_7 = \{7, 14, 28, 56, 49, 35\}, C_{11} = \{11, 22, 44, 25, 50, 37\},$ $C_{13} = \{13, 26, 52, 41, 19, 38\}, C_{15} = \{15, 30, 29, 27, 23\},$ $C_{21} = \{21, 42\}, C_{27} = \{27, 54, 45\}$
3	$C_{23} = \{23, 46, 29, 58, 53, 43\}, C_{31} = \{31, 62, 61, 59, 55, 47\}$
$n = 8$	
0	$C_0 = \{0\}, C_1 = \{1, 2, 4, 8, 16, 32, 64, 128\}$
1	$C_3 = \{3, 6, 12, 24, 48, 96, 192, 129\}, C_5 = \{5, 10, 20, 40, 80, 160, 65, 130\},$ $C_9 = \{9, 18, 36, 72, 144, 33, 66, 132\}, C_{17} = \{17, 34, 68, 136\}$
2	$C_7 = \{7, 14, 28, 56, 112, 224, 193, 131\}, C_{11} = \{11, 22, 44, 88, 176, 97, 194, 133\},$ $C_{13} = \{13, 26, 52, 104, 208, 161, 67, 134\}, C_{15} = \{15, 30, 60, 120, 240, 225, 195, 135\},$ $C_{19} = \{19, 38, 76, 152, 49, 98, 196, 137\}, C_{21} = \{21, 42, 84, 168, 81, 162, 69, 138\},$ $C_{25} = \{25, 50, 100, 200, 145, 35, 70, 140\}, C_{27} = \{27, 54, 108, 216, 177, 99, 198, 141\},$ $C_{37} = \{37, 74, 148, 41, 82, 164, 73, 146\}, C_{45} = \{45, 90, 180, 105, 210, 165, 75, 150\},$ $C_{51} = \{51, 102, 204, 153\}, C_{85} = \{85, 170\}$
3	$C_{23} = \{23, 46, 92, 184, 113, 226, 197, 139\}, C_{29} = \{29, 58, 116, 232, 209, 163, 71, 142\},$ $C_{31} = \{31, 62, 124, 248, 241, 227, 199, 143\}, C_{39} = \{39, 78, 156, 57, 114, 228, 201, 147\},$ $C_{43} = \{43, 86, 172, 89, 178, 101, 202, 149\}, C_{47} = \{47, 94, 188, 121, 242, 229, 203, 151\},$ $C_{53} = \{53, 106, 212, 169, 83, 166, 77, 154\}, C_{55} = \{55, 110, 220, 185, 115, 230, 205, 155\},$ $C_{59} = \{59, 118, 236, 217, 179, 103, 206, 157\}, C_{61} = \{61, 122, 244, 233, 211, 167, 79, 158\},$ $C_{63} = \{63, 126, 252, 249, 243, 231, 207, 159\}, C_{87} = \{87, 174, 93, 186, 117, 234, 213, 171\},$ $C_{91} = \{91, 182, 109, 218, 181, 107, 214, 173\}, C_{95} = \{95, 190, 125, 250, 245, 235, 215, 175\},$ $C_{111} = \{111, 222, 189, 123, 246, 237, 219, 183\}, C_{119} = \{119, 238, 221, 187\}$
4	$C_{127} = \{127, 254, 253, 251, 247, 239, 223, 191\}$

Table 2. Optimal 2-addition chains (in exponential notation) for cyclotomic classes for $n = 8$

k	2-addition chains with k nonlinear multiplications
1	$x^3 \leftarrow x \times x^2 - x^5 \leftarrow x \times x^4$ $x^9 \leftarrow x \times x^8 - x^{17} \leftarrow x \times x^{16}$
2	$x^7 \leftarrow x \times x^2 \times x^4 - x^{11} \leftarrow x \times x^2 \times x^8$ $x^{13} \leftarrow x \times x^4 \times x^8 - x^{15} \leftarrow x^3 \times (x^3)^4$ $x^{19} \leftarrow x \times x^2 \times x^{16} - x^{21} \leftarrow x \times x^4 \times x^{16}$ $x^{27} \leftarrow x^3 \times (x^3)^8 - x^{37} \leftarrow x \times x^4 \times x^{32}$ $x^{45} \leftarrow x^5 \times (x^5)^8 - x^{51} \leftarrow x^3 \times (x^3)^{16}$ $x^{85} \leftarrow x^5 \times (x^5)^{16}$
3	$x^{23} \leftarrow x \times x^2 \times x^4 \times x^{16} - x^{29} \leftarrow x \times x^4 \times x^8 \times x^{16}$ $x^{31} \leftarrow x^3 \times (x^3)^4 \times x^{16} - x^{29} \leftarrow x \times x^2 \times x^4 \times x^{32}$ $x^{43} \leftarrow x \times x^2 \times x^8 \times x^{32} - x^{47} \leftarrow x^3 \times (x^3)^4 \times x^{32}$ $x^{53} \leftarrow x \times x^2 \times x^{16} \times x^{32} - x^{55} \leftarrow x^3 \times x^4 \times (x^3)^{16}$ $x^{59} \leftarrow x^3 \times (x^3)^8 \times x^{32} - x^{59} \leftarrow x^5 \times x^{16} \times (x^5)^8$ $x^{63} \leftarrow x^7 \times (x^7)^8 - x^{87} \leftarrow x^2 \times x^5 \times (x^5)^{16}$ $x^{91} \leftarrow x^3 \times (x^3)^8 \times x^{64} - x^{95} \leftarrow x^5 \times (x^5)^2 \times (x^5)^{16}$ $x^{111} \leftarrow x^3 \times (x^3)^4 \times (x^3)^{32} - x^{63} \leftarrow x^7 \times (x^7)^{16}$
4	$x^{127} \leftarrow x^3 \times (x^3)^4 \times (x^3)^{16} \times x^{64}$

4.1 Cyclotomic Method

Let q denote the number of distinct cyclotomic classes of $[0; 2^n - 2]$. The polynomial representation of S can be written as:

$$S(x) = a_0 + \left(\sum_{i=1}^q Q_i(x) \right) + a_{2^n-1} x^{2^n-1} ,$$

where the Q_i are polynomials such that every Q_i has powers from a single cyclotomic class C_{α_i} , namely we can write $Q_i(x) = \sum_j a_{i,j} x^{\alpha_i 2^j}$ for some coefficients $a_{i,j}$ in \mathbb{F}_{2^n} . Let us then denote L_i the linearized polynomial $L_i(x) = \sum_j a_{i,j} x^{2^j}$ which is a \mathbb{F}_2^n -linear function of x . We have $Q_i(x) = L_i(x^{\alpha_i})$ by definition. The *cyclotomic method* simply consists in deriving the powers x^{α_i} for each cyclotomic class C_{α_i} as well as x^{2^n-1} if $a_{2^n-1} \neq 0$, and in evaluating $S(x) = a_0 + \left(\sum_{i=1}^q L_i(x^{\alpha_i}) \right) + a_{2^n-1} x^{2^n-1}$. The powers x^{α_i} can each be derived with a single nonlinear multiplication. This is obvious for the α_i lying in \mathcal{M}_1^n . Then it is clear that every power x^{α_i} with $\alpha_i \in \mathcal{M}_{k+1}^n$ can be derived with a single multiplication from the powers $(x^{\alpha_i})_{\alpha_i \in \mathcal{M}_k^n}$. The power x^{2^n-1} can then be derived with a single nonlinear multiplication from the power x^{2^n-2} . The cyclotomic method hence involves a number of nonlinear multiplications equal to the number of cyclotomic classes, minus 2 (as x^0 and x^1 are obtained without nonlinear multiplication), plus 1 (to derive x^{2^n-1}). By Proposition 2, we then have the following result.

Proposition 3 (Cyclotomic Method). *Let m and n be two positive integers such that $m \leq n$. The masking complexity of every (n, m) s-box is upper-bounded by:*

$$\sum_{\delta | (2^n-1)} \frac{\varphi(\delta)}{\mu(\delta)} - 1 .$$

An (n, m) s-box S is said to be *balanced* if for every $y \in \{0, 1\}^m$, the number of preimages of y for S is constant to 2^{n-m} . The following lemma gives a well-known folklore result.

Lemma 1. *Let m and n be two positive integers such that $m \leq n$. The polynomial representation of every balanced (n, m) s-box has degree strictly lower than $2^n - 1$.*

Proof. Since Lagrange basis polynomials are all monic of degree $2^n - 1$, the coefficient a of the power to the $2^n - 1$ in the polynomial representation of S satisfies $a = \sum_{\alpha \in \mathbb{F}_{2^n}} S(\alpha)$, which equals 0 if S is balanced. \square

When the polynomial representation of the s-box has degree strictly lower than $2^n - 1$, the cyclotomic method saves one nonlinear multiplication since the power x^{2^n-1} is not required. Namely, we have the following corollary of Proposition 3.

Corollary 1 (Cyclotomic Method). *Let m and n be two positive integers such that $m \leq n$ and let S be a (n, m) s -box. If S is balanced, then the masking complexity of S is upper-bounded by:*

$$\sum_{\delta|(2^n-1)} \frac{\varphi(\delta)}{\mu(\delta)} - 2 .$$

4.2 Parity-Split Method

The *parity-split method* is composed of two stages. The first stage derives a set of powers $(x^j)_{j \leq q}$ for some q using the straightforward method described in the introduction of this section. The second stage essentially consists in an application of the Knuth-Eve polynomial evaluation algorithm [9, 17] which is based on a recursive use of the following lemma.

Lemma 2. *Let n and t be two positive integers and let Q be a polynomial of degree t over $\mathbb{F}_{2^n}[x]$. There exist two polynomials Q_1 and Q_2 of degree upper-bounded by $\lfloor t/2 \rfloor$ over $\mathbb{F}_{2^n}[x]$ such that:*

$$Q(x) = Q_1(x^2) + Q_2(x^2)x . \tag{4}$$

By applying Lemma 2 to the polynomial representation of S , we get $S(x) = Q_1(x^2) + Q_2(x^2)x$, where Q_1 and Q_2 are two polynomials of degrees upper-bounded by $2^{n-1} - 1$. We deduce that S can be computed based on the set of powers $(x^{2^j})_{j \leq 2^{n-1}-1}$ plus a single multiplication by x . Then, applying Lemma 2 again to the polynomials Q_1 and Q_2 both of degrees upper bounded by $2^{n-1} - 1$, we get two new pairs of polynomials (Q_{11}, Q_{12}) and (Q_{21}, Q_{22}) such that $Q_1(x^2) = Q_{11}(x^4) + Q_{12}(x^4)x^2$ and $Q_2(x^2) = Q_{21}(x^4) + Q_{22}(x^4)x^2$. The degrees of the new polynomials are upper bounded by $2^{n-2} - 1$. We then deduce that S can be computed based on the set of powers $(x^{4^j})_{j \leq 2^{n-2}-1}$ plus 1 multiplication by x and 2 multiplications by x^2 . Eventually, by applying Lemma 2 recursively r times, we get an evaluation of S involving evaluations in x^{2^r} of polynomials of degrees upper-bounded by $2^{n-r} - 1$, plus $\sum_{i=0}^{r-1} 2^i = 2^r - 1$ multiplications by powers of x of the form x^{2^i} with $i \leq r - 1$. The overall evaluation of S hence requires $2^r - 1$ nonlinear multiplications (the x^{2^i} being obtained with squares only) plus the evaluation in x^{2^r} of polynomials of degrees upper-bounded by $2^{n-r} - 1$. The latter evaluation can be performed by first deriving all the powers $(x^{2^r j})_{j \leq 2^{n-r}-1}$ and then evaluating the polynomials (which only involves scalar multiplications and additions once the powers have been derived). For every $j \leq 2^{n-r} - 1$, the powers $(x^{2^r j})_{j \leq 2^{n-r}-1}$ can be computed successively from $y = x^{2^r}$ by $y^j = (y^{j/2})^2$ if j is even and $y^j = y^{j-1}x$ if j is odd. This takes some squares plus $2^{n-r-1} - 1$ nonlinear multiplications (*i.e.* one per odd integer in $[3, 2^{n-r} - 1]$).

We then deduce the following proposition.

Proposition 4. *Let m and n be two positive integers such that $m \leq n$. The masking complexity of every (n, m) s-box is upper-bounded by:*

$$\min_{0 \leq r \leq n} (2^{n-r-1} + 2^r) - 2 = \begin{cases} 3 \cdot 2^{(n/2)-1} - 2 & \text{if } n \text{ is even,} \\ 2^{(n+1)/2} - 2 & \text{if } n \text{ is odd.} \end{cases} \quad (5)$$

Note that the value of r for which the minimum is reached in (5) is $r = \lfloor \frac{n}{2} \rfloor$.

4.3 Comparison

Table 3 summarizes the number of nonlinear multiplications obtained by the cyclotomic method (for balanced s-boxes) and by the parity-split method. We see that the cyclotomic method works better for small dimensions ($n \leq 5$) and the parity-split method for higher dimensions ($n \geq 6$). Furthermore, the superiority of the parity-split method becomes significant as n grows.

Table 3. Number of nonlinear multiplications w.r.t. the evaluation method

Method \ n	3	4	5	6	7	8	9	10	11
Cyclotomic	1	3	5	11	17	33	53	105	192
Parity-Split	2	4	6	10	14	22	30	46	62

We emphasize that these bounds may not be optimal, namely they may be higher than the maximum masking complexity of (n, m) s-boxes. We let open the issue of finding more efficient (or provably optimal) methods in the general case for further research.

5 Application to DES and PRESENT

In this section we apply the proposed methods to the s-boxes of two different block ciphers: the well-known and still widely used Data Encryption Standard (DES) [11], and the lightweight block cipher PRESENT [5]. The former uses eight different $(6, 4)$ s-boxes and the latter uses a single $(4, 4)$ s-box. According to Table 3, we shall prefer the parity-split method for the DES s-boxes (10 nonlinear multiplications), and the cyclotomic method for the PRESENT s-box (3 nonlinear multiplications).

5.1 Parity-Split Method on DES S-Boxes

The parity-split method on a DES s-box uses a polynomial representation of the s-box over \mathbb{F}_{64} which satisfies:

$$S : x \mapsto Q_0(x^8) + Q_1(x^8) \cdot x^4 + (Q_2(x^8) + Q_3(x^8) \cdot x^4) \cdot x^2 + (Q_4(x^8) + Q_5(x^8) \cdot x^4 + (Q_6(x^8) + Q_7(x^8) \cdot x^4) \cdot x^2) \cdot x \quad (6)$$

where the Q_i are degree-7 polynomials, namely, there exist coefficients $a_{i,j}$ for $0 \leq i, j \leq 7$ such that:

$$Q_i(x^8) = a_{i,0} + a_{i,1}x^8 + a_{i,2}x^{16} + a_{i,3}x^{24} + a_{i,4}x^{32} + a_{i,5}x^{40} + a_{i,6}x^{48} + a_{i,7}x^{56} .$$

We first derive the powers x^{8j} for $j = 1, 2, \dots, 7$, which is done at the cost of 3 nonlinear multiplications by:

$$\begin{aligned} x^8 &\leftarrow ((x^2)^2)^2; \quad x^{16} \leftarrow (x^8)^2; \quad x^{24} \leftarrow x^8 \cdot x^{16}; \quad x^{32} \leftarrow (x^{16})^2; \\ x^{40} &\leftarrow x^8 \cdot x^{32}; \quad x^{48} \leftarrow (x^{24})^2; \quad x^{56} \leftarrow x^8 \cdot x^{48}; \end{aligned}$$

Then we evaluate each polynomial $Q_i(x^8)$ as a linear combination of the above powers. Finally, we evaluate (6) at the cost of 7 nonlinear multiplications and a few additions. The nonlinear multiplications are computed using the ISW scheme over \mathbb{F}_{64} such as recalled in Section 2.1. A detailed implementation for the overall masked s-box evaluation is given in the extended version of this paper.

5.2 Cyclotomic Method on PRESENT S-Box

The cyclotomic method on the PRESENT s-box starts from the straightforward polynomial representation of the s-box over \mathbb{F}_{16} :

$$S : x \mapsto a_0 + a_1x + a_2x^2 + \dots + a_{14}x^{14} ,$$

(where the degree is indeed strictly lower than 15 by Lemma 1). We then have:

$$S(x) = a_0 + L_1(x) + L_3(x^3) + L_5(x^5) + L_7(x^7) . \tag{7}$$

where:

$$\begin{aligned} L_1 : x &\mapsto a_1x + a_2x^2 + a_4x^4 + a_8x^8 \\ L_3 : x &\mapsto a_3x + a_6x^2 + a_{12}x^4 + a_9x^8 \\ L_5 : x &\mapsto a_5x + a_{10}x^2 \\ L_7 : x &\mapsto a_7x + a_{14}x^2 + a_{13}x^4 + a_{11}x^8 \end{aligned}$$

and the L_i are \mathbb{F}_2^4 -linear.

We first derive the powers x^3 , x^5 , and x^7 , which is done at the cost of 3 nonlinear multiplications by: $x^3 \leftarrow x \cdot x^2$; $x^5 \leftarrow x^3 \cdot x^2$; $x^7 \leftarrow x^5 \cdot x^2$. Then we evaluate (7) which costs a few linear transformations and additions. A detailed implementation for the overall masked s-box evaluation is given in the extended version of this paper.

5.3 Implementation Results

In this section, we give implementation results for our scheme applied to DES and PRESENT s-boxes. For comparison, we also give performances of some

higher-order masking schemes for the AES s-box, as well as performances of existing schemes for DES and PRESENT s-boxes at orders 1 and 2. For the AES s-box processing, we implemented Rivain and Prouff's method [27] and its improvement by Kim *et al.* [15]. We did not implement Genelle *et al.*'s scheme [12] since it addresses the masking of an overall AES and is not interesting while focusing on a single s-box processing. Regarding existing schemes for DES and PRESENT s-boxes, we implemented the generic methods proposed in [25] (for $d = 1$) and in [26] (for $d = 2$). We also implemented the improvement of these schemes described in [26, §3.3] that consists in treating two 4-bit outputs at the same time.³ Note that we did not implement the table re-computation method (for $d = 1$) since it only makes sense for an overall cipher and not for a single s-box processing.

Table 4 lists the timing/memory performances of the different implementations. We wrote the codes in assembly language for an 8051 based 8-bit architecture with bit-addressable memory. ROM consumptions (*i.e.* code sizes) are not listed since they are not prohibitive.

As expected, the cyclotomic method is very efficient when applied to protect the PRESENT s-box. The small input dimension of the s-box indeed implies a low masking complexity (equal to 3). Moreover, it enables to tabulate the multiplication over \mathbb{F}_{16} . At first order, it is even slightly better than the method in [25] (or its improvement). At second order, the cost of the secure multiplications involved in the cyclotomic method is approximatively doubled, which explains that the overall cost is multiplied by 1.8. This makes it less efficient than [25] and [26], which are less impacted by the increase of the masking order from 1 to 2. At third order, our method is the only one. The number of cycles staying small (630), Table 4 shows that achieving resistance against 3rd-order side-channel analysis is realistic for an implementation of PRESENT on a 8051 architecture. For DES s-boxes, the parity-split method is less efficient than the state-of-the art methods for $d = 1, 2$. This is an expected consequence of the high number of nonlinear multiplications (here 10) achieved with the parity-split method in dimension 6 and of the fact that the field multiplications can no longer be tabulated (and must therefore be computed thanks to log/alog look-up tables). At third order, the timing efficiency of the method becomes very low. The masked s-box processing is 5 (resp. 10) times slower than the efficiency of the AES s-box protected thanks to [15] (resp. [27]), though its input dimension is smaller.

The ranking of the timing efficiencies for AES, DES and PRESENT s-boxes is correlated to the number of nonlinear multiplications in the used scheme (3, 4-5, and 10, for PRESENT, AES and DES respectively) which underline the soundness of the masking complexity criterion. Therefore, while selecting an s-box for a block cipher design, one should favor an s-box with small masking complexity if side-channel attacks are taken into account.

³ This improvement is only described in [26] for $d = 2$ but it can be applied likewise to the 1st-order scheme of [25].

Table 4. Comparison of secure s-box implementations

Method		Reference	cycles	RAM (bytes)
First Order Masking				
1.	AES s-box	[27]	533	10
2.	AES s-box	[15]	320	14
3.	DES s-box	Simple version [25]	1096	2
4.	DES s-box	Improved version [25] & [26]	439	14
5.	DES s-box	this paper	4100	50
6.	PRESENT s-box	Simple Version [25]	281	2
7.	PRESENT s-box	Improved Version [25] & [26]	231	14
4.	PRESENT s-box	this paper	220	18
Second Order Masking				
1.	AES s-box	[27]	832	18
2.	AES s-box	[15]	594	24
3.	DES s-box	Simple version [26]	1045	69
4.	DES s-box	Improved version [26]	652	39
5.	DES s-box	this paper	7000	78
6.	PRESENT s-box	Simple Version [26]	277	21
7.	PRESENT s-box	Improved Version [26]	284	15
8.	PRESENT s-box	this paper	400	31
Third Order Masking				
1.	AES s-box	[27]	1905	28
2.	AES s-box	[15]	965	38
3.	DES s-box	this paper	10500	108
4.	PRESENT s-box	this paper	630	44

6 Discussion

In previous sections we have introduced the first schemes that can be used to mask any s-box at any order with fair performances in software. In particular, these schemes enable to apply higher-order masking on random s-boxes (*e.g.* the DES s-boxes) which have no specific mathematical structure. Prior to our work, the only existing methods were the circuit-oriented proposals of Ishai *et al.* [14] and of Faust *et al.* [10]. The main purpose of these works was a proof of concept for applying higher-order masking to circuits with formal security proofs, but they did not address efficient implementation. A direct application of [14] or [10] to a block cipher consists in taking its Boolean representation and in replacing every XOR and AND with $O(d)$ and $O(d^2)$ logical operations respectively (where d is the masking order). Applying such a strategy in software leads to inefficient implementation as the Boolean representation of an s-box includes a huge number of nonlinear gates (with a $O(d^2)$ factor to be protected). Compared to these techniques, our schemes achieve significant improvements. These are obtained by starting from the field representation of the s-box and applying methods to significantly reduce the number of nonlinear multiplications compared to the Boolean representation of the s-box. For instance, we have shown that a DES

s-box can be computed with 10 nonlinear multiplications whereas its Boolean representation involves several dozens of logical AND operations.

We believe that our work opens up new avenues for research in block cipher implementations and side-channel security. In particular, the issue of designing s-boxes with low masking complexity and good cryptographic criteria is still to be investigated. On the other hand, our work could be extended to take into account more general definitions of the masking complexity. Indeed Definition 1 is software oriented and hence does not encompass the hardware case. As discussed above, the complexity of masking in hardware merely depends on the number of nonlinear gates [10, 14], that is on the number of nonlinear multiplications in the (n -variate) s-box representation over \mathbb{F}_2 , the so-called *algebraic normal form*. One may also want to minimize the number of nonlinear multiplications in the (ℓ -variate) s-box representation over \mathbb{F}_{2^k} for some k (and $\ell = \lceil n/k \rceil$). This approach has actually already been followed in [15], where Kim *et al.* speeds up the scheme in [27] by using the fact that the AES s-box can be processed with 5 nonlinear multiplications over \mathbb{F}_{16} rather than 4 nonlinear multiplications over \mathbb{F}_{256} . Although requiring an additional nonlinear multiplication, the resulting implementation is faster since multiplications over \mathbb{F}_{16} can be tabulated while multiplications over \mathbb{F}_{256} are computed based on the slower log/alog approach. These observations motivate the following — more general — definition of the masking complexity.

Definition 3 (Masking Complexity). *Let m , n and k be three integers such that $m, k \leq n$. The masking complexity of a (n, m) s-box over \mathbb{F}_{2^k} is the minimal number of nonlinear multiplications required to evaluate its polynomial representation over \mathbb{F}_{2^k} .*

Here again, the masking complexity is independent of the representation of \mathbb{F}_{2^k} since one can go from one representation to another without any nonlinear multiplication. The issue of finding efficient methods with respect to the masking complexity over a smaller field \mathbb{F}_{2^k} is left open for further researches.

7 Conclusion

In this paper we have introduced new generic higher-order masking schemes for s-boxes with efficient software implementation. Specifically, we have extended the Rivain and Prouff's approach for the AES s-box to any s-box. The method consists in masking the polynomial representation of the s-box over \mathbb{F}_{2^n} where n is the input dimension. As argued, the complexity of this method mainly depends on the number of nonlinear multiplications involved in the polynomial representation (*i.e.* multiplications which are not squares nor scalar multiplications). We have then introduced the masking complexity parameter for an s-box as the minimal number of nonlinear multiplications required for its evaluation. We have provided the exact values of this parameter for the set of power functions and upper bounds for all s-boxes. Namely, we have presented optimal methods to mask power functions and efficient heuristics for the general case. Eventually we

have applied our schemes to the DES s-boxes and to the PRESENT s-box and we have provided implementation results. Our work stresses interesting open issues for further research. Among them the design of s-boxes taking into account the masking complexity criterion and the extension of our approach to masking over \mathbb{F}_{2^k} with $k < n$ (e.g. for efficient hardware implementations) are of particular interest.

References

1. Akkar, M.-L., Courtois, N., Duteuil, R., Goubin, L.: A Fast and Secure Implementation of Sflash. In: Desmedt, Y.G. (ed.) PKC 2003. LNCS, vol. 2567, pp. 267–278. Springer, Heidelberg (2002)
2. Akkar, M.-L., Giraud, C.: An Implementation of DES and AES, Secure against Some Attacks. In: Koç, Ç.K., Naccache, D., Paar, C. (eds.) CHES 2001. LNCS, vol. 2162, pp. 309–318. Springer, Heidelberg (2001)
3. Blakley, G.: Safeguarding cryptographic keys. In: National Comp. Conf., June 1979, vol. 48, pp. 313–317. AFIPS Press, New York (1979)
4. Blömer, J., Guajardo, J., Krummel, V.: Provably Secure Masking of AES. In: Handschuh, H., Hasan, M.A. (eds.) SAC 2004. LNCS, vol. 3357, pp. 69–83. Springer, Heidelberg (2004)
5. Bogdanov, A., Knudsen, L.R., Leander, G., Paar, C., Poschmann, A., Robshaw, M.J.B., Seurin, Y., Vikkelsoe, C.: PRESENT: An Ultra-Lightweight Block Cipher. In: Paillier, P., Verbaauwhede, I. (eds.) CHES 2007. LNCS, vol. 4727, pp. 450–466. Springer, Heidelberg (2007)
6. Brier, E., Clavier, C., Olivier, F.: Correlation Power Analysis with a Leakage Model. In: Joye, M., Quisquater, J.-J. (eds.) CHES 2004. LNCS, vol. 3156, pp. 16–29. Springer, Heidelberg (2004)
7. Chari, S., Jutla, C.S., Rao, J.R., Rohatgi, P.: Towards Sound Approaches to Counteract Power-Analysis Attacks. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 398–412. Springer, Heidelberg (1999)
8. Coron, J.-S., Prouff, E., Rivain, M.: Side Channel Cryptanalysis of a Higher Order Masking Scheme. In: Paillier, P., Verbaauwhede, I. (eds.) CHES 2007. LNCS, vol. 4727, pp. 28–44. Springer, Heidelberg (2007)
9. Eve, J.: The evaluation of polynomials. *Comm. ACM* 6(1), 17–21 (1964)
10. Faust, S., Rabin, T., Reyzin, L., Tromer, E., Vaikuntanathan, V.: Protecting Circuits from Leakage: the Computationally-Bounded and Noisy Cases. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 135–156. Springer, Heidelberg (2010)
11. FIPS PUB 46. The Data Encryption Standard. National Bureau of Standards (January 1977)
12. Genelle, L., Prouff, E., Quisquater, M.: Thwarting Higher-Order Side Channel Analysis with Additive and Multiplicative Maskings. In: Preneel, B., Takagi, T. (eds.) CHES 2011. LNCS, vol. 6917, pp. 240–255. Springer, Heidelberg (2011)
13. Goubin, L., Patarin, J.: DES and Differential Power Analysis. In: Koç, Ç.K., Paar, C. (eds.) CHES 1999. LNCS, vol. 1717, pp. 158–172. Springer, Heidelberg (1999)
14. Ishai, Y., Sahai, A., Wagner, D.: Private Circuits: Securing Hardware against Probing Attacks. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 463–481. Springer, Heidelberg (2003)

15. Kim, H., Hong, S., Lim, J.: A Fast and Provably Secure Higher-Order Masking of AES S-Box. In: Preneel, B., Takagi, T. (eds.) CHES 2011. LNCS, vol. 6917, pp. 95–107. Springer, Heidelberg (2011)
16. Knuth, D.: The Art of Computer Programming, 3rd edn., vol. 2. Addison-Wesley (1988)
17. Knuth, D.E.: Evaluation of polynomials by computers. *Comm. ACM* 5(12), 595–599 (1962)
18. Kocher, P.C., Jaffe, J., Jun, B.: Differential Power Analysis. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 388–397. Springer, Heidelberg (1999)
19. Mangard, S., Popp, T., Gammel, B.M.: Side-Channel Leakage of Masked CMOS Gates. In: Menezes, A. (ed.) CT-RSA 2005. LNCS, vol. 3376, pp. 351–365. Springer, Heidelberg (2005)
20. Mangard, S., Pramstaller, N., Oswald, E.: Successfully Attacking Masked AES Hardware Implementations. In: Rao, J.R., Sunar, B. (eds.) CHES 2005. LNCS, vol. 3659, pp. 157–171. Springer, Heidelberg (2005)
21. 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)
22. 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)
23. Nikova, S., Rijmen, V., Schläffer, M.: Secure Hardware Implementation of Non-linear Functions in the Presence of Glitches. In: Lee, P.J., Cheon, J.H. (eds.) ICISC 2008. LNCS, vol. 5461, pp. 218–234. Springer, Heidelberg (2009)
24. Popp, T., Kirschbaum, M., Zefferer, T., Mangard, S.: Evaluation of the Masked Logic Style MDPL on a Prototype Chip. In: Paillier, P., Verbauwhede, I. (eds.) CHES 2007. LNCS, vol. 4727, pp. 81–94. Springer, Heidelberg (2007)
25. Prouff, E., Rivain, M.: A Generic Method for Secure SBox Implementation. In: Kim, S., Yung, M., Lee, H.-W. (eds.) WISA 2007. LNCS, vol. 4867, pp. 227–244. Springer, Heidelberg (2008)
26. Rivain, M., Dottax, E., Prouff, E.: Block Ciphers Implementations Provably Secure Against Second Order Side Channel Analysis. In: Nyberg, K. (ed.) FSE 2008. LNCS, vol. 5086, pp. 127–143. Springer, Heidelberg (2008)
27. Rivain, M., Prouff, E.: Provably Secure Higher-Order Masking of AES. In: Mangard, S., Standaert, F.-X. (eds.) CHES 2010. LNCS, vol. 6225, pp. 413–427. Springer, Heidelberg (2010)
28. Satoh, A., Morioka, S., Takano, K., Munetoh, S.: A Compact Rijndael Hardware Architecture with S-Box Optimization. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 239–254. Springer, Heidelberg (2001)
29. Schramm, K., Paar, C.: Higher Order Masking of the AES. In: Pointcheval, D. (ed.) CT-RSA 2006. LNCS, vol. 3860, pp. 208–225. Springer, Heidelberg (2006)
30. Shamir, A.: How to Share a Secret. *Commun. ACM* 22(11), 612–613 (1979)
31. von zur Gathen, J.: Efficient and Optimal Exponentiation in Finite Fields. *Computational Complexity* 1, 360–394 (1991)

A Masking Complexity of Power Functions

Table 5 summarizes the masking complexity classes $(\mathcal{M}_k^n)_k$ for dimensions n in the set $\{3, 5, 7, 9, 10, 11\}$.

Table 5. Cyclotomic classes for $n \in \{3, 5, 7, 9, 10, 11\}$ w.r.t. the masking complexity k

k	Cyclotomic classes in \mathcal{M}_k^n
$n = 3$	
0	$C_0 = \{0\}, C_1 = \{1, 2, 4\}$
1	$C_3 = \{3, 6, 5\}$
$n = 5$	
0	$C_0 = \{0\}, C_1 = \{1, 2, 4, 8, 16\}$
1	$C_3 = \{3, 6, 12, 24, 17\}, C_5 = \{5, 10, 20, 9, 18\}$
2	$C_7 = \{7, 14, 28, 25, 19\}, C_{11} = \{11, 22, 13, 26, 21\}, C_{15} = \{15, 30, 29, 27, 23\}$
$n = 7$	
0	$C_0 = \{0\}, C_1 = \{1, 2, 4, 8, 16, 32, 64\}$
1	$C_3 = \{3, 6, 12, 24, 48, 96, 65\}, C_5 = \{5, 10, 20, 40, 80, 33, 66\},$ $C_9 = \{9, 18, 36, 72, 17, 34, 68\}$
2	$C_7 = \{7, 14, 28, 56, 112, 97, 67\}, C_{11} = \{11, 22, 44, 88, 49, 98, 69\},$ $C_{13} = \{13, 26, 52, 104, 81, 35, 70\}, C_{15} = \{15, 30, 60, 120, 113, 99, 71\},$ $C_{19} = \{19, 38, 76, 25, 50, 100, 73\}, C_{21} = \{21, 42, 84, 41, 82, 37, 74\},$ $C_{27} = \{27, 54, 108, 89, 51, 102, 77\}, C_{43} = \{43, 86, 45, 90, 53, 106, 85\}$
3	$C_{23} = \{23, 46, 92, 57, 114, 101, 75\}, C_{29} = \{29, 58, 116, 105, 83, 39, 78\},$ $C_{31} = \{31, 62, 124, 121, 115, 103, 79\}, C_{47} = \{47, 94, 61, 122, 117, 107, 87\},$ $C_{55} = \{55, 110, 93, 59, 118, 109, 91\}, C_{63} = \{63, 126, 125, 123, 119, 111, 95\}$
$n = 9$	
0	C_0, C_1
1	C_3, C_5, C_9, C_{17}
2	$C_7, C_{11}, C_{13}, C_{15}, C_{19}, C_{21}, C_{25}, C_{27}, C_{35}, C_{37}, C_{41}, C_{45}, C_{51}, C_{73}, C_{75}, C_{83}, C_{85}$
3	$C_{23}, C_{29}, C_{31}, C_{39}, C_{43}, C_{47}, C_{53}, C_{55}, C_{57}, C_{59}, C_{61}, C_{63}, C_{71}, C_{75}, C_{77},$ $C_{63}, C_{75}, C_{77}, C_{79}, C_{87}, C_{91}, C_{93}, C_{95}, C_{101}, C_{103}, C_{95}, C_{103}, C_{105}, C_{107}, C_{109}, C_{111}, C_{115},$ $C_{117}, C_{119}, C_{121}, C_{123}, C_{125}, C_{149}, C_{151}, C_{155}, C_{157}, C_{167}, C_{171}, C_{173}, C_{175}, C_{179},$ $C_{181}, C_{183}, C_{187}, C_{189}, C_{205}, C_{207}, C_{213}, C_{215}, C_{219}, C_{221}, C_{231}, C_{235}, C_{237}, C_{245},$ $C_{255}, C_{341}, C_{347}, C_{363}, C_{447}, C_{495}$
4	$C_{191}, C_{223}, C_{239}$
$n = 10$	
0	C_0, C_1
1	$C_3, C_5, C_9, C_{17}, C_{33}$
2	$C_7, C_{11}, C_{13}, C_{15}, C_{19}, C_{21}, C_{25}, C_{27}, C_{35}, C_{37},$ $C_{41}, C_{45}, C_{49}, C_{51}, C_{69}, C_{73}, C_{85}, C_{99}, C_{147}, C_{165}$
3	$C_{23}, C_{29}, C_{31}, C_{39}, C_{43}, C_{47}, C_{53}, C_{55}, C_{57}, C_{59}, C_{61}, C_{63}, C_{71}, C_{75}, C_{77},$ $C_{79}, C_{83}, C_{87}, C_{89}, C_{91}, C_{93}, C_{95}, C_{101}, C_{103}, C_{105}, C_{107}, C_{109}, C_{111}, C_{113},$ $C_{115}, C_{117}, C_{119}, C_{121}, C_{123}, C_{125}, C_{149}, C_{151}, C_{155}, C_{157}, C_{167}, C_{171}, C_{173}, C_{175}, C_{179},$ $C_{181}, C_{183}, C_{187}, C_{189}, C_{205}, C_{207}, C_{213}, C_{215}, C_{219}, C_{221}, C_{231}, C_{235}, C_{237}, C_{245},$ $C_{255}, C_{341}, C_{347}, C_{363}, C_{447}, C_{495}$
4	$C_{127}, C_{159}, C_{191}, C_{223}, C_{239}, C_{247}, C_{251}, C_{253}, C_{343},$ $C_{351}, C_{367}, C_{375}, C_{379}, C_{383}, C_{439}, C_{479}, C_{511}$
$n = 11$	
0	C_0, C_1
1	$C_3, C_5, C_9, C_{17}, C_{33}$
2	$C_7, C_{11}, C_{13}, C_{15}, C_{19}, C_{21}, C_{25}, C_{27}, C_{35}, C_{37}, C_{41}, C_{45}, C_{49}, C_{51},$ $C_{67}, C_{69}, C_{73}, C_{81}, C_{85}, C_{99}, C_{137}, C_{153}, C_{163}, C_{165}, C_{293}$
3	$C_{23}, C_{29}, C_{31}, C_{39}, C_{43}, C_{47}, C_{53}, C_{55}, C_{57}, C_{59}, C_{61}, C_{63}, C_{71}, C_{75}, C_{77},$ $C_{79}, C_{83}, C_{87}, C_{89}, C_{91}, C_{93}, C_{95}, C_{101}, C_{103}, C_{105}, C_{107}, C_{109}, C_{111}, C_{113},$ $C_{115}, C_{117}, C_{119}, C_{121}, C_{123}, C_{125}, C_{139}, C_{141}, C_{143}, C_{147}, C_{149}, C_{151}, C_{155},$ $C_{157}, C_{167}, C_{169}, C_{171}, C_{173}, C_{175}, C_{179}, C_{181}, C_{185}, C_{187}, C_{189}, C_{199}, C_{201},$ $C_{203}, C_{205}, C_{207}, C_{211}, C_{213}, C_{217}, C_{219}, C_{221}, C_{229}, C_{231}, C_{243}, C_{245},$ $C_{255}, C_{295}, C_{299}, C_{301}, C_{307}, C_{309}, C_{311}, C_{315}, C_{317}, C_{331}, C_{333}, C_{335},$ $C_{343}, C_{347}, C_{359}, C_{363}, C_{365}, C_{379}, C_{411}, C_{423}, C_{427}, C_{429}, C_{339}, C_{341},$ $C_{437}, C_{439}, C_{469}, C_{495}, C_{683}, C_{703}, C_{879}, C_{887}$
4	$C_{127}, C_{159}, C_{183}, C_{191}, C_{215}, C_{223}, C_{233}, C_{235}, C_{237}, C_{239}, C_{247}, C_{249}, C_{251},$ $C_{253}, C_{303}, C_{319}, C_{349}, C_{351}, C_{367}, C_{371}, C_{373}, C_{375}, C_{381}, C_{383},$ $C_{413}, C_{415}, C_{431}, C_{443}, C_{445}, C_{447}, C_{463}, C_{471}, C_{475}, C_{477}, C_{479}, C_{491},$ $C_{493}, C_{501}, C_{503}, C_{507}, C_{509}, C_{511}, C_{687}, C_{695}, C_{699}, C_{727}, C_{731}, C_{735}, C_{751},$ $C_{759}, C_{763}, C_{767}, C_{895}, C_{959}, C_{991}, C_{1023}$

Recursive Diffusion Layers for Block Ciphers and Hash Functions

Mahdi Sajadieh¹, Mohammad Dakhilalian¹,
Hamid Mala², and Pouyan Sepehrdad^{3,*}

¹ Cryptography & System Security Research Laboratory,
Department of Electrical and Computer Engineering,
Isfahan University of Technology, Isfahan, Iran
`sadjadieh@ec.iut.ac.ir`, `mdalian@cc.iut.ac.ir`

² Department of Information Technology Engineering,
University of Isfahan, Isfahan, Iran

`h.mala@eng.ui.ac.ir`

³ EPFL, Lausanne, Switzerland
`pouyan.sepehrdad@epfl.ch`

Abstract. Many modern block ciphers use maximum distance separable (MDS) matrices as the main part of their diffusion layers. In this paper, we propose a new class of diffusion layers constructed from several rounds of Feistel-like structures whose round functions are linear. We investigate the requirements of the underlying linear functions to achieve the maximal branch number for the proposed 4×4 words diffusion layer. The proposed diffusion layers only require word-level XORs, rotations, and they have simple inverses. They can be replaced in the diffusion layer of the block ciphers MMB and Hierocrypt to increase their security and performance, respectively. Finally, we try to extend our results for up to 8×8 words diffusion layers.

Keywords: Block ciphers, Diffusion layer, Branch number, Provable security.

1 Introduction

Block ciphers are one of the most important building blocks in many security protocols. Modern block ciphers are cascades of several rounds and each round consists of confusion and diffusion layers. In many block ciphers, non-linear substitution boxes (S-boxes) form the confusion layer, and a linear transformation provides the required diffusion. The diffusion layer plays an efficacious role in providing resistance against the most well-known attacks on block ciphers, such as differential cryptanalysis (DC) [2] and linear cryptanalysis (LC) [10].

In 1994, Vaudenay [15,16] suggested using MDS matrices in cryptographic primitives to produce what he called multipermutations, not-necessarily linear

* This work has been supported in part by the European Commission through the ICT program under contract ICT-2007-216646 ECRYPT II.

functions with this same property. These functions have what he called perfect diffusion. He showed how to exploit imperfect diffusion to cryptanalyze functions that are not multipermutations. This notion was later used by Daemen named as the branch number. Block ciphers exploiting diffusion layers with small branch number may suffer from critical weaknesses against DC and LC, even though their substitution layers consist of S-boxes with strong non-linear properties.

Two main strategies for designing block ciphers are Feistel-like and substitution permutation network (SPN) structures. In the last 2 decades, from these two families several structures have been proposed with provable security against DC and LC. Three rounds of Feistel structure [11,12], five rounds of RC6-like structure [9] and SDS (substitution-diffusion-substitution) structure with a perfect or almost perfect diffusion layer are examples of such structures [8].

1.1 Notations

Let \mathbf{x} be an array of s n -bit elements $\mathbf{x} = [x_{0(n)}, x_{1(n)}, \dots, x_{s-1(n)}]$. The number of non-zero elements in \mathbf{x} is denoted by $w(\mathbf{x})$ and is known as the Hamming weight of \mathbf{x} . The following notations are used throughout this paper:

- \oplus : The bit-wise XOR operation
- $\&$: The bit-wise AND operation
- L_i : Any linear function
- ℓ_i : The linear operator corresponding to the linear function L_i
- $(L_1 \oplus L_2)(x)$: $L_1(x) \oplus L_2(x)$
- $L_1 L_2(x)$: $L_1(L_2(x))$
- $L_1^2(x)$: $L_1(L_1(x))$
- $I(\cdot)$ function : Identity function, $I(x) = x$
- $x \gg m$ ($x \ll m$) : Shift of a bit string x by m bits to the right (left)
- $x \ggg m$ ($x \lll m$) : Circular shift of a bit string x by m bits to the right (left)
- $|\cdot|$: Determinant of a matrix in $\mathbf{GF}(2)$
- $a|b$: Concatenation of two bit strings a and b
- $x_{(n)}$: An n -bit value x

For a diffusion layer D applicable on \mathbf{x} , we have the following definitions:

Definition 1 ([4]). *The differential branch number of a linear diffusion layer D is defined as:*

$$\beta_d(D) = \min_{\mathbf{x} \neq 0} \{w(\mathbf{x}) + w(D(\mathbf{x}))\}$$

We know that the linear function D can be shown as a binary matrix \mathbf{B} , and D^t is a linear function obtained from \mathbf{B}^t , where \mathbf{B}^t is the transposition of \mathbf{B} .

Definition 2 ([4]). *The linear branch number of a linear diffusion layer D is defined as:*

$$\beta_l(D) = \min_{\mathbf{x} \neq 0} \{w(\mathbf{x}) + w(D^t(\mathbf{x}))\}$$

It is well known that for a diffusion layer acting on s -word inputs, the maximal β_d and β_l are $s + 1$ [4]. A diffusion layer D taking its maximal β_d and β_l is called a perfect or MDS diffusion layer. Furthermore, a diffusion layer with $\beta_d = \beta_l = s$ is called an almost perfect diffusion layer [8].

1.2 Our Contribution

In this paper, we define the notion of a *recursive diffusion layer* and propose a method to construct such perfect diffusion layers.

Definition 3. A diffusion layer D with s words x_i as the input, and s words y_i as the output is called a recursive diffusion layer if it can be represented in the following form:

$$D : \begin{cases} y_0 = x_0 \oplus F_0(x_1, x_2, \dots, x_{s-1}) \\ y_1 = x_1 \oplus F_1(x_2, x_3, \dots, x_{s-1}, y_0) \\ \vdots \\ y_{s-1} = x_{s-1} \oplus F_{s-1}(y_0, y_1, \dots, y_{s-2}) \end{cases} \tag{1}$$

where F_0, F_1, \dots, F_{s-1} are arbitrary functions.

As an example, consider a 2-round Feistel structure with a linear round function L as a recursive diffusion layer with $s = 2$. The input-output relation for this diffusion layer is:

$$D : \begin{cases} y_0 = x_0 \oplus L(x_1) \\ y_1 = x_1 \oplus L(y_0) \end{cases}$$

The quarter-round function of Salsa20 is also an example of a non-linear recursive diffusion layer [1].

$$D : \begin{cases} y_1 = x_1 \oplus ((x_0 + x_3) \lll 7) \\ y_2 = x_2 \oplus ((x_0 + y_1) \lll 9) \\ y_3 = x_3 \oplus ((y_1 + y_2) \lll 13) \\ y_0 = x_0 \oplus ((y_2 + y_3) \lll 18) \end{cases}$$

Also, the lightweight hash function PHOTON [5] and the block cipher LED [6] use MDS matrices based on Eq. (1). In these ciphers, an $m \times m$ MDS matrix \mathbf{B}^m was designed based on the following matrix \mathbf{B} for the performance purposes:

$$\mathbf{B} = \begin{pmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & & \ddots & & \\ 0 & 0 & 0 & \dots & 1 \\ Z_0 & Z_1 & Z_2 & \dots & Z_{m-1} \end{pmatrix}$$

By matrix \mathbf{B} , one elements of m inputs is updated and other elements are shifted. If we use \mathbf{B}^m , all inputs are updated, but we must check if this matrix is MDS. One example for $m = 4$ is the PHOTON matrix working over $\text{GF}(2^8)$:

$$\mathbf{B} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 2 & 1 & 4 \end{pmatrix} \Rightarrow \mathbf{B}^4 = \begin{pmatrix} 1 & 2 & 1 & 4 \\ 4 & 9 & 6 & 17 \\ 17 & 38 & 24 & 66 \\ 66 & 149 & 100 & 11 \end{pmatrix}$$

In this paper, we propose a new approach to design linear recursive diffusion layers with the maximal branch number in which F_i 's are composed of one or two linear functions and a number of XOR operations. The design of the proposed diffusion layer is based on the invertibility of some simple linear functions in $\text{GF}(2)$. Linear functions in this diffusion layer can be designed to be low-cost for different sizes of the input words, thus the proposed diffusion layer might be appropriate for resource-constrained devices, such as RFID tags. Although these recursive diffusion layers are not involutory, they have similar inverses with the same computational complexity. Another approach which is not recursive was picked by Junod and Vaudenay in [7] to design efficient MDS matrices.

This paper proceeds as follows: In Section 2, we introduce the general structure of our proposed recursive diffusion layer. Then, for one of its instances, we systematically investigate the required conditions for the underlying linear function to achieve the maximal branch number. In Section 3, we propose some other recursive diffusion layers with less than 8 input words and only one linear function. We use two linear functions to have a perfect recursive diffusion layer for $s > 4$ in Section 4. Finally, we conclude the paper in Section 5.

2 The Proposed Diffusion Layer

In this section, we introduce a new perfect linear diffusion layer with a recursive structure. The diffusion layer D takes s words x_i for $i = \{0, 1, \dots, s-1\}$ as input, and returns s words y_i for $i = \{0, 1, \dots, s-1\}$ as output. So, we can represent this diffusion layer as:

$$y_0|y_1|\dots|y_{s-1} = D(x_0|x_1|\dots|x_{s-1})$$

The first class of the proposed diffusion layer D is represented in Fig. 1, where L is a linear function, $\alpha_k, \beta_k \in \{0, 1\}$, $\alpha_0 = 1$, and $\beta_0 = 0$.

This diffusion layer can be represented in the form of Eq. (1) in which the F_i functions are all the same and can be represented as

$$F_i(x_1, x_2, \dots, x_{s-1}) = \bigoplus_{j=1}^{s-1} \alpha_j x_j \oplus L \left(\bigoplus_{j=1}^{s-1} \beta_j x_j \right)$$

To guarantee the maximal branch number for D , the linear function L and the coefficients α_j and β_j must satisfy some necessary conditions. Conditions on L are expressed in this section and those of α_j 's and β_j 's are expressed in Section 3. The diffusion layer described by Eq. (2) is an instance that satisfies the necessary conditions on α_j and β_j with $s = 4$. In the rest of this section, we concentrate

```

1: Input :  $s$   $n$ -bit words  $x_0, \dots, x_{s-1}$ 
2: Output :  $s$   $n$ -bit words  $y_0, \dots, x_{s-1}$ 
3: for  $i = 0$  to  $s - 1$  do
4:    $y_i = x_i$ 
5: end for
6: for  $i = 0$  to  $s - 1$  do
7:    $y_i = y_i \oplus \left( \bigoplus_{j=0, j \neq i}^{s-1} \alpha_{[(j-i) \bmod s]} y_j \right) \oplus L \left( \bigoplus_{j=0, j \neq i}^{s-1} \beta_{[(j-i) \bmod s]} y_j \right)$ 
8: end for
    
```

Fig. 1. The first class of the recursive diffusion layers

on the diffusion layers of this form and show that we can find invertible linear functions L such that D becomes a perfect diffusion layer.

$$D : \begin{cases} y_0 = x_0 \oplus x_2 \oplus x_3 \oplus L(x_1 \oplus x_3) \\ y_1 = x_1 \oplus x_3 \oplus y_0 \oplus L(x_2 \oplus y_0) \\ y_2 = x_2 \oplus y_0 \oplus y_1 \oplus L(x_3 \oplus y_1) \\ y_3 = x_3 \oplus y_1 \oplus y_2 \oplus L(y_0 \oplus y_2) \end{cases} \quad (2)$$

As shown in Fig. 2, this diffusion layer has a Feistel-like (GFN) structure, i.e.,

$$F_0(x_1, x_2, x_3) = x_2 \oplus x_3 \oplus L(x_1 \oplus x_3)$$

and for each $i > 0$, y_i is obtained by $(x_i, x_{i+1}, \dots, x_{s-1})$ and $(y_0, y_1, \dots, y_{i-1})$.

The inverse transformation, D^{-1} , has a very simple structure and does not require the inversion of the linear function L . Based on the recursive nature of D , if we start from the last equation of Eq. (2), x_3 is immediately obtained from y_i 's. Then knowing x_3 and y_i 's, we immediately obtain x_2 from the third line of Eq. (2). x_1 and x_0 can be obtained in the same way. Thus, the inverse of D is:

$$D^{-1} : \begin{cases} x_3 = y_3 \oplus y_1 \oplus y_2 \oplus L(y_0 \oplus y_2) \\ x_2 = y_2 \oplus y_0 \oplus y_1 \oplus L(x_3 \oplus y_1) \\ x_1 = y_1 \oplus x_3 \oplus y_0 \oplus L(x_2 \oplus y_0) \\ x_0 = y_0 \oplus x_2 \oplus x_3 \oplus L(x_1 \oplus x_3) \end{cases}$$

D and D^{-1} are different, but they have the same structure and properties. To show that D has the maximal branch number, first we introduce some lemmas and theorems.

Theorem 4 ([4]). *A Boolean function F has maximal differential branch number if and only if it has maximal linear branch number.*

As a result of Theorem 4, if we prove that the diffusion layer D represented in Eq. (2) has the maximal differential branch number, its linear branch number will be maximal too. Thus, in the following, we focus on the differential branch number.




Fig. 2. The proposed recursive diffusion layer of Eq. (2)

Lemma 5. For m linear functions L_1, L_2, \dots, L_m , the proposition

$$a \neq 0 \Rightarrow L_1(a) \oplus L_2(a) \oplus \dots \oplus L_m(a) \neq 0$$

implies that the linear function $L_1 \oplus L_2 \oplus \dots \oplus L_m$ is invertible.

Proof. We know that $(L_1 \oplus L_2 \oplus \dots \oplus L_m)(x)$ is a linear function and it can be represented as a binary matrix \mathbf{M} . So, \mathbf{M} is invertible if and only if $|\mathbf{M}| \neq 0$. \square

Lemma 6. Assume the linear operator ℓ_i corresponds to the linear function $L_i(x)$. If the linear operator ℓ_3 can be represented as the multiplication of two operators ℓ_1 and ℓ_2 , then the corresponding linear function $L_3(x) = L_2(L_1(x))$ is invertible if and only if the linear functions $L_1(x)$ and $L_2(x)$ are invertible.

Proof. If $L_1(x)$ and $L_2(x)$ are invertible, clearly $L_3(x)$ is invertible too. On the other hand, if $L_3(x)$ is invertible then $L_1(x)$ must be invertible, otherwise there are distinct x_1 and x_2 such that $L_1(x_1) = L_1(x_2)$. Thus, $L_3(x_1) = L_2(L_1(x_1)) = L_2(L_1(x_2)) = L_3(x_2)$ which contradicts the invertibility of $L_3(x)$. The invertibility of $L_2(x)$ is proved in the same way. \square

Example 1. We can rewrite the linear function $L_3(x) = L^3(x) \oplus x$ ($\ell_3 = \ell^3 \oplus I$) as $L_3(x) = L_1(L_2(x))$, where $L_1(x) = L(x) \oplus x$ ($\ell_1 = \ell \oplus I$) and $L_2(x) = L^2(x) \oplus L(x) \oplus x$ ($\ell_2 = \ell^2 \oplus \ell \oplus I$). Thus, the invertibility of $L_3(x)$ is equivalent to the invertibility of the two linear functions $L_1(x)$ and $L_2(x)$.

Theorem 7. For the diffusion layer represented in Eq. (2), if the four linear functions $L(x)$, $x \oplus L(x)$, $x \oplus L^3(x)$, and $x \oplus L^7(x)$ are invertible, then this diffusion layer is perfect.

Proof. We show that the differential branch number of this diffusion layer is 5. First, the 4 words of the output are directly represented as functions of the 4 words of the input:

$$D : \begin{cases} y_0 = x_0 \oplus L(x_1) \oplus x_2 \oplus x_3 \oplus L(x_3) \\ y_1 = x_0 \oplus L(x_0) \oplus x_1 \oplus L(x_1) \oplus L^2(x_1) \oplus x_2 \oplus L^2(x_3) \\ y_2 = L^2(x_0) \oplus x_1 \oplus L(x_1) \oplus L^3(x_1) \oplus x_2 \oplus L(x_2) \oplus x_3 \oplus L^2(x_3) \oplus L^3(x_3) \\ y_3 = x_0 \oplus L^2(x_0) \oplus L^3(x_0) \oplus L(x_1) \oplus L^2(x_1) \oplus L^3(x_1) \oplus L^4(x_1) \\ \quad \oplus L(x_2) \oplus L^2(x_2) \oplus L^2(x_3) \oplus L^4(x_3) \end{cases} \tag{3}$$

Now, we show that if the number of active (non-zero) words in the input is m , where $m = 1, 2, 3, 4$, then the number of non-zero words in the output is greater than or equal to $5 - m$. The diffusion layer represented in Eq. (2) is invertible. Consider $m = 4$, then all of the 4 words in the input are active, and we are sure at least one of the output words is active too. Thus the theorem is correct for $m = 4$. The remainder of the proof is performed for the 3 cases of $w(\Delta(\mathbf{x})) = m$, for $m = 1, 2, 3$ separately. In each of these cases, some conditions are forced on the linear function L .

Case 1: $w(\Delta\mathbf{x}) = 1$

To study this case, first the subcase

$$(\Delta x_0 \neq 0, \Delta x_1 = \Delta x_2 = \Delta x_3 = 0 \text{ or } \Delta \mathbf{x} = \Delta x_0 | 0 | 0 | 0)$$

is analyzed. For this subcase, Eq. (3) is simplified to:

$$D : \begin{cases} \Delta y_0 = \Delta x_0 \\ \Delta y_1 = (I \oplus L)(\Delta x_0) \\ \Delta y_2 = L^2(\Delta x_0) \\ \Delta y_3 = (I \oplus L^2 \oplus L^3)(\Delta x_0) \end{cases}$$

If D is a perfect diffusion layer then Δy_0 , Δy_1 , Δy_2 and Δy_3 must be non-zero. Clearly, Δy_0 is non-zero, and based on Lemma 5, the conditions for Δy_1 , Δy_2 and Δy_3 to be non-zero are that the linear functions $I \oplus L$, L^2 and $I \oplus L^2 \oplus L^3$ must be invertible. Note that based on Lemma 6, the invertibility of L^2 yields the invertibility of L . Considering Lemma 6, if the other three sub-cases are studied,

it is induced that the linear functions $x \oplus L(x) \oplus L^2(x)$ and $x \oplus L(x) \oplus L^3(x)$ must also be invertible.

Case 2: $w(\Delta \mathbf{x}) = 2$

In this case, there exist exactly two active words in the input difference and we obtain some conditions on the linear function L to guarantee the branch number 5 for D . In the following, we only analyze the subcase

$$(\Delta x_0, \Delta x_1 \neq 0 \text{ and } \Delta x_2 = \Delta x_3 = 0 \text{ or } \Delta \mathbf{x} = \Delta x_0 | \Delta x_1 | 0 | 0)$$

With this assumption, Eq. (3) is simplified to:

$$D : \begin{cases} \Delta y_0 = \Delta x_0 \oplus L(\Delta x_1) \\ \Delta y_1 = (I \oplus L)(\Delta x_0) \oplus (I \oplus L \oplus L^2)(\Delta x_1) \\ \Delta y_2 = L^2(\Delta x_0) \oplus (I \oplus L \oplus L^3)(\Delta x_1) \\ \Delta y_3 = (I \oplus L^2 \oplus L^3)(\Delta x_0) \oplus (L \oplus L^2 \oplus L^3 \oplus L^4)(\Delta x_1) \end{cases} \quad (4)$$

To show that $w(\Delta \mathbf{y})$ is greater than or equal to 3, we must find some conditions on L such that if one of the Δy_i 's is zero, then the other three Δy_j 's cannot be zero. Let $\Delta y_0 = 0$, then:

$$\Delta x_0 \oplus L(\Delta x_1) = 0 \Rightarrow \Delta x_0 = L(\Delta x_1)$$

If Δx_0 is replaced in the last three equations of Eq. (4), we obtain Δy_1 , Δy_2 and Δy_3 as follows:

$$\begin{cases} \Delta y_1 = \Delta x_1 \\ \Delta y_2 = \Delta x_1 \oplus L(\Delta x_1) \\ \Delta y_3 = L^2(\Delta x_1) \end{cases}$$

Obviously, Δy_1 is not zero. Furthermore, for Δy_2 and Δy_3 to be non-zero, considering Lemma 5, we conclude that the functions $x \oplus L(x)$ and $L^2(x)$ must be invertible. This condition was already obtained in the Case 1. We continue this procedure for $\Delta y_1 = 0$.

$$\begin{aligned} \Delta y_1 = \Delta x_0 \oplus L(\Delta x_0) \oplus x_1 \oplus L(\Delta x_1) \oplus L^2(\Delta x_1) = 0 \Rightarrow \\ \Delta x_0 \oplus L(\Delta x_0) = x_1 \oplus L(\Delta x_1) \oplus L^2(\Delta x_1) \end{aligned}$$

From the previous subcase, we know that if $\Delta y_0 = 0$ then $\Delta y_1 \neq 0$. Thus we conclude that, Δy_0 and Δy_1 cannot be simultaneously zero. Therefore, by contraposition we obtain that if $\Delta y_1 = 0$ then $\Delta y_0 \neq 0$. So, we only check Δy_2 and Δy_3 . From the third equation in Eq. (4), we have:

$$\begin{aligned} (I \oplus L)(\Delta y_2) &= L^2(\Delta x_1) \oplus L^3(\Delta x_1) \oplus L^4(\Delta x_1) \oplus \Delta x_1 \\ &\quad \oplus L^2(\Delta x_1) \oplus L^3(\Delta x_1) \oplus L^4(\Delta x_1) \\ &= \Delta x_1 \end{aligned}$$

$x \oplus L(x)$ is invertible, thus we conclude that with the two active words Δx_0 and Δx_1 in the input, Δy_1 and Δy_2 cannot be zero simultaneously. With the same procedure, we can prove that Δy_1 and Δy_3 cannot be zero simultaneously.

Here we only gave the proof for the case $(\Delta x_0, \Delta x_1 \neq 0, \Delta x_2 = \Delta x_3 = 0)$. We performed the proof procedure for the other cases and no new condition was added to the previous set of conditions in Case 1.

Case 3: $w(\Delta \mathbf{x}) = 3$

In this case, assuming three active words in the input, we show that the output has at least 2 non-zero words. Here, only the case

$$(\Delta x_0, \Delta x_1, \Delta x_2 \neq 0 \text{ and } \Delta x_3 = 0 \text{ or } \Delta \mathbf{x} = \Delta x_0 | \Delta x_1 | \Delta x_2 | 0)$$

is analyzed. The result holds for the other three cases with $w(\Delta \mathbf{x}) = 3$. Let rewrite the Eq. (3) for $\Delta x_3 = 0$ as follows:

$$D : \begin{cases} \Delta y_0 = \Delta x_0 \oplus L(\Delta x_1) \oplus \Delta x_2 \\ \Delta y_1 = (I \oplus L)(\Delta x_0) \oplus (I \oplus L \oplus L^2)(\Delta x_1) \oplus \Delta x_2 \\ \Delta y_2 = L^2(\Delta x_0) \oplus (I \oplus L \oplus L^3)(\Delta x_1) \oplus (I \oplus L)(\Delta x_2) \\ \Delta y_3 = (I \oplus L^2 \oplus L^3)(\Delta x_0) \oplus (L \oplus L^2 \oplus L^3 \oplus L^4)(\Delta x_1) \oplus (L \oplus L^2)(\Delta x_2) \end{cases} \tag{5}$$

When $\Delta y_0 = \Delta y_1 = 0$, from the first 2 lines of Eq. (5), Δx_0 and Δx_1 are obtained as the function of Δx_2 .

$$\begin{cases} \Delta y_0 = \Delta x_0 \oplus L(\Delta x_1) \oplus \Delta x_2 = 0 \\ \Delta y_1 = \Delta x_0 \oplus L(\Delta x_0) \oplus \Delta x_1 \oplus L(\Delta x_1) \oplus L^2(\Delta x_1) \oplus \Delta x_2 = 0 \end{cases} \Rightarrow \begin{cases} \Delta x_1 = L(\Delta x_2) \\ \Delta x_0 = \Delta x_2 \oplus L^2(\Delta x_2) \end{cases}$$

Now, replacing $\Delta x_0 = \Delta x_2 \oplus L^2(\Delta x_2)$ and $\Delta x_1 = L(\Delta x_2)$ into Δy_2 and Δy_3 yields:

$$\begin{cases} \Delta y_2 = L^2(\Delta x_0) \oplus (I \oplus L \oplus L^3)(\Delta x_1) \oplus (I \oplus L)(\Delta x_2) = \Delta x_2 \\ \Delta y_3 = (I \oplus L^2 \oplus L^3)(\Delta x_0) \oplus (L \oplus L^2 \oplus L^3 \oplus L^4)(\Delta x_1) \oplus (L \oplus L^2)(\Delta x_2) \\ \quad = (I \oplus L)(\Delta x_2) \end{cases}$$

From Case 1, we know that the functions $x \oplus L(x)$ and $x \oplus L(x) \oplus L^2(x)$ are invertible. Therefore, Δy_2 and Δy_3 are non-zero. If the other sub-cases with three active words in the input are investigated, it is easy to see that no new condition is added to the present conditions on L .

Finally, we conclude that the diffusion layer D presented in Fig. 1 is perfect if the linear functions

$$\begin{cases} L_1(x) = L(x) \\ L_2(x) = x \oplus L(x) \\ L_3(x) = x \oplus L(x) \oplus L^2(x) \\ L_4(x) = x \oplus L(x) \oplus L^3(x) \\ L_5(x) = x \oplus L^2(x) \oplus L^3(x) \end{cases}$$

are invertible. We know that $L_3(L_2(x)) = x \oplus L^3(x)$ and $L_5(L_4(L_2(x))) = x \oplus L^7(x)$. Thus, by Lemma 6, we can summarize the necessary conditions on the linear function L as the invertibility of $L(x)$, $(I \oplus L)(x)$, $(I \oplus L^3)(x)$ and $(I \oplus L^7)(x)$.

□

Next, we need a simple method to check whether a linear function L satisfies the conditions of Theorem 7 or not. For this purpose, we use the binary matrix representation of L . Assume that x_i is an n -bit word. Hence, we can represent a linear function L with an $n \times n$ matrix \mathbf{A} with elements in $\text{GF}(2)$. By using Lemma 5, if L is invertible, \mathbf{A} is not singular over $\text{GF}(2)$ ($|\mathbf{A}| \neq 0$). To investigate whether a linear function L satisfies the conditions of Theorem 7, we construct the corresponding matrix $\mathbf{A}_{n \times n}$ from L and check the non-singularity of the matrices \mathbf{A} , $\mathbf{I} \oplus \mathbf{A}$, $\mathbf{I} \oplus \mathbf{A}^3$ and $\mathbf{I} \oplus \mathbf{A}^7$ in $\text{GF}(2)$. We introduce some lightweight linear functions with n -bit inputs/outputs in Table 1 that satisfy the above conditions. Note that there exist many linear functions which satisfy the conditions of Theorem 7.

Table 1. Some instances of the linear function L satisfying Theorem 7

n	Some linear functions L
4	$L(x) = (x \oplus x \ll 3) \lll 1$
8	$L(x) = (x \oplus (x \& 0x2) \ll 1) \lll 1$
16	$L(x) = (x \oplus x \ll 15) \lll 1$
32	$L(x) = (x \oplus x \ll 31) \lll 15$ or $L(x) = (x \lll 24) \oplus (x \& 0xFF)$
64	$L(x) = (x \oplus x \ll 63) \lll 1$ or $L(x) = (x \lll 8) \oplus (x \& 0xFFFF)$

Unlike the shift and XOR operations, rotation cannot be implemented as a single instruction on many processors. So, to have more efficient diffusion layers, we introduce new L functions for 32-bit and 64-bit inputs in Table 2 that only use shift and XOR operations.

Table 2. Some examples for the linear function L satisfying Theorem 7 without a circular shift

n	Sample linear functions L
32	$L(x) = (x \ll 3) \oplus (x \gg 1)$
64	$L(x) = (x \ll 15) \oplus (x \gg 1)$

We can use this diffusion layer with $L(x) = (x \ll 3) \oplus (x \gg 1)$ instead of the diffusion layers used in the block ciphers MMB [3] or Hierocrypt [13]. In MMB, the diffusion layer is a 4×4 binary matrix with branch number 4. If we use the proposed diffusion layer in this cipher, it becomes stronger against LC and DC attacks. This change also prevents the attacks presented on this block cipher in

[17]. By computer simulations, we observed that this modification reduces the performance of MMB by about 10%. Also, if we use our proposed diffusion layer with the same $L(x)$, instead of the binary matrix of the block cipher Hierocrypt (called MDS_H [13]), we can achieve a 2 times faster implementation with the same level of security.

Moreover, in the nested SPN structure of Hierocrypt, we replaced the MDS matrix of AES in $GF(2^{32})$ (because inputs of MDS_H are 4 32-bit words) with irreducible polynomial $x^{32} + x^7 + x^5 + x^3 + x^2 + x + 1$ [14] instead of the binary matrix MDS_H . We observed that the replacement of our proposed diffusion layer instead of MDS_H yields 5% better performance than the replacement of the AES matrix in $GF(2^{32})$.

In Eq. (1), if $F_i(x_1, x_2, x_3) = F_0(x_1, x_2, x_3) = L(x_1) \oplus x_2 \oplus L^2(x_3)$, where $L(x) = 2x$ and $x \in GF(2^4)$, PHOTON MDS matrix is obtained [5]. If we change \mathbf{B} to Eq. (2) and define $L(x) = 2x$, we have:

$$\mathbf{B} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 2 & 1 & 3 \end{pmatrix} \Rightarrow \mathbf{B}^4 = \begin{pmatrix} 1 & 2 & 1 & 3 \\ 3 & 7 & 1 & 4 \\ 4 & 11 & 3 & 13 \\ 13 & 30 & 6 & 20 \end{pmatrix}$$

3 Other Desirable Structures for the Proposed Diffusion Layer

In Section 2, the general form of the proposed diffusion layer was introduced in Fig. 1. Then by assuming a special case of α_i 's and β_i 's, an instance of this diffusion layer was given in Eq. (2). In this section, we obtain all sets of α_i 's and β_i 's such that the diffusion layer of Fig. 1 becomes perfect. We know some properties of α_i 's and β_i 's; for instance if all the words of the output are directly represented as the function of input words, a function of each x_i ($0 \leq i \leq s - 1$) must appear in each equation. Another necessary condition is obtained for two active words of the input. Assume there exist only two indices i, j such that $x_i, x_j \neq 0$. If we write each two output words y_p, y_q in a direct form as a function of x_i and x_j , we obtain:

$$\begin{cases} y_p = L_{p_i}(x_i) \oplus L_{p_j}(x_j) \\ y_q = L_{q_i}(x_i) \oplus L_{q_j}(x_j) \end{cases}$$

If

$$\frac{\ell_{p_i}}{\ell_{q_i}} = \frac{\ell_{p_j}}{\ell_{q_j}} \quad \text{or} \quad \left| \begin{matrix} \ell_{p_i} & \ell_{p_j} \\ \ell_{q_i} & \ell_{q_j} \end{matrix} \right| = 0$$

then, $y_p = 0$ is equivalent to $y_q = 0$. Thus, the minimum number of active words in the input and output is less than or equal to s , and the branch number will not reach the maximal value $s + 1$. This procedure must be repeated for 3 and more active words in the input. As an extension, we can use Lemma 3 of [14].

Table 3. Perfect regular recursive diffusion layers for $s < 8$ with only one linear function L

s	Diffusion Layer D	Function that must be invertible
2	$\begin{cases} y_0 = x_0 \oplus L(x_1) \\ y_1 = x_1 \oplus L(y_0) \end{cases}$	$L(x)$ and $x \oplus L(x)$
3	$\begin{cases} y_0 = x_0 \oplus L(x_1 \oplus x_2) \\ y_1 = x_1 \oplus L(x_2 \oplus y_0) \\ y_2 = x_2 \oplus L(y_0 \oplus y_1) \end{cases}$	$L(x)$, $x \oplus L(x)$ and $x \oplus L^3(x)$
3	$\begin{cases} y_0 = x_0 \oplus x_1 \oplus L(x_1 \oplus x_2) \\ y_1 = x_1 \oplus x_2 \oplus L(x_2 \oplus y_0) \\ y_2 = x_2 \oplus y_0 \oplus L(y_0 \oplus y_1) \end{cases}$	$L(x)$, $x \oplus L(x)$, $x \oplus L^3(x)$ and $x \oplus L^7(x)$
3	$\begin{cases} y_0 = x_0 \oplus x_2 \oplus L(x_1 \oplus x_2) \\ y_1 = x_1 \oplus y_0 \oplus L(x_2 \oplus y_0) \\ y_2 = x_2 \oplus y_1 \oplus L(y_0 \oplus y_1) \end{cases}$	$L(x)$, $x \oplus L(x)$, $x \oplus L^3(x)$ and $x \oplus L^7(x)$
3	$\begin{cases} y_0 = x_0 \oplus x_1 \oplus x_2 \oplus L(x_1 \oplus x_2) \\ y_1 = x_1 \oplus x_2 \oplus y_0 \oplus L(x_2 \oplus y_0) \\ y_2 = x_2 \oplus y_0 \oplus y_1 \oplus L(y_0 \oplus y_1) \end{cases}$	$L(x)$, $x \oplus L(x)$, and $x \oplus L^3(x)$
4	$\begin{cases} y_0 = x_0 \oplus x_2 \oplus x_3 \oplus L(x_1 \oplus x_3) \\ y_1 = x_1 \oplus x_3 \oplus y_0 \oplus L(x_2 \oplus y_0) \\ y_2 = x_2 \oplus y_0 \oplus y_1 \oplus L(x_3 \oplus y_1) \\ y_3 = x_3 \oplus y_1 \oplus y_2 \oplus L(y_0 \oplus y_2) \end{cases}$	$L(x)$, $x \oplus L(x)$, $x \oplus L^3(x)$ and $x \oplus L^7(x)$
4	$\begin{cases} y_0 = x_0 \oplus x_1 \oplus x_2 \oplus L(x_1 \oplus x_3) \\ y_1 = x_1 \oplus x_2 \oplus x_3 \oplus L(x_2 \oplus y_0) \\ y_2 = x_2 \oplus x_3 \oplus y_0 \oplus L(x_3 \oplus y_1) \\ y_3 = x_3 \oplus y_0 \oplus y_1 \oplus L(y_0 \oplus y_2) \end{cases}$	$L(x)$, $x \oplus L(x)$, $x \oplus L^3(x)$ and $x \oplus L^7(x)$
4	$\begin{cases} y_0 = x_0 \oplus x_2 \oplus L(x_1 \oplus x_2 \oplus x_3) \\ y_1 = x_1 \oplus x_3 \oplus L(x_2 \oplus x_3 \oplus y_0) \\ y_2 = x_2 \oplus y_0 \oplus L(x_3 \oplus y_0 \oplus y_1) \\ y_3 = x_3 \oplus y_1 \oplus L(y_0 \oplus y_1 \oplus y_2) \end{cases}$	$L(x)$, $x \oplus L(x)$, $x \oplus L^3(x)$, $x \oplus L^7(x)$ and $x \oplus L^{15}(x)$
4	$\begin{cases} y_0 = x_0 \oplus x_1 \oplus x_3 \oplus L(x_1 \oplus x_2 \oplus x_3) \\ y_1 = x_1 \oplus x_2 \oplus y_0 \oplus L(x_2 \oplus x_3 \oplus y_0) \\ y_2 = x_2 \oplus x_3 \oplus y_1 \oplus L(x_3 \oplus y_0 \oplus y_1) \\ y_3 = x_3 \oplus y_0 \oplus y_2 \oplus L(y_0 \oplus y_1 \oplus y_2) \end{cases}$	$L(x)$, $x \oplus L(x)$, $x \oplus L^3(x)$, $x \oplus L^7(x)$ and $x \oplus L^{15}(x)$

Lemma 8. Assume the diffusion layer has m inputs/outputs bits and ℓ is the linear operator of $L(x)$ and I is the linear operator of $I(x)$. Moreover, \mathbf{ML}_D is an $m \times m$ matrix representation of the operator of the diffusion layer. If D is perfect, then all the sub-matrices of \mathbf{ML}_D is non-singular.

If we construct the \mathbf{ML}_D of Eq. (2), we have:

$$\mathbf{ML}_D = \begin{pmatrix} I & \ell & I & I \oplus \ell \\ I \oplus \ell & I \oplus \ell \oplus \ell^2 & I & \ell^2 \\ \ell^2 & I \oplus \ell \oplus \ell^3 & I \oplus \ell & I \oplus \ell^2 \oplus \ell^3 \\ I \oplus \ell^2 \oplus \ell^3 & \ell \oplus \ell^2 \oplus \ell^3 \oplus \ell^4 & \ell \oplus \ell^2 & \ell^2 \oplus \ell^4 \end{pmatrix}$$

If we calculate 69 sub-matrix determinant of \mathbf{ML}_D , we find the result of Theorem 7. However, by following this procedure, it is complicated to obtain all sets of α_i 's and β_i 's analytically. So, by systematizing the method based on Lemma 8, we performed a computer simulation to obtain all sets of α_i 's and β_i 's in the diffusion layer in Fig. 1 that yield a perfect diffusion. We searched for all α_i 's and β_i 's that make the diffusion layer of Fig. 1 a perfect diffusion layer. This procedure was repeated for $s = 2, 3, \dots, 8$. We found one set of (α_i, β_i) for $s = 2$, four sets for $s = 3$, and four sets for $s = 4$. The obtained diffusion layers along with the conditions on the underlying linear function L are reported in Table 3. We observed that for $s = 5, 6, 7$ the diffusion layer introduced in Fig. 1 cannot be perfect.

Note that some linear functions in Table 1 and Table 2 such as $L(x) = (x \ll 15) \oplus (x \gg 1)$ cannot be used in the diffusion layers for which $x \oplus L^{15}(x)$ must be invertible.

As we can see in Fig. 1 and its instances presented in Table 3, there exists some kind of regularity in the equations defining y_i 's, in the sense that the form of y_{i+1} is determined by the form of y_i and vice versa (F_i 's are all the same in Eq. (1)). However, we can present some non-regular recursive diffusion layers with the following more general form (F_i 's are different):

```

1: Input :  $s$   $n$ -bit words  $x_0, \dots, x_{s-1}$ 
2: Output :  $s$   $n$ -bit words  $y_0, \dots, x_{s-1}$ 
3: for  $i = 0$  to  $s - 1$  do
4:    $y_i = x_i$ 
5: end for
6: for  $i = 0$  to  $s - 1$  do
7:    $y_i = y_i \oplus \left( \bigoplus_{j=0, j \neq i}^{s-1} A_{i,j} y_j \right) \oplus L \left( \bigoplus_{j=0, j \neq i}^{s-1} B_{i,j} y_j \right)$ 
8: end for
    
```

Fig. 3. Non-regular recursive diffusion layers

where $A_{i,j}, B_{i,j} \in \{0, 1\}$. If $A_{i,j} = \alpha_{(j-i) \bmod s}$ and $B_{i,j} = \beta_{(j-i) \bmod s}$, then Fig. 3 is equivalent to Fig. 1. The main property of this new structure is that it still has one linear function L and a simple structure for the inverse. For example, if $s = 4$, then, the diffusion layer D is:

$$\begin{cases} y_0 = x_0 \oplus A_{0,1} \cdot x_1 \oplus A_{0,2} \cdot x_2 \oplus A_{0,3} \cdot x_3 \oplus L(B_{0,1} \cdot x_1 \oplus B_{0,2} \cdot x_2 \oplus B_{0,3} \cdot x_3) \\ y_1 = x_1 \oplus A_{1,0} \cdot y_0 \oplus A_{1,2} \cdot x_2 \oplus A_{1,3} \cdot x_3 \oplus L(B_{1,0} \cdot y_0 \oplus B_{1,2} \cdot x_2 \oplus B_{1,3} \cdot x_3) \\ y_2 = x_2 \oplus A_{2,0} \cdot y_0 \oplus A_{2,1} \cdot y_1 \oplus A_{2,3} \cdot x_3 \oplus L(B_{2,0} \cdot y_0 \oplus B_{2,1} \cdot y_1 \oplus B_{2,3} \cdot x_3) \\ y_3 = x_3 \oplus A_{3,0} \cdot y_0 \oplus A_{3,1} \cdot y_1 \oplus A_{3,2} \cdot y_2 \oplus L(B_{3,0} \cdot y_0 \oplus B_{3,1} \cdot y_1 \oplus B_{3,2} \cdot y_2) \end{cases}$$

We searched the whole space for $s = 3$ and $s = 4$ (the order of search spaces are 2^{12} and 2^{24} respectively). For $s = 3$, we found 196 structures with branch number 4 and for $s = 4$, 1634 structures with branch number 5. The linear functions that must be invertible for each case are different. Among the 196 structures for $s = 3$, the structure with the minimum number of operations (only 7 XORs and one L evaluation) is the following:

$$D : \begin{cases} y_0 = x_0 \oplus x_1 \oplus x_2 \\ y_1 = x_1 \oplus x_2 \oplus L(y_0 \oplus x_2) \\ y_2 = x_2 \oplus y_0 \oplus y_1 \end{cases}$$

where $L(x)$ and $x \oplus L(x)$ must be invertible.

This relation is useful to enlarge the first linear function of the new hash function JH for 3 inputs [18]. For $s = 4$, we did not find any D with the number of L evaluations less than four. However, the one with the minimum number of XORs is given as below:

$$D : \begin{cases} y_0 = x_0 \oplus x_1 \oplus x_2 \oplus L(x_3) \\ y_1 = x_1 \oplus x_3 \oplus y_0 \oplus L(x_2 \oplus y_0) \\ y_2 = x_2 \oplus x_3 \oplus y_0 \oplus L(x_3 \oplus y_1) \\ y_3 = x_3 \oplus y_1 \oplus y_2 \oplus L(y_0) \end{cases}$$

Searching the whole space for $s = 5, 6, \dots$ is too time consuming (note that for $s = 5$, the order of search has complexity 2^{40}) and we could not search all the space for $s \geq 5$.

4 Increasing the Number of Linear Functions

In Section 3, we observed that for $s > 4$ we cannot design a regular recursive diffusion layer in the form of Fig. 1 with only one linear function L . In this section, we increase the number of linear functions to overcome the regular structure of the diffusion layer of Eq. (2). A new structure is represented in Fig. 4, where $\alpha_k, \beta_k, \gamma_k \in \{0, 1\}$, $k \in \{0, 1, \dots, s - 1\}$, $\alpha_0 = 1, \beta_0 = 0$ and $\gamma_0 = 0$.

If L_1 and L_2 are two distinct linear functions, Fig. 4 is too complicated to easily obtain conditions on L_1 and L_2 that make it a perfect diffusion layer. To obtain simplified conditions for a maximal branch number, let L_1 and L_2 have a simple relation like $L_2(x) = L_1^2(x)$ or $L_2(x) = L_1^{-1}(x)$. For the linear functions in

```

1: Input :  $s$   $n$ -bit words  $x_0, \dots, x_{s-1}$ 
2: Output :  $s$   $n$ -bit words  $y_0, \dots, x_{s-1}$ 
3: for  $i = 0$  to  $s - 1$  do
4:    $y_i = x_i$ 
5: end for
6: for  $i = 0$  to  $s - 1$  do
7:    $y_i = \bigoplus_{j=0}^{s-1} \alpha_{[(j-i) \bmod s]} y_j \oplus L_1 \left( \bigoplus_{j=0}^{s-1} \beta_{[(j-i) \bmod s]} y_j \right) \oplus L_2 \left( \bigoplus_{j=0}^{s-1} \gamma_{[(j-i) \bmod s]} y_j \right)$ 
8: end for
    
```

Fig. 4. Regular recursive diffusion layers with two linear functions L

Table 2 and Table 3, $L^2(x)$ is more complex in comparison with $L(x)$. However, there exist some linear functions $L(x)$ such that $L^{-1}(x)$ is simpler than $L^2(x)$. As an example, for $L(x_{(n)}) = (x_{(n)} \oplus x_{(n)} \gg b) \lll a$, where $b > \frac{n}{2}$ we have $(x_{(n)} \gg 2b = 0)$:

$$L^{-1}(x_{(n)}) = ((x_{(n)} \ggg a) \oplus (x_{(n)} \ggg a) \gg b)$$

In Table 4, we introduce some recursive diffusion layers with $(L_1 = L$ and $L_2 = L^{-1})$ or $(L_1 = L$ and $L_2 = L^2)$ that have maximal branch numbers. These diffusion layers are obtained similar to that of Table 3. In this table, for each case only y_0 is presented. Other y_i 's can be easily obtained from Fig. 4, since F_i 's are all the same.

Table 4. Some perfect regular diffusion layers for $s = 5, 6, 7, 8$ with two linear functions

s	y_0 in a perfect diffusion Layer
5	$y_0 = x_0 \oplus x_2 \oplus x_3 \oplus L(x_4) \oplus L^2(x_1)$
5	$y_0 = L^{-1}(x_1 \oplus x_2) \oplus x_0 \oplus x_1 \oplus L(x_1 \oplus x_3 \oplus x_4)$
6	$y_0 = x_0 \oplus x_2 \oplus x_4 \oplus x_5 \oplus L(x_3 \oplus x_5) \oplus L^2(x_1 \oplus x_2 \oplus x_3)$
6	$y_0 = L^{-1}(x_1 \oplus x_3) \oplus x_0 \oplus x_2 \oplus x_3 \oplus x_4 \oplus x_5 \oplus L(x_1 \oplus x_3 \oplus x_4 \oplus x_5)$
7	$y_0 = x_0 \oplus x_2 \oplus L(x_3 \oplus x_4) \oplus L^2(x_1 \oplus x_2 \oplus x_4 \oplus x_5 \oplus x_6)$
7	$y_0 = L^{-1}(x_1 \oplus x_2 \oplus x_3 \oplus x_4 \oplus x_6) \oplus x_0 \oplus x_2 \oplus x_3 \oplus x_4 \oplus L(x_1 \oplus x_2 \oplus x_3 \oplus x_5)$
8	$y_0 = x_0 \oplus x_1 \oplus x_3 \oplus x_4 \oplus L(x_2 \oplus x_3 \oplus x_5) \oplus L^2(x_1 \oplus x_5 \oplus x_6 \oplus x_7)$
8	$y_0 = L^{-1}(x_1 \oplus x_2 \oplus x_3 \oplus x_5 \oplus x_7) \oplus x_0 \oplus x_2 \oplus x_3 \oplus x_4 \oplus x_5 \oplus x_7 \oplus L(x_1 \oplus x_3 \oplus x_5 \oplus x_6 \oplus x_7)$

If the 14 linear functions:

$$\begin{array}{lll}
 L(x) & I \oplus L(x) & I \oplus L^3(x) \\
 I \oplus L^7(x) & I \oplus L^{15}(x) & I \oplus L^{31}(x) \\
 I \oplus L^{63}(x) & I \oplus L^{127}(x) & I \oplus L^{255}(x) \\
 I \oplus L^{511}(x) & I \oplus L^{1023}(x) & I \oplus L^{2047} \\
 I \oplus L^{4095}(x) & I \oplus L^{8191}(x) &
 \end{array}$$

are invertible (all irreducible polynomials up to degree 13), then all the diffusion layers introduced in Table 4 are perfect. One example for a 32-bit linear function satisfying these conditions is:

$$L(x_{(32)}) = (x_{(32)} \oplus (x_{(32)} \ggg 31)) \lll 29$$

5 Conclusion

In this paper, we proposed a family of diffusion layers which are constructed using some rounds of Feistel-like structures whose round functions are linear. These diffusion layers are called recursive diffusion layers. First, for a fixed structure, we determined the required conditions for its underlying linear function to make it a perfect diffusion layer. Then, for the number of words in input (output) less than 8, we extended our approach and found all the instances of the perfect recursive diffusion layers with the general form of Fig. 1. Also, we proposed some other diffusion layers with non-regular forms which can be used for the design of lightweight block ciphers. Finally, diffusion layers with 2 linear functions were proposed. By using two linear functions, we designed perfect recursive diffusion layers for $s = 5, 6, 7, 8$ which cannot be designed based on Fig. 1, i.e., using only one linear function.

The proposed diffusion layers have simple inverses, thus they can be deployed in SPN structures. These proposed diffusion layers can be used to improve the security or performance of some of the current block ciphers and hash functions or in the design of the future block ciphers and hash functions (especially the block ciphers with provable security against DC and LC).

References

1. Bernstein, D.J.: The Salsa20 Stream Cipher. Symmetric Key Encryption Workshop, SKEW (2005), <http://www.ecrypt.eu.org/stream/salsa20p2.html>
2. 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)
3. Daemen, J.: Cipher and Hash Function Design Strategies Based on Linear and Differential Cryptanalysis. PhD thesis, Elektrotechniek Katholieke Universiteit Leuven, Belgium (1995)
4. Daemen, J., Rijmen, V.: The Design of Rijndael: AES - The Advanced Encryption Standard. Springer (2002)
5. Guo, J., Peyrin, T., Poschmann, A.: The PHOTON Family of Lightweight Hash Functions. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 222–239. Springer, Heidelberg (2011)
6. Guo, J., Peyrin, T., Poschmann, A., Robshaw, M.: The LED Block Cipher. In: Preneel, B., Takagi, T. (eds.) CHES 2011. LNCS, vol. 6917, pp. 326–341. Springer, Heidelberg (2011)
7. Junod, P., Vaudenay, S.: Perfect Diffusion Primitives for Block Ciphers. In: Handschuh, H., Hasan, M.A. (eds.) SAC 2004. LNCS, vol. 3357, pp. 84–99. Springer, Heidelberg (2004)

8. Kang, J., Hong, S., Lee, S., Yi, O., Park, C., Lim, J.: Practical and Provable Security Against Differential and Linear Cryptanalysis for Substitution-Permutation Networks. *ETRI Journal* 23(4), 158–167 (2001)
9. Lee, C., Kim, J., Sung, J., Hong, S., Lee, S.: Provable Security for an RC6-like Structure and a MISTY-FO-like Structure Against Differential Cryptanalysis. In: Gavrilova, M.L., Gervasi, O., Kumar, V., Tan, C.J.K., Taniar, D., Laganá, A., Mun, Y., Choo, H. (eds.) *ICCSA 2006*. LNCS, vol. 3982, pp. 446–455. Springer, Heidelberg (2006)
10. Matsui, M.: Linear Cryptanalysis Method for DES Cipher. In: Helleseth, T. (ed.) *EUROCRYPT 1993*. LNCS, vol. 765, pp. 386–397. Springer, Heidelberg (1994)
11. Matsui, M.: New Structure of Block Ciphers with Provable Security Against Differential and Linear Cryptanalysis. In: Gollmann, D. (ed.) *FSE 1996*. LNCS, vol. 1039, pp. 205–218. Springer, Heidelberg (1996)
12. Nyberg, K., Knudsen, L.: Provable Security Against a Differential Attack. *Journal of Cryptology* 8(1), 27–37 (1995)
13. Ohkuma, K., Muratani, H., Sano, F., Kawamura, S.: The Block Cipher Hierocrypt. In: Stinson, D.R., Tavares, S. (eds.) *SAC 2000*. LNCS, vol. 2012, pp. 72–88. Springer, Heidelberg (2001)
14. Sajadieh, M., Dakhilalian, M., Mala, H.: Perfect Involutory Diffusion Layers Based on Invertibility of Some Linear Functions. *IET Information Security Journal* 5(1), 228–236 (2011)
15. Schnorr, C.-P., Vaudenay, S.: Black Box Cryptanalysis of Hash Networks Based on Multipermutations. In: De Santis, A. (ed.) *EUROCRYPT 1994*. LNCS, vol. 950, pp. 47–57. Springer, Heidelberg (1995)
16. Vaudenay, S.: On the Need for Multipermutations: Cryptanalysis of MD4 and SAFER. In: Preneel, B. (ed.) *FSE 1994*. LNCS, vol. 1008, pp. 286–297. Springer, Heidelberg (1995)
17. Wang, M., Nakahara Jr., J., Sun, Y.: Cryptanalysis of the Full MMB Block Cipher. In: Jacobson Jr., M.J., Rijmen, V., Safavi-Naini, R. (eds.) *SAC 2009*. LNCS, vol. 5867, pp. 231–248. Springer, Heidelberg (2009)
18. Wu, H.: The Hash Function JH. Submission to NIST (2008)

Unaligned Rebound Attack: Application to Keccak

Alexandre Duc^{1,*}, Jian Guo^{2,**}, Thomas Peyrin^{3,***}, and Lei Wei^{3,†}

¹ Ecole Polytechnique Fédérale de Lausanne, Switzerland

`alexandre.duc@epfl.ch`

² Institute for Infocomm Research, Singapore

`{ntu.guo,thomas.peyrin}@gmail.com`

³ Nanyang Technological University, Singapore

`wl@pmail.ntu.edu.sg`

Abstract. We analyze the internal permutations of KECCAK, one of the NIST SHA-3 competition finalists, in regard to differential properties. By carefully studying the elements composing those permutations, we are able to derive most of the best known differential paths for up to 5 rounds. We use these differential paths in a rebound attack setting and adapt this powerful freedom degrees utilization in order to derive distinguishers for up to 8 rounds of the internal permutations of the submitted version of KECCAK. The complexity of the 8 round distinguisher is $2^{491.47}$. Our results have been implemented and verified experimentally on a small version of KECCAK.

Keywords: KECCAK, SHA-3, hash function, differential cryptanalysis, rebound attack.

1 Introduction

Cryptographic hash functions are used in many applications such as digital signatures, authentication schemes or message integrity and they are among the most important primitives in cryptography. Even if hash functions are traditionally used to simulate the behavior of a random oracle [3], classical security requirements are collision resistance and (second)-preimage resistance.

* Part of the work was done while the author was visiting Nanyang Technological University, supported by the NTU NAP Startup Grant M58110000. This work has also partially been supported by the European Commission through the ICT programme under contract ICT-2007-216676 ECRYPT II.

** Part of the work was done while the author was visiting Tsinghua University, supported by the National Natural Science Foundation of China under grant No. 61133013 and No. 60931160442.

*** The author is supported by the Lee Kuan Yew Postdoctoral Fellowship 2011 and the Singapore National Research Foundation Fellowship 2012.

† The author is supported by the Singapore National Research Foundation under Research Grant NRF-CRP2-2007-03, the Singapore Ministry of Education under Research Grant T206B2204 and by the NTU NAP Startup Grant M58110000.

Like any construction that builds a hash function from a subcomponent, the cryptographic quality of this internal permutation is very important for a sponge construction. Therefore, this permutation P should not present any structural flaw, or should not be distinguishable from a randomly chosen permutation. Zero-sum distinguishers [2] can reach an important number of rounds, but generally with a very high complexity. For example, the latest results [9] provide zero-sum partitions distinguishers for the full 24-round 1600-bit internal permutation with a complexity of 2^{1575} . When looking at smaller number of rounds the complexity would decrease, but it is unclear how one can describe the partition of a 1600-bit internal state without using the KECCAK round inside the definition of the partition. Moreover, such zero-sum properties seem very hard to exploit when the attacker aims at the whole hash function. On the other side, more classical preimage attack on 3 rounds using SAT-solvers have been demonstrated [18]. Finally, Bernstein published [4] a $2^{511.5}$ computations (second)-preimage attack on 8 rounds that allows a workload reduction of only half a bit over the generic complexity with an important memory cost of 2^{508} .

Our Contributions. In this article, we analyze the differential cryptanalysis resistance of the KECCAK internal permutation. More precisely, we first introduce a new and generic method that looks for good differential paths for all the KECCAK internal permutations, and we obtain the currently best known differential paths. We then describe a simple method to utilize the available freedom degrees which allows us to derive distinguishers for reduced variants of the KECCAK internal permutations with low complexity. Finally, we apply the idea of rebound attack [17] to KECCAK. This application is far from being trivial and requires a careful analysis of many technical details in order to model the behavior of the attack. This technique is in particular much more complicated to apply to KECCAK than to AES or to other 4-bit Sbox hash functions [21,13]. One reason for that is that KECCAK has *weak alignment* [6]. This is why we call our attack “unaligned rebound attack”. The model introduced has been verified experimentally on a small version of KECCAK and we eventually obtained differential distinguishers for up to 8 rounds of the submitted version of KECCAK to the SHA-3 competition. In order to demonstrate why differential analysis is in general more relevant than zero-sum ones in regards to the full hash function, we applied our techniques to the recent KECCAK challenges [23] and managed to obtain the currently best known practical collision attack for up to two rounds.

2 The KECCAK Hash Function Family

KECCAK [7,8] is a family of variable length output hash functions based on the sponge construction [5]. In KECCAK family, the underlying function is a permutation chosen from a set of seven KECCAK- f permutations, denoted as KECCAK- $f[b]$ where $b \in \{1600, 800, 400, 200, 100, 50, 25\}$ is the permutation width as well as the internal state size of the hash function. The KECCAK family is parametrized by an r -bit bitrate and c -bit capacity with $b = r + c$.

2.1 The KECCAK- f Permutations

The internal state of the KECCAK family can be viewed as a bit array of 5×5 lanes, each of length $w = 2^\ell$ where $\ell \in \{0, 1, 2, 3, 4, 5, 6\}$ and $b = 25w$. The state can also be described as a three dimensional array of bits defined by $a[5][5][w]$. A bit position (x, y, z) in the state is given by $a[x][y][z]$ where x and y coordinates are taken over modulo 5 and the z coordinate is taken over modulo w . A *lane* of the internal state at *column* x and *row* y is represented by $a[x][y][\cdot]$, while a *slice* of the internal state at *width* z is represented by $a[\cdot][\cdot][z]$.

KECCAK- $f[b]$ is an iterated permutation consisting of a sequence of n_r rounds indexed from 0 to $n_r - 1$ and the number of rounds are given by $n_r = 12 + 2\ell$. A round \mathbf{R} consists of a transformation of five step mappings and is defined by: $\mathbf{R} = \iota \circ \chi \circ \pi \circ \rho \circ \theta$. These step mappings are discussed below.

θ Mapping. This linear mapping intends to provide diffusion for the state and is defined for every x, y and z by: $\theta : a[x][y][z] \leftarrow a[x][y][z] + \bigoplus_{y'=0}^4 a[x-1][y'][z] + \bigoplus_{y'=0}^4 a[x+1][y'][z-1]$.

ρ Mapping. This linear mapping intends to provide diffusion between the slices of the state through intra-lane bit translations. For every x, y and z : $\rho : a[x][y][z] \leftarrow a[x][y][z + T(x, y)]$, where $T(x, y)$ is a translation constant.

π Mapping. This linear mapping intends to provide diffusion in the state through transposition of the lanes.

χ Mapping. This is the only non-linear mapping of KECCAK- f and is defined for every x, y and z by: $\chi : a[x][y][z] \leftarrow a[x][y][z] + ((\neg a[x+1][y][z]) \wedge a[x+2][y][z])$. This mapping is similar to an Sbox applied independently to each 5-bit row of the state and can be computed in parallel to all rows. We represent by $s = 5w$ the number of Sboxes/rows in KECCAK internal state. Here \neg denotes bit-wise complement, and \wedge the bit-wise AND.

ι Mapping. This mapping adds constants derived from an LFSR to the lane $a[0][0][\cdot]$. These constants are different in every round i . This mapping aims at destroying the symmetry introduced by the identical nature of the remaining mappings in a round.

3 Finding Differential Paths for KECCAK- f

We first study how to find “good” differential paths for all KECCAK variants. In this section, we describe our differential finding algorithms.

3.1 Special Properties of θ and χ

It is noted by the KECCAK designers [7, Section 2.4.3] that when every column of the state sums to 0, θ acts as identity. The set of such states is called *column*

parity kernel (CPK). Since θ is linear, this property applies not only to the state values, but also to differentials. While θ expands a single bit difference into at most 11 bits (2 columns and the bit itself), it acts as identity on differences in the CPK. This property will be intensively used in finding low Hamming weight bitwise differentials. Another interesting property is that θ^{-1} diffuses much faster than θ , i.e., a single bit difference can be propagated to about half state bits through θ^{-1} [7, Section 2.3.2]. However, the output of θ^{-1} is extremely regular when the Hamming weight of the input is low.

The χ layer updates is a row-wise operation and can also be viewed as a 5-bit Sbox. Similar to the analysis of other Sboxes, we build its differential distribution table (DDT). We remark that when a single difference is present, χ acts as identity with best probability 2^{-2} , while input differences with more active bits tend to lead to more possible output differences, but with lower probability. It is also interesting to note that given an input difference to χ , all possible output differences occur with same probability (however this is not the case for χ^{-1}).

3.2 First Tools

Our goal is to derive “good” bitwise differential paths by maintaining the bit difference Hamming weight as low as possible. The ι permutation adds predefined constants to the first lane, and hence has no essential influences when such differentials are considered. For the rest of the paper, we will ignore this layer. We note that θ , ρ and π are all linear mappings, while χ acts as a non-linear Sbox. Furthermore, ρ and π do not change the number of active bits in a differential path, but only bit positions. Hence, θ and χ are critical when looking for a “good” differential path. Since χ is followed by θ in the next round (ignoring ι), we consider these two mappings together by treating a slice of the state as a unit, and try to find the potential best mapping of the slice through χ with the following two rules.

1. Given an input difference of the slice, i.e., 5 row differences, find all possible output differences by looking into the DDT table. Then among all combinations of solutions of the 5 rows, choose the output combinations with minimum number of columns with odd parity.
2. In case of a draw, we select the state with the minimum number of active bits.

Rule 1 aims at reducing the amount of active bits after applying θ by choosing each slice of the output of the χ closest to the CPK (i.e., with even parity for most columns), and rule 2 further reduces the amount of active bits within the columns. Although this strategy may not lead to the minimum number of active bits after θ in the entire state, it finds the best slice-wise mappings with help of a table of size 2^{25} .

3.3 Algorithm for Differential Path Search

Denote $\lambda = \pi \circ \rho \circ \theta$ (all linear mappings), and the state at round i before (resp. after) applying the linear layer λ as a_i (resp. b_i). We start our search from a_1 ,

i.e., the input state to the second round, and compute backwards for one round, and few rounds forwards.

The forward part is longer than the backward part because the diffusion of θ^{-1} is much better than for θ , thus, it will be easier for us to control the bit differences Hamming weight for several forward rounds (instead of backward rounds).

We choose a_1 from the CPK. Since it is impossible to enumerate all combinations, we further restrict to a subset of the CPK with at most 8 active bits and each column having exactly 0 or 2 active bits. Note also that any bitwise differential path is invariant through position rotation along the z axis, so we have to run through a set of size about 2^{36} . An example of 4 round path is given in the full version of the paper [10]. We provide also in Table 1 some of the best differential path probabilities found for all KECCAK internal permutation sizes.

Table 1. Best differential path results for each version of KECCAK internal permutations, for 1 up to 5 rounds. The detailed differential paths for KECCAK-f[1600] are shown in the full version of the paper. Paths in bold are new results we found with the method presented in this paper.

b	best differential path probability (differential complexity of the rounds)		
	3 rds	4 rds	5 rds
400	2^{-24} (8 - 8 - 8)	2^{-84} (16 - 14 - 12 - 42)	2^{-216} (16 - 32 - 40 - 32 - 96)
800	2^{-32} (4 - 4 - 24)	2^{-109} (12 - 12 - 12 - 73)	2^{-432} (32 - 64 - 80 - 64 - 192)
1600	2^{-32} (4 - 4 - 24)	2^{-142} (12 - 12 - 12 - 106)	2^{-709} (16 - 16 - 16 - 114 - 547) A better path (2^{-510}) was found independently [20]

4 Simple Distinguishers for Reduced KECCAK- f

Once the differential paths obtained, we can concentrate our efforts on how to use at best the freedom degrees available in order to reduce the complexity required to find a valid pair for the differential trails or to increase the amount of rounds attacked. We present in this section a very simple method that allows to obtain low complexity distinguishers on a few rounds of the KECCAK internal permutations.

4.1 A Very Simple Freedom Degrees Fixing Method

We first describe an extremely simple way of using the available freedom degrees, which are exactly the b -bit value of the internal state (since we already fixed the differential path). For all the best differential paths found from Table 1, we can extend them by one round to the left or to the right by simply picking some valid Sboxes differential transitions. We can use our available freedom degrees specifically for this round so that its cost is null. One simply handles each of the

active Sboxes differential transitions for this round one by one, independently, by fixing a valid value for the active Sboxes. In terms of freedom degrees consumption for this extra round, in the worst case we have all s Sboxes active and a differential transition probability of 2^{-4} for each of them. Thus, we are ensured to have at least $2^{5s-4s} = 2^s$ freedom degrees remaining after handling this extra round.

Note that some more involved freedom degree methods (such as message modification [24]) might even allow to also control some of the conditions of the original differential path, thus further reducing the complexity.

4.2 Getting More Rounds

At the present time, we are able to find valid pairs of internal state values for some differential paths on a few rounds with a rather low complexity. Said in other words, we are able to compute internal state value pairs with a predetermined input/output difference. A direct application from this is to derive distinguishers. For a randomly chosen permutation of b bits, finding a pair of inputs with a predetermined difference that maps to a predetermined output difference costs 2^b computations. Indeed, since the input and the output differences are fixed, the attacker can not apply any birthday-paradox technique. Those distinguishers are called “limited-birthday distinguishers” and can be generalized in the following way (we refer to [11] for more details): for a randomly chosen b -bit permutation, the problem of mapping an input difference from a subset of size I to an output difference from a subset of size J requires $\max\{\sqrt{2^b/J}, 2^b/(I \cdot J)\}$ calls to permutation (while assuming without loss of generality since we are dealing with a permutation that $I \leq J$).

Using the freedom degrees technique from the previous section and reading Table 1, we are for example able to obtain a distinguisher for 5 rounds of the KECCAK- f [1600] internal permutation with complexity 2^{142} (while the generic case is 2^{1600}).

5 The Rebound Attack on KECCAK

The rebound attack is a freedom degrees utilization technique that was first proposed by Mendel et al. in [17] as an analysis of round-reduced Grøst1 and Whirlpool. It was then improved in [16,15,11,22] to analyze AES and AES-like permutations and also ARX ciphers [14].

With the help of rebound techniques, we show in this section how to extend the number of attacked rounds significantly, but for a higher complexity. We will see that the application of the rebound attack for KECCAK seems quite difficult. Indeed, the situation for KECCAK is not as pleasant as the AES-like permutations case where the utilization of truncated differential paths (i.e. path for which one only checks if one cell is active or inactive, without caring about the actual difference value) makes the application of rebound attacks very easy to handle.

5.1 The Original Rebound Attack

Let P denote a permutation, which can be divided into 3 sub-permutations, i.e., $P = E_F \circ E_I \circ E_B$. The rebound attack works in two phases.

- **Inbound phase or controlled rounds:** this phase usually starts with several chosen input/output differences of E_I that are propagated through linear layers forward and backward. Then, one can carry out meet-in-the-middle (MITM) match for differences through a single Sbox layer in E_I and generate all possible value pairs validating the matches.
- **Output phase or uncontrolled rounds:** With all solutions provided in the inbound phase, check if any pair validates as well the differential paths for both the backward part p_B and the forward part p_F .

The SuperSbox technique [15,11] extends the E_I from one Sbox layer to two Sbox layers for an AES-like permutation, by considering two consecutive AES-like rounds as one with column-wise SuperSboxes. This technique is possible due to the fact that one can swap few linear operations with the Sbox in AES, so that the two layers of Sboxes in two rounds become close enough to form one SuperSbox layer. However, in the case of KECCAK, it seems very hard to form any partition into independent SuperSboxes. For the same reason, using truncated differential paths seems very difficult for KECCAK, as it has recently been shown in [6].

5.2 Applying the Rebound Attack for KECCAK Internal Permutations

Assume that we know a set of n_B differential trails (called *backward trails*) on nr_B KECCAK rounds and whose DP is higher or equal to p_B . For the moment, we want all these backward paths to share the same input difference mask Δ_B^{in} and we denote by $\Delta_B^{\text{out}}[i]$ the output difference mask of the i -th trail of the set. Similarly, we consider that we also know a set of n_F differential trails (called *forward trails*) on nr_F KECCAK rounds and whose DP is higher or equal to p_F . We want all those forward paths to share the same output difference mask Δ_F^{out} and we denote by $\Delta_F^{\text{in}}[i]$ the input difference mask of the i -th trail of the set.

Our goal here is to build a differential path on $nr_B + nr_F + 1$ KECCAK rounds (thus one Sbox layer of inbound), by connecting a forward and a backward trail with the rebound technique, and eventually to find the corresponding solution values for the controlled round. We represent by p_{match} the probability that a match is possible from a given element of the backward set and a given element of the forward set, and we denote by N_{match} the number of solution values that can be generated once a match has been obtained.

For this connection to be possible, we need the inbound phase to be a valid differential path, that is we need to find a valid differential path from a $\Delta_B^{\text{out}*}$ to a Δ_F^{in} . By using random $\Delta_B^{\text{out}*}$ and Δ_F^{in} this will happen in general with very small probability, because we need the very same set of Sboxes to be active/inactive in both forward and backward difference masks to have a chance to get a match.




Fig. 1. Rebound attack on Keccak

Even if the set of active Sboxes matches, we still require the differential transitions through all the active Sboxes to be possible.

We can generalize a bit this approach by allowing a fixed set of differences Δ_B^{in} (resp. Δ_F^{out}) instead of just one. We call I_B^{in} (resp. I_B^{out}) the size of the set of possible Δ_B^{in} (resp. Δ_B^{out}) values for the backward paths. Similarly, we call I_F^{in} (resp. I_F^{out}) the size of the set of possible Δ_F^{in} (resp. Δ_F^{out}) values for the forward paths. In fact, the number of possible differences in the backward or forward parts will form a butterfly shape. We call I_B^{mid} (resp. I_F^{mid}) the minimum number of differences in the backward (resp. forward) part.

The total complexity C to find one valid internal state pair for the $(nr_B + nr_F + 1)$ -round path is

$$C = n_F + n_B + \frac{1}{p_{\text{match}}} \cdot \left\lceil \frac{1}{p_F \cdot p_B \cdot N_{\text{match}}} \right\rceil + \frac{1}{p_B \cdot p_F}, \tag{1}$$

with

$$I_B^{\text{out}} \cdot I_F^{\text{in}} = \frac{1}{p_{\text{match}}} \cdot \left\lceil \frac{1}{p_F \cdot p_B \cdot N_{\text{match}}} \right\rceil. \tag{2}$$

The first two terms are the costs to generate the backward and forward paths. The term $\left\lceil \frac{1}{p_F \cdot p_B \cdot N_{\text{match}}} \right\rceil$ denotes the number of time we will need to perform the inbound and each inbound costs $1/p_{\text{match}}$. The last term is the cost for actually performing the outbound phase. The condition (2) is needed since we need enough differences to perform the inbound phase.

5.3 An Ordered Buckets and Balls Problem

We model the active/inactive Sboxes match as a **limited capacity ordered buckets and balls problem**: the $s = 5w$ ordered buckets ($s = 320$ for Keccak- $f[1600]$) limited to capacity 5 will represent the s 5-bit Sboxes and the x_B (resp. x_F) balls will stand for the Hamming weight of the difference in $\Delta_B^{\text{out*}}$ (resp. in Δ_F^{in}). Given a set B of s buckets in which we randomly throw x_B balls and a set F of s buckets in which we randomly throw x_F balls, we call the result a **pattern-match** when the set of empty buckets in B and F after the experiment are the same. Before computing the probability of having a pattern-match, we need the following lemma.

Lemma 1. *The number of possible combinations $b_{\text{bucket}}(n, s)$ to place n balls into s buckets of capacity 5 such that no bucket is empty is*

$$b_{\text{bucket}}(n, s) := \sum_{i=\lceil n/5 \rceil}^s (-1)^{s-i} \binom{s}{i} \binom{5i}{n} \quad \text{if } s \leq n \leq 5s \quad (3)$$

and 0 otherwise.

The proof of this lemma is given in the full version of the paper.

Using (3), we can derive the probability $p_{\text{bucket}}(n, s)$ that every bucket contains at least one ball when n balls are thrown into s buckets with capacity 5 and the expected number of active buckets when n balls are thrown into s buckets. We can now relate this lemma to the more general pattern-match problem. This model tells us that when the number of balls (i.e., active bits) is not too small on both sides, most of the matches happen when (almost) all the Sboxes are active. We analyze this behavior in more details in the full version of the paper.

A More General Problem. We can also look into a more general problem, i.e., we characterize more precisely how the bits are distributed into the Sboxes.

Lemma 2. *The probability p_{dist} of distributing randomly n active bits into s 5-bit Sboxes such that exactly A_i Sboxes contain i bits, for $i \in [1, 5]$ is*

$$p_{\text{dist}}(A_1, \dots, A_5) := \frac{s! \binom{5}{1}^{A_1} \binom{5}{2}^{A_2} \binom{5}{3}^{A_3} \binom{5}{4}^{A_4} \binom{5}{5}^{A_5}}{(s - A_1 - A_2 - A_3 - A_4 - A_5)! A_1! A_2! A_3! A_4! A_5! \binom{5s}{n}}, \quad (4)$$

with $n = A_1 + 2A_2 + 3A_3 + 4A_4 + 5A_5$.

Important Remark. Since most matches happen when all the Sboxes are active, in order to simplify the analysis, we will use from now on only forward and backward paths such that *all Sboxes are active in the χ layer of the inbound phase*.

5.4 The Differential Paths Sets

In this section, we explain how we generate the forward and backward paths, since this will have an impact on the derivation of p_{match} and N_{match} (this will be handled in the next two sections).

The Forward Paths. For the forward paths set (see Fig. 2), we start by choosing a differential trail computed from the previous section and we derive a set from it by exhausting all the possible Sbox differential transitions for the inverse of the χ layer in its first round (all the paths will be the same except the differences on their input and on the input of the χ layer in the first round). For example, we can use the 2 first rounds of the 4-round differential path we found (see full version) which have a total success probability 2^{-24} and present 6 active Sboxes during the χ layer of the first round. We randomize the χ^{-1} layer




Fig. 2. The forward trails we are using. The distance between the two lines reflects the number of differences.

differential transitions for the 6 active Sboxes of the first round, and we obtain about 2^{19} distinct trails in total. We analyzed that all the trails of this set have a success probability of at least $2^{-24} \cdot 2^{-2.6} = 2^{-36}$ (this is easily obtained with the χ^{-1} DDT). Moreover, note that they will all have the same output difference mask (at the third round), but distinct input masks (at the first round). Since we previously forced the requirement that all Sboxes must be active for the inbound match, we check experimentally that $2^{17.3}$ of the 2^{19} members of the set fulfill this condition.¹ We call τ_F the ratio of paths that verify this condition over the total number of paths, i.e., $\tau_F = 2^{-1.7}$. Overall, we built a set of $2^{17.3}$ forward differential paths on $nr_F = 2$ KECCAK- $f[1600]$ rounds, all with DP higher or equal to $p_F = 2^{-36}$. We can actually generate 64 times more paths by observing that they are equivalent by translation along the KECCAK lane (the z axis). However, these paths will have distinct output difference masks (the same difference mask rotated along the z axis), and we have $\Gamma_F^{\text{mid}} = 2^6$. The total amount of input differences is $\Gamma_F^{\text{in}} := \Gamma_F^{\text{mid}} \cdot 2^{17.3} = 2^{23.3}$ and we have to generate in total $n_F = \tau_F \cdot \Gamma_F^{\text{in}} = 2^{25}$ forward differential paths. We discuss the amount of output differences in Section 5.8, since we extend there the path with two free additional rounds.

The Backward Paths. Applying the same technique to the backward case does not generate a sufficient amount of output differences Γ_B^{out} , crucial for a rebound-like attack. Thus, concerning the backward paths set, we build another type of 2-round trails. We need first to ensure that we have enough differential paths to be able to find a match in the inbound phase, i.e., we want $\Gamma_B^{\text{out}} \cdot \Gamma_F^{\text{in}} = 1/p_{\text{match}} \cdot \lceil \frac{1}{p_F \cdot p_B \cdot N_{\text{match}}} \rceil$ following (2). Moreover, we will require these paths to verify two conditions:

1. First, we need to filter paths that have not all Sboxes active in the χ layer of the inbound phase. This happens with a probability about $\tau_B^{\text{full}} : = b_{\text{bucket}}(800, 320) / \binom{5 \cdot 320}{n} = 2^{-15.9}$ if we assume that about half of the bits

¹ The small amount of filtered forward paths (a factor $2^{1.7}$) is due to the regularity of the output of θ inverse. Thus, most of the paths have all Sboxes active when the Hamming weight of the input is low.




Fig. 3. The backward trails we are using. The distance between the two arrows reflects the number of differences.

are active. This assumption will be verified in our case (and was verified in practice) since our control on the diffusion of the active bits will be reduced greatly.

- Moreover, *all* the paths we collect should have a DP of at least p_B such that the number of solutions N_{match} generated in the inbound phase is sufficient. Indeed, we must have $N_{\text{match}} \geq 1/(p_F \cdot p_B)$ in order to have a good success probability to find one solution for the entire path. We call τ_B^{DP} the probability that a path verifies this property. Hence, we need $p_B \geq p_B^{\text{min}} = 1/(p_F \cdot N_{\text{match}})$. We will show in Section 5.7 that $N_{\text{match}} = 2^{486}$ and we previously showed that $p_F = 2^{-36}$. Hence, $p_B^{\text{min}} = 2^{36-486} = 2^{-450}$.

These two filters induce a ratio $\tau_B := \tau_B^{\text{full}} \cdot \tau_B^{\text{DP}}$ of “good” paths. We have $n_B \cdot \tau_B = \Gamma_B^{\text{out}}$, where n_B is the number of paths we need to generate. Thus, we need to generate $n_B^{\text{min}} := 1/(p_{\text{match}} \cdot \lceil \frac{1}{p_F \cdot p_B \cdot N_{\text{match}}} \rceil \cdot \Gamma_F^{\text{in}} \cdot \tau_B)$ trails to perform the rebound. We will show in Section 5.7 that $p_{\text{match}} = 2^{-491.47}$, that $\lceil \frac{1}{p_F \cdot p_B \cdot N_{\text{match}}} \rceil = 1$ and that $\tau_B = 2^{-15.9}$. We also know that $\Gamma_F^{\text{in}} = 2^{23.3}$. Hence, $n_B^{\text{min}} = 2^{491.47+15.9-23.3} = 2^{484.07}$.

We show now how we generated these paths. Fig. 3 can help the reading. We start at the beginning of the second round by forcing X columns of the internal state to be active and each active column will contain only 2 active bits (thus a total of $2X$ active bits). Therefore, we can generate $\binom{5}{2}^X \cdot \binom{s}{X}$ distinct starting differences and each of them will lead to a distinct input difference of the backward path. Note also that all active columns are in the CPK and thus applying the θ function on this internal state will leave all bit-differences at the same place. Then, applying the ρ and π layers will move the $2X$ active bits to random locations before the Sbox layer of the second round. If X is not too large, we can assume that for a good fraction of the paths, all active bits are mapped to distinct Sboxes and thus we obtain $2X$ active Sboxes, each with one active bit on its input. We call ϵ this fraction of paths which is given by

$$\epsilon := p_{\text{dist}}(2X, 0, 0, 0, 0), \tag{5}$$

where p_{dist} is given by Lemma 2.² We will need to take ϵ into account when we count the total number of paths we can generate. This position in the paths, i.e., after the linear layer of the second round, is the part with the lowest amount of distinct differences. Hence, we call the number of differences at this point $\Gamma_B^{\text{mid}}(X) := \binom{5}{2}^X \cdot \binom{s}{X} \cdot \epsilon$.

Looking at the DDT from χ , one can check that with one active input bit in an Sbox, there always exists:

- two distinct transitions with probability 2^{-2} for the KECCAK Sbox such that we observe 2 active bits on its output (we call it a $1 \mapsto 2$ transition)
- one single transition with probability 2^{-2} and one single active bit on its output (a $1 \mapsto 1$ transition). This transition is in fact the identity.

We need to define how many $1 \mapsto 1$ and $1 \mapsto 2$ transitions we have to use, since there is a tradeoff between the number of paths obtained and the DP of these paths. Whatever choices we make, we always have that the success probability of this χ transition (in the second round) is 2^{-4X} . Let k be the number of $1 \mapsto 2$ transitions among the $2X$ possible ones. We will observe $2X + k$ active bits after χ . Before the χ transition, we have $\Gamma_B^{\text{mid}}(X)$ different paths from the initial choice. For each of these paths, we can now select $\binom{2X}{k}$ distinct sets of $1 \mapsto 2$ transitions each of which can generate 2^k different paths. These $2X + k$ bits are expanded through θ to *at most* $11 \cdot (2X + k) = 22X + 11k$ bits. However, this expansion factor (every active bit produces 11 one) is smaller when the number of bits increases. Let n be the number of obtained active bits at the input of the Sboxes in the third round. At the beginning of the third round, we have $2X + k$ active bits. For KECCAK- f [1600], given $2X + k$ active bits at the input of θ , we get $n \approx u - (u \cdot (u - 1))/1600$ bits at the output, with $u := 11(2X + k)$ for X small enough. Indeed, the $2X + k$ bits are first multiplied by 11 due to the property of θ . We suppose now that these u active bits are thrown randomly and we check for collisions. Given u bits, we can form $u \cdot (u - 1)/2$ different pairs of bits. The probability that a pair collides is 2^{-1600} , hence, we have about $u \cdot (u - 1)/(2 \cdot 1600)$ collisions of two bits. In a 2-collision, two active bits are wasted (they become inactive). Hence, we can remove $u \cdot (u - 1)/1600$ from the overall number of active bits. For small X , we can neglect collisions between three, four and five active bits, since the bits before θ are most likely separated and will not collide. Hence, we verify the equation for n . This model has been verified in simulations for the values we are using.

We need now to evaluate the number of active Sboxes in the χ layer of the third round. However, in order to precisely evaluate the DP of this layer (that we want to be higher than p_B^{min}) and the expansion factor we get on the amount of distinct differential paths, we also need to look at how the bits are distributed into the input of the Sboxes. The probability p_{dist} of distributing randomly n active bits into s 5-bit Sboxes such that exactly A_i Sboxes contain i bits, for $i \in [1, 5]$ is given by Lemma 2.

² Simulations verified this behavior in practice for the parameters we use in our attack.

Lemma 3. *Suppose that we have n active bits before χ in the third round. Then, if $n \leq s$, the expected number of useful (i.e., which have $\text{DP} \geq p_B^{\min}$) paths $G_B(n)$ we can generate verifies*

$$G_B(n) \geq \sum_{A_5=0}^{\lfloor \frac{n}{5} \rfloor} \sum_{A_4=0}^{N_4} \sum_{A_3=0}^{N_3} \sum_{A_2=0}^{N_2} F(X, A_1, \dots, A_5) \cdot 2^{2A_1+3A_2+3.58A_3+4(A_4+A_5)}, \quad (6)$$

where $N_4 := \lfloor (n - 5A_5)/4 \rfloor$, $N_3 := \lfloor (n - 5A_5 - 4A_4)/3 \rfloor$,
 $N_2 := \lfloor (n - 5A_5 - 4A_4 - 3A_3)/2 \rfloor$, $A_1 := n - 5A_5 - 4A_4 - 3A_3 - 2A_2$, and

$$F(X, A_1, \dots, A_5) := \begin{cases} p^{\text{dist}}(A_1, \dots, A_5) & \text{if } 2^{-4X-2A_1-3A_2-4(A_3+A_4+A_5)} \geq p_B^{\min} \\ 0 & \text{else.} \end{cases} \quad (7)$$

Note that we use $F(\dots)$ to filter the paths that have a too low DP.

The proof is given in the full version of the paper.

In practice, we compute $G_B(n)$ by summing over all possible values of A_1, \dots, A_5 , such that $n = A_1 + 2A_2 + 3A_3 + 4A_4 + 5A_5$.

We have now reached the inbound round and we discard all the paths that do not have all Sboxes active. Hence, we keep only a fraction of $\tau_B^{\text{full}} = 2^{-15.9}$ paths.

It is now easy to see that

$$\tau_B^{\text{DP}} := \sum_{A_5=0}^{\lfloor n/5 \rfloor} \sum_{A_4=0}^{\lfloor (n-5A_5)/4 \rfloor} \sum_{A_3=0}^{\lfloor (n-5A_5-4A_4)/3 \rfloor} \sum_{A_2=0}^{\lfloor (n-5A_5-4A_4-3A_3)/2 \rfloor} F(X, A_1, A_2, A_3, A_4, A_5) \quad (8)$$

with $F(\dots)$ defined in (7) since this is exactly the fraction of path we keep.

To summarize, we have now reached the inbound round and we are able to generate

$$\Gamma_B^{\text{out}} = \epsilon \cdot \binom{5}{2}^X \cdot \binom{s}{X} \cdot \binom{2X}{k} \cdot 2^k \cdot G_B(n) \cdot \tau_B^{\text{full}} \quad (9)$$

differences that have a good DP and all Sboxes active and the total number of paths we have to generate is $n_B = \Gamma_B^{\text{out}}/\tau_B = \Gamma_B^{\text{out}}/(\tau_B^{\text{full}} \cdot \tau_B^{\text{DP}})$.

By playing with the filter bound, we noticed the following behavior. The stronger the filter is (i.e., the higher we set the bound on the DP), the higher the expected value of the Hamming weight at the input of the Sboxes of the inbound phase will be. This behavior will allow us to reduce the complexity of our attack in Section 5.7, where we discuss the numerical application. Hence, instead of filtering at p_B^{\min} , we will filter at a higher value to get better results.

5.5 The Inbound Phase

Now that we have our forward and backward sets of differential paths, we need to estimate the average probability p_{match} that two trails can match during the inbound phase of the rebound attack. We recall that we already enforced all Sboxes

to be active during this match, so p_{match} only takes into account the probability that the differential transitions through all the s Sboxes of the internal state are possible.

A trivial method to estimate p_{match} would be to simply consider an average case on the KECCAK Sbox. More precisely, the average probability that a differential transition is possible through the KECCAK Sbox, given two random non-zero 5-bit input/output differences is equal to $2^{-1.605}$. Thus, one is tempted to derive $p_{\text{match}} = 2^{-1.605 \cdot s}$. However, we observed experimentally that the event of a match greatly depends on the **Hamming weight of the input of the Sboxes**. Note that *this effect is only strong regarding the input of the Sbox*, but there is no strong bias on the differential matching probability concerning the output Hamming weight.

Therefore, in order to model more accurately the input Hamming weight effect on the matching event, we first divide the backward paths *depending on their Hamming weight* and treat each class separately. More precisely, we look at each possible input Hamming weight division among the s Sboxes. To represent this division, we only need to look at the number of Sboxes having a specific input Hamming weight (their relative position do not matter). We denote by c_i the number of Sboxes having an input Hamming weight i and we need the following equations to hold $\sum_{i=1}^5 c_i = s$ since we forced that all Sboxes are active during a match. Moreover, for a Hamming weight value w , we have $\sum_{i=1}^5 i \cdot c_i = w$. The set of divisions c_i verifying the above mentioned equations is denoted by C_w . The number of possible $5s$ -bit vectors satisfying (c_1, \dots, c_5) (i.e., c_1 Sboxes with 1 active bit, c_2 with 2 etc.) is denoted $b_c(c_1, \dots, c_5)$ and

$$b_c(c_1, \dots, c_5) = \frac{s!}{c_1!c_2!\dots c_5!} \cdot 5^{c_1+c_4} \cdot 10^{c_2+c_3} . \tag{10}$$

We can now compute the probability of having a match p_{match} depending on the input Hamming weight divisions:

Theorem 1. *The probability p_{match} of having a match is*

$$\sum_{w=s}^{5s} \Pr[\text{Hw}_{\text{tot}} = w | \text{full}] \sum_{(c_1, \dots, c_5) \in C_w} \frac{b_c(c_1, \dots, c_5)}{b_{\text{bucket}}(w, s)} \prod_{i=1}^5 \left(\sum_{y \in \{0,1\}^5} \sum_{\substack{v \in \{0,1\}^5: \\ \text{Hw}(v)=i}} \frac{P_{\text{out}}(y) \mathbb{1}_{\text{DDT}[v][y]}}{\binom{5}{i}} \right)^{c_i} \tag{11}$$

where $P_{\text{out}}(y)$ is the measured probability distribution of having y at the output of an Sbox when we enforce all Sboxes to be active, $\Pr[\text{Hw}_{\text{tot}} = w | \text{full}]$ is the measured probability distribution of the Hamming weight of the input of the Sboxes when all Sboxes are active, $b_c(\dots)$ is given by (10), $b_{\text{bucket}}(w, s)$ by Lemma 1 and $\mathbb{1}_{\text{DDT}[v][y]}$ is set to one if the entry $[v][y]$ is non-zero in the DDT of the χ layer and zero otherwise.³

³ Note that $\Pr[\text{Hw}_{\text{tot}} = w | \text{full}]$ greatly depends on the backward paths we choose and that these paths depends on p_{match} . We explain how to solve this cyclic dependency in Section 5.7.

We leave the proof to the complete version of the paper. However, we define the following intermediate result: $p_{\text{match}}(w) := \Pr[\text{match} | \text{Hw}_{\text{tot}} = w, \text{full}]$ which can be written as

$$\sum_{(c_1, \dots, c_5) \in C_w} \Pr[\text{match} | (c_1, \dots, c_5), \text{Hw}_{\text{tot}} = w, \text{full}] \cdot \Pr[(c_1, \dots, c_5) | \text{Hw}_{\text{tot}} = w, \text{full}]. \quad (12)$$

The measured distributions along with some intermediate values will be given in the extended version of the paper.

We require to test $1/p_{\text{match}}$ backward/forward paths combinations in order to have a good chance for a match. Note that in the next section, we will actually put an extra condition on the match in order to be able to generate enough values in the worst case during the outbound phase.

5.6 The Outbound Phase

Now that we managed to obtain a match with complexity $1/p_{\text{match}}$, we need to estimate how many solutions can be generated from this match. Again, one is tempted to consider an average case on the KECCAK Sbox: the average number of Sbox values verifying a non-zero random input/output difference such that the transition is possible is equal to $2^{1.65}$. The overall number of solutions would then be $2^{1.65 \cdot s}$. However, as for p_{match} , this number highly depends on the Hamming weight of the input of the Sboxes and this can be easily observed from the DDT of the χ layer (for example with an input Hamming weight of one the average number of solutions is 2^3 , while for an input Hamming weight of four the average number of solutions is 2^1).

In order to obtain the expected number of values N_{match} we can get from a match, we proceed like in the previous section and divide according to the input Hamming weight.

Theorem 2. *Let N be a random variable denoting the number of values we can generate. Let also full be the event denoting that all the Sboxes are active for the inbound phase. Given a Hamming weight of w at the input of the Sboxes, we can get $N_w := \mathbb{E}[N | \text{match}, \text{Hw}_{\text{tot}} = w, \text{full}]$ values from a match, with*

$$N_w = \frac{1}{p_{\text{match}}(w)} \sum_{(c_1, \dots, c_5) \in C_w} \prod_{i=1}^5 Z^{c_i} \cdot \frac{b_c(c_1, \dots, c_5)}{b_{\text{bucket}}(w, s)}, \quad (13)$$

with

$$Z := \frac{1}{\binom{5}{i}^2} \left(\sum_{\substack{v \in \{0,1\}^5: \\ \text{Hw}(v)=i}} \text{DDT}[v] \right) \sum_{y \in \{0,1\}^5} \sum_{\substack{v \in \{0,1\}^5: \\ \text{Hw}(v)=i}} P_{\text{out}}(y) \cdot \mathbb{1}_{\text{DDT}[v][y]},$$

where $\text{DDT}[v]$ is the value of the non-zero entries in line v of the DDT, $P_{\text{out}}(y)$ is the measured probability distribution of having y at the output of an Sbox when we enforce all Sboxes to be active, $p_{\text{match}}(w)$ is given by (12), $b_c(\dots)$ is given by (10), $b_{\text{bucket}}(w, s)$ is given by Lemma 1 and $\mathbb{1}_{\text{DDT}[v][y]}$ is set to one if the entry $[v][y]$ is non-zero in the DDT of the χ layer and zero otherwise.

The proof is given in the complete version of the paper.

One would be tempted to take the expected value of all the N_w and compute N_{match} as

$$\sum_w \mathbb{E}[N | \text{match}, \text{Hw}_{\text{tot}} = w, \text{full}] \cdot \Pr[\text{Hw}_{\text{tot}} = w | \text{match}, \text{full}] .$$

This expectancy would be fine if we were expecting a high number of matches. This is however not necessarily our case. Hence, we need to ensure that the number of values we can generate from the inbound is sufficient. To do this, first note that N_w decreases exponentially while w increases. Similarly, $p_{\text{match}}(w)$ increases exponentially while w increases. Thus, we are more likely to obtain a match at a high Hamming weight which will lead to an insufficient N_{match} .

To solve this issue, we proceed as follows. First, we compute N_w for every w . We look then for the maximum Hamming weight w_{max} we can afford, i.e., such that $N_{w_{\text{max}}} \geq 1/(p_B \cdot p_F)$. This way, we are ensured to obtain enough solutions from the match. However, we need to update our definition of a match: a match occurs only when the Hamming weight of the input is lower than w_{max} . Hence, instead of summing over all possible values of w , we sum only up to w_{max} and need to update (11). The number of values we can then obtain from the inbound is $N_{\text{match}} \geq N_{w_{\text{max}}}$.

We can now apply this model to the KECCAK- $f[1600]$ internal permutation. Some useful intermediate results and relevant $N_{w_{\text{max}}}$ (with their associated p_{match}) will be given in the extended version of the paper.

5.7 Finalizing the Attack and Improvements

In Section 5.4, we showed how to choose the backward paths given the probability of having a match in the inbound phase (p_{match}) and the number of solution we can generate from this match (N_{match}). In Sections 5.5 and 5.6, we showed how to compute p_{match} and N_{match} . However, in these computations, we needed the probability distribution of the Hamming weight of the input of the Sbox, $\Pr[\text{Hw}_{\text{tot}} = w | \text{full}]$. This probability depends greatly on the paths we select in Section 5.4.

To solve this circular dependency, we performed several iterations of the following algorithm until we found some parameters that verify all equations. First, we estimated roughly $\Pr[\text{Hw}_{\text{tot}} = w | \text{full}]$ by taking some random backward paths with limited complexity. Using the worst case cost of these paths, we were able to select w_{max} such that the number of values generated from the inbound is sufficient. Then, we computed p_{match} and N_{match} . With this first guess, we searched for an X and a k such that the we can find a match with a good probability and such that we can generate enough values from the inbound. Then, we computed $\Pr[\text{Hw}_{\text{tot}} = w | \text{full}]$ using these new paths generated by X , k and p_B and started our algorithm again with this new distribution. After some iterations, we found a set of filtered backwards paths that provided a sufficient p_{match} and N_{match} .

When $(X, k) = (8, 8)$, we have $\epsilon = 0.736$ and $\Gamma_B^{\text{mid}} = \binom{5}{2}^X \cdot \binom{s}{X} \cdot \epsilon = 2^{77.26}$. If we filter all paths that have a DP smaller than 2^{-450} , i.e., we set $p_B = 2^{-450}$, we get for $(X, k) = (8, 8)$ at least $\epsilon \cdot \binom{5}{2}^X \cdot \binom{s}{X} \cdot \binom{2X}{k} \cdot 2^k \cdot G_B(n) \cdot \tau_B^{\text{full}} = 2^{493.88-15.9} = 2^{477.98}$ distinct differences using (9) for the inbound (for these parameters, the difference Hamming weight at the input of the χ layer in the third round is $n = 227.9$). With these parameters, since we remove the paths with a DP lower than p_B , we keep $\tau_B^{\text{DP}} \approx 1 - 10^{-10}$ of the paths, following (8), i.e., we have almost no filtering on the DP. Hence, we filter the backward paths with a ratio $\tau_B = \tau_B^{\text{full}} \cdot \tau_B^{\text{DP}} \approx 2^{-15.9} \cdot (1 - 10^{-10}) = 2^{-15.9}$. We have also $p_B = 2^{-450}$ and $p_F = 2^{-36}$. Therefore, we need $N_{\text{match}} \geq 2^{486}$. Computations show that we have to set $w_{\text{max}} = 1000$. This leads to $p_{\text{match}} = 2^{-491.47}$. This implies that the minimum total number of backward paths we need to generate is $n_B^{\text{min}} = 1/(p_{\text{match}} \cdot \Gamma_F^{\text{in}} \cdot \tau_B) = 1/(p_{\text{match}} \cdot \Gamma_F^{\text{in}} \cdot \tau_B^{\text{full}}) = 2^{484.07}$. All these paths apply on $nr_B = 2$ KECCAK- f [1600] rounds, all with DP higher or equal to $p_B^{\text{min}} = 2^{36-486} = 2^{-450}$.

To summarize, we have that the number of backward output differences is $\Gamma_B^{\text{out}} = n_B^{\text{min}} \cdot \tau_B = 2^{484.07-15.9} = 2^{468.17}$ and that the number of forward input differences is $\Gamma_F^{\text{in}} = 2^{23.3}$. Hence, there is a total of $2^{491.47}$ couples of $(\Delta_B^{\text{out}}, \Delta_F^{\text{in}})$ for the inbound phase, which is enough since it is equal to $1/p_{\text{match}}$. Once a match is found, the worst case complexity of the connected path is $1/(p_B \cdot p_F) \leq 2^{450+48} = 2^{498}$ which is lower or equal to N_{match} . Hence, we can generate enough values from the inbound phase to find with a good probability values verifying the differential path.

The overall complexity for the rebound attack given by (1) is $C = 2^{491.47}$.

This model was verified on the KECCAK- f [100] internal permutation. By applying this attack on it, we found a 4-round result together with solution pairs. This gives a 6-round distinguisher with complexity $2^{28.76}$.

5.8 The Distinguisher

We will use exactly the same type of limited-birthday distinguishers as in Section 4.

Relaxing the Forward Paths. We analyze now the impact of this two additional paths on Γ_F^{out} , the set of reachable output differences. At the entrance of the third round, every Sbox has one single active bit. Hence, according to the DDT, there are only 4 different possibilities at the output of the Sboxes. Since we have 6 active Sboxes in the third round, the number of possible differences at the output of the third round is multiplied by $4^6 = 2^{12}$. Thus, the number of differences at the output of the third round is $\Gamma_F^{\text{mid}} \cdot 2^{12} = 2^6 \cdot 2^{12} = 2^{18}$.

We need now to look at the fourth round to obtain Γ_F^{out} and compute the generic complexity of the distinguisher. In the third round, every active Sbox can produce at most 3 active bits at its output, since each active Sbox has only one single active bit at its input. Hence, the maximum Hamming weight at the output is $3 \cdot 6 = 18$. Each of these active bits can be expanded to at most 11 bits through θ and hence, we have at most $11 \cdot 18 = 198$ active bits at the input

of the Sboxes of the fourth round. In the worst case, each of these bits will be in a different Sbox and will produce four possible differences. Hence, we have $\Gamma_F^{\text{out}} \leq \Gamma_F^{\text{mid}} \cdot 2^{12} \cdot 4^{198} = 2^{18} \cdot 2^{396} = 2^{414}$.

Relaxing the Backward Paths. Each Sbox with one single active bit at its output can have 9 possible input differences and the maximum possible of input differences that can occur for a given input difference is 12 (see χ^{-1} DDT). Since we have $2X$ active Sboxes, the number of possible input differences is increased by a factor of at most 9^{2X} . Therefore, $\Gamma_B^{\text{in}} \leq \Gamma_B^{\text{mid}} \cdot 9^{2X}/\epsilon$ and we reduced the complexity by a factor 2^{4X} .

We have $\Gamma_B^{\text{in}} \leq \Gamma_B^{\text{mid}}(8) \cdot 9^{2 \cdot 8}/\epsilon = 2^{77.7+50.7} = 2^{128.4}$ and $\Gamma_F^{\text{out}} \leq 2^{414}$. The generic complexity of the distinguisher is, hence, greater than $2^{1057.6}$. This is much greater than the complexity of the rebound attack $C = 2^{491.47}$.

6 Results and Conclusion

In this article, we analysed the internal permutations used in the KECCAK family of hash functions in regards to differential cryptanalysis. We first proposed a generic method that looks for the best differential paths using CPK considerations and better χ mapping. This new method provides some of the best known differential paths for the KECCAK internal permutations and we derived distinguishers with rather low complexity exploiting these trails. In particular we were able to obtain a practical distinguisher for 6 rounds of the KECCAK- f [1600] permutation. Then, aiming for attacks reaching more rounds, we adapted the rebound attack to the KECCAK case. This adaptation is far from trivial and contains many technical details. Our model was verified by applying the attack on the reduced version KECCAK- f [100]. The main final result is a 8-round distinguisher for the KECCAK- f [1600] internal permutation with a complexity of $2^{491.47}$. Our distinguisher results are summarized in Table 2. Note that our attack does not endanger the security of the full KECCAK. We believe that this work will also help to apply the rebound attack on a much larger set of primitives.

This work might be extended in many ways, in particular by further refining the differential path search or by improving the inbound phase of the rebound

Table 2. Best differential distinguishers complexities for each version of KECCAK internal permutations, for 1 up to 8 rounds. Note that due to its technical complexity when applied on KECCAK, the rebound attack has only been applied to KECCAK- f [100] and KECCAK- f [1600].

b	best differential distinguishers complexity							
	1 rd	2 rds	3 rds	4 rds	5 rds	6 rds	7 rds	8 rds
400	1	1	1	2^2	2^8	2^{24}	2^{84}	-
800	1	1	1	2^2	2^8	2^{32}	2^{109}	-
1600	1	1	1	2^2	2^8	2^{32}	2^{142}	$2^{491.47}$

attack such that the overall cost is reduced. Moreover, another research direction would be to analyse how the differential paths derived in this article can lead to collision attacks against reduced versions of the KECCAK hash functions. Using the techniques presented in [19] could help reducing the complexity of it.

Acknowledgements. The authors would like to thank the anonymous referees for their helpful comments. We are extremely grateful to Guido Bertoni, Joan Daemen, Michaël Peeters and Gilles Van Assche for their remarks on the first drafts of this paper. Finally, we are very grateful to Praveen Gauravaram, Tao Huang, Phuong Ha Nguyen, Wun-She Yap, Przemyslaw Sokolowski and Wenling Wu for useful discussions.

References

1. Abe, M. (ed.): ASIACRYPT 2010. LNCS, vol. 6477. Springer, Heidelberg (2010)
2. Aumasson, J.-P., Meier, W.: Zero-sum distinguishers for reduced Keccak-f and for the core functions of Luffa and Hamsi. Presented at the Rump Session of CHES 2009 (2009)
3. Bellare, M., Rogaway, P.: Random Oracles are Practical: A Paradigm for Designing Efficient Protocols. In: CCS, pp. 62–73. ACM (1993)
4. Bernstein, D.J.: Second preimages for 6 (?? (8??)) rounds of Keccak? (November 2010), http://ehash.iaik.tugraz.at/uploads/6/65/NIST-mailing-list_Bernstein-Daemen.txt
5. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: Sponge functions. In: ECRYPT Hash Workshop 2007 (May 2007)
6. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: On alignment in Keccak. In: ECRYPT II Hash Workshop (2011)
7. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: The KECCAK Reference. Submission to NIST (Round 3) (2011)
8. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: The KECCAK SHA-3 Submission. Submission to NIST (Round 3) (2011)
9. Boura, C., Canteaut, A., De Cannière, C.: Higher-Order Differential Properties of KECCAK and *Luffa*. In: Joux, A. (ed.) FSE 2011. LNCS, vol. 6733, pp. 252–269. Springer, Heidelberg (2011)
10. Duc, A., Guo, J., Peyrin, T., Wei, L.: Unaligned Rebound Attack - Application to Keccak. Cryptology ePrint Archive, Report 2011/420 (2011), <http://eprint.iacr.org/>
11. Gilbert, H., Peyrin, T.: Super-Sbox Cryptanalysis: Improved Attacks for AES-Like Permutations. In: Hong and Iwata [12], pp. 365–383
12. Hong, S., Iwata, T. (eds.): FSE 2010. LNCS, vol. 6147. Springer, Heidelberg (2010)
13. Khovratovich, D., Naya-Plasencia, M., Röck, A., Schläffer, M.: Cryptanalysis of *Luffa v2* Components. In: Biryukov, A., Gong, G., Stinson, D.R. (eds.) SAC 2010. LNCS, vol. 6544, pp. 388–409. Springer, Heidelberg (2011)
14. Khovratovich, D., Nikolic, I., Rechberger, C.: Rotational Rebound Attacks on Reduced Skein. In: Abe [1], pp. 1–19
15. Lamberger, M., Mendel, F., Rechberger, C., Rijmen, V., Schläffer, M.: Rebound Distinguishers: Results on the Full Whirlpool Compression Function. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 126–143. Springer, Heidelberg (2009)

16. Mendel, F., Peyrin, T., Rechberger, C., Schl affer, M.: Improved Cryptanalysis of the Reduced `Gr ostl` Compression Function, `ECHO` Permutation and AES Block Cipher. In: Jacobson Jr., M.J., Rijmen, V., Safavi-Naini, R. (eds.) SAC 2009. LNCS, vol. 5867, pp. 16–35. Springer, Heidelberg (2009)
17. Mendel, F., Rechberger, C., Schl affer, M., Thomsen, S.S.: The Rebound Attack: Cryptanalysis of Reduced Whirlpool and `Gr ostl`. In: Dunkelman, O. (ed.) FSE 2009. LNCS, vol. 5665, pp. 260–276. Springer, Heidelberg (2009)
18. Morawiecki, P., Srebrny, M.: A SAT-based preimage analysis of reduced Keccak hash functions. Presented at Second SHA-3 Candidate Conference, Santa Barbara (2010)
19. Naya-Plasencia, M.: How to Improve Rebound Attacks. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 188–205. Springer, Heidelberg (2011)
20. Naya-Plasencia, M., R ock, A., Meier, W.: Practical Analysis of Reduced-Round `KECCAK`. In: Bernstein, D.J., Chatterjee, S. (eds.) INDOCRYPT 2011. LNCS, vol. 7107, pp. 236–254. Springer, Heidelberg (2011)
21. Rijmen, V., Toz, D., Varici, K.: Rebound Attack on Reduced-Round Versions of JH. In: Hong and Iwata [12], pp. 286–303
22. Sasaki, Y., Li, Y., Wang, L., Sakiyama, K., Ohta, K.: Non-full-active Super-Sbox Analysis: Applications to `ECHO` and `Gr ostl`. In: Abe [1], pp. 38–55
23. Keccak team. Keccak Crunchy Crypto Collision and Pre-image Contest (2011), http://keccak.noekeon.org/crunchy_contest.html
24. Wang, X., Yin, Y.L., Yu, H.: Finding Collisions in the Full SHA-1. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 17–36. Springer, Heidelberg (2005)

Differential Propagation Analysis of Keccak

Joan Daemen and Gilles Van Assche

STMicroelectronics

Abstract. In this paper we introduce new concepts that help read and understand low-weight differential trails in KECCAK. We then propose efficient techniques to exhaustively generate all 3-round trails in its largest permutation below a given weight. This allows us to prove that any 6-round differential trail in KECCAK- f [1600] has weight at least 74. In the worst-case diffusion scenario where the mixing layer acts as the identity, we refine the lower bound to 82 by systematically constructing trails using a specific representation of states.

Keywords: cryptographic hash function, KECCAK, differential cryptanalysis, computer-aided proof.

1 Introduction

The goal of cryptanalysis is to assess the security of cryptographic primitives. Finding attacks or properties not present in ideal instances typically contributes to the cryptanalysis of a given primitive. Building upon previous results, attacks can be improved over time, possibly up to a point where the security of the primitive is severely questioned.

In contrast, cryptanalysis can also benefit from positive results that exclude classes of attacks, thereby allowing research to focus on potentially weaker aspects of the primitive. Interestingly, weaknesses are sometimes revealed by challenging the assumptions underlying positive results. Nevertheless, both attacks and positive results can be improved over time and together contribute to the understanding and estimation of the security of a primitive by narrowing the gap between what is possible to attack and what is not.

Differential cryptanalysis (DC) is a discipline that attempts to find and exploit predictable difference propagation patterns to break iterative cryptographic primitives [6]. For ciphers, this typically means key retrieval, while for hash functions, this is the generation of collisions or of second preimages. The basic version makes use of differential trails (also called characteristics or differential paths) that consist of a sequence of differences through the rounds of the primitive. Given such a trail, one can estimate its differential probability (DP), namely, the fraction of all possible input pairs with the initial trail difference that also exhibit all intermediate and final difference when going through the rounds.

A more natural way to characterize the power of trails in unkeyed primitives is by their weight w . In general the weight of a trail is the sum of the weight of its round differentials, where the latter is the negative of its binary logarithm. For

many round functions, including that of KECCAK- f and Rijndael, the weight equals the number of binary equations that a pair must satisfy to follow the specified differences. Assuming that these conditions are independent, the weight of the trail relates to its DP as $DP = 2^{-w}$ and exploiting such a trail becomes harder as the weight increases. For a primitive with, say, b input and output bits, the number of pairs that satisfy these conditions is then 2^{b-w} . The assumption of independence does not always apply. For instance, a trail with $w > b$ implies redundant or contradictory conditions on pairs, for which satisfying pairs may or may not exist. Another example where this independence assumption breaks down are the plateau trails that occur in Rijndael [10]. These trails, with weight starting from $w = 30$ for 2 rounds, have a DP equal to 2^{z-w} with $z > 0$ for a fraction 2^{-z} of the keys and zero for the remaining part. In general, they occur in primitives with strong alignment [3] and a mixing layer based on maximum-distance separable (MDS) codes.

In the scope of DC, positive results can be established by finding a lower bound on the weight of any trail over a specified number of rounds. For instance, the structure of Rijndael and the properties of its diffusion operations allow to analytically derive such lower bounds [9]. Such results can be transposed to the permutations underlying the hash function Grøstl [13]. Other examples include a lower bound on the number of active S-boxes in JH [18] or computer-aided proofs on the weight of trails in NOEKEON [8] and on the minimum number of active AND gates in MD6 [17,14].

KECCAK is a sponge function submitted to the SHA-3 contest [16,4,2]. Recently, new results were published on the differential resistance of this function and among those heuristic techniques were proposed to build low-weight differential trails [12,15]. These gave the currently best trails for 3, 4 and 5 rounds of the underlying permutation KECCAK- f [1600]. In particular, Duc et al. found a trail of weight 32 for 3 rounds, and this motivated us to systematically investigate whether trails of lower weight exist. Also, there are some similarities between KECCAK and MD6, but unlike MD6, the permutation used in the proposed SHA-3 candidate KECCAK has no significant lower bounds on the weight of trails. So the philosophy behind [17,14] was another source of inspiration and motivation for our research.

Lower bounds on symmetric trails were already proven in [4]. They provide lower bounds with weight above the permutation width on KECCAK- f [25] to KECCAK- f [200] but only partial bounds in the case of KECCAK- f [1600]. Thanks to the Matryoshka structure [4], a lower bound W on trails in KECCAK- f [25 w] implies a lower bound $W' = W \frac{w'}{w}$ on w -symmetric trail in KECCAK- f [25 w'] for $w' > w$. These are summarized in Table 1.

In this paper, we report on techniques to efficiently generate all the trails in KECCAK- f [1600] up to a given weight. We implemented these techniques in a computer program, which allowed us at this point to completely scan the space of 3-round differential trails up to weight 36. This confirmed that the trail found by Duc et al. has minimum weight and allowed us to demonstrate that there are no 6-round trails with weight below 74. These results are summarized in Table 2.

Table 1. Lower bounds above the permutation width on 1- to 8-symmetric trails [4]

w	Lower bound for KECCAK- $f[25w]$	Lower bound for KECCAK- $f[1600]$	
1	30 per 5 rounds	1920 per 5 rounds	tight
2	54 per 6 rounds	1728 per 6 rounds	tight
4	146 per 16 rounds	2336 per 16 rounds	non-tight
8	206 per 18 rounds	1648 per 18 rounds	non-tight

The source code of the classes and methods used in this paper is available in the KECCAKTOOLS package [5].

As a by-product of this trail search, this paper proposes new techniques to relate the properties of the θ mapping in KECCAK to the weight of differential trails. In the worst-case diffusion scenario where θ acts as the identity, we build upon the results of [4] and [15] to systematically construct so-called in-kernel trails using an efficient representation of states.

Table 2. Weight of differential trails in KECCAK- $f[1600]$

Rounds	Lower bound	Best known
3	32 (this work)	32 [12]
4	-	134 [7]
5	-	510 [15]
6	74 (this work)	1360 [4]
24	296 (this work)	-

Further discussions on how to exploit differential trails in KECCAK can be found in [3]. Also, the attacks in [11] combine algebraic techniques with a differential trail.

The paper is organized as follows. In Section 2, we recall the structure of KECCAK and mappings inside its round function. Section 3 focuses on how to represent and extend the differential trails of KECCAK. Section 4 sets up the overall strategy and Section 5 introduces a basic trail generation technique. The advanced techniques are covered in Sections 6 and 7, which address two complementary cases. Finally, Section 8 extends the results from 3 to 6 rounds.

2 KECCAK

KECCAK combines the sponge construction with a set of seven permutations denoted KECCAK- $f[b]$, with b ranging from 25 to 1600 bits [1,4]. In this paper, we concentrate on the permutation used in the SHA-3 submission, namely, KECCAK- $f[1600]$.

The state of KECCAK- $f[1600]$ is organized as a set of $5 \times 5 \times 64$ bits with (x, y, z) coordinates. The coordinates are always considered modulo 5 for x and

y and modulo 64 for z . A *row* is a set of 5 bits with given (y, z) coordinates, a *column* is a set of 5 bits with given (x, z) coordinates and a *slice* is a set of 25 bits with given z coordinate.

The round function of KECCAK- f [1600] consists of the following steps, which are only briefly summarized here. For more details, we refer to the specifications [4].

- θ is a linear mixing layer, which adds a pattern that depends solely on the parity of the columns of the state. Its properties with respect to differential propagation will be detailed and exploited in Section 6.
- ρ and π displace bits without altering their value. Jointly, their effect is denoted by $(x, y, z) \xrightarrow{\pi \circ \rho} (x', y', z')$, with (x, y, z) a bit position before ρ and π and (x', y', z') its coordinates afterward.
- χ is a degree-2 non-linear mapping that processes each row independently. It can be seen as the application of a translation-invariant 5-bit S-box. The differential propagation properties will be detailed below.
- ι adds a round constant. As it has no effect on difference propagation, we will ignore it in the sequel.

3 Representing and Extending Trails

In general, for a function f with domain \mathbb{Z}_2^b , we define the *weight* of a differential (u', v') as

$$w(u' \xrightarrow{f} v') = b - \log_2 |\{u : f(u) \oplus f(u \oplus u') = v'\}|.$$

If the argument of the logarithm is non-zero (i.e., the DP is non-zero), we say that u' and v' are *compatible*. Otherwise, the weight is undefined.

The weight of a trail is the sum of the weight of the differentials that compose this trail. In KECCAK- f , we specify differential trails with the differences before each round function. For clarity, we adopt a redundant description by also specifying the differences before and after the linear steps $\lambda = \pi \circ \rho \circ \theta$. An n -round trail is of the following form, where each b_i must be equal to $\lambda(a_i)$,

$$Q = a_0 \xrightarrow{\pi \circ \rho \circ \theta} b_0 \xrightarrow{\chi} a_1 \xrightarrow{\pi \circ \rho \circ \theta} b_1 \xrightarrow{\chi} a_2 \xrightarrow{\pi \circ \rho \circ \theta} \dots \xrightarrow{\chi} a_n, \tag{1}$$

and has weight $w(Q) = \sum_i w(a_i \xrightarrow{\chi \circ \pi \circ \rho \circ \theta} a_{i+1})$. Since $b_i = \lambda(a_i)$, this expression simplifies to $w(Q) = \sum_i w(b_i \xrightarrow{\chi} a_{i+1})$.

3.1 Extending Forward and Trail Prefixes

Given a trail as in (1), it is possible to characterize all states that are compatible with $b_n = \lambda(a_n)$ through χ and thus to find all $n + 1$ -round trails Q' that have Q as its leading part. This process is called *forward extension*.

The χ mapping has algebraic degree 2 and, for a given input difference b_n , the space of compatible output differences forms a linear affine variety $\mathcal{A}(b_n)$ with

$|\mathcal{A}(b_n)|$ elements [4]. For a compatible a_{n+1} , the weight $w(b_n \xrightarrow{\chi} a_{n+1})$ depends only on b_n and is equal to $w(b_n) \triangleq \log_2 |\mathcal{A}(b_n)|$, with the symbol \triangleq denoting a definition. As χ operates on each row independently, the weight $w(b)$ can also be computed on each row independently and summed. To construct $\mathcal{A}(b)$, the bases resulting from each active row are gathered. Table 3 displays offsets and bases for the affine spaces of all single-row differences.

Table 3. Space of possible output differences, weight, minimum reverse weight and Hamming weight of all row differences, up to cyclic shifts

Difference	forward propagation					$w(\cdot)$	$w^{\text{rev}}(\cdot)$	$\ \cdot\ $
	offset	base elements						
00000	00000					0	0	0
00001	00001	00010	00100			2	2	1
00011	00001	00010	00100	01000		3	2	2
00101	00001	00010	01100	10000		3	2	2
10101	00001	00010	01100	10001		3	3	3
00111	00001	00010	00100	01000	10000	4	2	3
01111	00001	00011	00100	01000	10000	4	3	4
11111	00001	00011	00110	01100	11000	4	3	5

As a consequence, the weight of a n -round trail Q is $w(Q) = \sum_{i=0}^{n-1} w(b_i)$ and depends only on the n -tuple (b_0, \dots, b_{n-1}) . We call the latter a *trail prefix*. All n -round trails sharing this trail prefix and with a_n compatible with b_{n-1} through χ have the same weight.

3.2 Extending Backward and Trail Cores

Similarly, given a trail as in (1), it is possible to construct all states that are compatible with a_0 through χ^{-1} and thus to find all $n + 1$ -round trails Q' that have Q as its trailing part. This process is called *backward extension*. In contrast to χ , its inverse has algebraic degree 3 and the space of compatible differences is not an affine variety in general. Yet, compatible values can be identified per active row and combined.

For a difference a after χ , we define the *minimum reverse weight* $w^{\text{rev}}(a)$ as the minimum weight over all compatible b before χ . Namely,

$$w^{\text{rev}}(a) \triangleq \min_{b : a \in \mathcal{A}(b)} w(b).$$

Like for the restriction weight, the minimum reverse weight $w^{\text{rev}}(a)$ can be computed on each row independently and summed. Values are also shown in Table 3.

Given a $n - 1$ -round trail prefix $Q = (b_1, \dots, b_{n-1})$, it is easy to construct a difference b_0 such that the trail prefix $Q' = b_0 || Q$ has weight given by $w(Q') = w(Q) + w^{\text{rev}}(\lambda^{-1}(b_1))$. This is the smallest possible weight a n -round trail can

have with Q as its trailing part. It follows that a sequence of $n - 1$ state values $\tilde{Q} = (b_1, \dots, b_{n-1})$ defines a set of n -round trails with a weight at least

$$\tilde{w}(\tilde{Q}) \triangleq w^{\text{rev}}(\lambda^{-1}(b_1)) + \sum_{i=1}^{n-1} w(b_i).$$

We denote the former by the term *trail core* and the latter by its weight. Note that a n -round trail core is determined by only $n - 1$ states, although its weight takes n individual weights into account.

KECCAKTOOLS implements the representation of trails, trail prefixes and trail cores (see the `Trail` class), as well as the forward and backward extension (see the `KeccakFTrailExtension` class) [5].

4 Towards a Bound for Trails in KECCAK- f [1600]

To find a lower bound on differential trail weights in KECCAK- f [1600], our strategy is the following.

- First, we exhaustively generate all 3-round trails up to a given weight T_3 . There exists a trail of weight 32 as found by Duc et al. [12]. So by scanning the space of trails up to weight $T_3 \geq 32$, we are sure to hit at least one trail and the trail with minimum weight yields a tight lower bound on 3-round trails.
- Second, we derive a lower bound, not necessarily tight, on the weight of 6-round trails by using the 3-round trails found. Any 6-round trail of weight $2T_3 + 1$ or less satisfies either $w(b_0) + w(b_1) + w(b_2) \leq T_3$ or $w(b_3) + w(b_4) + w(b_5) \leq T_3$. We thus use forward and backward extension from 3-round trails up to weight $2T_3 + 1$. If such trails are found, the one with the smallest weight defines the lower bound, which is naturally tight. Otherwise, this establishes a lower bound for the weight of 6-round trails to $2T_3 + 2$. In the latter case no trail with weight $2T_3 + 2$ is known so the bound is not necessarily tight.

The reason for targeting 3-round trails in the first phase is the following. The minimum weight of a 1-round trail is 2, with a single active bit in b_0 . For the 24 rounds of KECCAK- f [1600], this amounts to a lower bound of $24 \times 2 = 48$. Constructing a state a with only two active bits in the same column leads to 2-round trail core with weight 8. Hence, if we base ourselves only on 2-round trail, we reach a lower bound of $12 \times 8 = 96$. If the 3-round trail of weight 32 found by Duc et al. [12] has minimum weight, this would mean that a 24-round trail has weight at least $8 \times 32 = 256$. Also, 3-round trail cores can be constructed by taking into account conditions across one layer of χ . Generating exhaustively trails of 4 rounds or more up to some weight would probably yield better bounds, but at the same time it is more difficult as several layers of χ must be dealt with. Instead, the two-step approach described above can take advantage of the exhaustive set of trails covered (i.e., all up to weight T_3) to derive a bound based on T_3 instead of on the minimum weight over 3 rounds.

4.1 Generating all 3-Round Trails Up to a Given Weight

In our approach we generate all 3-round differential trails of the form

$$Q = a_0 \xrightarrow{\pi \circ \rho \circ \theta} b_0 \xrightarrow{\chi} a_1 \xrightarrow{\pi \circ \rho \circ \theta} b_1 \xrightarrow{\chi} a_2 \xrightarrow{\pi \circ \rho \circ \theta} b_2 \xrightarrow{\chi} a_3, \quad (2)$$

up to some weight limit $w(Q) \leq T_3$. We call this the target space. We do this by searching for all trail cores (b_1, b_2) with weight below T_3 . Each such trail core (b_1, b_2) thus represents a set 3-round trails of the form of Eq. (2) with weight not below that of its core. In the scope of this paper, we limited ourselves to $T_3 = 36$.

We covered the set of all 3-round trails up to weight T_3 in three sub-phases:

1. In Section 5, we start with all cores such that $w^{\text{rev}}(\lambda^{-1}(b_1)) \leq 7$, $w(b_1) \leq 7$ or $w(b_2) \leq 7$.
2. In Section 6, we generate all remaining cores, except where both a_1 and a_2 are in the kernel.
3. In Section 7, we finish by generating all cores where both a_1 and a_2 are in the kernel.

4.2 Too Many States to Generate and Extend, Even When Exploiting Symmetry

A way to generate all trails in the target space is to first generate all states up to a given weight and then do backward and forward extensions to obtain trail cores. If we define $T_1 \triangleq \lfloor \frac{T_3}{3} \rfloor$, then for $\tilde{w}(b_1, b_2) \leq T_3$ either $w^{\text{rev}}(\lambda^{-1}(b_1)) \leq T_1$, $w(b_1) \leq T_1$ or $w(b_2) \leq T_1$. To cover the target space, we need to consider these cases:

- $w^{\text{rev}}(\lambda^{-1}(b_1)) \leq T_1$, so we have to generate all states a_1 with $w^{\text{rev}}(a_1) \leq T_1$, compute $b_1 = \lambda(a_1)$ and extend forward the 2-round trail cores (b_1) to get 3-round trail cores.
- $w(b_1) \leq T_1$, so we have to generate all states b_1 with $w(b_1) \leq T_1$ and extend forward the 2-round trail cores (b_1) .
- $w(b_2) \leq T_1$, so we have generate all states b_2 with $w(b_2) \leq T_1$ and extend backward the 2-round trail cores (b_2) .

Unfortunately, this brute-force strategy requires a high number of states to cover the whole space for an interesting target weight. E.g., if $T_3 = 36$, then $T_1 = 12$ and there are about $1.42 \times 10^{15} \approx 2^{50}$ states with weight up to 12 in KECCAK- f [1600].

We can reduce this number by taking the z symmetry into account. Except for ι , which does not influence difference propagation, all the step mappings of KECCAK- f are invariant when translated along z . Hence, for each trail $Q = (b_0, b_1, \dots, b_n)$ there exists a trail $Q' = (z(b_0), z(b_1), \dots, z(b_n))$ of same weight, with z the translation operator along the z axis. In the sequel, we will always consider trails up to translations in z . This reduces the search space by approximately a factor $w = 64$ —not exactly a factor w because of states that are periodic in z . Yet, the number of states to extend forward and backward is still about 2^{44} .

5 Generating Trails with a Low Number of Active Rows

In this section, we generate and extend states with weight up to $T'_1 = 7$. This does not cover the whole target space with $T_3 = 36$ but the remaining portion of the target space is limited to trails with a more flat weight profile, i.e., they satisfy $w(b_i) \geq T'_1 + 1 = 8$ for all $i \in \{0, 1, 2\}$ and $w(b_i) + w(b_{i+1}) \leq T'_2 = T_3 - (T'_1 + 1) = 28$ for all $i \in \{0, 1\}$.

More specifically, in this phase we look at the number of active rows in order to generate all trail cores such that $w^{\text{rev}}(\lambda^{-1}(b_1)) \leq T'_1$, $w(b_1) \leq T'_1$ or $w(b_2) \leq T'_1$, for $T'_1 = 7$. According to Table 3, each active row contributes for at least 2 to the weight. Hence,

$$w(b) \geq 2\|b\|_{\text{row}} \quad \text{and} \quad w^{\text{rev}}(b) \geq 2\|b\|_{\text{row}},$$

and we can cover all the states up to weight 7 by generating all states with up to $\lfloor \frac{T'_1}{2} \rfloor = 3$ active rows.

This approach can be refined by looking at the number of active rows not only for one state but for two consecutive states. With χ , the minimum weight a round differential can have is 2. So, $w^{\text{rev}}(\lambda^{-1}(b_1)) \geq 2$ implies that $w^{\text{rev}}(\lambda^{-1}(b_2)) + w(b_2) \leq w(b_1) + w(b_2) \leq T_3 - 2 = 34$ and similarly $w(b_2) \geq 2$ implies that $w^{\text{rev}}(\lambda^{-1}(b_1)) + w(b_1) \leq T_3 - 2 = 34$. Hence,

$$w^{\text{rev}}(\lambda^{-1}(b_i)) + w(b_i) \leq T_3 - 2 = 34 \Rightarrow \|\lambda^{-1}(b_i)\|_{\text{row}} + \|b_i\|_{\text{row}} \leq \left\lfloor \frac{T_3 - 2}{2} \right\rfloor = 17.$$

In practice, what we did was the following.

- Generate $\mathcal{B} = \{b : (\|b\|_{\text{row}} \leq 3 \text{ or } \|\lambda^{-1}(b)\|_{\text{row}} \leq 3) \text{ and } \|\lambda^{-1}(b)\|_{\text{row}} + \|b\|_{\text{row}} \leq 17\}$. This is done by first generating all states b with up to 3 active rows and filter on $\|\lambda^{-1}(b)\|_{\text{row}}$, and then generate all states a with up to 3 active rows, compute $b = \lambda(a)$ and filter on $\|b\|_{\text{row}}$.
- Do forward extension of all $b_1 \in \mathcal{B}$ and keep the cores $\tilde{Q} = (b_1, b_2)$ with $\tilde{w}(\tilde{Q}) \leq T_3$.
- Do backward extension of all $b_2 \in \mathcal{B}$ and keep the cores $\tilde{Q} = (b_1, b_2)$ with $\tilde{w}(\tilde{Q}) \leq T_3$.

We found a trail core (b_1, b_2) with $w^{\text{rev}}(\lambda^{-1}(b_1)) + w(b_1) + w(b_2) = 4 + 4 + 24 = 32$ (see also Table 4). It contains the 3-round trail found by Duc et al. [12].

There are $\binom{320}{n}(31)^n$ states with n active rows. As this function grows very quickly, it was not reasonable to extend this search beyond 3 active rows.

The generation of trail cores based on a small number of active rows is implemented in the `KeccakFTrailCoreRows` class [5].

6 Generating Trails Using the Properties of θ

To investigate the remaining part of the target space, we look at the properties of states a with respect to θ , and specifically the parity of its columns, to limit

the weight of two-round trails. An important parameter to classify the states a is their column parity, so as to study states in sets of parities. From the column parity, we derive the θ -gap, defined below. With θ -gap g , the effect of θ is to flip $10g$ bits. There are thus at least $10g$ active bits, each either in a or in $\theta(a)$. So, the higher the θ -gap the higher $w^{\text{rev}}(a) + w(\lambda(a))$ is likely to be. We can efficiently compute a lower bound for $w^{\text{rev}}(a) + w(\lambda(a))$ over all a with a given parity. For the target weights considered in this paper, this allows us to limit the states to consider to those with a parity belonging to a mere handful of values.

We then use the generated states a to build trail cores by forward and backward extension. As the θ -gap increases, the number of states a to consider decreases since more states a can immediately be excluded. An important case is when all the columns of a have even parity, i.e., a is in the kernel. In this case, the θ -gap is zero and a high number of states must be generated and extended. For this reason, this section focuses only the case where either a_1 or a_2 is not in the kernel. The complementary case is covered in Section 7.

6.1 Properties of θ

As θ is a linear function, its properties are the same whether applied on a state absolute value or on a difference, so we just write “value”. The following definitions are from [4].

The *column parity* (or *parity* for short) $P(a)$ of a value a is defined as the parity of the columns of a , namely $P(a)[x][z] = \sum_y a[x][y][z]$. A column is *even* (resp. *odd*) if its parity is 0 (resp. 1). The parity can also be defined on a slice, namely $P(a_z)[x] = \sum_y a[x][y][z]$. When the parity of a state or of a slice is zero (i.e., all its columns are even), we say it is in the *column-parity kernel* (or *kernel* for short).

The mapping θ consists in adding a pattern to the state, which we call the θ -effect. The θ -effect of a value a is $E(a)[x][z] = P(a)[x-1][z] + P(a)[x+1][z-1]$. For a fixed θ -effect $e[x][z]$, θ comes down to adding the y -symmetric pattern $e[x][y][z] \triangleq e[x][z](\forall y)$. So θ depends only on column parities and always affects columns symmetrically in y .

A column of coordinates (x, z) is *affected* iff $E(a)[x][z] = 1$; otherwise, it is *unaffected*. Note that the θ -effect always has an even Hamming weight so the number of affected columns is even.

The θ -gap is defined as the Hamming weight of the θ -effect divided by two. Hence, if the θ -gap of a value at the input of θ is g , the number of affected columns is $2g$ and applying θ to it results in $10g$ bits being flipped.

We have introduced the θ -gap via the θ -effect, but it can be defined directly using the parity itself. For this we introduce an alternative, single-dimensional, representation of a parity $p[x][z]$. We map the (x, z) coordinates to a single coordinate t as $t \rightarrow (x, z) = (-2t, t)$ and denote the result by $p[t]$. In this representation a *run* is a sequence of ones delimited by zeroes. As illustrated on Figure 1, each run induces two affected columns. First, if it starts in coordinates (x, z) , it implies an affected column in its right neighbor $(x+1, z)$. And if it ends in (x', z') it implies an affected column in its top-left neighbor $(x'-1, z'+1)$. The following lemma links the number of runs to the θ -gap.




Fig. 1. Example of parity pattern. Each square represents a column. An odd column contains a circle, while an affected column is denoted by a dot. A column can be both odd and affected. The odd columns of a run are connected with a line. The affected columns due to a run are located at the right (resp. top left) of the start (resp. end) column of the run.

Lemma 1. *The parity p has θ -gap g iff $p[t]$ has g runs.*

6.2 The Propagation Branch Number

The *propagation branch number* of a parity p is the minimum weight of the 2-round trail core (b) among states with this parity. More formally,

$$B(p) \triangleq \min\{\tilde{w}(b) : P(\lambda^{-1}(b)) = p\}.$$

Owing to the portion of the target space already covered in Section 5, we can limit the propagation branch number to $T'_2 = 28$. The strategy is as follows:

- First, we identify and exclude parity patterns p such that the propagation branch number can be proven to exceed $T'_2 = 28$.
- Then, for the remaining parity patterns p we look for all states $b = \lambda(a)$ with $P(a) = p$ and $\tilde{w}(b) \leq T'_2 = 28$.
- Finally, we forward and backward extend the states seen as 2-round trail cores up to weight $T_3 = 36$.

Clearly, the kernel states, i.e., states such that $P(a) = 0$ must be considered. For instance, a state a with just two active bits in the same column will have $w^{\text{rev}}(a) = 4$. Then, $b = \lambda(a) = \pi(\rho(a))$ since θ has no effect in this case, and b

also has two active bits. For KECCAK- f [1600], all the rotation constants in ρ are different and these two bits will not be in the same slice, so not in the same row and $w^{\text{rev}}(a) + w(b) = 8$. Hence, the propagation branch number of the all-zero parity is at least 8 and thus the all-zero parity pattern must be included.

States that are out of the kernel are likely to have a higher propagation branch number. We now concentrate on how to find a lower bound on the propagation branch number of a given parity pattern.

6.3 Bounding the Row Branch Number

The *row branch number* of a parity p is the minimum number of active rows before and after λ among states with this parity. More formally,

$$B_{\text{rows}}(p) \triangleq \min\{\|\lambda^{-1}(b)\|_{\text{row}} + \|b\|_{\text{row}} : P(\lambda^{-1}(b)) = p\}.$$

Since an active row has at least propagation weight 2, this means that $B(p) \geq 2B_{\text{rows}}(p)$. We can thus use the row branch number as a way to limit the search to parity patterns for which $\tilde{w}(b) \leq T'_2$.

For a given parity pattern, we classify the columns as either affected, unaffected odd or unaffected even. We make use of the following properties to find a lower bound on the row branch number.

Lemma 2. *In terms of active rows, θ satisfies the following properties:*

- An active bit in an affected column before θ will be passive after θ , and vice-versa. So, for each bit $(x, y, z) \xrightarrow{\pi \circ \rho} (x', y', z')$ of an affected column, at least one of row (y, z) in $\lambda^{-1}(b)$ and row (y', z') in b will be active.
- An odd unaffected column always contains at least one active bit and this bit stays active after θ . So, for at least one bit $(x, y, z) \xrightarrow{\pi \circ \rho} (x', y', z')$ of an odd unaffected column, both rows (y, z) in $\lambda^{-1}(b)$ and (y', z') in b will be active.

These properties are translated into Algorithm 1, which returns a lower bound of $B_{\text{rows}}(p)$. The algorithm avoids counting twice an active row by marking (in the sets a and b) the row positions already encountered.

6.4 Looking for Candidate Parity Patterns

To find trails such that any two consecutive rounds have weight up to $T'_2 = 28$, we have to consider the parity patterns listed in Lemma 3.

Lemma 3. *A 2-round differential trail $Q = (b_0, b_1, b_2)$ in KECCAK- f [1600] with $w(Q) \leq 28$ necessarily satisfies one of the following properties on the parity of $a_1 = \lambda^{-1}(b_1)$:*

- a_1 is in the kernel, i.e., $P(a_1) = 0$;
- the θ -gap of a_1 is 1 with a single run of length 1 or 2; or
- the θ -gap of a_1 is 2 or 3 with runs of length 1 each, all starting in the same slice.

Algorithm 1. Computing a lower bound of $B_{\text{rows}}(p)$

```

Let  $a$  and  $b$  be sets of row positions, which are initially empty
 $B \leftarrow 0$ 
for each affected column  $(x, z)$  do
  for  $y \in \mathbb{Z}_5$  do
    Let  $(x, y, z) \xrightarrow{\pi \circ \rho} (x', y', z')$ 
    if  $(y, z) \notin a$  and  $(y', z') \notin b$  then
       $B \leftarrow B + 1$ 
       $a \leftarrow a \cup \{(y, z)\}$  and  $b \leftarrow b \cup \{(y', z')\}$ 
    end if
  end for
end for
for each unaffected odd column  $(x, z)$  do
  Let  $(x, i, z) \xrightarrow{\pi \circ \rho} (x'_i, y'_i, z'_i)$  for  $i \in \mathbb{Z}_5$ 
  if  $\{(i, z), i \in \mathbb{Z}_5\} \cap a = \emptyset$  then
     $B \leftarrow B + 1$ 
     $a \leftarrow a \cup \{(i, z), i \in \mathbb{Z}_5\}$ 
  end if
  if  $\{(y'_i, z'_i), i \in \mathbb{Z}_5\} \cap b = \emptyset$  then
     $B \leftarrow B + 1$ 
     $b \leftarrow b \cup \{(y'_i, z'_i), i \in \mathbb{Z}_5\}$ 
  end if
end for
return  $B$ 

```

If parities are considered up to translation along z , we can restrict ourselves to parity patterns with runs starting in slice $z = 0$.

To prove this result, we conducted a recursive search as follows. Each parity is represented as a set of runs. First, all parity patterns p with a single run (so θ -gap 1) are investigated. All p with $B_{\text{rows}}(p) \leq \frac{T'_2}{2} = 14$ are stored into a set S . Then, we recursively add runs not overlapping the already added ones (so as to cover θ -gaps higher than 1), and all found p with $B_{\text{rows}}(p) \leq \frac{T'_2}{2} = 14$ are stored into a set S .

To limit the search, we use the following monotonicity property on the number of active rows. Using Lemma 2, changing an unaffected even column into either an unaffected odd or an affected column cannot decrease the number of active rows.

In the recursive search described above, adding a run to a parity pattern p can turn an unaffected odd column into an affected column. Hence, we cannot use the monotonicity property directly on the runs. However, adding a run never turns an affected column back into an unaffected one. So, before recursively adding a run to p , we apply a modified version of Algorithm 1 that does not take unaffected odd columns into account; this modified algorithm is monotonic in the runs. If the value returned by this modified algorithm is already above

$\frac{T'_2}{2} = 14$, then there is no need to further add runs. This efficiently cuts the search.

Before being added to the candidate set S , the parity pattern p is tested with the unmodified Algorithm 1. For the remaining parity patterns, we explicitly generated all states a with these parities up to $\tilde{w}(\lambda(a)) \leq T'_2 = 28$. This allowed us to prove Lemma 3.

Algorithm 1 is implemented in the `getLowerBoundTotalActiveRows` function and the recursive search in `lookForRunsBelowTargetWeight` [5].

6.5 Starting from Out-of-Kernel States

For a given parity pattern p , we can construct all states $b = \lambda(a)$ with $P(a) = p$ and $\tilde{w}(b) \leq T'_2 = 28$. We proceed in two phases.

- In a first phase, we generate all states a such that $P(a) = p$ by assigning all possible 16 values to affected (odd or even) columns and by assigning a single active bit in each unaffected odd column. These states are such that $\|a\| + \|\lambda(a)\|$ is exactly $10g + 2c$, with g the θ -gap and c the number of unaffected odd columns.
- In a second phase, we take the states generated in the first phase and add pairs of bits to all unaffected columns. By adding a pair of bits, we do not alter $P(a)$.

In both phases, we keep only the states $b = \lambda(a)$ for which $\tilde{w}(b) \leq T'_2 = 28$. As can be seen in Table 3, both the weight and the reverse minimum weight are *monotonic*, i.e., adding an active bit to the state cannot decrease them. We can therefore limit the search by stopping adding pairs of bits when $\tilde{w}(b)$ is above $T'_2 = 28$.

In practice, what we did was the following.

- Let \mathcal{P} be the set of parity patterns satisfying one of the conditions of Lemma 3 except $p = 0$.
- By the method described above, we construct all states in the set $\mathcal{B} = \{b : P(\lambda^{-1}(b)) \in \mathcal{P} \text{ and } \tilde{w}(b) \leq T'_2 = 28\}$.
- Finally, we forward and backward extend the states in \mathcal{B} to 3-round trail cores up to weight $T_3 = 36$.

We again found the same trail core as in Section 5. The trail prefix of weight 32 has $P(a_1) = 0$ (so a_1 is in the kernel) and $P(a_2)$ has one run of length 2 (so a_2 has θ -gap 1). No other trail cores were found.

When extending the states in \mathcal{B} , we exhaustively scan all compatible states, thereby including cases where $P(a_1) = 0$ or $P(a_2) = 0$. Hence, we covered the whole target space, except for trails such that both $P(a_1) = 0$ and $P(a_2) = 0$.

7 Generating In-Kernel Trails

To close the target space, we must look at in-kernel trails of the form in Eq. (2) with both $P(a_1) = 0$ and $P(a_2) = 0$. In the case of in-kernel trails, we were able to

be completely cover the space up to weight $T_3 = 40$, and we expect the techniques presented here can cover trails of higher weight. As $P(a_1) = P(a_2) = 0$, the θ operation has no effect and therefore $b_i = \pi(\rho(a_i))$. So this comes down to looking for states $a = a_1$, $b = b_1$, $c = a_2$ and $d = b_2$ connected as:

$$a \xrightarrow{\pi \circ \rho} b \xrightarrow{\chi} c \xrightarrow{\pi \circ \rho} d, \text{ with } P(a) = P(c) = 0. \quad (3)$$

We now summarize how we can efficiently generate all in-kernel three-round trail cores up to some weight and provide more details in following subsections. The key element in our method is the observation that any state b with $P(a) = 0$ and for which there exists a state c with $P(c) = 0$ can be represented in a specific way. The states a and b are iteratively constructed by adding active bits in the form of bit sequences called chains and vortices, defined in Section 7.2 below. Chains and vortices have an even number of active bits per column in a by construction and hence ensure $P(a) = 0$.

In b , there can be zero, one or more slices called knots, which contain three or more active bits. Each of these active bits is the end point of a chain that leads to another knot or that connects back to the same knot. The intermediate active bits of a chain appear pairwise in slices holding exactly two active bits in one column (called orbital slices, see Section 7.1). On top of chains connecting knots, a state b can exhibit a vortex, i.e., a cyclic sequence of active bits that appear pairwise both in the columns of a and in the columns of b .

By starting with an empty state and progressively adding chains, knots and vortices, one can quickly build states a and b that satisfy $P(a) = 0$ and for which there exist c with $P(c) = 0$, leading to 3-round in-kernel trail cores. Any state leading to a in-kernel trail can be represented in this way, and care is taken so that all possible states are generated, up to a given target weight. At each step, a lower bound on the weight of 3-round trail cores containing a and b is computed so as to efficiently limit the search.

As a final step, the generated states a and b are forward-extended to states c and d , limiting to c values in the kernel. Thanks to the properties of χ (see Section 3.1), the compatible states c can be expressed as a linear affine space. It is thereby easy to take the intersection of this affine space with the set of states such that $P(c) = 0$.

7.1 Characterizing the Slices in b

Definition 1. *A state b is tame if $P(\lambda^{-1}(b)) = 0$ and such that there exists at least one state c compatible with b through χ such that $P(c) = 0$.*

To characterize states b such that $P(c) = 0$, we can reason on the slices b_z of b since χ and P can be jointly described in terms of slices. In particular, each slice c_z of c must be in the kernel, namely, $P(c_z) = 0$, and we have to characterize the slices b_z under that constraint. First, if $b_z = 0$ then $c_z = 0$ and $P(c_z) = 0$. Then, a slice b_z with a single active bit cannot be in the kernel after χ , as at least one column of c_z will have a single active bit. Finally, a slice b_z with two

active bits must have its two active bits in the same column for c_z to be in the kernel. By inspection of Table 3, a row with a single active bit at coordinate x , e.g., 00100 transforms into an active row of the form $uv100$ with $u, v \in \{0, 1\}$, so the active bit stays active at x and zero, one or two active bits can appear at $x - 2$ and $x - 1$ of the same row. So, if the two bits are not in the same column, one of the active bits that stays after χ will not find another active bit in the same column. We summarize this in the next lemma.

Lemma 4. *If b is tame, then each of its slices has either*

- *no active bit,*
- *two active bits in the same column, or*
- *three or more active bits.*

We call an *empty slice* a slice with no active bit, and an *orbital slice* is a slice with two active bits in the same column. A slice that is neither empty nor an orbital slice is called a *knot*. We say that a knot is *tame* if it can transform after χ into a slice in the kernel. According to Lemma 4, a tame knot has at least three active bits.

7.2 Characterizing the Set of Active Bits

Since in the kernel θ acts as the identity, the active bits of a are just moved to other positions in b and their number remains the same, i.e., $||a|| = ||b||$. We can therefore represent a and b by a list of active bit positions $(p_i)_{i=1 \dots ||a||}$ in either the coordinates (x_i, y_i, z_i) in a or the coordinates (x'_i, y'_i, z'_i) in b , with $(x_i, y_i, z_i) \xrightarrow{\pi \circ \rho} (x'_i, y'_i, z'_i)$.

First, we start with the active bits in a . We say that active bits p_i and p_j are *peer* if they are in the same column in a , i.e., $x_i = x_j$ and $z_i = z_j$. Since each column has an even number of active bits when $P(a) = 0$, an active bit thus always has a peer.¹

Then, we move to the active bits in b . We say that the two active bits p_i and p_j are *chained* if they both lie in the same orbital slice in b . So $x'_i = x'_j$ and $z'_i = z'_j$ and no other active bit is in slice z'_i .

A *chain* is a sequence of bit positions of even length $(p_0, p_1, p_2, \dots, p_{2n-1})$ such that p_{2k} and p_{2k+1} are peer ($\forall k \in \{0, \dots, n-1\}$) and that p_{2k+1} and p_{2k+2} are chained ($\forall k \in \{0, \dots, n-2\}$). In addition, the first and last active bits p_0 and p_{2n-1} must be in knots (either the same one or different ones). The simplest possible chain has length 2 and consists only in two peer active bits. Figure 2 depicts the concept of chain.

¹ While for columns with two active bits, the peer relationship is unambiguous, in the case of columns with four active bits, we choose which pairs of active bits are peer. Thus we can see the representation of the states as being augmented with additional attributes specifying the peer relationship and there may be several ways to represent the same state. By generating states via this representation, the only risk is to generate more states than necessary.




Fig. 2. Schematic example of a chain. An active bit position is represented by a circle with its index. Two active bits connected by a plain line (resp. dashed line) are peer (resp. chained).

The definition of a *vortex* is the same as that of a chain $(p_0, p_1, p_2, \dots, p_{2n-1})$, except that the first and last active bits p_0 and p_{2n-1} must be chained. In other words, a vortex forms a cycle of bit positions linked alternatively by peer and chained relationships, all in orbital slices.

In a tame state, each active bit position has exactly one peer position. The active bit positions in knots are the end points of chains, while the active bits in orbital slices are chained and belong to chains or vortices. Therefore, any tame state can be represented as a set of vortices and chains connecting knots.

7.3 Generating All Tame States

To generate all tame states up to a target weight T_3 , we generate states a and b by representing them using the concepts of Sections 7.1 and 7.2. The generation builds (initially empty) states a and b by iterating the following nested loops:

- In the outer loop, we add chains to the existing state. When adding a chain $(p_0, p_1, p_2, \dots, p_{2n-1})$, the slices that receive the end points p_0 and p_{2n-1} must become knots if they are not already. If $n > 1$, the pairs of (chained) active bits (i_{2k+1}, i_{2k+2}) are added to empty slices, which become orbital slices. Active bits cannot be added to already constructed orbital slices, as it would contradict the definition of an orbital slice. Enough chains must be added such that each knot contains at least 3 active bits (see Lemma 4).
- For a fixed set of chains produced in the previous step, the inner loop iterates on the number and position of vortices. In a vortex, all active bits are chained, so they must be added to empty slices, which become orbital slices.

With the monotonic lower bound function defined in the next section, we add chains and vortices until this lower bound exceeds T_3 .

7.4 Lower-Bounding the Weight of In-Kernel Trails

We wish to determine a lower bound on the weight of 3-round in-kernel trail cores (b, d) , namely, on $w^{\text{rev}}(a) + w(b) + w(d)$ with $a = \lambda^{-1}(b)$, from a and b only, for use in our trail generation. Since only d is unknown, this implies finding a lower bound on $w(d)$. This can be done by first determining a lower bound on the Hamming weight $\|d\|$ and then bounding the weight of any state with given Hamming weight.

To determine a lower-bound on $\|d\|$, we work on each slice of b . If slice b_z has $u = \|b_z\|_{\text{row}}$ active rows, then the slice c_z has at least u active bits. In addition, $P(c_z) = 0$ implies that the number of active bits must be even, so $\|c_z\| \geq 2\lceil \frac{u}{2} \rceil$. Finally, we have $\|d\| = \|c\|$ so

$$\|d\| \geq 2 \sum_z \left\lceil \frac{\|b_z\|_{\text{row}}}{2} \right\rceil.$$

From Table 3, it is easy to verify the following lower bound:

$$w(d) \geq \hat{w}(\|d\|) \triangleq \left\lceil \frac{4\|d\|}{5} \right\rceil + [1 \text{ if } \|d\| = 1 \text{ or } 2 \pmod{5}].$$

Hence, we define the *lower weight* of b as

$$L(b) \triangleq w^{\text{rev}}(\lambda^{-1}(b)) + w(b) + \hat{w} \left(2 \sum_z \left\lceil \frac{\|b_z\|_{\text{row}}}{2} \right\rceil \right).$$

The lower weight yields a lower bound on the weight of 3-round in-kernel trail cores (b, d) regardless of d .

7.5 Limiting the Search by Lower-Bounding the Weight

At each level of the loop described in Section 7.3, the corresponding iteration is aborted, and elements are not further added, if we can be sure that the lower weight $L(b)$ will become larger than the target weight T_3 . Adding a chain to the state can potentially bring new knots and/or new orbital slices. Adding a vortex necessarily brings new orbital slices. Therefore, there is a limit in the number of knots and orbital slices that must be considered for the generation to be complete up to the target weight.

As a preliminary step, the minimum reverse weight satisfies the following inequality (see Table 3):

$$w^{\text{rev}}(a) \geq \hat{w}^{\text{rev}}(\|a\|) \triangleq \left\lceil \frac{3\|a\|}{5} \right\rceil.$$

We see from Lemma 4 that each tame knot contributes to at least 3 active bits in a and in b . Furthermore, the number of bits in each slice of a must be even

Table 4. Summary of all 3-round differential trail cores found in KECCAK- f [1600] up to weight 36, and up to weight 40 for in-kernel trails. The number indicates the number of cores with the same properties indicated in the other columns.

Number	$\tilde{w}(\cdot)$	$w^{\text{rev}}(a_1)$	$w(b_1)$	$w(b_2)$	$P(a_1)$	$P(a_2)$	Structure of a_1, b_1
1	32	4	4	24	kernel	θ -gap 1	
1	35	12	12	11	kernel	kernel	vortex of length 6
7	36	12	12	12	kernel	kernel	vortex of length 6
7	39	12	12	15	kernel	kernel	vortex of length 6
2	39	12	11	16	kernel	kernel	2 knots connected by 3 chains
41	40	12	12	16	kernel	kernel	vortex of length 6
4	40	12	12	16	kernel	kernel	2 knots connected by 3 chains

($P(a) = 0$), so $\|a\| \geq 2 \lceil \frac{3k}{2} \rceil$ and $w^{\text{rev}}(a) \geq \hat{w}^{\text{rev}}(\|a\|)$, with k the number of knots. In b , each tame knot has at least 3 active bits on at least 2 different active rows, hence contributing at least 5 to the weight, and so $w(b) \geq 5k$. Each active row in b contributes to at least one active bit in d so $\|d\| \geq 2k$ and $w(d) \geq \hat{w}(\|d\|)$.

For instance, $k = 5$ knots implies that $\|a\| \geq 16$ and $w^{\text{rev}}(a) \geq \hat{w}^{\text{rev}}(16) = 10$, that $w(b) \geq 25$ and that $\|d\| \geq 10$ and $w(d) \geq \hat{w}(10) = 8$, so a lower weight of at least 43. If $T_3 \leq 42$, looking for configurations with from 0 to 4 knots is therefore sufficient, not even counting the orbital slices that also compose chains.

We found cores of weight 35, 36, 39 and 40, as detailed in Table 4. For illustration purposes, examples of trail prefixes are shown in [7]. The search described in this section is implemented in the `TrailCore3Rounds` and `TrailCoreInKernel-AtC` classes [5].

8 Extension to Six-Round Trails

Table 4 summarizes all the 3-round cores found. These trail cores completely represent all the 3-round trails up to weight 36 (or 40 for in-kernel trails). They can be found in [5].

The second phase introduced in Section 4 consists in exhaustively extending forward and backward all the 3-round trail cores into 6-round trails cores. As no 6-round trail of weight up to 73 were found, we conclude that a 6-round differential trail in KECCAK- f [1600] has at least weight 74. In the specific case of in-kernel trails, no 6-round trail of weight up to 81 were found and we conclude that a 6-round in-kernel differential trail in KECCAK- f [1600] has at least weight 82.

For the 24 rounds of KECCAK- f [1600], a differential trail has at least weight 296, and an in-kernel trail has at least weight 328.

9 Conclusions

We studied and implemented the exhaustive generation of 3-round differential trails in the KECCAK- $f[1600]$ permutation, which allowed us to prove a lower bound on the weight of differential trails. The techniques developed in this paper exploit the properties of the mixing layer in its round function to provide better bounds than what a brute-force method could provide. Table 2 shows that there remains a gap between the best known trails and the lower bound beyond three rounds that calls for future work. Finally, the concepts introduced in this paper, such as chains, vortices, knots and parity runs, help read trails and understand them.

References

1. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: On the Indifferentiability of the Sponge Construction. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 181–197. Springer, Heidelberg (2008), <http://sponge.noekeon.org/>
2. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: Cryptographic sponge functions (January 2011), <http://sponge.noekeon.org/>
3. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: On alignment in KECCAK. In: ECRYPT II Hash Workshop 2011 (2011)
4. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: The KECCAK reference (January 2011), <http://keccak.noekeon.org/>
5. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: KECCAKTOOLS software (April 2012), <http://keccak.noekeon.org/>
6. 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)
7. Daemen, J., Van Assche, G.: Differential propagation analysis of KECCAK. Cryptology ePrint Archive, Report 2012/163 (2012), <http://eprint.iacr.org/>
8. Daemen, J., Peeters, M., Van Assche, G., Rijmen, V.: Nessie proposal: the block cipher Noekeon, Nessie submission (2000), <http://gro.noekeon.org/>
9. Daemen, J., Rijmen, V.: The design of Rijndael — AES, the advanced encryption standard. Springer (2002)
10. Daemen, J., Rijmen, V.: Plateau characteristics and AES. IET Information Security 1(1), 11–17 (2007)
11. Dinur, I., Dunkelman, O., Shamir, A.: New Attacks on Keccak-224 and Keccak-256. In: Canteaut, A. (ed.) FSE 2012. LNCS, vol. 7549, pp. 447–463. Springer, Heidelberg (2012)
12. Duc, A., Guo, J., Peyrin, T., Wei, L.: Unaligned rebound attack: Application to Keccak. In: Canteaut, A. (ed.) FSE 2012. LNCS, vol. 7549, pp. 407–426. Springer, Heidelberg (2012)
13. Gauravaram, P., Knudsen, L.R., Matusiewicz, K., Mendel, F., Rechberger, C., Schläffer, M., Thomsen, S.S.: Grøstl – a SHA-3 candidate. Submission to NIST (round 3) (2011)

14. Heilman, E.: Restoring the differential security of MD6. In: ECRYPT II Hash Workshop 2011 (2011)
15. Naya-Plasencia, M., Röck, A., Meier, W.: Practical Analysis of Reduced-Round KECCAK. In: Bernstein, D.J., Chatterjee, S. (eds.) INDOCRYPT 2011. LNCS, vol. 7107, pp. 236–254. Springer, Heidelberg (2011)
16. NIST, Announcing request for candidate algorithm nominations for a new cryptographic hash algorithm (SHA-3) family. Federal Register Notices 72(212), 62212–62220 (2007), <http://csrc.nist.gov/groups/ST/hash/index.html>
17. Rivest, R., Agre, B., Bailey, D.V., Cheng, S., Crutchfield, C., Dodis, Y., Fleming, K.E., Khan, A., Krishnamurthy, J., Lin, Y., Reyzin, L., Shen, E., Sukha, J., Sutherland, D., Tromer, E., Yin, Y.L.: The MD6 hash function – a proposal to NIST for SHA-3. Submission to NIST (2008), <http://groups.csail.mit.edu/cis/md6/>
18. Wu, H.: The hash function JH. Submission to NIST (round 3) (2011)

New Attacks on Keccak-224 and Keccak-256

Itai Dinur¹, Orr Dunkelman^{1,2}, and Adi Shamir¹

¹ Computer Science Department, The Weizmann Institute, Rehovot, Israel

² Computer Science Department, University of Haifa, Israel

Abstract. The Keccak hash function is one of the five finalists in NIST’s SHA-3 competition, and so far it showed remarkable resistance against practical collision finding attacks: After several years of cryptanalysis and a lot of effort, the largest number of Keccak rounds for which actual collisions were found was only 2. In this paper we develop improved collision finding techniques which enable us to double this number. More precisely, we can now find within a few minutes on a single PC actual collisions in standard Keccak-224 and Keccak-256, where the only modification is to reduce their number of rounds to 4. When we apply our techniques to 5-round Keccak, we can get in a few days excellent near collisions, where the Hamming distance is 5 in the case of Keccak-224 and 10 in the case of Keccak-256. Our new attack combines differential and algebraic techniques, and uses the fact that each round of Keccak is only a quadratic mapping in order to efficiently find pairs of messages which follow a high probability differential characteristic.

Keywords: Cryptanalysis, SHA-3, Keccak, collision, near-collision, practical attack.

1 Introduction

The Keccak hash function [4] uses the sponge construction [3] to map arbitrary long inputs into fixed length outputs, and is one of the five finalists of NIST’s SHA-3 competition. The Keccak versions submitted to the SHA-3 competition have an internal state size of $b = 1600$ bits, and an output size n of either 224, 256, 384 or 512 bits. The internal permutation of Keccak consists of 24 application of a non-linear round function, applied to the 1600-bit state. Previous papers on Keccak, such as [14], include analysis of Keccak versions with a reduced internal state size, or with different output sizes. However, in this paper, we concentrate on the standard Keccak versions submitted to the SHA-3 competition, and the only way in which we modify them is by reducing their number of rounds.

Previous results on Keccak’s internal permutation include zero-sum distinguishers presented in [1], and later improved in [5,6,9]. Although zero-sum distinguishers reach a significant number of rounds of Keccak’s internal permutation, they have very high complexities, and they seem unlikely to threaten the core security properties of Keccak (namely, collision resistance, preimage resistance and second-preimage resistance). Other results on Keccak’s internal permutation

include a differential analysis given in [10]. Using techniques adapted from the rebound attack [13], the authors construct differential characteristics which give distinguishers on up to 8 rounds of the permutation, with complexity of about 2^{491} . However, in their method it is not clear how to reach the starting state differences of these characteristics from valid initial states of Keccak’s internal permutation, since in sponge constructions a large portion of the initial state of the permutation is fixed and cannot be chosen by the cryptanalyst. Thus, although the results of [10] seem to be more closely related to the core security properties of Keccak than zero-sum distinguishers, they still do not lead to any attacks on the Keccak hash function itself.

Currently, there are very few results that analyze reduced-round variants of the full Keccak (rather than its building blocks): in [2], Bernstein described preimage attacks which extend up to 8 rounds of Keccak, but are only marginally faster than exhaustive search, and use a huge amount of memory. More recently, Naya-Plasencia, Röck and Meier presented practical attacks on Keccak-224 and Keccak-256 with a very small number of rounds [15]. These attacks include a preimage attack on 2 rounds, as well as collisions on 2 rounds and near-collisions on 3 rounds. In this paper, we extend these collision attacks on Keccak-224 and Keccak-256 by 2 additional rounds: we find actual collisions in 4 rounds and actual near-collisions in 5 rounds of Keccak-224 and Keccak-256, with Hamming distance 5 and 10, respectively.

The collisions and near-collisions of [15] were obtained using low Hamming weight differential characteristics, starting from the initial state of Keccak’s permutation. Such low Hamming weight characteristics are also the starting point of our new attacks, but we do not require the characteristics to start from the initial state of the permutation. Given a low Hamming weight starting state difference of a characteristic, we can easily extend it backwards by one round, and maintain its high probability (as done in [10]). However, due to the very fast diffusion of the inverse linear mapping used by Keccak’s permutation, the new starting state difference of the extended characteristic has a very high Hamming weight. We call this starting state difference a *target difference*, since our goal is to find message pairs which have this difference after one round of the Keccak permutation (after the fixed round, this difference will evolve according to the characteristic with high probability).¹ One of the main tools we develop in this paper is an algorithm that aims to achieve this goal, namely, to find message pairs which satisfy a given target difference after one Keccak permutation round. We call this algorithm a *target difference algorithm*, and it allows us to extend our initial characteristic by two additional rounds (as shown in Figure 1): we first extend the characteristic backwards by one round to obtain the target difference (while maintaining the characteristic’s high probability). Then, we use the target difference algorithm to link the characteristic to the initial state of

¹ We note that the target difference is not a valid initial difference of the permutation, which fixes many of the state bits to pre-defined values. As a result, the high probability characteristic cannot be used to extend the results of [15] by an additional round.

Keccak's permutation, through an additional round. We note that the final link, which efficiently bypasses Keccak's first Sbox layer, uses algebraic techniques rather than standard probabilistic techniques.

The target difference algorithm is related to several hash function cryptanalytic techniques that were developed in recent years. In particular, it is related to the work of Khovratovich, Biryukov and Nikolic [12], where, similarly to our algorithm, the authors use linear algebra to quickly satisfy many conditions of a differential characteristic. However, these techniques seem to work best on byte-oriented hash functions, whose internal structure can be described using a few sparse equations, which is not the case for Keccak. Our algorithm is also closely related to the work of Khovratovich [11] that exploits structures (which aggregate internal states of the hash function) in order to reduce the amortized complexity of collision attacks: the attacker first finds a truncated differential characteristic and searches for a few pairs of initial states that satisfy it. Then, using the structures and the initially found pairs, the attacker efficiently obtains many additional pairs that satisfy the truncated characteristic. However, in the case of Keccak, there are very few characteristics that can lead to a collision with high probability, and it seems unlikely that they can be joined in order to form the truncated differential characteristic required in order to organize the state differences into such structures. Moreover, it seems difficult to find even one pair of initial states that satisfy the target difference for Keccak. Another attack related to the target difference algorithm is the rebound attack [13]. In this attack, the cryptanalyst uses the available degrees of freedom to efficiently link and extend two truncated differential characteristics, both forwards and backwards, from an intermediate state of the hash function. However, once again, such high probability truncated characteristics are unlikely to exist for Keccak. Moreover, it is not clear how to use the rebound attack to link the backward characteristic to the initial state of the permutation. Thus, our target difference algorithm can be viewed as an asymmetric rebound attack, where one side of the characteristic is fixed.

Our full attacks have two parts, where in the first part we execute the target difference algorithm in order to obtain a sufficiently large set of message pairs that satisfy the target difference after the first round. In the second part of the attack, we try different message pairs in this set in order to find a pair whose difference evolves according to a characteristic whose starting state is the target difference. Since the target difference algorithm does not control the differences beyond the first round, the second part of the attack is a standard probabilistic differential attack (which only searches for collisions or near-collisions obtained from message pairs within a specific set). The high probability differential characteristic beyond the first round ensures that the time complexity of the second part of the attack is relatively low.

Although the target difference algorithm is heuristic, and there is no provable bound on its running time, it was successfully applied with its expected complexity to many target differences defined by the high probability differential characteristics. Consequently, we were able to find actual collisions for 4

rounds of Keccak-224 and Keccak-256 within minutes on a standard PC. By using good differential characteristics for an additional round, we found near-collisions for 5 rounds of Keccak-224 and Keccak-256. However, this required more computational effort (namely, a few days on a single PC), since the extended characteristics have lower probabilities.

The paper is organized as follows. In Section 2, we briefly describe Keccak, and in Section 3 we introduce our notations. In Section 4, we give a comprehensive overview of the target difference algorithm and describe the properties of Keccak that it exploits. In Section 5, we present our results on round-reduced Keccak. In the full version of the paper [8], we describe the full details of the target difference algorithm, and propose an alternative algorithm, which has a better understood time complexity. Since the original algorithm gave us very good results in practice, we did not use this alternative version. However, it may be more efficient in some cases, especially if someone finds longer high probability characteristics for Keccak’s permutation.




Fig. 1. Extending a 2-Round Differential Characteristic by Two Additional Rounds

2 Description of Keccak

In this section we give short descriptions of the sponge construction and the Keccak hash function. More details can be found in the Keccak specification [4].

The sponge construction [3] works on a state of b bits, which is split into two parts: the first part contains the first r bits of the state (called the outer part of the state) and the second part contains the last $c = b - r$ bits of the state (called the inner part of the state).

Given a message, it is first padded and cut into r -bit blocks, and the b state bits are initialized to zero. The sponge construction then processes the message in two phases: In the absorbing phase, the message blocks are processed iteratively by XORing each block into the first r bits of the current state, and then applying a fixed permutation on the value of the b -bit state. After processing all the blocks, the sponge construction switches to the squeezing phase. In this phase, n output bits are produced iteratively, where in each iteration the first r bits of the state are returned as output and the permutation is applied.

The Keccak hash function uses multi-rate padding: given a message, it first appends a single 1 bit. Then, it appends the minimum number of 0 bits followed by a single 1 bit, such that the length of the result is a multiple of r . Thus, multi-rate padding appends at least 2 bits and at most $r + 1$ bits.

The Keccak versions submitted to the SHA-3 competition have $b = 1600$ and $c = 2n$, where $n \in \{224, 256, 384, 512\}$. The 1600-bit state can be viewed as a 3-dimensional array of bits, $a[5][5][64]$, and each state bit is associated with 3 integer coordinates, $a[x][y][z]$, where x and y are taken modulo 5, and z is taken modulo 64.

The Keccak permutation consists of 24 rounds, which operate on the 1600 state bits. Each round of the permutation consists of five mappings $R = \iota \circ \chi \circ \pi \circ \rho \circ \theta$. Keccak uses the following naming conventions, which are helpful in describing these mappings:

- A row is a set of 5 bits with constant y and z coordinates, i.e. $a[*][y][z]$.
- A column is a set of 5 bits with constant x and z coordinates, i.e. $a[x][*][z]$.
- A lane is a set of 64 bits with constant x and y coordinates, i.e. $a[x][y][*]$.
- A slice is a set of 25 bits with a constant z coordinate, i.e. $a[*][*][z]$.

The five mappings are given below, for each x, y , and z (where the state addition operations are over $GF(2)$):

1. θ is a linear map, which adds to each bit in a column, the parity of two other columns.

$$\theta: a[x][y][z] \leftarrow a[x][y][z] + \sum_{y'=0}^4 a[x-1][y'][z] + \sum_{y'=0}^4 a[x+1][y'][z-1]$$

In this paper, we also use the inverse mapping, θ^{-1} , which is more complicated and provides much faster diffusion: for θ^{-1} , flipping the value of any input bit, flips the value of more than half of the output bits.

2. ρ rotates the bits within each lane by $T(x, y)$, which is a predefined constant for each lane.

$$\rho: a[x][y][z] \leftarrow a[x][y][z + T(x, y)]$$

3. π reorders the lanes.

$$\pi: a[x][y][z] \leftarrow a[x'][y'][z], \text{ where } \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 2 & 3 \end{pmatrix} \cdot \begin{pmatrix} x' \\ y' \end{pmatrix}$$

4. χ is the only non-linear mapping of Keccak, working on each of the 320 rows independently.

$$\chi: a[x][y][z] \leftarrow a[x][y][z] + ((-a[x+1][y][z]) \wedge a[x+2][y][z])$$

Since χ works on each row independently, it can be viewed as an Sbox layer which simultaneously applies the same 5 bits to 5 bits Sbox to the 320 rows of the state. We note that the Sbox function is an invertible mapping, and we will use the extremely important observation that the algebraic degree of each output bit of χ as a polynomial in the five input bits is only 2. As noted in [4], the algebraic degree of the inverse mapping χ^{-1} is 3.

5. ι adds a round constant to the state.

$$\iota: a \leftarrow a + RC[i_r]$$

We omit the values of $RC[i_r]$, as they are not needed for our analysis.

3 Notations

Given a message M , we denote its length in bits by $|M|$. Unless specified otherwise, in this paper we assume that $|M| = r - 8$, namely we consider only single-block messages of maximal length such that $|M| \pmod{8} \equiv 0$ (which give us the maximal number of degrees of freedom, for single-block messages containing an integral number of bytes). Given M , we denote the initial state of the Keccak permutation as the 1600-bit word $\overline{M} \triangleq M \parallel p \parallel 0^{2n}$, where \parallel denotes concatenation, and p denotes the 8-bit pad 10000001.

The first three operations of Keccak's round function are linear mappings, and we denote their composition by $L \triangleq \rho \circ \pi \circ \theta$. We denote the Keccak nonlinear function on 5-bit words defined by varying the first index by $\chi_{|5}$. The difference distribution table (*DDT*) of this function is a two-dimensional 32×32 integer table, where all the differences are assumed to be over $GF(2)$. The entry $DDT(\delta^{in}, \delta^{out})$ specifies the number of input pairs to this Sbox with difference δ^{in} that give the output difference δ^{out} (i.e., the size of the set $\{x \in \{0, 1\}^5 \mid \chi_{|5}(x) + \chi_{|5}(x + \delta^{in}) = \delta^{out}\}$).

We denote the 1600-bit target difference, which is the input of the target difference algorithm, by Δ_T . The output of the algorithm is a subset of ordered pairs of single block messages $\{(M_1^1, M_1^2), (M_2^1, M_2^2), \dots, (M_k^1, M_k^2)\}$ that satisfy this difference after one round R , namely $R(\overline{M}_i^1) + R(\overline{M}_i^2) = \Delta_T \forall i \in \{1, 2, \dots, k\}$.

4 Overview of the Target Difference Algorithm

When designing the target difference algorithm, we face two problems: first, the target difference extends backwards, beyond the first Keccak Sbox layer, with very low probability (due to its high Hamming weight). The second problem is that the initial state of the permutation fixes many of the state bits to pre-defined values, and the initial states that we use must satisfy these constraints. On the other hand, Keccak has several useful properties that we can exploit in our target difference algorithm. In this section, we describe these properties in detail and give an overview of the algorithm.

4.1 The Properties of Keccak Exploited by the Target Difference Algorithm

The First Property. Keccak-224 and Keccak-256 allow the user to control many of the 1600 state bits of the initial state of the permutation. Thus, given a target difference, we expect many solutions to exist (namely, one-block message pairs which have the 1600-bit target difference after one permutation round): since we consider message pairs, where each message is of length $r - 8 = 1600 - 8 - 2n$ bits (1144 for Keccak-224, and 1080 for Keccak-256), given an arbitrary 1600-bit target difference, there is an expected number of $2^{2(1600-8-2n)-1600} = 2^{1584-4n}$ message pairs of this length that satisfy this difference (regardless of

the value of the inner part of the state). Thus, the algorithm has 704 and 560 degrees of freedom for Keccak-224 and Keccak-256, respectively.

Despite the large number of available degrees of freedom, the number of possible solutions varies significantly according to the target difference. To demonstrate this, we use the fact that L^{-1} has very fast diffusion (i.e., even an input with one non-zero bit is mapped by L^{-1} into a roughly balanced output). We consider the case where $t > 0$ out of the 320 Sboxes of the target difference are active (i.e., they have a non-zero output difference). Each one of the $320 - t$ non-active Sbox zero output differences is uniquely mapped backwards to a zero input difference into the first Sbox layer. Using the Keccak Sbox *DDT*, it is easy to see that each one of the t active Sbox output differences is mapped to more than 8 possible input differences. Thus, the number of possible state differences after the first linear layer (or before the first Sbox layer) is more than $8^t = 2^{3t}$. Since L is invertible and acts deterministically on the differences, the number of possible input differences to the Keccak compression function remains the same. We now recall from the first difference constraint in Section 4, that we require that the $2n + 8$ MSBs of Δ_I are zero. However, for t large enough, we still expect more than $2^{3t-2n-8}$ valid solutions. When the target difference is chosen at random, we have $t \approx 310$ (since the probability that an Sbox output difference is zero is $\frac{1}{32}$). This gives more than $2^{930-448-8} = 2^{474}$ expected solutions for Keccak-224, and more than $2^{930-512-8} = 2^{410}$ expected solutions for Keccak-256. On the other hand, consider the extreme case of $t = 1$ (i.e., the target difference has only one active Sbox). Clearly, this Sbox cannot contribute more than 31 possible differences after the first linear layer. Since L^{-1} has very fast diffusion, these possible differences are mapped to at most 31 roughly balanced non-zero possible input differences, and we do not expect the $2n + 8$ MSBs of any of them to be zero. To conclude, target differences with a small number of active Sboxes are likely to have no solutions at all. On the other hand, a majority of the target differences have a very large number of expected solutions for Keccak-224 and Keccak-256. Note that having a large number of solutions does not imply that it is easy to find any one of them, since their density is still minuscule.

The Second Property. The algebraic degree of the Keccak Sboxes is only 2. This implies that given a 5-bit input difference δ^{in} and a 5-bit output difference δ^{out} , the set of values $\{v_1, v_2, \dots, v_l\}$ such that $\chi_{|5}(v_i) + \chi_{|5}(v_i + \delta^{in}) = \delta^{out}$ is an affine subset. Since $(v_i + \delta^{in}) + \delta^{in} = v_i$, then $v_i + \delta^{in} \in \{v_1, v_2, \dots, v_l\}$, implying $\{v_1, v_2, \dots, v_l\} = \{v_1 + \delta^{in}, v_2 + \delta^{in}, \dots, v_l + \delta^{in}\}$. Thus, both coordinates of the ordered pairs give the same subset, and we denote it by $A(\delta^{in}, \delta^{out})$ (note that $|A(\delta^{in}, \delta^{out})| = DDT(\delta^{in}, \delta^{out})$). On the other hand, since the algebraic degree of the inverse Sbox is 3, which is reduced to 2 (rather than 1) after differentiation, the output values that satisfy an input and an output difference do not necessarily form an affine subset.

The Third Property. For any non-zero 5-bit output difference δ^{out} to a Keccak Sbox, the set of possible input differences, $\{\delta^{in} | DDT(\delta^{in}, \delta^{out}) > 0\}$, contains at least 5 (and up to 17) 2-dimensional affine subspaces. These affine subspaces can

be easily pre-computed using the *DDT*, for each one of the 31 possible non-zero output differences. However, we note that there is no output difference for which the set of possible input differences contains an affine subspace of dimension 3 or higher.

4.2 Formulating the Problem

Given Δ_T , an arbitrary message pair (M^1, M^2) in which $|M^1| = |M^2| = r - 8$ is a solution to our problem if $R(\overline{M}^1) + R(\overline{M}^2) = \Delta_T$. This can be formulated using two constraints on the 1600-bit words $(\overline{M}_1, \overline{M}_2)$:

1. The $2n + 8$ MSBs of \overline{M}^1 and \overline{M}^2 are equal to $p||0^{2n}$, where p denotes the 8-bit pad 10000001.
2. $R(\overline{M}^1) + R(\overline{M}^2) = \Delta_T$ (where R is the permutation round of Keccak).

We can easily formulate the first constraint using linear equations on the bits of \overline{M}_1 and \overline{M}_2 . Since Keccak's Sbox has an algebraic degree of 2 over $GF(2)$, we can formulate the second constraint as a system of quadratic equations on these bits. Standard heuristic techniques for solving such systems include using the available degrees of freedom to fix some message values (or values before the first Sbox layer) in order to linearize the system. However, these techniques require many more than the available number of degrees of freedom. For example, in order to get linear equations after one round of Keccak's permutation, we can fix 3 out of the 5 bits entering an Sbox (after the first linear layer), such that there are no two consecutive unknown input bits entering the Sbox. Using this technique reduces the single quadratic term in the symbolic form of each of the Sbox's output bits to a linear term. However, this requires fixing $320 \cdot 3 = 960$ bits per message, and $2 \cdot 960 = 1920$ bits in total, which is significantly more than the 704 available degrees of freedom for Keccak-224 (and clearly more than the available number of degrees of freedom for the other Keccak versions). Consequently, we have to repeat the linearization procedure a huge number of times, with different fixed values, in order to find a solution.

A Two-Phase Algorithm. Although we expect our quadratic system to have many solutions, solving all the equations at once seems difficult. Thus, we split the problem into easier tasks by exploiting the low algebraic degree of Keccak's Sbox to a greater extent than in the standard techniques: as described in the second property of Section 4.1 given an input difference and an output difference to an Sbox, all the pairs of input values that satisfy them form an affine subset.² This suggests an algorithm with two phases, where in the first phase (called the *difference phase*) we find an input difference to all the Sboxes, and in the second

² Similar observations were used in [7] to suggest that when $DDT(\delta^{in}, \delta^{out}) = 2$ or 4, the same holds. In the specific case of Keccak, we also use 3-dimensional affine subsets of pairs that satisfy the Sbox difference transition $(\delta^{in}, \delta^{out})$, for which $DDT(\delta^{in}, \delta^{out}) = 8$.

phase (called the *value phase*) we obtain the actual values of the message pairs that lead to the target difference.

Using this two-phase approach, the ordered pairs produced by our algorithm satisfy two additional properties: the 1600-bit *input difference* of the initial states is fixed to some 1600-bit value Δ_I (i.e. $\overline{M}_i^1 + \overline{M}_i^2 = \Delta_I \forall i \in \{1, 2, \dots, k\}$), and the set composed of all the initial states defined by the first message in each ordered pair (i.e. $\bigcup\{\overline{M}_i^1\} \forall i \in \{1, 2, \dots, k\}$), forms an affine subset. The algorithm outputs the ordered pairs as the fixed 1600-bit input difference Δ_I , and some basis for the affine subset $\bigcup\{\overline{M}_i^1\} \forall i \in \{1, 2, \dots, k\}$. We note that the large number of degrees of freedom allows us to restrict the set of solutions (i.e. the set of message pairs that satisfy the target difference) to a smaller subset (but still large enough for our purposes) that can be found relatively easily. In particular, the algorithm considers only message pairs with a fixed difference Δ_I , for which all the solutions can be found by solving linear equations.

The two constraints above, which define our quadratic equation system, are broken into two sets of constraints, since we have to simultaneously enforce two *difference constraints* (given as constraints on the 1600-bit word Δ_I):

Difference Constraint 1. *The $2n + 8$ most significant bits (MSBs) of Δ_I are equal to zero.*

Difference Constraint 2. *$L(\Delta_I)$ is a valid input difference to the Sbox layer, i.e. there exists some 1600-bit word W such that $\chi(W) + \chi(W + L(\Delta_I)) = \Delta_T$ (note that since L is a linear function, $L(\Delta_I)$ is well-defined).*

The first difference constraint simply equates bits of the input difference Δ_I to zero (456 bits for Keccak-224 and 520 bits for Keccak-256), while the second difference constraint assigns to every 5 bits of $L(\Delta_I)$ that enter an Sbox, several possible values which are not related by simple affine equations.

In the second phase, we enforce additional *value constraints* (given on the 1600-bit word \overline{M}^1):

Value Constraint 1. *The $2n + 8$ MSBs of \overline{M}^1 are equal to $p||0^{2n}$, where p denotes the 8-bit pad 10000001.*

Value Constraint 2. *$R(\overline{M}^1) + R(\overline{M}^1 + \Delta_I) = \Delta_T$.*

Note that the first difference constraint and the first value constraint on each \overline{M}_i^1 also ensure that the same value constraint holds for \overline{M}_i^2 (i.e., the $2n + 8$ MSBs of \overline{M}_i^2 are equal to $p||0^{2n}$).

Given a single 1600-bit Sbox layer input difference, the second property of Section 4.1 implied that enforcing the two value constraints simply reduces to solving a union of two sets of linear equations. On the other hand, it is not clear how to simultaneously enforce both of the difference constraints, since given an output difference to an Sbox δ^{out} , all the possible input differences δ^{in} such that $DDT(\delta^{in}, \delta^{out}) > 0$, are not related by simple affine relations.

4.3 The Difference Phase

Unsuccessful Attempts to Enforce the Difference Constraints. We can try to enforce both difference constraints by assigning the undetermined $1600 - 2n - 8$ bits of Δ_I , in such a way that the second difference constraint will hold. This usually involves iteratively constructing an assignment for Δ_I , by guessing several undetermined bits at a time, and filtering the guesses by verifying the second difference constraint. However, this is likely to have a very large time complexity, since L diffuses the bits of Δ_I in a way that forces us to guess many bits before we can start filtering the guesses. Moreover, for any Δ_T , the fraction of input differences satisfying the first difference constraint that also satisfy the second difference constraint is very small. Thus, most of the computational effort turns out to be useless, since the guesses are likely to be discarded at later stages of the algorithm. Another approach is to guess $L(\Delta_I)$ by iteratively guessing the 5-bit Sbox input differences, and filtering the guesses by verifying the first difference constraint. For similar reasons, this approach is likely to have a very large time complexity.

A Better Approach. Both of these approaches are very strict, since each guess made by the algorithm commits to a specific value for some of the bits of Δ_I , or $L(\Delta_I)$, and restricts the solution space significantly. Thus, we use the third property of Section 4.1, which gives us more flexibility, and significantly reduces the time complexity: given any non-zero 5-bit output difference to a Keccak Sbox, the set of possible input differences contains at least five 2-dimensional affine subspaces. Consequently, in order to enforce the second difference constraint, for each Sbox with a non-zero output difference (i.e., an active Sbox), we choose one of the affine subsets (which contains 4 potential values for the 5 Sbox input bits of $L(\Delta_I)$), instead of choosing specific values for these bits. This enables us to maintain an affine subspace of potential values for $L(\Delta_I)$, starting with the full 1600-dimensional space, and iteratively reducing its dimension by adding affine equations in order to enforce the second difference constraint for each Sbox. In addition to these affine equations that we add per active Sbox, we also have to add the linear equations for the non-active Sboxes (which equate their 5 input difference bits to zero), and the additional $2n + 8$ linear equations that enforce the first difference constraint. All of these equations are added to a linear system of equations that we denote by E_Δ .

Since the $2n + 8$ equations that enforce the first difference constraint do not depend on the target difference, we add them to E_Δ before we iterate the Sboxes. While iterating over the active Sboxes, we add equations on $L(\Delta_I)$ in order to enforce the second difference constraint and hope that for each Sbox, we can add equations such that E_Δ is consistent. Note that the equations in E_Δ in each stage of the algorithm depend on the order in which we consider the active Sboxes, and on the order in which we consider the possible affine subsets of input differences for each Sbox. Thus, if we reach an Sbox for which we cannot add equations

in order to enforce the second constraint (while maintaining the consistency of E_Δ), it does not imply that it is impossible to satisfy the difference constraints. In this case, we can simply change the order in which we consider the active Sboxes, or the order in which we consider the affine subsets for each Sbox, and try again. Since we cannot predict in advance the orderings that give the best result, we choose them heuristically, as described in the full version of the paper [8].

4.4 The Value Phase

In case the difference phase procedure described above succeeds, it actually outputs an affine subspace of candidate input differences, rather than a single value for Δ_I . Next, we can commit to a specific value for Δ_I and run the value phase, hoping that the set of all linear equations defined by the value constraints has a solution. Namely, we allocate another system of equations, which we denote by E_M , and add the equations on \overline{M}^1 that enforce the first value constraint. We then add the additional linear equations that enforce the second value constraints for all the Sboxes, and output the solution to the system, if it exists. However, once again, this approach is too strict, and may force us to repeat the value phase a huge number of times with different values for Δ_I , until we find a solution. Thus, we do not choose a single value for Δ_I in advance. Instead, we reduce the linear subset of candidates for Δ_I gradually by fixing the input difference to each one of the active Sboxes, until a single value for Δ_I remains. Thus, we continue to maintain E_Δ throughout the value phase, and iteratively add the additional 2 equations which are required to uniquely specify a 5-bit input difference for each active Sbox, among the 2-dimensional affine subsets chosen in the difference phase. Once we fix the input difference to an Sbox, we immediately obtain linear equations on \overline{M}^1 , and we can check their consistency with the current equations in E_M . In case the equations in E_M are not consistent for a certain Sbox, we can try to choose another input difference for it. This gives different equations on \overline{M}^1 , which may be consistent and allow us to continue the process.

Similarly to the difference phase, the equations in E_M in each stage of the algorithm depend on the order in which we consider the active Sboxes, and on the order in which we consider the possible input differences for each Sbox. Thus, once again, if at some stage of the value phase we cannot add any consistent equations to E_M , we can change one of these orderings and try again, hoping to obtain a valid solution.

We stress again that both phases of the algorithm are not guaranteed to succeed. The success of each phase depends on the target difference, and on orderings which are chosen heuristically. As a result, we may have to iterate both phases of the algorithm an undetermined number of times with modified orderings, hoping to obtain better results.

5 Application of the Target Difference Algorithm to Round-Reduced Keccak

Since we would like to use the target difference algorithm in order to find collisions and near-collisions in Keccak, it is crucial to verify the algorithm's success on target differences which lead to these results. Thus, before we run the algorithm, we have to find such high probability differential characteristics, and to obtain the target differences which are likely to be the most successful inputs to the algorithm. As described in the introduction, once we find a high probability differential characteristic with a low Hamming weight starting state difference, we extend it backwards to obtain the target difference (while maintaining its high probability). We then use the target difference algorithm to link the extended characteristic backwards to the initial state of Keccak's permutation, with an additional round. Thus, any low Hamming weight characteristic for r rounds of Keccak's permutation can be used to obtain results on a round-reduced version of Keccak with $r + 2$ round. Specifically, in this section we demonstrate how we use 2-round characteristics in order to find collisions for 4 rounds of Keccak-224 and Keccak-256, and how to use 3-round characteristics in order to find near-collisions for 5 rounds of these Keccak versions.

5.1 Searching for Differential Characteristics

We reuse the notion of a *column parity kernel* or *CP-kernel* that was defined in the Keccak submission document [4]: a 1600-bit state difference is in the CP-kernel if all of its columns have even parity. It is easy to see that such state differences are fixed points of the function θ , which does not increase their Hamming weight. Since ρ and π just reorder the bits of the state, the application of L to a CP-kernel does not change its total Hamming weight. In addition, there is a high probability that such low Hamming weight differential states are fixed points of χ . Thus, when we start a differential characteristic from a low Hamming weight CP-kernel, we can extend it beyond the Sbox layer, χ , to one additional round of the Keccak permutation, with relatively high probability and without increasing its Hamming weight. However, extending such a characteristic to more rounds in a similar way is more challenging, since we have to ensure that the state difference before the application of θ remains in the CP-kernel at the beginning of each round.

Using Previous Results. In [10] and [15], the authors propose algorithms for constructing low Hamming weight differential characteristics for Keccak. Both of these algorithms successfully find differential characteristics that stay in the CP-kernel for 2 rounds (named *double kernel trails* in [15]), some of which lead to collisions on the n -bit extract taken from the final state after 2 rounds, with high probability. However, when trying to extend each one of these characteristics by another round, the state difference is no longer in the CP-kernel and thus its Hamming weight increases significantly (from less than 10 to a few dozen bits). Nevertheless, the Hamming weight of the characteristics is still relatively

low, and they can lead with reasonably high probability to near-collisions on the n output bits extracted. Beyond 3 rounds, the Hamming weight of the characteristics becomes very high (more than 100), and it seems unlikely that they can be extended to give collisions or near-collisions with reasonable probability. The currently known double kernel differential trails only extend forward to at most three rounds with reasonably high probability (higher than 2^{-100}). Finding new high probability differential characteristics, starting from a low Hamming weight state difference and extending forwards more than 3 rounds, remains a challenging task, which we do not deal with in this paper.

Our attacks on round-reduced Keccak make use of the type of differential characteristics that were found in [10] and [15], namely low Hamming weight characteristics that stay in the CP-kernel for 2 rounds. The double kernel trails with the highest probability have Hamming weight of 6 at the input to the initial round, and due to their low hamming weight, we could easily find all these characteristics within a minute on a standard PC. There are 571 such characteristics out of which, 128 can give collisions for Keccak-224 and 64 can give collisions for Keccak-256. However, when trying to extend the characteristics by an additional round, we were not able to find any characteristic that gives collisions for Keccak-224 (or Keccak-256) with reasonable probability. Thus, our best 3-round characteristics lead only to near-collisions, rather than collisions. The characteristics that give the near-collisions with the smallest difference Hamming weight for Keccak-224 and Keccak-256 are, again, double kernel trails with 6 non-zero input bits. The best 3-round characteristics for Keccak-224 lead to near-collisions with a difference Hamming weight of 5, and for Keccak-256, the best 3-round characteristics leads to a near-collision difference Hamming weight of 8. Examples of these characteristics are found in Appendix A.

Extending the Characteristics Backwards. Since the characteristics that we use start with a low Hamming weight state difference, we can extend them backwards by one round without reducing their probability significantly (as done in [10]): we take this low Hamming weight initial state difference, and choose a valid state difference input to the previous Sbox layer which could produce it. We then apply L^{-1} , and obtain a new initial state difference for the extended characteristic, which serves as a target difference for our new algorithm. Note that the target difference is not in the CP-kernel (otherwise, we would have found a low Hamming weight differential characteristic that stays in the CP-kernel for 3 rounds). Thus, when we apply L^{-1} to the state difference that enters the Sbox layer, we usually obtain a roughly balanced target difference, with only a few non-active Sboxes. This is significant to the success of the target difference algorithm, which strongly depends on the number of active Sboxes in the target difference.³ In case the target difference obtained from a characteristic has too

³ As demonstrated in Section 4.1, we expect a large number of non-active Sboxes to foil the target difference algorithm. This should be contrasted to differential attacks, where the attacker searches for differential characteristics with many non-active Sboxes, which ensure that the differential transitions occur with high probability.

many non-active Sboxes, we can try to select another target difference for the characteristic, by tweaking the state difference input to the second Sbox layer.

Assuming that the algorithm succeeds and we obtain a sufficiently large linear subspace of message pairs (such that it contains at least one pair whose difference evolve according to the characteristic), we can find collisions for 4 rounds and near-collisions for 5 rounds of Keccak-224 and Keccak-256. For example, given an extended characteristic which results in collisions for 3 round of Keccak-256 with probability 2^{-24} , we need a linear subspace containing at least 2^{24} message pairs in order to find a collision on 4-round Keccak-256 with high probability.

5.2 Applying the Target Difference Algorithm to the Selected Differential Characteristics

We tested our target difference algorithm using a standard PC, on dozens of double-kernel trails with Hamming weight of 6. For each one of them, after tweaking the state difference input to the second Sbox layer at most once, we could easily compute a target difference where all of the 320 Sboxes are active. We then ran the target difference algorithm on each one of these targets. For both Keccak-224 and Keccak-256, the target difference algorithm eventually succeeded: the basic procedure of the difference phase always succeeded within the first two attempts (after changing the order in which we considered the Sboxes), while the value phase was more problematic, and we had to iterate its basic procedure dozens to thousands of times in order to find a good ordering of the Sboxes and obtain results. For Keccak-224, the algorithm typically returned an affine subspace of message pairs with a dimension of about 100 within one minute. For Keccak-256, the dimension of the affine subspaces of message pairs returned was typically between 35 and 50, which is smaller compared to the typical result size for Keccak-224 (as expected since we have fewer degrees of freedom). In addition, unlike Keccak-224, for Keccak-256 we had to rerun the algorithm (starting from the difference phase) a few times, when the value phase did not seem to succeed for the choice of candidate input difference subset. Hence, the running time of the algorithm was typically longer – between 3 and 5 minutes, which is completely practical.

5.3 Obtaining Actual Collisions and Near-Collisions for Round-Reduced Keccak-224 and Keccak-256

Obtaining Collisions. After successfully running the target difference algorithm, we were able to find collisions for 4-round Keccak for each tested double-kernel trail with Hamming weight of 6 (which leads to a collision). Since the probability of each one of these differential characteristics is greater than 2^{-30} , the probability that a random pair which satisfies its corresponding target difference leads to a collision, is greater than 2^{-30} . Thus, we expect to find collisions quickly for both Keccak-224 and Keccak-256, once the target difference algorithm returns a set of more than 2^{30} message pairs. However, even though the subsets we used contained more than 2^{30} message pairs, we were not able to find

collisions within several of these subsets for Keccak-224, and for many of the subsets for Keccak-256. As a result, we had to rerun the target difference algorithm and obtain additional sets of message pairs, until a collision was found. Thus, the entire process of finding a collision typically takes about 2–3 minutes for Keccak-224, and 15–30 minutes for Keccak-256. The reason that there were no 4-round collisions within many of the message pair subsets, is the incomplete diffusion of the Keccak permutation within the first two rounds. Since our subsets of message pairs are relatively small (especially for Keccak-256), and the values of all the message pairs within a subset are closely related, some close relations between a small number of bits still hold before the Sbox layer of the second round (e.g., the value of a certain bit is always 0, or the XOR of two bits is always 1). Some of these relations make the desired difference transition into the second Sbox layer impossible, for all the message pairs within a subset. We note that we were still able to find collisions rather quickly, since it is easy to detect the cases where the difference transition within the second Sbox layer is impossible⁴(which allowed us to immediately rerun the target difference algorithm). In addition, when this difference transition is possible, we were always able to find collisions within the subset. Two concrete examples of colliding message pairs for Keccak-224 and Keccak-256 are given in Appendix B.

Obtaining Near-Collisions. In order to obtain near-collisions on 5-round Keccak-224 and Keccak-256, we again start by choosing suitable differential characteristics. Out of all the characteristics that we searched, we chose the differential characteristics described in Appendix A, which lead to near-collisions of minimal Hamming weight for the two versions of Keccak. The results of the target difference algorithm when applied to targets chosen according to these characteristics, were similar to the results described in Section 5.2. However, compared to the probability of the characteristics leading to a collision, the probability of these longer characteristics is lower: the probability of the characteristics are 2^{-57} and 2^{-59} for Keccak-224 and Keccak-256, respectively. Thus, obtaining message pairs whose differences propagate according to these characteristics, and lead to 5-round near-collisions, is more difficult than obtaining collisions for 4 rounds of Keccak-224 and Keccak-256. However, for each such main characteristic, there are several secondary characteristics which diverge from the main one in final two rounds and give similar results. Thus, the probabilities of finding near collisions with a small Hamming distance for 5 rounds of Keccak-224 and Keccak-256, are higher than the ones stated above. In addition, by using some simple message modification techniques within the subsets returned by the target difference algorithm, we were able to improve these probabilities further. Thus, for Keccak-224, we obtained near-collisions with a Hamming distance of

⁴ In order to detect that the difference transition within the second Sbox layer is impossible for all the pairs in our subset, we try several arbitrary pairs in the subset, and observe if at least one has the desired difference after two rounds. Since we only need to check one Sbox layer transition, we expect that if this transition is indeed possible, we will find a corresponding message pair very quickly. Otherwise, we have to find a different set of message pairs by running the difference phase again.

5, which is the same as the output Hamming distance of the main characteristic that we used. For Keccak-256, the main characteristic that we used has an output Hamming distance of 8, but we were only able to find message pairs which give a near-collision with a slightly higher Hamming distance of 10. All of these near-collisions were found within a few days on a standard PC. Examples of such near-collisions are given in Appendix B.

6 Conclusions and Future Work

In this paper, we presented practical collision and near-collision attacks on reduced-round variants of Keccak-224 and Keccak-256. Our attacks are based on a novel target difference algorithm, which is used to link high probability differential characteristics for the Keccak internal permutation to legal initial states of the hash function. Consequently, we were able to significantly improve the best known previous results on Keccak, by doubling (from 2 to 4) the number of rounds for which collisions can be found in a practical amount of time.

Our target difference algorithm is clearly limited by the number of available degrees of freedom, and it seems difficult to extend it to reach target differences spanning 2 or more rounds of the Keccak permutation. However, it seems very likely that the algorithm will be useful in the future if longer high probability differential characteristics are found for the Keccak permutation.

References

1. Aumasson, J.-P., Meier, W.: Zero-sum distinguishers for reduced Keccak-f and for the core functions of Luffa and Hamsi. NIST Mailing List (2009)
2. Bernstein, D.J.: Second preimages for 6 (7 (8??)) rounds of keccak? NIST mailing list (2010)
3. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: Sponge functions. Presented at the ECRYPT Hash Workshop (2007)
4. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: The Keccak SHA-3 submission. Submission to NIST (Round 3) (2011)
5. Boura, C., Canteaut, A.: Zero-Sum Distinguishers for Iterated Permutations and Application to KECCAK- f and Hamsi-256. In: Biryukov, A., Gong, G., Stinson, D.R. (eds.) SAC 2010. LNCS, vol. 6544, pp. 1–17. Springer, Heidelberg (2011)
6. Boura, C., Canteaut, A., Cannière, C.D.: Higher-order differential properties of keccak and luffa. Cryptology ePrint Archive, Report 2010/589 (2010), <http://eprint.iacr.org/>
7. Daemen, J., Rijmen, V.: Plateau Characteristics. IET Information Security 1(1), 11–17 (2007)
8. Dinur, I., Dunkelman, O., Shamir, A.: New attacks on Keccak-224 and Keccak-256. Cryptology ePrint Archive, Report 2011/624 (2011), <http://eprint.iacr.org/>
9. Duan, M., Lai, X.: Improved zero-sum distinguisher for full round Keccak-f permutation. Cryptology ePrint Archive, Report 2011/023 (2011)
10. Duc, A., Guo, J., Peyrin, T., Wei, L.: Unaligned rebound attack - application to keccak. Cryptology ePrint Archive, Report 2011/420 (2011)

11. Khovratovich, D.: Cryptanalysis of Hash Functions with Structures. In: Jacobson Jr., M.J., Rijmen, V., Safavi-Naini, R. (eds.) SAC 2009. LNCS, vol. 5867, pp. 108–125. Springer, Heidelberg (2009)
12. Khovratovich, D., Biryukov, A., Nikolic, I.: Speeding up Collision Search for Byte-Oriented Hash Functions. In: Fischlin, M. (ed.) CT-RSA 2009. LNCS, vol. 5473, pp. 164–181. Springer, Heidelberg (2009)
13. Mendel, F., Rechberger, C., Schläffer, M., Thomsen, S.S.: The Rebound Attack: Cryptanalysis of Reduced Whirlpool and Grøstl. In: Dunkelman, O. (ed.) FSE 2009. LNCS, vol. 5665, pp. 260–276. Springer, Heidelberg (2009)
14. Morawiecki, P., Srebrny, M.: A SAT-based preimage analysis of reduced KECCAK hash functions. Cryptology ePrint Archive, Report 2010/285 (2010)
15. Naya-Plasencia, M., Röck, A., Meier, W.: Practical Analysis of Reduced-Round KECCAK. In: Bernstein, D.J., Chatterjee, S. (eds.) INDOCRYPT 2011. LNCS, vol. 7107, pp. 236–254. Springer, Heidelberg (2011)

A Appendix: Differential Characteristics for Keccak

In this section, we give examples of 3-round differential characteristics, which lead to collisions on 4-round Keccak-224 and Keccak-256, and 4-round characteristics, which lead to near-collisions on 5-round Keccak-224 and Keccak-256.

The differential characteristics are given as a sequence of the starting state differences in each round. In all the presented characteristics, all the active Sboxes get an input difference with a Hamming weight of 1, and we assume that they produce the same differences as outputs (which occurs with probability 2^{-2}). In order to calculate the probability of the final transition, we only consider active Sboxes which effect the output bits (since we truncate the final state to obtain the hashed output). Each state difference is given as a matrix of 5×5 lanes of 64 bits, ordered from left to right, where each lane is given in hexadecimal using the little-endian format. The symbol '·' is used in order to denote a zero 4-bit difference value. For example, consider the second state difference in Characteristic 1: each of the first two lanes has a zero difference, and only the LSB of the third lane contains a non-zero difference.

B Appendix: Actual Collisions and Near-Collisions for Round-Reduced Keccak-224 and Keccak-256

We give several examples of collisions and near-collisions for Keccak-224 and Keccak-256. The padded messages and output values are given in blocks of 32-bits ordered from left to right, where each block is given in hexadecimal using the little-endian format.

```
|26978AF134CB835E|AF224C4D78366789|C4DAE35E2656F26B|357C4789AF3-6AF1|78D3526BC6A74C4D|
|26978AF134CB835E|AF224C4D78366789|C4DAE35E2656F26B|357C4789AF3-6AF1|78D3526BC6A74C4D|
|26978AF134CB835E|AF224C4D78366789|C4DAE35E2676F26B|357C4789AF3-6AF1|78D3526BC4A74C4D|
|26978AF134CB835E|AF224C4D78366789|C4DAE35E265EF26B|357C4789AF3-4AF1|78D3526BC6A74C4D|
|26978AF134CB835E|AF226C4D78366789|C4DAE35E2656F26B|35FC4789AF3-6AF1|78D3526BC6A74C4D|
```

```
|-----|-----|-----1|-----4-----|-----|
|-----|-----|-----|-----|-----|
|-----|-----|-----|-----|-----|
|-----|-----|-----1|-----4-----|-----8-----|
|-----|-----|-----1|-----|-----8-----|
```

```
|-----|-----|-----|--8-----|2-----|
|4-----|-----|-----|-----|2-----|
|-----|-----|-----|--8-----|-----|
|-----|-----|-----|-----|-----|
|4-----|-----|-----|-----|-----|
```

```
|-----|-----|-----|-----|-----|
|-----8-----|-----2-----|-----|-----|-----|
|-----|-----|-----1-----|-----|-----1-----|
|-----1-----|-----4-----|-----|-----|-----|
|-----|-----|-----|-----|-----|
```

The probability of each one of the first two transitions is 2^{-12} . The probability of the third transition is 1, since there are no active Sboxes which affect the output.

Characteristic 1: A 3-round characteristic leading to collisions on Keccak-224 and Keccak-256 with probability 2^{-24}

```
|BD135E2FA6BD1346|12D789A92F12D78F|D7E26BC344D7E224|E69AF134B5E69AD5|98BC4D6BF898BC58|
|BD135E2FA6BD1346|12D789A82F12D78F|D7E26BC344D7E264|E69AF134B5E69AD5|98BC4D6BF898BC58|
|BD135E2FA6BD1346|12D789AB2F12D78F|D7E26BC344D7E224|E69AF134B5E29AD5|98BC4D6BF898BC58|
|BD135E2FA6BD1346|12D789A92F12D78F|D7E26BC344D7E224|E69AF134B5E69AD5|98BC4D6BF898BC58|
|BD135E2FA6BD1346|12D789A92F12D78F|D7E26BC344D7E224|E29AF134B5E69AD5|98BC4D7BF898BC58|
```

```
|-----|-----1-----|-----|-----|-----4-----|
|-----|-----|-----|-----|-----|
|-----|-----1-----|-----8-----|-----|-----|
|-----|-----|-----8-----|-----|-----4-----|
|-----|-----|-----|-----|-----|
```

```
|-----|-----|-----4-----|-----|-----|
|-----|-----|-----|-----4-----|-----|
|-----2-----|-----|-----4-----|-----4-----|
|-----2-----|-----|-----4-----|-----4-----|
|-----|-----|-----|-----|-----|
```

```
|-----|-----|-----|-----8-----|-----|
|-----|-----|-----1-----|-----|-----|
|-----|-----|-----8-----|-----|-----|
|-----|-----|-----2-----|-----|-----|
|-----1-----|-----|-----|-----4-----|-----|
```

```
|-----|2-----|48-----4--2-----|4--12-----|8--82-----1|
|98-----|2--2-8-----4-----|4-----|1--8-----2--|
|-----4-----|2--1-----12-----|4--2--2-8-----|4-----|
|1-4-----2--1-----|-----8-----|2--8-----4-----|9--|
|2--1--4-----|48--|1-4--2--1--|-----8-----|
```

The characteristic leads to near-collisions with a Hamming distance of 5 for Keccak-224, and 8 for Keccak-256. The probability of each one of the first three transitions is 2^{-12} . The probability of the final transition is 2^{-21} for Keccak-224 and 2^{-23} for Keccak-256. The total probability is 2^{-57} for Keccak-224 and 2^{-59} for Keccak-256.

Characteristic 2: A 4-round characteristic leading to near-collisions on Keccak-224 and Keccak-256

```

M1=
C4F31C32 4C59AE6D 5D19F0F4 25C4E44B D8853032 8D5E12F2 BB6E6EE2 27C33B1E 6C091058 EB9002D5
3BA4A06F 4A0CC7F1 CCB55E51 8D0DD983 2B0A0843 9B21D3B0 53679075 526DDED2 48294844 6FF4ED2C
1ACE2C15 471C1DC7 D4098568 F1EBF639 EAF7B257 09FDAE87 688878E6 4875EB30 C9C32D80 3C9E6FCB
3C2ABCFA E6A4807B 2AB281B8 812332B3

```

```

M2=
A4D30EF7 80BB8F69 90C048DF EB7213B9 A6650424 3A65F63E 8C268881 B651B81F AADAF3C EE2CA5C3
DB78EAC2 C8EAE779 442F9C35 3221E287 B3017A5A 90790712 1B1C8BDC E08B10A8 9A9D25CA 1BE7AAAC
4E2F3E9C 73717DAD 5566015A A198CFB9 5A1CA8C2 A0E3348A AE6C0BB1 3980F9E4 A4FA8B91 6E81A989
89A9BCAA E12BF1F1 30EF9595 812E8B45

```

```

Output=
61FB1891 F326B6D5 24DD94DF 73274984 05DA9A1D 3FD359B9 78B8393B F2E7990B

```

The messages were found using the target difference algorithm on the target difference given by Characteristic 1.

Collision 1: A collision for 4-round Keccak-256

```

M1=
FAF7AC69 2710BE04 8462C382 7ABF1BF9 D065CD30 DB64DEB8 1410CD30 C837D79B 22E446B7 31E9BD55
A6B2074C C86E32CC DE50A10A F7BAAA58 D96CBC88 9FBD75F6 5E0D735A D22AA663 16A574AA 7DB08692
558AB029 109B4D30 86CE5DCA 13A295C7 E7C9D94B 648794D2 62EE3CF8 69439337 8CAB9F15 AC7C3267
90F41CBE A20E6893 B4781F24 0BA37647 F29A67A0 81F628D0

```

```

M2=
CE5FBC81 47710FCC 462C92E0 48F5D2CF F92F6EC3 053E64E1 570780B9 F838EC54 8F74809F 66B4AC6F
70DD1843 BF34F0C5 5010C89A D8791148 D5CC073E 3239AEBC 7DF48D79 0EC7767B FB081604 AFA975B9
F8EFAE0F ED793473 479E931C F2F80A74 7192D08F 5EB5AB27 F1CAC04E F394232D 48656B2A A3471644
DB74E60A 05FB3B18 41DC27C3 8384BF53 32534C3E 811C00B5

```

```

Output=
826B10DC 0670E4E1 5B510CDA AB876AA8 B50557ED 267932FB AA4D38E8

```

The messages were found using the target difference algorithm on the target difference given by Characteristic 1.

Collision 2: A collision for 4-round Keccak-224

M1=
 23296F07 44536A2B 16E1E363 09B509F9 639CA324 2B834133 61457E6D 9CF07597 6797B3D4 D1715ABA
 6D8F4F9F 70D12920 E014BB37 54C32ADE 6117B3FB 30114566 4BA7D70A 00F055F0 71CFDD4 B53F2563
 E223A16D CC8DDAC4 7A59836B A53FBDDE 9FFEC45F 6A3476DC 7349BB92 56AF6E92 83866932 56624032
 A936E410 60AC00FA 7E7C61F9 81583CAC

M2=
 49D48DE2 9FA843CA 747C88E0 55425134 098CA5B3 C97DC68A B82BC6FD 0F864996 26B13425 D9F73B75
 932CD02F FB12E036 47706100 9DEFFFE4 79435F9C DA727EFO D9CA67C6 520BE2D1 19CF3933 3136D1A9
 EEBEA9DD 150CA247 D494BF4A 492EFB26 11CB4C8D F5A10A05 69128FF4 B142742F CA59FE32 4FE68436
 068F76AB 041A673E 461575B5 81AA2A54

Output1=
 407D4466 FEAE8231 EC968181 DF902165 23C219FF 54571D70 2800F506 E818644B

Output2=
 407D4466 FEAE8231 EC928181 FF902165 23C019FF 1C571D74 2800F516 E810656B

The messages were found using the target difference algorithm on the target difference given by Characteristic 2.

Near-Collision 1: A near collision with Hamming distance of 10 for 5-round Keccak-256

M1=
 7DBC1AA9 62A70E2A C2BDAF81 4A4D484B 672F6FAF ED312C83 24BC1974 16946039 6B46EDF6 1AE571A0
 EDA59D6E 7561766D 8F0B4C57 3C05C569 715B7DF9 53F81761 F6D43507 6E040495 9B5C08AB 5130BA66
 22AF7F5C 755840F2 2E893F59 4C4A730F 8C4F425D 182F8D00 E98515ED E29251AD 853AB863 DC46A7AC
 9FB7BB08 14767EFC 5345C7AF AA774E81 8A01A570 81D65453

M2=
 5659C936 AF3BA787 809C1CE6 B287F81B E0A5E769 ECCEB8A0 72506F44 1A1B2A02 EE9AE408 D16A9358
 BF03C4D6 90845C46 0C0441CC 8203EA8D 6D122EB1 9193F64F 55C3A6A7 47377ED6 D26E806F DEC2CBF8
 A3B8949E A91B248D 420B13BC BEAB4166 EE348CF6 DB6CCD82 122F6BDA 2FBFA7EA 75E8A429 F397BC46
 7E9DE824 6A973A22 371FD02D 92035083 267D1C7A 812EDE70

Output1=
 85373497 97D871C2 FBD0A823 584C0ED4 C1B3BF4F BC408766 0584B08D

Output2=
 85373497 97D871C2 FBD0A823 784C0ED4 E1B1BF5F BC408776 0584B08D

The messages were found using the target difference algorithm on the target difference given by Characteristic 2.

Near-Collision 2: A near collision with Hamming distance of 5 for 5-round Keccak-224

Author Index

- Bogdanov, Andrey 29
- Carlet, Claude 366
- Chen, Jiazhe 90
- Courtois, Nicolas T. 306
- Daemen, Joan 422
- Dakhilalian, Mohammad 385
- De Cannière, Christophe 287
- Dinur, Itai 9, 442
- Dong, Le 127
- Duc, Alexandre 402
- Dunkelman, Orr 9, 442
- Feng, Dengguo 127
- Fleischmann, Ewan 196
- Forler, Christian 196
- Goubin, Louis 366
- Gu, Dawu 90
- Guo, Jian 127, 402
- Heyse, Stefan 346
- Isobe, Takanori 264
- Jean, Jérémy 110
- Khovratovich, Dmitry 244
- Kiltz, Eike 346
- Li, Ji 264
- Li, Leibo 90
- Li, Wei 90
- Ling, San 163
- Liu, Ya 90
- Liu, Zhiqiang 90
- Lu, Jiqiang 69
- Lucks, Stefan 196
- Lyubashevsky, Vadim 346
- Mala, Hamid 385
- Mendel, Florian 226
- Mouha, Nicky 287
- Nad, Tomislav 226
- Naya-Plasencia, María 110, 146
- Nguyen, Long Hoang 326
- Nyberg, Kaisa 1
- Paar, Christof 346
- Peyrin, Thomas 110, 146, 163, 402
- Pieprzyk, Josef 163
- Pietrzak, Krzysztof 346
- Preneel, Bart 49, 287
- Prouff, Emmanuel 366
- Quisquater, Michael 366
- Rechberger, Christian 244
- Rivain, Matthieu 366
- Rogaway, Phillip 180
- Roscoe, A.W. 326
- Saarinen, Markku-Juhani Olavi 216
- Sajadieh, Mahdi 385
- Savelieva, Alexandra 244
- Schläffer, Martin 226
- Sepehrdad, Pouyan 306, 385
- Shamir, Adi 9, 442
- Shibutani, Kyoji 264
- Sokołowski, Przemysław 163
- Sun, Yue 49
- Sušil, Petr 306
- Tischhauser, Elmar 49
- Van Assche, Gilles 422
- Vaudenay, Serge 306
- Velichkov, Vesselin 287
- Wang, Huaxiong 163
- Wang, Meiqin 29, 49
- Wang, Xiaoyun 90
- Wei, Lei 163, 402
- Wooding, Mark 180
- Wu, Shuang 127
- Wu, Wenling 127
- Zhang, Haibin 180
- Zou, Jian 127