

A study of practical-time distinguishing attacks against round-reduced Threefish-256

Aron Gohr
Bundesamt für Sicherheit in der Informationstechnik
Bonn, Germany
aron.gohr@bsi.bund.de

February 16, 2012

Abstract

The present paper describes a practical-time distinguishing attack against up to 27 rounds of Threefish-256, the core cryptographic primitive of the Skein-256-256 hash function. Our main attack is a round-reduced variant of theoretical-time related-key boomerang attacks first studied in [1]. The attack distinguishes 27 rounds of Threefish-256 from an ideal cipher with a predicted workload of $\approx 2^{42}$ related-key decryption or encryption queries to round-reduced Threefish. We validate our complexity estimate by fully implementing the first 26 and the last 23 rounds of the attack, obtaining empirically verified distinguishers with workloads of $\approx 2^{29}$ and ≈ 2000 evaluations of the related-key rectangle test used on these rounds respectively.

We also present a similar attack in the single key, related tweak setting which distinguishes 19 rounds of Threefish-256 with a very practical effort of $\approx 2^{26.3}$ evaluations of a related-tweak rectangle attack.

1 Introduction

1.1 Background, motivation and focus of the present work

This paper is concerned with cryptanalytic properties of Threefish, the block cipher which forms the cryptographic core of the Skein hash function. Skein is one of the five candidates remaining in the final round of the SHA-3-competition. The development of a new advanced hash standard through public competition has become important due to advances in the cryptanalysis of hash functions related to the current standard, the SHA-2 family of hash functions. Most notable here are the practical breaking of collision resistance of MD5 [23] and the existence of collision attacks against SHA-1 which are believed to be executable with very large but still in principle practical computing effort [7]. While no attacks on the full SHA-2 are currently known, cryptanalysis of SHA-2 is progressing [6, 13, 14, 20], and having an efficient, standardized hash function not closely related to the SHA family with thoroughly reviewed security properties certainly still seems very desirable.

As the Skein construction is supported by security proofs if Threefish is replaced by an ideal cipher [3], the Threefish block cipher is a natural target in the evaluation of the Skein hash function. While security weaknesses in full Threefish would not necessarily immediately translate into attacks on classical hash function security properties, any discovery of such weaknesses would cast doubt upon the applicability of the security proofs. Also, an evaluation of cryptanalytic properties of Threefish is interesting in itself insofar as Threefish would, if Skein were chosen as SHA-3, certainly become a block cipher with relatively wide deployment. Finally, Threefish is a conceptually relatively simple ARX construction, so in advancing attacks against Threefish, one may hope to gain knowledge with general relevance to the cryptanalysis of ARX ciphers.

This paper focuses, in particular, on attacks which can be executed in practical time. We choose, somewhat arbitrarily, a workload equivalent to 2^{45} evaluations of the underlying primitive as the threshold of what we will consider a feasible attack. As all attacks we will study in this paper are based on related-key or related-tweak boomerang techniques and therefore require the querying of encryption and decryption oracles

besides computations which the adversary can perform by himself, this bound appears not unreasonably low to us. Our focus on practical attacks has three motivations. Firstly, for practical attacks, it is possible to compare precisely the attack complexities as given by theoretical models and actual measurements. Secondly, in quantifying resistance against practical attacks against Threefish, one gains some understanding of which security parameters for Skein will be needed at a minimum to provide protection against potentially implementable attacks. Thirdly, attack optimizations found for low-complexity attacks may prove useful also for improving high-complexity attacks.

1.2 Structure of this paper

In this section, we will provide an overview of the present work, a short review of related literature, point out our main contributions, and fix notations and conventions for the rest of the article. In the next section, we briefly recall the structure of Threefish. Section three of this article then presents an unknown-related-key distinguisher against 27 rounds of Threefish-256 with an estimated workload of 2^{42} encryption or decryption operations under four related-key/related-tweak pairs. This distinguisher is a reduced-round variant of a line of boomerang distinguishers against Threefish which was first introduced in [1] and further studied in [8, 17]. Also rebound-type attacks on the compression function of Skein in [21, 24] should be mentioned in this context.

Our main improvement over the literature on Threefish to obtain an unknown-key distinguisher with practical time requirements is a truncated differential from round four to the input on the decryption direction of the boomerang distinguisher which to the best of our knowledge has not been described before and which allows us to pass the backwards direction of the boomerang with very low cost. In section four we apply the same attack concept to the fixed-key setting, deriving a distinguisher against 19 rounds of Threefish-256 with workload $2^{26.3}$ boomerang evaluations which uses related-tweak differentials. To the best of our knowledge, this represents the first cryptanalysis of Threefish in a fixed-key (but still related-tweak) setting.

1.3 Short overview of related work

Pre-existing results on Threefish Related-key Boomerang attacks against Threefish were first studied in [10]. Similar attacks based on modular differentials (which are also used in the present paper) were first introduced in [8]. An attack based on similar principles, using bitwise differential paths and techniques from [15] for calculation of optimal differential transitions and their transition probabilities, was proposed in [17] against 31 rounds of Threefish-256 at an estimated cost of 2^{234} encryptions and decryptions. An attack of this kind against 33 rounds of Threefish-256 (with the old set of rotation constants) was also briefly mentioned in [1].

To the best of our knowledge, the best practical-time attack on any version of Threefish which has so far appeared in the literature was a related-key attack from [1] based on using a short differential followed by a local collision to cross 12 initial rounds which induces a practically detectable bias still after 21 rounds of the cipher. As far as we are aware, no low-complexity variants of the boomerang attacks from [1, 8, 17] have yet appeared.

The possibility to use truncated differentials to speed up boomerang attacks on Threefish was briefly mentioned in section 6.3 of [1]. However, no details were provided.

Bitwise biases induced by differentials were studied for Threefish in [1, 10]. We remark that the findings reported there make it seem plausible that truncated differentials such as those which we use should exist. Other notable attacks on Skein and Threefish have included rotational attacks [11, 12] which use the observation that rotational identities are destroyed relatively slowly by ARX constructions without large asymmetric constants. A tweak of the key schedule constant in Threefish has largely rendered these attacks obsolete [10].

Summary of main contributions of the present paper The present paper presents practical-time distinguishing attacks against 27 rounds of Threefish-256 in the related-key setting and 19 rounds in the related-tweak setting. To the best of our knowledge, these are the best practical-time attacks on Threefish published so far. We achieve the low complexity of our attacks by using very high probability truncated differentials on the first four rounds of Threefish-256 in the decryption step of the boomerang which do not

seem to have been described before. Also, we take advantage of a very efficient matching phase in the four middle rounds of the boomerang which allows us to cross four rounds in the middle of the cipher at no detectable cost. We used well-known techniques for calculating differential properties of bitwise addition [16] and of rotation [9] with respect to modular differences to model the modular differential behaviour of the Mix transform, the basic building block of Threefish, and present in this paper comparisons of the theoretical predictions of path complexity with actual attack cost measurements both on single rounds and short differential paths. In general, we find both to be in very good agreement, but there are a few exceptions.

1.4 Conventions and notations

Notations for tweakable block ciphers A tweakable block cipher is a family $E_{K,T}$ of bijective maps $\mathcal{P} \rightarrow \mathcal{C}$ where \mathcal{P} denotes the set of plaintext blocks, \mathcal{C} denotes the set of ciphertext blocks, \mathcal{K} denotes a set of secret keys and \mathcal{T} the set of public tweak values.

Round-reduced versions of Threefish The Threefish block cipher contains a subkey addition only once in every four rounds of the cipher. Therefore, if we talk about an n -round reduced version of Threefish, we always assume that a final subkey addition has been added in order to prevent potential attackers from simply inverting the last rounds of the cipher up to the last normally scheduled subkey addition. The final subkey addition that we add is always the next subkey addition that would take place in the full version of the cipher as described in the specification [10].

We denote the round transformations by R_0, R_1, \dots, R_n . In other words, a ten-round reduced version of Threefish will consist in the application of the transformations R_0, \dots, R_9 followed by a final subkey addition which is the one regularly occurring in R_{12} .

Basic operations and data representation conventions We will denote modular addition of 64-bit numbers by \boxplus and bitwise addition by \oplus . A *block* will uniformly mean a vector of four 64-bit words and all bitwise operations on words will transfer component-wise to the block level. Concrete word values will always be represented in big endian notation as unsigned hexadecimal numbers. As usual, \lll_n denotes left rotation of a 64-bit word by n bits and analogously \ggg_n the corresponding right rotation.

Differential path conventions For related-key and related-Tweak differences, we allow bitwise differences only. In messages and state values, we will mostly consider \boxplus -differential paths, as these cross subkey additions in Threefish at zero cost.

We will sometimes need to specify truncated differential states. These will always be bitwise differentials where only the states of some bits of the state are defined. As the defined bits will in the cases considered here always be the least significant bits of a 64 bit word, we fix the following notation: for a truncated differential state the hexadecimal number x is identified with the set of 64 bit numbers which are equal to x modulo $2^{\lceil \log_2(x) \rceil}$. In other words, x describes the set of words which are identical to x except for the positions corresponding to the leading zeroes of x . We will always give x in hexadecimal notation in this situation and drop leading zeroes.

A differential path for a tweakable block cipher E with transition probability p , message differences ΔM , key differences ΔK , tweak differences ΔT and output differences ΔC will be denoted by

$$E : (\Delta M, \Delta K, \Delta T) \xrightarrow{p} \Delta C.$$

If the differential does not depend upon differences in the key, tweak or message, we will suppress these inputs. A typical case where this occurs is a modular differential through a few keyless rounds of Threefish with possibly a subkey addition at the end.

In this article, a differential will usually be a differential with respect to wordwise modular differences, i.e. with respect to differences modulo 2^{64} . However, for differences in the key and tweak, we will allow only bitwise differences as we consider this a more realistic attack model. We remind the reader briefly of the fact that in the context of related key attacks, it is important to choose carefully which related-key operations should be allowed, as it is easy to create related-key attackers who are too powerful, i.e. who can obtain fast key recovery even against ideal ciphers [2]. However, our attacker, who can ask for encryptions and

decryptions of chosen plaintexts and ciphertexts under keys related to some secret key by arbitrary bitwise key and tweak differences, is to the best of our knowledge unable to distinguish an ideal block cipher with unknown key from a random permutation of the same block size with expected effort less than exhaustive key search.

For truncated differentials, we will use the same notations. In this case, ΔC and/or ΔM will be truncated states as defined above.

2 A short overview of Threefish-256

Threefish is a tweakable block cipher with supported block sizes of 256, 512, and 1024 bits. Key sizes and block sizes of all Threefish versions are equal to the block size. As an additional parameter, a 128-bit public tweak value is input to the key schedule along with the key.

In this article, we will focus entirely on the 256-bit version. Therefore, we will in the sequel only describe Threefish-256. Analogous attacks, however, are expected to be feasible at comparable costs also against the other block sizes.

We only recall the structure of Threefish very briefly. For a more detailed exposition, we refer to the Skein submission paper [10].

Basic structure Threefish-256 processes a 256-bit message by evolving an internal state of 256 words over 72 iterations of a simple round transformation. The round transformation consists of two layers, namely a non-linear Mix-layer applying a transform $Mix_r : \{0, 1\}^{128} \rightarrow \{0, 1\}^{128}$ to pairs of 64-bit words in the four-word state and a linear layer consisting of a word level permutation which exchanges words S_1 and S_3 of the state, where the state is (S_0, S_1, S_2, S_3) and the S_i are 64-bit integers.

At the beginning of every fourth round and after the final round of the cipher, a subkey addition is performed. This updates the state by component-wise modular addition of a subkey for current keying s to the state.

The Mix transform Mix_r is a bijective transform with input and output two 64-bit words. Its definition is

$$Mix_r(x, y) := (x \boxplus y, (x \boxplus y) \oplus (y \lll r)).$$

The Mix-layer of a Threefish round consists of the parallel application of Mix transformations to the two 128-bit halves of the state. The rotation constant r is chosen differently for each of the two Mix transformations involved in a round and depends on the round number. Precise specifications can be found in [10].

The key schedule The Threefish key schedule takes as input a 256-bit key $K = (K_0, K_1, K_2, K_3)$ and a 128-bit Tweak $T = (T_0, T_1)$. The subkeys $k_{s,i}$ for keying s are then determined as follows: first, both K and T are expanded by one 64-bit word by setting $T_2 := T_0 \oplus T_1$ and $K_4 := C \oplus (\oplus_i K_i)$, where C is a constant defined in [10]. One furthermore sets recursively $K_{i+5} := K_i, T_{i+3} := T_i$. The subkey k_s for keying s is then given by the following expression:

$$k_s := (K_s, K_{s+1} \boxplus T_s, K_{s+2} \boxplus T_{s+1}, K_{s+3} \boxplus s).$$

We remark that for related-key differentials with key difference ΔK and tweak difference ΔT with differences only appearing in the most significant bit of key and tweak words, this implies subkey differences

$$\Delta k_s = (\Delta K_s, \Delta K_{s+1} \oplus \Delta T_s, \Delta K_{s+2} \oplus \Delta T_{s+1}, \Delta K_{s+3})$$

where $\Delta K_4 := \oplus_{i=0}^3 \Delta K_i$.

For compact notation of key, tweak and subkey differences, we will write δ to denote the 64 bit word corresponding to 2^{63} .

3 Practical-time distinguishers against round-reduced Threefish

3.1 Basic attack structure

Qualitative overview Basically, the distinguishers we will study are reduced-round variants of boomerang attacks against Threefish first introduced in [1] and further studied in [8, 17]. In the forward direction of the boomerang, we use a related-key differential path very similar to the ones used already in [1, 8, 17] to induce a local collision in round four. This local collision then survives until round 12 due to a subkey collision in round eight. Between rounds 12 and 16, matching between the forward and backwards phases of the boomerang happens. For cost computations regarding this attack phase, summation over various very short high-probability bitwise differential paths can be used. Empirically, this matching phase succeeds with probability one. Due to the use of multiple differential paths in the boomerang matching phase and near the decryption output (through the use of a truncated differential) our attacks are best characterised as rectangle attacks as introduced in [5].

In the decryption direction, the same method of using a local collision is used to produce a high-probability modular differential path from round 27 to round 16. As in the forward direction, a short bitwise differential path can be used to obtain upper bounds for the cost of successfully passing the matching phase of the boomerang. A truncated differential on rounds four to zero helps us avoid having to pay the cost of the initial four-round differential again in the return path of the boomerang.

The related tweak attacks we will study share the same structure but span a lower number of rounds.

Generic cost calculation Denote by $E_{K,T}$ the encryption function of round-reduced Threefish under key K and tweak T . Denote by D the corresponding decryption function. For a generic cost model, assume that we have a decomposition

$$E = E^\gamma \circ E^\beta \circ E^\alpha$$

and assume that we have related-key differential paths π_α

$$E^\alpha : (\Delta M, \Delta K, \Delta T) \xrightarrow{p} \Delta C$$

and π_γ

$$D^\gamma : (\Delta C', \Delta K', \Delta T') \xrightarrow{q} \Delta M'$$

as well as a truncated differential π_t

$$D^\alpha : \Delta C \xrightarrow{p_t} \Delta M_t$$

where ΔM_t is a truncated differential state defined in r bits. In addition, we assume that the probability of obtaining a right quartet for a boomerang attack on E^β with plaintext difference ΔC and ciphertext difference $\Delta M'$ is p_b . Finally, we assume availability of an *initial path boosting transform* F which given a right pair P_1, P_2 for π_α will produce a list of some number n of pairs $F_i(P_1, P_2) \in \mathcal{P} \times \mathcal{P}$ such that the probability (averaged over all lists generated for all eligible P_1, P_2) of any such $F_i(P_1, P_2)$ being a satisfying instance of π_α is $p' > p$. A natural candidate for an initial path boosting transform may be a differential covering a few initial rounds of E^α with some high probability p_h , as such a differential can be used to create from a conforming pair P_1, P_2 for π_α another pair P'_1, P'_2 such that in the first few rounds the differences generated by P'_1, P'_2 will satisfy π_α with probability approximately p_h^2 which may be higher than the probability of satisfying these rounds of π_α randomly.

Assuming all these elements, the basic attack pattern runs like this: we randomly generate a plaintext P_1 and compute from it a second plaintext $P_2 := P_1 \boxplus \Delta M$. One then obtains $C_1 := E_{K,T}(P_1)$ and $C_2 := E_{K \oplus \Delta K, T \oplus \Delta T}(P_2)$. From these values one obtains $C_3 := C_1 \boxplus \Delta C'$, $C_4 := C_2 \boxplus \Delta C'$ and asks for decryptions

$$P_3 := D_{K,T}(C_3), P_4 := D_{K \oplus \Delta K \oplus \Delta K', T \oplus \Delta T \oplus \Delta T'}(C_4).$$

We are interested in the case that $P_4 \boxplus P_3$ satisfies ΔM_t and call P_1, P_2, P_3, P_4 a *right quartet candidate* in this case. For a random function, this would happen with probability 2^{-r} . For E , however, one notes that, due to reasoning exactly as in [22, 5], one gets additional right quartet candidates in all those cases where the following conditions were met:

1. P_1, P_2 was a right pair for π_α .
2. C_1, C_3 and C_2, C_4 were right pairs for π_γ .
3. The boomerang test on E^β induced by the partial encryptions of P_1, P_2 and the partial decryptions of C_3, C_4 in the input and output of E^β respectively succeeded.
4. The partial decryptions of C_3, C_4 at the input of D^α yielded a conforming pair for π_t .

Under an assumption of independence of these events, the probability that all of them will occur in any given trial is expected to be $pq^2p_b p_t$. If this is much larger than 2^{-r} , most candidate right quartets will indeed be right quartets for the related-key rectangle test and we obtain a distinguisher against E immediately. Otherwise, we may have to use the initial path boosting transform to see if it gives us another example of a candidate right quartet. If P_1, P_2, P_3, P_4 was not a true right quartet, this is expected to happen with probability 2^{-r} , whereas if it was, a new right quartet is expected to be found with probability $p_h^2 q^2 p_b p_t$ per trial. In the attacks considered in this article, this latter probability is much larger than 2^{-r} , which means that a relatively small number of verification tests using initial path boosting will distinguish true right quartets for our test from false right quartets with high reliability by finding at least one other right quartet candidate where for a random permutation none would be expected to be found. For a random function, the work factor of finding a P_1 that survives the tests just described is expected to be close to $k/2^{2r}$, where k is assumed small compared to 2^r and is the number of verification trials. For E the corresponding probability is $\approx (1 - (1 - p_h^2 q^2 p_b p_t)^k) pq^2 p_b p_t$ or $\approx (1 - 1/e) pq^2 p_b p_t$ when $k \approx \frac{1}{p_h^2 q^2 p_b p_t}$. As long as this k is small compared to 2^r and as long as $(1 - 1/e) pq^2 p_b p_t \gg k/2^{2r}$, a distinguisher against E with workload approximately $\frac{1}{(1-1/e)pq^2 p_b p_t}$ boomerang test executions is obtained. In this situation, the workload of the distinguisher can be improved towards approximately $w_{lim} := \frac{1}{pq^2 p_b p_t}$ by increasing the number of verification trials while staying below a small fraction of 2^r verifications. How close to w_{lim} one can actually get depends on the difference in magnitude between 2^r and $\frac{1}{pq^2 p_b p_t}$.

In the case of our related-key distinguisher against 27 rounds of Threefish-256, we will have empirically determined values of $p \approx 2^{-29}, q \approx 2^{-5.6}, p_b \approx 1, p_t \approx 0.98, p_h = 1$ and $r = 25$. Hence, we will obtain a distinguisher at a cost of approximately $2^{40.2}$ boomerang evaluations.

For our related-tweak distinguisher against 19 rounds of Threefish-256, we get $r = 49$ and $p \approx 2^{-15.1}, q \approx 2^{-5.6}, p_b \approx 1, p_t \approx 0.999$. Hence, a distinguisher with total complexity of about $2^{26.3}$ boomerang evaluations is obtained.

Calculations of path complexities In the following sections, we will give full details on the two attacks which are the subject of this paper. Together with the differential paths, we will give both theoretical predictions of round-by-round transition probabilities for the paths we consider and empirical measurements of path probabilities. The theoretical predictions were obtained by using the methods from [16] to obtain precise calculations of the correct passage of the modular differences appearing in our paths through the bitwise additions in the Mix layers of Threefish and the methods from [9] to do the same for the rotations. Correct transition of differences through rotations and bitwise additions within a Mix transformation were viewed as independent events. While this assumption is certainly not precisely correct and can significantly differ from observation in some cases (see e.g. [19] and for a more general discussion of exact computation of such probabilities also [18]), comparisons of round-by-round predictions and round-by-round measurements show it to be accurate within the margin of error of practical experimentation in the cases we are dealing with, with few exceptions.

3.2 A related-key boomerang distinguisher against 27 rounds of Threefish-256

We will now explain the details of the 27-round distinguisher.

Details for E^α E^α corresponds here to the first 12 rounds of Threefish-256. We use here a related-key differential very similar to that used in [17], except that we use \boxplus -differences for differential trails as [8] do, which helps us pass all subkey additions at zero cost. In the encryption direction of the boomerang, we use the same subkey differential as [17]: we set $\Delta K = (0, 0, 0, \delta)$ and $\Delta T = (\delta, 0)$. For the message and state, we

Round	ΔM				$\log_2(\mathbb{P})$
0	faff6ff5afdef7c0:	100100210210800:	ffbf79fbbdfc:	40000084004204	-20.63
1	fbff7ff7bffffc0:	800040000040:	ffffbfffdfc0000:	40002000000	-7.95
2	fbffff800000000:	800000000:	ffffffffffc0000:	40000	-3,21
3	fc00000000000000:	400000000000000:	0:	0	-1
4	0:	0:	0:	8000000000000000	-

Table 1: Differential path from R_0 to R_4 (\boxplus -differences) for Threefish-256. The differential does not take subkey additions into account and needs to be corrected according to the subkey differential induced by subkey differences $(0, 0, 0, \delta)$ and tweak difference $(\delta, 0)$. Exactly as in [1], a state collision at round 4 is then obtained.

use a four round \boxplus -differential path given in table 1. Theoretically, our model predicts a path probability of about $2^{-32.79}$, whereas empirically we have measured a transition probability of approximately $2^{-29.0}$ over $2 \cdot 10^{11}$ trials. This is however not due to a failure of the model used to predict the transition probabilities round by round, as table 2 shows. We therefore ascribe the failure of the model to accurately predict the

Round r	Number of trials	Right pairs	$\log_2(\mathbb{P}_r)$ empirically	$\log_2(\mathbb{P}_r)$ predicted
Rotations	10^6	467913	-1.098	-1.098
0	10^8	61	-20.65	-20.63
1	10^6	3999	-7.97	-7.95
2	10^6	110954	-3.17	-3,21
3	10^6	499713	-1.00	-1

Table 2: Empirically determined round-by-round transition probabilities, number of trials used in their measurement, and absolute number of right pairs detected for the differential path from table 1. The first line reports theoretically predicted and empirically observed probability of the entire path where everything except the rotations and the key schedule has been linearized over $\mathbb{Z}/(2^{64})$.

likelihood of fulfilling the entire initial four round path to dependencies between passing different rounds which are not modelled.

After round 4, a state collision is induced if the initial four-round path was followed. This state collision survives until the input to round 12. Therefore, a related-key differential path is obtained for the first twelve rounds of Threefish which has exactly the same transition probability as the initial four round path.

Details for $E^\gamma - D^\gamma$ covers, in this attack, the decryption direction of the boomerang from rounds 16 to 27. The key and tweak differences chosen for this direction are $(0, \delta, \delta, 0)$ and $(\delta, 0)$. In the subkey addition of round 24, this induces a subkey difference of $(\delta, 0, 0, 0)$. Rounds 27 to 24 are covered by a three-round differential path reproduced in table 3. Our theoretical model predicts a transition probability of $2^{-9.9}$ for

```

Input      : 8000000000000000:          0:          0:          0
Round 24: 8000000000000000:          0:          0:8000000000000000
Round 25: 8000000000000000:8100000000000000:8000000000000000:8000000000000000
Round 26: 1000000000000000:      8000000000:          0: 1000000000408000

```

Table 3: \boxplus -differential path starting with an input difference in round 24 and ending with an output difference at round 26.

this path, which means that standard modelling of boomerang attacks would predict a contribution of $2^{19.8}$ to right quartet finding cost from this path for the 27-round boomerang. Empirically, however, we find that a boomerang attack covering rounds R_4 to R_{26} of Threefish-256 based on this path has a transition probability of $\approx 2^{-11.3}$, based on finding 406 right quartets in 10^6 trials. The corresponding attack on R_4 to R_{25} has an empirically determined probability of success of $\approx \frac{1}{4}$ per trial, based on one million trials finding 249549 right pairs in one test of ours. In this case, we suspect a failure of an independence assumption for differential transitions occurring in the rotation and bitwise addition of one of the Mix transformations involved.

Details for E^β E^β covers four keyless rounds of Threefish-256 from the output of the subkey addition in round 12 to the input of the subkey addition in round 16 by a small rectangle attack. The input and output differences to this miniature rectangle are easily derived from the subkey differences used on E^α and E^γ and turn out to be $(\delta, 0, 0, 0)$ and $(0, 0, 0, \delta)$. Empirically, the rectangle on E^β is found to have success probability one. It seems likely that it is possible to prove this success probability analytically but we have not tried to do so.

The truncated differential on D^α The truncated differential on D^α is derived from a truncated differential on the four initial rounds starting from a difference at the input of R_3^{-1} of $(0, 0, 0, \delta)$ and ending with the truncated state

$$S_{t1} := (40, 800, 4, 4)$$

in the decryption output. Empirically, we find this truncated differential to have a success probability of around 98 percent.

The initial path boosting transform Basically, what our distinguisher is meant to detect is successful passing of the initial four round differential path. This is an event with probability 2^{-29} which we detect by seeing the truncated state S_{t1} with probability $0.98 \cdot 2^{-11.3}$ if it occurs. On the other hand, the truncated state S_{t1} is defined in 25 bits, so we expect to encounter it with a per-trial probability of 2^{-25} also if the initial path was not passed successfully. Therefore, an initial path boosting transform is needed to distinguish cases where the truncated state is seen randomly from cases where it was caused by having chosen a right pair for the initial path.

We suggest two ways to construct an initial path boosting transform in the situation at hand. The first is based on the observation in [1] that tweak values are publicly known and that therefore a related tweak attacker can compensate changes in the tweak by changes in the message and thereby obtain for any given plaintext P , key K and tweak T 2^{128} new triples (P', K, T') such that encryption under key K with tweak T' of message P' will give exactly the same states in the first four keyless rounds of Threefish-256 encryption as does encryption under key K with tweak T of message P . As was already pointed out in [1], this effect immediately yields a way to construct from one conforming pair for the initial differential path 2^{128} other conforming pairs under related keys. In other words, we get a very efficient initial path boosting transform, albeit at the cost of a drastic increase of the number of related-tweak queries necessary. In our context, ≈ 4000 new related tweaks will have to be queried in order to make the 27-round distinguisher work at the complexity claimed in this paper.

The second strategy we suggest is to use neutral bits [4]. The following gives a simple method and underlying reasoning to obtain a strong enough boosting of the initial path for our purposes. We note that in the path given in table 1, the first two rounds are by far the most costly. It seems reasonable to assume that one can find many high-probability differential paths for these first two rounds. Addition of the input difference of any such path with success probability p to a conforming pair P_1, P_2 for the path is expected to yield P'_1, P'_2 which will produce correct differences in the first two rounds of the initial path with probability around p^2 or likely higher if there are additional contributions from other high-probability end-states of the path. A simple reasonable idea for creating many suitable paths is to consider paths which have differences only in a few high-order bits of the state words.

We have empirically tested this path boosting strategy and found that trying all 2^{16} differences which are restricted to the top four bits of each word regularly yields thousands of new conforming pairs, easily enough to reliably detect true right quartets in the rectangle attack. Therefore, this approach provides a way to make our distinguisher work without requiring queries under a large number of related tweaks.

Empirical attack verification We have practically implemented the first 26 rounds of our attack as well as the last 23 rounds. Measurements of attack complexity from these implementations have been in good agreement with predictions. In practical tests of the 26-round attack, we have seen three distinguishing events in $5 \cdot 10^9$ trials. This is well in line with the predicted workload of 2^{31} rectangle tests for a detection of non-random behaviour on average.

3.3 A related-tweak boomerang distinguisher against 19 rounds of Threefish-256

We now give a detailed account of our 19-round related-tweak distinguisher against Threefish-256. The basic attack pattern is the same as before, but due to the restriction of being able to query encryptions and decryptions under one key only, the attacker loses some freedom in creating collisions: he can still create state collisions, but not subkey collisions, as every subkey depends injectively on the tweak. As a consequence, state collisions can at most survive four rounds now. On the other hand, the need to choose message, key and tweak in such a way as to create a subkey collision and have it survive another keying did greatly limit the attacker’s choice of initial four round differential in the preceding attack. We can therefore choose more efficient prepended and appended differentials.

Details for E^α In this attack, E^α covers the first eight rounds of the cipher. We use a tweak difference $\Delta T = (0, \delta)$. This induces a subkey difference of $(0, \delta, \delta, 0)$ in the round four subkey addition. Accordingly, we report in table 4 on a \boxplus -differential covering the first four rounds of Threefish-256 without key additions which leads to a state difference of $(0, \delta, \delta, 0)$ at the input of round four with high probability. To correct for the effect of the subkey addition in round zero, a factor of $(0, 0, \delta, 0)$ needs to be added to the given input difference.

The given theoretical estimates of round transition probabilities have been derived in the same way as

Runde	ΔM				$\log_2(\mathbb{P})$
0	feffff7fd0000:	100000000210000:	7fbf7ffbf77bf40:	40800000084080	-12.09
1	fffff800000000:	800000000:	7fffffbf7ffc0:	40000040	-4.17
2	0:	0:	7fffff800000:	800000	-1
3	0:	0:	8000000000000000:	0	0
4	0:	8000000000000000:	8000000000000000:	0	-

Table 4: Differential path from R_0 to R_4 (\boxplus -differences) for related-tweak attacks on Threefish-256. The given path does not yet account for the effects of subkey additions. For each round, the difference in the input to that round is displayed.

those which we used in the preceding study of related-key attacks on Threefish. The observed probabilities for single round transitions are in very good agreement with these predictions. For the whole path, the predicted transition probability of $2^{-17.26}$ is somewhat lower than the observed one of $\approx 2^{-15.1}$ which was observed in 10^7 trials.

Details for E^β E^β has in this attack exactly the same structure as in the previous attack but covers the keyless parts of rounds 8 to 11 here. As in the previous attack, we observe a matching probability of one. Input and output differences in this sub-rectangle are $(0, \delta, 0, 0)$ and $(0, \delta, 0, 0)$ respectively.

Details for E^γ For the decryption phase of the boomerang, we can take a tweak difference of $(\delta, 0)$. Table 5 gives a differential from the input to R_{20} to the output of the subkey addition in R_{16} which produces a difference of $(0, 0, \delta, 0)$ after this subkey addition, hence matching the subkey difference induced by the tweak and inducing a collision in the keyless parts of rounds 15 to 12. Upon applying during decryption the subkey addition of round 12, appropriate differences between the encryptions of P_1 and P_2 and the decryptions of C_3 and C_4 for the sub-boomerang in E^β are induced.

The empirically observed transition probabilities for this path differ significantly from theoretical predictions for rounds 17 and 18, where instead of the given values transition probabilities of respectively $\approx \frac{1}{2}$ and $2^{-5.17}$ have been observed in one million trials each. For a boomerang attack covering 19 rounds, the empirically observed contribution to success probability of this appended differential path is approximately $2^{-11.2}$.

The truncated differential on D^α We use a similar truncated differential in this attack to cover the last four rounds of the decryption direction of the boomerang as in the previous attack, but as also the forward differential path in this attack is of much higher probability, we get a truncated differential defined in many

Runde	ΔM				$\log_2(\mathbb{P})$
16	0:	0:	8000000000000000:	0	0
17	0:	8000000000000000:	8000000000000000:	0	-2.58
18	8000000000000000:	8000000000000000:	8000000000000000:	8000000000000000:	-8.34
19	0:	8008008000000:	8000000000000:	400000	-14.30
E^{-1}	8008008000000:	808000000400000:	8000000400000:	108108108000000	-

Table 5: Appended \boxplus -differential for related tweak attacks on up to 20 rounds of Threefish-256 using related tweak attacks. Corrections due to final subkey additions are not yet considered.

more bits than before. The differential in question is

$$(0, \delta, \delta, 0) \rightarrow (10000, 10000, 40, 80)$$

and is defined on 49 bits. Empirically, we see a probability of ≈ 99.9 percent for this truncated path to occur. Therefore, we detect an event which has a generic probability of $\approx 2^{-49}$ in this distinguisher and as far as the 19-round distinguisher is concerned do so with a likelihood of $2^{-26.3}$ per boomerang evaluation, which means that no retesting mechanism is required. As each boomerang evaluation requires two encryption queries and two decryption queries, the total expected number of queries required for the distinguisher is about $2^{28.3}$. This is well supported by tests carried out on a full implementation of this distinguisher.

4 Conclusions

The present paper reports on attempts to see how many rounds variants of the best known theoretical distinguishers against Threefish-256 can cover if we require practical time and memory complexity. The main result of this study is that up to 27 rounds of Threefish-256 have suboptimal cryptographic properties under related-key attack which can be demonstrated in practice and up to at least 19 rounds allow attacks of low complexity which use a single key and four tweak values. Also, we find that a straightforward model of the Mix transform can be used to derive reasonably good predictions on the probability of \boxplus -differential transitions of the Threefish round function. Finally, some of the optimisations used in our distinguishers may be of some limited use in improving theoretical distinguishing attacks on round-reduced versions of Threefish.

While in the present work we focused entirely on Threefish-256, it is absolutely natural to expect that similar attacks will, with similar or possibly slightly lower complexity, apply as well to similar numbers of rounds of Threefish variants with larger state.

That being said, the attacks described in this paper do not transfer to attacks against the Skein hash function when the underlying Threefish version is reduced to 27 or 19 rounds respectively.

Finally, I thank Ernst Schulte-Geers for useful discussion during the completion of the work reported in this paper.

References

- [1] J.P. Aumasson, C. Calik, W. Meier, O. Özen, R. Phan, K. Varici, *Improved Cryptanalysis of Skein*, Advances in Cryptology - Asiacrypt 2009, Lecture Notes in Computer Science 5912, Springer Verlag 2009, 542-559
- [2] M. Bellare, T. Kohno, *A Theoretical Treatment of Related-Key Attacks: RKA-PRPs, RKA-PRFs, and Applications*, Advances in Cryptology - Eurocrypt 2003, Lecture Notes in Computer Science 2656, Springer Verlag 2003, 491-506
- [3] M. Bellare, T. Kohno, S. Lucks, N. Ferguson, B. Schneier, D. Whiting, J. Callas, J. Walker, *Provable Security Support for the Skein Hash Function Family*, 2009, <http://www.skein-hash.info/sites/default/files/skein-proofs.pdf>

- [4] E. Biham, R. Chen, *Near-Collisions of SHA-0*, Advances in Cryptology - Crypto 2004, Lecture Notes in Computer Science 3152, Springer Verlag 2004, 199-214
- [5] E. Biham, O. Dunkelman, N. Keller, *The Rectangle Attack - Rectangling the Serpent*, Advances in Cryptology - Eurocrypt 2001, Lecture Notes in Computer Science 2045, Springer Verlag 2001, 340-357
- [6] A. Biryukov, M. Lamberger, F. Mendel, I. Nikolic, *Second-order Differential Characteristics for SHA-256*, Advances in Cryptology - Asiacrypt 2011, Lecture Notes in Computer Science 7073, Springer Verlag 2011, 270-287
- [7] C. de Canniere, C. Rechberger, *Finding SHA-1 characteristics: General Results and Applications*, Advances in Cryptology - Asiacrypt 2006, Lecture Notes in Computer Science 4284, Springer Verlag 2006, 1-20
- [8] J. Chen, K. Jia, *Improved Related-Key Boomerang Attacks on Round-Reduced Threefish-512*, Information Security - Practice and Experience, Lecture Notes in Computer Science 6047, Springer Verlag 2010, 1-18
- [9] M. Daum, *Cryptanalysis of Hash Functions of the MD4 family*, PhD thesis, Ruhr-Universität Bochum, 2005
- [10] N. Ferguson, S. Lucks, B. Schneier, D. Whiting, M. Bellare, T. Kohno, J. Callas, J. Walker, *The Skein Hash Function Family*, Version 1.3, 1.10.2010, third round submission for the SHA-3 competition
- [11] D. Khovratovich and Ivica Nikolic, *Rotational Cryptanalysis of ARX*, Advances in Cryptology - Fast Software Encryption, Lecture Notes in Computer Science 6147, Springer Verlag 2010, 333-346
- [12] D. Khovratovic, I. Nikolic, C. Rechberger, *Rotational Rebound Attacks on Reduced Skein*, Advances in Cryptology - Asiacrypt 2010, Lecture Notes in Computer Science 6477, Springer Verlag 2010, 1-19
- [13] D. Khovratovic, C. Rechberger, A. Savelieva, *Bicliques for Preimages: Attacks on Skein-512 and the SHA-2 family*, IACR e-print report 286/2011
- [14] M. Lamberger, F. Mendel, *Higher-Order Differential Attack on Reduced SHA-256*, IACR e-print report 037/2011
- [15] H. Lipmaa, S. Moriai, *Efficient Algorithms for Computing Differential Properties of Addition*, Advances in Cryptology - Fast Software Encryption, Lecture Notes in Computer Science 2355, Springer Verlag 2002, 35-45
- [16] H. Lipmaa, J. Wallén, P. Dumas, *On the Additive Differential Probability of Exclusive-Or*, Advances in Cryptology - Fast Software Encryption 2004, Lecture Notes in Computer Science 3017, Springer Verlag 2004, 317-331
- [17] S. Liu, L. Wang, Z. Gong, *Improved Related-Key Boomerang Distinguishing Attack of Threefish-256*, IACR e-print report 323/2011
- [18] N. Mouha, V. Velichkov, C. de Canniere, B. Preneel, *The Differential Analysis of S-Functions*, Lecture Notes in Computer Science 6544 - Selected Areas in Cryptography 2010, Springer Verlag 2011, 36-56
- [19] N. Mouha, V. Velichkov, C. de Canniere, B. Preneel, *The Additive Differential Probability of ARX*, Advances in Cryptology - Fast Software Encryption 2011, Lecture Notes in Computer Science 6733, Springer Verlag 2011, 342-358
- [20] F. Mendel, T. Nad, M. Schläffer, *Finding SHA-2 Characteristics: Searching through a Minefield of Contradictions*, Advances in Cryptology - Asiacrypt 2011, Lecture Notes in Computer Science 7073, 288-307
- [21] B. Su, W. Wu, S. Wu, L. Lodang, *Near-collisions on the Reduced-Round Compression Function of Skein and Blake*, Cryptology and Network Security 2010, Lecture Notes in Computer Science 6467, Springer Verlag 2010, 124-139

- [22] D. Wagner, *The boomerang attack*, Advances in Cryptology - Fast Software Encryption 1999, Lecture Notes in Computer Science 1636, Springer Verlag 1999, 156-170
- [23] X. Wang, H. Yu, *How to Break MD5 and Other Hash Functions*, Advances in Cryptology - Eurocrypt 2005, Lecture Notes in Computer Science 3494, Springer Verlag 2005, 19-35
- [24] H. Yu, J. Chen, K. Jia, X. Wang, *Near-Collision Attack on the Step-Reduced Compression Function of Skein-256*, IACR e-print report 148/2011