Linear Cryptanalysis of TFHE-Friendly Cipher FRAST in the Average Key Setting

Xingwei Ren[1,2], Zhenyu Xiong[1,2], Bo Xu[1,2], Yongqiang Li[1,2] ✉ and Mingsheng Wang[1,2]

[1] State Key Laboratory of Cyberspace Security Defense, Institute of Information Engineering, Chinese Academy of Sciences, Beijing, China
[2] School of Cyber Security, University of Chinese Academy of Sciences, Beijing, China

Abstract

In the average key setting, a linear attack of Bak et al. on the TFHE-friendly stream cipher FRAST let wrong guesses pass, which they left as an open problem. The S-boxes of FRAST are public and random for every keystream block. For the mask on the least significant bit, we show that key differences fall into five indistinguishable classes per round. The filter of Bak et al. gives some wrong classes of the right parity a nonzero expected score. We compute the trail correlation exactly through the real cipher and bound the success of our statistic for every fixed master key. We bound its false alarms for a fixed master key with the last round kept and no filter, and for every attack against an ideal keystream.

With known XOF inputs, our algorithm Contract recovered the classes of a branch pair of 15 rounds and the master key of ten rounds. A cost model gives class recovery for every sampled key at 26 rounds with known inputs and for a third at 38 rounds with chosen inputs. In the accounting of Bak et al., it reaches full FRAST on weak master keys of density $2^{-11.8}$ to $2^{-5.4}$ within the claimed security. A distinguisher of full FRAST has a proved advantage on these keys, and its time over the advantage stays below the claimed security as well. Independent branch S-boxes remove the cheap classes at a cost in TFHE.

Keywords: FRAST, TFHE-friendly stream cipher, linear cryptanalysis, average key setting, weak keys, classes of key differences

1 Introduction

Homomorphic encryption lets a server compute on data that it receives only in encrypted form. Since the first fully homomorphic scheme of Gentry [Gen09], several schemes have been proposed, among them BGV [BGV12], BFV [FV12, Bra12], CKKS [CKKS17], FHEW [DM15] and TFHE [CGGI20]. Their ciphertexts are much larger than the data, and encrypting under them is heavy for a client. Naehrig et al. [NLV11] proposed transciphering, also called hybrid homomorphic encryption, to address both costs. The client encrypts its data with a symmetric cipher and sends the homomorphic encryption of the key once, and the server evaluates the decryption of the cipher homomorphically. The data travel without expansion, and the cost moves to the server. A cipher for this framework must therefore be cheap to evaluate in the homomorphic scheme.

Emails: renxingwei@iie.ac.cn (Xingwei Ren), xiongzhenyu@iie.ac.cn (Zhenyu Xiong), xubo2024@iie.ac.cn (Bo Xu), liyongqiang@iie.ac.cn (Yongqiang Li), wangmingsheng@iie.ac.cn (Mingsheng Wang)

Date of this document: 2026-10-22.

Many symmetric ciphers have been designed for this evaluation. Among them are the block ciphers LowMC [ARS+15] and Chaghri [AMT22] and the stream ciphers Kreyvium [CCF+18], FLIP [MJSC16], Rasta [DEG+18] and its variants Dasta [HL20], Fasta [CIR22] and Pasta [DGH+23, GLR+24]. Further designs are Masta [HKC+20], HERA [CHK+21], Rubato [HKL+22], Elisabeth-4 [CHMS22], Transistor [BBB+25] and FRAST [CCH+24]. For BGV and BFV, a cipher should use few multiplications. Several of the designs draw a component at random for each encryption: the affine layers in Rasta, Masta and Pasta, the permutations in FLIP and Elisabeth-4, the key schedule in HERA and Rubato and the S-boxes in FRAST. Kreyvium, Chaghri and Dasta have no such component, and the second version of Pasta keeps only an initial random matrix. Masta, Pasta, HERA, Rubato and Transistor work over non-binary fields or rings. Elisabeth-4 and FRAST work on the ring $\mathbb{Z}_{16}$ with negacyclic S-boxes, which TFHE evaluates by table lookups.

Of these designs, Rasta and Dasta [LSMI21, LSMI22], LowMC [LSW+22] and HERA for some parameter sets [LKSM24] have been attacked algebraically, and coefficient grouping broke the full Chaghri [LAW+23]. Gilbert et al. [GBJR23] broke the full Elisabeth-4 by linearisation. Grassi et al. [GAH+23] recovered the key of Rubato for a fraction of the moduli, and Duval et al. [DLR16] attacked FLIP through the structure of its filter function. Beyne et al. [BGØ+26] studied symmetric primitives over integer rings. In a physical model, Wang and Tang [WT26] gave differential fault attacks on FRAST, which lie outside our models. Cho et al. and Bak et al. applied the linear cryptanalysis of Matsui [Mat94] to $\mathbb{Z}_{16}$ in the approach of Baignères et al. [BSV07].

Cho et al. [CCH+24] designed the TFHE-friendly stream cipher FRAST for transciphering. The S-boxes of FRAST are public and random for every keystream block. An extendable output function (XOF) derives them from the nonce and the counter of the block in 32 of the 40 rounds, which Bak et al. call random rounds. The other eight rounds are fixed rounds with one fixed S-box. A linear key schedule derives the round keys from the master key, and FRAST claims 128 bits of security with at most 2^{64} blocks under one key. The designers gave a linearisation attack on 20 random rounds and chose 32 random rounds as a precaution against unknown attacks that guess round keys. In their analysis, a trail through more than two S-box evaluations needs the differences of their round key words. They called guessing them infeasible, as it needs more than 128 bits for 32 random rounds.

Against this argument, Bak et al. [BGL+26] found an alternative representation of the round function, in which three linear combinations of the state evolve on their own. They used it for a distinguisher on full FRAST for a class of weak keys. They also turned an iterated trail of this representation into a linear attack in the average key setting that is not restricted to weak keys. The correlation of the trail depends on the key only through the differences of the round key words of two branches. The attack keeps the blocks whose S-boxes have a large expected linear potential (ELP). On five rounds, the right guess passed their test, but a non-negligible fraction of the wrong guesses passed as well. In their Section 4.4, Bak et al. left the explanation of these wrong guesses as an open problem, and their conclusion names a working linear attack in this setting as a further one. This paper gives the explanation and working average key attacks on reduced-round FRAST, in the known-nonce and the chosen-nonce models of Section 2.4.

Our Contributions. To explain the wrong guesses, we first show that the key differences of a round fall into five classes for the mask on the least significant bit. In the observation model of Section 3.1, no test on this mask separates two differences of one class. For a uniformly random negacyclic S-box, the autocorrelations at differences of different classes have the expected product zero. The ELP filter of Bak et al. merges classes at its strongest thresholds. Below them, it gives some wrong classes of the right parity a nonzero expected

score, while by the parity rule a parity mismatch keeps it zero. An attack therefore targets the class of each round rather than the key difference. Two of the five classes cost little or no data, because the branches share one expanding S-box per round.

We then compute the correlation of the iterated trail through the real cipher, for every fixed master key. Inside a segment of random rounds, the contracting S-box makes the next input uniform. The expected observation of a block given its expanding S-boxes is therefore the exact product of their autocorrelations. A fixed round acts by a transfer matrix, and its side trails, which are as large as the main trail block by block, add nothing to the expected score under conditions that our weights and filters satisfy. Our matched filter statistic weights each block by the autocorrelations that a guess predicts. We prove bounds on its success side for all our attacks and on its false alarms with the last round kept and no filter. Against an ideal keystream, we bound the false alarms of every attack without a symmetry of the wrong guesses, which gives a distinguisher of full FRAST with a proved advantage for weak master keys. Experiments reproduce the observation of Bak et al. and check the correlation and success models and our implementation.

Our online algorithm Contract multiplies the classes into the blocks round by round and merges blocks that agree on the remaining rounds, in storage linear in the data. With it, we executed known-nonce attacks that recover, on 15 rounds, the classes of a branch pair chosen with knowledge of the key and, on ten rounds, the master key from the accepted branch pairs. Since $\mathbb{Z}_{16}$ is not a field, we solve the key schedule bit by bit, with the top bit from the fixed rounds.

Consequences for FRAST. In a cost model with the data limit of the designers, our class recovery reaches 26 of the 40 rounds of FRAST with known nonces for every sampled master key. With chosen nonces that make some rounds free for all branch pairs, it reaches 38 rounds for a third of the keys. On full FRAST, in the accounting of Bak et al., it covers weak master keys of density $2^{-11.8}$ to $2^{-5.4}$ with time over density below the claimed security. A distinguisher for these keys keeps its time over its proved advantage below the claimed security as well, and in bit operations alone this holds by 0.37 bits at the model constant, for one choice of five restricted rounds.

These densities are about 2^{49} to 2^{56} times that of the weak keys of their distinguisher. Behind this reach are the cheap classes and the rounds freed for all pairs by the choice of the nonce, and both come from the expanding S-box that the branches of a round share, on which the double blind rotation of the designers also rests. This technique lets TFHE evaluate the S-boxes of a round at about the cost of one, and independent branch S-boxes remove the cheap classes but lose it.

Related Work. For one round, the spectral identity follows from the composition rule and the diagonal action of a key addition in the geometric approach of Beyne [Bey21]. Our statements are fixed-key in his sense, except those over a random key or key difference. Kranz et al. [KLW17] showed that round constants on a linear key schedule remove the cross terms between trails on average. Averaging over the public S-boxes gives a wrong class zero expected score in the same way. Röck and Nyberg [RN13] divided the keys into classes of equal correlation, and our statistic has a weighted form of the variance of Li et al. [LMSW22]. Bogdanov and Tischhauser [BT14] let a wrong key keep a bias of its own. Beyne et al. [BLMR26] found classes of wrong keys that pass a differential filter almost as often as the right key and resemble our wrong guesses of the right parity. Civino [Civ26] built distinguishers for weak keys of CRAFT, and Hadipour et al. [HMAV26] reported correlations as distributions over master keys, as our cost model reports fractions of keys.

Outline. Section 2 recalls FRAST, the attack of Bak et al. and our models. We build our attacks in Sections 3 to 5 by explaining the wrong guesses, computing the correlations

and bounding the errors of the statistic. Sections 6 to 8 give CONTRACT, the attacks with their reach and their experimental validation, and the variant, and Section 9 concludes.

2 Preliminaries

In this section, we fix the notation, recall the specification of FRAST and the alternative representation of Bak et al. with its iterated trail, and recall their attack before we define the two models of our attacks.

2.1 Notation

The words of FRAST are elements of the ring $\mathbb{Z}_{16}$ of integers modulo 16, and all arithmetic on words is in this ring. Since $\mathbb{Z}_{16}$ is not a field, every linear system over $\mathbb{Z}_{16}$ in this paper is solved by modular arithmetic over the ring. We write $\zeta = e^{2\pi i/16}$, and for a mask $a \in \mathbb{Z}_{16}$ we write $\chi_a(x) = \zeta^{ax}$ for the corresponding character of $\mathbb{Z}_{16}$. For the mask 8, the character $\chi_8(x) = (-1)^x$ depends only on the least significant bit $\mathrm{LSB}(x)$ of x. We number the rounds from 1 to 40 and the branches, which are the positions of the state, from 1 to 32, as Cho et al. [CCH+24] do. A superscript (r) marks a quantity of round r, and $y^{(r)}$ is the state after round r, with $y^{(0)}$ the input of the block cipher. We write $\mathbf{1}[E]$ for the indicator of a condition E, $\mathbb{E}$ for an expectation and log for the logarithm to the base 2.

2.2 Specification of FRAST

Cho et al. [CCH+24] specified FRAST as a stream cipher with a master key $k \in \mathbb{Z}_{16}^{64}$ of 256 bits and a claimed security level of 128 bits. A keystream block has 32 words, that is, 128 bits. It is the image of the fixed input $\mathrm{ic} = (0, 1, \ldots, 15, 0, 1, \ldots, 15) \in \mathbb{Z}_{16}^{32}$ under a block cipher of 40 rounds, whose S-boxes depend on a nonce nc of 64 bits and a counter ctr of 64 bits. The block of the nonce nc with the counter value ctr is

$$z = R_{40} \circ R_{39} \circ \cdots \circ R_1(\mathrm{ic}) \in \mathbb{Z}_{16}^{32}, \tag{1}$$

where the counter runs through $0, 1, 2, \ldots$ along the keystream of one nonce. Round r maps the state $y^{(r-1)}$ to the state $y^{(r)}$ by

$$\begin{aligned} y_j^{(r)} &= y_j^{(r-1)} + S_{\mathrm{erf}}^{(r)}\big(y_1^{(r-1)} + \mathrm{rk}_j^{(r)}\big), \qquad j = 2, \ldots, 32, \\ y_1^{(r)} &= y_1^{(r-1)} + S_{\mathrm{crf}}^{(r)}\Big(\mathrm{rk}_1^{(r)} + \textstyle\sum_{j=2}^{32} y_j^{(r)}\Big), \end{aligned} \tag{2}$$

where the sum in the second line is over the updated words. Each round uses a round key $\mathrm{rk}^{(r)} = (\mathrm{rk}_1^{(r)}, \ldots, \mathrm{rk}_{32}^{(r)}) \in \mathbb{Z}_{16}^{32}$ and two S-boxes from $\mathbb{Z}_{16}$ to $\mathbb{Z}_{16}$, which are the expanding S-box $S_{\mathrm{erf}}^{(r)}$ and the contracting S-box $S_{\mathrm{crf}}^{(r)}$, as shown in Figure 1. The figure draws the words y_1, y_2 and y_{32}, and every $+$ in it is an addition in $\mathbb{Z}_{16}$. Its node Σ adds the updated words $y_2^{(r)}, \ldots, y_{32}^{(r)}$.

Following Bak et al. [BGL+26], we call round r a fixed round when $r \equiv 0 \pmod 5$ and a random round otherwise. FRAST therefore has eight fixed rounds and 32 random rounds, and we call a maximal sequence of consecutive random rounds a segment. The segments are the rounds $5b - 4$ to $5b - 1$ for $b = 1, \ldots, 8$. Both S-boxes of a fixed round are the permutation

$$\big(S_{\mathrm{fix}}(0), S_{\mathrm{fix}}(1), \ldots, S_{\mathrm{fix}}(15)\big) = (0, 3, 5, 8, 6, 9, 12, 7, 13, 10, 14, 4, 1, 15, 11, 2). \tag{3}$$

In a random round, both S-boxes are negacyclic, which means that

$$S(x + 8) = -S(x) \qquad \text{for all } x \in \mathbb{Z}_{16}. \tag{4}$$

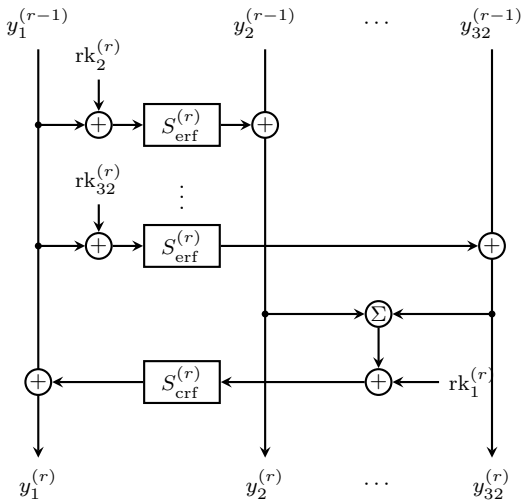

Figure 1: The round function R_r of FRAST.

A negacyclic S-box is therefore determined by its eight values $S(0), \ldots, S(7)$, and S_{fix} is not negacyclic. Cho et al. took these values from the output of an extendable output function on the input $\mathrm{nc} \,\|\, \mathrm{ctr}$ of 128 bits, which does not contain the key. They instantiated the function with SHAKE256 [Nat15], and each random S-box consumes 32 bits of the output. As Bak et al. noted, no value is rejected, so a random S-box need not be a permutation. Appendix A gives the order in which the output bits fill the S-boxes.

The key schedule is linear over $\mathbb{Z}_{16}$ and adds no constant. An invertible matrix $M \in \mathbb{Z}_{16}^{64\times 64}$, given in Supplementary Material G of the specification [CCH+24], defines the round keys by

$$\mathrm{rk}^{(2i-1)} \,\|\, \mathrm{rk}^{(2i)} = M^i k, \qquad i = 1, \ldots, 20. \tag{5}$$

The round keys of different rounds are therefore linear functions of one master key and not independent. Cho et al. limited the number of keystream blocks under one key to 2^{64} and modelled the extendable output function as a random oracle. Since the input ic of the block cipher is a constant, the only input an attacker can influence is the nonce, which selects the S-boxes and cannot fix them.

2.3 The Alternative Representation and the Iterated Trail

Bak et al. [BGL+26] showed that three linear combinations of the state evolve on their own. For a branch pair (i, j) with $2 \le i < j \le 32$, round r acts on $x_1 = y_1$, $x_2 = \sum_{l=2}^{32} y_l$ and $x_3 = y_i - y_j$ by

$$x_2 \leftarrow x_2 + F_r(x_1), \qquad x_3 \leftarrow x_3 + \Delta_r(x_1), \qquad x_1 \leftarrow x_1 + S_{\mathrm{crf}}^{(r)}\big(x_2 + \mathrm{rk}_1^{(r)}\big), \tag{6}$$

where the update of x_1 uses the new value of x_2, and where

$$F_r(z) = \sum_{l=2}^{32} S_{\mathrm{erf}}^{(r)}\big(z + \mathrm{rk}_l^{(r)}\big), \qquad \Delta_r(z) = S_{\mathrm{erf}}^{(r)}\big(z + \mathrm{rk}_i^{(r)}\big) - S_{\mathrm{erf}}^{(r)}\big(z + \mathrm{rk}_j^{(r)}\big). \tag{7}$$

We call this system the alternative representation, as Bak et al. do, and there is one such system for each of the 465 branch pairs formed by the 31 branches from 2 to 32. The word x_3 changes only by addition, so a linear mask on x_3 is preserved by every round. Both

S-boxes of Δ_r receive the same word x_1, shifted by round key words whose difference is the key difference

$$t'_r = \mathrm{rk}_i^{(r)} - \mathrm{rk}_j^{(r)} \in \mathbb{Z}_{16}. \tag{8}$$

With $u = x_1 + \mathrm{rk}_j^{(r)}$, the increment of x_3 is $S_{\mathrm{erf}}^{(r)}(u + t'_r) - S_{\mathrm{erf}}^{(r)}(u)$, which is a derivative of $S_{\mathrm{erf}}^{(r)}$ in the direction t'_r. For a function $S\colon \mathbb{Z}_{16} \to \mathbb{Z}_{16}$, a mask $a \in \mathbb{Z}_{16}$ and a difference $t \in \mathbb{Z}_{16}$, we define the autocorrelation

$$C_{a,S}(t) = \frac{1}{16} \sum_{x \in \mathbb{Z}_{16}} \zeta^{a(S(x+t)-S(x))}. \tag{9}$$

It is the correlation of the derivative $x \mapsto S(x+t) - S(x)$ for the input mask 0 and the output mask a, which Bak et al. write as $C^{\Delta S}(a, 0)$.

With the iterated trail $(0,0,a) \rightsquigarrow (0,0,a) \rightsquigarrow \cdots \rightsquigarrow (0,0,a)$ of the alternative representation, Bak et al. attacked FRAST. They wrote its correlation over R rounds as the product of $C_{a,S_{\mathrm{erf}}^{(r)}}(t'_r)$ over $r = 1, \ldots, R$, which depends on the key only through the key differences. Whether the expected observation of a block equals this product for the real cipher is the subject of Section 4. For the mask $a = 8$, the trail predicts the sign

$$(-1)^{x_3^{(R)} - x_3^{(0)}} = (-1)^{\mathrm{LSB}(z_i) \oplus \mathrm{LSB}(z_j) \oplus \mathrm{LSB}(\mathrm{ic}_i) \oplus \mathrm{LSB}(\mathrm{ic}_j)}, \tag{10}$$

which the attacker computes from two words of a keystream block and the public input. This paper uses the mask $a = 8$ throughout, except where a statement names another mask.

2.4 The Attack of Bak et al. and Our Models

Bak et al. turned the iterated trail into an attack on R rounds. In a precomputation, they computed the S-boxes of many XOF inputs and kept the inputs whose inner rounds have a large expected linear potential,

$$\mathrm{ELP}_a(S) = \frac{1}{16} \sum_{t \in \mathbb{Z}_{16}} \left| C_{a,S}(t) \right|^2, \qquad \prod_{r=2}^{R-1} \mathrm{ELP}_a\left(S_{\mathrm{erf}}^{(r)}\right) \geq \tau^2, \tag{11}$$

for a threshold τ^2, and the precomputation costs the number of kept inputs divided by the probability that a random input is kept. They then guessed the key material of the first and the last round, together with guesses $t_2, \ldots, t_{R-1}$ of the key differences of the inner rounds. For each guess, they computed the empirical correlation

$$\hat{c} = \frac{1}{N} \sum_{u=1}^{N} g_u \, \zeta^{a\tilde{x}_u}, \qquad g_u = \mathrm{sgn} \prod_{r=2}^{R-1} C_{a,S_{\mathrm{erf},u}^{(r)}}(t_r), \tag{12}$$

where N is the number of kept blocks and $S_{\mathrm{erf},u}^{(r)}$ is the expanding S-box of round r in block u. The value $\tilde{x}_u$ is the increment of x_3 over the inner rounds under the guess of the outer rounds, and a guess passed the test when $|\hat{c}|^2 \geq (1 - \varepsilon)\tau^2$.

On five rounds, Bak et al. implemented the attack and set the thresholds by the methodology of Baignères et al. [BSV07]. Their right key hypothesis assumed that the empirical correlation of the right guess converges to its expected value, and their wrong key hypothesis assumed that every wrong guess behaves as noise around zero. The right guess passed the test in all their experiments, but a non-negligible fraction of wrong guesses passed as well. They concluded that the wrong key hypothesis was the most plausible point of failure. When only the guess of round 2 was wrong, they observed nontrivial biases for some $t_2 \neq t'_2$.

Our attacks are in the known-nonce model, where the attacker receives keystream blocks under one unknown master key with their XOF inputs, and in the chosen-nonce model, where it also chooses these inputs. They follow the model that Cho et al. stated in their Section 5, in which the adversary chooses nonces and receives the keystream without information on the secret key. Cho et al. defined the cipher on an input of 128 bits, which is $\mathrm{nc} \,\|\, \mathrm{ctr}$ along a keystream. We read the nonce of their model as this whole XOF input.

A chosen-nonce attack therefore selects its blocks by a filter on their S-boxes and pays one XOF call for every input it examines, as the precomputation of Bak et al. does. A counter that had to run sequentially would require recomputing the chosen-nonce costs of Section 7 and leave the known-nonce results unchanged. With 2^{64} nonces and the counter 0 alone, a filter of ten restricted rounds of the first kind of Section 5.1 keeps about $2^{64-60} = 2^4$ inputs, so the chosen-nonce attacks need the whole input of 128 bits. Following Bak et al., we speak of the average key setting when an attack is not restricted to a class of weak keys.

The offline time T_{off} counts XOF calls, the online time T_{on} counts bit operations as Bak et al. do, or the word-level multiply-adds of Section 6.1, and the storage counts words. Only the criterion of Section 7.3 adds XOF calls and bit operations, in their accounting. Unless a statement says otherwise, the master key is fixed and every expectation is over the public S-boxes, so an expected value is an average over the blocks and does not hold block by block. Reduced-round FRAST is the first R rounds of the real cipher with the real key schedule and the XOF as a random oracle, and full FRAST is the case $R = 40$. Sections 3 and 4 also use the single S-box model, in which one negacyclic S-box is uniformly random and the key difference is fixed, and the ideal multi-round model, in which the inner rounds have independent uniformly random negacyclic S-boxes.

3 Spectra, Classes of Key Differences and the Parity Rule

In this section, we show that the key differences of a round fall into classes that no statistic of the mask 8 separates, and that a wrong class has zero expected score without a filter. The classes are therefore the objects that our attacks guess. A filter on the ELP keeps this zero in general only for a parity mismatch, and the classes and the wrong guesses of the same parity explain what Bak et al. observed. Every statement is in the single S-box model or in the ideal multi-round model of Section 2.4. We use a finite abelian group G of order n and the group $\widehat{G}$ of its characters $\chi\colon G \to \mathbb{T}$ into the complex unit circle $\mathbb{T}$, as presented by Terras [Ter99].

3.1 Spectra and Classes of Key Differences

For $f\colon G \to \mathbb{T}$, the *autocorrelation*, the *Fourier coefficients*, the *spectrum* and the *expected linear potential* of f are

$$C_f(t) = \frac{1}{n}\sum_{x\in G} f(x+t)\overline{f(x)}, \qquad \hat{f}(\chi) = \frac{1}{n}\sum_{x\in G} f(x)\overline{\chi(x)},$$

$$P_\chi^f = \left|\hat{f}(\chi)\right|^2, \qquad \mathrm{ELP}(f) = \frac{1}{n}\sum_{t\in G}\left|C_f(t)\right|^2,$$

for $t \in G$ and $\chi \in \widehat{G}$, and the *support* of the spectrum is $\operatorname{supp} P^f = \{\chi \in \widehat{G} : P_\chi^f \neq 0\}$. For FRAST, $G = \mathbb{Z}_{16}$, every character is χ_k for one $k \in \mathbb{Z}_{16}$, and a mask a turns an S-box S into the function $f = \chi_a \circ S$. Then C_f is $C_{a,S}$ of equation (9), $\mathrm{ELP}(f)$ is $\mathrm{ELP}_a(S)$ of equation (11), and $\hat{f}(\chi_k) = C^S(a,k)$ is the correlation of S for the input mask k and the output mask a, and we write P_k for $P_{\chi_k}^f$. Let $\mathcal{F}$ be a nonempty family of functions $G \to \mathbb{T}$, for instance the functions $\chi_a \circ S$ of the S-boxes that a filter admits. Two elements $t, t' \in G$

are *indistinguishable* for $\mathcal{F}$, written $t \approx_{\mathcal{F}} t'$, if $C_f(t) = C_f(t')$ for every $f \in \mathcal{F}$, and the equivalence classes of $\approx_{\mathcal{F}}$ are the *classes of key differences* for $\mathcal{F}$.

Theorem 1 (Spectral Identity and ELP). *For every $f\colon G \to \mathbb{T}$ and every $t \in G$,*

$$C_f(t) = \sum_{\chi \in \widehat{G}} P_\chi^f \, \chi(t), \qquad \sum_{\chi \in \widehat{G}} P_\chi^f = 1.$$

In particular, $C_f(0) = 1$, $|C_f(t)| \le 1$ and $C_f(-t) = \overline{C_f(t)}$. The ELP is the collision probability of the distribution P^f on $\widehat{G}$, and it satisfies

$$\mathrm{ELP}(f) = \sum_{\chi} (P_\chi^f)^2 \quad \textit{and} \quad 1/|\operatorname{supp} P^f| \le \mathrm{ELP}(f) \le \max_\chi P_\chi^f \le 1.$$

The equality $\mathrm{ELP}(f) = 1$ holds if and only if P^f is a point mass at some $\eta \in \widehat{G}$, and if and only if $|C_f(t)| = 1$ for all t. In this case, $f = \gamma\eta$ for some $\gamma \in \mathbb{T}$, $C_f = \eta$, and $C_f(t) = C_f(t')$ holds if and only if $t - t' \in \ker \eta$.

Proof. Substituting the expansion $f = \sum_\chi \hat{f}(\chi)\chi$ in the orthonormal basis of characters and $\chi(x+t) = \chi(x)\chi(t)$ into C_f gives

$$C_f(t) = \sum_{\chi, \chi' \in \widehat{G}} \hat{f}(\chi)\overline{\hat{f}(\chi')}\, \chi(t)\, \frac{1}{n} \sum_{x \in G} \chi(x)\overline{\chi'(x)} = \sum_{\chi \in \widehat{G}} \left|\hat{f}(\chi)\right|^2 \chi(t).$$

At $t = 0$, $\sum_\chi P_\chi^f = C_f(0) = \frac{1}{n}\sum_x |f(x)|^2 = 1$, and $P_\chi^f \ge 0$ and $|\chi(t)| = 1$ give the bound and the symmetry. Parseval's identity for the coefficients P_χ^f of C_f gives the formula of the ELP. The Cauchy–Schwarz inequality gives $1 = \left(\sum_{\chi \in \operatorname{supp} P^f} P_\chi^f\right)^2 \le |\operatorname{supp} P^f| \sum_\chi (P_\chi^f)^2$, and $\sum_\chi (P_\chi^f)^2 \le \max_\chi P_\chi^f \sum_\chi P_\chi^f = \max_\chi P_\chi^f$. If $\mathrm{ELP}(f) = 1$, then $\max_\chi P_\chi^f = 1$, and the point mass at η gives $f = \hat{f}(\eta)\eta$ with $|\hat{f}(\eta)| = 1$ and $C_f = \eta$. Conversely, a character has modulus 1, and $|C_f| = 1$ gives $\mathrm{ELP}(f) = 1$. The last claim holds because $\eta(t) = \eta(t')$ is equivalent to $\eta(t - t') = 1$. □

For the mask 8, the value $(-1)^{S(x)}$ depends only on x mod 8, because $S(x+8) = -S(x)$ has the parity of $S(x)$. We call $b = (b_x)_{x \in \mathbb{Z}_8}$ with $b_x = (-1)^{S(x)}$ the *parity pattern* of S, and we write $c(t) = C_{8,S}(t)$. By equation (9),

$$c(t) = \frac{1}{8} \sum_{x \in \mathbb{Z}_8} b_x \, b_{x+t}, \tag{13}$$

with indices modulo 8. A uniformly random negacyclic S-box has a parity pattern that is uniform among the 256 sign vectors of length 8. Negacyclicity makes the spectrum symmetric, $P_{-k} = P_k$, and every autocorrelation $C_{a,S}$ real and even. For the mask 8, $P_k = 0$ for every odd k and c has the period 8. The two constant and the two alternating parity patterns are the only ones with $\mathrm{ELP}_8(S) = 1$. These facts are proved in Appendix B.2, together with Theorem 2.

Lemma 1 (No Test Separates Indistinguishable Differences). *Let an attacker process blocks adaptively and possibly with randomness, the j-th block giving a public function $f_j \in \mathcal{F}$ and an observation Y_j. Assume that, given the history, the distribution of f_j does not depend on the key difference (H1), and the distribution of Y_j given also f_j depends on the key difference only through the value of C_{f_j} at it (H2). If $t \approx_{\mathcal{F}} t'$, then the record $(f_1, Y_1, f_2, Y_2, \dots)$ has the same distribution under the key differences t and t', and every test has advantage zero in deciding between them.*

Proof. The distribution of the record is the product of the conditional distributions of its entries. By (H1), the conditional distribution of f_j is the same under both key differences. By (H2) and $C_{f_j}(t) = C_{f_j}(t')$, so is the conditional distribution of Y_j, and induction on j gives the claim. □

Theorem 2 (Classes of Key Differences and Orthogonality). *Let S be a uniformly random negacyclic S-box.*

(i) For the mask 8, the key differences fall into the five classes $[0] = \{0, 8\}$, $[1] = \{1, 7, 9, 15\}$, $[2] = \{2, 6, 10, 14\}$, $[3] = \{3, 5, 11, 13\}$ and $[4] = \{4, 12\}$. A difference t lies in $[h]$ for the $h \in \{0, \dots, 4\}$ with $t \equiv \pm h \pmod 8$. For $t' \notin [t]$, $\mathbb{E}[c(t)c(t')] = 0$, and $\mathbb{E}[c(t)^2]$ is 1 on $[0]$, $\frac{1}{8}$ on $[1]$, $[2]$ and $[3]$, and $\frac{1}{4}$ on $[4]$.

(ii) For a mask $a \notin \{0, 8\}$, the key differences fall into the nine classes $\{t, -t\}$, and $\mathbb{E}[C_{a,S}(t)C_{a,S}(t')] = 0$ for $t' \notin \{t, -t\}$. The second moments are $\mathbb{E}[C_{a,S}(t)^2] = \frac{1}{16}$ for $t \notin \{0, 8\}$, and $\mathbb{E}[C_{a,S}(8)^2]$ is $\frac{1}{8}$ for $a \in \{4, 12\}$ and $\frac{1}{16}$ for the other masks.

(iii) In the ideal multi-round model with the mask 8, let a guess (t_r) differ from the key differences (t'_r) in class in at least one inner round. Then $\mathbb{E}\big[\prod_r c_r(t_r)\, c_r(t'_r)\big] = 0$, where c_r is the autocorrelation of equation (13) for round r.

Assumption (H1) holds when the XOF is a random oracle and blocks are selected by their public S-boxes, and (H2) for the mask 8 in the ideal multi-round model, but not across a fixed round, as discussed in Appendix B.2. Theorem 2 follows from the kernel $K(t, t') = \mathbb{E}\big[\overline{C_f(t)}\, C_f(t')\big]$ of a random function f, which Theorem 10 computes. Its zeros are averages over the S-box and do not hold block by block. For $f = \chi_a \circ S$, the kernel is the expected contribution of one block to the score of the guess t under the key difference t' with the *matched filter weight* $\overline{C_f(t)}$ of t for f. The *phase weight* of t for f is $\overline{C_f(t)}/|C_f(t)|$, set to 0 where $C_f(t) = 0$, and for a real C_f, the *sign weight* is $\operatorname{sgn} C_f(t)$.

3.2 Filters and the Parity Rule

A *filter* is an event Φ of positive probability on the S-boxes of a block, and an attacker with a filter uses only the blocks in Φ. For a family $\mathcal{F}$, the *ELP filter* with the threshold τ^2 keeps $\mathcal{F}_\tau = \{f \in \mathcal{F} : \mathrm{ELP}(f) \geq \tau^2\}$, and the *product filter* of Bak et al. keeps the blocks with $\prod_r \mathrm{ELP}(f_r) \geq \tau^2$ over the inner rounds, as in equation (11). By Theorem 1, an ELP filter keeps the S-boxes whose spectrum is concentrated on few characters. A stronger ELP filter keeps fewer functions and can only merge classes. The proof of the following parity rule in Appendix B.2 modulates the function of one round by a character and applies Lemma 3.

Theorem 3 (Parity Rule). *Let $a \neq 0$ be a mask of FRAST, and consider the ideal multi-round model. Let the filter Φ be the absence of a filter, a single-round ELP filter or the product filter. Let w be a product over the inner rounds of matched filter, phase or sign weights at a guess (t_r). If $t_r \not\equiv t'_r \pmod 2$ for some inner round r, then*

$$\mathbb{E}\Big[w \prod_r C_{a,S_{\mathrm{erf}}^{(r)}}(t'_r) \;\Big|\; \Phi\Big] = 0.$$

The entries of Table 1 are the matched filter kernel $K(t, t') = \mathbb{E}[c(t)c(t') \mid \Phi]$ of one random round under each filter on the parity pattern, at the class representatives $t, t' \in \{0, \dots, 4\}$. Kept is the number of the 256 parity patterns that the filter admits, and Classes is the number of classes for statistics of the mask 8 alone. The entries are exact over the kept patterns, and the kernel is symmetric in t and t'. The last row keeps the two constant parity patterns, which are the weak S-boxes of Bak et al.

Table 1: Kernel of one random round for the mask 8 under each filter.

Filter Φ	Kept	Classes	$K(0,1)$	$K(0,2)$	$K(0,4)$	$K(1,3)$	$K(2,4)$	$K(1,1)$	$K(2,2)$	$K(4,4)$
none	256	5	0	0	0	0	0	1/8	1/8	1/4
$\mathrm{ELP}_8 \geq 3/16$	224	5	0	0	1/14	0	0	1/7	1/7	1/4
$\mathrm{ELP}_8 \geq 7/32$	128	5	0	0	1/8	0	0	3/16	3/16	7/16
$\mathrm{ELP}_8 \geq 1/4$	96	5	0	1/6	0	0	1/12	1/4	1/6	1/2
$\mathrm{ELP}_8 \geq 9/32$	88	5	0	2/11	−1/11	0	1/11	3/11	2/11	5/11
$\mathrm{ELP}_8 \geq 11/32$	56	5	0	2/7	1/7	1/7	1/7	2/7	2/7	4/7
$\mathrm{ELP}_8 \geq 3/8$	24	5	0	0	−1/3	0	0	1/3	1/3	1
$\mathrm{ELP}_8 \geq 1/2$	8	3	0	0	1	1/2	0	1/2	1	1
$\mathrm{ELP}_8 = 1$	4	2	0	1	1	1	1	1	1	1
weak S-boxes	2	1	1	1	1	1	1	1	1	1

Table 1 shows two effects of a filter, and the first is that the exact classes merge only at the top thresholds, where the five classes become three at $\mathrm{ELP}_8 \geq \frac{1}{2}$ and the even and the odd differences at $\mathrm{ELP}_8 = 1$. Below these thresholds, the five classes remain, but some kernel entries between classes of the same parity become nonzero. For instance, at $\mathrm{ELP}_8 \geq \frac{11}{32}$ and under the key difference 2, the wrong guess 0 has the expected score $K(0,2) = \frac{2}{7}$, which is as large as the score $K(2,2) = \frac{2}{7}$ of the right guess. The entries with an odd difference, such as $K(0,1)$, vanish under every ELP filter, as required by Theorem 3. Only the filter of the weak S-boxes, which is not invariant under the modulation by χ_8, violates the parity rule. The product filter of Bak et al. behaves in the same way.

These kernels explain what Bak et al. observed. Their mask and thresholds are not published, and we use the mask 8 and those of Table 1. The three inner rounds of their experiment lie in the first segment, where reduced-round FRAST satisfies the ideal multi-round model by Theorem 4 when the outer guesses are right and the filter looks only at the expanding S-boxes, as theirs does. A wrong guess of round 2 with a parity mismatch therefore passes their test only at the rate of noise, and the wrong guesses of the outer rounds that they also counted lie outside Theorem 3.

Two causes account for the wrong guesses that passed. The first cause is the classes, which no statistic of the mask 8 separates by Lemma 1 and which the strongest filters merge further. The second cause is the ELP filter, which gives some wrong guesses of the same parity a nonzero expected score. At the threshold $\mathrm{ELP}_8 \geq 11/32$ of Table 1, $K(0,2) = K(2,2) = 2/7$, so a wrong guess of the class [0] has the same expected score as the right class [2]. A second mask separates merged classes: with the masks 8 and 4 together, the filtered families with two, three and five classes all have the nine classes $\{t, -t\}$ of Theorem 2(ii). Section 7.4 reproduces the observation on real five-round FRAST and checks the parity rule there.

4 Exact Correlations of the Iterated Trail

In this section, we compute the expected observation of a block of reduced-round FRAST given its expanding S-boxes, by a product of autocorrelations inside a segment and a transfer matrix across each fixed round. This expected observation gives the expected score of a guess in our attacks. We then show that the side trails of the fixed rounds average to zero against the weights of our guesses and determine the usable branch pairs. Throughout, the master key, the branch pair (i, j) and the mask 8 are fixed, and the proofs are in Appendix B.3. The first round contributes a sign that a guess strips, and the last round belongs to the key guess of Section 5.

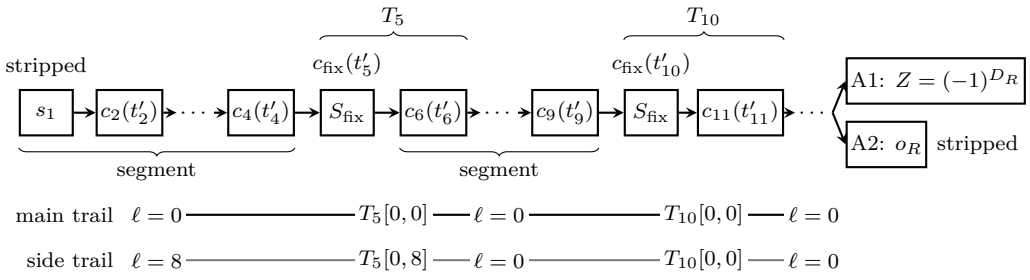

Figure 2: The approximation of R rounds: the stripped first round, the segments of random rounds, the fixed rounds with their transfer matrices, the two treatments of the last round, and the value of ℓ of the main trail and of one side trail.

4.1 Segments and Fixed Rounds

In the alternative representation of equation (6), round r reads the word $u_r = x_1^{(r-1)}$, and we write $D_r = x_3^{(r)} - x_3^{(0)}$. These words satisfy $x_2^{(r)} = x_2^{(r-1)} + F_r(u_r)$, $D_r = D_{r-1} + \Delta_r(u_r)$ and $x_1^{(r)} = u_r + S_{\text{crf}}^{(r)}\big(x_2^{(r)} + \text{rk}_1^{(r)}\big)$, and without a guess of the last round, the observation of equation (10) is $Z = (-1)^{D_R}$. We write $\mathcal{E}$ for the tuple of the expanding S-boxes of a block, $b^{(r)}$ for the parity pattern of $S_{\text{erf}}^{(r)}$, c_r for its autocorrelation in equation (13) and $c_{\text{fix}}(t) = C_{8,S_{\text{fix}}}(t)$. With the XOF as a random oracle, the S-boxes of different rounds and blocks are independent, and both S-boxes of a random round are uniformly random negacyclic functions.

Theorem 4 (Exact Product Inside a Segment). *Let the XOF be a random oracle, and let the filter depend only on the expanding S-boxes. Let $q \le q'$ be rounds such that the rounds $q-1, q, \dots, q'$ are random. For every block that the filter keeps, conditioned on $\mathcal{E}$ and on D_{q-1}, the words $u_q, \dots, u_{q'}$ are independent and uniform on $\mathbb{Z}_{16}$, and*

$$\mathbb{E}\big[(-1)^{D_{q'} - D_{q-1}} \mid \mathcal{E}\big] = \prod_{r=q}^{q'} c_r(t'_r).$$

In particular, if $R \le 4$, then $\mathbb{E}[Z \mid \mathcal{E}] = s_1 \prod_{r=2}^{R} c_r(t'_r)$ with the first-round sign $s_1 = (-1)^{\Delta_1(0)}$.

The uniform inputs come from the contracting S-boxes, and inside a segment, reduced-round FRAST therefore satisfies the ideal multi-round model for every filter on the expanding S-boxes. The first-round sign is $s_1 = b^{(1)}_{v'_i} b^{(1)}_{v'_j}$ with $v'_i = \text{rk}_i^{(1)} \bmod 8$ and $v'_j = \text{rk}_j^{(1)} \bmod 8$. Across a fixed round, the input of the next random round is not uniform, and we follow the Fourier coefficients of the state. For $\alpha, \ell \in \mathbb{Z}_{16}$, let

$$\Psi_r(\alpha, \ell) = \mathbb{E}\big[\zeta^{\alpha x_1^{(r)} + \ell x_2^{(r)}} (-1)^{D_r} \mid \mathcal{E}\big], \tag{14}$$

where the expectation is over the contracting S-boxes of the random rounds, and $\mathbb{E}[Z \mid \mathcal{E}] = \Psi_R(0,0)$. For a function $S\colon \mathbb{Z}_{16} \to \mathbb{Z}_{16}$ used in round r, let F and Δ be the functions of equation (7) with S in place of $S_{\text{erf}}^{(r)}$, and let

$$m_S(\ell; \beta) = \frac{1}{16} \sum_{z \in \mathbb{Z}_{16}} \zeta^{\ell F(z)} (-1)^{\Delta(z)} \zeta^{-\beta z}. \tag{15}$$

We write m_r for $S = S_{\text{erf}}^{(r)}$ and m_r^{fix} for $S = S_{\text{fix}}$. The substitution in the proof of Theorem 4 gives $m_r(0;0) = c_r(t'_r)$ and $m_r^{\text{fix}}(0;0) = c_{\text{fix}}(t'_r)$. The fixed S-box enters

through its correlations $C^{S_{\text{fix}}}(\alpha, \mu)$, which satisfy $C^{S_{\text{fix}}}(\beta, 0) = \mathbf{1}[\beta = 0]$ because S_{fix} is a permutation. For every function S, the identity $C^{S}(0, \mu) = \mathbf{1}[\mu = 0]$ holds, since the sum of a nontrivial character over $\mathbb{Z}_{16}$ vanishes.

Theorem 5 (Transfer Matrices Across Fixed Rounds). *Let the XOF be a random oracle, fix $\mathcal{E}$, and let $1 \le R \le 40$. Start with $\psi \in \mathbb{C}^{16}$ given by $\psi(\ell) = \zeta^{\ell(\sigma_0 + F_1(0))} s_1$, where $\sigma_0 = \sum_{l=2}^{32} \mathrm{ic}_l$. For $r = 2, \dots, R$ in increasing order, replace $\psi(\ell)$ by $m_r(\ell; 0)\, \psi(\ell)$ if the rounds $r-1$ and r are random. If round $r < R$ is fixed, replace ψ by $T_r \psi$, where*

$$T_r[\ell, \ell''] = \sum_{\beta \in \mathbb{Z}_{16}} m_{r+1}(\ell; \beta)\, C^{S_{\text{fix}}}(\beta, \ell'' - \ell)\, \zeta^{(\ell'' - \ell)\mathrm{rk}_1^{(r)}}\, m_r^{\text{fix}}(\ell''; -\beta),$$

and skip round $r+1$. At the end, $\mathbb{E}[Z \mid \mathcal{E}] = \psi(0)$ if round R is random, and $\mathbb{E}[Z \mid \mathcal{E}] = c_{\text{fix}}(t'_R)\, \psi(0)$ if round R is fixed.

Let $1, \dots, B$ be the segments that contain a round $r \le R$. Expanding the products of Theorem 5 by the value ℓ_b of ℓ in each segment b writes $\Psi_R(0, 0)$ as a sum of terms $\mathcal{P}(\ell_1, \dots, \ell_B)$. The *main trail* is the term $\mathcal{P}(0, \dots, 0)$, and a *side trail* is any other term, as drawn in Figure 2. With the random rounds $2 \le r \le R$ in the first product and the fixed rounds $r \le R$ in the second, the main trail is the product of the correlations of the iterated trail of Bak et al. and the first-round sign,

$$\mathcal{P}(0, \dots, 0) = s_1\, \Gamma \prod_r c_r(t'_r), \qquad \Gamma = \prod_r c_{\text{fix}}(t'_r). \tag{16}$$

4.2 Side Trails and Usable Branch Pairs

A guess weights each block by a function W of the parity patterns of its random rounds, and the filter keeps a block according to the same patterns and may couple the rounds. We write $\mathcal{G}_{-r}$ for the expanding S-boxes of the random rounds other than r, ν_r for the conditional distribution of $b^{(r)}$ given $\mathcal{G}_{-r}$ and the filter, and $d_i = \mathrm{rk}_i^{(r)} \bmod 8$ and $d_j = \mathrm{rk}_j^{(r)} \bmod 8$ for a random round r. Condition (F_r) holds if ν_r and W are invariant under $b^{(r)} \mapsto -b^{(r)}$ for every value of $\mathcal{G}_{-r}$. Condition (S_r) holds if $W = W_r W'_r$ with W_r a function of $b^{(r)}$ and W'_r independent of $b^{(r)}$, and if $\mathcal{L}_r(u) = \mathbb{E}_{\nu_r}\big[W_r\, b^{(r)}_{u+d_i} b^{(r)}_{u+d_j}\big]$ satisfies $\mathcal{L}_r(u+2) = \mathcal{L}_r(u)$ for all $u \in \mathbb{Z}_8$ and every value of $\mathcal{G}_{-r}$.

Theorem 6 (Side Trails Orthogonal to the Weights). *Let the XOF be a random oracle, let the filter depend only on the parity patterns of the random rounds, and let W be a function of these parity patterns. A side trail $\mathcal{P}$ with $\ell_b \notin \{0, 8\}$ for some segment b satisfies $\mathbb{E}[W\mathcal{P} \mid \Phi] = 0$. Any other side trail $\mathcal{P}$ satisfies $\mathbb{E}[W\mathcal{P} \mid \Phi] = 0$ if (F_r) holds for a random round r of a segment with $\ell_b = 8$, or if (S_{f+1}) holds for the return round $f+1$, where f is the fixed round after the last segment b^* with $\ell_{b^*} = 8$. If every side trail satisfies one of these conditions, then $\mathbb{E}[WZ \mid \Phi] = \Gamma\, \mathbb{E}\big[W s_1 \prod_r c_r(t'_r) \;\big|\; \Phi\big]$.*

Every filter and weight of this paper satisfies both conditions in every random round $r \ge 2$, and the first-round weight $s_1(o_1)$ satisfies (F_1). The uniform distribution of an unfiltered round and the ELP filters, which depend on $b^{(r)}$ through ELP_8, are invariant under $b^{(r)} \mapsto -b^{(r)}$ and under the shift $x \mapsto x+1$ of indices. Both the weights $c_r(h_r)$ and $\operatorname{sgn} c_r(h_r)$ are even and invariant under the shift, so $\mathcal{L}_r$ is constant. The filters that keep the weak S-boxes or the constant and the alternating patterns are invariant under $b^{(r)} \mapsto -b^{(r)}$ and make $b^{(r)}_{u+d_i} b^{(r)}_{u+d_j}$ independent of u. Hence (S_r) holds for every weight W_r and (F_r) for every even one. A wrong class therefore has zero expected score in a random round on which the filter does not depend and whose weight is $c_r(h_r)$, and in round 1 without a filter, as derived from Theorem 6 in Appendix B.3.

Side trails that avoid a fixed round with $c_{\mathrm{fix}}(t'_r) = 0$ exist for weights that depend on the key and are not invariant under the shift of indices. By Theorem 6, the expected score of the right guess is proportional to Γ, which vanishes when $c_{\mathrm{fix}}(t'_r) = 0$ in a fixed round of the approximation. We call a branch pair *usable* for an approximation when $c_{\mathrm{fix}}(t'_r) \neq 0$ in each of its fixed rounds. The autocorrelation of the fixed S-box is

$$\big(c_{\mathrm{fix}}(0), c_{\mathrm{fix}}(1), \ldots, c_{\mathrm{fix}}(15)\big) = \big(1, 0, -\tfrac{1}{2}, -\tfrac{1}{4}, 0, \tfrac{1}{4}, \tfrac{1}{4}, 0, -\tfrac{1}{2}, 0, \tfrac{1}{4}, \tfrac{1}{4}, 0, -\tfrac{1}{4}, -\tfrac{1}{2}, 0\big). \quad (17)$$

It is real and even and vanishes at six of the 16 key differences. For a uniformly random master key, the key differences $t'_1, \ldots, t'_{40}$ of every branch pair are independent and uniform on $\mathbb{Z}_{16}$. Hence a fixed round has $c_{\mathrm{fix}}(t'_r) = 0$ with probability $\frac{3}{8}$. All of f given fixed rounds have $c_{\mathrm{fix}}(t'_r) \neq 0$ with probability $(\frac{5}{8})^f$. This holds for one branch pair and says nothing about the joint distribution of the key differences of two pairs, so Section 7 evaluates the best of the 465 branch pairs on sampled master keys.

5 Key Guesses, the Statistic and Its Error Bounds

In this section, we define the key guesses and the matched filter statistic of our attacks, give the cost of a branch pair and state when an attack succeeds. We then bound the probabilities of a false alarm and of a missed right guess for every fixed master key. For every fixed master key, the bound on false alarms is proved for variant A1 in the known-nonce model without a filter, and the bound on the success side covers all our attacks, including variant A2 and the filters of the chosen-nonce attacks. Against an ideal keystream, the false alarms of every attack are bounded, which gives a distinguisher with a proved advantage.

5.1 Key Guesses and the Statistic

An attack on a branch pair (i, j) guesses the key material of the weight and of the observation of Section 4, which we write as a triple $H = (o_1, h, o_R)$. The first-round guess $o_1 = \{v_i, v_j\}$ stands for $\{\mathrm{rk}_i^{(1)} \bmod 8, \mathrm{rk}_j^{(1)} \bmod 8\}$ and gives the sign $s_1(o_1) = b_{v_i}^{(1)} b_{v_j}^{(1)}$. It takes 29 values, that is, one for $v_i = v_j$, where $s_1(o_1) = 1$, and one for each of the 28 pairs with $v_i \neq v_j$. The class guess $h = (h_r)_r$ gives a class $h_r \in \{0, \ldots, 4\}$ of Theorem 2 for each of the x random rounds $r \geq 2$ of the approximation without a filter. Fixed rounds are not guessed, because their factors of Γ do not depend on the block, and statistics are compared by absolute value, since the sign of Γ is unknown.

In the chosen-nonce attacks, a filter may restrict a random round to the constant or the alternating parity patterns, and the guess has a parity bit for each of these w rounds. We write Φ_{c} for the filter that keeps only the constant patterns of a round. The filter Φ_{ca} keeps the constant and the alternating patterns, with one parity bit guessed, and a round restricted by Φ_{ca} is of the first kind and one restricted by Φ_{c} of the second kind. As drawn in Figure 2, variant A1 keeps the last round R, with $Z = (-1)^{D_R}$, and variant A2 strips it with a guess o_R of $\mathrm{rk}_1^{(R)}$ and the unordered pair $\{\mathrm{rk}_i^{(R)}, \mathrm{rk}_j^{(R)}\}$. This guess takes at most 2176 values, gives $x_1^{(R-1)}$ from the block and leaves the observation $Z = (-1)^{D_{R-1}}$ of $R - 1$ rounds. With every guess of all 465 branch pairs, counted by distinct statistics, the set $\mathcal{H}$ of guesses has $|\mathcal{H}| = 465 \cdot 29 \cdot 5^x \cdot n_o \cdot 2^w$ elements. A pair (o_1, h) has $n_o = 1$ statistic in variant A1 and at most 1921 in variant A2, where the 256 last-round guesses with $\mathrm{rk}_i^{(R)} = \mathrm{rk}_j^{(R)}$ give the statistic of variant A1.

Let Z_u be the observation of a block u under H. The matched filter weight is $W_u = s_1(o_1) \prod_r c_r^u(h_r)$ over the random rounds $r \geq 2$ whose class is guessed, with the value $c_r^u(h_r)$ of the autocorrelation of round r in block u on the class $[h_r]$. In the chosen-nonce attacks, W_u also has a sign for each parity bit, defined in Section 6.1. For N blocks,

Table 2: Cost of one round of the approximation. The cases are the first-round guess o_1, the last-round guess o_R of variant A2, the class of the key difference of an unfiltered random round, and the chosen-nonce filters Φ_{c} and Φ_{ca}.

Round	Case	Probability	Data	Offline
first	o_1	—	0	0
last	o_R	—	0	0
random	$t'_r \in [0]$	1/8	0	0
random	$t'_r \in [4]$	1/8	2	0
random	$t'_r \in [1] \cup [2] \cup [3]$	3/4	3	0
random	Φ_{c}	—	0	7
random	Φ_{ca}	—	0	6
fixed	$c_{\mathrm{fix}}(t'_r) = \pm 1$	1/16	0	0
fixed	$c_{\mathrm{fix}}(t'_r) = \pm \frac{1}{2}$	3/16	2	0
fixed	$c_{\mathrm{fix}}(t'_r) = \pm \frac{1}{4}$	3/8	4	0
fixed	$c_{\mathrm{fix}}(t'_r) = 0$	3/8	∞	0

the statistic of H is

$$\xi(H) = \frac{\Sigma_1}{\sqrt{\Sigma_2}}, \qquad \Sigma_1 = \sum_{u=1}^{N} W_u Z_u, \qquad \Sigma_2 = \sum_{u=1}^{N} W_u^2, \tag{18}$$

with $\xi(H) = 0$ when $\Sigma_2 = 0$. Since $Z_u^2 = s_1(o_1)^2 = 1$, the denominator depends only on h.

Lemma 2 (The Right Guess). *Fix the master key. Let H be the right guess of variant A1, without a filter or with a filter in which each restricted random round $r \geq 2$ keeps either the constant patterns or the constant and the alternating patterns, while the other random rounds stay unfiltered. Then $\mathbb{E}[Z_u \mid W_u] = \Gamma W_u$ for every block u. Given the weights $(W_u)_u$, the observations are independent, and*

$$\mathbb{E}\big[\xi(H) \;\big|\; (W_u)_u\big] = \Gamma\sqrt{\Sigma_2}, \qquad \mathrm{Var}\big[\xi(H) \;\big|\; (W_u)_u\big] = 1 - \Gamma^2 \frac{\sum_u W_u^4}{\Sigma_2}.$$

The same holds in variant A2 with the right last-round guess, with $R-1$ rounds in place of R.

By Lemma 2, proved in Appendix B.4, the mean $\mu = |\Gamma|\sqrt{\Sigma_2}$ of the right guess satisfies $\mathbb{E}[\mu^2] = N \cdot \mathrm{ELP}_{\mathrm{pair}}$ with the expected linear potential $\mathrm{ELP}_{\mathrm{pair}} = \Gamma^2 \prod_r \mathbb{E}\big[c_r(t'_r)^2\big]$ of the branch pair, whose cost is $-\log \mathrm{ELP}_{\mathrm{pair}}$. The cost is the sum of the costs per round in Table 2. There, Data is $-\log$ of the factor of the round in $\mathrm{ELP}_{\mathrm{pair}}$ and Offline is log of the XOF calls of its filter per kept block, and both are in bits. Probability is over a uniform key difference, which a uniform master key provides by Section 4.2, and a dash marks a cost independent of it.

5.2 Success and the Error Bounds

The *score* of a pair (o_1, h) is the largest $|\xi(o_1, h, o_R)|$ over its last-round guesses, and $|\xi(o_1, h)|$ in variant A1. An attack on a branch pair *succeeds* when the score of the right pair is at least that of every other pair (o_1, h) of the branch pair, and for a threshold θ, a *false alarm* is a wrong guess with zero expected score and $|\xi| \geq \theta$. In variant A1 without a filter, every wrong pair (o_1, h) has zero expected score by Section 4.2, and the attack succeeds if the right guess has $|\xi| \geq \theta$ and no false alarm occurs.

Proposition 1 (Rank-One Probability and Threshold in the Gaussian Model). *Let the statistics of the guesses in $\mathcal{H}$ be independent, let ξ of the right guess be normal with mean μ and standard deviation $\sigma > 0$, and let every other ξ be standard normal. Write*

$Q(y)$ for the upper tail of the standard normal distribution and φ for its density. Then the right guess has a larger $|\xi|$ than every other guess with probability

$$P_{\text{rank1}} = \int_{-\infty}^{\infty} \frac{1}{\sigma}\,\varphi\Big(\frac{y-\mu}{\sigma}\Big)\big(1-2Q(|y|)\big)^{|\mathcal{H}|-1}\,dy. \tag{19}$$

For a threshold θ, the right guess has $|\xi| \ge \theta$ with probability $Q\big(\frac{\theta-\mu}{\sigma}\big) + Q\big(\frac{\theta+\mu}{\sigma}\big)$, and the expected number of other guesses with $|\xi| \ge \theta$ is $2(|\mathcal{H}|-1)\,Q(\theta)$.

Proof. Given the value y of ξ for the right guess, each other guess has $|\xi| < |y|$ with probability $1-2Q(|y|)$, independently, and integrating over y gives equation (19). The right guess has $\xi \ge \theta$ with probability $Q(\frac{\theta-\mu}{\sigma})$ and $\xi \le -\theta$ with probability $Q(\frac{\theta+\mu}{\sigma})$, and each other guess has $|\xi| \ge \theta$ with probability $2Q(\theta)$. □

The Gaussian model, which is an estimate, gives the right guess the moments of Lemma 2 and treats the wrong guesses as independent standard normal variables. Their sums Σ_1 are uncorrelated across classes by Theorem 2(i) except for last-round guesses of one pair. The cost model takes $\sigma = 1$, $\mu = \sqrt{N \cdot \mathrm{ELP}_{\text{pair}}}$ and $\theta = \sqrt{2\ln|\mathcal{H}|}$, and the right guess passes with probability at least p_s when $N \cdot \mathrm{ELP}_{\text{pair}} \ge \kappa = (\theta + Q^{-1}(1-p_s))^2$. With $p_s = 0.9$, $Q^{-1}(1-p_s) = 1.28$. At the thresholds $\sqrt{2\ln|\mathcal{H}|}$ of 10.2 to 10.9 for $|\mathcal{H}|$ between 2^{75} and 2^{85}, $Q(10.9)$ is about 2^{-90}, whereas the central limit theorem holds for ξ only up to an error of order $N^{-1/2}$. We therefore bound the expected number of false alarms for every fixed master key. If it is at most the target $\varepsilon_{\text{fa}} = 0.1$, Markov's inequality bounds the probability of a false alarm by ε_{fa}, and the success probability is at least that of the right guess passing minus ε_{fa}. A rigorous threshold θ^* is the smallest θ at which a proved bound meets the target, and the Gaussian threshold is $\theta_G = Q^{-1}(\varepsilon_{\text{fa}}/(2|\mathcal{H}|))$.

Theorem 11, in Appendix B.4, bounds $\Pr[|\xi(H)| \ge \theta]$ for every wrong guess H of variant A1 in the known-nonce model without a filter and does not need $\mathbb{E}[Z_u \mid W_u] = 0$. A symmetry of the wrong guesses gives a bound close to the Gaussian tail, and Theorem 7 uses Theorem 11 only for a set $\mathcal{H}_{13}$ of wrong guesses. A guess of $\mathcal{H}_{13}$ has its first-round guess in the right class, and in each of its wrong rounds, the key difference and the guess have the classes [1] and [3]. The proofs of Theorem 7 and of Proposition 2 and Theorem 8 below are in Appendix B.4, and the first uses the inequality of Pinelis [Pin12]. Two extensions of Theorem 7, to variant A2 and to the chosen-nonce filters, are not proved. Its proof and an exact computation in Section 7.4 support them, and Appendix C.2 states them. The bound on the success side without the Gaussian model uses the ratios $\rho = \mathbb{E}[c^4]/\mathbb{E}[c^2]$ of a uniformly random negacyclic S-box, which are

$$\big(\rho_{[0]}, \rho_{[1]}, \rho_{[2]}, \rho_{[3]}, \rho_{[4]}\big) = \big(1, \tfrac{11}{32}, \tfrac{7}{16}, \tfrac{11}{32}, \tfrac{5}{8}\big),$$

and ρ_r is the ratio of the class of t'_r in round r.

Theorem 7 (False Alarms of Variant A1). *Fix the master key, and consider variant A1 in the known-nonce model without a filter.*

(i) If $\Gamma \ne 0$, then a wrong guess H has a distribution of $W_u Z_u$ that is symmetric about zero if and only if $H \notin \mathcal{H}_{13}$. If $\Gamma = 0$, then every wrong guess has this property.

(ii) Every wrong guess $H \notin \mathcal{H}_{13}$ satisfies, for all N and all $y > 0$,

$$\Pr\big[|\xi(H)| \ge y\big] \le B_{\text{P}}(y) = 2\Big(Q(y) + \frac{C_{\text{P}}\,\varphi(y)}{9+y^2}\Big), \tag{20}$$

where $C_{\text{P}} = 5\sqrt{2\pi e}\,\big(1-2Q(1)\big) = 14.1069$.

(iii) The expected number of false alarms at the threshold θ is at most $(|\mathcal{H}|-|\mathcal{H}_{13}|)\,B_{\mathrm{P}}(\theta)+\sum_{H\in\mathcal{H}_{13}} B_{\mathrm{T}}(H)$, where $B_{\mathrm{T}}(H)$ is the bound of Theorem 11 for H, minimised over the parameters of that theorem.

Proposition 2 (Success Side). *Fix the master key, let H be a right guess as in Lemma 2, and let $\mu_W = |\Gamma|\sqrt{\Sigma_2}$. Given the weights $(W_u)_u$, if $\mu_W \geq \theta$, then*

$$\Pr\big[|\xi(H)| < \theta \;\big|\; (W_u)_u\big] \leq \exp\big(-(\mu_W-\theta)^2/2\big).$$

For $V = N\,\mathbb{E}[W_u^2]$ and $\delta \in (0,1)$, the denominator satisfies $\Pr[\Sigma_2 \leq (1-\delta)V] \leq \exp\big(-\delta^2 V/(2\prod_r \rho_r)\big)$, with the product over the rounds whose class is guessed.

Given the weights, the right guess therefore passes with probability at least $1-\exp(-(\mu_W-\theta^*)^2/2)$, and for every fixed master key the success probability is at least this bound at the realised weights minus $\varepsilon_{\mathrm{fa}}$ on average over the blocks. A success probability of 0.9 needs $\mu_W-\theta^* \geq \sqrt{2\ln 10} = 2.146$, against 1.28 in the Gaussian model. The *strict condition* replaces κ by $\kappa''/(1-\delta)$ with $\kappa'' = (\theta^*+2.146)^2$, where δ makes the lower tail of Σ_2 at most 0.01, and Section 7.2 and Appendix C report the reach under it beside the Gaussian condition. The success probability is then at least $0.9-0.01-\varepsilon_{\mathrm{fa}}$.

In the *ideal world*, the keystream block of every XOF input is uniform in $\mathbb{Z}_{16}^{32}$ and independent of the other blocks and of the S-boxes, whereas the *real world* is FRAST under a uniformly random master key. On N blocks and for a set $\mathcal{H}$ of guesses, the *distinguisher* D_θ outputs FRAST when some guess $H \in \mathcal{H}$ has $|\xi(H)| \geq \theta$, and its *advantage* $\mathrm{Adv}(D_\theta)$ is the probability of this output in the real world minus that in the ideal world. The *success bound* p_{hit} is the bound of Proposition 2 on the probability that a right guess passes θ, with the lower tail of Σ_2 subtracted.

Theorem 8 (False Alarms in the Ideal World). *Let the N XOF inputs be distinct and chosen without the keystream, as functions of the S-boxes and of independent randomness, and let the filter and the weights W_u of every guess depend only on the S-boxes of the blocks and on the guess. Parts (i) and (ii) hold in the ideal world for every guess H of variant A1 or A2, with any last-round guess o_R.*

(i) Given the weights, the observations Z_u of H are independent and uniform on $\{\pm 1\}$.

(ii) For all $\theta > 0$, $\Pr\big[|\xi(H)| \geq \theta \;\big|\; (W_u)_u\big] \leq 2\exp(-\theta^2/2)$, and for every set $\mathcal{H}$ of guesses, $\Pr\big[\max_{H\in\mathcal{H}} |\xi(H)| \geq \theta\big] \leq 2|\mathcal{H}|\exp(-\theta^2/2)$.

(iii) Let the filter be as in Lemma 2, and let D_θ output FRAST when $\max_{H\in\mathcal{H}} |\xi(H)| \geq \theta$ on N blocks and the ideal world otherwise. For a set of master keys of density d in each of which some branch pair has its right guess in $\mathcal{H}$ with $\mathbb{E}[\mu_W^2] = N\cdot \mathrm{ELP}_{\mathrm{pair}} \geq m^2$, the advantage of D_θ over a uniform master key satisfies $\mathrm{Adv}(D_\theta) \geq d\,p_{\mathrm{hit}} - 2|\mathcal{H}|\exp(-\theta^2/2)$. Here $p_{\mathrm{hit}} = 1-\exp\big(-(m\sqrt{1-\delta}-\theta)^2/2\big)-\exp\big(-\delta^2 V/(2\prod_r \rho_r)\big)$ is the bound of Proposition 2 at the smallest $V/\prod_r \rho_r$ over the keys of the set, for any $\delta \in (0,1)$ with $m\sqrt{1-\delta} > \theta$.

Theorem 8 bounds the advantage of telling FRAST from an ideal keystream, while, on full FRAST, class recovery and key recovery stay in the cost model. Its part (ii) needs neither independence between the guesses nor a symmetry of the wrong guesses, because every observation of the ideal world is a fair sign given the weights. This bound therefore holds for variant A2 and for the chosen-nonce filters, which lie outside the hypotheses of Theorem 7. We apply part (iii) to full FRAST in Section 7.3, with the weak master keys of the cost model as the set of keys.

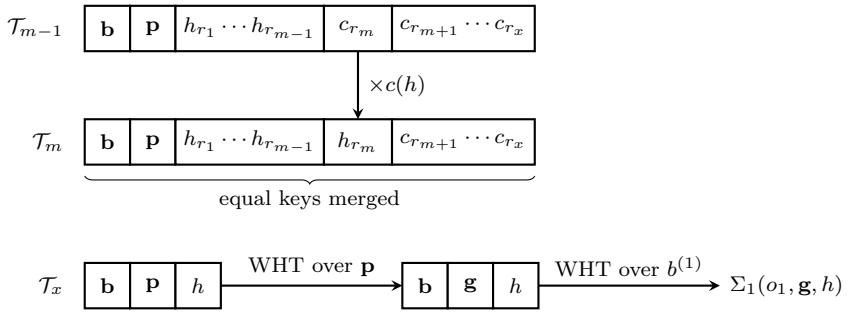

Figure 3: The keys of the tables $\mathcal{T}_{m-1}$ and $\mathcal{T}_m$ of CONTRACT, in which c_{r_l} is the class vector of round r_l that a key holds, and the Walsh–Hadamard transforms WHT from the keys of $\mathcal{T}_x$ to Σ_1 in variant A1.

6 The Online Algorithm Contract and Key Recovery

In this section, we describe the online algorithm CONTRACT for the statistics of Section 5 and bound the sizes of its tables, its time and its storage. We then recover the master key from the accepted guesses of all branch pairs through the linear key schedule, layer by layer over the bits of its words. CONTRACT computes every statistic of a branch pair, also in variant A2 and in the chosen-nonce attacks. With prefix enumeration, its storage is linear in the data for the attacks of this paper. The proofs of this section are in Appendix B.5.

6.1 The Algorithm Contract

A restricted round is a random round whose parity pattern a chosen-nonce filter restricts to the constant or the alternating patterns, and $r_1 < \cdots < r_x$ are the random rounds whose class is guessed. CONTRACT reads four data of every block u: the sign $Z_u^{(R)} = (-1)^{D_R}$ of equation (10), the outer cell $\mathbf{b}_u \in \mathcal{B}$, the vector $\mathbf{p}_u \in \mathbb{Z}_2^w$ of restricted bits and the class vectors of the rounds $r_1, \ldots, r_x$. In variant A1, the sign is the observation Z_u. The outer cell holds the parity pattern of round 1 and, in variant A2, the data of round R that the last-round guess needs, and there are $|\mathcal{B}| = 2^8$ outer cells in variant A1 and 2^{16} in variant A2.

The restricted bits record the restricted rounds with an alternating pattern, and the parity bits $\mathbf{g} \in \mathbb{Z}_2^w$ of a guess contribute the sign $(-1)^{\mathbf{p}_u \cdot \mathbf{g}}$ to W_u. The class vector $c_r^u = \big(c_r^u(0), \ldots, c_r^u(4)\big)$ lists the values of the autocorrelation of round r in block u on the five classes. It takes 18 values, and its entry $c(0)$ is always 1. For a round without a filter, the other entries are nonzero with the probabilities below, which give the mean $\bar{n}$ of the number $n(c_r^u)$ of nonzero entries:

$$\big(\Pr[c(1) \neq 0], \Pr[c(2) \neq 0], \Pr[c(3) \neq 0], \Pr[c(4) \neq 0]\big) = \big(\tfrac{29}{64}, \tfrac{13}{32}, \tfrac{29}{64}, \tfrac{5}{8}\big),$$
$$\bar{n} = \mathbb{E}\big[n(c_r^u)\big] = \tfrac{47}{16}.$$

Expanding every block over the class guesses with nonzero values costs $N\bar{n}^x$ multiply-adds for one branch pair in expectation, and an algorithm that outputs every statistic needs at least $|\mathcal{H}|$ steps. The classes cannot be tested round by round, because a guess with a wrong class in one round has zero expected score by Section 4.2. CONTRACT relies instead on the small value of $\bar{n}$, multiplies the classes into the values one round at a time and merges the blocks that agree on the rounds not yet treated. Algorithm 1 lists its steps, with its tables drawn in Figure 3. The keys and the values of these tables are described in Appendix B.5, which also proves Proposition 3 below.

Algorithm 1 The online algorithm CONTRACT for one branch pair.

Input: the N blocks u with $Z_u^{(R)}$, $\mathbf{b}_u$, $\mathbf{p}_u$ and the class vectors $c^u_{r_1}, \dots, c^u_{r_x}$
Output: Σ_1 of every guess and Σ_2 of every class guess

▷ in variant A2 with a random round R, once for each of the 16 guesses of $\mathrm{rk}_1^{(R)}$

$\mathcal{T}_0 \leftarrow$ the sums of $Z_u^{(R)}$ over the blocks u with equal keys ▷ step 1
for $m = 1, \dots, x$ **do** ▷ step 2
  for all entries of $\mathcal{T}_{m-1}$, with the value v and the class vector c of round r_m in the key **do**
    for all classes h with $c(h) \neq 0$ **do**
      add $v\, c(h)$ to the entry of $\mathcal{T}_m$ whose key has h in place of c
    end for
  end for
end for
for all class guesses h **do** ▷ step 3
  a Walsh–Hadamard transform over $\mathbf{p}$ gives the 2^w guesses $\mathbf{g}$ of the parity bits
  a Walsh–Hadamard transform over the parity pattern $b^{(1)}$, read at the mask $e_{v_i} \oplus e_{v_j}$ of every $o_1 = \{v_i, v_j\}$ with the unit vectors e_v of $\mathbb{Z}_2^8$, gives the sums over the outer cells that the last-round step turns into Σ_1, or $\Sigma_1(o_1, \mathbf{g}, h)$ itself in variant A1
  if variant A2 with a fixed round R **then**
    for all 16 guesses of $\mathrm{rk}_1^{(R)}$ **do**
      sum the values by the word $x_1^{(R-1)}$ that the guess gives
      map these 16 sums to the 136 guesses of $\{\mathrm{rk}_i^{(R)}, \mathrm{rk}_j^{(R)}\}$ by a sign matrix
    end for
  else if variant A2 with a random round R **then**
    a Walsh–Hadamard transform over the parity pattern of round R, shifted by $x_1^{(R-1)}$, gives the guesses of $\{\mathrm{rk}_i^{(R)}, \mathrm{rk}_j^{(R)}\}$ modulo 8 at the 29 masks $e_{v_i} \oplus e_{v_j}$
  end if
end for
Σ_2 of every class guess: steps 1 and 2 with every sign replaced by 1, the squared class values, one outer cell and no restricted bits ▷ step 4

Proposition 3 (Sizes of the Tables). *Let the XOF be a random oracle, and let the filter depend only on the restricted rounds. Then, for $m = 0, \dots, x$,*

$$\mathbb{E}|\mathcal{T}_m| \le \min\bigl(N\bar{n}^m,\ |\mathcal{B}|\, 2^w 5^m 18^{x-m}\bigr).$$

The two bounds of Proposition 3 cross at

$$m^* = \frac{\log|\mathcal{B}| + w + x\log 18 - \log N}{\log(18\bar{n}/5)}, \qquad N\bar{n}^{m^*} = N^{0.543}\bigl(|\mathcal{B}|\, 2^w 18^x\bigr)^{0.457}, \tag{21}$$

and the expected size of every table is at most $N\bar{n}^{m^*}$. We estimate the time of step 2 in a model that charges $\bar{n}$ multiply-adds to every entry of $\mathcal{T}_{m-1}$. The model assumes that the class vector of round r_m in a key of $\mathcal{T}_{m-1}$ has the distribution of that of a block, and Proposition 5 bounds the work without this assumption. For all branch pairs in variant A1, the model gives the online time

$$T_{\mathrm{on}} \approx 465 \cdot 2^5 \cdot \bigl(c_{\mathrm{on}} N\bar{n}^{m^*} + 5^x 2^w 2^{11}\bigr)$$

word-level multiply-adds. Variant A2 replaces 2^{11} by at most $2^{21.1}$, and with a random last round it multiplies the time by 16. The constant c_{on} of step 2, which is at most 5.6 in the model, is fixed in Section 7.2, and Appendix C.1 derives the model and the cost of step 3.

The storage of Contract is that of its largest table, which can exceed N, and prefix enumeration splits the computation by the classes of the first j rounds. For $1 \le j \le x$ and each of the 5^j prefixes $(h_{r_1}, \ldots, h_{r_j})$, it runs steps 2 and 3 of Contract for the rounds $r_{j+1}, \ldots, r_x$ on the blocks with $c^u_{r_l}(h_{r_l}) \neq 0$ for all $l \le j$, starting from the values $Z^{(R)}_u \prod_{l \le j} c^u_{r_l}(h_{r_l})$. The tables of all prefixes at a step $m \ge j$ together form $\mathcal{T}_m$ and give the statistics of Contract. After one sort of the blocks by their class vectors in the rounds $r_1, \ldots, r_j$, gathering the blocks of all prefixes takes $N\bar{n}^j$ accesses in expectation. With $j = \lfloor m^* \rfloor$, the gathering costs at most $N\bar{n}^{m^*}$ accesses, the time is that of Contract up to this term, and the storage is $O(N)$ words for the attacks of this paper.

The online time has the two units of Section 2.4, which are the bit operations of Bak et al. and the word-level multiply-adds of Contract with a margin 2^5 for sorting, hashing, merges and constants. The model omits the cost of random access to a memory of up to 2^{64} words. To compare with the mixed unit of Bak et al., we charge one word-level multiply-add on 64-bit words as about $2^{6.6}$ bit operations, by convention, so 2^t word-level multiply-adds with the margin are $2^{t+6.6}$ bit operations in the main unit.

6.2 Key Recovery

Key recovery runs the class recovery of variant A1 on all 465 branch pairs with the same blocks and the steps of Algorithm 2. An accepted right pair constrains the round keys modulo 8. Its first-round guess gives $\{\mathrm{rk}^{(1)}_i \bmod 8, \mathrm{rk}^{(1)}_j \bmod 8\} = \{v_i, v_j\}$, or $\mathrm{rk}^{(1)}_i \equiv \mathrm{rk}^{(1)}_j \pmod 8$ for the guess with $v_i = v_j$, and its classes give $t'_r \equiv \pm h_r \pmod 8$ by Theorem 2(i). By equation (5), every round key word is a linear form in the 64 words of k over $\mathbb{Z}_{16}$. Since the parity patterns of the random rounds have period 8 by equation (4), the guesses depend on these forms only modulo 8, and the bit 3 of the key words does not enter the constraints.

Since $\mathbb{Z}_{16}$ is not a field, we solve the constraints layer by layer over the bits of $k = \sum_{e=0}^{3} 2^e k_{[e]}$, where $k_{[e]} \in \mathbb{Z}_2^{64}$ holds the bits e of the 64 key words. A linear form $y = \langle \mathbf{f}, k \rangle$ with $\mathbf{f} \in \mathbb{Z}_{16}^{64}$ satisfies $y \equiv \langle \mathbf{f}, k \bmod 2^e \rangle + 2^e \langle \mathbf{f}, k_{[e]} \rangle \pmod{2^{e+1}}$. Given $k \bmod 2^e$, the bit e of y is therefore affine over $\mathbb{Z}_2$ in $k_{[e]}$, with the coefficients $\mathbf{f} \bmod 2$. The values of $k \bmod 8$ that satisfy the constraints are the solutions of the layers 0 to 2, whose equations are checked exhaustively in Appendix C.3, while layer 3 uses the fixed rounds. Once $k \bmod 8$ is known, the bit 3 of t'_r fixes the sign of $c_{\mathrm{fix}}(t'_r)$ in every fixed round with $c_{\mathrm{fix}}(t'_r) \neq 0$, by equation (17). Since $\mathbb{E}\big[\Sigma_1 \,\big|\, (W_u)_u\big] = \Gamma\Sigma_2$ by Lemma 2, the sign of Σ_1 of an accepted pair gives one equation over $\mathbb{Z}_2$ on $k_{[3]}$. These equations have rank at most 30, and the signs alone leave at least 2^{34} values of $k_{[3]}$, so the size of Σ_1 adds equations of a larger rank through the estimate $\hat{g}$ of $|\Gamma|$, with the tolerance of Appendix B.5.

If a set of accepted pairs contains a wrong pair, then the true $k \bmod 8$ violates a constraint of the set. Since the layers 0 to 2 impose every accepted constraint, a wrong accepted pair removes the true $k \bmod 8$ from their solutions, and layer 3 never sees a wrong pair together with the true $k \bmod 8$. A wrong pair can therefore make the recovery fail, and the failure shows as a failed majority check or keys that all fail the keystream test. Without a filter, every wrong guess has zero expected score by Section 4.2, so a wrong pair is accepted only through a false alarm. For a threshold $\theta \ge \theta^*$ of Section 5.2, none occurs with probability at least $1 - \varepsilon_{\mathrm{fa}}$.

We write rank_3 for the rank of layer 3, which leaves $2^{64-\mathrm{rank}_3}$ values of $k_{[3]}$. Section 7 reports the recovery of a ten-round master key from real data and computes the rounds at which the data limit leaves enough branch pairs above the threshold. With N blocks, key recovery costs the online time of class recovery on all 465 branch pairs plus at most $2^{64-\mathrm{rank}_3}$ keystream tests for each solution $k \bmod 8$ of the layers 0 to 2, and $O(N)$ words of storage. Each keystream test is one evaluation of FRAST with 2^{12} word-level operations.

Algorithm 2 Key recovery.

Input: the blocks of one master key and a threshold θ
Output: the master key k, or failure

1: **for all** 465 branch pairs **do**
2: run Algorithm 1 in variant A1, and take the pair (o_1, h) of largest score
3: if its score is at least θ, accept it, and keep $\operatorname{sgn}\Sigma_1$ and $\hat{g} = |\Sigma_1|/\Sigma_2$
4: **end for**
5: **for all** accepted pairs, in decreasing $|\xi|$ **do**
6: add all its constraints to the layers 0 to 2 unless one becomes inconsistent
7: form layer $e = 1, 2$ from the solution of the layers below once they have rank 64
8: **end for**
9: if a pair was rejected, move single pairs accepted before the layers 0 to 2 had rank 64 to the end of the order, and keep the order that accepts the most pairs
10: if the majority check fails, that is, a layer has rank below 64 or the accepted pairs do not outnumber the rejected ones, repeat Lines 5 to 8 with random orders and keep the first order that passes the check
11: solve the layers 0 to 2 by elimination over $\mathbb{Z}_2$, enumerating several solutions
12: **for all** values of k mod 8 that satisfy the layers 0 to 2 **do**
13: **for all** accepted pairs, in decreasing $|\xi|$ **do**
14: allow the bits 3 of t'_r in its fixed rounds r that give Γ the sign of Σ_1 and a size within the tolerance of Appendix B.5 around $\hat{g}$
15: add the equations on $k_{[3]}$ of their constant combinations, unless inconsistent
16: **end for**
17: **for all** $k_{[3]}$ that satisfy these equations and give every pair allowed bits **do**
18: test k on the observations of a few blocks, then on whole keystream blocks
19: **return** k if it passes both tests
20: **end for**
21: **end for**
22: widen the tolerance of Appendix B.5 once, and repeat the loop over k mod 8
23: **return** failure

7 Attacks on Reduced-Round FRAST

In this section, we report the attacks that we executed on reduced-round FRAST and compute in a cost model how far they reach, up to weak master keys of full FRAST. The last subsection validates the models by experiments. Every executed attack is in the known-nonce model without a filter, on master keys drawn uniformly and expanded by the real key schedule. An attacker runs all 465 branch pairs on the same blocks, because it does not know which pair is best, and the data are set by the best pair. The class recoveries on ten and 15 rounds treated branch pairs chosen with knowledge of the key, and key recovery treated all 465 pairs.

In Tables 3 to 5, a row marked † is an estimate of the cost model and was not run. Data counts blocks per master key, T_{off} counts XOF calls, and T_{on}^{b} and T_{on}^{w} count all 465 branch pairs in the bit operations and in the word-level multiply-adds of Section 2.4. The accounting in bit operations alone, in the last column of Table 5, charges $2^{18.5}$ bit operations per XOF call. The storage is $O(N)$ words in every row with the prefix enumeration of Section 6.1, and Guesses counts the guesses of one branch pair, in distinct statistics for variant A2 on ten rounds. In Table 3, Model is the known-nonce or the chosen-nonce model, and w counts the restricted rounds of Section 5.1.

In an executed row, Success is the fraction of repeated runs in which the right guess ranked first. Its T_{on} and Fraction come from the cost model at the data of the run, because the runs recorded elapsed times and not operation counts. For a † row, Data is that of the

Table 3: Attacks on reduced-round FRAST. Target h is the class recovery of one branch pair and k the master key, and a dash marks a missing value.

Rounds	Target	Model	Variant	w	Guesses	Data	T_{off}	T_{on}^{b}	T_{on}^{w}	Success	Fraction
5	h	known	A2	0	$29 \cdot 5^3 \cdot 2176$	2^{15}	2^{15}	$2^{48.1}$	$2^{41.5}$	1.00	1.000
10	h	known	A1	0	$29 \cdot 5^7$	2^{22}	2^{22}	$2^{51.1}$	$2^{44.5}$	1.00	0.993
10	h	known	A2	0	$29 \cdot 5^7 \cdot 1921$	2^{22}	2^{22}	$2^{57.6}$	$2^{51.0}$	1.00	1.000
15	h	known	A1	0	$29 \cdot 5^{11}$	2^{29}	2^{29}	$2^{62.4}$	$2^{55.8}$	—	0.130
10	k	known	A1	0	$29 \cdot 5^7$	2^{28}	2^{28}	$2^{54.3}$	$2^{47.7}$	—	1.000
20†	h	known	A2	0	—	$2^{40.8}$	$2^{40.8}$	$2^{80.2}$	$2^{73.6}$	0.9	0.5
26†	h	known	A2	0	—	2^{64}	2^{64}	$2^{104.0}$	$2^{97.4}$	0.9	1.000
26†	h	known	A1	0	—	2^{64}	2^{64}	$2^{98.2}$	$2^{91.6}$	0.9	0.983
24†	k	known	A1	0	—	$2^{62.7}$	$2^{62.7}$	$2^{95.6}$	$2^{89.0}$	0.56	0.99
35†	h	chosen	A2	10	—	2^{64}	2^{124}	$2^{101.1}$	$2^{94.5}$	0.9	0.996
36†	h	chosen	A2	10	—	2^{64}	2^{124}	$2^{105.1}$	$2^{98.5}$	0.9	0.83
38†	h	chosen	A2	10	—	2^{64}	2^{124}	$2^{108.9}$	$2^{102.3}$	0.9	0.33

median key, that of 99% of the keys for key recovery, or 2^{64} at the data limit. In a † row of class recovery, Success is the probability in the Gaussian model that the right guess passes the threshold. In the row of key recovery on 24 rounds, Success is the smallest probability over the sampled keys that at least ten pairs pass when the pairs are treated as independent. Fraction is over uniform master keys under the real key schedule, and a dash marks a missing value.

7.1 Executed Attacks

The attack on five rounds uses variant A2, whose last-round guess strips the fixed round 5. Its approximation therefore consists of the random rounds 2 to 4, and every branch pair is usable. In 20 runs on random master keys and random branch pairs, the right guess ranked first in every run with 2^{15} blocks. In variant A1 on ten rounds, the approximation contains the random rounds 2 to 4 and 6 to 9 and the fixed rounds 5 and 10, which enter only through Γ. The right guess ranked first with 2^{22} blocks on the best branch pair of the master key, and in all 275 repeated runs of Appendix C.3. Variant A2 on ten rounds strips the fixed round 10. With 2^{22} blocks of one master key, the right guess ranked first on each of its 13 branch pairs of cost at most 14 bits, and with 2^{20} blocks twelve of the 13 right guesses ranked first.

In variant A1 on 15 rounds, the approximation contains eleven random rounds and the fixed rounds 5, 10 and 15. With 2^{29} blocks and the best branch pair of the master key, whose cost is 23 bits, the right guess ranked first. For this branch pair, the attack needs 2^{29} blocks, 2^{29} XOF calls and $O(N)$ words. The model of Section 6.1 gives its online time as 2^{47} word-level multiply-adds, or $2^{53.6}$ bit operations. At 2^{29} blocks, the cost model with κ for the guesses of all 465 pairs covers the master keys whose best usable pair costs at most 22.77 bits, which form the fraction 0.130 of Table 3 and do not include the attacked key. With κ for the guesses of the one pair that ran, the threshold is 23.08 bits, which covers the attacked key and a fraction 0.232 of the sampled keys.

Key recovery on ten rounds ran the class recovery of variant A1 on all 465 branch pairs with 2^{28} blocks of one master key and the threshold $\sqrt{2 \ln |\mathcal{H}|}$ of Section 5.2. It accepted 77 branch pairs, and the accepted pair of each was the right one. The layers 0 to 2 had rank 64 and gave k mod 8 uniquely. Layer 3 had rank 60 and left 16 values of $k_{[3]}$, and the keystream test kept one key, which was the master key. The recovery needs 2^{28} blocks, 2^{28} XOF calls, the online time of class recovery on the 465 pairs plus 16 keystream tests, and $O(N)$ words.

7.2 Reach in the Cost Model

The cost model computes the cost of every branch pair of a master key by Table 2 from the real key schedule. It assigns the data $N = \kappa\, 2^c$ to the best usable pair, where c is the cost of this pair and κ is the constant of the Gaussian model for the guesses of the attack. A master key is covered when $N \leq 2^{64}$, and since the costs are integers, the covered keys are those whose best pair costs at most an integer cap. The fractions are over 4000 sampled master keys with known nonces, 1500 with chosen nonces and 1000 for key recovery. Table 3 lists the executed attacks and the estimates of the cost model.

Every online time is that of the model of Section 6.1 with the margin 2^5. Its model constant $c_{\mathrm{on}} = 3$ is taken from the exact expectation of the table sizes on eight parameter sets with 15 to 40 rounds and restricted rounds of the first kind, where step 2 costs 2.3 to 3.4 times $N\bar{n}^{m^*}$ multiply-adds. The bound 5.6 of the model lies within the margin above it. Without the assumption, Proposition 5 bounds the work of step 2 by 5.99 to 6.70 times $N\bar{n}^{m^*}$ on these parameter sets, and the margin at $w = 5$ in Section 7.3 is then 1.28 bits.

With known nonces, the median master key needs $2^{40.8}$ blocks at 20 rounds in variant A2. At 26 rounds, it needs 2^{57} blocks in variant A2 and $2^{59.8}$ blocks in variant A1, and these medians depend on the key sample by about one bit. With 2^{64} blocks, the attack on 26 rounds covers every sampled key in variant A2 and 98.3% of the keys in variant A1. At 28 rounds, it covers 72% of the keys in variant A2. Under the strict condition of Section 5.2, with the extensions of Appendix C.2 for variant A2, the cap falls from 57 to 56 at 26 rounds and at 28 rounds in variant A1. This leaves 99.8% of the sampled keys covered at 26 rounds in variant A2 and 94.2% in variant A1.

With chosen nonces, the attacks on 35, 36 and 38 rounds use variant A2 with $w = 10$ restricted rounds that keep the constant or the alternating patterns. The filter costs six bits of offline time per restricted round by Table 2, so $N = 2^{64}$ blocks cost $T_{\mathrm{off}} = 2^{124}$ XOF calls. Within $T_{\mathrm{off}} \leq 2^{126}$, the attacks cover 99.6% of the sampled keys at 35 rounds and 33% at 38 rounds, and the strict condition under the extensions of Theorem 7 leaves these fractions unchanged. Their online times stay below T_{off} in both units.

Key recovery needs enough accepted branch pairs to determine k mod 8. In simulations of the layers with the cheapest pairs, ten pairs suffice from 15 rounds on, and ten rounds need 20 pairs. For the 2^{28} blocks of the executed recovery on ten rounds, the twentieth best usable pair must cost at most 22.10 bits, which holds for every sampled key. With 20 pairs, the median key needs $2^{24.90}$ blocks on ten rounds. At 24 rounds, the model gives $2^{60.7}$ blocks for half of the sampled keys and $2^{62.7}$ blocks for 99% of them.

With 0.9 per pair, and treating the pairs as independent, which their shared blocks and key schedule do not guarantee, at least ten pairs pass at 24 rounds with probability 0.98 for the median key and at least 0.56 for every sampled key. A probability of 0.99 per pair raises this minimum to 0.95 and adds 0.28 bits to the data of the median key. At 2^{56} word-level operations, the keystream tests stay at least 19 bits below the online time. Within 2^{64} blocks, the recovery covers 21.7% of the keys at 25 rounds and 0.7% at 26 rounds.

7.3 Full FRAST

For full FRAST, the chosen-nonce attack uses variant A2 with the fixed round 40 stripped, $N = 2^{64}$ blocks and w restricted rounds. A restricted round keeps either the constant or the alternating patterns and guesses one parity bit, or it keeps only the constant patterns, at the offline costs of Table 2. We write w_a and w_c for the numbers of restricted rounds of the two kinds. A round of the second kind adds no restricted bit and no entry to the tables of Contract, so 2^{w_a} replaces 2^w in the guesses and in Proposition 3. A master key is weak when its best branch pair has a cost of at most the cap 56, and the fraction d_w of weak master keys does not depend on the partition. We count an attack on a fraction

Table 4: Weak key spectrum of full FRAST in the cost model.

w	(w_a, w_c)	Data	T_{off}	T^{b}_{on}	T^{w}_{on}	d_w	Criterion	Threshold
$4^\dagger$	$(0,4)$	2^{64}	2^{92}	$2^{115.56}$	$2^{108.96}$	$2^{-13.6}$	$2^{129.15}$	56.337
$5^\dagger$	$(0,5)$	2^{64}	2^{99}	$2^{113.65}$	$2^{107.05}$	$2^{-11.8}$	$2^{125.49}$	56.373
$6^\dagger$	$(0,6)$	2^{64}	2^{106}	$2^{111.75}$	$2^{105.15}$	$2^{-10.1}$	$2^{121.91}$	56.410
$7^\dagger$	$(3,4)$	2^{64}	2^{110}	$2^{111.21}$	$2^{104.61}$	$2^{-8.5}$	$2^{120.22}$	56.446
$8^\dagger$	$(8,0)$	2^{64}	2^{112}	$2^{111.59}$	$2^{104.99}$	$2^{-6.9}$	$2^{119.72}$	56.483
$9^\dagger$	$(9,0)$	2^{64}	2^{118}	$2^{110.15}$	$2^{103.55}$	$2^{-5.4}$	$2^{123.41}$	56.520
$10^\dagger$	$(10,0)$	2^{64}	2^{124}	$2^{108.70}$	$2^{102.10}$	$2^{-4.0}$	$2^{128.00}$	56.556
Bak et al. [BGL+26]		2^7	$\le 2^{126}$	$\le 2^{31.86}$	—	$\le 2^{-61.14}$	$2^{175.1}$–$2^{194.1}$	—

d_w of the keys when $(T_{\text{off}} + T_{\text{on}})/d_w < 2^{128}$, with T_{on} in the main unit. This criterion adds XOF calls to bit operations, as Bak et al. did when they chose w by equating T_{precomp} with $T_{\text{online}}/\text{Adv}(D)$ in their Section 3.5. Unlike their criterion $\text{Adv}(D) > T_{\text{online}}/2^{128}$, it charges the offline time.

The results on full FRAST in Table 4 are estimates of the cost model, and no attack on full FRAST was run. For each w, the table takes the partition with the smallest criterion, and the attack guesses the classes of $x = 31 - w$ random rounds. In the Gaussian model, the right guess of a weak branch pair passes with probability 0.9. Criterion is $(T_{\text{off}} + T^{\text{b}}_{\text{on}})/d_w$, and Threshold is the largest detectable cost under the strict condition and the extensions of Appendix C.2, with the target $d_w/10$ for false alarms. The integer part 56 of every threshold leaves the densities unchanged. The criterion holds for $w \in [5, 9]$, with densities from $2^{-11.8}$ to $2^{-5.4}$ and offline times from 2^{99} to 2^{118} XOF calls. It fails for $w = 4$ by 1.15 bits, and $w = 10$ reaches 2^{128} and is not counted. The margin below 2^{128} is smallest at $w = 5$, with 2.51 bits, or 2.86 bits in the exact expectation model. With the extrapolation of the constant of the implementation in Appendix C.1, the margin is 1.5 to 2.4 bits, or about 0.4 bits from the least favourable point of the server.

Table 5 applies Theorem 8(iii) to the weak master keys of Table 4 with $d = d_w$ and $N = 2^{64}$. Its set $\mathcal{H}$ holds the $465 \cdot 29 \cdot 5^x \cdot 1921 \cdot 2^{w_a}$ guesses of all 465 branch pairs, and $m^2 = 2^{64-56}$ at the cap. In every row, the target $\varepsilon = 2|\mathcal{H}|e^{-\theta^2/2}$ for false alarms with the largest advantage among $d_w/10$, $d_w/100$ and $d_w/1000$ is $d_w/1000$, and p_{hit} subtracts the lower tail of Σ_2 at ε/d_w for the worst composition of classes at cost 56. In the accounting of Bak et al., the criterion with Adv in place of d_w is below 2^{128} for $w \in [5, 9]$ and exceeds that of Table 4 by less than 0.01 bits. Its smallest margin is 2.51 bits at $w = 5$, and $w = 10$ reaches 2^{128} and is not counted.

In bit operations alone, no row of Table 4 has its criterion below 2^{128}. Of the partitions that minimise this criterion, which differ from Table 4 at $w = 5$, 6 and 7, only $(w_a, w_c) = (3, 2)$ at $w = 5$ keeps it 0.37 bits below 2^{128} at the model constant. The margin vanishes when the constant of step 2 of Contract exceeds 4.50, which is 0.59 bits above the model constant and within the 0.1 to 1.0 bits by which the extrapolated constant of Appendix C.1 exceeds it.

The advantage bound is proved, while the densities d_w remain figures of the cost model. The densities of Table 4 for $w = 5$ to 9 are about 2^{49} to 2^{56} times the advantage of the distinguisher of Bak et al. in its last row, which approximates their fraction of weak master keys. In the columns T_{off}, T^{b}_{on} and d_w, this row gives the largest T_{precomp}, T_{online} and $\text{Adv}(D)$ over the four rows of their Table 2, and its Data are the 2^7 queries of the distinguisher. Its Criterion is the range of the four rows under our criterion. Our densities treat the branch pairs as independent, and a Monte Carlo over 10^9 master keys with the real key schedule agrees with them within 0.3% for $w = 4$ to 8. For a weak master key, the attack recovers the classes of one branch pair. Joint detection over several branch pairs and side trails through a fixed round with $c_{\text{fix}}(t'_r) = 0$ add little, as detailed in Appendix C.4.

Table 5: Distinguisher D_θ on full FRAST for the weak master keys of Table 4, the last three rows with the partition that minimises the bit-operation criterion. The threshold θ is that of the target $\varepsilon = d_w/1000$ for false alarms, p_{hit} the success bound of Proposition 2 at the cap 56, Adv the advantage bound of Theorem 8, and the last two columns count the time in the accounting of Bak et al. and in bit operations.

w	(w_a, w_c)	$\log \lvert\mathcal{H}\rvert$	θ	p_{hit}	$\log \text{Adv}$	$\log \frac{T_{\text{off}}+T_{\text{on}}^{\text{b}}}{\text{Adv}}$	$\log \frac{2^{18.5}T_{\text{off}}+T_{\text{on}}^{\text{b}}}{\text{Adv}}$
$4^\dagger$	$(0,4)$	87.32	12.454	0.9971	−13.60	129.16	129.20
$5^\dagger$	$(0,5)$	85.00	12.225	0.9982	−11.84	125.49	129.44
$6^\dagger$	$(0,6)$	82.67	11.994	0.9987	−10.14	121.91	134.64
$7^\dagger$	$(3,4)$	83.35	11.938	0.9987	−8.49	120.22	136.99
$8^\dagger$	$(8,0)$	86.03	12.002	0.9987	−6.91	119.73	137.41
$9^\dagger$	$(9,0)$	84.71	11.838	0.9988	−5.41	123.42	141.91
$10^\dagger$	$(10,0)$	83.39	11.677	0.9989	−4.00	128.00	146.50
$5^\dagger$	$(3,2)$	88.00	12.394	0.9975	−11.84	126.86	127.63
$6^\dagger$	$(6,0)$	88.67	12.336	0.9978	−10.14	124.63	128.73
$7^\dagger$	$(7,0)$	87.35	12.168	0.9983	−8.49	121.55	132.99

Table 6: Runs of three executed attacks on the server.

Target	Variant	Rounds	Pairs	Blocks	Proc.	Time	CPU	Mem.
h	A2	10	13	2^{20}, 2^{22}	48	176	2.24	0.45
h	A1	15	1	2^{29}	48	709	9.14	20.21
k	A1	10	465	2^{28}	48	7 872	57.76	2.95

7.4 Experimental Validation

Every run is registered with its command, seeds, threads and code version, and the runs of Table 6 ran on one server with two Intel Xeon Gold 6230R processors at 2.1 GHz, 256 GB of memory and at most 48 threads. In the runs of Table 6 and in the statistical experiments, a pseudorandom generator replaced the XOF, and on five rounds, the attack with S-boxes from SHAKE256 agreed with it within sampling error on 20 master keys. Table 6 lists the runs on the server of three of the executed attacks with the data of Table 3, and its row of variant A2 on ten rounds sums 26 runs. Target and Variant are as in Table 3, and Pairs and Blocks give the size of the runs. Proc. is the number of threads, Time is the elapsed time in seconds, CPU is the processor time in core-hours, and Mem. is the peak resident memory in gigabytes of 10^9 bytes. Appendix C.3 reports the checks of the success model, of the implementation and of key recovery and the calibration of the constant of Contract.

We reproduced the observation of Bak et al. on five rounds of FRAST for one master key and one branch pair. With the product filter, the sign statistic and right guesses of the rounds 1 and 5, 63 of the 4095 wrong guesses of the rounds 2 to 4 gave the statistic of the right guess. These are the guesses in the classes of the right key differences. Without a filter, the matched filter ranked the right triple of classes first among the 125 triples in all 40 runs with 2^{14} blocks. With 2^{22} blocks at each of three thresholds, we checked the parity rule on the 448 triples of residues modulo 8 with a parity mismatch. Their exact expected score is zero, and their statistics deviated from zero by at most 2.91 standard errors.

Simulations of FRAST on 10, 15, 20 and 40 rounds, with 2^{24} to 2^{36} blocks and 27 restricted rounds on 40 rounds, checked the correlations of Section 4. In every run, the measured $|\xi|$ of the right guess differed from its expected value $|\Gamma|\sqrt{\Sigma_2}$ by at most 1.89, against a standard deviation of at most 1 by Lemma 2. On 20 rounds with 16 unfiltered random rounds and 2^{36} blocks, the largest $|\xi|$ over the 60 wrong classes of single rounds was 2.07. Exact computations agreed with Theorem 5 and with the full cipher, and the

side trails averaged to zero against the matched filter weight. An exact dynamic program over the parity patterns of the guessed rounds under the filter tested the extensions of Appendix C.2 on ten instances of variant A2 and seven for the chosen-nonce and ELP filters. Every category of guesses agreed with the extensions, and under the ELP filters, which the extensions do not cover, the program found asymmetric wrong guesses.

8 A Variant with Independent Branch S-Boxes

In this section, we define a variant of FRAST in which the branches of a random round have independent expanding S-boxes, and we recompute the classes, the side trails and the reach. The branches 2 to 32 of a random round of FRAST share one expanding S-box, so the correlation of a round is an autocorrelation of it. In the variant, which changes only this choice, every random round r and every branch $l \geq 2$ have their own uniformly random negacyclic S-box $S_l^{(r)}$ from the XOF, and the round updates the branches by

$$y_l^{(r)} = y_l^{(r-1)} + S_l^{(r)}\big(y_1^{(r-1)} + \mathrm{rk}_l^{(r)}\big), \qquad l = 2, \ldots, 32.$$

We write $b^{(r,l)}$ for the parity pattern of $S_l^{(r)}$ and define, for $s \in \mathbb{Z}_{16}$ and with indices modulo 8,

$$c_r(s) = \frac{1}{8} \sum_{x \in \mathbb{Z}_8} b_{x+s}^{(r,i)}\, b_x^{(r,j)}.$$

The substitution $x = z + \mathrm{rk}_j^{(r)}$ of the proof of Theorem 4 shows that the correlation of round r for the mask 8 is $c_r(t'_r)$.

Theorem 9 (Classes of a Round of the Variant). *Let round r be random, and let $b^{(r,i)}$ and $b^{(r,j)}$ be independent and uniform on $\{\pm 1\}^8$.*

(i) For every $s \in \mathbb{Z}_{16}$, $c_r(s+8) = c_r(s)$, so the key differences fall into the eight classes of their residues modulo 8. For all $s, t \in \mathbb{Z}_{16}$, $\mathbb{E}[c_r(s)\, c_r(t)] = \frac{1}{8}\mathbf{1}[s \equiv t \pmod 8]$. In particular, $\mathbb{E}[c_r(s)^2] = \frac{1}{8}$ for every s.

(ii) For every s, $|c_r(s)| \in \{0, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1\}$ and $\Pr[|c_r(s)| = 1] = 2^{-7}$, and $|c_r(t'_r)| = 1$ if and only if $b^{(r,i)} = \pm \mathrm{shift}_{t'_r}(b^{(r,j)})$, where $\mathrm{shift}_s(b)_y = b_{y-s}$.

(iii) $|c_r(s)| = 1$ holds for all s if and only if both patterns are constant or both are alternating, and this has the probability 2^{-13}.

Theorem 9 and Proposition 4 below are proved in Appendix B.6. By Theorem 9(i), every class of the variant costs three bits for every key difference, whereas Theorem 2 gives FRAST the class [0] at no cost and the class [4] at two bits. The side trails of the variant average to zero for every weight and every filter that depend only on the patterns of the two branches. The reason is that the S-box of a branch outside the pair takes uniform values, whereas Theorem 6 needs the condition (F_r) or (S_r) in FRAST.

Proposition 4 (Side Trails and the First Round of the Variant). *Let the XOF be a random oracle, fix the master key and a branch pair (i, j), and let the filter depend only on the patterns $b^{(r,i)}$ and $b^{(r,j)}$ of the random rounds r, which we write as $\mathcal{C}$. Then the observation $Z = (-1)^{D_R}$ of the variant satisfies*

$$\mathbb{E}[Z \mid \mathcal{C}, \Phi] = s_1\, \Gamma \prod_r c_r(t'_r), \qquad s_1 = b_{v'_i}^{(1,i)} b_{v'_j}^{(1,j)},$$

with $v'_i = \mathrm{rk}_i^{(1)} \bmod 8$, $v'_j = \mathrm{rk}_j^{(1)} \bmod 8$, Γ of equation (16) and the product over the random rounds $2 \leq r \leq R$. Without a filter, the weight $W = b_{v_i}^{(1,i)} b_{v_j}^{(1,j)} \prod_r c_r(h_r)$ of a

guess with $(v_i, v_j) \in \mathbb{Z}_8^2$ *and* $h_r \in \mathbb{Z}_8$ *satisfies* $\mathbb{E}[WZ] = 0$ *unless* $(v_i, v_j) = (v'_i, v'_j)$ *and* $h_r \equiv t'_r \pmod 8$ *in every random round* $2 \le r \le R$.

In the cost model with known nonces and 2^{64} blocks, variant A2 covers every sampled master key at 22 rounds and 8.6% of them at 24 rounds, against 26 rounds for every key of FRAST. At 40 rounds, the best branch pair of every sampled key needs at least $2^{104.64}$ blocks. With chosen nonces, a filter of the variant acts on one branch pair and makes a round free for every key difference only with the probability 2^{-13} of Theorem 9(iii). The constant and the alternating patterns of FRAST have the probability 2^{-6} and serve all 465 branch pairs at once. Under the criterion of Section 7.3, the chosen-nonce attacks therefore do not reach 40 rounds of the variant either, as computed in Appendix C.5.

The variant has a cost in TFHE, because it loses the double blind rotation of Cho et al. [CCH+24], in which one programmable bootstrapping and CMux gates evaluate one negacyclic S-box at x_1 plus each round key word. This works because all branches share the blind rotation by x_1. With a different S-box in each branch, the rotated encodings differ, and the sharing is lost. The shared expanding S-box is therefore both the source of the cheap classes [0] and [4] of our attacks and the property on which the double blind rotation rests.

9 Conclusion

We explained why wrong guesses passed the linear attack of Bak et al. on FRAST and turned it into working average key attacks. The key differences of a round fall into five classes that no test on the least significant bit separates. The ELP filter gives some wrong classes of the right parity a nonzero expected score. The trail correlation of the real cipher is an exact product inside a segment, with a transfer matrix across each fixed round. With known nonces, our attacks recovered the classes of a branch pair of 15 rounds and the master key of ten rounds.

The cost model gives class recovery for every sampled key at 26 rounds with known nonces, for a third at 38 rounds with chosen nonces and, on full FRAST, for weak keys of density $2^{-11.8}$ to $2^{-5.4}$ within the claimed security. A distinguisher for these keys has a proved advantage. The expanding S-box shared by the branches of a round makes two classes cheap and lets a chosen nonce make a round free for all pairs. How a public random component is shared among the branches is therefore part of a design's security analysis. One open question is whether the bound of Theorem 7 for a fixed master key holds with a stripped last round and with the chosen-nonce filters, as exact computations indicate. The variant with independent branch S-boxes removes the cheap classes, and its reach, its security against other attacks and its cost in TFHE are open.

References

[AMT22] Tomer Ashur, Mohammad Mahzoun, and Dilara Toprakhisar. Chaghri - A FHE-friendly block cipher. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 139–150. ACM Press, November 2022. doi:10.1145/3548606.3559364.

[ARS+15] Martin R. Albrecht, Christian Rechberger, Thomas Schneider, Tyge Tiessen, and Michael Zohner. Ciphers for MPC and FHE. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 430–454. Springer, Berlin, Heidelberg, April 2015. doi:10.1007/978-3-662-46800-5_17.

[BBB+25] Jules Baudrin, Sonia Belaïd, Nicolas Bon, Christina Boura, Anne Canteaut, Gaëtan Leurent, Pascal Paillier, Léo Perrin, Matthieu Rivain, Yann Rotella, and Samuel Tap. Transistor: a TFHE-friendly stream cipher. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part V*, volume 16004 of *LNCS*, pages 530–565. Springer, Cham, August 2025. doi:10.1007/978-3-032-01901-1_17.

[Bey21] Tim Beyne. A geometric approach to linear cryptanalysis. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part I*, volume 13090 of *LNCS*, pages 36–66. Springer, Cham, December 2021. Statement numbers refer to the version in Cryptology ePrint Archive, Paper 2021/1247. doi:10.1007/978-3-030-92062-3_2.

[BGL+26] Antoine Bak, Shibam Ghosh, Fukang Liu, Willi Meier, Jianqiang Ni, and Léo Perrin. Cryptanalysis of TFHE-friendly cipher FRAST. *IACR Trans. Symm. Cryptol.*, 2026(1):119–147, 2026. doi:10.46586/tosc.v2026.i1.119-147.

[BGØ+26] Tim Beyne, Lorenzo Grassi, Morten Øygarden, Berenika Richterová, and Arne Sandrib. The ABC of symmetric primitives over integer rings: Milk before meat. In Nadia Heninger and Mike Rosulek, editors, *CRYPTO 2026, Part VI*, LNCS, pages 438–470. Springer, Cham, August 2026. doi:10.1007/978-3-032-35412-9_15.

[BGV12] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In Shafi Goldwasser, editor, *ITCS 2012*, pages 309–325. ACM, January 2012. doi:10.1145/2090236.2090262.

[BLM13] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, Oxford, 2013. doi:10.1093/acprof:oso/9780199535255.001.0001.

[BLMR26] Tim Beyne, Gregor Leander, Mariia Mutkovina, and Ricardo Rodriguez Reveco. A quasidifferential analysis of the wrong-key randomization hypothesis. Cryptology ePrint Archive, Paper 2026/1848, 2026.

[Bra12] Zvika Brakerski. Fully homomorphic encryption without modulus switching from classical GapSVP. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 868–886. Springer, Berlin, Heidelberg, August 2012. doi:10.1007/978-3-642-32009-5_50.

[BSV07] Thomas Baignères, Jacques Stern, and Serge Vaudenay. Linear cryptanalysis of non binary ciphers. In Carlisle M. Adams, Ali Miri, and Michael J. Wiener, editors, *SAC 2007*, volume 4876 of *LNCS*, pages 184–211. Springer, Berlin, Heidelberg, August 2007. doi:10.1007/978-3-540-77360-3_13.

[BT14] Andrey Bogdanov and Elmar Tischhauser. On the wrong key randomisation and key equivalence hypotheses in Matsui's Algorithm 2. In Shiho Moriai, editor, *FSE 2013*, volume 8424 of *LNCS*, pages 19–38. Springer, Berlin, Heidelberg, March 2014. doi:10.1007/978-3-662-43933-3_2.

[CCF+18] Anne Canteaut, Sergiu Carpov, Caroline Fontaine, Tancrède Lepoint, María Naya-Plasencia, Pascal Paillier, and Renaud Sirdey. Stream ciphers: A practical solution for efficient homomorphic-ciphertext compression. *Journal of Cryptology*, 31(3):885–916, July 2018. doi:10.1007/s00145-017-9273-9.

[CCH+24] Mingyu Cho, Woohyuk Chung, Jincheol Ha, Jooyoung Lee, Eun-Gyeol Oh, and Mincheol Son. FRAST: TFHE-friendly cipher based on random S-boxes. *IACR Trans. Symm. Cryptol.*, 2024(3):1–43, 2024. doi:10.46586/tosc.v2024.i3.1-43.

[CGGI20] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. TFHE: Fast fully homomorphic encryption over the torus. *Journal of Cryptology*, 33(1):34–91, January 2020. doi:10.1007/s00145-019-09319-x.

[CHK+21] Jihoon Cho, Jincheol Ha, Seongkwang Kim, ByeongHak Lee, Joohee Lee, Jooyoung Lee, Dukjae Moon, and Hyojin Yoon. Transciphering framework for approximate homomorphic encryption. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part III*, volume 13092 of *LNCS*, pages 640–669. Springer, Cham, December 2021. doi:10.1007/978-3-030-92078-4_22.

[CHMS22] Orel Cosseron, Clément Hoffmann, Pierrick Méaux, and François-Xavier Standaert. Towards case-optimized hybrid homomorphic encryption – featuring the Elisabeth stream cipher. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022, Part III*, volume 13793 of *LNCS*, pages 32–67. Springer, Cham, December 2022. doi:10.1007/978-3-031-22969-5_2.

[CIR22] Carlos Cid, John Petter Indrøy, and Håvard Raddum. FASTA - A stream cipher for fast FHE evaluation. In Steven D. Galbraith, editor, *CT-RSA 2022*, volume 13161 of *LNCS*, pages 451–483. Springer, Cham, March 2022. doi:10.1007/978-3-030-95312-6_19.

[Civ26] Roberto Civino. Circle-linear cryptanalysis: Bibrace characters and weak-key linear distinguishers for CRAFT. Cryptology ePrint Archive, Paper 2026/1767, 2026.

[CKKS17] Jung Hee Cheon, Andrey Kim, Miran Kim, and Yong Soo Song. Homomorphic encryption for arithmetic of approximate numbers. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 409–437. Springer, Cham, December 2017. doi:10.1007/978-3-319-70694-8_15.

[DEG+18] Christoph Dobraunig, Maria Eichlseder, Lorenzo Grassi, Virginie Lallemand, Gregor Leander, Eik List, Florian Mendel, and Christian Rechberger. Rasta: A cipher with low ANDdepth and few ANDs per bit. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 662–692. Springer, Cham, August 2018. doi:10.1007/978-3-319-96884-1_22.

[DGH+23] Christoph Dobraunig, Lorenzo Grassi, Lukas Helminger, Christian Rechberger, Markus Schofnegger, and Roman Walch. Pasta: A case for hybrid homomorphic encryption. *IACR TCHES*, 2023(3):30–73, 2023. doi:10.46586/tches.v2023.i3.30-73.

[DLR16] Sébastien Duval, Virginie Lallemand, and Yann Rotella. Cryptanalysis of the FLIP family of stream ciphers. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 457–475. Springer, Berlin, Heidelberg, August 2016. doi:10.1007/978-3-662-53018-4_17.

[DM15] Léo Ducas and Daniele Micciancio. FHEW: Bootstrapping homomorphic encryption in less than a second. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 617–640. Springer, Berlin, Heidelberg, April 2015. doi:10.1007/978-3-662-46800-5_24.

[FV12] Junfeng Fan and Frederik Vercauteren. Somewhat practical fully homomorphic encryption. Cryptology ePrint Archive, Paper 2012/144, 2012. URL: https://eprint.iacr.org/2012/144.

[GAH+23] Lorenzo Grassi, Irati Manterola Ayala, Martha Norberg Hovd, Morten Øygarden, Håvard Raddum, and Qingju Wang. Cryptanalysis of symmetric primitives over rings and a key recovery attack on Rubato. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part III*, volume 14083 of *LNCS*, pages 305–339. Springer, Cham, August 2023. doi:10.1007/978-3-031-38548-3_11.

[GBJR23] Henri Gilbert, Rachelle Heim Boissier, Jérémy Jean, and Jean-René Reinhard. Cryptanalysis of Elisabeth-4. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part III*, volume 14440 of *LNCS*, pages 256–284. Springer, Singapore, December 2023. doi:10.1007/978-981-99-8727-6_9.

[Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009. doi:10.1145/1536414.1536440.

[GLR+24] Lorenzo Grassi, Fukang Liu, Christian Rechberger, Fabian Schmid, Roman Walch, and Qingju Wang. Minimize the randomness in Rasta-like designs: How far can we go? - Application to Pasta. In Maria Eichlseder and Sébastien Gambs, editors, *SAC 2024, Part II*, volume 15517 of *LNCS*, pages 207–238. Springer, Cham, August 2024. doi:10.1007/978-3-031-82841-6_9.

[HKC+20] Jincheol Ha, Seongkwang Kim, Wonseok Choi, Jooyoung Lee, Dukjae Moon, Hyojin Yoon, and Jihoon Cho. Masta: An HE-friendly cipher using modular arithmetic. *IEEE Access*, 8:194741–194751, 2020. doi:10.1109/ACCESS.2020.3033564.

[HKL+22] Jincheol Ha, Seongkwang Kim, ByeongHak Lee, Jooyoung Lee, and Mincheol Son. Rubato: Noisy ciphers for approximate homomorphic encryption. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part I*, volume 13275 of *LNCS*, pages 581–610. Springer, Cham, May / June 2022. doi:10.1007/978-3-031-06944-4_20.

[HL20] Phil Hebborn and Gregor Leander. Dasta – alternative linear layer for Rasta. *IACR Trans. Symm. Cryptol.*, 2020(3):46–86, 2020. doi:10.13154/tosc.v2020.i3.46-86.

[HMAV26] Hosein Hadipour, Saleh Khalaj Monfared, Jens Alich, and Jan Vorloeper. LibFWHT: From exact Walsh spectra to key dependence in differential-linear correlations. Cryptology ePrint Archive, Paper 2026/1997, 2026.

[KLW17] Thorsten Kranz, Gregor Leander, and Friedrich Wiemer. Linear cryptanalysis: Key schedules and tweakable block ciphers. *IACR Trans. Symm. Cryptol.*, 2017(1):474–505, 2017. Also Cryptology ePrint Archive, Paper 2017/154. doi:10.46586/tosc.v2017.i1.474-505.

[LAW+23] Fukang Liu, Ravi Anand, Libo Wang, Willi Meier, and Takanori Isobe. Coefficient grouping: Breaking Chaghri and more. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part IV*, volume 14007 of *LNCS*, pages 287–317. Springer, Cham, April 2023. doi:10.1007/978-3-031-30634-1_10.

[LKSM24] Fukang Liu, Abul Kalam, Santanu Sarkar, and Willi Meier. Algebraic attack on FHE-friendly cipher HERA using multiple collisions. *IACR Trans. Symm. Cryptol.*, 2024(1):214–233, 2024. doi:10.46586/tosc.v2024.i1.214-233.

[LMSW22] Muzhou Li, Nicky Mouha, Ling Sun, and Meiqin Wang. Revisiting the extension of Matsui's Algorithm 1 to linear hulls: Application to TinyJAMBU. *IACR Trans. Symm. Cryptol.*, 2022(2):161–200, 2022. doi:10.46586/tosc.v2022.i2.161-200.

[LSMI21] Fukang Liu, Santanu Sarkar, Willi Meier, and Takanori Isobe. Algebraic attacks on Rasta and Dasta using low-degree equations. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part I*, volume 13090 of *LNCS*, pages 214–240. Springer, Cham, December 2021. doi:10.1007/978-3-030-92062-3_8.

[LSMI22] Fukang Liu, Santanu Sarkar, Willi Meier, and Takanori Isobe. The inverse of χ and its applications to Rasta-like ciphers. *Journal of Cryptology*, 35(4):art. 28, October 2022. doi:10.1007/s00145-022-09439-x.

[LSW+22] Fukang Liu, Santanu Sarkar, Gaoli Wang, Willi Meier, and Takanori Isobe. Algebraic meet-in-the-middle attack on LowMC. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022, Part I*, volume 13791 of *LNCS*, pages 225–255. Springer, Cham, December 2022. doi:10.1007/978-3-031-22963-3_8.

[Mat94] Mitsuru Matsui. Linear cryptanalysis method for DES cipher. In Tor Helleseth, editor, *EUROCRYPT'93*, volume 765 of *LNCS*, pages 386–397. Springer, Berlin, Heidelberg, May 1994. doi:10.1007/3-540-48285-7_33.

[MJSC16] Pierrick Méaux, Anthony Journault, François-Xavier Standaert, and Claude Carlet. Towards stream ciphers for efficient FHE with low-noise ciphertexts. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 311–343. Springer, Berlin, Heidelberg, May 2016. doi:10.1007/978-3-662-49890-3_13.

[Nat15] National Institute of Standards and Technology. SHA-3 standard: Permutation-based hash and extendable-output functions. Federal Information Processing Standards Publication 202, National Institute of Standards and Technology, August 2015. doi:10.6028/NIST.FIPS.202.

[NLV11] Michael Naehrig, Kristin Lauter, and Vinod Vaikuntanathan. Can homomorphic encryption be practical? In *Proceedings of the 3rd ACM workshop on Cloud computing security workshop*, pages 113–124, New York, NY, USA, October 2011. ACM. doi:10.1145/2046660.2046682.

[Pin12] Iosif Pinelis. An asymptotically Gaussian bound on the Rademacher tails. *Electron. J. Probab.*, 17:no. 35, 1–22, 2012. arXiv:1007.2137. doi:10.1214/EJP.v17-2026.

[RN13] Andrea Röck and Kaisa Nyberg. Generalization of Matsui's Algorithm 1 to linear hull for key-alternating block ciphers. *DCC*, 66(1–3):175–193, 2013. doi:10.1007/s10623-012-9679-1.

[Ter99] Audrey Terras. *Fourier Analysis on Finite Groups and Applications*, volume 43 of *London Mathematical Society Student Texts*. Cambridge University Press, 1999. doi:10.1017/CBO9780511626265.

[WT26] Weizhe Wang and Deng Tang. Differential fault attacks on TFHE-friendly cipher FRAST. *IEEE Trans. Inf. Forensics Secur.*, 21:7980–7991, 2026. Also Cryptology ePrint Archive, Paper 2025/771. doi:10.1109/TIFS.2026.3725560.

A Specification Notes

SHAKE256 [Nat15] absorbs for each block the nonce and then the counter, as strings of eight bytes with the least significant byte first and without the key or a separator, and produces 256 bytes $\beta_0, \dots, \beta_{255}$. The random rounds take eight bytes each, in the order of the rounds, and the fixed rounds take none. For the random round r, let $q = r - \lfloor r/5 \rfloor$ be its position among the random rounds. Each byte gives two words with its low four bits first, and

$$S_{\mathrm{erf}}^{(r)}(x) = \left(\beta_{8(q-1)+\lfloor x/2 \rfloor} \gg 4(x \bmod 2)\right) \bmod 16,$$
$$S_{\mathrm{crf}}^{(r)}(x) = \left(\beta_{8(q-1)+4+\lfloor x/2 \rfloor} \gg 4(x \bmod 2)\right) \bmod 16$$

for $x = 0, \dots, 7$. Equation (4) then gives the values at $x = 8, \dots, 15$, and no value is rejected. We compared our implementation with the reference implementation of the designers at commit `ccee1d9` on eight triples of master key, nonce and counter. Their master keys were zero, the words 0 to 15 repeated and the first 64 hexadecimal digits of the fractional part of π. The keystream blocks, the S-boxes and the round keys of all 40 rounds agreed on every triple.

B Deferred Proofs

In this appendix, we give the proofs that the body defers and the general statements that only these proofs use, with the notation of Section 3.1. Appendix B.1 collects those on random functions, and Appendices B.2 to B.6 give the proofs for Sections 3 to 6 and 8.

B.1 General Statements on Random Functions

Let $m \geq 2$, and let μ_m be the group of the m-th roots of unity. In the *model* (N_c), for an element $c \in G$ of order 2 and a set T of representatives of the cosets of $\{0, c\}$, the values $f(x)$ for $x \in T$ are independent and uniform on μ_m, and $f(x + c) = \overline{f(x)}$ for every $x \in G$.

Theorem 10 (Kernel of a Random Function). *In the model* (N_c), $K(0,0) = 1$, *and for* $(t, t') \neq (0, 0)$ *the kernel is*

$$K(t,t') = \begin{cases} \frac{1}{n}\left(\mathbf{1}[t' = t] + \mathbf{1}[t' = -t] - \mathbf{1}[m \neq 4]\,\mathbf{1}[t = t' = c]\right) & \text{if } m \geq 3, \\ \frac{2}{n}\left(\mathbf{1}[\bar{t}' = \bar{t}] + \mathbf{1}[\bar{t}' = -\bar{t}]\right) & \text{if } m = 2,\ (\bar{t}, \bar{t}') \neq (0,0), \\ 1 & \text{if } m = 2,\ \bar{t} = \bar{t}' = 0, \end{cases}$$

where $\bar{t}$ *is the image of* t *in* $G/\{0, c\}$. *The class* $[t]$ *is* $\{t, -t\}$ *for* $m \geq 3$ *and* $\{\pm t, \pm t + c\}$ *for* $m = 2$. *In both cases,* $K(t,t') = 0$ *for* $t' \notin [t]$ *and* $K(t,t) \geq 1/n$, *and for the set* $\mathcal{F}$ *of all functions of the model, the classes of* $\approx_{\mathcal{F}}$ *are the sets* $[t]$.

Proof of Theorem 10. We expand $\overline{C_f(t)}\, C_f(t') = \frac{1}{n^2} \sum_{x,y \in G} \overline{f(x+t)} f(x) f(y+t') \overline{f(y)}$ and count the pairs (x, y) whose term has mean 1. Since $C_f(0) = 1$, we may assume $(t, t') \neq (0, 0)$. In the case $m \geq 3$, we use $\overline{f(z)} = f(z + c)$ to write the term as $f(z_1) f(z_2) f(z_3) f(z_4)$ with $z_1 = x + t + c$, $z_2 = x$, $z_3 = y + t'$ and $z_4 = y + c$. Let $N(z)$ be the number of indices i with $z_i = z$. With $M_x = N(x) + N(x + c)$ and $E_x = N(x) - N(x + c)$ for $x \in T$, independence, $f(x + c) = f(x)^{-1}$ and $\mathbb{E}[g^E] = \mathbf{1}[m \mid E]$ for g uniform on μ_m give the term

the mean $\prod_{x\in T} \mathbf{1}[m \mid E_x]$. This mean is 1 if and only if (A) $N(z) = N(z+c)$ for every $z \in G$, or (B) $m = 4$ and $z_1 = z_2 = z_3 = z_4$. Condition (A) makes every E_x zero, and (B) gives one $|E_x| = 4$, which m divides. Conversely, if m divides every E_x and some $E_x \neq 0$, then $|E_x| \geq m \geq 3$, and $M_x \geq 3$ because $|E_x| \leq M_x$ and $E_x \equiv M_x \pmod 2$. If $M_x = 3$, then $m = |E_x| = 3$, and the fourth point gives some $M_{x'} = 1$ with $|E_{x'}| = 1$, which m does not divide. If $M_x = 4$, then $m = |E_x| = 4$, and all four points coincide.

Condition (A) holds if and only if the indices $1, 2, 3, 4$ split into two pairs $\{i, i'\}$ with $z_i - z_{i'} = c$. Since $2c = 0$, the pairs $\{1,2\}, \{3,4\}$ need $t = 0$ and $t' = 0$, which is excluded. The pairs $\{1,3\}, \{2,4\}$ give $y = x$ and $t' = t$, and the pairs $\{1,4\}, \{2,3\}$ give $y = x + t + c$ and $t' = -t$. These two sets coincide when $t = t' = c$ and are disjoint otherwise, and (B) gives $t = t' = c$ and $y = x + c$, disjoint from both. Counting the pairs (x, y) gives $n(\mathbf{1}[t' = t] + \mathbf{1}[t' = -t] - \mathbf{1}[t = t' = c]) + n\,\mathbf{1}[m = 4]\mathbf{1}[t = t' = c]$.

For $m = 2$, f takes the values ± 1, so $f(x + c) = \overline{f(x)} = f(x)$. Hence f is the pullback of a function $\bar{f}$ of the model (U) with $m = 2$ on $\bar{G} = G/\{0, c\}$, in which the values are independent and uniform on $\{1, -1\}$, and $C_f(t) = \frac{2}{n} \sum_{\bar{x}\in\bar{G}} \bar{f}(\bar{x} + \bar{t})\bar{f}(\bar{x}) = C_{\bar{f}}(\bar{t})$ is 1 for $\bar{t} = 0$. For $(\bar{t}, \bar{t}') \neq (0, 0)$, a term of $\overline{C_{\bar{f}}(\bar{t})}\, C_{\bar{f}}(\bar{t}')$ has mean 1 if every point of $\{\bar{x}, \bar{x} + \bar{t}, \bar{y}, \bar{y} + \bar{t}'\}$ has an even multiplicity, and 0 otherwise. Pairing $\bar{x}$ with $\bar{x} + \bar{t}$ needs $\bar{t} = \bar{t}' = 0$. Hence the mean is 1 if and only if $\bar{y} = \bar{x}$ and $\bar{t}' = \bar{t}$, or $\bar{y} = \bar{x} + \bar{t}$ and $\bar{t}' = -\bar{t}$. These sets of pairs, $\bar{x} = \bar{y}$ with $\bar{x} + \bar{t} = \bar{y} + \bar{t}'$ and $\bar{x} = \bar{y} + \bar{t}'$ with $\bar{x} + \bar{t} = \bar{y}$, are disjoint for $\bar{t} \neq 0$ and empty for $\bar{t} = 0$. Each has $|\bar{G}| = n/2$ pairs, and dividing their count by $|\bar{G}|^2$ gives the case $m = 2$.

Every indicator vanishes for $t' \notin [t]$, since $\mathbf{1}[t = t' = c]$ implies $t' = t$, and the diagonal is at least $\frac{1}{n}$, since $\mathbf{1}[t = -t] = 1$ whenever $t = c$. Turning to the classes, every C_f is real, by $\overline{f(z)} = f(z + c)$ for $m \geq 3$ and because f is real for $m = 2$, and Theorem 1 gives $C_f(-t) = C_f(t)$. When $m = 2$, f and C_f also have the period c, and in both cases $t \approx_{\mathcal{F}} t'$ for $t' \in [t]$. For $t' \notin [t]$, $\mathbb{E}|C_f(t) - C_f(t')|^2 = K(t,t) + K(t',t') - 2\,\mathrm{Re}\,K(t,t')$ is positive, and some $f \in \mathcal{F}$ separates t from t'. □

Lemma 3 (Modulation Covariance). *Let f be a random function with values in $\mathbb{T}^G$, and let $\lambda \in \widehat{G}$ satisfy $f\lambda \overset{d}{=} f$. Let Φ be an event with $\Pr[\Phi] > 0$ and $\mathbf{1}_\Phi(f\lambda) = \mathbf{1}_\Phi(f)$, and let w_t be a weight with $w_t(f\lambda) = \overline{\lambda(t)}\, w_t(f)$. Then*

$$\mathbb{E}\big[w_t(f)\, C_f(t') \;\big|\; \Phi\big] = \overline{\lambda(t)}\, \lambda(t')\, \mathbb{E}\big[w_t(f)\, C_f(t') \;\big|\; \Phi\big],$$

which is zero when $\lambda(t' - t) \neq 1$. For a tuple $(f_r)_r$ of functions of several rounds, let multiplying the function of one round r_0 by λ leave the distribution of the tuple and the event unchanged, and let it multiply the weight by $\overline{\lambda(t_{r_0})}$. Then the same holds with $\prod_r C_{f_r}(t'_r)$ in place of $C_f(t')$, and the expectation is zero when $\lambda(t'_{r_0} - t_{r_0}) \neq 1$.

Proof. For every f, $C_{f\lambda}(t') = \frac{1}{n}\sum_x f(x + t')\lambda(x + t')\overline{f(x)\lambda(x)} = \lambda(t')\, C_f(t')$, and since $f\lambda \overset{d}{=} f$ and Φ is invariant,

$$\mathbb{E}\big[w_t(f) C_f(t') \mathbf{1}_\Phi(f)\big] = \mathbb{E}\big[w_t(f\lambda) C_{f\lambda}(t') \mathbf{1}_\Phi(f\lambda)\big] = \overline{\lambda(t)}\, \lambda(t')\, \mathbb{E}\big[w_t(f) C_f(t') \mathbf{1}_\Phi(f)\big].$$

Dividing by $\Pr[\Phi]$ gives the claim, and for a tuple only the factor of the round r_0 changes. □

The ELP filters, the product filter and the absence of a filter are invariant, since $\mathrm{ELP}(f\lambda) = \mathrm{ELP}(f)$. The matched filter and phase weights of Section 3.1 transform as required, and so does the sign weight for a real C_f and a λ with the values ± 1. So do their products over the rounds, such as g_u of Bak et al. in equation (12).

B.2 Proofs for Section 3

We first state when the assumptions of Lemma 1 hold. Assumption (H1) holds when the XOF is a random oracle and the attacker selects blocks by their public S-boxes. Assumption (H2) holds for the mask 8 and one round, because the observation $(-1)^{S(x+t)-S(x)}$ at a uniform hidden input x equals 1 with probability $(1+c(t))/2$. Across several rounds, (H2) needs the product of the ideal multi-round model, and Lemma 1 does not apply across a fixed round, where side trails involve other round keys.

Proof of the properties of negacyclic S-boxes of Section 3.1. For $f = \chi_a \circ S$, negacyclicity gives $f(x+8) = \overline{f(x)}$, and $\zeta^{-8k} = (-1)^k$. Splitting the sum over x into the terms of x and $x+8$ for $x < 8$ gives $\hat{f}(\chi_k) = \frac{1}{8}\sum_{x<8} \mathrm{Re} f(x)\,\zeta^{-kx}$ for even k and $\hat{f}(\chi_k) = \frac{i}{8}\sum_{x<8} \mathrm{Im} f(x)\,\zeta^{-kx}$ for odd k. In both cases, $\hat{f}(\chi_{-k})$ has the modulus of $\hat{f}(\chi_k)$ because the coefficients of the sums are real. Hence $P_{-k} = P_k$, and Theorem 1 makes $C_{a,S}$ real and even. For the mask 8, f is real, the odd coefficients vanish, and $\zeta^{8k} = 1$ for every even k gives the period 8. For the last claim and the mask 8, $f = (-1)^S$ is real, and by Theorem 1, $\mathrm{ELP}_8(S) = 1$ if and only if $f = \gamma\chi_k$ with $\gamma = f(0) \in \{1, -1\}$ and a real character χ_k. A character χ_k of $\mathbb{Z}_{16}$ is real if and only if $2k \equiv 0 \pmod{16}$, that is, $k \in \{0, 8\}$. Hence, f is ± 1 or $\pm(-1)^x$, which are the two constant and the two alternating parity patterns. □

Proof of Theorem 2. By the properties of negacyclic S-boxes of Section 3.1, $\chi_a \circ S$ follows the model (N_8) of Appendix B.1 on $\mathbb{Z}_{16}$ with $n = 16$ and $m = m_a = 16/\gcd(a, 16)$, since $x \mapsto ax$ maps the uniform distribution on $\mathbb{Z}_{16}$ to that on $a\mathbb{Z}_{16}$. The classes of (i) and (ii) are the sets $[t]$ that the last claim of Theorem 10 gives for this model, generated by the negation and, in (i), by the period 8. For the mask 8, $m = 2$, and Theorem 10 gives $K(t,t) = \frac{1}{8}(1 + \mathbf{1}[\bar{t} = -\bar{t}])$ for $\bar{t} \neq 0$ in $\mathbb{Z}_8$. This value is $\frac{1}{4}$ for $\bar{t} = 4$ and $\frac{1}{8}$ for the other $\bar{t} \neq 0$, and the class $[0]$ has $c(t) = 1$. For $a \notin \{0, 8\}$, $m \geq 4$, and the case $m \geq 3$ with $c = 8$ gives the values of (ii), where $m = 4$ if and only if $\gcd(a, 16) = 4$, that is, $a \in \{4, 12\}$. Part (iii) follows from (i), because the S-boxes of different rounds are independent and the expectation of the product is the product of their kernels. □

A wrong guess with the parity of the key difference in every round has zero expected score without a filter by Theorem 2, but that zero is a cancellation of moments in Theorem 10, which a filter can destroy. For the mask 8 in the ideal multi-round model, a wrong guess therefore has zero expected score without a filter when it differs from the key differences in class in some inner round. With an ELP filter or the product filter, a difference in parity suffices.

Proof of Theorem 3. Let $f = \chi_a \circ S^{(r)}_{\mathrm{erf}}$ for an inner round r with a parity mismatch, and let $g_u = f(u)$ for $u < 8$. The map $(g_u)_{u<8} \mapsto ((-1)^u g_u)_{u<8}$ is a bijection of $\mu^8_{m_a}$, because m_a is even and $-1 \in \mu_{m_a}$, and it preserves the uniform distribution. Since $(f\chi_8)(u+8) = (-1)^u\overline{g_u} = \overline{(-1)^u g_u}$, the function $f\chi_8$ is the negacyclic extension of the image. Hence $f\chi_8 \overset{d}{=} f$, and the independence of the rounds extends this to the joint distribution. Finally, $\chi_8(t'_r - t_r) = (-1)^{t'_r - t_r}$ differs from 1 if and only if the parities differ, and Lemma 3 gives the claim. □

B.3 Proofs for Section 4

Lemma 4 (Uniformisation by the Contracting S-box). *Let the XOF be a random oracle, and let round r be random. Conditioned on $\mathcal{E}$, on u_r, $x^{(r)}_2$ and D_r, and on the contracting S-boxes of the rounds before r, the word $x^{(r)}_1$ is uniform on $\mathbb{Z}_{16}$.*

Proof. The contracting S-box $G = S^{(r)}_{\mathrm{crf}}$ is evaluated once in the block, at $v = x^{(r)}_2 + \mathrm{rk}^{(r)}_1$. The words u_r and v depend only on the key, on $\mathcal{E}$ and on earlier contracting S-boxes, so G

is independent of them. Given v, equation (4) gives $G(v) = \pm G(v \bmod 8)$ with a sign fixed by v, where $G(v \bmod 8)$ is uniform on $\mathbb{Z}_{16}$ and so is its negation. Hence $x_1^{(r)} = u_r + G(v)$ is uniform. □

Proof of Theorem 4. For $r = q, \ldots, q'$, Lemma 4 applied to round $r-1$ makes $u_r = x_1^{(r-1)}$ uniform given u_{r-1}, D_{r-1} and everything before, also under the filter, which is a function of $\mathcal{E}$. Hence $u_q, \ldots, u_{q'}$ are independent, uniform and independent of D_{q-1}. For a uniform u, the substitution $x = u + \mathrm{rk}_j^{(r)}$ and equation (9) give $\mathbb{E}(-1)^{\Delta_r(u)} = c_r(t'_r)$, and $D_{q'} - D_{q-1} = \sum_{r=q}^{q'} \Delta_r(u_r)$ gives the product. If $R \le 4$, all rounds up to R are random, and round 1 has the input $u_1 = \mathrm{ic}_1 = 0$, so $D_1 = \Delta_1(0)$ and the first part with $q = 2$ and $q' = R$ gives the claim. □

Proof of Theorem 5. The recursion of Ψ_r has one form in a random round and another in a fixed round. In a random round r, Lemma 4 makes $x_1^{(r)}$ uniform given the other variables, so $\Psi_r(\alpha, \ell) = 0$ for $\alpha \neq 0$. For $\alpha = 0$, the recursions of Section 4.1 give $\zeta^{\ell x_2^{(r)}}(-1)^{D_r} = \zeta^{\ell x_2^{(r-1)}}(-1)^{D_{r-1}} g(u_r)$ with $g(z) = \zeta^{\ell F_r(z)}(-1)^{\Delta_r(z)} = \sum_\beta m_r(\ell; \beta)\zeta^{\beta z}$. Since g is fixed by $\mathcal{E}$ and $u_r = x_1^{(r-1)}$,

$$\Psi_r(0, \ell) = \sum_{\beta \in \mathbb{Z}_{16}} m_r(\ell; \beta)\, \Psi_{r-1}(\beta, \ell). \tag{22}$$

In a fixed round r, $x_1^{(r)} = u_r + S_{\mathrm{fix}}(y)$ with $y = x_2^{(r)} + \mathrm{rk}_1^{(r)}$, and expanding $\zeta^{\alpha S_{\mathrm{fix}}(y)} = \sum_\mu C^{S_{\mathrm{fix}}}(\alpha, \mu)\zeta^{\mu y}$ gives $\zeta^{\alpha x_1^{(r)} + \ell x_2^{(r)}} = \sum_\mu C^{S_{\mathrm{fix}}}(\alpha, \mu)\zeta^{\mu \mathrm{rk}_1^{(r)}} \zeta^{\alpha u_r + (\ell+\mu) x_2^{(r)}}$. We insert $x_2^{(r)} = x_2^{(r-1)} + F(u_r)$ and $D_r = D_{r-1} + \Delta(u_r)$ for the functions F and Δ of S_{fix} and expand $\zeta^{(\ell+\mu)F(z)}(-1)^{\Delta(z)}$ by equation (15). Each term then has the expectation $\Psi_{r-1}(\alpha+\beta, \ell+\mu)$, since a fixed round adds no randomness, and

$$\Psi_r(\alpha, \ell) = \sum_{\mu, \beta \in \mathbb{Z}_{16}} C^{S_{\mathrm{fix}}}(\alpha, \mu)\, \zeta^{\mu \mathrm{rk}_1^{(r)}}\, m_r^{\mathrm{fix}}(\ell + \mu; \beta)\, \Psi_{r-1}(\alpha + \beta, \ell + \mu).$$

Round $r-1$ is random, so only $\beta = -\alpha$ remains, and $\ell'' = \ell + \mu$ gives

$$\Psi_r(\alpha, \ell) = \sum_{\ell'' \in \mathbb{Z}_{16}} C^{S_{\mathrm{fix}}}(\alpha, \ell'' - \ell)\, \zeta^{(\ell'' - \ell)\mathrm{rk}_1^{(r)}}\, m_r^{\mathrm{fix}}(\ell''; -\alpha)\, \Psi_{r-1}(0, \ell''). \tag{23}$$

Round 1 is random with the public input $u_1 = 0$, so $x_2^{(1)} = \sigma_0 + F_1(0)$ and $D_1 = \Delta_1(0)$ are fixed by $\mathcal{E}$, and $\Psi_1(\alpha, \ell) = \mathbf{1}[\alpha = 0]\, \zeta^{\ell(\sigma_0 + F_1(0))} s_1$. By induction, $\psi = \Psi_r(0, \cdot)$ after every random round that the procedure reaches. If the rounds $r-1$ and r are random, then $\Psi_{r-1}(\beta, \ell) = 0$ for $\beta \neq 0$, and equation (22) is the multiplication by $m_r(\ell; 0)$. If round $r < R$ is fixed, equation (23) and equation (22) for round $r+1$ give $\Psi_{r+1}(0, \ell) = \sum_\beta m_{r+1}(\ell; \beta)\Psi_r(\beta, \ell) = (T_r\psi)(\ell)$. When round R is fixed, equation (23) with $\alpha = \ell = 0$ and $C^{S_{\mathrm{fix}}}(0, \mu) = \mathbf{1}[\mu = 0]$ gives $\Psi_R(0, 0) = m_R^{\mathrm{fix}}(0; 0)\, \Psi_{R-1}(0, 0) = c_{\mathrm{fix}}(t'_R)\, \psi(0)$. □

Two facts on the trails of Section 4.1 follow from Theorem 5. Every term has $\ell_B = 0$, because $\mathbb{E}[Z \mid \mathcal{E}]$ reads the component $\ell = 0$. Since $C^{S_{\mathrm{fix}}}(\beta, 0) = \mathbf{1}[\beta = 0]$, the entry $T_r[0, 0]$ is $c_{r+1}(t'_{r+1})\, c_{\mathrm{fix}}(t'_r)$.

Proof of Theorem 6. In a term $\mathcal{P}$, the factor of a random round r of segment b is $\psi(\ell_1)$ if $r = 1$, $m_r(\ell_b; \beta)$ if round $r-1$ is fixed and $m_r(\ell_b; 0)$ otherwise, and it depends only on the expanding S-box of round r. A fixed round contributes a constant and a sum over its β. Conditioning on $\mathcal{G}_{-r}$ and on the filter therefore fixes all factors but that of a random round r. We write $S(x) = p_x + 2T_x$ for $x \in \mathbb{Z}_8$, with the parity $p_x \in \{0, 1\}$ and $T_x \in \mathbb{Z}_8$.

Since the filter sees only the parities, the words T_x remain independent and uniform given $\mathcal{G}_{-r}$, the filter and $b^{(r)}$. We treat in turn a side trail with $\ell_b \notin \{0, 8\}$ for some segment b, one with (F_r) in a random round r of a segment with $\ell_b = 8$, and one with (S_{f+1}) in its return round.

First, let r be a random round of a segment with $\ell = \ell_b \notin \{0, 8\}$, whose factor is a sum over z of terms $\zeta^{\ell F_r(z)}$ times functions of $b^{(r)}$. By equation (4), $F_r(z) = \sum_{x \in \mathbb{Z}_8} n_x(z) S(x)$ with integers $n_x(z)$, where $n_x(z)$ counts the $l \geq 2$ with $z + \mathrm{rk}_l^{(r)} \equiv x$ minus those with $z + \mathrm{rk}_l^{(r)} \equiv x + 8$. Their sum has the parity of 31, so some $n_{x_0}(z)$ is odd. Then $\zeta^{\ell F_r(z)}$ contains the factor $\zeta^{2\ell n_{x_0}(z) T_{x_0}}$, whose mean over a uniform T_{x_0} is zero because $\ell n_{x_0}(z) \not\equiv 0 \pmod 8$. Since the other factors of $W\mathcal{P}$ are fixed by $\mathcal{G}_{-r}$, the filter and $b^{(r)}$, $\mathbb{E}[W\mathcal{P} \mid \Phi] = 0$.

Second, let (F_r) hold for a random round r of a segment with $\ell_b = 8$, whose factor is a sum over z of constants times $(-1)^{F_r(z) + \Delta_r(z)}$. Since $(-1)^{S(x+8)} = (-1)^{S(x)}$, the sign $(-1)^{F_r(z)}$ is a product of 31 entries of $b^{(r)}$, and $(-1)^{\Delta_r(z)} = b^{(r)}_{z+d_i} b^{(r)}_{z+d_j}$ with indices modulo 8. Hence every term changes sign under $b^{(r)} \mapsto -b^{(r)}$, whereas W and ν_r are invariant under this map by (F_r), and the conditional expectation of $W\mathcal{P}$ is zero.

Finally, let (S_{f+1}) hold for the return round, and let the side trail pass the fixed round f from $\ell = 8$ to $\ell = 0$ by $T_f[0, 8] = \sum_\beta m_{f+1}(0; \beta)\, C^{S_{\mathrm{fix}}}(\beta, 8)\, \zeta^{8\mathrm{rk}_1^{(f)}}\, m_f^{\mathrm{fix}}(8; -\beta)$. Only $\beta \notin \{0, 8\}$ contribute, because $C^S(0, 8) = 0$ for every function S and a direct computation from equation (3) gives $C^{S_{\mathrm{fix}}}(8, 8) = 0$. With $b = b^{(f+1)}$, the sign $(-1)^{\Delta_{f+1}(z)} = b_{z+d_i} b_{z+d_j}$ has period 8 in z, and

$$m_{f+1}(0; \beta) = \frac{1 + (-1)^\beta}{16} \sum_{u \in \mathbb{Z}_8} b_{u+d_i} b_{u+d_j}\, \zeta^{-\beta u},$$

which vanishes for odd β. For $\beta = 2\beta'$ with $\beta' \in \{1, 2, 3, 5, 6, 7\}$, the conditional expectation of $W_{f+1} m_{f+1}(0; \beta)$ is $\frac{1}{8} \sum_u \mathcal{L}_{f+1}(u) \zeta^{-2\beta' u}$, which is zero because $\mathcal{L}_{f+1}$ has period 2 and $\zeta^{4\beta'} \neq 1$. The other factors of $W\mathcal{P}$, including the constants of round f, are fixed by $\mathcal{G}_{-(f+1)}$ and the filter, and the sum over β is linear. If no random round follows segment b^*, then $f = R$, and the read-out needs $\ell_{b^*} = 0$, which is excluded. The last claim follows from Theorem 5 and equation (16), because W and the filter are functions of $\mathcal{E}$. □

For a weight $W = s_1(o_1) \prod_r W_r$ with even factors W_r of $b^{(r)}$ and distributions ν_r invariant under $b^{(r)} \mapsto -b^{(r)}$, (F_r) holds in every round, and Theorem 6 reduces $\mathbb{E}[WZ \mid \Phi]$ to the main trail. In a random round r on which the filter does not depend and with $W_r = c_r(h_r)$, the main trail has the factor $\mathbb{E}[c_r(h_r) c_r(t'_r)]$, which vanishes for $t'_r \notin [h_r]$ by Theorem 2(i). In round 1 without a filter, $\mathbb{E}[s_1(o_1) s_1]$ is the mean of a product of four signs of a uniform parity pattern, in which some index occurs an odd number of times when o_1 is in a wrong class. A wrong class in such a round or in round 1 therefore has zero expected score. These zeros are averages over the expanding S-boxes, and for a single block, the side trails have in general the size of the main trail.

For the key differences of a branch pair in Section 4.2, equation (5) makes $k \mapsto (t'_r)_{r \leq 40}$ a homomorphism $\mathbb{Z}_{16}^{64} \to \mathbb{Z}_{16}^{40}$ with a matrix A_{ij} over $\mathbb{Z}_{16}$. For each of the 465 branch pairs, a computation from the matrix M of equation (5) gives a 40×40 minor of A_{ij} whose determinant is odd, a unit of $\mathbb{Z}_{16}$. Hence A_{ij} is surjective, and since its fibres are the cosets of its kernel, it maps the uniform distribution to the uniform distribution. The six zeros in equation (17) give the probability $\frac{6}{16} = \frac{3}{8}$.

B.4 Proofs for Section 5

Proof of Lemma 2. The right guess has $s_1(o_1) = s_1$ and $c_r^u(h_r) = c_r^u(t'_r)$. With a filter, a round restricted to the constant patterns has $c_r^u \equiv 1$. For a round restricted to the constant or the alternating patterns, the sign of the right parity bit is $c_r^u(t'_r)$, because an alternating

pattern has $c(t) = (-1)^t$. Hence W_u is the main trail of equation (16) divided by Γ. For a value v, $\mathbf{1}[W_u = v]$ is a function of the parity patterns invariant under $b^{(r)} \mapsto -b^{(r)}$ in every round, because every factor of W_u is quadratic in the pattern of its round or depends only on whether this pattern is alternating.

The distribution ν_r is uniform without a filter and in an unfiltered round of a filter of the lemma, and in a restricted round it is invariant under this map by Section 4.2. Hence (F_r) holds for the indicator in every random round $r \geq 2$. Every segment of a side trail with $\ell_b = 8$ contains such a round, and Theorem 6 with this indicator in place of W gives $\mathbb{E}[Z_u \mathbf{1}[W_u = v]] = \Gamma v \Pr[W_u = v]$, the first claim. Given all weights, the Z_u are independent, since the blocks have independent S-boxes, with means ΓW_u and variances $1 - \Gamma^2 W_u^2$, which gives the moments of $\xi(H)$. In variant A2 with the right o_R, the observation is that of variant A1 on $R - 1$ rounds. □

Theorem 11 (Unconditional Bound on False Alarms). *Fix the master key, and let H be a wrong guess of variant A1 in the known-nonce model without a filter. Let the pairs (W_u, Z_u) of the N blocks be independent and identically distributed with $|W_u| \leq 1$ and $\mathbb{E}[W_u Z_u] = 0$. For $L > 0$, let $p_L = \Pr[|W_u| > L]$, $q_L = \mathbb{E}\big[|W_u| \, \mathbf{1}[|W_u| > L]\big]$ and $V_L = N \, \mathbb{E}\big[W_u^2 \, \mathbf{1}[|W_u| \leq L]\big]$, and let $\omega(y) = (1 + y)\ln(1 + y) - y$. For an integer J with $1 \leq J \leq \theta^2 + 1$ and $\delta \in (0, 1)$, let $\theta_J = \sqrt{\theta^2 - J + 1}$ and $\gamma = \theta_J \sqrt{(1 - \delta) V_L} - N q_L$. If $\gamma > 0$, then*

$$\Pr\big[|\xi(H)| \geq \theta\big] \leq \Big(\frac{eNp_L}{J}\Big)^J + 2\exp\Big(-\frac{V_L}{L^2}\,\omega\Big(\frac{\gamma L}{V_L}\Big)\Big) + \exp\Big(-\frac{\delta^2 V_L}{2L^2}\Big).$$

Proof. Let $X_u = W_u Z_u$, and split Σ_1 and Σ_2 into the sums $\Sigma_1^>, \Sigma_2^>$ over the set $U_>$ of blocks with $|W_u| > L$ and $\Sigma_1^\leq, \Sigma_2^\leq$ over the others. Let $\xi_\leq = \Sigma_1^\leq / \sqrt{\Sigma_2^\leq}$, with $\xi_\leq = 0$ when $\Sigma_2^\leq = 0$. The Cauchy–Schwarz inequality gives $|\Sigma_1^>| \leq \sqrt{|U_>| \, \Sigma_2^>}$ because $|X_u| = |W_u|$, and a second application to the vectors $(|\xi_\leq|, \sqrt{|U_>|})$ and $(\sqrt{\Sigma_2^\leq}, \sqrt{\Sigma_2^>})$ gives $|\Sigma_1| \leq \sqrt{\xi_\leq^2 + |U_>|}\,\sqrt{\Sigma_2}$. Hence $|\xi(H)| \leq \sqrt{\xi_\leq^2 + |U_>|}$, and since $|U_>|$ is an integer, $|\xi(H)| \geq \theta$ implies $|U_>| \geq J$ or $|\xi_\leq| \geq \theta_J$.

Since $|U_>|$ is binomial with the parameters N and p_L, $\Pr[|U_>| \geq J] \leq \binom{N}{J} p_L^J \leq (eNp_L/J)^J$. If $|\xi_\leq| \geq \theta_J$ and $\Sigma_2^\leq > (1 - \delta) V_L$, then $|\Sigma_1^\leq| \geq \theta_J \sqrt{(1 - \delta) V_L}$. The terms $Y_u = X_u \mathbf{1}[|W_u| \leq L]$ of $\Sigma_1^\leq$ are independent, with $|Y_u| \leq L$, $\sum_u \mathbb{E}[Y_u^2] \leq V_L$ and means $-\mathbb{E}[X_u \mathbf{1}[|W_u| > L]]$ of absolute value at most q_L, since $\mathbb{E}[X_u] = 0$. Hence $\pm\Sigma_1^\leq$ deviates from its mean by at least γ. Bennett's inequality, in the form given by Boucheron et al. [BLM13], bounds each sign by $\exp(-\frac{V_L}{L^2}\omega(\frac{\gamma L}{V_L}))$. It applies with the sum of the second moments in place of the variance, and its bound does not decrease when this sum grows to V_L. For the terms $T_u = W_u^2 \mathbf{1}[|W_u| \leq L] \in [0, L^2]$ of $\Sigma_2^\leq$, whose means sum to V_L, $e^{-y} \leq 1 - y + y^2/2$ for $y \geq 0$ gives $\ln \mathbb{E} e^{-\lambda(T_u - \mathbb{E}T_u)} \leq \lambda^2 \mathbb{E}[T_u^2]/2$ for $\lambda \geq 0$. The Chernoff bound then gives $\Pr[\Sigma_2^\leq \leq (1 - \delta) V_L] \leq \exp\big(-\delta^2 V_L^2 / (2 \sum_u \mathbb{E}[T_u^2])\big) \leq \exp(-\delta^2 V_L / (2L^2))$, and the three bounds give the claim. □

Theorem 11 bounds $\Pr[|\xi(H)| \geq \theta]$ for every wrong guess H of variant A1 in the known-nonce model without a filter and does not need $\mathbb{E}[Z_u \mid W_u] = 0$. By Section 4.2, every such wrong guess satisfies its hypotheses. The bound is far from the Gaussian tail for guesses with a small $V = N \, \mathbb{E}[W_u^2]$, which dominate the union over $\mathcal{H}$.

Proof of Theorem 7. For (i), let $W^\star = s_1 \prod_r c_r(t'_r)$ be the weight of the right guess, let $h_r^\star$ be the class of t'_r, and let $X = WZ$ for one block. Since $Z = \pm 1$ and $\mathbb{E}[Z \mid \mathcal{E}] = \Psi_R(0, 0)$, every v has $\Pr[X = v] - \Pr[X = -v] = \mathbb{E}\big[\Psi_R(0, 0)\big(\mathbf{1}[W = v] - \mathbf{1}[W = -v]\big)\big]$. Without a filter, every ν_r is uniform, and $\mathbf{1}[W = v] - \mathbf{1}[W = -v]$ is invariant under $b^{(r)} \mapsto -b^{(r)}$

in every round, because each factor $s_1(o_1)$ and $c_r(h_r)$ of W is quadratic in the pattern of its round. Hence (F_r) holds in every round, and since Theorem 6 removes all side trails, the main trail $\Gamma W^\star$ gives the difference $\Gamma(\lambda(v) - \lambda(-v))$ with the signed measure $\lambda(v) = \mathbb{E}[W^\star \mathbf{1}[W = v]]$.

With $\lambda_r(v) = \mathbb{E}[W_r^\star \mathbf{1}[W_r = v]]$ for the factors $W_r^\star$ and W_r of round r in $W^\star$ and W, the independence of the rounds makes λ the multiplicative convolution of the λ_r. It is even or odd, $\lambda(-v) = \pm\lambda(v)$, if one λ_r is even or odd, and zero if both occur. A map of the patterns of round r that preserves their distribution makes λ_r even, zero or odd when it changes the sign of W_r alone, of $W_r^\star$ alone or of both. The sign flips $b \mapsto (\epsilon_x b_x)_x$ with $\epsilon \in \{\pm 1\}^8$ preserve the uniform distribution, and by equation (13), a flip changes the sign of $c(t)$ if $\epsilon_x \epsilon_{x+t} = -1$ for all x and keeps $c(t)$ if $\epsilon_x \epsilon_{x+t} = 1$ for all x. Three flips act on $c(0) = 1$ and on $c(1), \dots, c(4)$ as follows. The flip $\epsilon = (-1)^x$ changes the signs of $c(1)$ and $c(3)$ and keeps $c(0)$, $c(2)$ and $c(4)$. The flip $\epsilon = (1, 1, -1, -1, 1, 1, -1, -1)$ changes the sign of $c(2)$ and keeps $c(0)$ and $c(4)$. Finally, the flip $\epsilon = (1, 1, 1, 1, -1, -1, -1, -1)$ changes the sign of $c(4)$ and keeps $c(0)$.

In a right round $r \geq 2$, a flip that changes the sign of $c(h_r)$ makes λ_r odd for $h_r \neq 0$. For $h_r = 0$, λ_r is the point mass at 1, the unit of the convolution. In a wrong round with $\{h_r, h_r^\star\} \neq \{1, 3\}$, one of the three flips changes the sign of c at one of the two classes and keeps it at the other. This flip makes λ_r even if it changes $c(h_r)$ and zero if it changes $c(h_r^\star)$. For $\{h_r, h_r^\star\} = \{1, 3\}$, the flip $(-1)^x$ changes both signs, and λ_r is odd. In round 1, let o_1 be wrong. If $v_i \neq v_j$, the flip of $b_{x_0}^{(1)}$ alone, for some $x_0 \in \{v_i, v_j\} \setminus \{v_i', v_j'\}$, changes the sign of W_1 and keeps $W_1^\star$, so λ_1 is even. If $v_i = v_j$, then $W_1 = 1$ and $v_i' \neq v_j'$, and the flip of $b_{v_i'}^{(1)}$ makes λ_1 zero. A right o_1 gives an odd λ_1 when $v_i' \neq v_j'$ and the point mass at 1 otherwise. Hence, for a wrong guess $H \notin \mathcal{H}_{13}$, some λ_r and therefore λ are even or zero, and X is symmetric.

For $H \in \mathcal{H}_{13}$, every λ_r is odd or the point mass at 1, and at least one is odd, so λ is odd. Since $|c(t)| \leq 1$, the weight $W = 1$ needs $W_r = \pm 1$ in every round, and the n odd factors give $\lambda(1) = 2^{n-1} \prod_r \lambda_r(1)$ over these factors. A right o_1 with $v_i' \neq v_j'$ has $\lambda_1(1) = \frac{1}{2}$. In a round $r \geq 2$, $\lambda_r(1) > 0$: for a right class, $\lambda_r(1) = \Pr[c(h_r) = 1]$, and for $\{h_r, h_r^\star\} = \{1, 3\}$, the constant patterns are the only patterns with $c(h_r) = 1$ and have $c(h_r^\star) = 1$ as well. Hence $\lambda(1) - \lambda(-1) = 2\lambda(1) \neq 0$, and X is not symmetric when $\Gamma \neq 0$. When $\Gamma = 0$, the difference vanishes for every guess.

For (ii), let $H \notin \mathcal{H}_{13}$. The variables $X_u = W_u Z_u$ are independent and symmetric by (i), and $\xi(H) = \sum_u X_u / \sqrt{\sum_u X_u^2}$ because $Z_u^2 = 1$. Pinelis [Pin12] proved that such a self-normalised sum satisfies $\Pr[\xi \geq y] \leq Q(y) + C_{\mathrm{P}} \varphi(y)/(9 + y^2)$, and the symmetry of ξ gives the factor 2. With $\xi = 0$ for a zero denominator, the bound follows by conditioning on $|X_u|$. The guesses of $\mathcal{H}_{13}$ have zero expected score by Section 4.2, so Theorem 11 applies to them, and the union over $\mathcal{H}$ gives (iii). □

Proof of Proposition 2. With $\pm$ the sign of Γ and given the weights, $\pm\xi(H) - \mu_W = \sum_u (\pm W_u Z_u - |\Gamma| W_u^2)/\sqrt{\Sigma_2}$ is a sum of independent terms with mean zero by Lemma 2, in intervals of the lengths $2|W_u|/\sqrt{\Sigma_2}$, whose squares sum to 4. Hoeffding's inequality, in the form given by Boucheron et al. [BLM13], bounds $\Pr[|\xi(H)| < \theta] \leq \Pr[\pm\xi(H) < \theta]$ by $\exp(-2(\mu_W - \theta)^2/4)$. The terms $W_u^2 \in [0, 1]$ of Σ_2 are independent. The rounds whose class is guessed have no filter and are independent of each other and of the filter, and the other factors of W_u have the absolute value 1, so $\mathbb{E}[W_u^4] = \prod_r \rho_r \, \mathbb{E}[W_u^2]$. With $t = \delta V$, the lower tail $\Pr[\Sigma_2 \leq V - t] \leq \exp\bigl(-t^2/(2N\mathbb{E}[W_u^4])\bigr)$ of the proof of Theorem 11 gives the second claim. □

Proof of Theorem 8. In variant A1, $Z_u = (-1)^{D_R}$ with $D_R = z_i - z_j - \mathrm{ic}_i + \mathrm{ic}_j$ by equations (6) and (10), and $z_i - z_j$ is uniform on $\mathbb{Z}_{16}$ for a uniform block z. In variant A2, the stripped value $x_1^{(R-1)} = z_1 - S_{\mathrm{crf}}^{(R)}\bigl(\mathrm{rk}_1^{(R)} + \sum_{l \geq 2} z_l\bigr)$ of equation (2), with the

guessed $\mathrm{rk}_1^{(R)}$, is a bijection of z_1 for fixed $z_2, \ldots, z_{32}$. Hence $(x_1^{(R-1)}, z_2, \ldots, z_{32})$ is uniform, and since z_i is uniform given $x_1^{(R-1)}$ and the other words, the stripped difference $D_{R-1} = D_R - S_{\mathrm{erf}}^{(R)}\big(x_1^{(R-1)} + \mathrm{rk}_i^{(R)}\big) + S_{\mathrm{erf}}^{(R)}\big(x_1^{(R-1)} + \mathrm{rk}_j^{(R)}\big)$ is uniform for every guessed $\mathrm{rk}_1^{(R)}$, $\mathrm{rk}_i^{(R)}$ and $\mathrm{rk}_j^{(R)}$ and every $2 \le i < j \le 32$. Fix the S-boxes of the N blocks, those of round R included, and the guess, which fixes the inputs, the filter, the weights and the signs of the parity bits by the hypothesis. In the ideal world, the blocks of these distinct inputs are then independent and uniform, so the Z_u are independent fair signs given the weights, which is (i).

Given the weights, $\Sigma_1 = \sum_u W_u Z_u$ is then a sum of independent terms with mean zero, in intervals of the lengths $2|W_u|$ whose squares sum to $4\Sigma_2$. Hoeffding's inequality, in the form given by Boucheron et al. [BLM13], gives $\Pr\big[|\Sigma_1| \ge \theta\sqrt{\Sigma_2} \;\big|\; (W_u)_u\big] \le 2\exp\big(-2\theta^2\Sigma_2/(4\Sigma_2)\big) = 2\exp(-\theta^2/2)$ when $\Sigma_2 > 0$. Since $\xi(H) = 0 < \theta$ when $\Sigma_2 = 0$, the first claim of (ii) follows. After averaging over the weights, this bound holds for each guess of $\mathcal{H}$. The union bound over $\mathcal{H}$ then gives the second claim, with no assumption on the joint distribution of the statistics of different guesses.

Finally, let the master key be uniform. For a key of the set, Proposition 2 applies to the right guess in $\mathcal{H}$ of its branch pair, since the filter is as in Lemma 2, and this guess has $\mathbb{E}[\mu_W^2] = \Gamma^2 V \ge m^2$ and a ratio $V/\prod_r \rho_r$ at least the smallest over the set. Outside the lower tail $\Sigma_2 \le (1-\delta)V$, whose probability is at most $\exp\big(-\delta^2 V/(2\prod_r \rho_r)\big)$ and hence at most the second term of p_{hit}, the value $\mu_W = |\Gamma|\sqrt{\Sigma_2}$ exceeds $m\sqrt{1-\delta} > \theta$. Given such weights, the first claim of Proposition 2 bounds the probability that $|\xi(H)| < \theta$ by $\exp\big(-(m\sqrt{1-\delta} - \theta)^2/2\big)$. The right guess therefore passes θ with probability at least p_{hit}, and since D_θ outputs FRAST whenever it passes, $\Pr_{\mathrm{real}}[D_\theta = FRAST] \ge d\, p_{\mathrm{hit}}$. By (ii), $\Pr_{\mathrm{ideal}}[D_\theta = FRAST] \le 2|\mathcal{H}|\exp(-\theta^2/2)$, and $\mathrm{Adv}(D_\theta)$ is the difference of the two probabilities. □

B.5 Contract and Key Recovery

A key of the table $\mathcal{T}_m$ of Contract consists of an outer cell, restricted bits, classes $h_{r_1}, \ldots, h_{r_m}$ and the class vectors of the rounds $r_{m+1}, \ldots, r_x$. For a key of $\mathcal{T}_m$ with the classes $h_{r_1}, \ldots, h_{r_m}$, the value is the sum of $Z_u^{(R)} \prod_{l \le m} c_{r_l}^u(h_{r_l})$ over the blocks u that agree with the key in the outer cell, the restricted bits and the class vectors of $r_{m+1}, \ldots, r_x$. This follows by induction on m, since step 2 skips only the classes with zero values and sums equal keys linearly. Lines 9 to 21 of Algorithm 1 multiply these values by the signs of the first round, of the parity bits and of round R in variant A2, and give Σ_1 for every guess and Σ_2 for every class guess. These signs depend on a block only through its outer cell and restricted bits.

Proof of Proposition 3. A block contributes to at most $\prod_{l \le m} n(c_{r_l}^u)$ keys of $\mathcal{T}_m$. Its class vectors in different rounds are independent, because the S-boxes of different rounds are independent and the filter does not depend on the rounds $r_1, \ldots, r_x$, so the expectation of this product is $\bar{n}^m$. The second bound counts the keys, each of which consists of one of $|\mathcal{B}|$ outer cells, w restricted bits, m classes and $x - m$ class vectors. □

Proposition 5 bounds the work of the steps of Contract without the model of Section 6.1. Section 7.2 applies it to step 2 on the parameter sets that fix the model constant and gives the margin at $w = 5$ in Section 7.3 under this bound.

Proposition 5 (Work of the Contraction Without the Model). *Under the assumptions of Proposition 3, the expected number of multiply-adds of step m of Contract is at most*

$$\min\big(N\bar{n}^m,\ 18\,|\mathcal{B}|\,2^w 5^m 18^{x-m}\big).$$

Proof. Every entry of $\mathcal{T}_{m-1}$ has a block that contributes to it, and the class vector of round r_m in the key is that of the block. We map every multiply-add of step m to a block of its entry and the classes of the rounds $r_1, \dots, r_m$. This map is injective, because for given classes a block belongs to one key of $\mathcal{T}_{m-1}$. The products of step m are therefore at most as many as the pairs of a block and classes of the rounds $r_1, \dots, r_m$ with nonzero values for this block. By the proof of Proposition 3, the expected number of these pairs is $N\bar{n}^m$. In addition, every entry produces at most five products, and $\mathcal{T}_{m-1}$ has at most $|\mathcal{B}|\, 2^w 5^{m-1} 18^{x-m+1}$ keys. □

Prefix enumeration in Section 6.1 rests on the fact that, in Contract, a key of $\mathcal{T}_m$ with $m \geq j$ receives only the blocks with nonzero class values at its classes of the rounds $r_1, \dots, r_j$, each with the factor $\prod_{l \leq j} c^u_{r_l}(h_{r_l})$. The run of a prefix starts from these blocks and factors and applies step 2 to the rounds $r_{j+1}, \dots, r_x$ in the order of Contract. Its table at step m is therefore the part of $\mathcal{T}_m$ with the classes of the prefix. Step 3 treats every class guess separately and gives the same statistics. After the sort, a block is read once for each of the $\prod_{l \leq j} n(c^u_{r_l})$ prefixes with nonzero values for its class vectors, whose expected number is $\bar{n}^j$ by the independence in the proof of Proposition 3.

The storage of prefix enumeration is bounded as follows. The keys of the run of a prefix do not contain the classes of the rounds $r_1, \dots, r_j$, and the prefix $(0, \dots, 0)$ keeps every block, because $c(0) = 1$. At every step $m \geq j$, the table of a prefix is therefore a subset of that of $(0, \dots, 0)$. With $j = \lfloor m^* \rfloor$, every block gives one key at step j, so this table has at most N entries. By equation (21), its keys number $|\mathcal{B}|\, 2^w 18^{x-j} = N\, 18^{m^*-j} (\bar{n}/5)^{m^*} < 18N$, and since every later step multiplies their number by $\frac{5}{18}$, every later table has fewer than $5N$ entries. Prefix enumeration therefore needs $O(N)$ words, with a constant of at most five, for every parameter set and every attack in Tables 3 and 4.

A wrong pair differs from the right pair of its branch pair in the first-round guess or in the class of some round, and by the constraints of these guesses, the claim of Section 6.2 on wrong pairs follows. The constraint of o_1 holds for the true key if and only if $o_1 = \{\mathrm{rk}_i^{(1)} \bmod 8, \mathrm{rk}_j^{(1)} \bmod 8\}$, or $\mathrm{rk}_i^{(1)} \equiv \mathrm{rk}_j^{(1)} \pmod 8$ for the guess with $v_i = v_j$, and the constraint of h_r holds if and only if $t'_r \in [h_r]$, by Theorem 2(i).

The size of $c_{\mathrm{fix}}(t'_r)$ depends on the bit 3 of t'_r when $t'_r \bmod 8 \in \{0, 2, 6\}$, and $\hat{g}$ has a standard deviation of at most $\sigma_{\hat{g}} = 1/\sqrt{\Sigma_2}$ by Lemma 2. An accepted pair keeps the vectors of the bits 3 of t'_r in its fixed rounds that give Γ the sign of Σ_1 and a size within $3\sigma_{\hat{g}}$ of $\hat{g}$. For $\delta \in (0, 1)$ and the right pairs of K branch pairs fixed before the data, the tolerance $t_K \sigma_{\hat{g}}$ with $t_K = \sqrt{2 \ln(2K/\delta)}$ covers their true sizes $|\Gamma|$ at once with probability at least $1 - \delta$, without independence between the pairs. Given the weights, $\Sigma_1/\Sigma_2 - \Gamma$ is a sum of independent terms with mean zero by Lemma 2, in intervals whose squared lengths sum to $4\sigma_{\hat{g}}^2$. Hoeffding's inequality therefore bounds the probability that $\left|\hat{g} - |\Gamma|\right| \geq t_K \sigma_{\hat{g}}$ for one pair by $2\exp(-t_K^2/2) = \delta/K$, and the union bound over the K pairs gives the claim. With all 465 branch pairs, the tolerance covers every set of accepted right pairs, whichever pairs the data select.

A bit 3 of a form, or a sum of two such bits, that is constant on these vectors gives an equation over $\mathbb{Z}_2$ on $k_{[3]}$. These equations have different coefficients in different fixed rounds, so their rank is not bounded by 30. When no value of $k_{[3]}$ remains or the keystream test of Section 6.2 fails, the tolerance is widened to $4\sigma_{\hat{g}}$. The executed recovery of Section 7.1 used $3\sigma_{\hat{g}}$ for its 77 accepted pairs without widening it, and at $M = 77$, independent normal estimates would miss this tolerance at least once with probability about 0.2.

B.6 The Variant

Proof of Theorem 9. In (i), the period holds because a parity pattern has the period 8. For the second moments, the independence of the two patterns gives $\mathbb{E}[c_r(s)c_r(t)] =$

$\frac{1}{64}\sum_{x,y}\mathbb{E}[b_{x+s}^{(r,i)}b_{y+t}^{(r,i)}]\,\mathbb{E}[b_x^{(r,j)}b_y^{(r,j)}]$, and only the terms with $y = x$ remain, which give $\frac{1}{8}\mathbf{1}[s \equiv t \pmod 8]$. For (ii), given $b^{(r,i)}$, the eight products $b_{x+s}^{(r,i)}b_x^{(r,j)}$ are independent uniform signs with the sum $8c_r(s)$. This sum is ± 8 with the probability 2^{-7}, that is, if and only if $b_{x+s}^{(r,i)} = \pm b_x^{(r,j)}$ for all x with one sign. For (iii), let $|c_r(s)| = 1$ for all s. By the proof of (ii), $b_{x+s}^{(r,i)} = \epsilon_s b_x^{(r,j)}$ for all x and s, with signs ϵ_s. Comparing the relations for $s+1$ and for s at $x+1$ gives $b_{x+1}^{(r,j)} = \epsilon_s\epsilon_{s+1}b_x^{(r,j)}$. Hence $b^{(r,j)}$ is constant or alternating, and so is $b^{(r,i)} = \epsilon_0 b^{(r,j)}$. Conversely, two constant patterns give $c_r(s) = \pm 1$, two alternating ones give $c_r(s) = \pm(-1)^s$, and a constant and an alternating one give $c_r \equiv 0$. A pattern is constant or alternating with the probability 2^{-6}, and two such patterns have the same type with the probability $\frac{1}{2}$, which gives 2^{-13}. □

The proof of Proposition 4 uses the round of the variant in the alternative representation. The contracting S-boxes, the fixed rounds with S_{fix} and the key schedule are those of FRAST. In the variant, the alternative representation of equation (6) keeps its form with $F_r(z) = \sum_{l=2}^{32} S_l^{(r)}(z + \mathrm{rk}_l^{(r)})$ and $\Delta_r(z) = S_i^{(r)}(z + \mathrm{rk}_i^{(r)}) - S_j^{(r)}(z + \mathrm{rk}_j^{(r)})$.

Proof of Proposition 4. Lemma 4 holds in the variant with the same proof, because the contracting S-boxes are those of FRAST. The proof of Theorem 5 uses the expanding S-boxes only through the functions $z \mapsto \zeta^{\ell F_r(z)}(-1)^{\Delta_r(z)}$, which they fix. Hence Theorem 5, the main and the side trails of Section 4.1 and equation (16) hold for the variant, with $\mathcal{E}$ the tuple of all S-boxes $S_l^{(r)}$ and with m_r computed from its functions F_r and Δ_r. The substitution of Section 8 gives $m_r(0;0) = c_r(t'_r)$, and round 1 has the input $u_1 = 0$ and the sign $s_1 = (-1)^{\Delta_1(0)} = b_{v'_i}^{(1,i)}b_{v'_j}^{(1,j)}$.

Expanding the products of Theorem 5 writes each trail $\mathcal{P}(\ell_1, \dots, \ell_B)$ as a sum, over the indices β of the fixed rounds, of terms that are a constant of S_{fix}, the round keys and the indices times one factor f_r for each random round $r \le R$. For r in the segment b, f_r is $\psi(\ell_1)$, $m_r(\ell_b;\beta)$ or $m_r(\ell_b;0)$ as in the proof of Theorem 6, a linear combination with constant coefficients of the values $g_r(z) = \zeta^{\ell_b F_r(z)}(-1)^{\Delta_r(z)}$ for $z \in \mathbb{Z}_{16}$ by equation (15). Given $\mathcal{C}$, the remaining randomness of the expanding S-boxes consists of the S-boxes $S_l^{(r)}$ with $l \notin \{i,j\}$ and of the values of $S_i^{(r)}$ and $S_j^{(r)}$ beyond their parities. This randomness is independent between the rounds and of $\mathcal{C}$, because the S-boxes are independent and the values of a uniformly random negacyclic S-box at $0, \dots, 7$ are independent and uniform. The filter is a function of $\mathcal{C}$, so the conditional expectation of each product is $\prod_r \mathbb{E}[f_r \mid \mathcal{C}]$ times its constant factor.

Let $\mathcal{P}$ be a side trail, and let b be a segment with $\ell_b \ne 0$. Since $\ell_B = 0$, the segment b precedes the last segment, and all its rounds are at most R. For a random round r of segment b and a branch $l_0 \notin \{i,j\}$, the value $g_r(z)$ is the product of $\zeta^{\ell_b S_{l_0}^{(r)}(z+\mathrm{rk}_{l_0}^{(r)})}$ and of factors that are independent of $S_{l_0}^{(r)}$ given $\mathcal{C}$. A uniformly random negacyclic S-box takes a uniform value at every point by equation (4), and $\sum_{y\in\mathbb{Z}_{16}}\zeta^{\ell_b y} = 0$ for $\ell_b \ne 0$. Hence $\mathbb{E}[g_r(z) \mid \mathcal{C}] = 0$ for every z, so $\mathbb{E}[f_r \mid \mathcal{C}] = 0$, and the side trail has the conditional expectation zero. Only the main trail remains, and it is a function of $\mathcal{C}$. Since $\mathbb{E}[Z \mid \mathcal{E}] = \Psi_R(0,0)$ is the sum of the trails, and $\mathcal{C}$ and the filter are functions of $\mathcal{E}$, this gives the first claim. Without a filter, the patterns of different rounds and branches are independent and uniform, and the first claim gives $\mathbb{E}[WZ] = \Gamma\,\mathbb{E}[b_{v_i}^{(1,i)}b_{v'_i}^{(1,i)}]\,\mathbb{E}[b_{v_j}^{(1,j)}b_{v'_j}^{(1,j)}]\prod_r \mathbb{E}[c_r(h_r)\,c_r(t'_r)]$. The first two factors are $\mathbf{1}[v_i = v'_i]$ and $\mathbf{1}[v_j = v'_j]$, and Theorem 9(i) gives $\frac{1}{8}\mathbf{1}[h_r \equiv t'_r \pmod 8]$ for each of the others. □

C Numerical Tables and Further Computations

In this appendix, we give the dependence of the weak key spectrum of full FRAST on the unit of the online time and the extrapolation of the constant of the implementation of Contract. We then give the rigorous thresholds and the fractions under the strict condition. The figures of Appendices C.1, C.2, C.4 and C.5 are computed, in the cost model or from the bounds of Section 5, apart from the measured timings of Appendix C.1, and no attack on full FRAST was run. Appendices C.4 and C.5 give the figures of joint detection, of side trails through a fixed round with $c_{\mathrm{fix}}(t'_r) = 0$ and of the chosen-nonce filters of the variant, and Appendix C.3 lists further experiments of Section 7.4.

With the strict condition of Section 5.2, the success probability is at least $0.9 - 0.01 - \varepsilon_{\mathrm{fa}}$ for every fixed master key, under the extensions of Appendix C.2 for variant A2 and the chosen-nonce attacks. The cap then falls from 57 to 56 at 26 rounds and at 28 rounds in variant A1. The word-level unit drops the charge of $2^{6.6}$ bit operations per multiply-add. The constant of the implementation, in time per unit $N\bar{n}^{m^*}$, was measured for $x \le 14$ in Appendix C.3 and is extrapolated to $x \ge 15$ in Appendix C.1, up to $x = 26$ for full FRAST with $w = 5$.

C.1 Units and Constant of the Online Time

We first derive the online time of Section 6.1 under the assumption of its model of step 2. Under the assumption, summing the bound of Proposition 3 over the steps gives at most $5.6\, N\bar{n}^{m^*}$ multiply-adds for step 2 and one branch pair. Step 3 costs $5^x 2^w 2^{11}$ multiply-adds for one branch pair in variant A1, and step 4 is negligible.

The weak key spectrum of Section 7.3 is stated in the main unit, in which the criterion holds for $w \in [5, 9]$. For four units of the online and the offline time, we give the restricted rounds w for which the criterion of the spectrum holds, with the partition (w_a, w_c) that minimises the criterion in that unit. Here, T_{on} is the online time of the model of Section 6.1 with $c_{\mathrm{on}} = 3$ for all 465 branch pairs, d_w is the density of Table 4, and a range $[a, b]$ contains the integers from a to b. With the factor 1, that is, in word-level multiply-adds with the margin, the criterion holds for $w \in [3, 9]$. It holds for $w \in [0, 9]$ with T_{on} in evaluations of FRAST. With the best partition for that unit, it holds only for $w = 5$ with both times in bit operations, at $2^{18.5}$ bit operations per XOF call, and $w = 10$ satisfies it in no unit.

The factor $2^{6.6}$ enters only T_{on}, which dominates the criterion for $w \le 7$. With the partitions of Table 4, the row $w = 5$ satisfies the criterion as long as one word-level multiply-add is charged at most $2^{9.11}$ bit operations, or $2^{9.46}$ in the exact expectation model. For $w = 4$, the factor must be at most $2^{5.45}$ in the model and $2^{5.81}$ in the exact expectation model, which are both below $2^{6.6}$. For $w \ge 8$, the offline time already exceeds the online time at the factor $2^{6.6}$, by 0.41 bits at $w = 8$, 7.85 bits at $w = 9$ and 15.30 bits at $w = 10$.

Without the counting of the table sizes in the calibration of Appendix C.3, a unit $N\bar{n}^{m^*}$ of step 2 of Contract at $x = 12$ took 9.49 and 9.68 ns in two runs with 2^{22} blocks. At 2.1 and 4.0 GHz, this is 2.3 and 1.3 bits below the 96 multiply-adds of the model, and a trial with 2^{20} blocks took 20.83 ns, which is 1.1 and 0.2 bits below. At equal load, a unit took 1.123 times as long at $x = 13$ as at $x = 12$, in one measurement. Extrapolated by the factor $1.123^{14} = 2^{2.34}$ of the 14 rounds from $x = 12$ to $x = 26$, the two runs give a constant 0.1 to 1.0 bits above the model constant. The trial at 4.0 GHz, which is the least favourable point of the server, gives a constant 2.14 bits above it.

Table 7: Thresholds and largest detectable costs at 2^{64} blocks, with Basis the statement or the extension that gives the row.

Rounds	Model	Variant	Basis	x	w	θ_G	θ^*	Data	$64 - \log \kappa'$	$64 - \log \kappa''_\delta$
26	known	A1	Theorem 7	20	0	9.116 3	9.210 8	0.026	57.218	56.989
26	known	A2	extension	19	0	9.753 2	9.838 5	0.022	57.050	56.834
28	known	A1	Theorem 7	22	0	9.459 0	9.548 3	0.024	57.127	56.904
28	known	A2	extension	21	0	10.074 7	10.155 8	0.021	56.969	56.758
35	chosen	A2	extension	17	10	10.110 9	10.191 6	0.020	56.960	56.750
36	chosen	A2	extension	17	10	10.123 1	10.203 7	0.020	56.957	56.747
38	chosen	A2	extension	19	10	10.433 4	10.510 3	0.019	56.881	56.676

C.2 Thresholds of the Strict Condition

Two extensions of Theorem 7 are not proved, and its proof and an exact computation over the parity patterns, reported in Section 7.4, support them. In variant A2, the last-round guesses of a wrong pair $(o_1, h) \notin \mathcal{H}_{13}$ give symmetric distributions. For $(o_1, h) \in \mathcal{H}_{13}$, the asymmetry is proportional to the coefficient $K_0(o_R)$ of the main trail of the first $R-1$ rounds, so the second term of Theorem 7(iii) counts only the o_R with $K_0(o_R) \neq 0$. The conclusions also hold for chosen-nonce filters that keep the constant, or the constant and the alternating patterns, with $\mathcal{H}_{13}$ replaced by its guesses with every parity bit right. This needs independent filters in the rounds, no filter in round 1, uniform kept patterns and independent blocks.

At a design point, a branch pair of cost c of variant A1 in the known-nonce model receives the data $N = \kappa' 2^c$ with $\kappa' = (\theta^* + 1.28)^2$. Here θ^* is the rigorous threshold at which the bound on false alarms meets the target $\varepsilon_{\text{fa}} = 0.1$ for the worst master key, and θ_G is the Gaussian threshold, with $\kappa_G = (\theta_G + 1.28)^2$. At the six design points, with $x = 7$ to 23 and the guesses of all 465 branch pairs, $\log(\kappa'/\kappa_G)$ is 0.16 to 1.41 bits with Theorem 11 alone and 0.023 to 0.739 bits with Theorem 7. The strict condition has the constant $\kappa''_\delta = \kappa''/(1-\delta)$ of Section 5.2, and $\log(\kappa''_\delta/\kappa_G)$ is 0.242 to 1.015 bits.

Table 7 gives the thresholds at $N = 2^{64}$ for the reach of Section 7.2, for the same target and the worst master key. The largest detectable cost is $64 - \log \kappa'$ under the Gaussian success condition and $64 - \log \kappa''_\delta$ under the strict condition, and its integer part is the cap of the cost model. Model is the known-nonce or the chosen-nonce model, and the attacks of variant A2 on 26, 28, 36 and 38 rounds strip a random last round. The number w counts the restricted rounds with constant or alternating patterns, which have one parity guess each. The threshold θ^* is that of Theorem 7, and Data is $\log(\kappa'/\kappa_G)$ in bits. For variant A2 and the chosen-nonce attacks, θ^* uses the extensions of Theorem 7 stated above.

Under the strict condition, 2^{64} blocks cover 99.8% of the sampled keys at 26 rounds in variant A2 and 94.2% in variant A1. At 28 rounds, they cover 41.4% in variant A1, and the 72% of variant A2 do not change. With 0.99 per pair, the strict condition covers 99.3% of the keys in key recovery at 24 rounds. For the chosen-nonce attacks and full FRAST, the cap is 56 under both conditions, so their fractions and their densities do not change.

C.3 Further Experiments

We report the further experiments of Section 7.4, with the scale of each and its result. The Gaussian model of Proposition 1 was compared with 1000 runs on five rounds in variant A2 and with 700 runs in variant A1 on ten rounds, 100 master keys and seven data sizes. In every bin of $\log(N \cdot \text{ELP}_{\text{pair}})$, the fraction of runs that ranked the right guess first lay within three standard errors of the model. On ten rounds, all 275 runs with $N \cdot \text{ELP}_{\text{pair}} \geq 2^6$ ranked it first. The largest deviation, which was 2.96 standard errors on five rounds, was not reproduced with new seeds. In variant A2 on ten rounds, 500 runs on 100 master keys ranked the right guess first in every run with $\log(N \cdot \text{ELP}_{\text{pair}}) \geq 7$. The

largest difference from the model, in the bin $[6, 7)$, was 0.028 against a standard error of 0.025.

Our implementation of the specification numbers the branches from 0. With S-boxes from SHAKE256, it reproduced the reference implementation of the designers on eight triples of master key, nonce and counter, as detailed in Appendix A. The implementation of Contract computes in integers, and in variant A1 it gave Σ_1, Σ_2 and $|\xi|$ of a direct evaluation of every guess bit for bit in six comparisons on ten and twelve rounds with up to 56 640 625 guesses. In variant A2, it gave the statistics of a direct computation on five to seven rounds. On five to eleven rounds, it gave the checksums of an implementation without Contract, which cover all $4.35 \cdot 10^9$ guesses on ten rounds.

The rule that turns a class into equations of the layers 0 to 2 was checked exhaustively for each of the 64 types of constraint. We ran the executed recovery 480 times on its observations with one, two or four injected wrong pairs. The layers 0 to 2 removed all 1120 injected pairs, and the keystream test always kept the master key. On the server, a unit $N\bar{n}^{m^*}$ of step 2 of Contract took 14.7 to 24.4 ns at $x = 11$ to 14, or 31 to 51 cycles at 2.1 GHz and 59 to 98 cycles at the turbo frequency of 4.0 GHz. The model of Section 6.1 charges $3 \cdot 2^5 = 96$ multiply-adds per unit, counted at one cycle each.

C.4 Joint Detection and Side Trails

For a master key in variant A1, let $c_{\min}$ be the cost of its best usable pair and P_Δ the set of its m usable pairs with costs $c_p \le c_{\min} + \Delta$. Merging their capacities gives the effective cost $-\log \sum_{p \in P_\Delta} 2^{-c_p}$, and the gain is $c_{\min}$ minus this cost. Their $|\mathcal{H}|^m$ joint guesses raise the data constant from κ to $\kappa_m = (\sqrt{2m \ln |\mathcal{H}|} + 1.28)^2$, and the net gain subtracts $\log(\kappa_m/\kappa)$ from the gain. At 26 rounds, the median net gain is -0.309 bits for $\Delta = 2$ and -0.647 bits for $\Delta = 3$, and the fraction of the 4000 sampled keys whose effective cost needs at most 2^{64} blocks with κ_m falls from 0.9832 at $\Delta = 0$ to 0.9345 at $\Delta = 3$.

In variant A2, a side trail with $\ell_b = 8$ before a fixed round r strictly inside the approximation with $c_{\mathrm{fix}}(t'_r) = 0$, and with $\ell = 0$ after it, passes r through the entry $T_r[0, 8]$. By the condition (S_{r+1}), it contributes only with a weight of round $r + 1$ that depends on the key and is not invariant under the shift of indices. With a first segment that keeps the constant patterns, it costs five more bits of data than a round with $c_{\mathrm{fix}}(t'_r) = \pm 1$ for $t'_{r+1} \in [1] \cup [2] \cup [3]$ and seven for $[4]$, and its weight takes 504 functions. Through the fixed round 5, side trails raise the coverage by 2.38 percentage points at 28 rounds and by at most 2.00 with chosen nonces. With restricted rounds of the first kind, they raise the density of weak master keys by 0.005 bits at $w = 0$ to 0.041 bits at $w = 10$.

C.5 Chosen-Nonce Filters of the Variant

At 40 rounds of the variant, for a branch pair with $\Gamma^2 = 1$, filters that act round by round without guessing the key difference save at most 17.80 bits of data within $T_{\mathrm{off}} \le 2^{128}$, against the 36.64 bits needed. A joint filter on all rounds needs at least $2^{73.07}$ blocks within the same offline time. A filter that keeps the blocks with $b^{(r,i)} = \pm \mathrm{shift}_g(b^{(r,j)})$ for a guess g of $t'_r \bmod 8$ makes round r free when the guess is right. With $w \ge 8$ such rounds, $N \le 2^{64}$ and $T_{\mathrm{off}} \le 2^{128}$ hold only for a branch pair with $\Gamma^2 = 1$ that matches every guess, on master keys of density at most 8^{-w} for one branch pair. None of the 4000 sampled keys has a branch pair with $\Gamma^2 = 1$ at 40 rounds. The cost divided by the density is at least $2^{148.8}$. Hence, the conclusion of Section 8 holds for savings averaged over the key difference, or for the cost divided by the density.