

FRAST: TFHE-Friendly Cipher Based on Random S-Boxes

Mingyu Cho^{1*}, Woohyuk Chung², Jincheol Ha², Jooyoung Lee^{2†}, Eun-Gyeol Oh² and Mincheol Son²

¹ Mobilint, Inc., Seoul, Korea

mingyu@mobilint.com

² Korea Advanced Institute of Science and Technology (KAIST), Daejeon, Korea

{hephaistus,smilecjf,hicalf,eun-gyeol.oh,encrypted.def}@kaist.ac.kr

Abstract. A transciphering framework, also known as hybrid homomorphic encryption, is a practical method of combining a homomorphic encryption (HE) scheme with a symmetric cipher in the client-server model to reduce computational and communication overload on the client side. As a server homomorphically evaluates a symmetric cipher in this framework, new design rationales are required for “HE-friendly” ciphers that take into account the specific properties of the HE schemes.

In this paper, we propose a new TFHE-friendly cipher, dubbed FRAST, with a TFHE-friendly round function based on a random S-box to minimize the number of rounds. The round function of FRAST can be efficiently evaluated in TFHE by a new optimization technique, dubbed double blind rotation. Combined with our new WoP-PBS method, the double blind rotation allows computing multiple S-box calls in the round function of FRAST at the cost of a single S-box call. In this way, FRAST enjoys 2.768 (resp. 10.57) times higher throughput compared to Kreyvium (resp. Elisabeth) for TFHE keystream evaluation in the offline phase of the transciphering framework at the cost of slightly larger communication overload.

Keywords: homomorphic encryption · programmable bootstrapping · transciphering framework · stream cipher · HE-friendly cipher

1 Introduction

HOMOMORPHIC ENCRYPTION. In order to achieve both privacy and usability of data in use, various types of homomorphic encryption (HE) schemes have been proposed and widely studied. HE schemes are typically targeted at the client-server model, where a large amount of data from clients is processed by an untrusted server. In this scenario, each client encrypts its data using an HE scheme and sends the HE-ciphertexts to the server so that the server is able to perform meaningful computations on the encrypted data without decrypting it. However, this approach of homomorphically encrypting all the clients’ data in the client side has two technical problems. One is that a client should encrypt all its data with an HE scheme, which is heavier than traditional symmetric encryption schemes. Typically, even public key schemes such as RSA are only used for key-sharing and symmetric key schemes are used for actual data encryption in practice, so it would be desirable to reduce the computational cost required to the client with the help

*This work was done while M. Cho was a Master’s student at KAIST.

†This research was sponsored and supported by the collaboration project between LG Electronics, Inc. and KAIST(Korea Advanced Institute of Science and Technology).



of the server. The other problem, which seems inevitable for HE schemes, is ciphertext expansion that increases the amount of data to transfer to the server.

TRANSCIPHERING FRAMEWORK. In order to address the problems mentioned above, a transciphering framework, also known as hybrid homomorphic encryption, has been proposed [NLV11]. In this framework, the client encrypts its data using a symmetric key cipher, and the server homomorphically recovers the data to get the HE-ciphertext of the original data.

More precisely, the client chooses a symmetric key \mathbf{k} of a symmetric cipher E , computes the HE-ciphertext of \mathbf{k} , namely, $\text{Enc}^{\text{HE}}(\mathbf{k})$ using an HE scheme HE, and sends it to the server only once. After that, when the client wants to send data \mathbf{m} to the server, the client encrypts \mathbf{m} using E with key \mathbf{k} and sends the ciphertext \mathbf{c} to the server. Then the server homomorphically encrypts it to obtain $\text{Enc}^{\text{HE}}(\mathbf{c})$, and then homomorphically evaluates the decryption circuit of E using $\text{Enc}^{\text{HE}}(\mathbf{k})$, obtaining the HE-ciphertext $\text{Enc}^{\text{HE}}(\mathbf{m})$ of the original data \mathbf{m} . Figure 1 describes the transciphering framework using a stream cipher, where the server precomputes $\text{Enc}^{\text{HE}}(\mathbf{z})$ for a keystream \mathbf{z} from the stream cipher in the offline phase, and $\text{Enc}^{\text{HE}}(\mathbf{m})$ is obtained by homomorphically computing $\text{Enc}^{\text{HE}}(\mathbf{c}) - \text{Enc}^{\text{HE}}(\mathbf{z})$ in the online phase. In the transciphering framework, the client takes the following advantages.

- A client does not need to encrypt all its data using an HE scheme (except the symmetric key). All the data can be encrypted using only a symmetric cipher, significantly saving computational resources in terms of time and memory.
- Symmetric encryption does not result in ciphertext expansion, so the communication overload between the client and the server will be significantly lower compared to using any homomorphic encryption scheme alone.

All these merits come at the cost of computational overload on the server-side. That said, this trade-off would be worth considering in practice since servers are typically more powerful than clients.

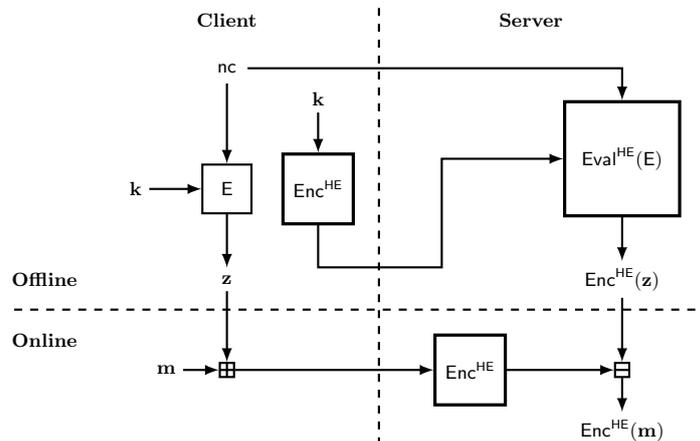


Figure 1: The transciphering framework with a stream cipher. Homomorphic operations are performed in the boxes with thick lines. The vertical dashed line distinguishes the client and server, and the horizontal dashed line distinguishes the offline and online computations. The encryption Enc^{HE} performed by the server in the online phase can be done trivially. nc is a public nonce.

HE-FRIENDLY CIPHERS. For most HE schemes, linear operations such as addition and constant multiplication are way cheaper than nonlinear operations such as multiplication.

With this observation, the efficiency of an HE-friendly cipher has been known to depend on its nonlinear operations. That said, more specific design rationale might differ according to the specific properties of the HE scheme. When combined with the BGV [BGV12, BGV14] and BFV [Bra12, FV12] schemes, HE schemes supporting homomorphic addition and multiplication over finite fields, it would be desirable to use as few multiplications as possible. However, multiplication is needed to construct nonlinear layers, which are essential to achieve security against various attacks including differential, linear, and algebraic attacks. So designing a symmetric cipher with simple nonlinear layers might require a large number of rounds. To address this issue, a novel design strategy has been proposed [MJSC16]: using randomly generated linear layers and simple nonlinear layers. Many HE-friendly ciphers (for the BGV and BFV schemes) follow this design rationale.

However, when it comes to the TFHE scheme [CGGI20, CJP21], one might take a different strategy of designing TFHE-friendly ciphers due to the difference of TFHE from BGV and BFV. We note that when TFHE [CGGI20] was first proposed, it supported homomorphic XOR and AND operations on its ciphertexts encrypting a single bit of message, both of which require the same computational cost of one bootstrapping. Later, the leveled version of TFHE [CJP21] was proposed¹ by expanding the plaintext encoding to \mathbb{Z}_{2^t} for a small integer t such as $t = 4$. In particular, it enjoys fast linear operations and nonlinear operations by means of table lookup. In this paper, TFHE refers to the leveled version of TFHE.

The table lookup operation, homomorphically performed by programmable bootstrapping (PBS), is a distinctive feature of TFHE. Since the table lookup cost does not depend on the algebraic representation of a function to evaluate, it is not required to approximate nonlinear functions into polynomials, as in the case of BGV and BFV. Applying PBS to nonlinear operations such as ReLU and max, [CJP21] enables efficient homomorphic inference of deep neural networks.

In terms of designing TFHE-friendly ciphers, homomorphic nonlinear operation by PBS removes the requirement of the algebraic simplicity of the nonlinear layers. On the other hand, programmable bootstrapping requires that either one padding bit is empty in the most significant bit (MSB) of the value hidden in the ciphertext,² or the table is negacyclic, which means $f(x + p/2) = -f(x)$ for all x where the domain of f is \mathbb{Z}_p for some even p . For example, the Elisabeth cipher [CHMS22], first proposed as a TFHE-friendly cipher, uses 4-bit negacyclic S-boxes in the nonlinear layers.

Regardless of which HE scheme is used, an HE-friendly cipher is typically designed as a stream cipher. The authors of Kreyvium [CCF⁺18], a variant of Trivium [DCP08] with 128-bit security, first claimed that with a stream cipher, a keystream can be precomputed independently of a message in the transciphering framework, leading to simple homomorphic decryption in the online phase.³ Since then, most HE-friendly ciphers have been designed as stream ciphers. As the online phase can be made simple with a stream cipher, high throughput for keystream evaluation in the offline phase becomes of practical relevance, in particular, in an environment where a large amount of data is transferred.

1.1 Our Contribution

In this paper, we propose a new TFHE-friendly stream cipher, dubbed FRAST (Feistel with RAndomly generated S-boxes for Tfhe), which enjoys high throughput for its keystream evaluation on TFHE. The main design rationale of FRAST is twofold.

First, the round function is designed to allow multiple S-box evaluations at the cost of almost a single S-box evaluation. More precisely, it computes $S(x + rk_i)$ for a common

¹It was named Concrete in [CJP21], while recently it is also simply called TFHE.

²Precisely, the MSB should be known to perform PBS correctly.

³In the case of TFHE, one PBS operation might be required to clear a carry bit after the subtraction. See Section 6 for details.

input x and multiple round keys rk_i for an S-box used in the round function. Since the inputs to the S-box are distinct, a naive way of evaluating the round function of FRAS would be to evaluate each $S(x + rk_i)$ separately.

On the other hand, in order to optimize the TFHE evaluation of FRAS, we opt for the structure such that the inputs to the S-box share a common value x , while the round keys rk_i 's are fixed once the master key is chosen. Then $S(x + rk_i)$ can be evaluated for multiple round keys rk_i by sharing the internal state of computing $S(x)$ at the cost of some precomputation on the round keys. We call this type of optimization technique *double blind rotation*.

As mentioned before, the PBS operation basically requires either one empty padding bit in the ciphertext or the negacyclicity of the function to evaluate. To address this issue, advanced PBS techniques called programmable bootstrapping without padding (WoP-PBS)⁴ are proposed [CLOT21, BBB⁺23, YXS⁺21, LMP22, KS22, CBSZ23]. Since FRAS uses a non-negacyclic S-box for its component, we also propose a new WoP-PBS method supporting the double blind rotation.

Our WoP-PBS uses three GenPBS operations in a naive evaluation without requiring additional evaluation keys or larger TFHE parameters for a larger PBS precision. To the best of our knowledge, ComBo [CBSZ23] is the only WoP-PBS method satisfying the above constraints, while it uses four GenPBS operations in its naive evaluation. Combined with the multi-value PBS [CIM19] or PBSmanyLUT [CLOT21], it is possible to reduce one GenPBS for both our WoP-PBS and ComBo. Using parallel computation with $\log p$ threads, the latency of our WoP-PBS becomes almost the same as a single GenPBS, where the plaintext space is \mathbb{Z}_p for some power-of-two p . For ComBo, the latency can be reduced to two GenPBS operations using parallel computation with two threads.

The second feature of our design is that the round function is based on randomly generated S-boxes for some rounds. We note that the TFHE evaluation of an S-box is independent of its structure unless it is constant or negacyclic. Exploiting this property of the TFHE operation, random S-boxes efficiently mitigate various attacks using multiple input-output pairs from a fixed function. On the other hand, some rounds of FRAS are still based on fixed S-boxes to guarantee concrete security against algebraic attacks.

We implement FRAS using the `tfhe-rs` library [Zam22]. FRAS achieves 2.768 (resp. 10.57) times higher throughput compared to Kreyvium (resp. Elisabeth) on the server-side offline phase.

1.2 Related Work

In this section, we briefly review some existing TFHE-friendly ciphers. The FLIP stream cipher [MJSC16] is based on a filter permutator that randomly permutes key bits and computes a nonlinear function, called a filter function, generating a single keystream bit. The filter function is chosen to have a simple algebraic representation for its efficient homomorphic evaluation, and its low security is enhanced by the randomly generated permutation layer. On the other hand, the permutation layer is publicly generated, so one can homomorphically evaluate this layer with almost no cost.

Later, an improved filter permutator, dubbed FiLIP, has been proposed [MCJS19]. As both FLIP and FiLIP have been proposed for the 3rd generation FHE such as TFHE, Hoffmann et al. [HMR20] proposed an efficient evaluation of FiLIP with TFHE. Cong et al. [CDPP22] adopt a transciphering using the FiLIP cipher in private decision tree evaluation, but they use the FINAL [BIP⁺22] scheme for the transciphering, which is an HE scheme based on the NTRU assumption. Méaux et al. [MPP24] also proposed a FINAL-based transciphering with FiLIP that supports setup-independent plaintext space.

⁴It is also called the full domain functional bootstrapping (FDFB).

The (improved) filter permutator originally works on the Boolean space \mathbb{F}_2 . By generalizing this space to a group, Cosseron et al. [CHMS22] proposed a TFHE-friendly cipher Elisabeth. In this construction, \mathbb{Z}_{16} -addition is used for the group operation and the filter function consists of \mathbb{Z}_{16} -addition and evaluation of some negacyclic S-boxes. This exploits the negacyclic-friendly property of the PBS operation in TFHE while its overall structure still follows the strategy of randomly generating its permutation layer. Recently, the Elisabeth cipher has been broken by an algebraic attack [GHBJR23]. After that, Hoffmann et al. [HMS24] proposed several patches for Elisabeth, named Elisabeth-b, Gabriel and Margrethe, whose TFHE evaluation cost is at least twice more than Elisabeth under a single thread.

For the other type of TFHE-friendly ciphers, Balenbois et al. [BOS23] proposed to use Trivium and Kreyvium in the transciphering framework with TFHE. Kreyvium is a variant of a stream cipher Trivium of 80-bit security, that supports a larger key to achieve 128-bit security. Once an initial vector IV is chosen, the key and IV are loaded on the registers. The registers are updated by a nonlinear function, which also generates keystream bits after some initialization rounds. They presented an efficient TFHE evaluation of Trivium and Kreyvium keystreams by the multithreading technique.

Very recently, Deo et al. [DJL⁺24] proposed LWR-based PRFs and applied them to transciphering. The security of their transciphering method is based on the LWR assumption instead of the security of symmetric ciphers.

2 Preliminaries

2.1 Notations

Throughout the paper, bold lowercase letters (resp. bold uppercase letters) denote vectors (resp. matrices). For two vectors (bit strings) \mathbf{a} and \mathbf{b} , their concatenation is denoted by $\mathbf{a}\|\mathbf{b}$. $\lceil r \rceil$ denotes the nearest integer to r , rounding upwards in case of a tie. A real interval $[a, b)$ has an alternative notation: $[a, b[$. For two integers a and b , $\mathbb{Z} \cap [a, b[$ is denoted by $\llbracket a, b \rrbracket$. For an integer q , we identify $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$ with $\llbracket -q/2, q/2 \rrbracket$ in the context of TFHE. The set \mathbb{B} and $[n]$ denote $\{0, 1\}$ and $\{1, 2, \dots, n\}$, respectively, for a positive integer n . For a set S , we will write $a \leftarrow S$ to denote that a is chosen from S uniformly at random. For a probability distribution \mathcal{D} , $a \leftarrow \mathcal{D}$ denotes that a is sampled according to the distribution \mathcal{D} . Unless stated otherwise, all logarithms are to the base 2.

In the context of TFHE, we use p and q for the moduli of messages and ciphertexts, respectively. We only consider the case where p and q are powers of two. For a power-of-two N , we denote the cyclotomic ring $\mathbb{Z}[X]/(X^N + 1)$ by $\mathbb{Z}_N[X]$. For the polynomial ring over \mathbb{Z}_q , we write $\mathcal{R}_{q,N} = \mathbb{Z}_q[X]/(X^N + 1)$. Similarly, we write $\mathbb{B}_N[X] = \mathbb{B}[X]/(X^N + 1)$.

2.2 TFHE

In this section, we briefly review the core concepts of the TFHE scheme. Although TFHE itself is mathematically defined over the real torus $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ [CGGI20], it is common to use the discretized torus $\frac{1}{q}\mathbb{Z}/\mathbb{Z}$ for $q = 2^{32}$ or $q = 2^{64}$ considering its implementation. Hence, we identify the discretized torus $\frac{1}{q}\mathbb{Z}/\mathbb{Z}$ as \mathbb{Z}_q , which is commonly used in the recent descriptions of TFHE [CJP21, CLOT21, BBB⁺23].

LWE, RLWE, AND GLWE CIPHERTEXTS. Under a secret key $\mathbf{S} \in \mathcal{R}_{q,N}^k$, a message $M \in \mathcal{R}_{p,N}$ is encrypted into a GLWE ciphertext $\mathbf{c} \in \mathcal{R}_{q,N}^{k+1}$ with a scaling factor Δ such that $\Delta \leq q/p$ as follows [BGV12, BGV14].

$$\mathbf{c} = \text{GLWE}_{\mathbf{S}}(\Delta \cdot M) = (A_1, \dots, A_k, B = \sum_{i=1}^k A_i \cdot S_i + [M \cdot \Delta]_q + E)$$

where $\mathbf{S} = (S_1, \dots, S_k)$, $A_i \leftarrow \mathcal{R}_{q,N}$ for $i = 1, 2, \dots, k$, and $E \leftarrow \chi_\sigma$ for some Gaussian distribution χ_σ denoting the error distribution. (A_1, \dots, A_k) and B are called the mask and the body of the GLWE ciphertext \mathbf{c} , respectively, and k is called the GLWE dimension. It is common to use a binary secret key in the TFHE scheme, so we only deal with binary secret keys in this paper.

A GLWE ciphertext with $N = 1$ is called an LWE ciphertext. In this case, it is common to use n to denote the LWE dimension instead of k , so that an LWE ciphertext is usually denoted $(a_1, \dots, a_n, b) \in \mathbb{Z}_q^{n+1}$. When $k = 1$, a GLWE ciphertext is called an RLWE ciphertext. In this paper, we refer LWE ciphertexts separately from GLWE ciphertexts of $N > 1$.

The decryption of a GLWE ciphertext is computing its phase, which is defined as $B - \langle (A_1, \dots, A_k), \mathbf{S} \rangle$, followed by rounding the phase by the scaling factor Δ . The decryption works correctly if the error contained in the ciphertext is small enough to be eliminated during the rounding by Δ .

From the definition of the GLWE ciphertext, the addition of two GLWE ciphertexts under the same secret key results in the addition of their internal plaintexts in $\mathcal{R}_{q,N}$. Multiplying the ciphertext by a scalar plaintext is possible by iterating the addition several times. Both the addition and the scalar multiplication increase the error of the resulting ciphertext linearly.

GGSW CIPHERTEXTS. In the case of nonlinear operations such as multiplication, TFHE uses another type of ciphertext called GGSW [GSW13]. Let $B \in \mathbb{N}$ be a power-of-two and $\ell \in \mathbb{N}$. A GGSW ciphertext $\mathbf{C} \in \mathcal{R}_{q,N}^{\ell(k+1) \times (k+1)}$ of a message $M \in \mathbb{Z}_N[X]$ under a secret key $\mathbf{S} \in \mathbb{B}_N[X]^k$ is an $\ell(k+1) \times (k+1)$ matrix over $\mathcal{R}_{q,N}$ defined as follows.

$$\mathbf{C} = \left(\text{GLWE}_{\mathbf{S}} \left(-S_i \cdot \frac{q}{B^j} M \right) \right)_{(i,j) \in [k+1] \times [\ell]}$$

where $\mathbf{S} = (S_1, \dots, S_k)$, $S_{k+1} = -1$, and each GLWE ciphertext is considered as a row having $k+1$ columns of polynomials in $\mathcal{R}_{q,N}$. B and ℓ are called the decomposition base and the decomposition level of the GGSW ciphertext \mathbf{C} , respectively.

EXTERNAL PRODUCT AND CMUX GATE. The external product \boxtimes between GGSW ciphertext \mathbf{C} and GLWE ciphertext \mathbf{c} is defined as

$$\mathbf{C} \boxtimes \mathbf{c} = \text{GadgetDecomp}(\mathbf{c}) \cdot \mathbf{C}$$

where $\text{GadgetDecomp}(\mathbf{c}) \in \mathcal{R}_{q,N}^{\ell(k+1)}$ is the gadget decomposition of \mathbf{c} [GSW13, CLOT21] of which coefficients are lying in $\llbracket -B/2, B/2 \rrbracket$. The external product between GGSW and GLWE ciphertexts defines homomorphic module scalar multiplication on the discretized torus $\frac{1}{q}\mathbb{Z}/\mathbb{Z}$. Roughly speaking, the external product increases the error by the magnitude of the plaintext in the GGSW ciphertext. Thus it is common to use GGSW ciphertext encrypting a single bit of message in the external product.

The controlled mux gate, dubbed CMux, is the key operation used in TFHE. Suppose two GLWE ciphertexts \mathbf{c}_0 and \mathbf{c}_1 are given along with a secret Boolean value b encrypted to a GGSW ciphertext \mathbf{C} , where all three ciphertexts are encrypted with the same key \mathbf{S} . Then one may select \mathbf{c}_b without knowing b by

$$\text{CMux}(\mathbf{C}, \mathbf{c}_0, \mathbf{c}_1) = (\mathbf{c}_1 - \mathbf{c}_0) \boxtimes \mathbf{C} + \mathbf{c}_0.$$

PROGRAMMABLE BOOTSTRAPPING. The programmable bootstrapping (PBS) of TFHE supports an extra functionality that evaluates a function for free during the bootstrapping. Suppose an LWE ciphertext $\mathbf{c} = (a_1, \dots, a_n, b) \in \mathbb{Z}_q^{n+1}$ of $\mu = \Delta m + e$ under a secret key $\mathbf{s} = (s_1, \dots, s_n) \in \mathbb{B}^n$ is given. The PBS operation outputs a refreshed LWE ciphertext $\mathbf{c}' \in \mathbb{Z}_q^{kN}$ of the message $f(m)$ under a secret key $\mathbf{s}' \in \mathbb{B}^{kN}$ by the following steps.

1. Encode the function f on a new GLWE ciphertext under a different secret key $\mathbf{S}' \in \mathbb{B}_N[X]^k$. The half of the function values of f are redundantly encoded in the coefficients of the plaintext of the (trivial) GLWE ciphertext.
2. (Modulus switch) Compute $\tilde{\mathbf{c}} = (\tilde{a}_1, \dots, \tilde{a}_n, \tilde{b}) \in \mathbb{Z}_{2N}^{n+1}$ where

$$\tilde{a}_i = \lfloor a_i \cdot (2N)/q \rfloor \text{ and } \tilde{b} = \lfloor b \cdot (2N)/q \rfloor,$$
 obtaining an LWE ciphertext encrypting $\tilde{\mu} \approx \lfloor \mu \cdot (2N)/q \rfloor$.
3. (Blind rotation) Multiply $X^{-\tilde{b} + \sum_{i=1}^n \tilde{a}_i s_i} = X^{-\tilde{\mu}}$ to the GLWE ciphertext encoding the function using a bootstrapping key $\{\text{GGSW}_{\mathbf{S}'}(s_i)\}_{i=1}^n$; multiply either 1 or $X^{-\tilde{a}_i}$ according to $s_i \in \{0, 1\}$ by the CMux gate.
4. (Sample extraction) Extract the constant term of the GLWE ciphertext, obtaining an LWE ciphertext of $f(m)$ under the secret key $\mathbf{s}' \in \mathbb{B}^{kN}$ which is a reordering of the coefficients of \mathbf{S}' .

In this paper, we call the above PBS operation GenPBS as in [CLOT21].

Since $X^N = -1$ in the ring $\mathcal{R}_{q,N}$, it is only possible to evaluate a negacyclic function $f: \mathbb{Z}_p \rightarrow \mathbb{Z}_q$ such that $f(x + p/2) = -f(x)$ by encoding only half of the function values. To evaluate an arbitrary function, TFHE requires one padding bit of zero in the MSB of m to guarantee $\tilde{m} < N$.

In this paper, we define the precision of GenPBS as the log of the number of the function values encoded in the GLWE ciphertext that the GenPBS operation evaluates. Then, a GenPBS operation of t -bit precision can compute an arbitrary function of t -bit precision, i.e., a function on $\llbracket 0, 2^t \llbracket$, or an arbitrary negacyclic function of $(t + 1)$ -bit precision without the padding bit. A GenPBS operation with a larger precision requires a larger polynomial size N , increasing computational cost for it.

KEYSWITCHING. Since GenPBS outputs an LWE ciphertext under a different secret key, it has to be switched back to a ciphertext under the original secret key. This step is called the *keyswitching*. In practice, the keyswitching operation is performed only once just before the GenPBS operation to match the LWE dimension rather than after every GenPBS operation. To denote the keyswitching operation before the GenPBS operation, we denote a GenPBS after a keyswitching as KSthenGenPBS.

PLAINTEXT ENCODING IN TFHE. To keep the padding bit zero, Bergerat et al. [BBB⁺23] proposed a new encoding method for TFHE splitting the traditional plaintext space into three parts: one (or more) bit of padding at the MSB, the carry subspace after the padding bits, and the message subspace at the LSBs. By tracking the maximum possible value in the ciphertext, it clears the carry space before the padding bit is filled. For example, the default parameters of `tfhe-rs` library for `shortint` type, called `PARAM_MESSAGE_2_CARRY_2`, uses the encoding that consists of two message bits, two carry bits, and one padding bit.

2.3 Algebraic Attacks

In this section, we briefly review algebraic attacks applicable to FRAST.

2.3.1 Trivial Linearization and Extended Linearization

TRIVIAL LINEARIZATION. Given input-output pairs of a symmetric cipher, it is possible to construct a system of polynomial equations with respect to the key variables \mathbf{k} . Trivial linearization is to make the system linear by introducing new variables for all monomials of degree greater than one and to solve it. For the linearized system to be solved, the

number of equations should be greater than or equal to the number of variables including newly introduced ones.

If a system of Boolean equations in n variables is of degree d and almost all the monomials of degree up to d appear in the system, then the complexity of the trivial linearization attack is given by

$$\left(\sum_{i=1}^d \binom{n}{i} \right)^\omega$$

ignoring the constant factor, where $2 \leq \omega \leq 3$ denotes the linear algebra constant.

EXTENDED LINEARIZATION. Courtois et al. [CKPS00] proposed the eXtended Linearization algorithm that can be used when the number of equations is less than the number of monomials. Given a system of m equations of degree d in n variables over \mathbb{F}_2 , the XL algorithm extends the system by multiplying all the monomials of degree at most $D - d$ for some $D (> d)$ to obtain a larger number of (linearly independent) equations of degrees at most D . As the number of equations grows faster than the number of monomials, it is possible to solve the system for a sufficiently large D . The problem is that it is hard to determine the smallest D , called the *solving degree*.⁵

When designing a symmetric cipher, we can assume that all the resulting equations are linearly independent, which is in favor of an adversary. Then it is possible to estimate the solving degree D as the smallest one satisfying

$$\left(\sum_{i=0}^{D-d} \binom{n}{i} \right) m \geq \sum_{i=1}^D \binom{n}{i}$$

assuming that all the monomials appear in the extended system of equations. Then the complexity of the XL algorithm is given by

$$\left(\sum_{i=1}^D \binom{n}{i} \right)^\omega \quad (1)$$

ignoring the constant factor. There are some optimized variants such as the Wiedemann XL algorithm [YCBC07], while using (1) with $\omega = 2$ still gives a lower bound on its complexity.

HYBRID STRATEGY. A hybrid strategy that guesses the values of some variables can be applied to the linearization attacks. The complexity of the hybrid trivial linearization after guessing k variables is given by

$$\min_k 2^k \left(\sum_{i=1}^d \binom{n-k}{i} \right)^\omega \quad (2)$$

ignoring the constant factor.

For the XL attack, the solving degree might differ according to the number of guessed variables. Let $D_{n,k}$ be the solving degree of the system after guessing k variables. Then the complexity of the hybrid XL attack over \mathbb{F}_2 is given by

$$\min_k 2^k \left(\sum_{i=1}^{D_{n,k}} \binom{n-k}{i} \right)^\omega \quad (3)$$

ignoring the constant factor.

⁵The recent results show that the solving degree is the same as the degree of regularity [YC04, AFI+04].

2.3.2 Gröbner Basis Attack

The Gröbner basis attack is to solve a system of equations by computing its Gröbner basis. The complexity of Gröbner basis computation can be estimated using the *degree of regularity* of the system of equations [BFS04], upper bounding the degree of polynomials that occur during the computation of a Gröbner basis using algorithms such as F_4 [Fau99] and F_5 [Fau02].

Consider a system $\{f_i\}_{i=1}^m$ of m equations in n variables. Let d_i denote the degree of f_i for $i = 1, 2, \dots, m$. When working on \mathbb{F}_2 , Bardet et al. [BFSS13] proposed the following Hilbert series to estimate the degree of regularity taking homogenization into account.⁶

$$T_{m,n}(z) = \frac{(1+z)^n}{(1-z) \prod_{i=1}^m (1+z^{d_i})}. \quad (4)$$

The degree of regularity d_{reg} is determined by the smallest degree of the term with a non-positive coefficient in the series $T_{m,n}$. Given the degree of regularity d_{reg} over \mathbb{F}_2 , the complexity of computing a Gröbner basis is known to be

$$O\left(\binom{n}{d_{\text{reg}}}\right)^\omega. \quad (5)$$

The degree of regularity estimated by the Hilbert series is commonly used to theoretically estimate the complexity of the Gröbner basis attack. However, it requires an assumption that the system of equations is semi-regular. It is known that almost all polynomial sequences are semi-regular [Frö85], but for a system obtained from a symmetric cipher it might not be the case since it has a certain structure.

3 The FRAST Cipher

3.1 Specification

A stream cipher FRAST with 128-bit security takes as input a 256-bit key $\mathbf{k} \in \mathbb{Z}_{16}^{64}$ and a 128-bit nonce $\text{nc} \in \{0, 1\}^{128}$, and returns a 128-bit keystream block $\mathbf{k}_{\text{nc}} \in \mathbb{Z}_{16}^{32}$. The FRAST cipher has two types of round functions: the randomized one and the fixed one. Both types of round functions have the same structure except for the underlying S-boxes.

In a nutshell, FRAST consists of 40 rounds, namely,

$$\text{FRAST}[\mathbf{k}, \text{nc}] = \text{RF}[\mathbf{k}, \text{nc}, 40] \circ \text{RF}[\mathbf{k}, \text{nc}, 39] \circ \dots \circ \text{RF}[\mathbf{k}, \text{nc}, 1]$$

where $\text{RF}[\mathbf{k}, \text{nc}, i]$ is the i -th round function using secret key \mathbf{k} and nonce nc . For $(x_1, \dots, x_{32}) \in \mathbb{Z}_{16}^{32}$, $\text{RF}[\mathbf{k}, \text{nc}, i](x_1, \dots, x_{32}) = (y_1, \dots, y_{32}) \in \mathbb{Z}_{16}^{32}$ is defined as

$$\begin{aligned} y_j &= x_j + S_{\text{erf}}^{(\text{nc}; i)}(x_1 + rk_j^{(i)}) \text{ for } j = 2, 3, \dots, 32, \\ y_1 &= x_1 + S_{\text{crf}}^{(\text{nc}; i)}(y_2 + y_3 + \dots + y_{32} + rk_1^{(i)}) \end{aligned}$$

where $\mathbf{rk}^{(i)} = (rk_1^{(i)}, rk_2^{(i)}, \dots, rk_{32}^{(i)}) \in \mathbb{Z}_{16}^{32}$ is the i -th round key (derived from the master key \mathbf{k}), and $S_{\text{erf}}^{(\text{nc}; i)}$ and $S_{\text{crf}}^{(\text{nc}; i)}$ are the S-boxes used in the i -th round function⁷ (see Figure 2).

FRAST repeats 4 random round functions followed by 1 fixed round function. Hence, the i -th round is a random round if i is not a multiple of 5, and is a fixed round if i is a multiple of 5.

⁶It already takes into account the field equations of the form $x^2 - x = 0$ over \mathbb{F}_2 .

⁷The subscript erf (resp. crf) stands for expanding round function (resp. contracting round function), a term used in the generalized Feistel networks [AGP⁺19].

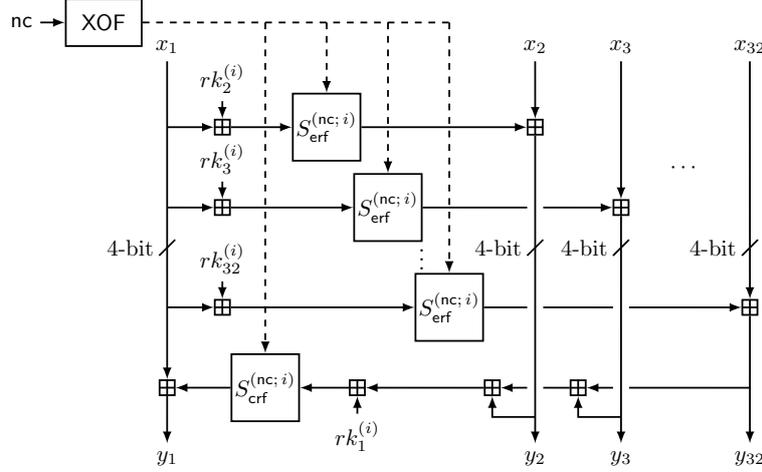


Figure 2: The i -th round function of FRAST. $rk_1^{(i)}, \dots, rk_{32}^{(i)}$ are the i -th round keys.

S-BOX. The random rounds of FRAST generate their S-boxes $S_{\text{erf}}^{(\text{nc}; i)}$ and $S_{\text{crf}}^{(\text{nc}; i)}$ as negacyclic functions; the function values $(S_{\text{erf}}^{(\text{nc}; i)}(0), \dots, S_{\text{erf}}^{(\text{nc}; i)}(7))$ and $(S_{\text{crf}}^{(\text{nc}; i)}(0), \dots, S_{\text{crf}}^{(\text{nc}; i)}(7))$ are sampled by the output from the underlying extendable output function XOF with input nc , and the other function values are determined by the negacyclic property of the S-boxes.

The fixed rounds of FRAST use the same fixed S-box defined in Table 1 for their S-boxes $S_{\text{erf}}^{(\text{nc}; i)}$ and $S_{\text{crf}}^{(\text{nc}; i)}$. The S-box is one of the golden S-boxes of size 4-bit proposed in [Saa12].

Table 1: The fixed S-box used in the fixed rounds of FRAST.

x	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
$S(x)$	0	3	5	8	6	9	12	7	13	10	14	4	1	15	11	2

KEY SCHEDULE. The round keys $\mathbf{rk}^{(i)}$ for $i = 1, 2, \dots, 40$ are obtained by multiplying 64×64 invertible matrices over \mathbb{Z}_{16} to the master key $\mathbf{k} \in \mathbb{Z}_{16}^{64}$. More precisely, an invertible matrix \mathbf{M} is chosen from the output of SHAKE256 with a fixed public input.⁸ Then, round keys \mathbf{rk}_{2i-1} and \mathbf{rk}_{2i} are defined by $\mathbf{M}^i \cdot \mathbf{k} = \mathbf{rk}_{2i-1} \parallel \mathbf{rk}_{2i}$ for $i = 1, 2, \dots, 20$. The exact specification of \mathbf{M} is given in Supplementary Material G.

ENCRYPTION MODE. When a keystream of m blocks in $(\mathbb{Z}_{16}^{32})^m$ is needed for some $m > 0$, the “inner-counter mode” can be used: for $\text{ctr} = 0, 1, \dots, m - 1$, one computes

$$\mathbf{z}[\text{ctr}] = \text{FRAST}[\mathbf{k}, \text{nc} \parallel \text{ctr}](\mathbf{ic}),$$

where \mathbf{ic} denotes a constant $(0, 1, \dots, 15, 0, 1, \dots, 15) \in \mathbb{Z}_{16}^{32}$. Figure 3 shows the overall structure of FRAST in the counter mode.

3.2 Design Rationale

Each round of FRAST uses a common input branch to the S-boxes, while each S-box input is masked with a distinct round key. The motivation for using the common input branch lies in the multi-value PBS [CIM19] that evaluates multiple functions on the same input

⁸We used the string `Feistel with Randomly generated S-boxes for Tfhe` in our implementation, obtaining an invertible matrix over both \mathbb{Z}_{16} and \mathbb{F}_{2^4} without rejection.

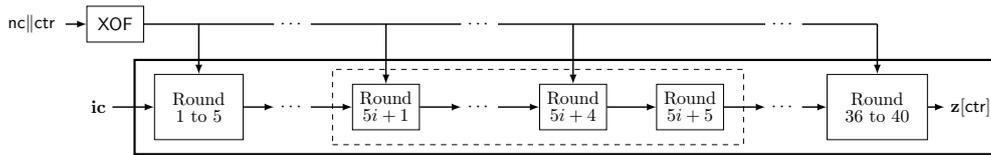


Figure 3: The overall structure of FRAST in the counter mode, where \mathbf{ic} is the public input constant and $\mathbf{z}[\text{ctr}]$ is the keystream. Homomorphic operations are performed in the box with thick lines. The i -th round is a fixed round if i is a multiple of 5, and a random round otherwise.

at the cost of almost one PBS call. To evaluate multiple functions efficiently, it computes a common state by the blind rotation, the most expensive part of the PBS operation, and induces the desired function values by polynomial multiplications. Similarly to this method, by sharing the result of the blind rotation on the common input value and using the precomputed GGSW ciphertexts of the round key bits, we can evaluate multiple S-box calls on a common input with multiple round keys added at the cost of almost one S-box evaluation. This optimization technique is dubbed *double blind rotation* since one blind rotation by the common input is followed by another by the round key (see Section 4.2 for the details).

The round function of FRAST, which can be seen as a generalized Feistel network, has an issue on the error growth in its TFHE evaluation. Unlike the substitution-permutation network, the branches are not refreshed during the evaluation of S-boxes, so linear layers, linearly increasing the internal noise, are not “free” anymore. With a simple permutation layer, the number of rounds might be much higher than the number of branches due to the weak diffusion of the permutation layer.

It is a common strategy to design HE-friendly ciphers with random linear layers (in the substitution-permutation network) in order to reduce the required number of rounds [DEG⁺18, HKC⁺20, HL20, DGH⁺23]. However, TFHE evaluation of such linear layers will lead to significant error growth. So we opt for random S-box generation to randomize the FRAST encryption function. Since TFHE evaluates S-boxes by PBS operations, the random structure of the S-box will not affect the evaluation cost of FRAST.

The size of the S-boxes in FRAST is 4 bits, which is common in symmetric ciphers. As the performance of PBS operation with 4-bit precision is sufficiently efficient,⁹ we choose 4-bit S-boxes for FRAST. For the efficiency of the homomorphic S-box evaluation, the random S-boxes are defined by negacyclic functions whose half of the function values are sampled uniformly at random from \mathbb{Z}_{16} , determining the others by the negacyclic property. The negacyclic property is the only condition for the random S-boxes, while one might consider algebraically structured S-boxes with simple randomness such as affine equivalence. Unlike the cases of previous HE-friendly ciphers using large S-boxes with random linear layers, the small S-box size of FRAST requires high randomness for its random S-box generation to prevent linear attacks with a small number of rounds (see Section 5.2).

On the other hand, the fixed S-box of FRAST is non-negacyclic. It is one of the golden S-boxes proposed in [Saa12], which is known to have no polynomial representation over \mathbb{Z}_{16} (see Lemma 1). Although FRAST uses addition over \mathbb{Z}_{16} , we analyze its security against algebraic attack with the XOR-variant of FRAST as done in the Elisabeth cipher, focusing on classical algebraic properties over \mathbb{F}_2 . Over \mathbb{F}_2 , all the output bits of the S-box have algebraic degree 3 and all the input bits affect the output bits nonlinearly. Its inverse also has the same algebraic properties. The number of quadratic equations induced by

⁹The default parameter of `tfhe-rs` library for `shortint` type supports GenPBS with 4-bit precision.

the S-box is 21, which is the minimum for 4-bit S-boxes. The algebraic representations of the S-box are given in Supplementary Material A. Zhang et al. [ZBRL15] argue that the golden S-boxes proposed in [Saa12] might be vulnerable to differential and linear attacks due to certain properties of the linear layers combined with the S-boxes. When it comes to FRAST, randomly generated S-boxes mitigate such vulnerabilities.

As mentioned above, the *random* rounds and the *fixed* rounds mitigate linear attacks and algebraic attacks, respectively, so the number of random and fixed rounds can be determined by separate criteria. On the other hand, since only the fixed round refreshes the internal state of the linear sum at no additional cost while evaluating S_{crf} (see Supplementary Material E), FRAST places one fixed round in every four rounds.

FRAST uses 256-bit keys, providing 128-bit security. The key length typically affects the communication overload of the transciphering since it should be sent to the server as a TFHE-ciphertext. When it comes to FRAST, most of the round keys are sent to the server to enable the double blind rotation, so that the length of the master key itself does not affect the communication overload. Hence, to achieve stronger security against algebraic attacks, FRAST uses master keys longer than the target level of security.

4 PBS Techniques for FRAST

In the transciphering framework with TFHE, each S-box of the cipher should be evaluated by the GenPBS operation. However, it requires either one bit of padding in the input ciphertext or the S-box to be negacyclic. Although the random S-boxes of FRAST are negacyclic, the fixed S-box of FRAST is non-negacyclic.

To address the issue of padding bit, PBS techniques called without-padding PBS (WoP-PBS) have been proposed [CLOT21, YXS+21, LMP22, BBB+23, KS22]. Most of the previous WoP-PBS methods require either additional evaluation keys, such as LWE-to-GLWE keyswitching keys [CLOT21, KS22] and circuit bootstrapping keys [BBB+23], or parameters supporting GenPBS of t -bit precision to compute functions of t -bit input [YXS+21, LMP22] (while a negacyclic function of t -bit input can be computed by GenPBS of $(t - 1)$ -bit precision). Using the additional evaluation keys of large size requires an additional communication overload, so it might weaken the purpose of the transciphering framework. WoP-PBS of t -bit precision using GenPBS of t -bit precision is not compatible with the double blind rotation, a technique for the FRAST evaluation described in Section 4.2, since the double blind rotation requires one more padding bit to evaluate a non-negacyclic function. ComBo, proposed by Clet et al. [CBSZ23], is the first WoP-PBS method that can compute an arbitrary function of t -bit precision using GenPBS of $(t - 1)$ -bit precision without additional evaluation keys. However, ComBo is also not compatible with the double blind rotation since it uses GenPBS operations in depth 2.

4.1 Our New WoP-PBS

In this section, we introduce our new WoP-PBS that can compute an arbitrary function of t -bit precision using at most 3 GenPBS operations of $(t - 1)$ -bit precision with no additional evaluation key. If it is possible to use the multi-value PBS [CIM19] or PBSmanyLUT [CLOT21] together, our WoP-PBS can compute arbitrary functions in 2 GenPBS operations. In parallel computation, a variant of our WoP-PBS can fully parallelize the required GenPBS operations, obtaining the latency almost the same as that of a single GenPBS operation. Most importantly, our WoP-PBS method supports the double blind rotation technique, described in Section 4.2, to evaluate FRAST efficiently.

For an input message $m \in \llbracket 0, p \rrbracket$ with $m = \beta \llbracket m' = \beta \cdot \frac{p}{2} + m' \rrbracket$ for $\beta \in \{0, 1\}$ and $m' \in \llbracket 0, p/2 \rrbracket$, let $\text{ClearMSB}(m) = 0 \llbracket m' \rrbracket$, which can be computed in a single GenPBS operation by extracting the MSB of the input and subtracting it from the input on the MSB position.

The resulting ciphertext can be regarded as a ciphertext of $\text{ClearMSB}(m) \in \llbracket 0, p/2 \rrbracket$ with one bit of padding.

The function f to compute is decomposed into $f_{\text{msb-odd}}$ and $f_{\text{msb-even}}$, which are defined as follows.

$$\begin{aligned} f_{\text{msb-odd}}(m) &= \frac{1}{2} \left(f(m) - f\left(m + \frac{p}{2}\right) \right) = \frac{1}{2} \left(f(\beta \| m') - f(\bar{\beta} \| m') \right) \\ f_{\text{msb-even}}(m) &= \frac{1}{2} \left(f(m) + f\left(m + \frac{p}{2}\right) \right) = \frac{1}{2} \left(f(\beta \| m') + f(\bar{\beta} \| m') \right) \end{aligned}$$

where $\bar{\beta} = 1 - \beta \in \{0, 1\}$, the flipped bit of β . From $f = f_{\text{msb-odd}} + f_{\text{msb-even}}$, one can compute f by computing $f_{\text{msb-odd}}$ and $f_{\text{msb-even}}$.

We note that $f_{\text{msb-odd}}$ becomes a negacyclic function on $\llbracket 0, p \rrbracket$ from its definition, i.e., $f_{\text{msb-odd}}(m + \frac{p}{2}) = -f_{\text{msb-odd}}(m)$, so it can be evaluated by a single GenPBS operation. In the case of $f_{\text{msb-even}}$, it can be regarded as a function on $\llbracket 0, p/2 \rrbracket$ since $f_{\text{msb-even}}(m + \frac{p}{2}) = f_{\text{msb-even}}(m)$. Observing that the MSB of the input does not affect the output of $f_{\text{msb-even}}$, one can compute $f_{\text{msb-even}}$ in a single GenPBS operation with an input of the ciphertext of $\text{ClearMSB}(m)$ with a padding bit. All the GenPBS operations are of $(\log p - 1)$ -bit precision. Since the output ciphertext of our WoP-PBS is sum of two outputs of GenPBS operations, the error variance of it is two times larger than that of the GenPBS operation in the same parameters.

Our WoP-PBS requires 3 GenPBS operations in its naive evaluation: one for each of $f_{\text{msb-odd}}$, ClearMSB , and $f_{\text{msb-even}}$. When the multi-value PBS [CIM19] or PBSmany-LUT [CLOT21] can be used together, one can compute $f_{\text{msb-odd}}$ and ClearMSB in a single GenPBS operation since they compute negacyclic functions on the same input.

Using parallel computation, a variant of our WoP-PBS can achieve almost the same latency with a single GenPBS operation as follows. After decomposing f into $f_{\text{msb-odd}}$ and $f_{\text{msb-even}}$, one can further decompose $f_{\text{msb-even}}$ using the same method recursively by regarding $f_{\text{msb-even}}$ as a function on $\llbracket 0, p/2 \rrbracket$. For example, when $p = 16$, an arbitrary function f on $\llbracket 0, p \rrbracket$ is decomposed into 4 negacyclic functions f_0, f_1, f_2, f_3 and one constant f_4 such that

$$f(m) = f_0(m) + f_1(m \bmod 8) + f_2(m \bmod 4) + f_3(m \bmod 2) + f_4$$

where f_i is negacyclic on $\llbracket 0, p/2^i \rrbracket$ for $i = 0, \dots, 3$. All the f_i can be computed by a single GenPBS of $(\log p - i - 1)$ -bit precision without the padding bit. The input ciphertext of $m \bmod 2^{4-i}$ without padding can be obtained by using the plaintext modulus switching ignoring some of the first MSB bits [CLOT21], which is the same as multiplying 2^i to the original input. This decomposition can be generalized to the function on $\llbracket 0, p \rrbracket$ for arbitrary power-of-two p , resulting in $\log p$ negacyclic functions and one constant. All the decomposed functions can be computed in parallel, achieving the latency of a single GenPBS operation using $\log p$ threads. The error variance of the output ciphertext is $(\log p)^2$ times larger than the output of GenPBS in the worst-case.¹⁰ For the exact specification of the decomposed functions, see Supplementary Material D.

4.2 Double Blind Rotation

The round function of FRAST requires multiple S-box calls on the inputs of the form $x_1 + rk_j$ for $j = 2, \dots, \ell$ where x_1 is the value of the first branch, rk_j 's are the round keys, and ℓ is the number of branches. In the GenPBS operation computing a function f on $x_1 + rk_j$, the blind rotation step rotates a GLWE ciphertext encoding f by a factor of

¹⁰Estimating the error variance as $\log p$ times larger one in the average-case requires the heuristic assumption that outputs of the GenPBS operations on the inputs that differ only by constant factors have independent errors.

$x_1 + rk_j$. In a naive way, evaluating $f(x_1 + rk_j)$ for all $j = 2, \dots, \ell$ requires rotating a GLWE ciphertext encoding f by $x_1 + rk_j$ independently for all $j = 2, \dots, \ell$.

The idea of the double blind rotation is that the result of the rotation by x_1 can be shared. Suppose an LWE ciphertext of x_1 without padding and GGSW ciphertexts of the round key bits $rk_{j,b}$ are given where $rk_j = rk_{j,4} \parallel \dots \parallel rk_{j,1}$ for $j = 2, 3, \dots, \ell$ and $b = 1, \dots, 4$. For a negacyclic function f , let $\text{GLWE}(P_f)$ be a GLWE ciphertext of $P_f \in \mathcal{R}_{q,N}$ that encodes f on its coefficients. The blind rotation on $\text{GLWE}(P_f)$ by $\text{LWE}(x_1)$ outputs $\text{GLWE}(P_f \cdot X^{-\hat{x}_1})$ where \hat{x}_1 is a scaled value of x_1 such that the constant term of $P_f \cdot X^{-\hat{x}_1}$ becomes $f(x_1)$. Then one can compute $\text{GLWE}(P_f \cdot X^{-(\hat{x}_1 + \hat{r}k_j)})$ by multiplying $X^{-\hat{r}k_j}$ homomorphically where $\hat{r}k_j$ is the scaled value of rk_j . Given the GGSW ciphertexts of the round key bits, it can be computed by additional 4 CMux gates. Therefore, one can compute $\text{LWE}(f(x_1 + rk_j))$ for all $j = 2, \dots, \ell$ by a single GenPBS followed by $4(\ell - 1)$ CMux gates.

Two issues remain for applying the double blind rotation on FRAST. First, the fixed S-box in the FRAST round function is not negacyclic, requiring WoP-PBS for its evaluation instead of GenPBS. In this case, a WoP-PBS method that does not perform GenPBS operations in depth 2 is required. Evaluating a non-negacyclic function with the help of padding bits is also possible, but the double blind rotation requires two bits of padding to guarantee $\hat{x}_1 + \hat{r}k_i < N$. Our WoP-PBS method can resolve this issue for evaluating FRAST with TFHE parameters supporting a GenPBS operation in 4-bit precision. When S is decomposed into $S_{\text{msb-odd}}$ and $S_{\text{msb-even}}$, the double blind rotation can be applied to $S_{\text{msb-odd}}$ as it is negacyclic. In the case of $S_{\text{msb-even}}$, one can make two bits of padding in the ciphertext of $\text{ClearMSB}(x_1)$ by one more GenPBS operation to change its scaling factor, allowing the double blind rotation using a GenPBS operation of 4-bit precision.¹¹ See Supplementary Material E for details of FRAST evaluation by the double blind rotation.

The other is computing the GGSW ciphertexts of the round key bits. Since the round key bits are fixed, we directly transfer the round key bits used for the double blind rotation packed in the GLWE ciphertexts once, and convert it into GGSW ciphertexts on the server-side by the GLWEtoGGSW conversion proposed in [CCR19].¹² The communication overload for the round keys and the evaluation keys for the conversion is only a few MBs.

5 Security Analysis

In this work, we will consider the standard “secret-key model”, where an adversary arbitrarily chooses a nonce, and obtains the corresponding keystream without any information on the secret key. The related-key and the known-key models are beyond the scope of this paper. We also limit the number of encryptions under the same key up to 2^{64} blocks since otherwise one would not be able to avoid a nonce collision (when nonces are chosen uniformly at random).

The extendable output function whose output determines the random S-boxes is modeled as a random oracle, so an adversary is not able to freely choose the S-boxes. The input to the FRAST is also fixed as the known constant \mathbf{ic} . Therefore, in this model, we believe that FRAST is secure against any type of chosen-plaintext attacks such as (higher-order) differential, truncated differential, invariant subspace trail, and cube attacks. On the other hand, we assume that the specifications of the random S-boxes are given to the attacker.

¹¹Using the variant of our WoP-PBS that fully decomposes a function into negacyclic functions enables the double blind rotation using a GenPBS operation of 3-bit precision at the cost of more CMux gate operations.

¹²It only deals with the case of converting RLWE to RGSW, but converting GLWE to GGSW for $k > 1$ is also possible using the same idea.

Overall, in this section, our focus will be mainly put on algebraic and linear attacks, which are possible in the known-plaintext models. We analyze the security of FRAST against algebraic (resp. linear) attacks based on the fixed (resp. random) round functions of FRAST.

SUMMARY OF THE ANALYSIS. For random rounds, we found a linearization attack using nonzero input masks on the outputs from XOF works for 20 rounds with 2^{126} data. The trail used for the attack is a simple one that does not depend on the round keys. For more complicated trails, a certain relation between round keys might be exploited in each round. Hence, we opted for 32 random rounds to mitigate unknown attacks that might be made possible by guessing round keys. So 12 rounds can be seen as a security margin.

For fixed rounds, a hybrid trivial linearization attack works for 5 fixed rounds with time complexity 2^{125} . So we opted for 8 fixed rounds (including 2 marginal rounds), placing them between the random rounds. Each fixed round has been placed after 4 random rounds for noise management as described in Section 3.2. Overall, FRAST consists of 32 random rounds and 8 fixed rounds.

5.1 Algebraic Attacks

An algebraic attack is to build a system of polynomial equations and solve it to recover the secret key. For simplicity, we ignore the random rounds of FRAST. Instead, we assume that a distinct input is given to FRAST since the first 4 rounds are randomized ones. This can be considered as a case such that all the random S-boxes are linear, which is in favor of an attacker, so that the algebraic degree does not increase in the random rounds. The negacyclicity of random S-boxes does not affect the internal states since they are updated by adding S-box outputs in the Feistel structure of FRAST, not evaluating the S-boxes. So random rounds can be seen as marginal against algebraic attacks although it is hard to estimate their algebraic resistance due to the randomness.

ALGEBRAIC REPRESENTATION OVER \mathbb{Z}_{16} . Recently, Grassi et. al. [GMAH⁺23] showed that the brute-force attack on a polynomial system of n secret elements over \mathbb{Z}_q , where $q = p^y$ for a prime p and $y > 1$, requires only $O(y \cdot p^n)$ computation instead of $O(p^{y \cdot n})$. This implies that if FRAST has a polynomial representation over \mathbb{Z}_{16} then it becomes vulnerable to the brute-force search attack. On the other hand, for an S-box to have a polynomial representation over \mathbb{Z}_{16} , the following condition should be satisfied.

Lemma 1. *Let $f : \mathbb{Z}_{16} \rightarrow \mathbb{Z}_{16}$ be a function over \mathbb{Z}_{16} . If f has a polynomial representation, then it should satisfy $f(i+8) - f(i) \in \{0, 8\}$ for all $i = 0, 1, \dots, 7$.*

Proof. Suppose that f is represented by $f(x) = a_0 + a_1x + \dots + a_{15}x^{15}$ where $a_0, a_1, \dots, a_{15} \in \mathbb{Z}_{16}$. Then, for all $i = 0, 1, \dots, 7$, one obtains

$$f(i+8) - f(i) = 8 \left(\sum_{j=1}^{15} a_j \left(\sum_{m=0}^{j-1} (i+8)^{j-m-1} i^m \right) \right).$$

Since $8z \in \{0, 8\}$ for every $z \in \mathbb{Z}_{16}$, $f(i+8) - f(i) \in \{0, 8\}$. □

As the fixed S-box of FRAST does not satisfy the necessary condition above, we can conclude that the FRAST round function cannot be represented as a polynomial over \mathbb{Z}_{16} .

ALGEBRAIC ANALYSIS ON THE XOR VARIANT OF FRAST. We consider an XOR-variant of FRAST with addition and constant multiplication over \mathbb{Z}_{16} replaced by XOR and multiplication over \mathbb{F}_{2^4} . Such an approach has also been taken for the algebraic analysis of the Elisabeth cipher by introducing an XOR-variant of Elisabeth, dubbed Beth.

The XOR-variant of FRAST can be represented by a system of Boolean equations. In this section, we analyze FRAST using various systems of Boolean equations for its

XOR-variant. For the complexity of linearization attacks, we will use $\omega = 2$ in (2) and (3). In particular, for the XL attack, the independence assumption will be used to estimate the solving degree as mentioned in Section 2.3.1.

5.1.1 Trivial Linearization

One can build a system of equations using only the key variables as unknowns and apply trivial linearization attack. The attack cost depends on the number of monomials appearing in the system, determined by the degree of the system.

Consider a single round function of FRAST with input $(x_1, x_2, \dots, x_\ell)$, output $(y_1, y_2, \dots, y_\ell)$, and round key $(rk_1, rk_2, \dots, rk_\ell)$ where $x_j, y_j, rk_j \in \mathbb{F}_2^4$ for $j = 1, 2, \dots, \ell$.¹³ Then we have

$$\begin{aligned} y_1 &= x_1 + S(y_2 + y_3 + \dots + y_\ell + rk_1), \\ y_j &= x_j + S(x_1 + rk_j) \end{aligned} \quad \text{for } j = 2, 3, \dots, \ell. \quad (6)$$

Since all the outputs of S are of degree 3, the degree of $S(x + rk)$ is at least $\deg x + 2$, assuming that rk_j is a dense linear combination of the master key. Then, we have $\deg y_j \geq \deg x_1 + 2$ for $j = 2, \dots, \ell$ and $\deg y_1 \geq \deg x_1 + 4$. Let $(x_1^{(i-1)}, \dots, x_\ell^{(i-1)})$ (resp. $(x_1^{(i)}, \dots, x_\ell^{(i)})$) be the input (resp. output) of the i -th round function. From $\deg x_j^{(1)} = 3$ for $j = 2, \dots, \ell$, we obtain

$$\begin{aligned} \deg x_1^{(i)} &\geq 4i + 1, \\ \deg x_j^{(i)} &\geq 4i - 1 \end{aligned} \quad \text{for } j = 2, \dots, \ell$$

for $i = 1, 2, \dots, r$ where r is the number of rounds.

Considering the meet-in-the-middle attack, we also need to consider the backward direction. Using the same argument, one can obtain a system of degree at least $4\lfloor r/2 \rfloor + 1$, which is 17 when $r = 8$.

Since the round keys are dense linear combinations of the master key, one can expect almost all the monomials of degree up to the lower bound to appear in the system. Experimental results on the number of the appearing monomials with toy parameters are given in Supplementary Material C. Assuming the system of FRAST contains almost all the monomials of degree 17, the time complexity of the (hybrid) trivial linearization attack is $2^{173.76}$.

Remark 1. The attack on Elisabeth [GHBJR23] is one of the trivial linearization attacks combined with an optimization technique on the Gaussian elimination. Hence, FRAST is also secure against the attack.

5.1.2 XL Attack

Other than the equations only in the key variables, one can build a system of equations of low degrees by introducing new variables other than the key variables or guessing some bits of the internal state. Then it is possible to apply algebraic attacks such as the XL attack and the Gröbner basis attack. We consider the following three kinds of systems for FRAST.

1. A system of equations by introducing new intermediate variables for each round.
2. A system of equations by introducing new variables for the first branch of each round.
3. A system of equations by guessing the values of the first branch for each round.

¹³ $\ell = 32$ in the actual specification of FRAST.

NEW VARIABLES FOR THE INTERMEDIATE STATE. One can build a system of equations by introducing new intermediate variables for each round. In this case, using implicit relations induced by S leads to a larger number of equations of lower degrees than using explicit relations. Let $(x_j^{(i-1)})_{j=1}^\ell$ and $(x_j^{(i)})_{j=1}^\ell$ denote the input and the output of the i -th round function, respectively. From the implicit relation $S'(x, y) = 0$ induced by S , one can obtain the following equations.

$$\begin{aligned} S'(x_2^{(i)} + x_3^{(i)} + \cdots + x_\ell^{(i)} + rk_1^{(i)}, x_1^{(i)} - x_1^{(i-1)}) &= 0 \\ S'(x_1^{(i-1)} + rk_j^{(i)}, x_j^{(i)} - x_j^{(i-1)}) &= 0 \quad \text{for } j = 2, 3, \dots, \ell \end{aligned} \quad (7)$$

where $(rk_j^{(i)})_{j=1}^\ell$ denotes the i -th round key which is linear to the master key.

Since S-box S has 21 implicit quadratic equations over \mathbb{F}_2 , each round induces 21 quadratic equations. For r rounds of FRAST, one obtains $21\ell r$ quadratic equations in $256 + 128(r-1)$ variables. When m keystream blocks are used, one obtains $21\ell r m$ quadratic equations in $256 + 128(r-1)m$ variables.

NEW VARIABLES FOR THE FIRST BRANCH. Since the first branch is common to all the S-box evaluations, introducing new variables for $(x_1^{(i)})_{i=1}^{r-1}$ might significantly reduce the degree of the keystream $x_j^{(r)}$ for all $j = 2, 3, \dots, \ell$. Regarding $(x_1^{(i)})_{i=1}^{r-1}$ as new variables, one obtains two types of equations: one from the keystream $(x_j^{(r)})_{j=2}^\ell$ and the other from $(x_1^{(i)})_{i=1}^r$. The first type of equations are

$$x_j^{(r)} = x_j^{(0)} + \sum_{i=1}^r S(x_1^{(i-1)} + rk_j^{(i)}) \quad (8)$$

for $j = 2, 3, \dots, \ell$, and the other type of equations are

$$rk_1^{(i)} + \sum_{j=2}^\ell x_j^{(i)} = S^{-1}(x_1^{(i)} - x_1^{(i-1)}) \quad (9)$$

where $x_j^{(i)} = x_j^{(0)} + \sum_{t=1}^i S(x_1^{(t-1)} + rk_j^{(t)})$ for $i = 1, 2, \dots, r$. Hence, we obtain $4(\ell + r - 1)$ equations of degree 3 in $256 + 4(r-1)$ variables. From m keystream blocks, one obtains $4(\ell + r - 1)m$ equations of degree 3 in $256 + 4(r-1)m$ variables.

Remark 2. Both equations (8) and (9) use the explicit representation of S . If the implicit representation is used, then the resulting system of equations is of degree 4. Since the XL attacks are based on the strategy of extending equations of lower degrees to a larger number of equations of higher degrees by multiplying polynomials (or monomials), we only consider the system of equations of a lower degree if the same variables are used.

GUESSING THE FIRST BRANCH. It is possible to guess the intermediate states of the first branch $x_1^{(i)}$ for $i = 1, 2, \dots, r-1$, in which case we obtain equations of the same degree except that equation (9) becomes linear for $i = r$ since $x_j^{(r)}$ are known for $j = 1, 2, \dots, \ell$. This linear equation determines the value of $rk_1^{(r)}$, reducing the number of variables by 4. Hence, by guessing $4(r-1)$ bits of the intermediate state, one obtains a system of $4(\ell + r - 2)$ equations of degree 3 in 252 variables.

When m keystream blocks are used, the round key $rk_1^{(1)}$ fixed in the first block determines $x_1^{(r-1)}$ for all the other blocks. Therefore, one obtains a system of $4(\ell + r - 2)m$ equations of degree 3 in 252 variables by guessing $4(r-1) + 4(r-2)(m-1)$ bits of the intermediate state of the first branch.

XL ATTACK COMPLEXITY. Table 2 summarizes the complexity of the XL attacks for the above systems according to the number of keystream blocks used to build the systems.

One can find that FRAS is secure against the XL attacks under 128-bit security even with the independent assumption.

Table 2: Complexity of the hybrid XL attacks for each system. g_{opt} denotes the optimal number of guessing for the hybrid XL attack, D denotes the solving degree, and ‘Cost’ denotes the attack complexity in bits. For the system guessing the first branch, guessing $32m$ bits of the first branch is considered in its cost where m is the number of blocks.

# Blocks	Intermediate Variables					First Branch Variables					Guessing First Branch				
	# Var	# Eqs	g_{opt}	D	Cost	# Var	# Eqs	g_{opt}	D	Cost	# Var	# Eqs	g_{opt}	D	Cost
1	1152	5376	3	16	239.56	284	156	127	26	323.85	252	152	97	26	320.82
2	2045	10752	10	20	327.32	312	312	120	26	333.02	252	304	70	25	335.01
3	2944	16128	0	24	394.81	340	468	124	26	346.39	252	456	46	25	335.01
4	3840	21504	0	27	456.45	368	624	8	39	356.90	252	760	18	26	359.92

5.1.3 Gröbner Basis Attack

We experimentally computed Gröbner basis for the systems described in Section 5.1.2 on toy parameters and found that the highest degree reached during the computation is not well estimated by the degree of regularity from (4). Instead, we experimentally verified that the actual Gröbner basis computation time grows exponentially according to the number of rounds and the number of branches. Figure 4 shows the Gröbner basis computation time of the systems obtained from a single input-output pair. For the systems obtained from multiple input-output pairs, see Supplementary Material B.

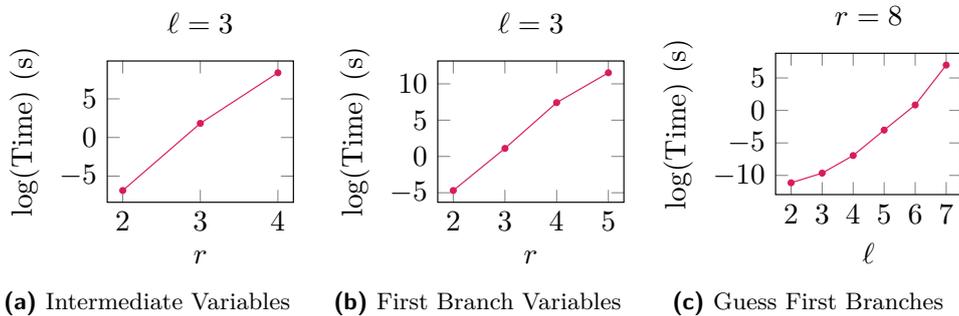


Figure 4: Gröbner basis computation time of the systems on toy parameters. The number of rounds and branches are denoted by r and ℓ , respectively. The key size is set to 4ℓ bits, which is half the actual key size.

5.1.4 Other Algebraic Attacks

The fast exhaustive search attack [BCC⁺10] can be seen as an optimized variant of the brute-force search. It evaluates Boolean equations of degree d in n variables using $d \cdot 2^n$ bit operations. However, the key length of 256 bits mitigates the fast exhaustive search attack against FRAS.

The polynomial method [Bei93] is a technique for solving a system of multivariate polynomial equations. For a system of Boolean equations of degree d in n variables, its time complexity is estimated as $n^2 \cdot 2^{(1-1/(2.7d))n}$ in bit operations [Din21]. The polynomial method is also infeasible for FRAS using a secret key of 256 bits.

The interpolation attack [JK97] is to establish an encryption polynomial only in plaintext variables by considering the secret key as an (unknown) constant. The input to

FRAST is fixed as constant \mathbf{ic} and the encryption function is distinct for every encryption, so it seems infeasible to apply the interpolation attack in a straightforward manner. On the other hand, one might try to establish a polynomial in the XOF outputs (that determine the random S-boxes). In this case, however, the number of variables is much larger and the degree growth by lookup operation on the random S-box is much faster compared to the systems only in key variables described in Section 5.1.1.

5.2 Linear Attacks

The linear attack was originally introduced for binary spaces [Mat94], but it can also be applied to non-binary spaces [BSV07]. The linear probability of a function $E : \mathbb{Z}_p^\ell \rightarrow \mathbb{Z}_p^\ell$ with input and output masks $\mathbf{a}, \mathbf{b} \in \mathbb{Z}_p^\ell$ is defined as follows:

$$\text{LP}^E(\mathbf{a}, \mathbf{b}) = \left| \mathbb{E}_{\mathbf{x}} \left[\exp \left\{ \frac{2\pi i}{p} (-\langle \mathbf{a}, \mathbf{x} \rangle + \langle \mathbf{b}, E(\mathbf{x}) \rangle) \right\} \right] \right|^2$$

where \mathbf{x} is uniformly distributed over \mathbb{Z}_p^ℓ . We refer to [BSV07] for the details.

Before going into the details, we recap the condition on input/output linear masks for some linear operations. For a branching operation $x \mapsto (x, x)$ with input mask u and output mask (v_1, v_2) , $u = v_1 + v_2$ should be satisfied. For an addition operation $(x_1, x_2) \mapsto x_1 + x_2$ with input mask (u_1, u_2) and output mask v , $u_1 = u_2 = v$ should be satisfied. See Figure 5.

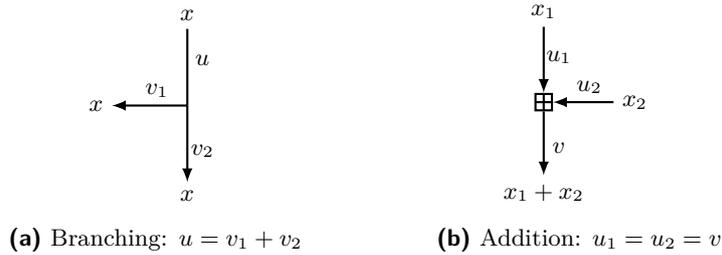


Figure 5: Linear masks for branching and addition

Traditional linear cryptanalysis that requires many input-output pairs of a fixed function does not apply to FRAST since the keystream generating function of FRAST changes for each keystream block. So in order to analyze the resistance of FRAST against linear cryptanalysis, we consider the following three strategies to be combined with linear attacks.

1. Applying nonzero linear masks on the XOF outputs generating the random S-boxes of FRAST.
2. Collecting input-output pairs that can be analyzed by the same linear trail.
3. Collecting input-output pairs whose random S-boxes have linear relations.

In this section, we show how the above strategies are mitigated by the randomly generated S-boxes of FRAST.

5.2.1 Analysis with Linear Masks on XOF Outputs

The first approach is to apply nonzero linear masks on the XOF outputs that determine random S-boxes of FRAST. Although the XOF outputs themselves are not controllable by an attacker, they can be considered as additional inputs in the KPA model since they are publicly open. Using this approach, one can apply the linear cryptanalysis to FRAST

by considering it as a fixed function whose input size is larger than its output size. A negacyclic random S-box S itself can be described by its function values $(S(0), \dots, S(7))$, and the function $\text{LookUp}_k(x, S) = S(x+k)$ can be described as a function from $\mathbb{Z}_{16} \times \mathbb{Z}_{16}^8$ to \mathbb{Z}_{16} defined by

$$\text{LookUp}_k(x, S) = \sum_{i=0}^7 (\mathbf{1}\{x+k=i\} - \mathbf{1}\{x+k=i+8\}) S(i)$$

where $\mathbf{1}\{x+k=i\} = 1$ if $x+k=i$ and 0 otherwise. From this point of view, the linear probability of FRAST with additional linear masks on the XOF outputs is well-defined.

For $(\mathbf{x}, S_{\text{erf}}, S_{\text{crf}}) \in \mathbb{Z}_{16}^\ell \times \mathbb{Z}_{16}^8 \times \mathbb{Z}_{16}^8$, define $\text{RF}[\mathbf{k}](\mathbf{x}, S_{\text{erf}}, S_{\text{crf}}) = \mathbf{y} \in \mathbb{Z}_{16}^\ell$ as follows.

$$\begin{aligned} y_j &= x_j + \text{LookUp}_{rk_j}(x_1, S_{\text{erf}}) \text{ for } j = 2, \dots, \ell, \\ y_1 &= x_1 + \text{LookUp}_{rk_1}(y_2 + \dots + y_\ell, S_{\text{crf}}), \end{aligned}$$

where $\mathbf{rk} = (rk_1, \dots, rk_\ell) \in \mathbb{Z}_{16}^\ell$ is a round key derived from the master key \mathbf{k} , $\mathbf{x} = (x_1, \dots, x_\ell) \in \mathbb{Z}_{16}^\ell$, $\mathbf{y} = (y_1, \dots, y_\ell) \in \mathbb{Z}_{16}^\ell$, and S_{erf} and S_{crf} are the negacyclic S-boxes derived from the XOF. As S_{erf} and S_{crf} are independently sampled, we separate the round function RF into two parts for simplicity: RF_{erf} and RF_{crf} .

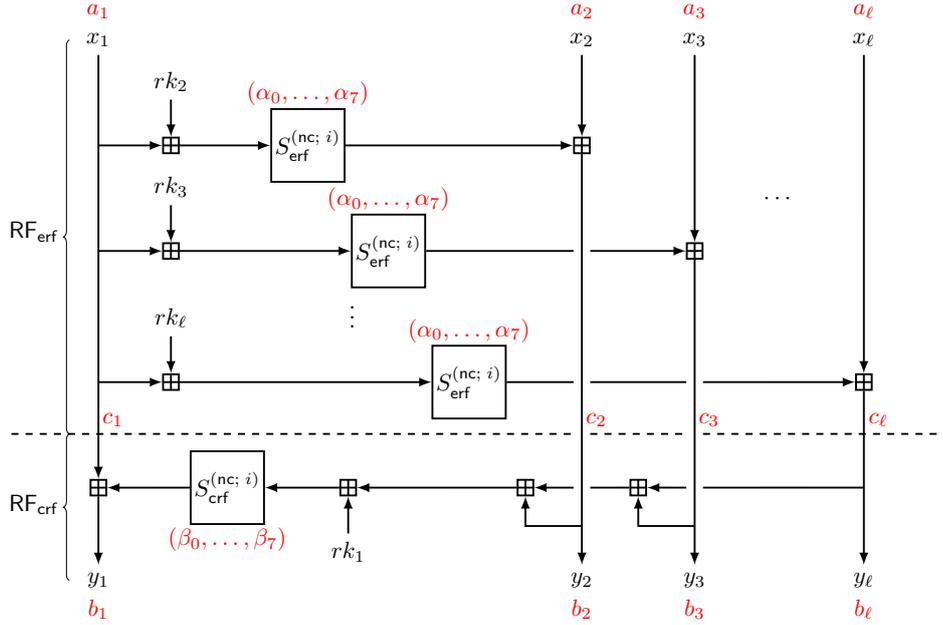


Figure 6: Linear masks in a single round of the FRAST considering XOF.

RF_{erf} is the expanding part of the round function. For $(\mathbf{x}, S_{\text{erf}}) \in \mathbb{Z}_{16}^\ell \times \mathbb{Z}_{16}^8$, $\text{RF}_{\text{erf}}[\mathbf{k}](\mathbf{x}, S_{\text{erf}}) = \mathbf{y} \in \mathbb{Z}_{16}^\ell$ is defined as follows.

$$\begin{aligned} y_j &= x_j + \text{LookUp}_{rk_j}(x_1, S_{\text{erf}}) \text{ for } j = 2, \dots, \ell, \\ y_1 &= x_1, \end{aligned}$$

where $(rk_2, \dots, rk_\ell) \in \mathbb{Z}_{16}^{\ell-1}$ is a round key derived from the master key \mathbf{k} , $\mathbf{x} = (x_1, \dots, x_\ell) \in \mathbb{Z}_{16}^\ell$, $\mathbf{y} = (y_1, \dots, y_\ell) \in \mathbb{Z}_{16}^\ell$, and S_{erf} is the negacyclic S-box derived from the XOF.

RF_{crf} is the contracting part of the round function. For $(\mathbf{x}, S_{\text{crf}}) \in \mathbb{Z}_{16}^\ell \times \mathbb{Z}_{16}^8$, $\text{RF}_{\text{crf}}[\mathbf{k}](\mathbf{x}, S_{\text{crf}}) = \mathbf{y} \in \mathbb{Z}_{16}^\ell$ is defined as follows.

$$\begin{aligned} y_j &= x_j \text{ for } j = 2, \dots, \ell, \\ y_1 &= x_1 + \text{LookUp}_{rk_1}(x_2 + \dots + x_\ell, S_{\text{crf}}), \end{aligned}$$

where $rk_1 \in \mathbb{Z}_{16}$ is a round key derived from the master key \mathbf{k} , $\mathbf{x} = (x_1, \dots, x_\ell) \in \mathbb{Z}_{16}^\ell$, $\mathbf{y} = (y_1, \dots, y_\ell) \in \mathbb{Z}_{16}^\ell$, and S_{crf} is the negacyclic S-box derived from the XOF.

Then, the round function RF can be described as a composition of RF_{erf} and RF_{crf} as follows.

$$\text{RF}[\mathbf{k}](\mathbf{x}, S_{\text{erf}}, S_{\text{crf}}) = \text{RF}_{\text{crf}}[\mathbf{k}](\text{RF}_{\text{erf}}[\mathbf{k}](\mathbf{x}, S_{\text{erf}}), S_{\text{crf}}).$$

We depict the relation between RF, RF_{erf} and RF_{crf} in Figure 6. By separating RF into RF_{erf} and RF_{crf} , we can compute the linear probability of RF from that of RF_{erf} and RF_{crf} .

Let $\mathbf{a} = (a_1, \dots, a_\ell)$ be an input mask to \mathbf{x} , $\boldsymbol{\alpha} = (\alpha_0, \dots, \alpha_7)$ (resp. $\boldsymbol{\beta} = (\beta_0, \dots, \beta_7)$) be an input mask to S_{erf} (resp. S_{crf}), and $\mathbf{b} = (b_1, \dots, b_\ell)$ be an output mask to \mathbf{y} . To represent the linear probability of RF with respect to those of RF_{erf} and RF_{crf} , let $\mathbf{c} = (c_1, \dots, c_\ell)$ be an output (resp. input) mask of RF_{erf} (resp. RF_{crf}) (see Figure 6). Then, the linear relation on RF forces $c_j = a_j$ for all $j = 2, \dots, \ell$, $c_1 = b_1$, and $c_2 - b_2 = \dots = c_\ell - b_\ell$. For such \mathbf{c} , we obtain the following.

$$\text{LP}^{\text{RF}[\mathbf{k}]}((\mathbf{a}, \boldsymbol{\alpha}, \boldsymbol{\beta}), \mathbf{c}) = \text{LP}^{\text{RF}_{\text{erf}}[\mathbf{k}]}((\mathbf{a}, \boldsymbol{\alpha}), \mathbf{c}) \cdot \text{LP}^{\text{RF}_{\text{crf}}[\mathbf{k}]}((\mathbf{c}, \boldsymbol{\beta}), \mathbf{b}).$$

LINEAR PROBABILITY OF RF_{erf} . Let $(\mathbf{a}, \boldsymbol{\alpha}) \in \mathbb{Z}_{16}^\ell \times \mathbb{Z}_{16}^8$ be the input mask and $\mathbf{c} \in \mathbb{Z}_{16}^\ell$ be the output mask to $\text{RF}_{\text{erf}}[\mathbf{k}]$ where $c_j = a_j$ for $j = 2, \dots, \ell$. Then, the linear probability of $\text{RF}_{\text{erf}}[\mathbf{k}]$ is given as follows.

$$\begin{aligned} \text{LP}^{\text{RF}_{\text{erf}}[\mathbf{k}]}((\mathbf{a}, \boldsymbol{\alpha}), \mathbf{c}) &= \frac{1}{16^2} \left| \sum_{x_1=0}^{15} \exp\left(\frac{2\pi i}{16}(c_1 - a_1)x_1\right) \right. \\ &\quad \left. \times \mathbf{1} \left\{ \sum_{i=2}^{\ell} c_i (\mathbf{1}_{\{x_1 + rk_i = j\}} - \mathbf{1}_{\{x_1 + rk_i = j+8\}}) = \alpha_j \forall j \in \{0, \dots, 7\} \right\} \right|^2 \end{aligned}$$

where (rk_2, \dots, rk_ℓ) is a part of the round key derived from \mathbf{k} . One can see that the above linear probability depends on the relation between the masks and the keys, which is not the case in traditional linear cryptanalysis. Hence, the masks should be chosen carefully to satisfy the following relation.

$$\exists x_1 \in \mathbb{Z}_{16}; \sum_{i=2}^{\ell} c_i (\mathbf{1}_{\{x_1 + rk_i = j\}} - \mathbf{1}_{\{x_1 + rk_i = j+8\}}) = \alpha_j \forall j = 0, \dots, 7. \quad (10)$$

Otherwise, the linear probability would be zero.

For an attacker who does not know the round key, there are two possible ways to build a trail of nonzero linear probability: a trivial linear trail such that $\boldsymbol{\alpha} = \mathbf{0}$, and a linear trail such that only one component of $\boldsymbol{\alpha}$ is nonzero.

The first approach is to build a linear trail that does not activate the LookUp function by setting $\boldsymbol{\alpha} = \mathbf{0}$ and $c_2 = \dots = c_\ell = 0$, which also implies that $a_2 = \dots = a_\ell = 0$. Then, by setting $c_1 = a_1 \neq 0$, one obtain $\text{LP}^{\text{RF}_{\text{erf}}[\mathbf{k}]}((\mathbf{a}, \boldsymbol{\alpha}), \mathbf{c}) = 1$ for nonzero input/output masks. This is the trivial linear trail of linear probability 1 on RF_{erf} .

The other approach is to set only one component of (c_2, \dots, c_ℓ) and $\boldsymbol{\alpha}$ to 8 and the others to 0. Then, regardless of the round key, there exists a unique $z \in \{0, \dots, 7\}$ such that (10) holds for $x_1 = z$ and $x_1 = z + 8$. By setting $2 \mid (c_1 - a_1)$, one obtain $\text{LP}^{\text{RF}_{\text{erf}}[\mathbf{k}]}((\mathbf{a}, \boldsymbol{\alpha}), \mathbf{c}) = 2^{-6}$.

LINEAR PROBABILITY OF RF_{crf} . Let $(\mathbf{c}, \beta) \in \mathbb{Z}_{16}^\ell \times \mathbb{Z}_{16}^8$ be the input mask and $\mathbf{b} \in \mathbb{Z}_{16}^\ell$ be the output mask to $\text{RF}_{\text{crf}}[\mathbf{k}]$ where $c_1 = b_1$ and $c_2 - b_2 = \dots = c_\ell - b_\ell$. Then the linear probability of $\text{RF}_{\text{crf}}[\mathbf{k}]$ is given as follows.

$$\text{LP}^{\text{RF}_{\text{crf}}[\mathbf{k}]}((\mathbf{c}, \beta), \mathbf{b}) = \frac{1}{16^2} \left| \sum_{x=0}^{15} \exp\left(\frac{2\pi i}{16}(b_2 - c_2)x\right) \right. \\ \left. \times \mathbf{1} \left\{ b_1 (\mathbf{1}_{\{x+rk_1=j\}} - \mathbf{1}_{\{x+rk_1=j+8\}}) = \beta_j \forall j \in \{0, \dots, 7\} \right\} \right|^2$$

where rk_1 is the first component of the round key derived from \mathbf{k} . The masks should be chosen carefully to satisfy the following to build a linear trail of nonzero linear probability.

$$\exists x \in \mathbb{Z}_{16}; b_1 (\mathbf{1}_{\{x+rk_1=j\}} - \mathbf{1}_{\{x+rk_1=j+8\}}) = \beta_j \forall j = 0, \dots, 7. \quad (11)$$

Similar to the case of RF_{erf} , an attacker can build two kinds of linear trails of nonzero linear probability without knowing the round key. One is the trivial linear trail to set $\beta = \mathbf{0}$ and $b_1 = c_1 = 0$, obtaining $\text{LP}^{\text{RF}_{\text{crf}}[\mathbf{k}]}((\mathbf{b}, \beta), \mathbf{b}) = 1$. The other nontrivial trail is to set $b_1 = 8$, only one component of β by 8 and the others by 0, obtaining the linear probability of 2^{-6} provided that $2 \mid (c_2 - b_2)$.

COMBINING TWO RESULTS. One can build a linear trail on RF by combining those on RF_{erf} and RF_{crf} . However, combining two trivial trails on RF_{erf} and RF_{crf} is impossible since it implies that all the input/output masks are zero. Instead, it is possible to combine one of the trivial trails and the other nontrivial trail, resulting in the linear trail of linear probability 2^{-6} . Such trail activates only one `LookUp` function.

If more than two `LookUp` functions are activated, then the attacker should know the difference between the round keys used in the activated `LookUp` functions. The attacker might try to guess them, but it is infeasible since more than 128 bits need to be guessed for 32 random rounds of FRAST.

5.2.2 Analysis on Compatible Data

The linear attack is a kind of statistical attack that requires many input-output pairs of a fixed function. In FRAST, negacyclic S-boxes are independent randomly selected. Hence, we need to compute the probability that two input-output pairs of FRAST from independent random round functions can be used together to measure the linear bias for a given linear approximation.

LINEAR MASKS FOR THE FRAST ROUND FUNCTION. Suppose that an input linear mask $\mathbf{a} = (a_1, a_2, \dots, a_\ell)$ and an output linear mask $\mathbf{b} = (b_1, b_2, \dots, b_\ell)$ are used for a single round of FRAST. Let u_i denote an input mask and let v_i denote an output mask of an S-box whose output is added to x_i (see Figure 7). From the properties of the branching and addition operations, the following conditions must hold for the input and output masks to have a nonzero linear probability.

$$v_1 = b_1 \text{ and } v_i = a_i \text{ for } i = 2, 3, \dots, \ell, \quad (12)$$

$$u_1 = w_j = a_j - b_j \text{ for all } j = 2, 3, \dots, \ell. \quad (13)$$

Suppose that there are nonzero input and output masks $\mathbf{a} = (a_1, a_2, \dots, a_\ell)$ and $\mathbf{b} = (b_1, b_2, \dots, b_\ell)$ activating no S-box. Since all the S-boxes are not activated, we have $b_1 = 0$ and $a_2 = a_3 = \dots = a_\ell = 0$ by (12). We also have $b_2 = b_3 = \dots = b_\ell$ by (13). Then the linear probability $\text{LP}^{\text{RF}}(\mathbf{a}, \mathbf{b})$ of the round function for the masks \mathbf{a} and \mathbf{b} is given as

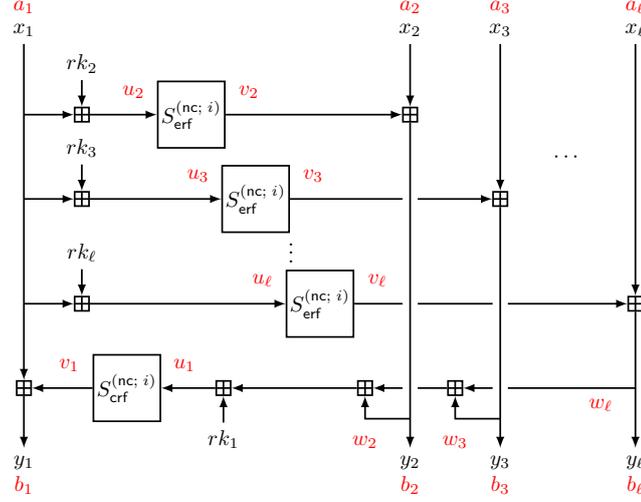


Figure 7: Linear masks in a single round of the FRAST without considering XOF.

follows.

$$\begin{aligned} \text{LP}^{\text{RF}}(\mathbf{a}, \mathbf{b}) &= \frac{1}{16^{2\ell}} \left| \sum_{x_2, \dots, x_\ell \in \mathbb{Z}_{16}} \exp \left\{ \frac{2\pi i}{16} (b_2 (x_2 + \dots + x_\ell)) \right\} \right|^2 \\ &\quad \times \left| \sum_{x_1 \in \mathbb{Z}_{16}} \exp \left\{ \frac{2\pi i}{16} \left(b_2 \sum_{j=2}^{\ell} S_{\text{erf}}^{(\text{nc}; i)}(x_1 + rk_j) - a_1 x_1 \right) \right\} \right|^2 \\ &= \mathbf{1} \{a_1 = b_2 = 0\} \end{aligned}$$

where $\mathbf{1} \{a_1 = b_2 = 0\}$ is 1 if $a_1 = b_2 = 0$ and 0 otherwise, $S_{\text{erf}}^{(\text{nc}; i)}$ is the randomly generated S-box used in the round function and rk_j is the round key added to the input of the S-box whose output is added to the j -th branch. So we conclude that the input and output masks activating no S-box cannot be both nonzero and trivial.

COMPATIBILITY OF A LINEAR TRAIL. Even for a round function using different negacyclic S-boxes, we can estimate its linear bias using the following condition: given two input-output pairs, suppose that their i -th round functions use S-boxes S_1 and S_2 , respectively. Since S_1 and S_2 work identically with respect to the output masks, one can use both pairs together to see if $b(S_1(x) - S_2(x))$ is constant for all $x \in \{0, 1, \dots, 7\} \subset \mathbb{Z}_{16}$ and for every nonzero output mask b for these negacyclic S-boxes. When S_1 and S_2 satisfy this condition, we say that S_1 and S_2 are *compatible* with respect to the output mask $b \neq 0$. If all the active S-boxes in the round functions are compatible with respect to their output masks, then we say that the round functions are compatible.

For independently sampled S_1 and S_2 , we can compute the probability that these two S-boxes are compatible with respect to b . The probability is maximized when $b = 8$, which is 2^{-7} . The compatibility with respect to the output mask of 8 means that the LSBs of the outputs of the two S-boxes are either identical or all different. Since every pair of nontrivial input and output masks for the FRAST round function has at least one active S-box, the probability that two round functions are compatible is also upper bounded by 2^{-7} .

Now, consider a linear trail $T = (\mathbf{a}^{(i)})_{i=0}^r$ for r random rounds of FRAST, where $\mathbf{a}^{(i)} = (a_1^{(i)}, a_2^{(i)}, \dots, a_\ell^{(i)}) \in \mathbb{Z}_{16}^\ell$ for $i = 0, 1, \dots, r$. So $\mathbf{a}^{(i-1)}$ and $\mathbf{a}^{(i)}$ become an input and

an output mask for the i -th round function, respectively, and the linear probability of the trail is given by the product of $\text{LP}(\mathbf{a}^{(i-1)}, \mathbf{a}^{(i)})$ for all $i = 1, 2, \dots, r$. To apply the trail on two input-output pairs, every round function for the two pairs should be compatible with respect to the output masks obtained by the trail. In this case, we say the two pairs are compatible with respect to the trail. The probability that the two pairs are compatible is upper bounded by 2^{-7r} , hence FRAST achieves 128-bit security against the linear attack if it has at least 19 random rounds.

5.2.3 Analysis with Random S-boxes Having Linear Relations

Suppose a negacyclic S-box S on \mathbb{Z}_{16} has a linear relation $ax + bS(x) = c$ for all $x = 0, 1, \dots, 15$. Then, the following holds from the negacyclic property of S .

$$a(x + 8) + bS(x + 8) = ax + 8a - bS(x) = c$$

over \mathbb{Z}_{16} for all $x = 0, 1, \dots, 7$. Combined with the original linear relation, one obtain

$$(2x + 8)a = 2c$$

over \mathbb{Z}_{16} for all $x = 0, 1, \dots, 7$. It can hold only if $a, c \in \{0, 8\}$. In this case, the function $bS(x)$ should be identical to one of $0, 8, 8x$, and $8x + 8$, which implies that $bS(0)$ and $bS(1)$ determines $bS(i)$ for $i = 2, 3, \dots, 7$. Hence, a negacyclic S-box over \mathbb{Z}_{16} has a linear relation with a probability at most 2^{-6} .

5.2.4 Zero-Correlation Attack

In contrast to the classical linear attack finding a high linear correlation, Bogdanov and Rijmen [BR14] proposed a variant of linear attack using linear hulls with correlation zero. This attack is based on the assumption that there might be a linear hull of correlation zero for every secret key (due to a certain specific structure of the block cipher), while it is not the case for a truly random permutation. Hence if one knows such a linear hull of correlation zero and collects 2^{n-1} input-output pairs under the same key, then the block cipher can be distinguished from a random permutation. Later, Bogdanov and Wang [BW12] reduced the data complexity down to $O(2^n/\sqrt{\ell})$, where ℓ is the number of the linear hulls with correlation zero.

The zero-correlation attack is not applicable to FRAST since the output keystream blocks of FRAST are not produced by a fixed permutation. To generate keystream blocks of FRAST, we feed a fixed input \mathbf{ic} to different encryption functions based on random S-boxes. If each keystream block of FRAST is regarded as an output of an independent permutation, then there will be no correlation between the keystream blocks, giving no distinguishing advantage to an adversary.

6 Performance Evaluation

In this section, we evaluate the server-side performance of FRAST in the transciphering framework with TFHE and compare it with Elisabeth (and their patches) and Kreyvium. We omit the previous result for FiLIP proposed by Hoffmann et al. [HMR20] since it takes more than a second to evaluate a single bit, which is far slower than Elisabeth and Kreyvium. Méaux et al. [MPP24] proposed a FINAL-based FiLIP evaluation method with a notable performance in terms of computation time, while we believe that it is hard to make an apples-to-apples comparison of the benchmark since it requires much larger key size of about 1 GB (even might be larger in the TFHE scheme).¹⁴ For the patches of Elisabeth, we

¹⁴The PARAM_MESSAGE_2_CARRY_2 parameters correspond to the case of $p = 2^5$ with Set-II.

only consider *Elisabeth-b* and *Gabriel* since *Margrethe* requires to evaluate a lookup table of 18-bit inputs, which is impractical.¹⁵ We also note that Deo et al. [DJL⁺24] recently proposed LWR-based transciphering with a notable performance, while we only compared our result with symmetric cipher-based transciphering.

The source codes of the server-side computation are developed in Rust with `tfhe-rs` library [Zam22] which supports the TFHE scheme. The extendable output function XOF has been instantiated with SHAKE256.¹⁶ Our experiments are executed in Intel i5-13600K @ 3.90 GHz.¹⁷

For homomorphic evaluation of FRAST, the default parameters of `tfhe-rs` library for `shortint` type, named `PARAM_MESSAGE_2_CARRY_2`, are used along with some chosen parameters for the GLWEtoGGSW conversion. Specifically, the following parameters have been used.

- GenPBS parameter
 - LWE parameters: $n = 742$, $\sigma_{\text{LWE}} = 7.06984 \times 10^{-6}$
 - GLWE parameters: $k = 1$, $N = 2048$, $\sigma_{\text{GLWE}} = 2.94036 \times 10^{-16}$
 - PBS parameters: $\log B_{\text{PBS}} = 23$, $\ell_{\text{PBS}} = 1$
 - Keyswitching parameters: $\log B_{\text{KS}} = 3$, $\ell_{\text{KS}} = 5$
- GLWEtoGGSW parameter
 - GGSW parameters of the GLWE secret key: $\log B_{\text{SK}} = 9$, $\ell_{\text{SK}} = 5$
 - GLWE keyswitching parameters: $\log B_{\text{subs}} = 9$, $\ell_{\text{subs}} = 5$
 - GGSW parameters of the round key bits: $\log B_{\text{rk}} = 7$, $\ell_{\text{rk}} = 3$

With the above parameters, 128-bit security is achieved and the error probability is upper bounded by 2^{-80} . See Supplementary Material E for detailed error analysis. For *Elisabeth* and *Kreyvium*, the TwoKS parameters in [CHMS22] and the parameters in [BOS23] are used, respectively. For the performance evaluation, we consider the case where the actual parameters after the transciphering are the default parameters of `tfhe-rs` library for `shortint` type as in [BOS23].

6.1 Benchmark and Comparison

The transciphering framework with a stream cipher requires only simple subtraction in the online phase, while, when it comes to TFHE, additional computation is required after the subtraction for plaintext encoding such as clearing carry bits and matching the plaintext encoding. *Kreyvium* clears the carry bit after subtraction by a single KSthenGenPBS operation [BOS23]. When it comes to *Elisabeth* producing ciphertexts of 4-bit keystream blocks without padding, Cosseron et al. proposed to use one-bit smaller plaintext space with padding for the resulting ciphertext [CHMS22]. Although the online phase only requires homomorphic subtraction for this plaintext encoding, it is not commonly used in TFHE. Most importantly, the ciphertext expansion ratio becomes greater than 1 since each ciphertext of 4 bits only contains a plaintext of 3 bits.

FRAST produces ciphertexts of 4-bit keystream blocks without padding, too. In order to support various types of plaintext encoding without ciphertext expansion, we add a bit extraction process at the end of the offline phase; a ciphertext of 4 bits is decomposed into 4 ciphertexts of a single bit whose plaintext encoding supports 1-bit plaintext space without

¹⁵The authors also left the homomorphic evaluation of *Margrethe* as an open problem.

¹⁶Our source-code is publicly available on <https://github.com/KAIST-CryptLab/FRAST>.

¹⁷It has 6 P-cores @ 5.30 GHz and 8 E-cores @ 3.90 GHz, and we only used the 8 E-cores for the benchmark.

carry and padding space. This plaintext encoding is also called 2-encoding in [BPR24]. This extraction can be done homomorphically at the cost of almost one PBS by the multi-value PBS [CIM19] since the extraction functions are negacyclic under the 2-encoding.

The online phase computes homomorphic subtraction, a single PBS operation to match the plaintext encoding of the resulting ciphertext, and a keyswitching to the final TFHE parameters. We note that the online phase performance does not depend on the cipher, but on the TFHE parameters it uses. Hence, by using a faster bootstrapping key for the online phase, the online performance can be improved. For FRAST, we use the bootstrapping key of Kreyvium for the bit extraction and the online phase. The bootstrapping key of Elisabeth has a similar performance to that of Kreyvium, so we can say that all the ciphers have almost the same online performance.

Table 3: Server-side offline phase performance. The setup time of Kreyvium is optimized in 4 threads. Keystream evaluation is performed in a single thread.

Cipher	Setup (s)	Keystream Evaluation	
		Lat. (ms)	Thrp. (bit/s)
Elisabeth	-	2049 (per 4-bit)	1.955
Elisabeth-b	-	5538 (per 4-bit)	0.749
Gabriel	-	4662 (per 4-bit)	0.858
Kreyvium	44.09 (4 threads)	134.0 (per 1-bit)	7.465
FRAST	24.99 (8 threads)	6194 (per 128-bit)	20.66

OFFLINE PHASE. The offline phase consists of the setup phase and the keystream evaluation phase. The setup phase is performed only once so that the latency (optimized by multithreading) should be seen as a more appropriate metric than the throughput. The setup time of Kreyvium is estimated as 1152 cycles for the main loop, where each cycle is estimated by 2 KSthenGenPBS operations using 4 threads.¹⁸ The setup time of FRAST is spent to convert the GLWE ciphertexts packing round key bits into GGSW ciphertexts of each round key bit, which can be optimized by using multiple threads. On the other hand, Elisabeth and its patches have no setup phase.

The latency for the keystream evaluation can be reduced by using multiple threads, while it is more efficient to evaluate each keystream block independently by each thread in terms of throughput. The offline performance of Kreyvium is estimated by only 7 KSthenGenPBS operations, evaluating a single keystream bit. The offline performance of Elisabeth is estimated by 96 KSthenGenPBS operations, evaluating a 4-bit keystream block, followed by 1 KSthenGenPBS operations for bit extraction. In case of Elisabeth-b (resp. Gabriel), evaluating 4-bit keystream block requires 252 (resp. 220) KSthenGenPBS operations, and the bit extraction requires 1 KSthenGenPBS. The result is summarized in Table 3. One can see that FRAST outperforms Elisabeth and Kreyvium in terms of throughput by factors of 10.57 and 2.768, respectively.

ONLINE PHASE. For all the ciphers, the online performance is estimated by two GenPBS operations followed by a single keyswitching to the default parameters for a ciphertext of a 2-bit message, and it only depends on the online phase parameters. In our setting, we obtain the latency of 46.08 ms for a ciphertext of a 2-bit message, which implies the throughput of 43.40 bit/s.

COMMUNICATION OVERLOAD. Communication overload is mainly due to homomorphic ciphertexts of the secret keys and the TFHE evaluation keys for the transciphering. The

¹⁸The setup time for Kreyvium in [BOS23] (which is called warm-up time) is obtained by dividing the time for 1152 cycles by 64, considering the bit size of `FheUInt64` type it computes. In this paper, we consider the total time required for initialization.

communication overload of Kreyvium (resp. Elisabeth) is estimated as 10.72 MB (resp. 12.29 MB).

When it comes to FRAST, the bootstrapping keys for the default parameters can be recycled in the actual usecase after transcribing, reducing the overall communication overload. Instead, the bootstrapping keys of Kreyvium are used for the bit extraction and the online phase, and additional evaluation keys for the double blind rotation are required. In this way, the communication overload for FRAST is estimated as 11.88 MB. See Supplementary Material F for details.

References

- [AFI⁺04] Gwénolé Ars, Jean-Charles Faugère, Hideki Imai, Mitsuru Kawazoe, and Makoto Sugita. Comparison Between XL and Gröbner Basis Algorithms. In Pil Joong Lee, editor, *ASIACRYPT 2004*, pages 338–353, Heidelberg, 2004. Springer.
- [AGP⁺19] Martin R. Albrecht, Lorenzo Grassi, Léo Perrin, Sebastian Ramacher, Christian Rechberger, Dragos Rotaru, Arnab Roy, and Markus Schofnegger. Feistel Structures for MPC, and More. In Kazue Sako, Steve Schneider, and Peter Y. A. Ryan, editors, *Computer Security – ESORICS 2019*, pages 151–171, Cham, 2019. Springer International Publishing.
- [BBB⁺23] Loris Bergerat, Anas Boudi, Quentin Bourgerie, Ilaria Chillotti, Damien Ligier, Jean-Baptiste Orfila, and Samuel Tap. Parameter Optimization and Larger Precision for (T)FHE. *Journal of Cryptology*, 36:28, 2023.
- [BCC⁺10] Charles Bouillaguet, Hsieh-Chung Chen, Chen-Mou Cheng, Tung Chou, Ruben Niederhagen, Adi Shamir, and Bo-Yin Yang. Fast Exhaustive Search for Polynomial Systems in \mathbb{F}_2 . In *CHES 2010*, pages 203–218. Springer, 2010.
- [BCP97] Wieb Bosma, John Cannon, and Catherine Playoust. The Magma algebra system. I. The user language. *J. Symbolic Comput.*, 24(3-4):235–265, 1997. Computational algebra and number theory (London, 1993).
- [Bei93] R. Beigel. The polynomial method in circuit complexity. In *[1993] Proceedings of the Eighth Annual Structure in Complexity Theory Conference*, pages 82–95, 1993.
- [BFS04] Magali Bardet, Jean-Charles Faugere, and Bruno Salvy. On the complexity of Gröbner basis computation of semi-regular overdetermined algebraic equations. In *Proceedings of the International Conference on Polynomial System Solving*, pages 71–74, 2004.
- [BFSS13] Magali Bardet, Jean-Charles Faugère, Bruno Salvy, and Pierre-Jean Spaenlehauer. On the complexity of solving quadratic Boolean systems. *Journal of Complexity*, 29(1):53–75, 2013.
- [BGV12] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, page 309–325. ACM, 2012.
- [BGV14] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) Fully Homomorphic Encryption without Bootstrapping. *ACM Trans. Comput. Theory*, 6(3), jul 2014.

- [BIP⁺22] Charlotte Bonte, Iliia Iliashenko, Jeongeun Park, Hilder V. L. Pereira, and Nigel P. Smart. FINAL: Faster FHE Instantiated with NTRU and LWE. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022*, pages 188–215, Cham, 2022. Springer.
- [BOS23] Thibault Balenbois, Jean-Baptiste Orfila, and Nigel Smart. Trivial Transciphering With Trivium and TFHE. In *Proceedings of the 11th Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, WAHC '23, page 69–78, New York, NY, USA, 2023. Association for Computing Machinery.
- [BPR24] Nicolas Bon, David Pointcheval, and Matthieu Rivain. Optimized Homomorphic Evaluation of Boolean Functions. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2024(3):302–341, Jul. 2024.
- [BR14] Andrey Bogdanov and Vincent Rijmen. Linear hulls with correlation zero and linear cryptanalysis of block ciphers. *Designs, Codes and Cryptography*, 70:369–383, 2014.
- [Bra12] Zvika Brakerski. Fully Homomorphic Encryption without Modulus Switching from Classical GapSVP. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417, pages 868–886. Springer, 2012.
- [BSV07] Thomas Baignères, Jacques Stern, and Serge Vaudenay. Linear Cryptanalysis of Non Binary Ciphers. In Carlisle Adams, Ali Miri, and Michael Wiener, editors, *Selected Areas in Cryptography*, volume 4876, pages 184–211. Springer, 2007.
- [BW12] Andrey Bogdanov and Meiqin Wang. Zero Correlation Linear Cryptanalysis with Reduced Data Complexity. In *FSE 2012, Washington, DC, USA, March 19–21, 2012. Revised Selected Papers*, pages 29–48. Springer, 2012.
- [CBSZ23] Pierre-Emmanuel Clet, Aymen Boudguiga, Renaud Sirdey, and Martin Zuber. ComBo: A Novel Functional Bootstrapping Method for Efficient Evaluation of Nonlinear Functions in the Encrypted Domain. In Nadia El Mrabet, Luca De Feo, and Sylvain Duquesne, editors, *AFRICACRYPT 2023*, pages 317–343. Springer, 2023.
- [CCF⁺18] Anne Canteaut, Sergiu Carpov, Caroline Fontaine, Tancrede 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, 2018.
- [CCR19] Hao Chen, Iliaria Chillotti, and Ling Ren. Onion Ring ORAM: Efficient Constant Bandwidth Oblivious RAM from (Leveled) TFHE. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, CCS '19, page 345–360. ACM, 2019.
- [CDPP22] Kelong Cong, Debajyoti Das, Jeongeun Park, and Hilder V.L. Pereira. Sorting-Hat: Efficient Private Decision Tree Evaluation via Homomorphic Encryption and Transciphering. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, CCS '22, page 563–577. ACM, 2022.
- [CGGI20] Iliaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. TFHE: Fast Fully Homomorphic Encryption Over the Torus. *Journal of Cryptology*, 33:34–91, 2020.

- [CHMS22] Orel Cosserson, Clément Hoffmann, Pierrick Méaux, and François-Xavier Standaert. Towards Case-Optimized Hybrid Homomorphic Encryption. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022*, pages 32–67. Springer, 2022.
- [CIM19] Sergiu Carpov, Malika Izabachène, and Victor Mollimard. New Techniques for Multi-value Input Homomorphic Evaluation and Applications. In Mitsuru Matsui, editor, *CT-RSA 2019*, pages 106–126. Springer, 2019.
- [CJP21] Ilaria Chillotti, Marc Joye, and Pascal Paillier. Programmable Bootstrapping Enables Efficient Homomorphic Inference of Deep Neural Networks. In Shlomi Dolev, Oded Margalit, Benny Pinkas, and Alexander Schwarzmann, editors, *Cyber Security Cryptography and Machine Learning*, pages 1–19, Cham, 2021. Springer International Publishing.
- [CKPS00] Nicolas Courtois, Alexander Klimov, Jacques Patarin, and Adi Shamir. Efficient Algorithms for Solving Overdefined Systems of Multivariate Polynomial Equations. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 392–407. Springer, 2000.
- [CLOT21] Ilaria Chillotti, Damien Ligier, Jean-Baptiste Orfila, and Samuel Tap. Improved Programmable Bootstrapping with Larger Precision and Efficient Arithmetic Circuits for TFHE. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021*, pages 670–699. Springer, 2021.
- [DCP08] Christophe De Cannière and Bart Preneel. Trivium. In Matthew Robshaw and Olivier Billet, editors, *New Stream Cipher Designs: The eSTREAM Finalists*, pages 244–266. Springer, 2008.
- [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, *Advances in Cryptology – CRYPTO 2018*, volume 10991, pages 662–692. Springer, 2018.
- [DGH⁺23] Christoph Dobraunig, Lorenzo Grassi, Lukas Helming, Christian Rechberger, Markus Schafneger, and Roman Walch. Pasta: A Case for Hybrid Homomorphic Encryption. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2023(3):30–73, Jun. 2023.
- [Din21] Itai Dinur. Cryptanalytic Applications of the Polynomial Method for Solving Multivariate Equation Systems over $\text{GF}(2)$. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021*, pages 374–403, Cham, 2021. Springer.
- [DJL⁺24] Amit Deo, Marc Joye, Benoit Libert, Benjamin R. Curtis, and Mayeul de Bellabre. Homomorphic Evaluation of LWR-based PRFs and Application to Transciphering. *Cryptology ePrint Archive*, Paper 2024/665, 2024.
- [Fau99] Jean-Charles Faugère. A new efficient algorithm for computing Gröbner bases (F_4). *Journal of Pure and Applied Algebra*, 139(1-3):61–88, 1999.
- [Fau02] Jean Charles Faugère. A new efficient algorithm for computing Gröbner bases without reduction to zero (F_5). In *Proceedings of the 2002 international symposium on Symbolic and algebraic computation*, pages 75–83. ACM, 2002.

- [Frö85] Ralf Fröberg. An Inequality for Hilbert Series of Graded Algebras. *MATHEMATICA SCANDINAVICA*, 56, Dec. 1985.
- [FV12] Junfeng Fan and Frederik Vercauteren. Somewhat Practical Fully Homomorphic Encryption. IACR Cryptology ePrint Archive, Report 2012/144, 2012. <https://eprint.iacr.org/2012/144>.
- [GHBJR23] Henri Gilbert, Rachele Heim Boissier, Jérémy Jean, and Jean-René Reinhard. Cryptanalysis of Elisabeth-4. In Jian Guo and Ron Steinfeld, editors, *Advances in Cryptology – ASIACRYPT 2023*, pages 256–284, Singapore, 2023. Springer Nature Singapore.
- [GMAH⁺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*, pages 305–339, Cham, 2023. Springer.
- [GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic Encryption from Learning with Errors: Conceptually-Simpler, Asymptotically-Faster, Attribute-Based. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013*, volume 8042, pages 75–92. Springer, 2013.
- [HKC⁺20] J. Ha, S. Kim, W. Choi, J. Lee, D. Moon, H. Yoon, and J. Cho. Masta: An HE-Friendly Cipher Using Modular Arithmetic. *IEEE Access*, 8:194741–194751, 2020.
- [HL20] Phil Hebborn and Gregor Leander. Dasta – Alternative Linear Layer for Rasta. *IACR Transactions on Symmetric Cryptology*, 2020(3):46–86, Sep. 2020.
- [HMR20] Clément Hoffmann, Pierrick Méaux, and Thomas Ricosset. Transciphering, Using FiLIP and TFHE for an Efficient Delegation of Computation. In Karthikeyan Bhargavan, Elisabeth Oswald, and Manoj Prabhakaran, editors, *INDOCRYPT 2020*, pages 39–61, Cham, 2020. Springer.
- [HMS24] Clément Hoffmann, Pierrick Méaux, and François-Xavier Standaert. The Patching Landscape of Elisabeth-4 and the Mixed Filter Permutator Paradigm. In Anupam Chattopadhyay, Shivam Bhasin, Stjepan Picek, and Chester Rebeiro, editors, *Progress in Cryptology – INDOCRYPT 2023*, pages 134–156, Cham, 2024. Springer Nature Switzerland.
- [JK97] Thomas Jakobsen and Lars R Knudsen. The Interpolation Attack on Block Ciphers. In *FSE’97*, pages 28–40. Springer, 1997.
- [KS22] Kamil Kluczniak and Leonard Schild. FDFB: Full Domain Functional Bootstrapping Towards Practical Fully Homomorphic Encryption. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2023(1):501–537, Nov. 2022.
- [LMP22] Zeyu Liu, Daniele Micciancio, and Yuriy Polyakov. Large-Precision Homomorphic Sign Evaluation Using FHEW/TFHE Bootstrapping. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022*, pages 130–160. Springer, 2022.
- [Mat94] Mitsuru Matsui. Linear Cryptanalysis Method for DES Cipher. In Tor Hellesest, editor, *EUROCRYPT ’93*, volume 765, pages 386–397. Springer, 1994.

- [MCJS19] Pierrick Méaux, Claude Carlet, Anthony Journault, and François-Xavier Standaert. Improved Filter Permutators for Efficient FHE: Better Instances and Implementations. In Feng Hao, Sushmita Ruj, and Sourav Sen Gupta, editors, *INDOCRYPT 2019*, volume 11898, pages 68–91. Springer, 2019.
- [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*, volume 9665, pages 311–343. Springer, 2016.
- [MPP24] Pierrick Méaux, Jeongeun Park, and Hilder V. L. Pereira. Towards Practical Transciphering for FHE with Setup Independent of the Plaintext Space. *IACR Communications in Cryptology*, 1(1), 2024.
- [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*, page 113–124. ACM, 2011.
- [Saa12] Markku-Juhani O. Saarinen. Cryptographic Analysis of All 4×4 -Bit S-Boxes. In Ali Miri and Serge Vaudenay, editors, *Selected Areas in Cryptography*, pages 118–133, Heidelberg, 2012. Springer.
- [YC04] Bo-Yin Yang and Jiun-Ming Chen. Theoretical Analysis of XL over Small Fields. In Huaxiong Wang, Josef Pieprzyk, and Vijay Varadharajan, editors, *Information Security and Privacy*, pages 277–288, Heidelberg, 2004. Springer.
- [YCBC07] Bo-Yin Yang, Owen Chia-Hsin Chen, Daniel J. Bernstein, and Jiun-Ming Chen. Analysis of QUAD. In Alex Biryukov, editor, *FSE 2007*, pages 290–308, Heidelberg, 2007. Springer.
- [YXS⁺21] Zhaomin Yang, Xiang Xie, Huajie Shen, Shiyong Chen, and Jun Zhou. Total: Fully homomorphic encryption with smaller parameters and stronger security. *Cryptology ePrint Archive*, Paper 2021/1347, 2021.
- [Zam22] Zama. TFHE-rs: A Pure Rust Implementation of the TFHE Scheme for Boolean and Integer Arithmetics Over Encrypted Data, 2022. <https://github.com/zama-ai/tfhe-rs>.
- [ZBRL15] Wentao Zhang, Zhenzhen Bao, Vincent Rijmen, and Meicheng Liu. A New Classification of 4-bit Optimal S-boxes and Its Application to PRESENT, RECTANGLE and SPONGENT. In Gregor Leander, editor, *FSE 2015*, pages 494–515, Heidelberg, 2015. Springer.

A Algebraic Representations of the FRAST Fixed S-box

Let $(x_0, \dots, x_3) \in \mathbb{F}_2^4$ (resp. $(y_0, \dots, y_3) \in \mathbb{F}_2^4$) be the input (resp. output) of S where x_0 (resp. y_0) is the LSB of the input (resp. output). The explicit representation of S over \mathbb{F}_2 is given as follows.

$$\begin{cases} y_0 = x_0x_1x_3 + x_0x_2x_3 + x_0 + x_1x_2 + x_1 + x_3 \\ y_1 = x_0x_1x_2 + x_0x_1x_3 + x_0x_1 + x_0 + x_1x_2x_3 + x_1x_2 + x_1x_3 + x_2x_3 + x_2 \\ y_2 = x_0x_1 + x_0x_2x_3 + x_0x_2 + x_0x_3 + x_1x_2x_3 + x_1x_2 + x_1x_3 + x_1 + x_2 + x_3 \\ y_3 = x_0x_1x_2 + x_0x_1 + x_0x_2 + x_1x_2 + x_2x_3 + x_3 \end{cases}$$

All the output bits are of degree 3 with respect to the input bits, and all the input bits work nonlinearly to the output as described in [Saa12]. For the backward direction, we obtain the following equations which have the same properties with the forward ones.

$$\begin{cases} x_0 = y_0y_1y_3 + y_0y_2y_3 + y_0y_2 + y_1y_3 + y_1 + y_2 + y_3 \\ x_1 = y_0y_1y_3 + y_0y_1 + y_0y_3 + y_1y_2y_3 + y_1 + y_2y_3 + y_2 + y_3 \\ x_2 = y_0y_2y_3 + y_0y_2 + y_0 + y_1y_2y_3 + y_1y_3 + y_1 + y_2y_3 \\ x_3 = y_0y_1y_2 + y_0y_2y_3 + y_0y_3 + y_0 + y_1 + y_2y_3 + y_2 \end{cases}$$

The implicit representations over \mathbb{F}_2 are given as follows.

$$\begin{cases} x_3x_2 + x_3y_3 + x_3y_1 + x_3y_0 = 0 \\ x_3x_2 + x_2x_1 + x_2x_0 + x_3y_2 + x_3y_0 + x_2 + x_1 + y_3 + y_1 + y_0 = 0 \\ x_3x_2 + x_2y_3 + x_3y_2 + x_3y_0 + x_2 + x_1 + y_3 + y_1 + y_0 = 0 \\ x_2x_1 + x_3x_0 + x_2y_2 + x_2 + y_2 + y_1 + y_0 = 0 \\ x_3x_2 + x_3x_1 + x_3x_0 + x_3y_3 + x_3y_2 + x_2y_1 + x_1 + y_3 + y_1 + y_0 = 0 \\ x_3x_1 + x_3y_3 + x_2y_0 + x_2 + x_0 + y_3 + y_1 = 0 \\ x_3x_0 + x_1x_0 + x_3 + y_3 + y_2 + y_1 + y_0 = 0 \\ x_1y_3 + x_3y_0 + x_3 + x_1 + x_0 + y_3 + y_2 + y_1 = 0 \\ x_2x_1 + x_3x_0 + x_3y_3 + x_1y_2 + x_3y_0 + x_2 + x_1 + x_0 + y_1 = 0 \\ x_3x_2 + x_3x_1 + x_1y_1 + x_3y_0 + x_3 + x_2 + x_1 + y_3 + y_2 = 0 \\ x_2x_1 + x_3x_0 + x_3y_3 + x_3y_2 + x_3y_0 + x_1y_0 + x_1 + y_3 + y_2 + y_1 + y_0 = 0 \\ x_3x_0 + x_0y_3 + x_3y_0 + x_3 + x_2 + x_0 + y_2 + y_0 = 0 \\ x_3x_2 + x_3x_1 + x_2x_1 + x_3y_3 + x_3y_2 + x_0y_2 + x_3 + x_2 + x_0 + y_3 + y_2 + y_0 = 0 \\ x_3x_2 + x_3x_1 + x_3x_0 + x_3y_3 + x_0y_1 + x_3 + x_2 + y_2 + y_0 = 0 \\ x_2x_1 + x_3x_0 + x_3y_2 + x_3y_0 + x_0y_0 + x_2 + y_1 = 0 \\ x_3x_2 + x_3x_1 + x_3x_0 + x_3y_2 + y_3y_2 + x_3 + x_2 + x_0 + y_3 + y_2 + y_0 = 0 \\ x_3x_1 + x_3x_0 + y_3y_1 = 0 \\ x_3x_1 + y_3y_0 + x_1 + x_0 + y_2 + y_1 = 0 \\ x_3x_2 + x_3x_1 + x_3x_0 + y_2y_1 + x_3y_0 + x_3 + x_1 + y_3 + y_1 + y_0 = 0 \\ x_3x_2 + x_2x_1 + y_2y_0 + x_1 + y_3 + y_2 + y_1 + y_0 = 0 \\ x_3x_1 + y_1y_0 + x_2 + y_2 + y_0 = 0 \end{cases}$$

There are 21 linearly independent quadratic equations for S , which is the minimum for 4-bit S-boxes.

B Gröbner Basis Computation on Toy Parameters

In this section, we summarize the experimental result of the Gröbner basis computation time on toy parameters. The source codes of the experiment are developed in MAGMA [BCP97], and are executed in AMD Ryzen 7 2700X @ 3.70 GHz with 128 GB memory. Let r , ℓ , and m denote the number of rounds, branches, and input/output pairs used to build a system, respectively.

Figure 8 shows the Gröbner basis computation time according to r when $\ell = 3$ for the system introducing new variables for all the intermediate states and the system introducing new variables for the states of the first branch. The key size is set to 4ℓ , which is half of the actual key size, to run the experiment on various parameters. One can see that the computation time grows exponentially according to r , and using larger m increases computation time.

Figure 9 shows the Gröbner basis computation time for the system by guessing the values of the first branch according to ℓ when $r = 8$. The key size is set to 8ℓ , which is the

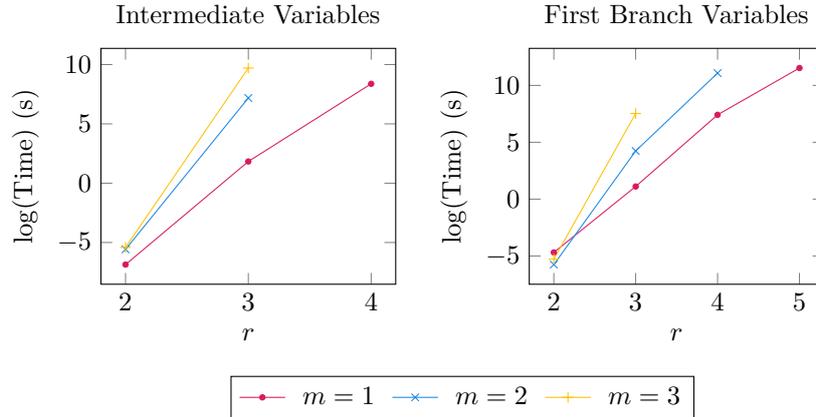


Figure 8: Gröbner basis computation time of the systems introducing new variables according to the number of rounds r . The number of branches ℓ is 3 and the key size is 4ℓ bits.

actual key size.¹⁹ One can see that the computation time grows exponentially according to ℓ . The peak at $(\ell, m) = (3, 3)$ is caused by some outliers in the data.

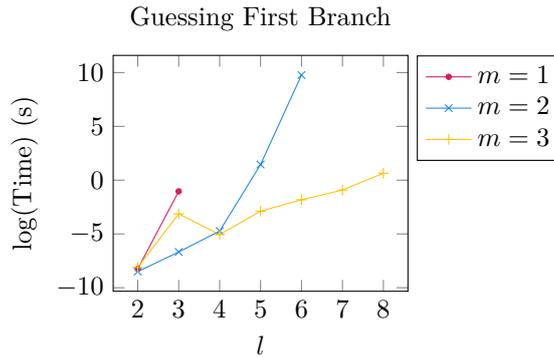


Figure 9: Gröbner basis computation time of the systems with guessing the values of the first branches. The number of round r is 8 and the key size is 8ℓ bits.

Conversely to the previous systems, using larger m tends to decrease the Gröbner basis computation time. The reason is that the number of variables does not change while the number of equations increases according to m . That said, this does not imply using larger m is advantageous for the attack since the number of guessing 2^{4rm} increases much faster.

Remark 3. The whole running time of the Gröbner basis computing program is much longer than the Gröbner basis computing time denoted in the graph, so we could not run the experiment on larger parameters.

C The Number of Monomials on Toy Parameters

In this section, we show the experimental result of the number of monomials appearing in the system of equations for the XOR-variant of FRAST in the key variables. We used toy parameters of $\ell = 4$ and the key size of 4ℓ bits where ℓ is the number of branches.

¹⁹The graph in Figure 4 uses the key size of 4ℓ .

We chose 100 random inputs, represented the corresponding outputs for a single and two fixed rounds of the XOR-variant of FRAS as polynomials in the key variables, and counted the number of monomials according to their degrees. Table 4 summarize the results. One can see that more than 90% of monomials of degree d appear in the system where d is the lower bound proposed in Section 5.1.1.

We also experimented in the backward direction considering the meet-in-the-middle attack, and summarized the results in Table 5. One can see that more than 90% of monomials of degree d appear after two rounds in the backward direction where d is the lower bound proposed in Section 5.1.1.

Table 4: The (average) number of monomials appearing in the fixed round functions of FRAS in the forward direction according to degrees. ‘Total’ denotes the total number of monomials, ‘First’ (resp. ‘Others’) denotes the number of monomials appearing in the first branch (resp. the other branches), and ‘Ratio’ denotes its ratio to the total number of monomials. The bold fonts denote the lower bound of the degree proposed in Section 5.1.1.

Degree	Total	Single Round				Two Rounds			
		First	Ratio	Others	Ratio	First	Ratio	Others	Ratio
1	16	15.24	0.953	16.00	1.000	14.91	0.932	16.00	1.000
2	120	112.74	0.940	119.80	0.998	112.52	0.938	119.93	0.999
3	560	524.99	0.937	542.68	0.969	525.45	0.938	559.78	0.999
4	1820	1707.71	0.938	-	-	1705.61	0.937	1819.62	0.999
5	4368	4091.44	0.937	-	-	4093.88	0.937	4367.08	0.999
6	8008	7485.90	0.935	-	-	7506.56	0.937	8006.04	0.999
7	11440	9989.81	0.873	-	-	10725.53	0.938	11437.21	0.999
8	12870	10852.44	0.843	-	-	12068.20	0.938	12866.73	0.999
9	11440	5254.97	0.459	-	-	10725.86	0.938	11428.58	0.999
10	8008	-	-	-	-	7505.68	0.937	7863.41	0.982
11	4368	-	-	-	-	4095.01	0.938	3930.39	0.900
12	1820	-	-	-	-	1705.20	0.937	-	-
13	560	-	-	-	-	524.46	0.937	-	-
14	120	-	-	-	-	112.95	0.941	-	-
15	16	-	-	-	-	14.17	0.886	-	-
16	1	-	-	-	-	0.89	0.890	-	-

D Function Decomposition of New WoP-PBS

For a t -bit integer $m \in \llbracket 0, 2^t \rrbracket$, let $\text{FlipMSB}_t(m) = (m + 2^{t-1}) \bmod 2^t \in \llbracket 0, 2^t \rrbracket$. Given an arbitrary function f on $\llbracket 0, p \rrbracket$ for a power-of-two p , f can be decomposed into $\log p$ functions $f_0, f_1, \dots, f_{\log p - 1}$ and one constant $f_{\log p}$ such that

$$f(x) = \sum_{j=0}^{\log p - 1} f_j(x \bmod p/2^j) + f_{\log p}$$

where f_j is a negacyclic function on $\llbracket 0, p/2^j \rrbracket$ defined as

$$f_j(x) = \frac{1}{2^{j+1}} \sum_{a=0}^{2^j - 1} (f(a \cdot 2^j + x) - f(a \cdot 2^j + \text{FlipMSB}_{\log p - j}(x)))$$

Table 5: The (average) number of monomials appearing in the fixed round functions of FRAST in the backward direction according to degrees. ‘Total’ denotes the total number of monomials, ‘First’ (resp. ‘Others’) denotes the number of monomials appearing in the first branch (resp. the other branches), and ‘Ratio’ denotes its ratio to the total number of monomials. The bold fonts denote the lower bound of the degree proposed in Section 5.1.1.

Degree	Total	Single Round				Two Rounds			
		First	Ratio	Others	Ratio	First	Ratio	Others	Ratio
1	16	16.00	1.000	16.00	1.000	14.97	0.936	16.00	1.000
2	120	108.07	0.901	119.92	0.999	112.53	0.938	119.96	0.999
3	560	409.45	0.731	558.90	0.998	524.48	0.937	559.88	0.999
4	1820	-	-	1815.19	0.997	1706.24	0.937	1819.42	0.999
5	4368	-	-	4213.67	0.965	4096.11	0.938	4366.93	0.999
6	8008	-	-	-	-	7504.14	0.937	8005.58	0.999
7	11440	-	-	-	-	10718.37	0.937	11437.00	0.999
8	12870	-	-	-	-	12049.91	0.936	12866.47	0.999
9	11440	-	-	-	-	9945.95	0.869	11436.99	0.999
10	8008	-	-	-	-	6671.91	0.833	8006.18	0.999
11	4368	-	-	-	-	1987.14	0.455	4366.63	0.999
12	1820	-	-	-	-	-	-	1810.18	0.995
13	560	-	-	-	-	-	-	529.75	0.946

for $j = 0, \dots, \log p - 1$, and $f_{\log p}$ is the constant given by

$$f_{\log p} = \frac{1}{p} \sum_{a=0}^{p-1} f(a).$$

This decomposition can be obtained by applying the decomposition of f into $\frac{1}{2}(f_{\text{msb-odd}} + f_{\text{msb-even}})$ described in Section 4.1 recursively on $f_{\text{msb-even}}$.

For example, let f be a function defined on $\llbracket 0, 4 \llbracket$. Then f is decomposed into two functions f_0, f_1 and one constant f_2 . The function f_0 on $\llbracket 0, 4 \llbracket$ is given by

$$\begin{aligned} f_0(0) &= \frac{1}{2}(f(0) - f(2)), \\ f_0(1) &= \frac{1}{2}(f(1) - f(3)), \\ f_0(2) &= \frac{1}{2}(f(2) - f(0)), \\ f_0(3) &= \frac{1}{2}(f(3) - f(1)), \end{aligned}$$

the function f_1 on $\llbracket 0, 2 \llbracket$ is given by

$$\begin{aligned} f_1(0) &= \frac{1}{4}(f(0) + f(2) - f(1) - f(3)), \\ f_1(1) &= \frac{1}{4}(f(1) + f(3) - f(0) - f(2)), \end{aligned}$$

and the constant f_2 is given by

$$f_2 = \frac{1}{4}(f(0) + f(1) + f(2) + f(3)).$$

E TFHE Evaluation Methods of FRAST with Error Analysis

In this section, we describe the TFHE evaluation of FRAST using the default parameters described in Section 6, and the error probability of it. For simplicity, we only describe how to evaluate the fixed round functions as evaluating the random round functions using negacyclic S-boxes is simpler.

E.1 TFHE Evaluation of FRAST

The FRAST round function consists of two parts: evaluating multiple S-boxes fed with the first branch added by multiple round keys and evaluating the last S-box of which output is added to the first branch. These two parts are called the expanding and the contracting parts of the round function, respectively. After evaluating FRAST, each bit of the keystream is extracted.

Let (x_1, \dots, x_ℓ) be input, (y_1, \dots, y_ℓ) be output of the FRAST fixed round function of which S-box is S . We consider the ciphertext modulus of $q = 2^{64}$, the message modulus of $p = 16$, and the scaling factor of $\Delta = q/p = 2^{60}$.

The input $x_i \in \llbracket 0, p \rrbracket$ is given as an LWE ciphertext scaled by Δ without padding, namely, $\text{LWE}(\Delta \cdot x_i)$, for $i = 1, \dots, \ell$. It can be also regarded as a ciphertext of $2x_i \in \llbracket 0, 2p \rrbracket$ with a scaling factor of $\Delta/2$ since

$$[\Delta \cdot x_i]_q = [(\Delta/2) \cdot (2x_i)]_q.$$

E.1.1 Expanding Part

The expanding part of the FRAST round function is evaluated by the double blind rotation as mentioned in Section 4.2. That said, for the first round, it is possible to evaluate the S-boxes without computing GenPBS on x_1 since the input to the first round is the known constant \mathbf{ic} . By giving GGSW ciphertexts of the round key bits used in the expanding part of the first round, one can evaluate the expanding part of the first round in a much smaller number of the CMux gates.

For the other rounds, decompose S into $S_{\text{msb-odd}}$ and $S_{\text{msb-even}}$. Since $S_{\text{msb-odd}}$ is negacyclic, it is possible to compute $\text{LWE}(\Delta \cdot S_{\text{msb-odd}}(x_1 + rk_j))$ for all $j = 2, \dots, \ell$ in a single GenPBS operation using the double blind rotation. Although the range of $S_{\text{msb-odd}}$ itself is not $\llbracket 0, p \rrbracket$, one can regard $\text{LWE}(\Delta \cdot S_{\text{msb-odd}}(x))$ as an LWE ciphertext of $2S_{\text{msb-odd}}(x) \in \llbracket 0, 2p \rrbracket$ with a scaling factor of $\Delta/2$ as mentioned above.

During the evaluation of $S_{\text{msb-odd}}$ on the input x_1 , it is also possible to extract the MSB bit of x_1 for free using PBSmanyLUT [CLOT21].²⁰ The extracted MSB bit of x_1 is subtracted from x_1 , obtaining an LWE ciphertext of $\text{ClearMSB}(x_1) \in \llbracket 0, p/2 \rrbracket$ that will be fed into $S_{\text{msb-even}}$.

To compute $S_{\text{msb-even}}$ on $\text{ClearMSB}(x_1)$ is possible with a single GenPBS operation using the cleared MSB bit of x_1 , while one more padding bit is required for the double blind rotation since $\Delta \cdot (\text{ClearMSB}(x_1) + \text{ClearMSB}(rk_i))$ might exceed $q/2$, filling the padding bit of the ciphertext. Since $S_{\text{msb-even}}$ is of only 3-bit precision, we address this issue by giving one more padding bit to the ciphertext of $\text{ClearMSB}(x_1)$. Computing $\text{LWE}((\Delta/2) \cdot \text{ClearMSB}(x_1))$ using one more GenPBS operation, one can apply the double blind rotation using GenPBS of 4-bit precision since $\frac{\Delta}{2}(\text{ClearMSB}(x_1) + \text{ClearMSB}(rk_i)) \in \llbracket q/2^5, q/2 \rrbracket$.

It is also possible to refresh the ciphertext of $\text{ClearMSB}(x_1)$ simultaneously during the double blind rotation. Adding the refreshed ciphertext of the MSB of x_1 and $\text{ClearMSB}(x_1)$,

²⁰The success probability of PBSmanyLUT is sensitive to the parameter and the error contained in the input when N is small, so that it should be used carefully. For FRAST, we have checked that PBSmanyLUT can be used with negligible failure probability.

which is obtained during adjusting its scaling factor from Δ to $\Delta/2$ to evaluate $S_{\text{msb-even}}$, one can refresh the first state before adding $S(y_2 + \dots + y_\ell + rk_1)$, making it possible to use PBSmanyLUT on the input of the first branch.

The number of the CMux gates also can be reduced with a simple tweak: to compute the blind rotation on $S(x_1 + rk_2)$ instead of $S(x_1)$. Then $S(x_1 + rk_j)$ is computed by the second blind rotation by the bits of $rk_j - rk_2$ for $j = 3, \dots, \ell$. It reduces the number of round key bits to be packed on the GLWE ciphertext, but rk_2 (for each round) should be sent to the server as an LWE ciphertext. Refreshing the state of the first branch also requires additional subtraction by rk_2 since the ciphertext of $x_1 + rk_2$ is refreshed. In summary, evaluating the expanding part of the FRAST round function requires 2 GenPBS operations followed by $7(\ell - 2)$ CMux gates.

E.1.2 Contracting Part

The contracting part computes $S(y_2 + \dots + y_\ell + rk_1)$. If the summation is directly computed from y_2, \dots, y_ℓ , the magnitude of the error inside the summation increases with the round. Instead, we use another variable, dubbed *crfsum*, to manage the noise for the summation.

At first, *crfsum* is initialized by the trivial encryption of $\sum_{j=2}^{\ell} \mathbf{ic}[j]$. In the expanding part, the output ciphertexts of $S(x_1 + rk_j)$ for $j = 2, \dots, \ell$ are added to *crfsum*. Then *crfsum* added by rk_1 becomes the input of the S-box S in the contracting part, which can be evaluated in 3 GenPBS operations using our WoP-PBS. Using the same idea to refresh the first branch, one can refresh *crfsum* by one more GenPBS operation. By the help of PBSmanyLUT, both evaluating the S-box and refreshing *crfsum* can be done using 2 GenPBS operations.

E.1.3 Bit Extraction

A ciphertext containing a FRAST keystream word of 4 bits is decomposed into 4 ciphertexts containing each keystream bit scaled by $q/2$ using the multi-value PBS [CIM19]. Let $x = b_0 + 2b_1 + 2^2b_2 + 2^3b_3$ be a keystream word where $b_i \in \{0, 1\}$ for $i = 0, \dots, 3$. The MSB b_3 can be extracted by the negacyclic function

$$x \mapsto (-1)^{b_3+1} \cdot \frac{q}{2},$$

followed by an addition of $q/2$. The other bits b_i can be extracted by the negacyclic functions

$$x \mapsto b_i \cdot \frac{q}{2}$$

for $i = 0, 1, 2$. All the extraction functions evaluate negacyclic functions on the same input x , so they can be evaluated at the cost of almost one PBS by the multi-value PBS [CIM19].

E.2 Error Analysis

In this section, we analyze the error growth in homomorphic keystream evaluation of FRAST. We use the following conventions as in [CLOT21].

- $\text{Var}(E) \leq \sigma^2$ for $E \in \mathcal{R}_{q,N}$ if all coefficients of E have variances at most σ^2 .
- $\text{Var}(\mathbf{c}) \leq \sigma^2$ for a GLWE ciphertext \mathbf{c} if its phase E satisfies $\text{Var}(E) \leq \sigma^2$.
- $\text{Var}(\mathbf{C}) \leq \sigma^2$ for a GGSW ciphertext $\mathbf{C} = (\mathbf{C}^{(i,j)})_{(i,j) \in [k+1] \times [\ell]}$ if all of its GLWE ciphertext components $\mathbf{C}^{(i,j)}$ satisfies $\text{Var}(\mathbf{C}^{(i,j)}) \leq \sigma^2$.

Every error associated with those ciphertexts follows a zero-mean distribution in this section. We note that the error analysis in [CCR19] uses a different convention in terms of describing the error, so we properly translated its result to the convention used in this paper with slightly improved bound used in [CLOT21]. The formal proof described in [CLOT21] is way complicated, so we only give a sketch of how the results in [CLOT21] can be applied in the following lemmas.

Lemma 2 (Theorem 4.1 in [CCR19]). *Suppose GLWetoGLWE keyswitching keys*

$$\text{KS}_m = \left\{ \text{GLWE}_{\mathbf{S}(X^m)} \left(S_i \cdot \frac{q}{B_{\text{subs}}^j} \right) \right\}_{(i,j) \in [k] \times [\ell_{\text{subs}}]}$$

such that $\text{Var}(\text{KS}_m) \leq \sigma_{\text{subs}}^2$ for all $m = N/2^{i-1} + 1$, $i = 1, \dots, \log N$ are given. For an input $\mathbf{c} = \text{GLWE}_{\mathbf{S}}(\sum_{i=0}^{N-1} b_i X^i)$, Algorithm 3 in [CCR19] outputs a set of ciphertexts $\{\mathbf{c}_j = \text{GLWE}_{\mathbf{S}}(Nb_j)\}_{j=0}^{N-1}$ with noise variance

$$\text{Var}(\mathbf{c}_j) \leq N^2 \text{Var}(\mathbf{c}) + \frac{N^2 - 1}{3} V_{\text{sk}}$$

where

$$V_{\text{sk}} = \frac{kN}{2} \left(\frac{q^2}{12B_{\text{subs}}^{2\ell_{\text{subs}}}} - \frac{1}{12} \right) + \frac{kN}{16} + kN\ell_{\text{subs}}\sigma_{\text{subs}}^2 \frac{B_{\text{subs}}^2 + 2}{12}.$$

Proof. The noise increment V_{sk} of the GLWetoGLWE keyswitching can be estimated by the result of Appendix D in [CLOT21] with a slight modification. For each iteration on i in Algorithm 3 in [CCR19], the noise increases by V_{sk} for each GLWetoGLWE keyswitching. Then one can obtain the above upper bound following the proof of Theorem 4.1 in [CCR19] (considering the difference of the convention as mentioned before). \square

Lemma 3 (Theorem 4.2 in [CCR19]). *Let $\mathbf{A} = \{\text{GGSW}_{\mathbf{S}}(-S_i)\}_{i=1}^k$ be a set of GGSW ciphertexts of $-S_i$ for $i = 1, \dots, k$ with the base B_{SK} and level ℓ_{SK} . For inputs \mathbf{A} and GLWE ciphertexts $\{\mathbf{c}_j = \text{GLWE}_{\mathbf{S}}(\sum_{i=0}^N \frac{b_i X^i}{NB_{\text{rk}}^{i+1}})\}_{j=1}^{\ell_{\text{rk}}}$, Algorithm 4 in [CCR19] outputs GGSW ciphertexts $\mathbf{C}_i = \text{GGSW}_{\mathbf{S}}(b_i)$ with the base B_{rk} and level ℓ_{rk} of which noise variance is given by*

$$\begin{aligned} \text{Var}(\mathbf{C}_i) &\leq N^2 \text{Var}(\mathbf{c}) + \frac{N^2 - 1}{3} V_{\text{sk}} + \ell_{\text{SK}}(k+1)N \frac{B_{\text{SK}}^2 + 2}{12} \text{Var}(\mathbf{A}) \\ &\quad + \frac{q^2 - B_{\text{SK}}^{2\ell_{\text{SK}}}}{12B_{\text{SK}}^{2\ell_{\text{SK}}}} \left(1 + \frac{kN}{2} \right) + \frac{kN}{8} + \frac{1}{4} \left(1 - \frac{kN}{2} \right)^2 \end{aligned} \quad (14)$$

for all $i = 0, \dots, N-1$.

Proof. One can obtain the above bound following the proof of Theorem 4.2 in [CCR19] directly (considering the difference of the convention as mentioned before), except that the GGSW ciphertexts \mathbf{A} and \mathbf{C}_i have different decomposition base and level. \square

We denote the upper bound in (14) by σ_{dbr}^2 for the rest of this section.

Lemma 4 (Theorem 4 in [CLOT21]). *The noise variance of the GenPBS output is*

$$\begin{aligned} \text{Var}(\text{PBS}) &= n\ell(k+1)N \frac{B_{\text{PBS}}^2 + 2}{12} \sigma_{\text{PBS}}^2 \\ &\quad + n \frac{q^2 - B_{\text{PBS}}^{2\ell_{\text{PBS}}}}{24B_{\text{PBS}}^{2\ell_{\text{PBS}}}} \left(1 + \frac{kN}{2} \right) + \frac{nkN}{32} + \frac{n}{16} \left(1 - \frac{kN}{2} \right)^2 \end{aligned}$$

where σ_{PBS} is the noise variance of the bootstrapping key, and B_{PBS} and ℓ_{PBS} are the base and level of GenPBS, respectively.

Lemma 5. *Let \mathbf{C}_i be a GGSW ciphertext of noise variance at most σ^2 of the base B and the level of ℓ for $i = 1, \dots, t$. Let \mathbf{c} be a GLWE ciphertext of noise variance at most σ_{prev}^2 . When the nested external products on \mathbf{c} by $\mathbf{c}' = \mathbf{C}_t \boxtimes \dots \boxtimes (\mathbf{C}_2 \boxtimes (\mathbf{C}_1 \boxtimes \mathbf{c}))$ is well-defined, the noise variance of \mathbf{c}' is bounded by*

$$\begin{aligned} \text{Var}(\mathbf{c}') &\leq \sigma_{\text{prev}}^2 + t^2 \ell (k+1) N \frac{B^2 + 2}{12} \sigma^2 \\ &\quad + t \left(\frac{q^2 - B^{2\ell}}{12B^{2\ell}} \left(1 + \frac{kN}{2} \right) + \frac{kN}{8} + \frac{1}{4} \left(1 - \frac{kN}{2} \right)^2 \right). \end{aligned}$$

Proof. Following the analysis of Appendix B in [CLOT21] similarly with a partial Assumption 3.11 (Independence heuristic) in [CGGI20], we may claim the term appearing in Step (1) of Appendix B has at most quadratic growth over the iteration of external products, and the term appearing in Step (2) has a linear growth because of their relative independence. \square

Lemma 6. *In the double blind rotation, t nested CMux gates by GGSW ciphertexts of noise variance at most σ_{dbr} after the common blind rotation increases the output noise variance by*

$$\begin{aligned} V_{\text{DBR},t} &\leq t^2 \ell_{\text{rk}} (k+1) N \frac{B_{\text{rk}}^2 + 2}{12} \sigma_{\text{dbr}}^2 \\ &\quad + t \left(\frac{q^2 - B_{\text{rk}}^{2\ell_{\text{rk}}}}{12B_{\text{rk}}^{2\ell_{\text{rk}}}} \left(1 + \frac{kN}{2} \right) + \frac{kN}{8} + \frac{1}{4} \left(1 - \frac{kN}{2} \right)^2 \right). \end{aligned}$$

where B_{rk} and ℓ_{rk} are the base and level of the GGSW ciphertext, respectively.

Proof. The noise increment by the CMux gate can be estimated by the result of Appendix B in [CLOT21] along with Lemma 5. The above upper bound is for the worst-case such that the distribution of the plaintext of the GGSW ciphertext is unknown over \mathbb{B} , considering the dependency of the round key bits in FRAST obtained from its master key. \square

Using the above lemmas, one can upper bound the noise variance of the S-box outputs obtained by the double blind rotation as described in Section E.1.1. After decomposing the S-box S into $S_{\text{msb-odd}}$ and $S_{\text{msb-even}}$, one computes the common blind rotation for $S_{\text{msb-odd}}$ and $S_{\text{msb-even}}$ on the input $x_1 + rk_2$ and $\text{ClearMSB}(x_1 + rk_2)$, respectively. Then, 4 (resp. 3) CMux gates follow using GGSW ciphertexts of each bit of $rk_j - rk_2$ (resp. $\text{ClearMSB}(rk_j - rk_2)$) for $j = 3, \dots, \ell$. From Lemma 6, the noise variance of the output ciphertext of $S_{\text{msb-odd}}(x_1 + rk_j)$ (resp. $S_{\text{msb-even}}(\text{ClearMSB}(x_1 + rk_j))$) by the double blind rotation is given as $\text{Var}(\text{PBS})$ increased by $V_{\text{DBR},4}$ (reps. $V_{\text{DBR},3}$).

In the evaluation of the FRAST round function, the noisiest ciphertext is one for the input to the S-box in the contracting part of the fixed round function, which is named crfsum in Section E.1.1. Considering the output noise variance of the resulting ciphertexts of the blind rotation, one can upper bound the noise increment of crfsum in each round by

$$2(\ell - 1)^2 \text{Var}(\text{PBS}) + (\ell - 2)(V_{\text{DBR},3} + V_{\text{DBR},4})$$

assuming that there is no correlation between the noises resulting from the common blind rotation and the GGSW ciphertexts of the round key bits, while there is clear dependence on the noises from the common blind rotation.

Considering additional noise increments such as the initial noise in refreshed crfsum , the noise in rk_1 and the keyswitching noise before the GenPBS operation,²¹ all of which

²¹It can be computed by Theorem 2 in [CLOT21].

do not affect the failure probability significantly, we obtain noise of standard deviation $2^{54.17}$ for `crfsum`, small enough compared to the scaling factor $\Delta = 2^{60}$. We observed the upper bound of noise coincides empirically with our result as Figure 10 shows the result of noise measurement to `crfsum` in 1000 evaluations of FRAST for each round. The failure probability of the GenPBS operation on the input of `crfsum` computed by Theorem 3 in [CLOT21] is negligible, and even PBSmanyLUT with $\vartheta = 1$, i.e., computing two functions on the same input in a single GenPBS call, can be used with failure probability less than $2^{-85.64}$. As FRAST has 40 rounds, the failure probability of FRAST evaluation is less than $2^{-80.32}$.

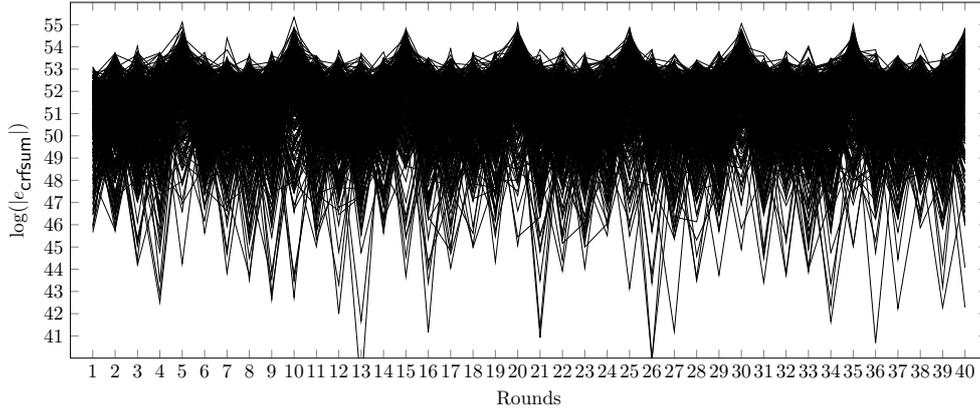


Figure 10: The magnitude of error of `crfsum` in each round of FRAST evaluation. The experiment is performed 1000 times.

F Communication Overload

In this section, we describe the communication overload of FRAST in the transciphering framework. The communication overload of the transciphering with TFHE consists of two parts: one for the homomorphic ciphertexts of the secret key, and the other for the TFHE evaluation keys only used for the transciphering. Since FRAST uses GLWEtoGGSW conversion in the setup phase, the evaluation keys for the conversion become additional communication overload.

CIPHERTEXT SIZE. An LWE ciphertext $(a_1, \dots, a_n, b) \in \mathbb{Z}_q^{n+1}$ consists of $n + 1$ elements in \mathbb{Z}_q , so its size is given by $(n + 1) \log q$ bits. If the LWE ciphertext is a fresh ciphertext such that no homomorphic operation is performed on it yet, the one can compress the random mask a into a seed for generating it. Such LWE ciphertexts are called seeded LWE ciphertexts. Ignoring the seed size by assuming that one seed generates all the random masks for multiple seeded ciphertexts, the size of the seeded LWE ciphertext is only $\log q$.²² In case of a GLWE ciphertext $(A_1, \dots, A_k, B) \in \mathcal{R}_{q,N}^{k+1}$, it is size of $(k + 1)N \log q$ bits. When it is compressed similarly, the seeded GLWE ciphertext is of size $N \log q$ bits. For a GGSW ciphertext $\mathbf{C} \in \mathcal{R}_{q,N}^{\ell(k+1) \times (k+1)}$, it can be considered as a vector of $\ell(k + 1)$ GLWE ciphertexts. The size of a GGSW ciphertext is $\ell(k + 1)^2 N \log q$ bits, and that of a compressed GGSW ciphertext is $\ell(k + 1)N \log q$ bits. Table 6 summarizes the size of each type of TFHE ciphertexts.

²²In the `tfhe-rs` library, auxiliary information such as the LWE dimension or ciphertext modulus type is saved together. We ignore such additional data size assuming that it is fixed in the transciphering framework.

Table 6: Size of TFHE ciphertexts in bits. The size of seeds or auxiliary information is ignored.

	LWE	GLWE	GGSW
Normal	$(n+1)\log q$	$(k+1)N\log q$	$\ell(k+1)^2N\log q$
Seeded	$\log q$	$N\log q$	$\ell(k+1)N\log q$

Table 7: Size of the GenPBS evaluation keys in bits. The size of seeds or auxiliary information is omitted.

	Bootstrapping Key	Keyswitching Key
Normal	$\ell_{\text{PBS}}(k+1)^2nN\log q$	$\ell_{\text{KS}}k(n+1)N\log q$
Seeded	$\ell_{\text{PBS}}(k+1)nN\log q$	$\ell_{\text{KS}}kN\log q$

GENPBS KEYSIZE. The evaluation keys for the GenPBS operation consist of the bootstrapping key and the keyswitching key. Given an LWE secret key $\mathbf{s} = (s_1, \dots, s_n) \in \mathbb{B}^n$ and a GLWE secret key $\mathbf{S}' = (S'_1, \dots, S'_k) \in \mathbb{B}_N[X]^k$, the bootstrapping key is a set of GGSW ciphertexts $\{\text{GGSW}_{\mathbf{S}'}(s_i)\}_{i=1}^n$ with the decomposition base B_{PBS} and level ℓ_{PBS} . Since the GGSW ciphertexts are fresh, the bootstrapping key can be compressed into seeded GGSW ciphertexts, resulting in the size of $\ell_{\text{PBS}}(k+1)nN\log q$ bits. Let $\mathbf{s}' = (s'_1, \dots, s'_{kN}) \in \mathbb{B}^{kN}$ be the LWE secret key induced from \mathbf{S}' . The keyswitching key is a set of LWE ciphertexts $\{\text{LWE}_{\mathbf{s}}(s'_i \cdot q/B_{\text{KS}}^j)\}_{(i,j) \in [kN] \times [\ell_{\text{KS}}]}$ where B_{KS} and ℓ_{KS} are keyswitching decomposition base and level, respectively. As a set of seeded LWE ciphertexts, the keyswitching key is of size $\ell_{\text{KS}}kN\log q$ bits. Table 7 summarizes the size of the evaluation keys for the GenPBS operation.

GLWETO GGSW KEYSIZE. For the GLWetoGGSW conversion proposed in [CCR19], two types of evaluation keys are needed; one is the GLWetoGLWE keyswitching keys, and the other is the GGSW ciphertext $\{\text{GGSW}_{\mathbf{S}}(-S_i)\}_{i=1}^k$ where $\mathbf{S} = (S_1, \dots, S_k)$. The GLWetoGGSW conversion rotates the coefficients in a GLWE ciphertext by switching a GLWE ciphertext under the secret key $S(X^m)$ to the GLWE ciphertext of the same plaintext under a secret key $S(X)$ for $m = N/2^{i-1} + 1$ where $i = 1, \dots, \log N$. A GLWetoGLWE keyswitching key that switches a GLWE ciphertext under $\mathbf{S} = (S_1, \dots, S_k)$ to the GLWE ciphertext under $\mathbf{S}' = (S'_1, \dots, S'_k)$ is a set of GLWE ciphertexts $\{\text{GLWE}_{\mathbf{S}'}(S_i \cdot q/B_{\text{subS}}^j)\}_{(i,j) \in [k] \times [\ell_{\text{subS}}]}$ where B_{subS} and ℓ_{subS} are the decomposition base and level of the GLWetoGLWE keyswitching in the GLWetoGGSW conversion.²³ Since the GLWetoGGSW conversion requires $\log N$ keyswitching keys, the size for them is $\ell_{\text{subS}}kN\log N\log q$ bits as seeded GLWE ciphertexts. The seeded GGSW ciphertexts $\{\text{GGSW}_{\mathbf{S}}(-S_i)\}_{i=1}^k$ with the decomposition base B_{SK} and level ℓ_{SK} are of size $\ell_{\text{SK}}k(k+1)N\log q$ bits. Table 8 summarizes the size of the evaluation keys for the GLWetoGGSW conversion.

COMMUNICATION OVERLOAD FOR FRAST. Transciphering with FRAST requires two bootstrapping keys: one of the default parameters of size 23.19 MB for FRAST keystream evaluation, and the other of the Kreyvium parameters of size 10.72 MB for the bit extraction

²³In general, \mathbf{S}' may have different GLWE dimension.

Table 8: Size of the GLWetoGGSW evaluation keys in bits. The size of seeds or auxiliary information is omitted.

	GLWetoGLWE Keyswitching Key	$\{\text{GGSW}_{\mathbf{S}}(-S_i)\}_{i=1}^k$
Normal	$\ell_{\text{subS}}k(k+1)N\log N\log q$	$\ell_{\text{SK}}k(k+1)^2N\log q$
Seeded	$\ell_{\text{subS}}kN\log N\log q$	$\ell_{\text{SK}}k(k+1)N\log q$

and the online phase.²⁴ The bootstrapping key of the default parameters is not taken into account for the communication overload since it is used in the actual usecase after the transciphering. To use double blind rotation technique, FRAST requires additional evaluation keys: the GLWEtoGLWE keyswitching keys of 880 KB, and the GGSW ciphertexts $\{\text{GGSW}_S(-S_i)\}_{i=1}^{k+1}$ of 160 KB. The round keys of 4808 bits are packed in $\lceil 4808/N \rceil \cdot \ell_{rk}$ GLWE ciphertexts of 144 KB, and the remaining round keys of 312 bits are sent to the server in 78 LWE ciphertexts of size 624 B (see Supplementary Material E). Hence, the total communication overload for FRAST is 11.88 MB.

G Keyschedule Matrix of FRAST

The following is the invertible 64×64 matrix \mathbf{M} over \mathbb{Z}_{16} to generate the round keys for FRAST in hex. The j -th hex digit in the i -th row denotes $\mathbf{M}_{i,j}$.

[

```

123e2d4f0a68befe1b7d2a8a858a011cc4f75074a6112874974ea0ed3f003e97,
78c9a6ca9af337f6c16d9b8d9651790a672aff5a4ddf5ef53801e55d71111a3a,
688b7691e958404437ef1190b84aae2a289ce0aef6ddde8899e9db73770bfb0,
b9dfd3ad6350b7b68b59249681939dea07d3e1897052989f651d8f902cbd2e6d,
d88befa672f89f48ae8827afb8c57350d05d44f09f858ffb309c33e4fcbc2b7c,
3390d1a59bbc7345f617d5d1d242b7ca5d59fb237c28e91a4d1d9e03aa526ac,
cf01a1143779cdca5a2d909144dad777a5f1bdf10d0b4355b46950aea6810c28,
54db8d16b878595f09aaa5569293f1e16453923dd9d639df690d545cb34eb287,
e33d69fe80adfd18b16f0885ce5a83e2cf08a35d785c22755a2aee03c5364d19,
7733b2708ef22719a886ef7e83909fbff6ad122e08d876646f85ca5e737a9784,
6b56dda75fa664d45b8eed98bfc10aa0a73c756b2826c88a5734beb6d458a635,
7c764e4ee7ea41d2ab2e52e65dae1258633fab9f9824eeda7461bd1e5d9d5dfa,
274929e465b09f096a890853505f4422e732d34cc6ad11df6b9a318834567981,
43f45459c7eb531181f9069772106b4700a14264378d8a1188c519444dc412b4,
56030bb64d3e51e69a52d0e6d835fb6fde71190fe0990c50506269b4fb53c50e,
8ecece73caf1f42a311ed490388e62f5f51b878d1ad2c3cf49091b97e771428d,
f798b7780d3528223e64577a4e93067bd70b94c563caa1b02843942c7eab2e0d,
88208a47bd7348fd3711dc63fa4b51c3c78a124f147f13aeca38dcaec6434c6,
4bd72876ca5df876a90965429c591904bc46b2472ca6a8a5a8a71611fc3a38,
81145d688175a61195940fa86398e5b18f209595179454709aeade1e0206731a,
5056684dc7d6598511e9a67f5f131f11e096754fc7fd992f0c9428121c2cab55,
8045cd6a51f879e0ad1fac41eb4935f5393d64c4170b413de21a6862a162d5c8,
81c639033f98b85f1b0da21c235971e1329c7bbca555d3d948fb8ec788f2ea92,
d900297fad1be4ec5c1b941e486ad8d2378ab5150c31444665eda786407fafbe,
80deccafb5db083163d050859df830017953c9e39bb3ac1f0af14d519538d16f,
2ddb2af637f17c67ff5ecee151edde07a3cccfce1ce610ef843d6f131502812,
13bb3c0c389e5d1b64e51becd6b1a7ec80ba13eac4d90a81c87a316425249b7f,
bd2de0f717cf1d713430da20fc4b425ec482a79ef9b360f9e55b71dbd44ad6f4,
d16221d21e6889f3d4666f735aa2ca9bf02adee153cd79d64cd307015eab836a,
31d26500d305b0ec48a681aba0de1a51b4433c2d50f19e74dff72d2e4147a655,
64f97ff5ddf77ecc9af0b9a169158fba74928ad5ad8262dd300741d8ae686551,
5e94a56cc8643dec58df4605bcbe2e15564bd2b59735efbab6ec7b0955e3c24a,
6c5bcd35dd82eba1421f985324092acb190739956c86399f1bdcf1d060a8f583,
32b43d7d456a82b38cb09b1af9b1a927934b8abb1120e48c1823bb3ce41bc3e6,
67ef894f2d1c96525a8313fcae3ea323f7676a02d5c77276a62a9084a35171f8,

```

²⁴There are keyswitching keys used inside the bootstrapping keys and between the bootstrapping keys, but we ignore them since their sizes are of several KBs.

6d91fb6ec05b7ebf023cc5f194c89252a1a7c57a23fed84801c9aa8730daf1d3,
c22ea03a5e77a62e7a938d43b3aa3a2afb681fe7bbb585a3d0a13f4abfeb7235,
8d1b6b34141a0681f984a3c1a531a3f89b90ca8f9500ee34c6ba15e3ae102a5a,
68d0a550478363bc5f184798bd9eda8bead7fa2da1c8a05190619676138a9220,
cc5d2c6e0fc017d51dd55f15487505e6c8e49c8d5254552bf99e9f027a157d0d,
e46d369567e85730d69850b5a66381e8b351683333b7506ad2a7255b64d73d62,
28c59e68559ee74a28b004315653ee7b60201141119be7f6c9c5db6a9426aa84,
f8a619ee67f02c28240731427c7a73510de95c8d14bd535615124f7f2ea28536,
08cc2655daa1166a3e1e27e4d32e1964f4668906105704f16674e05f454674ac,
824b77fb8ebac0cfd6e501ebdd8399ffc6d3fdfaa252c546f07b2fcb3211c44b,
338ee3c9b964f458368b12fb5dda0a21d71b82895f344c84d79a209d2a8a2e9e,
11f895bcaa829a35766a9df798f76fdb75f52acf844b38e29843b3c4cba641e0,
76d3864ea625fa5162e954ee2e21de7d2b31f5c2a848b055f9a8d2ce3c4d602a,
8d2f7999216bb4ccb0ab4d2f3ea49216e3806f68911777716c5a3f57f39ad5ab,
932be23c972a02d874c896c8bdd9f9e8f25c7597255c785febb36f39f21e6047,
2af439cffb90a7e3cd449b98b4e41986c62feeeac893e948f6d39bccbaea882c,
b2ce76fa91b65528128fc2f75cb210d627de1a60405a99fe68049adf26ce6290,
90199dbacdbc4cd27d89078cb88488c953e8aafe40f9fdf145ac4126807ce2bb,
1e9edf31bc7674add23dc865f2b7e459fce2855b725168b709246198b67c28bc,
365c951b8f9157c59975fe4e2b2fe6d20559e8945cc951f2d66fa60189181705,
6ae54a50a73127b6e384b823cd5f7274c499db4357aa984134fcb8d7a2263a1c,
af24aab01b444a3a8d1959a200051f61b154411a9318d26f3c9fc9089dc5a15e,
a7c8d71ff32d30c5da29787bc37f7faf3ec301db0ee66623d73a44cf0a5221fc,
8db8ecc3b85043eb2c092b39393a501b286fe5133f2f13182b262c3fd5fb0ffe,
824d35b5f8343f86fcb4c850113c7b1c57c500cbe3c7d3dd57c2fffeaa472d73,
255bf17ef996b37e6a2031faeb3ebb5171db7b3a9edc299b1fc7b068acf41e67,
123e3ef3ce478c12792754b8fc94ba37dc7331415b0e1ad8fef70b4002500580,
fde6a170a8ba8349a6f366f26e5bdf1a4979c94ca8800c4875d3ef6e3314fac3,
cecab9b85ef7b9fbe7792d72c43b144ccc1cabcb24a3403c38e00d8763610b6bb,

]