



Practical Homomorphic Evaluation of Block-Cipher-Based Hash Functions with Applications

Adda Akram Bendoukha^{1(✉)}, Oana Stan¹, Renaud Sirdey¹, Nicolas Quero^{1,2},
and Luciano Freitas³

¹ Université Paris-Saclay, CEA-List, 91120 Palaiseau, France
{adda.bendoukha,oana.stan,renaud.sirdey}@cea.fr

² Expleo, Saint-Quentin-en-Yvelines, Montigny-le-Bretonneux, France
nicolas.quero@expleogroup.com

³ LTCI, Télécom Paris, Institut Polytechnique de Paris, Palaiseau, France
lfreitas@telecom-paris.fr

Abstract. Fully homomorphic encryption (FHE) is a powerful cryptographic technique allowing to perform computation directly over encrypted data. Motivated by the overhead induced by the homomorphic ciphertexts during encryption and transmission, the transcribing technique, consisting in switching from a symmetric encryption to FHE encrypted data was investigated in several papers. Different stream and block ciphers were evaluated in terms of their “FHE-friendliness”, meaning practical implementations costs while maintaining sufficient security levels. In this work, we present a first evaluation of hash functions in the homomorphic domain, based on well-chosen block ciphers. More precisely, we investigate the cost of transforming PRINCE, SIMON, SPECK, and LowMC, a set of lightweight block-ciphers into secure hash primitives using well-established hash functions constructions based on block-ciphers, and provide evaluation under bootstrappable FHE schemes. We also motivate the necessity of practical homomorphic evaluation of hash functions by providing several use cases in which the integrity of private data is also required. In particular, our hash constructions can be of significant use in a threshold-homomorphic based protocol for the single secret leader election problem occurring in blockchains with Proof-of-stake consensus. Our experiments showed that using a TFHE implementation of a hash function, we are able to achieve practical runtime, and appropriate security levels (e.g., for PRINCE it takes 1.28 minutes to obtain a 128 bits of hash).

Keywords: FHE · Hash functions

1 Introduction

Fully homomorphic encryption (FHE) allows in theory to compute any function over an encrypted input. A plethora of works [5, 16, 20, 26] investigated the

N. Quero—This author contribution to this work was done while at CEA LIST.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2023
G.-V. Jourdan et al. (Eds.): FPS 2022, LNCS 13877, pp. 88–103, 2023.

https://doi.org/10.1007/978-3-031-30122-3_6

evaluation of symmetric cryptographic primitives over FHE encrypted keys. The interest in this topic is mainly due to the advent of proxy-re-encryption or transciphering [12], which is a technique that partially solves transmission of massive FHE ciphertexts through limited bandwidth networks, by having the receiver computing an homomorphic decryption of a symmetric cryptosystem. Therefore, many stream and block-ciphers were designed to be efficiently evaluated using an FHE encryption of their key. All the above methods were designed mainly to protect data confidentiality, either through symmetric encryption (for the encryption step and the transmission), or through homomorphic encryption for their processing by an honest-but-curious entity. We argue that there are applications of FHE in which it is useful not only to have confidentiality guarantees but also an integrity check over homomorphically encrypted data. More precisely, in this work we discuss the evaluation of hash functions over an FHE encrypted message and provide several scenarios in which this application can be a solution to achieve integrity along with data privacy. Let us now present the major contributions of our paper.

1.1 Contribution and Motivation

In this paper, we present a set of FHE-friendly hash functions built on lightweight block-ciphers using provably-secure constructions, and with reasonable homomorphic execution times. Our choice for a block-cipher-based construction is well motivated and it is the result of investigating several other options, including the homomorphic execution of lightweight hash functions as well as the building of hash functions from FHE-friendly stream-ciphers. As discussed more in details in Sect. 1.2, the preliminary analysis of several lightweight hash functions candidate to the NIST competition on lightweight cryptography showed that they are not well suited for homomorphic execution. As for the second option, to the best of our knowledge, there is no known practical method to design a secure hash function directly from a stream-cipher (although universal constructions do exist based on Luby-Rackoff theory [31]). As such, we present here some hash function constructions from “FHE-friendly” block-ciphers such as PRINCE [11], LowMC [1] and SIMON [3]. These block-ciphers are interesting candidates to build hash functions, from the homomorphic evaluation point of view, since they have an appropriate design, and have already been implemented with second-generation homomorphic schemes in the context of transciphering.

First, we derive several constructions of hash functions from PRINCE by means of the double block length hash construction, which enables a 128-bits hash size taking into account that the original block size in the PRINCE design is only of 64 bits. We then look into more details and evaluate the performances of a TFHE [15] gate-bootstrapping implementation of these hash functions. Additionally, we leverage on SIMON and LowMC (in their 128-bits block size flavors) to obtain hash sizes of 256 bits via the same construction.

Finally, we describe several use-cases which may require to run our hash functions in the encrypted domain, including integrity checking of homomorphically

encrypted data, oblivious authentication, homomorphic database querying, and a FHE-based protocol for single secret leader election.

1.2 Why Block Cipher-Based Constructions?

Beside security considerations, when constructing our hash functions, another criteria we looked at was to have a relatively fast evaluation in the homomorphic domain (e.g. less than one minute for a 256-size digest). A first idea for the construction of secure hashes suitable for homomorphic evaluation was to investigate three of the NIST lightweight competition finalists [17]: SPARKLE [4], XOODYAK [18] and Photon-Beetle [38]. We analysed them in function of the type of homomorphic bitwise operations one should execute: “free” operations such as permutations and concatenations, relatively easy operations such as the XOR and the AND (recall we use mainly TFHE in this work), and more difficult operations such as the modulo. We found out that their underlying primitives (e.g. S-boxes, modulo) and the number of rounds they require makes their homomorphic evaluation too expensive even with a bootstrapping-based homomorphic scheme, like TFHE. We also analyzed SPONGENT [8], another lightweight hash function, imposing to execute ≈ 30000 S-box in homomorphic domain (which corresponds to 68 S-boxes per round, 140 required rounds and 32 absorbing and squeezing steps) for 256-bits of output. Taking into account that the execution of the S-box used takes ≈ 0.6 s under TFHE, it follows that an homomorphic implementation of the SPONGENT hash function would be too slow to be of practical interest. Further details are available in [36].

Another appealing path was to explore hash-based constructions inspired from “FHE-friendly” stream ciphers. This option was tempting since nowadays there are several practical solutions implementing stream-ciphers into homomorphic domain (e.g. Kreyvium [12], Grain128 [5], PASTA [20]). However, even if it seems possible to obtain hash functions with very interesting homomorphic performances, their security seems difficult to assess and this, thus, remains an interesting open question. In essence, although theoretical constructions do exist, the symmetric cryptography community has, to the best of our knowledge, only marginally followed this path for building hash functions. Still, the possibility of achieving better FHE evaluation performances may be a new motivation for further investigations along this line.

As a consequence, we decided to consider block-cipher algorithms which have been already considered for homomorphic evaluations and turn them into secure hash functions using generic methods such the ones described in Sect. 2.3.

2 Background

2.1 Transciphering

Transciphering is a technique that allows offloading massive data from client to server with the aim to perform server-side homomorphic computations. Indeed,

when a message m is encrypted under an FHE cryptosystem, the resulting size of the ciphertext $\text{FHE.Enc}_{\text{FHE.pk}}(m)$ is much larger than the size of the original message m , by an expansion factor which depends polynomially on the security parameter λ . In all modern FHE schemes for a λ large enough (in the 110–130 bits of security ranges) ciphertext sizes reach several kbytes or even megabytes (depending on the chosen cryptosystem and its security level). So, instead of encrypting m directly using an FHE scheme and sending $\text{FHE.Enc}_{\text{FHE.pk}}(m)$, a client will rather encrypt m using a symmetric cryptosystem and sends the encryption $\text{SYM.Enc}_{\text{SYM.sk}}(m)$ to the server along with $\text{FHE.Enc}_{\text{FHE.pk}}(\text{SYM.sk})$, the FHE encryption of the symmetric key SYM.sk . The server then homomorphically runs $\text{SYM.Dec}_{\text{FHE.Enc}(\text{SYM.sk})}(\text{SYM.Enc}(m))$ and recovers the message encrypted under the homomorphic public key $\text{FHE.Enc}_{\text{FHE.pk}}(m)$.

$\text{SYM.Enc}_{\text{SYM.sk}}(m)$ is roughly of the same size as m while SYM.sk , which is the only FHE encrypted and transmitted element, is of fixed size and often small enough to be homomorphically encrypted and sent (once and offline) through the network, whilst m can be arbitrarily large. Switching from a symmetric scheme to an FHE one allows a form of secure *compression* of the homomorphic ciphertexts. It requires however, the evaluation of SYM.Enc homomorphically, which introduces a non-negligible additional computational overhead on the server-side. In [5, 12], it is argued that the use of a stream-cipher is more suitable for transcribing in the case of both 2nd generations FHE schemes (e.g., BGV, BFV) as well as TFHE. In [20] authors discuss the semantic security of transcribing seen as Key encapsulation/Data encapsulation mechanism (KEM-DEM) depending on the semantic security of both the symmetric and homomorphic schemes involved, and provide also an FHE-friendly stream-cipher named Pasta, suited for levelled FHE schemes.

2.2 Hash Functions and Security Properties

A general definition of a hash function is a mapping of messages of arbitrary length to a fixed size digest. Additionally, a *cryptographic* hash function requires the following security properties.

Pre-image Resistance. Given $h \in \{0, 1\}^n$ the output of the hash function $H : \{0, 1\}^* \rightarrow \{0, 1\}^n$, it must be computationally hard to find $m \in \{0, 1\}^*$ such that $H(m) = h$.

Collision Resistance. It must be computationally hard to find two distinct messages m_1 and m_2 such that $H(m_1) = H(m_2)$.

Second Pre-image Resistance. Given m and h such that $H(m) = h$, it must be computationally hard to find $m' \neq m$ and $H(m) = H(m')$.

Since we only consider cryptographic hash functions, for simplicity sake, in the remaining of the paper we will refer to a “cryptographic hash function” as “a hash function”.

Black-Box Model. To prove the security of a block-cipher-based hash function independently of the underlying cipher's structure, it is used the black-box model, in which a block-cipher is modeled as an invertible random permutation defined by the key. An adversary is given access to encryption and decryption oracles, such that given m (resp. c) the encryption $E_k(m)$ (resp. the decryption $E_k^{-1}(c)$) is returned. The complexity of an attack is measured by the number of encryption and decryption queries that an optimal adversary performs. Since most attacks on block-cipher-based hash functions do not take advantage of the block-cipher's potential structural weaknesses or flaws, it is relevant to use a black-box model for security analysis.

2.3 Block-Cipher-Based Hash Functions

Among the most widely used constructions of hash functions are the iterated hash functions, in which a round function, also referred to as a compression function $F : \{0, 1\}^n \cdot \{0, 1\}^l \rightarrow \{0, 1\}^n$ is iterated over every message block, taking as input the current message block of size n and the previous hash value¹. The output of the final compression function call is the hash of the input message as shown in Algorithm 1. Due to its simplicity, this construction has been intensively studied in the state of the art [6, 30], giving birth to many hashing standards such as SHA-0, SHA-1, and SHA-2. A large part of the security of these hash functions can be attributed to the underlying compression function². In [19] authors demonstrate that the collision resistance of F implies collision resistance of the hash function built from F using the Merkle-Damgård construction.

These results raised interests in building secure compression functions from which it will be easy to build secure hash functions. A block-cipher is a primitive that already provides security properties by construction. Although the security requirements of an encryption algorithm are different by nature from those of a hash function, the question of how to build a secure compression function from a block-cipher quickly appeared and was intensively investigated, laying foundation for instance for the MDC family of hash functions [34] based on the block-cipher DES. The main motivation of this approach is to minimize design efforts, and use existing primitives. The task is to transform the security properties of a block-cipher into those of a cryptographic hash function, by carefully executing it over well-chosen linear combinations of the current message block, the chaining variable, or other conventional constants, taken as encryption keys or message blocks. This gave birth to a plethora of constructions, some of them were proven secure in the black-box model, others exhibited weaknesses regardless of the underlying block-cipher's potential weaknesses.

One important security element is the size of the digest. Due to the birthday paradox, collision security level of a hash function is upper-bounded by $O(2^{n/2})$, where n is the size of the hash. Thus, having a size for the hash equal to the size of the block for the cipher used to construct the compression function raised some

¹ A chaining value to provide dependency between successive hash values.

² The security under all aspects : Pre-image, second pre-image, and collision resistance.

issues. The size of some block-cipher's blocks can be too small to be considered as a secure hash size, and using a block-cipher with a large block length often results in higher execution times. Providing a secure construction that produces a hash twice larger than the block-cipher's block length was subject to several research efforts.

Algorithm 1. Merkle Damgård iterated hash function

```

input :  $m = (m_0, m_1, \dots, m_l)$ 
 $h_0$  is set to an initialization vector
for  $i = 0$  to  $l$  do
     $h_i = F(h_{i-1}, m_i)$ 
end for
return  $h_l$ 

```

Single Block Length (SBL) Hash Functions. One of the very first constructions of single-block-length hash functions is the Davies-Meyer construction where $H_{i+1} = E_{M_i}(H_i) \oplus M_i$ and the Mugiuchi-Prenel's scheme with $H_{i+1} = E_{M_i}(H_i) \oplus M_i \oplus H_i$, where H_i is the previous hash value and each block of the message (M_i) is the key to a block cipher E .

Later, in [35], Prenel, Govaerts and Vandewalle (PGV) provided an exhaustive analysis of iterated hash functions defined over $\{0, 1\}^* \rightarrow \{0, 1\}^n$ and based on a block-cipher. The compression function is in the form $F(a, b) = E_a(b) \oplus c$ where a, b and c are in $\{m_i, h_{i-1}, IV, m_i \oplus h_{i-1}\}$, and E is $\{0, 1\}^n \cdot \{0, 1\}^n \rightarrow \{0, 1\}^n$ block-cipher. There are $4^3 = 64$ such compression functions, among which 12 are presented as secure. Afterwards, Black, Rogaway and Shrimpton [7] provided formal security proofs in the black-box model of the 12 constructions analysed in [35]. They also demonstrated that among the remaining 52 constructions, 8 of them were actually secure with respect to collision and pre-image resistance. In this work we chose to evaluate Davies-Meyer's hash function under several block-ciphers, as it provides optimal security in the black-box model, and is equivalent in terms of computation complexity to other secure constructions from [7].

The security analysis and explicit constructions are provided in [7].

Double Block Length (DBL) Hash Functions. As mentioned before, constructions by PGV provide a hash of n -bits size when using a $\{0, 1\}^n \cdot \{0, 1\}^n \rightarrow \{0, 1\}^n$ underlying block-cipher in the compression function. Due to the birthday paradox, these hash functions require block-ciphers with a large enough block length in order to provide security against collision attacks.

A measure of the efficiency of a hash function is its rate, that is, the inverse of the number of calls to the compression function per iteration.

In [33] Merkle presents three optimally collision resistant double block length hash functions, based on the block-cipher DES. However, their rates are low compared to the next generation of DBL constructions.

Lai and Massey proposed TANDEM-DM [28] for a rate 1/2 hash construction, using a $(n, 2n)$ block-cipher. It was proven optimally collision and pre-image secure in [21]. It makes however two non-independent calls³ per iteration making it non-parallel. Abreast-DM [29] is another construction with a rate of 1/2 making two parallel calls to the block-cipher, and was proven to have optimal collision resistance in [22].

Lucks in [5] provides a first DBL construction of rate 1. Making a single block-cipher call per iteration comes at the cost of computing a heavy linear combination of the message block and the previous hash resulting in a significant overhead. Hirose in [25] provides a rate 1/2 construction with two distinct $(n, 2n)$ block-ciphers, then uses a tweak in order to use a single block-cipher. This construction provides optimal bounds for both collision and pre-image resistance in the black-box model, and is parallel. Indeed, the two calls to the compression function (and thus, to the block-cipher) are independent, making its performance comparable to rate 1 constructions.

Other works from [27, 32] studied the possibility to build DBL hash functions from an (n, n) -block-cipher. MDC-2 fails to provide optimal security, while MDC-4 [32] is near optimal, but has a rate smaller than 1/2.

In this work, we homomorphically evaluate the constructions of Hirose and Tandem-DM. The goal is to provide an idea of the runtime of two optimally secure hash functions of rate 1/2 from both the parallel and non-parallel types on top of an FHE encryption layer.

3 Applications of Homomorphic Hash Functions

3.1 Homomorphic Data Integrity Check

As described in Sect. 2.1, transciphering allows to transfer symmetrically encrypted data instead of homomorphically encrypted and thus reduces the required bandwidth. However, transciphering while preserving data privacy does not ensure data integrity during transmission. In [5] authors describe how to include data integrity check within transciphering, but their approach required an AEAD encryption scheme (Authenticated Encryption with Associated Data).

Indeed, all stream-ciphers suffer from malleability, i.e., the possibility for an adversary to create an encryption of $m + k$ where k is some constant, from an encryption of m ⁴. A malleable encryption scheme can be subject to man-in-the-middle attacks. Some modern stream-ciphers (e.g. [24]) come with the possibility to compute a MAC (Message Authentication Code) along with the encryption in an attempt to circumvent this issue. Another simple way to perform integrity check within transciphering when the chosen stream-cipher does not embed a MAC computation is to include a hash function. A client encrypts m concatenated to $H(m)$ using a symmetric encryption scheme. She then transmits these

³ The output of the first block-cipher call is used to build the key of the second block-cipher call.

⁴ $m \oplus \text{keystream} \oplus k = \text{SYM.Enc}(m \oplus k)$.

elements to the server along with $\text{FHE.Enc}(\text{SYM.sk})$ (once and for all). Once the server has finished transcribing both the message and the hash, it recovers $\text{FHE.Enc}(m')$ and $\text{FHE.Enc}(h')$, he computes $[h] = H(\text{FHE.Enc}(m'))$. If $m = m'$ then $h = h'$ (with overwhelming probability, of course). The server computes the homomorphically encrypted bit $[r]_{\text{FHE.pk}} = \prod_{i=0}^n (1 \oplus h_i \oplus h'_i)$ where n is the size of the hash. $[r]_{\text{FHE.pk}}$ is the output of the integrity check, such that :

$$[r] = \left[\begin{cases} \text{An encryption of 1} & \text{if } m = m' \\ \text{An encryption of 0} & \text{otherwise} \end{cases} \right], \quad (1)$$

This FHE encrypted bit could then be used in many ways. The server can simply transmit it to the client, in order to give him the ability to verify if his data was altered or corrupted during the transmission. Or the server could choose to reply with $[f(m)]$ or a *NIL* value outside of the range of f , according to the value of the bit r . In TFHE [15] for example, this can be realized using a homomorphic CMUX gate at roughly the cost of an extra homomorphic multiplication⁵.

3.2 Single Secret Leader Election (SSLE)

The problem of securely electing a single leader in a distributed system was formally defined by Boneh et al. in [9]. For a committee of peers which collaboratively elect a node to complete a task, the problem consists in electing a node in a way that only this elected peer is able to know that he was elected and the others learn only that they were not elected. Also, the elected peer must be able to provide a proof of his election when he decides to reveal himself once his task is done. In [23] a solution to the SSLE problem is proposed based on Threshold Fully Homomorphic encryption [10] for partially-synchronous systems. A very high level description is the following. Every peer P_i wishing to register to the election at a given height and cycle (low and high level steps in the leader election protocol), provides an FHE encryption of $p_i = H(h||t_i||c)$ called the proof, where h is the height of the blockchain, c the current cycle of elections, and t_i a locally generated number belonging to process P_i . Every participating peer performs a sampling circuit following a weighted distribution over the FHE encrypted list of proofs and ids of all registered peers, using collaboratively generated randomness from [37]. Then, each peer homomorphically selects⁶ a proof and the associated id from the set of all proofs. He then homomorphically hashes $(p_i||i)$ where i is the id of the elected peer, and p_i the corresponding proof. The next step is to broadcast a partial decryption of the voucher $v_{h,c,r} = H(p_i||i)$. Every honest peer samples the same p_i and i , and broadcasts his partial decryptions of $v_{h,c,r}$ using his secret key share. Assuming we have at least t honest peers in the system, where t is the decryption threshold, every peer must eventually receive enough partial decryptions and be able to perform a full decryption of $v_{h,c,r}$. The elected

⁵ $\text{CMUX}([r], [f(m)], \text{NIL}) = [r] \cdot [f(m)] + (1 - [r]) \cdot \text{NIL}$

⁶ I.e., homomorphically computes a one-hot encoding of an index in the proofs list and performs a dot-product to extract one such proof, thus without knowing which one.

peer recognizes his *voucher*, whereas other peers gain no information from plain $v_{h,c,r}$, nor can fake the election, since H is secure against pre-image and second pre-image attacks. Afterward, the leader is able to prove his election by submitting his plaintext proof $p_i = H(h||t_i||c)$. The verification is simply performed by running the test $H(p_i||i) == v_{h,c,r}$.

The homomorphically evaluated hash function plays a significant role in this protocol. It hides the sensitive elements from Byzantine peers providing the secrecy of the election and a simple proof mechanism, making the election easily verifiable, yet computationally hard to forge fake proofs.⁷

3.3 Homomorphic Database Querying

Suppose a server maintaining a database of elements DB such that query m has the answer A_m stored at index $H(m)$, where H is a hash function (with a small digest size w. r. t. to cryptographic standards). In this case H is not necessarily cryptographic. For instance pre-image resistance is not necessary since the query is already private under an FHE encryption layer. Nevertheless, we require from H to have balanced collisions⁸ and, for this sake, one can use Luby-Rackoff's universal hash functions from [31]. In this setting, the server is able to homomorphically answer FHE encrypted queries.

A client homomorphically encrypts a query x and sends $[x]_{\text{FHE.pk}}$ to the server. The server computes $[i]_{\text{FHE.pk}} = H([x]_{\text{FHE.pk}})$, which is an FHE encryption of the index of A_x inside his database. The server then computes a vector V which contains FHE encryptions of 0 everywhere except at index i in which an encryption of 1 is stored. V is computed as follows : $V[k] = ([i]_{\text{FHE.pk}} == k)$ with $k \in \llbracket 0, n-1 \rrbracket$. Lastly, to extract an FHE encryption of A_x , the server performs a homomorphic dot product between the vector V and his database of elements $\sum_{i=0}^n DB[i] \cdot V[i]$, and sends back to the client the result of this final dot product, which will be $[A_x]_{\text{FHE.pk}}$.

One remaining problem of this use-case is to homomorphically resolve collisions of H . A first approach is to have the server creating lists of answers to different queries which hash to the same index at the position $H(x)$ in DB , and provide a second hash function H' , whose output is smaller than the one of H , and which will compute the index of A_x inside the corresponding list. Thus, when an FHE encrypted query $[x]_{\text{FHE.pk}}$ is received, the position $(H([x]_{\text{FHE.pk}}), H'([x]_{\text{FHE.pk}}))$ provides an answer.

3.4 Oblivious Authenticated (Homomorphic) Calculations

It is well known that (keyed) hash functions are used in many authentication protocols where an entity (the user) can prove its knowledge of a secret (the key

⁷ Secrecy is granted by the pre-image resistance of the hash function. Having a single verifiable leader is due to the second pre-image resistance of the hash function.

⁸ $H : \{0, 1\}^* \rightarrow \{0, 1\}^n$ has balanced collisions if all elements in $\{0, 1\}^n$ have the same number of pre-images under H .

of the hash function) to another entity (the server). To do so, the server sends a random challenge to the user which replies with the hash of the challenge. Since the server can also perform the same calculation, it can check the correctness of the client replies which proves the latter knowledge of the secret. With the ability of running hash functions in the homomorphic domain, we can now provide the server with an FHE encryption of the secret key and have the server performing the authentication in the encrypted domain i.e., the server generates a challenge in the clear domain, sends it to the user and get its (non encrypted) reply. The server can then run the (keyed) hash function homomorphically on its challenge, and homomorphically compare the obtained (encrypted) result with the reply received from the user. As the end of this process, the server possesses an encrypted boolean, say β , indicating whether or not the client has successfully authenticated (but has by construction no knowledge of whether or not that authentication was successful).

One way of using this consists in providing a valid calculation only to successfully authenticated users. In essence, rather than computing $f(x)$ in the homomorphic domain, the server can now compute $\beta f(x) + (1 - \beta)\perp$ (where \perp denotes a constant value meaning, by convention, “not an answer”). As a consequence, (encrypted) valid calculation results are duly returned only to authenticated users, while other users receive only useless encryptions of \perp . This is then (nicely) done obliviously to the server, which cannot distinguish between ciphertexts of valid results and ciphertexts of \perp , and without revealing the secret hash function key (since it is only provided with an FHE-encryption of that secret key).

4 Adaptations of Block Ciphers for FHE-friendly Hashes

4.1 Block-Ciphers Considered

The Low-MC block-cipher [1] is part of a family of symmetric schemes designed for practical instantiations in homomorphic domain with the objectives of minimizing both the multiplicative complexity and the multiplicative depth, making it efficient for levelled homomorphic schemes. This design principle, had to be compensated with a large number of xor gates in order to ensure algebraic properties that provide an appropriate level of security. This latter fact makes it rather inefficient when ran under TFHE, since the cost of all Boolean homomorphic gates is the same within this FHE scheme (A bootstrapping operation is performed after every Boolean gate). It remains however a interesting candidate for hash constructions targeting efficient homomorphic evaluation in a levelled FHE setting. Even if the first variants of LowMC were successfully attacked, the subsequent proposed design is more secure and highly parametrizable. In particular, there is a closed-form formula to determine the minimal number of rounds to reach a given security target depending on the block size (128 or 256 bits), the key size, the number of S-boxes and the allowed data complexity.

PRINCE [11], SPECK and SIMON [3] are lightweight block-ciphers, with a relatively small block length. They were initially designed for constrained embedded execution environments. Their design approaches result in small gate counts⁹ which results in high performances when ran under TFHE. Due to its small block length, PRINCE is better suited with double block length constructions, resulting in a hash function which provides $O(2^{64})$ collision resistance, and $O(2^{128})$ for pre-image resistance. SPECK and SIMON can be instantiated in both the DBL and SBL settings since they both provide a double-key-size variants.

4.2 FHE Schemes Considered

We chose to run our experiments under the TFHE cryptosystem since it provides the possibility to evaluate (multiplicatively) unbounded homomorphic circuits thanks to its fast bootstrapping operation. This scheme is more suited for protocols where scalability is a requirement. For example, the secret single leader election protocol [9] described in Sect. 3.2 requires flexibility regarding the number of peers being able to disconnect or join the committee at different times. These variations in the number of peers linearly increase the multiplicative depth of the sampling circuit, which would be difficult to manage if a levelled homomorphic scheme were to be used¹⁰.

4.3 Tool: Cingulata Homomorphic Compiler

Cingulata, formerly known as Armadillo [14], is a toolchain and run-time environment (RTE) for implementing applications running over homomorphic encryption. Cingulata provides high-level abstractions and tools to facilitate the implementation and the execution of privacy-preserving applications.

Cingulata relies on instrumented C++ types to denote private variables, e.g., `CiInt` for integers and `CiBit` for Booleans. Integer variables are dynamically sized and are internally represented as arrays of `CiBit` objects. The Cingulata environment monitors/tracks each bit independently. Integer operations are performed using Boolean circuits, which are automatically generated by the toolchain. For example a full-adder circuit is employed to perform an integer addition. The Boolean circuit generation is configurable and two generators are available: focused on minimal circuit size or on small multiplicative-depth. More generally, it is possible to implement additional circuit generators or to combine them.

A `CiBit` object can be in either plain or encrypted state. Plain-plain and plain-encrypted bit operations are optimized out, in this way constant folding and propagation is automatically performed at the bit-level. Bit operations

⁹ A round of encryption of a block-cipher often includes a multiplication of the internal state with an $\mathbb{F}^{\mathbb{F}}$ matrix, this makes the number of operations quadratic with respect to its block size.

¹⁰ In this category of homomorphic schemes, the multiplicative depth of the homomorphic circuit to be evaluated has to be known in advance in order to generate a parameter set which allows homomorphic computations up to this depth.

between encrypted values are performed by a “bit execution” object implementing the `IBitExec` interface. This object can either be a HE library wrapper, simply a bit-tracker object or even a plaintext bit execution used for algorithm debugging purposes. When a HE library wrapper is used the Cingulata environment directly executes the application using the underlying HE library.

Another option is to use the bit-tracker in order to build a circuit representation of the application. This allows to use circuit optimization modules in order to further optimize the Boolean circuit representation. The hardware synthesis toolchain ABC¹¹ is used to minimize circuit size. It is an open-source environment providing implementations of state-of-the-art circuit optimization algorithms. These algorithms are mainly designed for minimizing circuit area or latency but, currently, none of them is designed for multiplicative depth minimization. In order to fill this gap, several heuristics for minimizing the multiplicative depth are available in Cingulata, refer to [2, 13] for more details.

The optimized Boolean circuit is then executed using Cingulata’s parallel run-time environment. The RTE is generic, meaning that it uses a HE library wrapper, i.e. a “bit execution” object as defined earlier, in order to execute the gates of the circuit. The scheduler of the run-time allows to fully take advantage of many-core processors. Besides, a set of utility applications are provided for parameter generation (given a target security level), key generation, encryption and decryption. These applications are also generic, in the same vein as the parallel RTE.

4.4 Experimental Results and Performances

We ran *multi core* performance tests on an Intel(R) Xeon(R) CPU E3-1240 v5 @ 3.50 GHz and 8 GB RAM using Cingulata in TFHE mode. We provide parallelism when possible using the OpenMP library.

For single block length construction, we implement Davies-Meyer’s compression function which requires a (n, n) -block-cipher. Therefore, we instantiate this construction with SPECK, SIMON¹², and the $(128, 128)$ variant of LowMC. In the double block length setting, since these constructions require an $(n, 2n)$ -block-cipher, we instantiate Hirose’s and Tandem-DM constructions with PRINCE, and the $(128, 256)$ variants of LowMC, and SIMON. The results are shown in Table 1 with the execution times in minutes when the hash functions are instantiated, and an “-” symbol when the construction is not compatible with the sizes of the key and the block of the cipher.

The obtained performances are as expected: lightweight ciphers provide better runtimes compared to LowMC. PRINCE is the most efficient cipher for DBL

¹¹ <http://people.eecs.berkeley.edu/alanmi/abc/>.

¹² For SIMON, these are estimations based on the gate count from [3] and the gate-bootstrapping time of TFHE.

constructions as it has the lowest gate-count, and is also the most parallelizable cipher. The number of rounds performed in every construction to produce the hash of a 128-bits message is $\lceil \frac{128}{blocklength} \rceil$. Thus, in the first row, DBL-PRINCE performs two iterations and produces a 128-bits hash. All the remaining constructions perform a single iteration.

Table 1. Evaluation of hash functions over a 128-bits TFHE encrypted message in minutes

Instantiation	Davies-Meyer (SBL)	Hirose (DBL)	Tandem-DM (DBL)
(64, 128)-PRINCE	–	1.28	2.98
(128, 128)-SPECK	3.78	–	–
(128, 256)-SPECK	–	4.91	8.16
(128, 128)-SIMON	2.14	–	–
(128, 256)-SIMON	–	3.64	7.05
(128, 128)-LowMC	6.12	–	–
(128, 256)-LowMC	–	8.58	17.32

5 Conclusion and Perspectives

In this work, we have investigated scenarios in which the ability to (efficiently) evaluate hash functions in the homomorphic domain is an interesting building block. To the best of our knowledge, this work is one of the first to address this issue, at least for the TFHE cryptosystem. We also explored various provably-secure constructions of “(T)FHE friendly” hash functions based on respected block-ciphers in order to achieve several digest sizes. Fully homomorphic encryption on its own opens perspectives towards a new set of applications. Then, combining it with the execution of hash functions in the homomorphic domain provides it with additional versatility which can serve in various scenarios and protocols.

References

1. Albrecht, M., Rechberger, C., Schneider, T., Tiessen, T., Zohner, M.: Ciphers for MPC and FHE. *Cryptology ePrint Archive, Paper 2016/687* (2016). <https://eprint.iacr.org/2016/687>
2. Aubry, P., Carpov, S., Sirdey, R.: Faster homomorphic encryption is not enough: improved heuristic for multiplicative depth minimization of Boolean circuits. In: Jarecki, S. (ed.) *CT-RSA 2020*. LNCS, vol. 12006, pp. 345–363. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-40186-3_15
3. Beaulieu, R., Shors, D., Smith, J., Treatman-Clark, S., Weeks, B., Wingers, L.: The SIMON and SPECK families of lightweight block ciphers. *Cryptology ePrint Archive, Paper 2013/404* (2013). <https://eprint.iacr.org/2013/404>
4. Beierle, C., et al.: Schwaemm and Esch: Lightweight authenticated encryption and hashing using the sparkle permutation family (2019)

5. Bendoukha, A.A., Boudguiga, A., Sirdey, R.: Revisiting stream-cipher-based homomorphic transciphering in the TFHE era. In: Aïmeur, E., Laurent, M., Yaich, R., Dupont, B., Garcia-Alfaro, J. (eds.) *Foundations and Practice of Security*. LNCS, vol. 13291, pp. 19–33. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-08147-7_2
6. Biham, E., Dunkelman, O.: A framework for iterative hash functions - HAIFA. *Cryptology ePrint Archive*, Paper 2007/278 (2007). <https://eprint.iacr.org/2007/278>
7. Black, J., Rogaway, P., Shrimpton, T.: Black-box analysis of the block-cipher-based hash-function constructions from PGV. In: Yung, M. (ed.) *CRYPTO 2002*. LNCS, vol. 2442, pp. 320–335. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-45708-9_21
8. Bogdanov, A., Knežević, M., Leander, G., Toz, D., Varıcı, K., Verbauwhede, I.: SPONGENT: a lightweight hash function. In: Preneel, B., Takagi, T. (eds.) *CHES 2011*. LNCS, vol. 6917, pp. 312–325. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-23951-9_21
9. Boneh, D., Eskandarian, S., Hanzlik, L., Greco, N.: Single secret leader election. *Cryptology ePrint Archive*, Paper 2020/025 (2020). <https://eprint.iacr.org/2020/025>
10. Boneh, D., et al.: Threshold cryptosystems from threshold fully homomorphic encryption. *Cryptology ePrint Archive*, Paper 2017/956 (2017). <https://eprint.iacr.org/2017/956>
11. Borghoff, J., et al.: Prince - a low-latency block cipher for pervasive computing applications (full version). *Cryptology ePrint Archive*, Paper 2012/529 (2012). <https://eprint.iacr.org/2012/529>
12. Canteaut, A., et al.: Stream ciphers: a practical solution for efficient homomorphic ciphertext compression. *J. Cryptol.* **31**(3), 885–916 (2018). <https://doi.org/10.1007/s00145-017-9273-9>
13. Carpov, S., Aubry, P., Sirdey, R.: A multi-start heuristic for multiplicative depth minimization of Boolean circuits. In: Brankovic, L., Ryan, J., Smyth, W.F. (eds.) *IWOCA 2017*. LNCS, vol. 10765, pp. 275–286. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-78825-8_23
14. Carpov, S., Dubrulle, P., Sirdey, R.: Armadillo: a compilation chain for privacy preserving applications. In: Bao, F., Miller, S., Chow, S.S.M., Yao, D. (eds.) *Proceedings of the 3rd International Workshop on Security in Cloud Computing, SCC@ASIACCS 2015*, Singapore, Republic of Singapore, 14 April 2015, pp. 13–19. ACM (2015). <https://doi.org/10.1145/2732516.2732520>
15. Chillotti, I., Gama, N., Georgieva, M., Izabachène, M.: TFHE: fast fully homomorphic encryption over the torus. *Cryptology ePrint Archive*, Paper 2018/421 (2018). <https://eprint.iacr.org/2018/421>
16. Cho, J., et al.: Transciphering framework for approximate homomorphic encryption (full version). *Cryptology ePrint Archive*, Paper 2020/1335 (2020). <https://eprint.iacr.org/2020/1335>
17. Cryptography: NIST lightweight cryptography. <https://csrc.nist.gov/Projects/Lightweight>
18. Daemen, J., Hoffert, S., Peeters, M., Assche, G.V., Keer, R.V.: Xoodyak and a lightweight cryptographic scheme (2020)
19. Damgård, Ivan Bjerre: A design principle for hash functions. In: Brassard, Gilles (ed.) *CRYPTO 1989*. LNCS, vol. 435, pp. 416–427. Springer, New York (1990). https://doi.org/10.1007/0-387-34805-0_39

20. Dobraunig, C., Grassi, L., Helminger, L., Rechberger, C., Schafneggler, M., Walch, R.: Pasta: a case for hybrid homomorphic encryption. Cryptology ePrint Archive, Paper 2021/731 (2021). <https://eprint.iacr.org/2021/731>
21. Fleischmann, E., Gorski, M., Lucks, S.: On the security of TANDEM-DM. In: Dunkelman, O. (ed.) FSE 2009. LNCS, vol. 5665, pp. 84–103. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-03317-9_6
22. Fleischmann, E., Gorski, M., Lucks, S.: Security of cyclic double block length hash functions. In: Parker, M.G. (ed.) IMACC 2009. LNCS, vol. 5921, pp. 153–175. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-10868-6_10
23. Freitas, L., et al.: Homomorphic sortition - secret leader election for blockchain (2022). <https://doi.org/10.48550/ARXIV.2206.11519>, <https://arxiv.org/abs/2206.11519>
24. Hell, M., Johansson, T., Meier, W., Sönnerup, J., Yoshida, H.: An AEAD variant of the grain stream cipher. In: Carlet, C., Guilley, S., Nitaj, A., Souidi, E.M. (eds.) C2SI 2019. LNCS, vol. 11445, pp. 55–71. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-16458-4_5
25. Hirose, S.: Provably secure double-block-length hash functions in a black-box model. In: Park, C., Chee, S. (eds.) ICISC 2004. LNCS, vol. 3506, pp. 330–342. Springer, Heidelberg (2005). https://doi.org/10.1007/11496618_24
26. Hoffmann, C., Méaux, P., Ricosset, T.: Transciphering, using FiLIP and TFHE for an efficient delegation of computation. In: Bhargavan, K., Oswald, E., Prabhakaran, M. (eds.) INDOCRYPT 2020. LNCS, vol. 12578, pp. 39–61. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-65277-7_3
27. Jetchev, D., Özen, O., Stam, M.: Collisions are not incidental: a compression function exploiting discrete geometry. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 303–320. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-28914-9_17
28. Lai, X., Massey, J.L.: Hash functions based on block ciphers. In: Rueppel, R.A. (ed.) EUROCRYPT 1992. LNCS, vol. 658, pp. 55–70. Springer, Heidelberg (1993). https://doi.org/10.1007/3-540-47555-9_5
29. Lee, J., Kwon, D.: The security of Abreast-DM in the ideal cipher model. Cryptology ePrint Archive, Paper 2009/225 (2009). <https://eprint.iacr.org/2009/225>
30. Lei, D., Lin, D., Chao, L., Feng, K., Qu, L.: The design principle of hash function with Merkle-Damgård construction. Cryptology ePrint Archive, Paper 2006/135 (2006). <https://eprint.iacr.org/2006/135>
31. Luby, M.: Pseudorandomness and Cryptographic Applications (1996). <https://doi.org/10.2307/j.ctvs32rpn>
32. Mennink, B.: On the collision and preimage security of MDC-4 in the ideal cipher model. Cryptology ePrint Archive, Paper 2012/113 (2012). <https://eprint.iacr.org/2012/113>
33. Merkle, R.C.: One way hash functions and DES. In: Brassard, G. (ed.) CRYPTO 1989. LNCS, vol. 435, pp. 428–446. Springer, New York (1990). https://doi.org/10.1007/0-387-34805-0_40
34. Preneel, B.: MDC-2 and MDC-4. In: van Tilborg, H.C.A., Jajodia, S. (eds.) Encyclopedia of Cryptography and Security, pp. 771–772. Springer, Boston, MA (2011). https://doi.org/10.1007/978-1-4419-5906-5_596
35. Preneel, B., Govaerts, R., Vandewalle, J.: Hash functions based on block ciphers: a synthetic approach. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 368–378. Springer, Heidelberg (1994). https://doi.org/10.1007/3-540-48329-2_31

36. Quero, N.: Etude des fonctions de hachage homomorphes pour un protocole d'élection secrète pour la blockchain (2022). Internship report. REF: LIST/DSCIN/21-0189/NQ
37. de Souza, L.F., Tucci Piergiovanni, S., Sirdey, R., Stan, O., Quero, N., Kuznetsov, P.: Randsolomon: optimally resilient multi-party random number generation protocol. CoRR abs/2109.04911 (2021). <https://arxiv.org/abs/2109.04911>
38. Guo, J., Peyrin, T., Poschmann, A.: The PHOTON family of lightweight hash functions. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 222–239. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22792-9_13