

Efficient FPGA Implementations of LowMC and Picnic^{*}

Daniel Kales¹, Sebastian Ramacher², Christian Rechberger¹, Roman Walch^{1,3},
and Mario Werner¹

¹ Graz University of Technology, Graz, Austria

{daniel.kales,christian.rechberger,roman.walch,mario.werner}@iaik.tugraz.at

² AIT Austrian Institute of Technology, Vienna, Austria

sebastian.ramacher@ait.ac.at

³ Know-Center GmbH, Graz, Austria

Abstract. Post-quantum cryptography has received increased attention in recent years, in particular, due to the standardization effort by NIST. One of the second-round candidates in the NIST post-quantum standardization project is PICNIC, a post-quantum secure signature scheme based on efficient zero-knowledge proofs of knowledge. In this work, we present the first FPGA implementation of PICNIC. We show how to efficiently calculate LOWMC, the block cipher used as a one-way function in PICNIC, in hardware despite the large number of constants needed during computation. We then combine our LOWMC implementation and efficient instantiations of KECCAK to build the full PICNIC algorithm. Additionally, we conform to recently proposed hardware interfaces for post-quantum schemes to enable easier comparisons with other designs. We provide evaluations of our PICNIC implementation for both, the standalone design and a version wrapped with a PCIe interface, and compare them to the state-of-the-art software implementations of PICNIC and similar hardware designs. Concretely, signing messages on our FPGA takes 0.25 ms for the L1 security level and 1.24 ms for the L5 security level, beating existing optimized software implementations by a factor of 4.

Keywords: LowMC · FPGA · digital signatures · NIST PQC · Picnic

1 Introduction

Cryptographic primitives with low multiplicative complexity have many interesting applications in higher-level protocols.⁴ Recently, the post-quantum secure

^{*} This is the full version of a paper which appears in CT-RSA 2020 – The Cryptographers’ Track at the RSA Conference 2020, San Francisco, CA, USA, February 24-28, 2020, Proceedings. © Springer, 2020.

⁴ For example, they find use as pseudo-random functions (PRF) in secure multi-party computation (MPC) [ARS⁺15,GRR⁺16,RSS17] to handle encrypted data, as oblivious PRFs in private set intersection (PSI) [FIPR05,KRS⁺19], but also enable the elimination of ciphertext expansion in homomorphic encryption schemes [NLV11,ARS⁺15].

digital signature scheme PICNIC [CDG⁺17,CDG⁺19] used zero-knowledge proof of knowledge schemes to build a signature based on the knowledge of a pre-image of a one-way function (OWF). Since the size of the proof of knowledge is directly related to the number of AND gates in the OWF, PICNIC employed block ciphers with low multiplicative complexity. In particular, the designers of PICNIC chose LOWMC [ARS⁺15], a very parameterizable block cipher design, to build the OWF. PICNIC is currently a round 2 candidate in the NIST post-quantum cryptography project [AASA⁺19] and, since the construction lends itself to design more complex statements, it has also been extended to other signature variants including ring and EPID signatures [DRS18b,KKW18,BEF19] as well as double-authentication preventing signatures [DRS18a]. All of those signatures only rely on symmetric-key primitives for their security guarantees.

LOWMC allows its users to select parameters suitable for the intended application. For the use in PICNIC, this means that one can select instances with a reduced data complexity and thereby reducing the number of required rounds. More importantly, it can also be parametrized in such a way that the number of multiplication gates – in this case AND gates – is minimized. For classical security of 128 bits, LOWMC only requires 861 AND gates for full security and 546 AND gates in the reduced data complexity case. In comparison to that, other candidates for lightweight cipher designs require significantly more AND gates to achieve full data security: Simon requires 4352 AND gates [BSS⁺13], Kreyvium requires 1537 AND gates [CCF⁺16], and Fantomas requires 2112 AND gates [GLSV14] (cf. [CDG⁺17, Section 6.1] and [AGR⁺16, Table 1]). None of them come close to the numbers of LOWMC. Only recently, GMiMC [AGP⁺19] was proposed, which can be reduced to 783 AND gates and can compete in the low data complexity scenario. For reference, AES-128 implemented over GF(2) requires more than 5000 AND gates [BMP13] and a round reduced version for the low data complexity case would amount to 3200 AND gates (including key schedule) [BDF11b].

While the choice of LOWMC with small multiplicative complexity significantly reduces the signature size of PICNIC, the number of LOWMC rounds has to be increased for security. Conversely, with higher multiplicative complexity, fewer rounds are required to achieve a secure design. Since each round consists of a matrix multiplication involving the full state, the number of rounds essentially define the runtime characteristics of PICNIC. Additionally, one matrix is sampled uniformly at random for each round during instance generation. In practice, this means that the size of the constants stored in implementations also grows linearly in the number of rounds. For the instances selected for NIST’s L5 security level of 128-bit post-quantum security, the constants sum up to 621 KB. Recent optimizations of the linear layer by Dinur et al. [DKP⁺19] reduced the storage requirements for the constants down to 129 KB. Even with these optimizations, the sheer size of involved matrices seems to prohibit an implementation on resource-constrained devices, like microcontrollers or FPGAs. While the size of constants is less of a problem for software implementations running on desktops, servers

or mobile phones, the size of these constants has a direct impact on the area of a hardware design or the hardware utilization of FPGA designs, respectively.

As the NIST post-quantum project progressed into the second round, the performance of the candidates is becoming a more important criterion. Consequently, NIST published targets for optimized implementations. Since optimized software implementations of LOWMC and PICNIC already exist, we focus on the implementation of several variants on FPGA platforms, including the Xilinx Kintex-7 as well as the Xilinx Artix-7. The latter is one of the optimization platforms recommended by NIST.

1.1 Contribution

Our contribution can be summarized as follows. We provide the first FPGA implementation of LOWMC using a state machine design. Due to the structure of LOWMC, the evaluation of the encryption algorithm requires a large number of constants in the form of uniformly random matrices and vectors. We adapt the recent result of Dinur et al. [DKP⁺19] to our FPGA implementation and are thereby able to significantly reduce the hardware utilization of our design compared to a naïve implementation.

We combine the LOWMC implementation and custom KECCAK modules to instantiate the complete PICNIC design on a Xilinx Kintex-7 board. This implementation conforms to the round 2 submission of the PICNIC signature scheme [CDG⁺19] and supports the L1 and L5 parameter sets. Additionally, we port our implementation of PICNIC-L1 to the Xilinx Artix-7 board. The implementation is flexible enough to support signing only and verification only versions besides the full version without significant overhead. However, our implementation focuses on the Fiat-Shamir transformed version of PICNIC due to recent results [DFMS19,Cha19] improving the confidence in its post-quantum security. Furthermore, the implementation also conforms to the proposed guidelines for hardware designs [FFD⁺18] of post-quantum schemes in the NIST standardization project to facilitate easier comparisons with other designs.

We evaluate the performance of our FPGA design and provide comparisons to optimized software implementations of PICNIC. Our design performs up to a factor 4 faster for signing and up to a factor 3 faster for verification than SIMD optimized software implementations. We also compare our design to other hardware designs of signature schemes including SPHINCS [BHH⁺15]. While the PICNIC design has a higher hardware utilization, it performs significantly better in run time than the SPHINCS design.

We discuss potential modifications to the LOWMC cipher design that would benefit the FPGA implementation of PICNIC. In particular, the suggested changes improve the hardware utilization by up to 30% and makes it possible to fit a PICNIC implementation for the 128-bit post-quantum security level on the NIST recommended Artix-7 evaluation board.

Finally, we also provide a pipeline design of LOWMC for high throughput scenarios. This design may also be of interest in other contexts such as PSI protocols [KRS⁺19], or fast database joins on secret shared data [MRR19].

1.2 Related Work

Efficient hardware implementations are a very active point of research, and the NIST post-quantum standardization project only amplifies this [AASA⁺19]. In the following, we discuss other works in this area, with a focus on hardware implementations of post-quantum algorithms. We still want to mention that for a wide variety of primitives and schemes, hardware implementations have been proposed over the years [MPL⁺11,ACZ16,SG15,SA14], among others.

Basu et al. [BSNK19] provide an evaluation of 11 of the second-round candidates of the NIST post-quantum standardization project. Their approach is based on an automated synthesis of post-quantum accelerators based on the C code provided by the submissions to NIST. They analyze the designs for both FPGA and ASIC targets and compare their runtime and hardware utilization. But since PICNIC is a fairly complex design with a large number of individual primitives and constants, an FPGA implementation is not straightforward and can not be synthesized easily from the software implementation. Therefore, Basu et al. do not include PICNIC in their evaluation.

Besides this generic approach, Amiet et al. [ACZ18] presented an FPGA accelerator for SPHINCS-256 [BHH⁺15], a predecessor of SPHINCS+ [BHK⁺19], another candidate in the NIST post-quantum standardization project. Their implementation provides a highly optimized CHACHA12 pipeline, which is the core of the SPHINCS design. They report signing times of 1.53 ms on the same FPGA we target in our work.

Wang et al. [WSN18] presented FPGA implementations of the Niederreiter cryptosystem using binary Goppa codes. For key encapsulation mechanisms and key exchange, Howe et al. [HOKG18] presented implementations of FrodoKEM, a post-quantum key encapsulation mechanism, for FPGAs and microcontrollers. Roy et al. [RM19] give an optimized implementation of SIKE, a post-quantum key exchange algorithm based on supersingular isogenies, for FPGA platforms. In a different approach, Albrecht et al. [AHH⁺19] repurpose existing RSA coprocessors to speed up computations of RLWE-based schemes, concretely for KEM Kyber.

2 Preliminaries

We will give an overview of LOWMC, PICNIC and its main building blocks.

2.1 LowMC

LowMC [ARS⁺15] is a block cipher designed to reduce the number of AND gates needed for symmetric encryption. The design of the cipher is based on the substitution-permutation network (SPN) design strategy, with the choice to move to a partial substitution layer instead of applying the Sbox on the full state. The parameters block size (n), key size (k), number of Sboxes per round (m), allowed data complexity (d) and the number of rounds (r) are parameterizable




Fig. 1: One round of encryption with LOWMC (modified from [ARS⁺15]).

according to the LOWMC v3 round formula [RST18]. This formula calculates the lowest number of rounds necessary to provide secure encryption for the given parameter set. The script for determining the number of rounds for a given set of LOWMC parameters can be found in the official GitHub repository [Low].

A LOWMC encryption starts with an initial whitening by XORing the first round key to the plaintext, followed by r rounds. As depicted in Figure 1, one round consists of four steps: (i) SBOXLAYER, (ii) LINEARLAYER, (iii) CONSTANTADDITION and (iv) KEYADDITION. In the SBOXLAYER, m 3-bit Sboxes are applied to the first $s = 3 \cdot m$ bits of the state. The remaining bits of the state are not affected by the SBOXLAYER. The Sbox is defined as

$$S(a, b, c) = (a \oplus b \cdot c, a \oplus b \oplus a \cdot c, a \oplus b \oplus c \oplus a \cdot b),$$

with three $\text{GF}(2)$ inputs and outputs. From this definition, it is obvious that only 3 AND gates are required per Sbox. In the LINEARLAYER, the state is multiplied with a pseudorandomly generated matrix $L_r \in \text{GF}(2)^{n \times n}$, where r is the current round. The matrices are chosen pseudorandomly from the set of all invertible binary $n \times n$ matrices during the instantiation of LOWMC. During the CONSTANTADDITION the vector $C_r \in \text{GF}(2)^n$ is XORed to the state, where r describes the current round. The vectors are chosen pseudorandomly during the instantiation of LOWMC. During KEYADDITION, the round key of the current round is XORed to the state. All round keys are generated as a result of the multiplication of the master key with the matrix $K_r \in \text{GF}(2)^{n \times k}$, where r is the current round. The matrices are chosen pseudorandomly from the set of all full-rank binary $n \times k$ matrices during the instantiation of LOWMC.

2.2 Picnic and ZKB++

The PICNIC signature scheme is based on zero-knowledge proofs of knowledge of pre-images of one-way functions. Currently, PICNIC supports two proof systems: ZKB++ [CDG⁺17] and KKW [KKW18]. They both improve on the “MPC-in-the-head” paradigm [IKOS07], which describes a generic way to turn MPC protocols into zero-knowledge proofs. In this work we focus on the parameter set using ZKB++, since it is more efficient in runtime.

ZKB++ builds the zero-knowledge proof system from (2, 3)-circuit decomposition, which we describe in more detail. Let ϕ be some circuit and $y = \phi(x)$, where

x is some secret input and y is the publicly-known output. Within (2,3)-circuit decomposition, the computation is decomposed in the following way [GMO16]:

- SHARE splits the input into three input shares.
- UPDATE advances the computation one gate at a time, computes the wire values for the next gate and returns the updated view.
- OUTPUT produces the output shares based on the final view.
- RECONSTRUCT recomputes the output from the three output shares.

The decomposition of the circuit has to satisfy correctness and 2-privacy [GMO16]:

Correctness: The reconstruction of the output shares y_i must always be the result of the original relation $y = \phi(x)$.

2-Privacy: It should not be possible to reveal information about the private key x by publishing any information on any two players.

In ZKB++, (2,3)-circuit decomposition is constructed as follows: Let R be an arbitrary finite ring and ϕ a function such that $\phi: R^m \rightarrow R^\ell$ can be expressed by an n -gate arithmetic circuit over the ring using addition (respectively multiplications) by constants, and binary addition and binary multiplication gates. A (2,3)-decomposition of ϕ is then given by:

Share(x, k_1, k_2, k_3): Samples random $x_1, x_2 \in R^m$ from k_1 and k_2 and computes x_3 such that $x_1 + x_2 + x_3 = x$. Returns views containing x_1, x_2, x_3 .

Update _{i} ^(j)($\text{view}_i^{(j)}, \text{view}_{i+1}^{(j)}, k_i, k_{i+1}$): Computes player P_i 's view of the output wire of gate g_j and appends it to the view. For the k -th wire w_k where $w_k^{(i)}$ denotes P_i 's view, the update operation is defined as follows:

Addition by constant ($w_b = w_a + c$): $w_b^{(i)} = w_a^{(i)} + c$ if $i = 1$ and $w_b^{(i)} = w_a^{(i)}$ otherwise.

Multiplication by constant ($w_b = c \cdot w_a$): $w_b^{(i)} = c \cdot w_a^{(i)}$

Binary addition ($w_c = w_a + w_b$): $w_c^{(i)} = w_a^{(i)} + w_b^{(i)}$

Binary multiplication ($w_c = w_a \cdot w_b$): $w_c^{(i)} = w_a^{(i)} \cdot w_b^{(i)} + w_a^{(i+1)} \cdot w_b^{(i)} + w_a^{(i)} \cdot w_b^{(i+1)} + R_i(c) - R_{i+1}(c)$ where $R_i(c)$ is the c -th output of a pseudo-random generator seeded with k_i .

Output _{i} ($\text{view}_i^{(n)}$): Return the ℓ output wires stored in the view $\text{view}_i^{(n)}$.

Reconstruct(y_1, y_2, y_3): Computes $y = y_1 + y_2 + y_3$ and returns y .

Note that the player P_i can compute all gate types except for binary multiplication gates locally as the latter requires inputs from P_{i+1} . In other words, only outputs of binary multiplication gates need to be serialized as part of the communication transcript, and thus the view size and consequentially the signature size of PICNIC depend on the size of the ring R and the number of these gates.

To create a proof the prover repeats the (2,3)-circuit decomposition protocol T times. For each run, the prover commits to the view of each player P_i consisting of the input share, a communication transcript, and the output share. After all T runs, the prover sends all the output shares and commitments for each run and player to the verifier, who responds with a challenge vector c . The challenge

tells the prover which two of the three players should be corrupted per run and therefore which views should be published as part of the proof. Since the decomposition satisfies the 2-privacy property, no information is leaked on the secret key by publishing the views of two players. The verifier then recalculates the two opened views and checks, (1) whether the opened views were calculated correctly, (2) and if the three output shares can be reconstructed to y .

Each run gives some assurance that the prover knows the secret key x , therefore increasing the number of runs T decreases the probability that the prover can cheat without the verifier catching him at least once. Due to the nature of the circuit decomposition, the prover could potentially cheat in 2 of the 3 possible challenges per run; therefore we calculate the probability for him to cheat without getting caught as $(2/3)^T$.

ZKB++ is a Σ -protocol, i.e., an interactive proof, that can be made non-interactive by applying standard techniques such as the Fiat-Shamir (FS) transformation [FS86]. FS transformed Σ -protocols compute the challenge c as the output of a random oracle on the first message from the prover to the verifier, which contains the commitments to the shares of the circuit evaluation. This results in a non-interactive zero-knowledge proof protocol secure in the random oracle model. To obtain a signature scheme the message is included in the call to the random oracle as well.

In PICNIC, the circuit used for the circuit decomposition is $C = \text{LOWMC}_k(p)$, where k is a LOWMC secret key and (p, C) , a corresponding plain-/ciphertext pair, which is known publicly and constitutes the public key. A signature is then a proof of knowledge of a k satisfying this relation.

2.3 Picnic Instances and Parameters

For each of the three security levels $S \in \{128, 192, 256\}$, there exist two variants of the PICNIC algorithms differing in the choice of transformations turning the Σ -protocol into a signature scheme: one variant is based on the Fiat-Shamir (FS) transformation, and the other is based on the Unruh (UR) transformation [Unr15]. In contrary to the FS transformation, which makes the resulting non-interactive Σ -protocol secure in the random oracle model, the UR transformation is provably secure in the quantum random oracle model (QROM) [BDF⁺11a], where an adversary can query the random oracle in quantum superposition. However, recent results by Don et al. [DFMS19] and Chailloux [Cha19] show that the specific use of the FS transformation in PICNIC is also secure in the quantum random oracle model. Therefore, we focus our implementations on the variants of PICNIC using the FS transformation and ignore the more costly variants based on the Unruh transformation, that would require additional KECCAK instances to be fitted into the FPGA design.

PICNIC uses LOWMC to reduce the size of the overall proof and thus the signature. The proof contains a transcript of a party, i.e., the view of the party for each AND gate in the circuit. Due to the additive secret-sharing used, XOR gates can be computed locally and do not influence the signature size. Therefore, LOWMC is used instead of other lightweight ciphers, since, as discussed

Table 1: PICNIC parameters with LOWMC instances (block size n , key size k , # of Sboxes m , rounds r), sizes of public key pk , secret key sk and signatures σ .

Parameter Set	S	T	LOWMC				Hash/KDF		Sizes		
			n	k	m	r	Algorithm	ℓ	pk	sk	σ
PICNIC-L1-FS	128	219	128	128	10	20	SHAKE128	256	32	16	≤ 34032
PICNIC-L3-FS	192	329	192	192	10	30	SHAKE256	384	48	24	≤ 76772
PICNIC-L5-FS	256	438	256	256	10	38	SHAKE256	512	64	32	≤ 132856

in Section 1, alternatives to LOWMC require significantly more AND gates. Table 1 shows the LOWMC instances that are used in PICNIC and their AND gate counts. We note that those instances are selected to provide a trade-off between signature size and runtime [CDG⁺17]. We want to note though, that even when the instances are selected based on this trade-off, they can still be represented with a lower number of AND gates than alternative cipher designs.

Table 1 shows the parameters of the different PICNIC versions based on the Fiat-Shamir transformation. The expected security of the various instances corresponds to S bits against classical attacks and $S/2$ bits against quantum attacks. The parameter T describes the number of repetitions of ZKB++ required to reduce the soundness error to the desired security level [CDG⁺19]. Additionally, Table 1 shows the different key and signature sizes for the PICNIC instances. One particular optimization of ZKB++ [CDG⁺17] has the result that the signature size is dependent on the challenge, because for one of the players in the MPC protocol we need to include some auxiliary information in the proof. Therefore, the expected signature size will be smaller than values specified in the table.

In the remainder of this work, we focus our implementation on the PICNIC instances with security levels $S \in \{128, 256\}$ based on the FS transformation, namely PICNIC-L1-FS and PICNIC-L5-FS. Similarly, for LOWMC the focus lies on instances with 128 and 256-bit block and key sizes.

PICNIC2 Instances. In the second round of the NIST post-quantum standardization project, the PICNIC team introduced an additional new parameter set, called PICNIC2. The main difference in the new parameter set is the choice of the underlying proof system. In PICNIC2, ZKB++ was complemented with KKW [KKW18], which is also based on the ‘‘MPC-in-the-head’’ paradigm, but uses a different MPC protocol with precomputation. The nature of the new proof system allows for shorter signatures, but it has an increased number of players in the simulated MPC protocol resulting in longer signing and verification times when compared to PICNIC. An evaluation of these additional parameter sets on FPGA platforms is an interesting topic for future work.

3 Implementation

We will now describe our implementations⁵ of LOWMC and the PICNIC signature algorithm on an FPGA platform. We first give insight into the design of the main module of PICNIC, the computation of LOWMC. Following that, we show how to combine this module with several SHAKE modules to instantiate the full PICNIC signature scheme for the L1 and L5 security levels. In our implementation, we use the dual-port RAM module available as in the open-source framework MEMSEC [WUS⁺17,WU].

Target Platform. For our design, we target a Xilinx Kintex-7 board – concretely we use a Xilinx Kintex-7 FPGA KC705 Evaluation Kit. Our target FPGA has 203800 lookup-tables (LUTs), 407600 flip-flops (FF) and 445 BRAMs available. In an announcement on the official mailing list of the NIST post-quantum standardization project, it was specified that implementors should target the Artix-7 platform due to its widespread use. Since the toolchain we use, Xilinx Vivado, also supports Artix-7 platforms as a target, adaptation to this platform is straightforward, and we discuss the resulting Artix-7 resource utilization in Section 4.1.

3.1 Optimized VHDL Implementation of LowMC

One of the major modules in our PICNIC design is the evaluation of the LOWMC block cipher. During PICNIC’s signing process, a proof of knowledge is generated by evaluating the LOWMC encryption function in an MPC protocol. As discussed in Section 2.2, this is done by applying the (2, 3)-circuit decomposition as defined by ZKB++ to the LOWMC circuit. In terms of the matrix multiplications, the sharing of the circuit requires the 3-fold evaluation for signing.

In this section, we discuss our design choices and the difficulty of a LOWMC VHDL implementation and compare a straightforward standalone implementation of LOWMC with a standalone implementation using the optimizations by Dinur et al. [DKP⁺19]. We shortly give an intuition of these optimizations in the following Sections and refer the reader to [DKP⁺19] for a more detailed explanation.

Design Choices. The difficulty in implementing LOWMC (and consequently PICNIC) in VHDL arises from the high number of constants involved in the matrix multiplications in LOWMC’s linear layer and round key schedule. For the LOWMC instance in PICNIC-L5, 621 kB of constants are required which can be reduced to 129 kB by using the optimizations in [DKP⁺19]. Usually, we consider using block RAM (BRAM), RAM cells directly located on FPGAs, for storing a large amount of constants. The Kintex-7 FPGA comes with dual-port BRAM cells with a capacity of 36 kB each, which are capable of providing at most 72 bits during one clock cycle at each port. During one round we multiply

⁵ All implementations are available at <https://github.com/IAIK/Picnic-FPGA>.

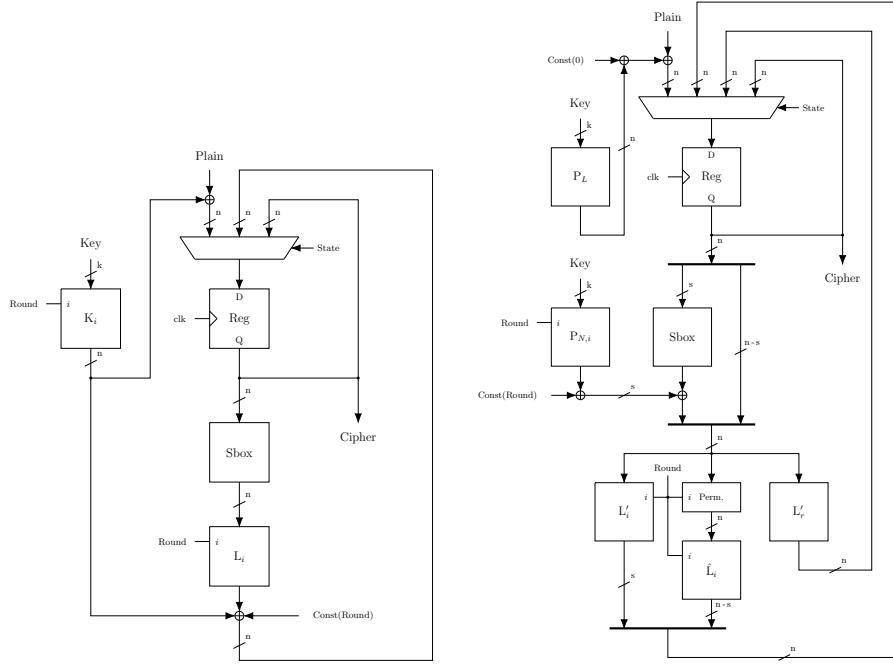
the inner state of LOWMC to an $S \times S$ bit (256×256 for PICNIC-L5) matrix. Considering a high-performance implementation, where we want to perform the matrix multiplication in one clock cycle, we would have to use ≈ 455 BRAM cells in parallel, which exceeds the number of available cells. The alternative multi-cycle approach would, therefore, necessarily lower the performance of the implementation. Furthermore, in the case of high BRAM cell usage, additional clock frequency penalties have to be expected due to increasing routing delays.

In our implementation, we decided to encode the constants for the matrices in lookup-tables (LUTs). This decision implies a high hardware utilization of our design, but comes with the advantage of fast matrix multiplications (1 clock cycle each) and therefore with the best performance. A low area implementation of LOWMC (and consequently PICNIC) using BRAMs instead of LUTs for constants could be an interesting topic for future work.

LOWMC Optimizations. The main idea behind the optimizations proposed in [DKP⁺19] is that all operations except the Sbox layer are linear. Furthermore, only part of the state is affected by the Sbox layer. For example, consider the key addition: we can swap the order of the linear layer and the key addition by multiplying the round key with the inverse of the linear layer. Subsequently, we can move part of the key addition through the identity part of the Sbox layer and combine this part of the key addition with the key addition of the previous round. This process can be repeated recursively until we have combined a large part of the key additions in the initial key addition before the first round. The same process can be repeated for the round constants. Using some more advanced linear algebra properties, we can also move parts of the linear layer matrix multiplication to the next or previous round, again repeating this process to combine parts of the linear layer in the last or first round.

Figure 2a shows the VHDL design of LOWMC without the optimizations. In Figure 2b, we present the design with the optimizations applied. Without the optimizations, there is only one implementation for all the rounds, and the matrix multiplications affect the entire state. In the optimized implementation, there are 5 different matrix multiplication modules, each for a matrix with different dimensions. The Sbox layer, round key, and constants of the new implementation only affect the first s bits of the state, and the linear layer matrix multiplication follows the algorithm in [DKP⁺19].

Optimized Hardware Utilization. The impact of the optimizations depends on the concrete LOWMC instance. It especially depends on the number of Sboxes m and the resulting size of the non-linear layer. The fewer Sboxes per round, the more significant is the effect of the optimizations. The concrete effect of the optimizations can be seen in Table 2, where the required lookup-tables (LUTs) of the LOWMC VHDL implementation are shown for the two different LOWMC instances used in PICNIC-L1-FS and PICNIC-L5-FS. The instance for security level L1 only requires about a third of the LUTs required before, and the instance for security level L5 only requires about a fifth of the LUTs of the straightforward version. Without the optimizations, it would not even be possible to synthesize



(a) LOWMC implementation without optimizations of the round key and linear layer computations. (b) LOWMC implementation with optimized round key computation and linear layer evaluation.

Fig. 2: State diagrams of different LOWMC implementations.

Table 2: LUTs of one LOWMC with/without optimizations (203800 available).

LOWMC instance	LOWMC				without opt.		with opt.		Improv. %
	n	k	m	r	LUTs	% LUTs	LUTs	% LUTs	
PICNIC-L1-FS	128	128	10	20	42395	20.80 %	13558	6.65 %	68.02 %
PICNIC-L5-FS	256	256	10	38	209348	102.72 %	44431	21.8 %	78.78 %

one LOWMC instance for security level L5 on our FPGA board, whereas we require several instances for the PICNIC implementation. The improvement for LOWMC instances with larger non-linear layers is smaller, though.

3.2 Pipeline versus State Machine

Besides the implementation of LOWMC using a simple state machine, we also provide an alternative implementation using a pipelined design. While both designs have a latency of r cycles to get a specific ciphertext, the pipelined design has a much higher throughput with 1 ciphertext per cycle. The state machine

design, on the other hand, has to wait for an encryption to be finished before it can process another plaintext and therefore has a throughput of 1 ciphertext per r rounds. However, the state machine design requires fewer lookup tables on an FPGA, because the LOWMC round only needs to be instantiated once. For the PICNIC coprocessor, we use the state machine design due to smaller hardware utilization. When interested in higher throughput, for example, when it is used as an oblivious pseudo-random function in a PSI protocol, the pipeline design is the better choice.

3.3 Optimized VHDL Implementation of Picnic

We now use the LOWMC implementation as a building block for our PICNIC coprocessors. In the following, we shortly describe the other different submodules and finally, the high-level design of the PICNIC coprocessors.

LOWMC-MPC. In PICNIC, three copies of the LOWMC encryption circuit are evaluated with three random additive shares of the secret key. Since the secret-sharing used is additive, XOR gates can be computed locally for each part, while some communication between the parties and randomness is required for computing an AND gate. While a straightforward implementation of this uses three copies of the LOWMC circuit, we present a further optimization. The nature of the secret-sharing and circuit decomposition used in ZKB++ ensures that for each wire w in the circuit, the equality $w = w_1 \oplus w_2 \oplus w_3$, holds, where w_i is the share of party i . If we evaluate the circuit once in plain and store all intermediate values w , we can use only two instances of LOWMC for signing and compute the shares of the third party $w_3 = w \oplus w_1 \oplus w_2$ whenever needed. This optimization allows us to implement the LOWMC-MPC module using resources equivalent to only two LOWMC circuit evaluations, while still being able to evaluate all players simultaneously. Additionally, we can precompute the plain evaluation of the LOWMC circuit in parallel to the Seeds calculation at the beginning of the PICNIC signing process and, therefore, do not slow down signing while using this optimization.

During signature verification, only two players perform the LOWMC-MPC circuit evaluation; therefore, we naturally only require resources of about two LOWMC circuit evaluations to perform all players in parallel.

SHAKE. In PICNIC, instances of SHAKE are used for different purposes, both as a hash function with fixed output or as an extensible output function to generate pseudorandom tapes of arbitrary size. Therefore, we implemented a custom, flexible KECCAK design, supporting many different configurations while maintaining efficiency and small hardware utilization.

Seeds, Tapes, and Commitments. In the beginning, one master seed is pseudo-randomly generated and expanded into seeds for each of the T runs. We use three instances of SHAKE to expand the seeds for each player’s current run into its random tape and three more instances to commit to the transcript of

the current run for each player. We are capable of calculating the randomness required for run $t+1$ of PICNIC’s circuit decomposition in parallel to calculating the commitments of run t , reducing the overall number of clock cycles for signing and verification. However, due to limited routing freedom due to high resource utilization of our synthesized PICNIC-L5 design, this optimization would significantly increase the critical path of the design and, therefore, this optimization is only used for PICNIC-L1.

Challenge Generation (H3). Based on the Fiat-Shamir transformation, we instantiate the random oracle for the challenge generation using SHAKE. All commitments for all T runs are hashed together with some additional parameters to produce the challenge vector. Since the challenge vector consists of entries in $\{0, 1, 2\}$ to denote the player that is not revealed for this run, the H3 module takes care to filter the output bits of the SHAKE call according to the PICNIC specification.

Serialization and Deserialization. We also implemented small submodules to assemble the final signature as a byte array conforming to the PICNIC specification. For verification, we parse incoming signatures and store all the intermediate values of the opened views in the block RAM cells of the FPGA. These modules are implemented to be able to handle the variable signature length of PICNIC internally.

High-Level Design. We developed several different VHDL designs for PICNIC-L1-FS and PICNIC-L5-FS. We implemented a standalone version for message signing or signature verification only, as well as a version which is capable of doing both.

The overall design of the implementations is a nested state machine, where the high-level design connects the inputs and outputs of all the described submodules. Figure 3 shows a diagram of the high-level design, with the signing process shown on the left side and the verification process on the right. In the designs which are capable of doing both signing and verification, both processes are implemented. Most of the submodules can be reused for both signing and verification, only the MPC module has to implement two different Sbox calculations, and the combined design has to include both the *Serialize* and *Deserialize* submodules. Therefore, the difference in hardware utilization between a sign-only design and a sign/verify design is quite low.

4 Evaluation

In the following, we evaluate and discuss the performance and hardware utilization of our design. We not only give cycle counts for signing and verification but also show the real-world performance of our designs by additionally synthesizing a PCIe wrapper around our PICNIC cores and using them from a C library.

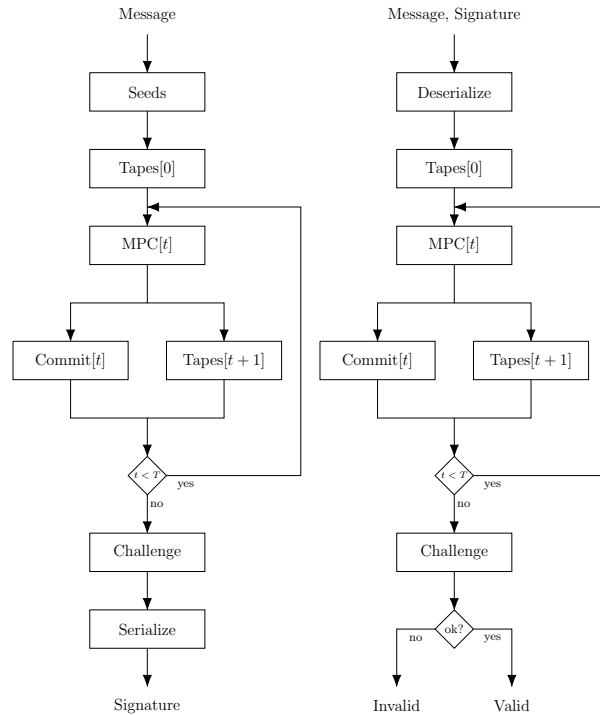


Fig. 3: High-level design of PICNIC signing (left) and verification (right).

4.1 Hardware Utilization

First, we give an overview of the required hardware utilization of the different PICNIC submodules, and then show the utilization of the developed coprocessors. The used FPGA, a Xilinx Kintex-7 board, has 203800 lookup-tables (LUTs), 407600 flip-flops (FF) and 445 BRAMs available.

PICNIC Submodules. To give an overview of the costs of the individual submodules, we present their hardware utilization for PICNIC-L1-FS and PICNIC-L5-FS in Table 3. This table shows that the LOWMC-MPC modules require by far the most hardware utilization. However, observe that the combined submodule which is able to do the LOWMC-MPC encryption for both signing and verification only requires less than one percent more LUTs than the submodule which can only be used for signing. This is because we can reuse large parts of the circuit for both signing and verification, the only difference is in the Sbox layer, where the AND gates are evaluated.

PICNIC Coprocessors. Table 4 compares the hardware utilization of the different submodules of the final coprocessors, including our 6 different PICNIC cores synthesized for the Kintex-7. Our PICNIC cores require a lot of LUTs on the

Table 3: Hardware utilization for different parts of the L1 and L5 designs on Kintex-7.

Design Part	L1				L5			
	LUTs	%	FF	%	LUTs	%	FF	%
KECCAK	3726	1.83 %	1606	0.39 %	3726	1.83 %	1606	0.39 %
Tapes (3× KECCAK)	9574	4.70 %	5589	1.37 %	9420	4.62 %	9621	2.36 %
Commits (3× KECCAK)	12221	6.00 %	5589	1.37 %	14160	6.95 %	6357	1.56 %
Seeds (1× KECCAK)	5867	2.88 %	1846	0.45 %	8974	4.40 %	2640	0.65 %
H3 (1× KECCAK)	7236	3.55 %	3641	0.89 %	8815	4.33 %	4085	1.00 %
Serialize	1962	0.96 %	125	0.03 %	1608	0.79 %	172	0.79 %
Deserialize	2025	0.99 %	125	0.03 %	2317	1.14 %	155	0.04 %
LowMC-MPC Sign	31837	15.62 %	3060	0.75 %	97066	47.63 %	5940	1.46 %
LowMC-MPC Verify	29756	14.60 %	1126	0.28 %	93959	46.10 %	2246	0.55 %
LowMC-MPC	32224	15.81 %	3061	0.75 %	98319	48.24 %	5958	1.46 %

Table 4: Hardware utilization for different parts of the coprocessor for Kintex-7.

Design Part	LUTs	%	FF	%	BRAM	%
PCIe/DMA	22216	10.90 %	22692	5.57 %	42.5	9.55 %
PICNIC-L1	90037	44.18 %	23105	5.67 %	52.5	11.80 %
PICNIC-L1-SIGN	76472	37.52 %	21061	5.17 %	52.5	11.80 %
PICNIC-L1-VERIFY	68614	33.67 %	16821	4.13 %	33.5	7.53 %
PICNIC-L5	167530	82.20 %	33164	8.14 %	98.5	22.13 %
PICNIC-L5-SIGN	149456	73.33 %	30441	7.47 %	98.5	22.13 %
PICNIC-L5-VERIFY	138547	67.98 %	24278	5.96 %	62.5	14.04 %

used FPGA, especially the PICNIC-L5-FS implementations. The PCIe/DMA Subsystem which connects the PICNIC cores to the PCIe port of the used FPGA board adds about 22 000 additional LUTs to the design.

On the Artix-7, the picture is quite different, as we only have 133 800 LUTs, 267 600 flip-flops and 365 BRAMs available. Consequently, neither PICNIC-L5 PICNIC-L5-SIGN, nor PICNIC-L5-VERIFY fit on this board. The hardware utilization of the Artix implementations of L1 are depicted in Table 5.

Critical Path. The critical path of the synthesized design is across the matrix multiplications in LOWMC’s linear layer and round key schedule, due to the high number of constants involved. But we also observed, that since the PICNIC-L5 design has a considerable hardware utilization, the synthesizer has much less freedom in routing the design and, therefore, naturally produces long paths

Table 5: Hardware utilization on Artix-7.

Design Part	LUTs	%	FF	%	BRAM	%
PICNIC-L1	90037	67.29 %	23105	8.63 %	52.5	14.38 %
PICNIC-L1-SIGN	76472	57.15 %	21061	7.87 %	52.5	14.38 %
PICNIC-L1-VERIFY	68614	51.28 %	16821	6.29 %	33.5	9.18 %

between registers. These long paths make it very difficult to optimize the design for high frequencies.

4.2 Clock Cycles

Table 6 lists the number of clock cycles each submodule of our PICNIC implementation requires. The LOWMC-MPC module performs the evaluation of a round in two clock cycles and, therefore, requires $2 \cdot r$ cycles in total. Evaluating a round in one cycle would have drastically increased the critical path of the design since two matrix multiplications (linear layer and round key schedule) would have been performed sequentially in this case.

Our KECCAK implementation performs one round of the state transformation function during one clock cycle, which leads to 24 cycles for one absorbing/squeezing phase. The number of absorbing/squeezing phases, therefore, determines the number of clock cycles required for the *Tapes*, *Commits*, *Seeds*, and the first part of the *H3* submodules. The duration of the second part of the *H3* submodule depends on the generated challenge and differs for every signature.

In PICNIC we have T runs of the FS transformed ZKB++ proof system, in contrary to the *Seeds*, *H3*, *Serialize* and *Deserialize* modules which are only required once. In Table 6 we, therefore, also show the overall runtime of each submodule involved in the proof creation and it can be seen, that the proof system dominates the overall runtime of the signature creation and verification process.

4.3 Benchmarks

To verify the performance characteristics of our implementation, we compared the runtime of the coprocessors running on a Kintex-7 board to the state-of-the-art optimized software implementations of PICNIC. The platforms used for the benchmarks are as follows:

Platform A Intel i7-960, 3.2 GHz with 16 GB RAM, Debian 9

Platform B Intel i7-4790, 3.6 GHz with 16 GB RAM, Ubuntu 18.04.1

Platform C Intel E31230, 3.2 GHz with 8 GB RAM, Ubuntu 18.04.2

We used platform A to test our coprocessors, platforms B and C were used in the PICNIC design document [CDG⁺19] to test their optimized software implementations.

Table 6: Clock Cycles per Submodule.

Design Part	PICNIC-L1-FS	PICNIC-L5-FS
LowMC-MPC	40	76
Tapes	51	75
Commits	51	100
Seeds	1 732	7 904
H3 (absorb)	6 490	26 220
Deserialize	1 per 128 bit	1 per 128 bit
Serialize	1 per 128 bit	1 per 128 bit
$T \times$ LowMC-MPC	8 760	31 844
$T \times$ Tapes	11 169	31 425
$T \times$ Commits	11 169	41 900

Table 7: Runtime comparison of the coprocessors on benchmark platform A.

Coprocessor	clock frequency (MHz)	clock cycles	FPGA runtime (ms)	C-Access runtime (ms)
PICNIC-L1-SIGN	125	\approx 31300	0.250	0.349
PICNIC-L1-VERIFY	125	\approx 29600	0.237	0.395
PICNIC-L5-SIGN	125	\approx 154500	1.236	1.383
PICNIC-L5-VERIFY	125	\approx 146600	1.173	2.128

Table 7 shows the average runtime of the developed coprocessors for signing and verification. The column *FPGA runtime* is the calculated time resulting from the clock frequency and the number of clock cycles (including 1 cycle per 128 bit of data transmission) and therefore is the actual runtime of the FPGA. The column *C-Access runtime* is the measured runtime using our developed C library on platform A.

As Table 7 shows, the C library developed to interface with the coprocessor adds some overhead to the signing and verification process. For signing, the overhead is about 0.1 ms in runtime, but for verification, the overhead is a bit larger. Especially for PICNIC-L5-FS the measured runtime is much bigger than the raw verification runtime of the coprocessor. We suspect that this is due to the driver for the PCIe/DMA Subsystem being slower for writing large amounts of data, like the PICNIC-L5-FS signature, from the PC to the FPGA board and that this overhead could be optimized further.

For comparison, Table 8 shows the runtime of the optimized implementation of PICNIC in C and an optimized version which uses processor-specific compiler intrinsics on two different benchmark platforms as described in the official PICNIC design document [CDG⁺19]. This table shows, that the runtime of PICNIC highly depends on the underlying hardware and if the CPU supports *single in-*

Table 8: Runtime comparison of optimized software implementations [CDG⁺19].

Platform	Parameters	using SIMD	Sign	Verify
B	PICNIC-L1-FS	✓	1.44 ms	1.15 ms
B	PICNIC-L5-FS	✓	5.87 ms	4.92 ms
B	PICNIC-L1-FS	✗	2.82 ms	2.34 ms
B	PICNIC-L5-FS	✗	12.37 ms	10.59 ms
C	PICNIC-L1-FS	✓	4.20 ms	3.40 ms
C	PICNIC-L5-FS	✓	17.67 ms	14.67 ms
C	PICNIC-L1-FS	✗	4.41 ms	3.56 ms
C	PICNIC-L5-FS	✗	19.52 ms	16.81 ms

struction, multiple data (SIMD) instruction sets, like SSE2 and AVX2, which further improve execution time. However, in any case, our developed coprocessors are faster than the corresponding software counterparts and do not rely on specific CPU instructions. For PICNIC-L1-FS signing is ≈ 4 times faster than the fastest software implementation, verification is ≈ 3 times faster. For PICNIC-L5-FS our implementations are ≈ 4 times faster for signing and ≈ 2.3 times faster for verification. For CPUs which do not support AVX2 instructions and for portable C-only implementations the speedup of our coprocessors is even more significant.

4.4 Comparison to FPGA Implementations of Other Signature Schemes

To put our FPGA implementation in context of other signature schemes, we compare our PICNIC coprocessors to implementations of ECDSA [ACZ16] and RSA coprocessors [SA14] as well as implementations of SPHINCS-256 [ACZ18] and BLISS-IV [PDG14]. Table 9 compares several different FPGA implementations of various signature schemes, the runtime for signing t is calculated from the clock frequency and the number of clock cycles and therefore does not take the overhead of any transmission of data via a C-program into account. Thus this value compares to the column *FPGA runtime* of Table 7.

As Table 9 shows, our PICNIC-L5-FS coprocessors, which have the same security level as a SPHINCS-256 [ACZ18] coprocessor, have a slightly better runtime for signing on the Kintex-7 (K7) FPGA. Similar, for the SPHINCS+ design obtained from the high-level synthesis design flow [BSNK19], our coprocessor has a significant better runtime at the cost of higher hardware utilization. The implementations of the traditional signature schemes RSA [SA14] and ECDSA [ACZ16] on a Virtex-7 (V7) FPGA are also slower than our coprocessors. The ECDSA implementation in [Gün11] occupies more area but uses high parallelism to drastically increase their throughput. The implementation of BLISS-IV [PDG14], another post-quantum signature scheme based on lattices,

Table 9: Comparison of FPGA implementations (modified from [ACZ18]).

Scheme	Security			Area	f	t
	Classic	PQ	FPGA	LUT/FF/BRAM	MHz	ms
SPHINCS-256 [ACZ18]	256	128	K7	19067/38132/36	525	1.53
SPHINCS+-128 [BSNK19]	128	64	V7	11438/3335/?	100	9.38
BLISS-IV [PDG14]	192	?	S6	6438/6198/7	135	0.35
ECDSA-256 [ACZ16]	128	✗	V7	6816/4442/0	225	1.49
ECDSA-256 [Gün11]	128	✗	V4	34869/32430/176	375	0.04
ECDSA-521 [ACZ16]	256	✗	V7	8273/7689/0	161	5.02
RSA-2048 [SA14]	112	✗	V7	3558 slices/0	399	5.68
PICNIC-L1-FS	128	64	K7	90037/23105/52.5	125	0.25
PICNIC-L5-FS	256	128	K7	167530/33164/98.5	125	1.24

on a Spartan-6 (S6) FPGA is very efficient regarding area and runtime for signing. However, it has a lower security level, and its security against a quantum adversary is not as well understood as for schemes based on symmetric primitives like SPHINCS and PICNIC.

However, even though our coprocessors are very competitive with regards to signing times, the hardware utilization is significantly higher in comparison to implementations of other signature schemes. This is due to the nature of PICNIC relying on a high number of different KECCAK and LOWMC primitives, where especially the LOWMC instances have a high hardware utilization on their own. In comparison, the coprocessor of SPHINCS-256, a hash-based post-quantum signature scheme, can be built efficiently using only one pipelined CHACHA12 instance and one instance of BLAKE-256 [ANWW13] and as a result, requires less hardware utilization [ACZ18].

4.5 Evaluation of the LowMC Pipeline Design

Finally, we evaluate our pipelined design. After r cycles, the design is capable of producing one ciphertext per cycle (cf. Section 3.2), a feature which is of particular interest for high throughput use cases. We compare our coprocessor ($f = 125$ MHz) for a LOWMC instance with 128 bit block size and full data complexity to AES-128 accelerated with the AES-NI instruction set [Int10]. For this comparison we choose a LOWMC instance with $n = 128$, $k = 128$, $m = 25$, $r = 11$. This instance provides a trade-off between costs in the linear layer and the number of AND gates. The comparison of the coprocessor, including C-access times, the raw FPGA runtime, the SIMD-optimized LOWMC software implementation and AES-NI is depicted in Table 10. These benchmarks were recorded on a PC running Ubuntu 16.05 with an Intel i7-4790 CPU, 3.6 GHz. As the table shows, the coprocessor speeds up encryption by a factor of ≈ 84 compared to the LOWMC software implementation, and when considering the C-access time, the improvement is still up to a factor of ≈ 14 . Compared to

Table 10: Performance of LOWMC($n = 128, k = 128, m = 25, r = 11$) implemented in software and in our pipeline coprocessor, as well as AES-NI.

# Encryptions	Size	LOWMC		AES	
		FPGA-Raw	FPGA-C	Software	AES-NI
2^{20}	16 MB	0.008s	0.046s	0.677s	0.022s
2^{24}	256 MB	0.134s	0.771s	10.78s	0.359s
2^{26}	1024 MB	0.537s	3.11s	43.31s	1.436s
2^{28}	4096 MB	2.15s	12.57s	182.65s	5.743s

AES-NI, the raw performance on the FPGA is better by a factor of ≈ 2.75 , but the access time adds significant overhead. Therefore, we expect this design to render LOWMC an alternative for PSI protocols [KRS⁺19] or database joins on secret shared data [MRR19].

This speed up direct translates to the same performance gain in the setup phase of the PSI protocol as proposed in [KRS⁺19]. Thus, the excellent performance of our coprocessor makes it feasible to use LOWMC in the PSI protocol. Thereby, the PSI protocol can profit from reduced communication overhead during the online phase due to the reduced multiplicative complexity of LOWMC without requiring the tradeoff of having a slower setup phase.

5 Reducing the Hardware Utilizations

The large size of the constants needed for LOWMC is one of the limiting factors to implement PICNIC on FPGAs. Even after applying the optimizations to the round key and linear layer computations, the constants are still too large to fit an implementation on an Artix-7 board. To fit an implementation of PICNIC suitable for the 128-bit post-quantum security level on this board, different LOWMC instances with fewer rounds could be selected. Conversely, as this change requires the number of Sboxes to be increased to retain the security guarantees, the signature size will increase. In addition to fitting PICNIC on smaller FPGA boards, the performance of the optimized implementations would also improve, since fewer rounds are required to achieve the same level of security when more Sboxes are used. Alternatively, further improvements are required to reduce the size of LOWMC constants. We envision multiple alternatives that could make this possible.

The current design of PICNIC was chosen to have an acceptable trade-off between area and runtime and, therefore, evaluates the LOWMC-MPC simulation concurrently for all three players by using two instances of the LOWMC matrices. By doing the MPC simulation consecutively, we would be able to reduce the instances used to only one and reduce the hardware utilization. However, this optimization would result in more clock cycles per LOWMC-MPC rounds and

Table 11: Hardware utilization (LUTs) with reduced LOWMC.

Design Part	LUTs	Improvement	Utilization	
			Kintex-7	Artix-7
LOWMC MPC-L1	17751	44.91 %	8.71 %	13.27 %
LOWMC MPC-L5	47615	51.57 %	23.36 %	35.59 %
PICNIC-L1	75662	15.97 %	37.13 %	56.55 %
PICNIC-L1-SIGN	62272	18.57 %	30.56 %	46.54 %
PICNIC-L1-VERIFY	55321	19.37 %	27.14 %	41.35 %
PICNIC-L5	121299	27.60 %	59.52 %	90.66 %
PICNIC-L5-SIGN	103688	30.62 %	50.88 %	77.49 %
PICNIC-L5-VERIFY	92910	32.94 %	45.59 %	69.44 %

longer critical paths after synthesis reducing the clock frequency, and, therefore, would produce a very high performance penalty of at least a factor 2 if not more.

Another possibility to reduce the hardware utilization by modifying our design would be to reuse KECCAK instances for different purposes in the design. However, this would again result in longer critical paths after synthesis, and since our KECCAK design is very small in comparison to the LOWMC design, the resulting performance penalty is too big in comparison to the actual reduction of the hardware utilization.

The use of LOWMC in PICNIC is relatively unique in the sense that it uses LOWMC instances with low data complexity. Only recently, LOWMC in this setting has seen more security analysis [RST18], leading to LOWMC version 3 with a higher number of rounds. While the higher number of rounds on its own is not a problem for the FPGA implementation, the size of the constants also increases as more and more unique matrices are required. However, new designs [GKK⁺19] that are also optimized for a low multiplicative complexity make use of a single matrix for the linear layer. We propose to apply this idea also to LOWMC, that is, the same uniformly random matrix is re-used for all linear layers.⁶ Thereby we can significantly reduce the hardware utilization as can be seen in Table 11. With this change, the PICNIC-L5 design fits on the Artix-7.

Furthermore, with this change in place, one could go a step further and remove the constants from the implementation altogether. The matrices and round constants could then be derived from the LFSR as specified in the LOWMC instance generation algorithm. It would be necessary to store the intermediate states of the LFSR where it is known to produce invertible matrices, though, but then no invertible checks would need to be implemented. While deriving the

⁶ The LOWMC designers confirmed in private communication that they do not expect this change to enable a new attack vector. However, more security analysis on this case would be required before this can be integrated into PICNIC itself.

matrices during runtime would come with a performance penalty, we expect it to reduce the hardware utilization significantly.

Alternatively, LOWMC could be replaced by recently proposed cipher designs such as GMiMC [AGP⁺19]. Similarly to LOWMC, GMiMC – and in particular its ERF variant – can be parameterized for the low data complexity scenario. It can be parameterized in a way leading to roughly similar sized signatures with better performance (in software). However, for an FPGA implementation, we expect it to use a lot less area since the size of the constants is significantly smaller. For the GMiMC instance over $\text{GF}(2^{17})$ with 63 rounds, the constants would consist of only 63 field elements in total. The additional multipliers required for $\text{GF}(2^{17})$ are cheap [ER15] when compared to the size of LOWMC matrices.

6 Conclusion

In this work, we presented two LOWMC designs for FPGAs. The first design relying on a simple state machine shows the feasibility of implementing the post-quantum signature scheme PICNIC on FPGA platforms. The resulting FPGA design can sign messages for the L1 security level in ≈ 31300 cycles and verify signatures in ≈ 29600 cycles. Using our concrete FPGA board, a Xilinx Kintex-7 FPGA KC705 evaluation kit, this allows signing of a message using a C library communicating with our board connected via PCIe in 0.35 ms.

Acknowledgements. This work was partially supported by the EU’s Horizon 2020 ECSEL Joint Undertaking project SECREDAS under grant agreement n°783119, by the European Research Council (ERC) under Horizon 2020 grant agreement n°681402, by EU’s Horizon 2020 project Safe-DEED under grant agreement n°825225, and by the IoT4CPS project which is partially funded by the “ICT of the Future” Program of the FFG and the BMVIT. D. Kales was supported by iov42 Ltd.

References

- AASA⁺19. Gorjan Alagic, Jacob Alperin-Sheriff, Daniel Apon, David Cooper, Quynh Dang, Carl Miller, Dustin Moody, Rene Peralta, Ray Perlner, Angela Robinson, Daniel Smith-Tone, and Yi-Kai Liu. Status report on the first round of the nist post-quantum cryptography standardization process, 2019.
- ACZ16. Dorian Amiet, Andreas Curiger, and Paul Zbinden. Flexible fpga-based architectures for curve point multiplication over $\text{gf}(p)$. In *DSD*, pages 107–114. IEEE Computer Society, 2016.
- ACZ18. Dorian Amiet, Andreas Curiger, and Paul Zbinden. Fpga-based accelerator for post-quantum signature scheme SPHINCS-256. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(1):18–39, 2018.
- 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 *ESORICS (2)*, volume 11736 of *LNCS*, pages 151–171. Springer, 2019.

- AGR⁺16. Martin R. Albrecht, Lorenzo Grassi, Christian Rechberger, Arnab Roy, and Tyge Tiessen. Mimc: Efficient encryption and cryptographic hashing with minimal multiplicative complexity. In *ASIACRYPT (1)*, volume 10031 of *LNCS*, pages 191–219, 2016.
- AHH⁺19. Martin R. Albrecht, Christian Hanser, Andrea Höller, Thomas Pöppelmann, Fernando Virdia, and Andreas Wallner. Implementing rlwe-based schemes using an RSA co-processor. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2019(1):169–208, 2019.
- ANWW13. Jean-Philippe Aumasson, Samuel Neves, Zooko Wilcox-O’Hearn, and Christian Winnerlein. BLAKE2: simpler, smaller, fast as MD5. In *ACNS*, volume 7954 of *LNCS*, pages 119–135. Springer, 2013.
- ARS⁺15. Martin R. Albrecht, Christian Rechberger, Thomas Schneider, Tyge Tiessen, and Michael Zohner. Ciphers for MPC and FHE. In *EUROCRYPT (1)*, volume 9056 of *LNCS*, pages 430–454. Springer, 2015.
- BDF⁺11a. Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In *ASIACRYPT*, volume 7073 of *LNCS*, pages 41–69. Springer, 2011.
- BDF11b. Charles Bouillaguet, Patrick Derbez, and Pierre-Alain Fouque. Automatic search of attacks on round-reduced AES and applications. In *CRYPTO*, volume 6841 of *LNCS*, pages 169–187. Springer, 2011.
- BEF19. Dan Boneh, Saba Eskandarian, and Ben Fisch. Post-quantum EPID signatures from symmetric primitives. In *CT-RSA*, volume 11405 of *LNCS*, pages 251–271. Springer, 2019.
- BHH⁺15. Daniel J. Bernstein, Daira Hopwood, Andreas Hülsing, Tanja Lange, Ruben Niederhagen, Louiza Papachristodoulou, Michael Schneider, Peter Schwabe, and Zooko Wilcox-O’Hearn. SPHINCS: practical stateless hash-based signatures. In *EUROCRYPT (1)*, volume 9056 of *LNCS*, pages 368–397. Springer, 2015.
- BHK⁺19. Daniel J. Bernstein, Andreas Hülsing, Stefan Kölbl, Ruben Niederhagen, Joost Rijneveld, and Peter Schwabe. The sphincs⁺ signature framework. In *CCS*, pages 2129–2146. ACM, 2019.
- BMP13. Joan Boyar, Philip Matthews, and René Peralta. Logic minimization techniques with applications to cryptology. *J. Cryptology*, 26(2):280–312, 2013.
- BSNK19. Kanad Basu, Deepraj Soni, Mohammed Nabeel, and Ramesh Karri. NIST post-quantum cryptography- A hardware evaluation study. *ePrint*, 2019:47, 2019.
- BSS⁺13. Ray Beaulieu, Douglas Shors, Jason Smith, Stefan Treatman-Clark, Bryan Weeks, and Louis Wingers. The SIMON and SPECK families of lightweight block ciphers. *ePrint*, 2013:404, 2013.
- CCF⁺16. 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. In *FSE*, volume 9783 of *LNCS*, pages 313–333. Springer, 2016.
- CDG⁺17. Melissa Chase, David Derler, Steven Goldfeder, Claudio Orlandi, Sebastian Ramacher, Christian Rechberger, Daniel Slamanig, and Greg Zaverucha. Post-quantum zero-knowledge and signatures from symmetric-key primitives. In *ACM CCS*, pages 1825–1842. ACM, 2017.
- CDG⁺19. Melissa Chase, David Derler, Steven Goldfeder, Jonathan Katz, Valdimir Kolesnikov, Claudio Orlandi, Sebastian Ramacher, Christian Rechberger, Daniel Slamanig, Xiao Wang, and Greg Zaverucha. The picnic signature scheme design document (version 2), 2019.

- Cha19. André Chailloux. Quantum security of the fiat-shamir transform of commit and open protocols. *ePrint*, 2019:699, 2019.
- DFMS19. Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Security of the fiat-shamir transformation in the quantum random-oracle model. In *CRYPTO (2)*, volume 11693 of *LNCS*, pages 356–383. Springer, 2019.
- DKP⁺19. Itai Dinur, Daniel Kales, Angela Promitzer, Sebastian Ramacher, and Christian Rechberger. Linear equivalence of block ciphers with partial non-linear layers: Application to lowmc. In *EUROCRYPT (1)*, volume 11476 of *LNCS*, pages 343–372. Springer, 2019.
- DRS18a. David Derler, Sebastian Ramacher, and Daniel Slamanig. Generic double-authentication preventing signatures and a post-quantum instantiation. In *ProvSec*, volume 11192 of *LNCS*, pages 258–276. Springer, 2018.
- DRS18b. David Derler, Sebastian Ramacher, and Daniel Slamanig. Post-quantum zero-knowledge proofs for accumulators with applications to ring signatures from symmetric-key primitives. In *PQCrypto*, volume 10786 of *LNCS*, pages 419–440. Springer, 2018.
- ER15. Hayssam El-Razouk and Arash Reyhani-Masoleh. New bit-level serial GF (2^m) multiplication using polynomial basis. In *ARITH*, pages 129–136. IEEE, 2015.
- FFD⁺18. Ahmed Ferozpur, Farnoud Farahmand, Viet Dang, Malik Umar Sharif, Jens-Peter Kaps, and Kris Gaj. Hardware API for Post-Quantum Public Key Cryptosystems, 2018.
- FIPR05. Michael J. Freedman, Yuval Ishai, Benny Pinkas, and Omer Reingold. Keyword search and oblivious pseudorandom functions. In *TCC*, volume 3378 of *LNCS*, pages 303–324. Springer, 2005.
- FS86. Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *CRYPTO*, volume 263 of *LNCS*, pages 186–194. Springer, 1986.
- GKK⁺19. Lorenzo Grassi, Daniel Kales, Dmitry Khovratovich, Arnab Roy, Christian Rechberger, and Markus Schofnegger. Starkad and poseidon: New hash functions for zero knowledge proof systems. *ePrint*, 2019:458, 2019.
- GLSV14. Vincent Grosso, Gaëtan Leurent, François-Xavier Standaert, and Kerem Varici. Ls-designs: Bitslice encryption for efficient masked software implementations. In *FSE*, volume 8540 of *LNCS*, pages 18–37. Springer, 2014.
- GMO16. Irene Giacomelli, Jesper Madsen, and Claudio Orlandi. Zkboo: Faster zero-knowledge for boolean circuits. In *USENIX Security Symposium*, pages 1069–1083. USENIX Association, 2016.
- GRR⁺16. Lorenzo Grassi, Christian Rechberger, Dragos Rotaru, Peter Scholl, and Nigel P. Smart. Mpc-friendly symmetric key primitives. In *ACM CCS*, pages 430–443. ACM, 2016.
- Gün11. Tim Güneysu. Utilizing hard cores of modern FPGA devices for high-performance cryptography. *J. Cryptographic Engineering*, 1(1):37–55, 2011.
- HOKG18. James Howe, Tobias Oder, Markus Krausz, and Tim Güneysu. Standard lattice-based key encapsulation on embedded devices. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(3):372–393, 2018.
- IKOS07. Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from secure multiparty computation. In *STOC*, pages 21–30. ACM, 2007.
- Int10. Intel Corporation. Securing the enterprise with intel® aes-ni, 2010.

- KKW18. Jonathan Katz, Vladimir Kolesnikov, and Xiao Wang. Improved non-interactive zero knowledge with applications to post-quantum signatures. In *ACM CCS*, pages 525–537. ACM, 2018.
- KRS⁺19. Daniel Kales, Christian Rechberger, Thomas Schneider, Matthias Senker, and Christian Weinert. Mobile private contact discovery at scale. In *USENIX Security Symposium*, pages 1447–1464. USENIX Association, 2019.
- Low. LowMC. Official LowMC Github Repository.
- MPL⁺11. Amir Moradi, Axel Poschmann, San Ling, Christof Paar, and Huaxiong Wang. Pushing the limits: A very compact and a threshold implementation of AES. In *EUROCRYPT*, volume 6632 of *LNCS*, pages 69–88. Springer, 2011.
- MRR19. Payman Mohassel, Peter Rindal, and Mike Rosulek. Fast database joins for secret shared data. *ePrint*, 2019:518, 2019.
- NLV11. Michael Naehrig, Kristin E. Lauter, and Vinod Vaikuntanathan. Can homomorphic encryption be practical? In *CCSW*, pages 113–124. ACM, 2011.
- PDG14. Thomas Pöppelmann, Léo Ducas, and Tim Güneysu. Enhanced lattice-based signatures on reconfigurable hardware. In *CHES*, volume 8731 of *LNCS*, pages 353–370. Springer, 2014.
- RM19. Debapriya Basu Roy and Debdeep Mukhopadhyay. Post quantum ECC on FPGA platform. *ePrint*, 2019:568, 2019.
- RSS17. Dragos Rotaru, Nigel P. Smart, and Martijn Stam. Modes of operation suitable for computing on encrypted data. *IACR Trans. Symmetric Cryptol.*, 2017(3):294–324, 2017.
- RST18. Christian Rechberger, Hadi Soleimany, and Tyge Tiessen. Cryptanalysis of low-data instances of full lowmcv2. *IACR Trans. Symmetric Cryptol.*, 2018(3):163–181, 2018.
- SA14. Ismail San and Nuray At. Improving the computational efficiency of modular operations for embedded systems. *Journal of Systems Architecture - Embedded Systems Design*, 60(5):440–451, 2014.
- SG15. Pascal Sasdrich and Tim Güneysu. Implementing curve25519 for side-channel-protected elliptic curve cryptography. *TRETS*, 9(1):3:1–3:15, 2015.
- Unr15. Dominique Unruh. Non-interactive zero-knowledge proofs in the quantum random oracle model. In *EUROCRYPT (2)*, volume 9057 of *LNCS*, pages 755–784. Springer, 2015.
- WSN18. Wen Wang, Jakub Szefer, and Ruben Niederhagen. Fpga-based niederreiter cryptosystem using binary goppa codes. In *PQCrypto*, volume 10786 of *LNCS*, pages 77–98. Springer, 2018.
- WU. Mario Werner and Thomas Unterluggauer. Transparent memory encryption and authentication.
- WUS⁺17. Mario Werner, Thomas Unterluggauer, Robert Schilling, David Schaffner, and Stefan Mangard. Transparent memory encryption and authentication. In *FPL*, pages 1–6. IEEE, 2017.