

The Picnic Signature Scheme

Design Document

Melissa Chase, David Derler, Steven Goldfeder,
Claudio Orlandi, Sebastian Ramacher, Christian Rechberger,
Daniel Slamanig, Greg Zaverucha

September 29, 2017
DRAFT Version 0.8

Contents

1	Introduction	3
1.1	The Picnic Design Team	3
2	Building Blocks	4
2.1	Commitments	4
2.2	Zero-Knowledge Proofs and Σ -Protocols	5
2.3	Non-interactive Zero-Knowledge Proofs of Knowledge	6
2.4	Fiat-Shamir Transform	7
2.5	Unruh Transform	8
2.6	(2,3)-Decomposition of Circuits	9
2.7	ZKB++	12
2.8	LowMC	17
2.8.1	Comparison to Other Primitives	18
2.9	Signature Schemes	19
3	The <code>picnic-FS</code> and <code>picnic-UR</code> Signature Schemes	21
3.1	Instantiation and Optimizations of Unruh’s Transform	21
3.2	Seed Generation	23
3.3	Random Tapes	23
3.4	Challenge Generation	23
3.5	Function G	23
4	Choice of Parameters	24
4.1	LowMC Parameters	24
4.2	Number of Parallel Repetitions	24
5	Formal Security Analysis	25
5.1	Security Analysis of ZKB++	25
5.2	Security Analysis of Picnic	26
5.2.1	Security Analysis of <code>picnic-FS</code>	27
5.2.2	Security Analysis of <code>picnic-UR</code>	29
5.3	Strong Unforgeability of <code>picnic-FS</code> and <code>picnic-UR</code>	36
6	Analysis with Respect to Known Attacks	37
6.1	Usage and security margin of LowMC	38
6.2	Attacks in the Single-User Setting	39
6.3	Attacks in the Multi-User Setting	39

7	Expected Security Strength	40
7.1	LowMC Parameter Selection	42
7.2	Hash Function Security	43
8	Advantages and Limitations	44
8.1	Compatibility with Existing Protocols	44
9	Additional Security Properties	44
9.1	Side-Channel Attacks	44
9.2	Security Impact of Using Weak Ephemeral Values	46
9.3	Parameter Integrity	47
10	Efficiency and Memory Usage	47
10.1	Description of the Benchmark Platforms	48
10.1.1	Platform A	48
10.1.2	Platform B	48
10.1.3	Platform C	49
10.2	Description of the benchmarking methodology	49
10.3	Benchmark Results: Timings	49
10.4	Memory Requirements	52
10.4.1	Reference Implementation Detailed Memory usage . . .	52
10.4.2	Optimized Implementation Detailed Memory usage . .	56
10.5	Size of precomputed constants and data	56

1 Introduction

Picnic is a signature scheme which is designed to provide security against attacks by quantum computers, in addition to attacks by classical computers. The building blocks are a zero-knowledge proof system (with post-quantum security), and symmetric key primitives like hash functions and block ciphers, with well-understood post-quantum security. In particular, Picnic does not require number-theoretic, or structured hardness assumptions.

In this document we first present the building blocks of the Picnic signature scheme in Section 2. Second, in Section 3 we present two variants of the Picnic signature scheme and various optimizations to the building blocks that we can employ in our schemes. Third we specify the parameters for some building blocks in Section 4 (and leave complete details of the parameters to the specification document). In the subsequent Section 5 of this document we include the formal security proofs for the proposed instantiations of Picnic. Section 6 presents an analysis of the algorithm with respect to known attacks and Section 7 provides a thorough description of the expected security strength. Finally, Section 8 discuss advantages and limitations, Section 9 discusses additional security properties and Section 10 presents results on efficiency and memory usage of the Picnic scheme.

1.1 The Picnic Design Team

Picnic was designed collaboratively by the following group of people.

Melissa Chase, Microsoft

David Derler, Graz University of Technology

Steven Goldfeder, Princeton

Claudio Orlandi, Aarhus University

Sebastian Ramacher, Graz University of Technology

Christian Rechberger, Graz University of Technology & DTU

Daniel Slamanig, AIT Austrian Institute of Technology

Greg Zaverucha, Microsoft

2 Building Blocks

2.1 Commitments

Formally a (non-interactive) commitment scheme consists of three algorithms KG , Com , Open with the following properties:

$\text{KG}(1^\kappa)$: The key generation algorithm, on input the security parameter κ it outputs a public key pk (we henceforth assume pk to be an implicit input to the subsequent algorithms).

$\text{Com}(M)$: On input of a message $M \in \{0, 1\}$, the commitment algorithm outputs $(C, D) \leftarrow \text{Com}(M; R)$, where R is a random value used when forming the commitment. C is the commitment string, while D is the decommitment string which is kept secret until opening time.

$\text{Open}(C, D)$: On input C, D , the verification algorithm either outputs a message M or \perp .

Computationally secure commitments must satisfy the following properties

Correctness If $\text{Com}(M)$ outputs (C, D) then $\text{Open}(C, D) = M$.

Hiding For every message pair M, M' the probability ensembles $\{C : (C, D) \leftarrow \text{Com}(M)\}_{\kappa \in \mathbb{N}}$ and $\{C : (C, D) \leftarrow \text{Com}(M')\}_{\kappa \in \mathbb{N}}$ are computationally indistinguishable for security parameter κ .

Binding We say that an adversary \mathcal{A} wins if it outputs C, D, D' such that $\text{Open}(C, D) = M$, $\text{Open}(C, D') = M'$ and $M \neq M'$. We require that for all efficient algorithms \mathcal{A} (running in time polynomial in κ), the probability that \mathcal{A} wins is a negligible function of κ .

Our implementation uses hash-based commitments, which requires modeling the hash function as a random oracle in our security analysis. Let H be a cryptographic hash function. The commitment scheme works as follows:

$\text{Com}(M)$: Sample $R \xleftarrow{R} \{0, 1\}^\kappa$ and set $C \leftarrow H(R, M)$ and return $(C, (R, M))$;

$\text{Open}(C, D)$: Parse D as (R, M) and return M if $H(R, M) = C$, and return \perp otherwise.

2.2 Zero-Knowledge Proofs and Σ -Protocols

A sigma protocol is a three-flow protocol between a prover and verifier, used to prove knowledge of a secret. A well-known class of sigma protocols are the so-called generalized Schnorr proofs, which allow the prover to prove knowledge of a discrete logarithm, and that it satisfies certain properties. In the present work we use a sigma protocol that allows one to prove knowledge of an input to an arbitrary binary circuit. Sigma protocols are usually zero-knowledge proofs, which informally means that the proof protocol does not reveal any information about the secret. We describe interactive protocols, but will show later how to make them non-interactive (so that signatures are non-interactive). Let L be an **NP**-language with associated witness relation R so that $L = \{x \mid \exists w : R(x, w) = 1\}$. A Σ -protocol for language L is defined as follows.

Definition 2.1. A Σ -protocol for language L is an interactive three-move protocol between a PPT prover $P = (\text{Commit}, \text{Prove})$ and a PPT verifier $V = (\text{Challenge}, \text{Verify})$, where P makes the first move and transcripts are of the form $(\mathbf{a}, \mathbf{e}, \mathbf{z}) \in A \times E \times Z$, where \mathbf{a} is output by **Commit**, \mathbf{e} is output by **Challenge** and \mathbf{z} is output by **Prove**. Additionally, Σ protocols satisfy the following properties

Completeness. A Σ -protocol for language L is complete, if for all security parameters κ , and for all $(x, w) \in R$, it holds that

$$\Pr[\langle P(1^\kappa, x, w), V(1^\kappa, x) \rangle = 1] = 1.$$

s -Special Soundness. A Σ -protocol for language L is s -special sound, if there exists a PPT extractor E so that for all x , and for all sets of accepting transcripts $\{(\mathbf{a}, \mathbf{e}_i, \mathbf{z}_i)\}_{i \in [s]}$ with respect to x , generated by any algorithm with polynomial runtime in κ , it holds that

$$\Pr \left[w \leftarrow E(1^\kappa, x, \{(\mathbf{a}, \mathbf{e}_i, \mathbf{z}_i)\}_{i \in [s]}) \quad : \right. \\ \left. \begin{array}{l} (x, w) \in R \wedge \\ \forall i, j \in [s], i \neq j : \mathbf{e}_i \neq \mathbf{e}_j \end{array} \right] \geq 1 - \epsilon(\kappa).$$

Special Honest-Verifier Zero-Knowledge. A Σ -protocol is special honest-verifier zero-knowledge, if there exists a PPT simulator S so that for

every $x \in L$ and every challenge \mathbf{e} from the challenge space, it holds that a transcript $(\mathbf{a}, \mathbf{e}, \mathbf{z})$, where $(\mathbf{a}, \mathbf{z}) \leftarrow \mathcal{S}(1^\kappa, x, \mathbf{e})$ is computationally indistinguishable from a transcript resulting from an honest execution of the protocol.

The s -special soundness property gives an immediate bound for the soundness of the protocol: if no witness exists then (ignoring a negligible error) the prover can successfully answer at most to $(s - 1)/t$ challenges, where $t = |\mathbf{E}|$ is the size of the challenge space. In case this value is too large, it is possible to reduce the soundness error using well known properties of Σ -protocols which we restate here for completeness (see [Dam10, CDS94] for details).

Lemma 2.2. *The properties of Σ -protocols are invariant under parallel repetition. In particular, the ℓ -fold parallel repetition of a Σ -protocol for relation R with challenge length t yields a new Σ -protocol with challenge length ℓt .*

2.3 Non-interactive Zero-Knowledge Proofs of Knowledge

For our signatures, we use non-interactive zero-knowledge proofs of knowledge, in which the proof is a single message. Here we define zero knowledge and the necessary notion of proof of knowledge. We present the definition in the random oracle model against classical adversaries and the definition in the quantum random oracle model against quantum adversaries.¹

Zero knowledge says that there is a simulator which can produce proofs that are indistinguishable from those produced by P without knowing the witnesses to an adversary who is given access to a simulated version of the random oracle.

Definition 2.3. A protocol P, V for relation R is a zero knowledge against (quantum) adversaries in the (quantum) random oracle model if there exist PPT algorithms $\mathcal{S}_R, \text{Sim}$ such that for all (quantum) PPT adversaries \mathcal{A}

$$|\Pr[b \leftarrow \mathcal{A}^{R(\cdot), P(\cdot, \cdot)}(1^\lambda) : b = 1] - \Pr[b \leftarrow \mathcal{A}^{\mathcal{S}_R(\cdot), \text{Sim}_P(\cdot, \cdot)}(1^\lambda) : b = 1]|$$

is negligible in λ , where R is a random function to which \mathcal{A} can provide (quantum) inputs, Sim_P takes a pair $(x, w) \in R$ and calls $\text{Sim}(x)$, and $\text{Sim}, \mathcal{S}_R$ share state.

¹These definition roughly combine the ROM definitions of [BPW12] and the QROM definitions of [Unr15].

Simulation sound extractability says that an adversary cannot produce a new proof for a statement for which he does not know the witness, even if he is allowed to see proofs produced by someone else for statements of his choice. More formally, we say that even when an adversary is given access to the zero knowledge simulator to obtain proofs for (true or false) statements of its choice, whenever it produces a new proof (not produced by the simulator) that is accepted by the verifier, there is an extractor that can look at the implementation of the adversary and extract a valid witness for that statement.

Definition 2.4. A protocol P, V for relation R satisfies simulation-sound extractability against (quantum) adversaries in the (quantum) random oracle model if there exist PPT algorithms S_R, Sim satisfying the zero-knowledge definition and a PPT (quantum) extractor such that for all (quantum) PPT adversaries \mathcal{A}

$$\Pr[(x, \pi) \leftarrow \mathcal{A}^{S_R(\cdot), \text{Sim}(\cdot)}(1^\lambda); w \leftarrow E(\mathcal{A}, x, \pi) : \\ V^{S_R}(x, \pi) = 1 \wedge (x, \pi) \notin \mathcal{Q} \wedge (x, w) \notin R]$$

is negligible in λ , where S_R, Sim , and E share state, \mathcal{Q} is the list of \mathcal{A} 's queries to Sim and the resulting responses, and passing \mathcal{A} as input to E means E is given access to a (quantum) implementation of \mathcal{A} .

2.4 Fiat-Shamir Transform

The Fiat-Shamir (FS) transform [FS86] is an elegant way to construct signature schemes from Σ -protocols. The basic idea is similar to constructing NIZK proofs from Σ -protocols, but the challenge e is generated by hashing the prover's first message \mathbf{a} and the message m to be signed, i.e., define a modified challenge algorithm $\text{Challenge}'$ that outputs $e \leftarrow H(\mathbf{a}, m)$. Then, the prover can locally obtain the challenge after computing the initial message. Starting a verifier $V' = (\text{Challenge}', \text{Verify})$ on the same initial message would yield the same challenge. The prover outputs (\mathbf{a}, \mathbf{z}) as the challenge.

More formally, by using the hash function $H : \mathbf{A} \times \mathbf{X} \rightarrow \mathbf{E}$, which we model as a random oracle, we obtain the non-interactive PPT algorithms $(\text{Prove}_H, \text{Verify}_H)$, defined as follows:

$\text{Prove}_H(1^\kappa, (x, m), w)$: Start P on input $(1^\kappa, x, w)$, obtain the first message \mathbf{a} , answer with $e \leftarrow H(\mathbf{a}, m)$, and finally obtain \mathbf{z} . Return $\pi \leftarrow (\mathbf{a}, \mathbf{z})$.

$\text{Verify}_H(1^\kappa, (x, m), \pi)$: Parse π as (\mathbf{a}, \mathbf{z}) . Start \mathbf{V}' on $(1^\kappa, x)$, send (\mathbf{a}, m) as the first message to the verifier. When \mathbf{V}' outputs \mathbf{e} , reply with \mathbf{z} and output 1 if \mathbf{V}' accepts and 0 otherwise.

2.5 Unruh Transform

Similar to the Fiat-Shamir transform, Unruh's transform [Unr12, Unr15, Unr16] allows one to construct NIZK proofs and signature schemes from Σ -protocols. In contrast to the FS transform, Unruh's transform can be proven secure in the QROM (quantum random oracle model), strengthening the security guarantee against quantum adversaries.

At a high level, Unruh's transform works as follows: Given a 2-special-sound Σ -protocol, integers t and M , a statement x and a random permutation G , the prover will repeat the first phase of the Σ -protocol t times. Then, for each of the t runs, it produces proofs to M different randomly selected challenges. All so-obtained responses are then permuted using G . The prover then selects the responses to publish for each round of the Σ -protocol by querying the random oracle on the message to be signed, all first rounds of the Σ -protocol and all permuted responses.

For a more formal treatment, we keep $t \in \mathbb{N}$ as above and let $M \in [2, |\mathbf{E}|]$. Let $H : \{0, 1\}^* \rightarrow [M]^t$ be a hash function we model as a random oracle. We obtain the non-interactive PPT algorithms $(\text{Prove}_H, \text{Verify}_H)$ defined as follows:

$\text{Prove}_H(1^\kappa, (x, m), w)$:

1. For $i \in [t]$:
 - (a) Start \mathbf{P} on $(1^\kappa, x, w)$ and obtain first message \mathbf{a}_i .
 - (b) For $j \in [M]$, set $\mathbf{e}_{i,j} \xleftarrow{R} \mathbf{E} \setminus \{\mathbf{e}_{i,1}, \dots, \mathbf{e}_{i,j-1}\}$ and obtain response $\mathbf{z}_{i,j}$ for challenge $\mathbf{e}_{i,j}$.
2. For $i, j \in [t] \times [M]$, set $g_{i,j} \leftarrow G(\mathbf{z}_{i,j})$.
3. Let $(J_1, \dots, J_t) \leftarrow H(m, (\mathbf{a}_i)_{i \in [t]}, (\mathbf{e}_{i,j})_{(i,j) \in [t] \times [M]}, (g_{i,j})_{(i,j) \in [t] \times [M]})$
4. Return $\pi \leftarrow ((\mathbf{a}_i)_{i \in [t]}, (\mathbf{e}_{i,j})_{(i,j) \in [t] \times [M]}, (g_{i,j})_{(i,j) \in [t] \times [M]}, (\mathbf{z}_{i,J_i})_{i \in [t]})$

$\text{Verify}_H(1^\kappa, (x, m), \pi)$: Parse π as $((\mathbf{a}_i)_{i \in [t]}, (\mathbf{e}_{i,j})_{(i,j) \in [t] \times [M]}, (g_{i,j})_{(i,j) \in [t] \times [M]}, (\mathbf{z}_i)_{i \in [t]})$.

1. Let $(J_1, \dots, J_t) \leftarrow H(m, (\mathbf{a}_i)_{i \in [t]}, (\mathbf{e}_{i,j})_{(i,j) \in [t] \times [M]}, (g_{i,j})_{(i,j) \in [t] \times [M]})$
2. For $i \in [t]$ check that all $\mathbf{e}_{i,1}, \dots, \mathbf{e}_{i,M}$ are pairwise distinct.
3. For $i \in [t]$ check whether V accepts the proof with respect to x , first message \mathbf{a}_i , challenge \mathbf{e}_{i,J_i} and response \mathbf{z}_i .
4. For $i \in [t]$ check $g_{i,J_i} = G(\mathbf{z}_i)$.
5. Output 1 if all checks succeeded and 0 otherwise.

We discuss a specialization of Unruh's transform to our Σ -protocol in Section 3.1.

2.6 (2,3)-Decomposition of Circuits

A circuit decomposition is a protocol for jointly computing a circuit, similar to an MPC protocol, but with greater efficiency. In a (2,3)-decomposition there are three players and the protocol has 2-privacy, i.e., it remains secure even if two of the three players are corrupted.

Definition 2.5 ((2,3)-decomposition). Let $f(\cdot)$ be a function that is computed by an n -gate circuit ϕ such that $f(x) = \phi(x) = y$. Let k_1, k_2 , and k_3 be tapes of length κ chosen uniformly at random from $\{0, 1\}^\kappa$ corresponding to players P_1, P_2 and P_3 , respectively. Consider the following set of functions, \mathcal{D} :

$$\begin{aligned}
 (\mathbf{view}_1^{(0)}, \mathbf{view}_2^{(0)}, \mathbf{view}_3^{(0)}) &\leftarrow \text{Share}(x, k_1, k_2, k_3) \\
 \mathbf{view}_i^{(j+1)} &\leftarrow \text{Update}(\mathbf{view}_i^{(j)}, \mathbf{view}_{i+1}^{(j)}, k_i, k_{i+1}) \\
 y_i &\leftarrow \text{Output}(\mathbf{View}_i) \\
 y &\leftarrow \text{Reconstruct}(y_1, y_2, y_3)
 \end{aligned}$$

such that **Share** is a potentially randomized invertible function that takes x as input and outputs the initial view for each player containing the secret share of x_i of x - i.e. $\mathbf{view}_i^{(0)} = x_i$. The function **Update** computes the wire values for the next gate and updates the view accordingly. The function **Output** _{i} takes as input the final view, $\mathbf{View}_i \equiv \mathbf{view}_i^{(n)}$ after all gates have been computed and outputs player P_i 's *output share*, denoted y_i .

Correctness requires that reconstructing a (2,3)-decomposed evaluation of a circuit ϕ yields the same value as directly evaluating ϕ on the input value.

The 2-privacy property requires that revealing the values from two shares reveals nothing about the input value. More formally, these two properties are defined as follows: We define the experiment $\text{EXP}_{\text{decomp}}^{(\phi,x)}$ in Experiment 1, which runs the decomposition over a circuit ϕ on input x : We say that \mathcal{D} is

$\text{EXP}_{\text{decomp}}^{(\phi,x)}$:

1. First run the **Share** function on x : $\text{view}_1^{(0)}, \text{view}_2^{(0)}, \text{view}_3^{(0)} \leftarrow \text{Share}(x, k_1, k_2, k_3)$
2. For each of the three views, call the update function successively for every gate in the circuit: $\text{view}_i^{(j)} = \text{Update}(\text{view}_i^{(j-1)}, \text{view}_{i+1}^{(j-1)}, k_i, k_{i+1})$ for $i \in [1, 3], j \in [1, n]$
3. From the final views, compute the output share of each view: $y_i \leftarrow \text{output}(\text{View}_i)$

Experiment 1: Decomposition Experiment

a $(2, 3)$ -decomposition of ϕ if the following two properties hold when running $\text{EXP}_{\text{decomp}}^{(\phi,x)}$:

Correctness: For all circuits ϕ , for all inputs x and for the y_i 's produced by , for all circuits ϕ , for all inputs x ,

$$\Pr[\phi(x) = \text{Reconstruct}(y_1, y_2, y_3)] = 1.$$

2-Privacy: Let \mathcal{D} be correct. Then for all $e \in \{1, 2, 3\}$ there exists a PPT simulator \mathcal{S}_e such that for any probabilistic polynomial-time (PPT) algorithm \mathcal{A} , for all circuits ϕ , for all inputs x , and for the distribution of views and k_i 's produced by $\text{EXP}_{\text{decomp}}^{(\phi,x)}$ we have that

$$\left| \Pr[\mathcal{A}(x, y, k_e, \text{View}_e, k_{e+1}, \text{View}_{e+1}, y_{e+2}) = 1] - \Pr[\mathcal{A}(x, y, \mathcal{S}_e(\phi, y)) = 1] \right|$$

is negligible.

We now discuss the $(2,3)$ -decomposition used by ZKB++. 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 by constant,

multiplication by constant, binary addition and binary multiplication gates. A $(2, 3)$ -decomposition of ϕ is given by the following functions. In the notation below, arithmetic operations are done in R^s where the operands are elements of R^s):

- $(x_1, x_2, x_3) \leftarrow \text{Share}(x, k_1, k_2, k_3)$ samples random $x_1, x_2, x_3 \in R^m$ such that $x_1 + x_2 + x_3 = x$.
- $y_i \leftarrow \text{Output}_i(\text{view}_i^{(n)})$ selects the ℓ output wires of the circuit as stored in the view $\text{view}_i^{(n)}$.
- $y \leftarrow \text{Reconstruct}(y_1, y_2, y_3) = y_1 + y_2 + y_3$
- $\text{view}_i^{(j+1)} \leftarrow \text{Update}_i^{(j)}(\text{view}_i^{(j)}, \text{view}_{i+1}^{(j)}, k_i, k_{i+1})$ computes P_i 's view of the output wire of gate g_j and appends it to the view. Notice that it takes as input the views and random tapes of both party P_i as well as party P_{i+1} . We use w_k to refer to the k -th wire, and we use $w_k^{(i)}$ to refer to the value of w_k in party P_i 's view. The update operation depends on the type of gate g_j .

The gate-specific operations are defined as follows.

Addition by Constant ($w_b = w_a + k$)

$$w_b^{(i)} = \begin{cases} w_a^{(i)} + k & \text{if } i = 1, \\ w_a^{(i)} & \text{otherwise.} \end{cases}$$

Multiplication by Constant ($w_b = w_a \cdot k$)

$$w_b^{(i)} = k \cdot w_a^{(i)}$$

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 pseudorandom generator seeded with k_i .

Note that with the exception of the constant addition gate, the gates are symmetric for all players. Also note that P_i can compute all gate types locally with the exception of binary multiplication gates as this requires inputs from P_{i+1} . In other words, for every operation except binary multiplication, the `Update` function does not use the inputs from the second party, i.e., $\text{view}_{i+1}^{(j)}$ and k_{i+1} .

While we do not give the details here, [GMO16a] shows that this decomposition meets the correctness and 2-privacy requirements of Definition 2.5. In other words, for every operation except binary multiplication, the `Update` function does not use the inputs from the second party, i.e., $\text{view}_{i+1}^{(j)}$ and R_{i+1} .

2.7 ZKB++

ZKB++, an optimized version of ZKBOO [GMO16a], is a proof system for zero-knowledge proofs on arbitrary circuits. ZKBOO and ZKB++ build on the MPC-in-the-head paradigm of Ishai *et al.* [IKOS09], that we describe only informally here. The multiparty computation protocol (MPC) will implement the relation, and the input is the witness. For example, the MPC could compute $y = \text{SHA-256}(x)$ where players each have a share of x and y is public. The idea is to have the prover simulate a multiparty computation protocol “in their head”, commit to the state and transcripts of all players, then have the verifier “corrupt” a random subset of the simulated players by seeing their complete state. The verifier then checks that the computation was done correctly from the perspective of the corrupted players, and if so, he has some assurance that the output is correct. Iterating this for many rounds then gives the verifier high assurance.

ZKBOO generalizes the idea of [IKOS09] by replacing MPC with circuit decompositions. In Scheme 1 and Scheme 2 we present the prover and the verifier of the ZKB++ Σ -protocol.

We now discuss various instantiation aspects of ZKB++ that make the optimizations with respect to ZKBOO possible. To highlight the difference, we also present the Fiat-Shamir transformed ZKBOO proof system in Scheme 3 and the Fiat-Shamir transformed ZKB++ in Scheme 4.

The Share Function. We make the `Share` function sample the shares

P.Commit : 1. For each iteration $i \in [t]$: Sample random tapes $k_1^{(i)}, k_2^{(i)}, k_3^{(i)}$ obtain view $\text{View}_j^{(i)}$ and output share $y_j^{(i)}$. For each player P_j compute

- (a) $(x_1^{(i)}, x_2^{(i)}, x_3^{(i)}) \leftarrow \text{Share}(x, k_1^{(i)}, k_2^{(i)}, k_3^{(i)})$
- (b) $\text{View}_j^{(i)} \leftarrow \text{Update}(\dots \text{Update}(x_j^{(i)}, x_{j+1}^{(i)}, k_j^{(i)}, k_{j+1}^{(i)}) \dots)$
- (c) $y_j^{(i)} \leftarrow \text{Output}(\text{View}_j^{(i)})$
- (d) Commit $C_j^{(i)} \leftarrow \text{Com}(k_j^{(i)}, x_j^{(i)}, \text{View}_j^{(i)}, y_j^{(i)})$, and let $\mathbf{a}^{(i)} \leftarrow (y_1^{(i)}, y_2^{(i)}, y_3^{(i)}, C_1^{(i)}, C_2^{(i)}, C_3^{(i)})$.

2. Return $(\mathbf{a}^{(i)})_{i \in [t]}$.

P.Prove : On input of a challenge $(\mathbf{e}^{(i)})_{i \in [t]}$, set for each iteration $i \in [1, t]$

$$\mathbf{z}^{(i)} \leftarrow \begin{cases} (\text{View}_2^{(i)}, k_1^{(i)}, k_2^{(i)}) & \text{if } \mathbf{e}^{(i)} = 1, \\ (\text{View}_3^{(i)}, k_2^{(i)}, k_3^{(i)}, x_3^{(i)}) & \text{if } \mathbf{e}^{(i)} = 2, \\ (\text{View}_1^{(i)}, k_3^{(i)}, k_1^{(i)}, x_3^{(i)}) & \text{if } \mathbf{e}^{(i)} = 3. \end{cases}$$

and return $(\mathbf{z}^{(i)})_{i \in [t]}$.

Scheme 1: The prover of the ZKB++ Σ -protocol.

pseudorandomly as:

$$(x_1, x_2, x_3) \leftarrow \text{Share}(x, k_1, k_2, k_3) := \\ x_1 = R_1(0), \quad x_2 = R_2(0), \quad x_3 = x - x_1 - x_2.$$

R_i is a pseudorandom generator seeded with k_i . We specify the **Share** function in this manner as it will lead to more compact proofs. Moving now to the ZKBOO protocol, for each round, the prover is required to “open” two views. In order to verify the proof, the verifier must be given both the random tape and the input share for each opened view. If these values are generated independently of one another, then the prover will have to explicitly include both of them in the proof. However, with our sampling method, in View_1 and View_2 , the prover only needs to include k_i , as x_i can be deterministically computed by the verifier.

The exact savings depends on which views the prover must open, and thus depends on the challenge.

V.Challenge : Store $(\mathbf{a}^{(i)})_{i \in [t]}$ and return $(\mathbf{e}^{(i)})_{i \in [t]} \leftarrow^R \mathbf{E}^t$.

V.Verify : 1. For each iteration $i \in [t]$ reconstruct the views, input and output shares that were not explicitly given as part of the proof response $\mathbf{z}^{(i)}$:

(a) Set

$$x_{\mathbf{e}^{(i)}}^{(i)} \leftarrow \begin{cases} R_{\mathbf{e}^{(i)}}(0) & \text{if } \mathbf{e}^{(i)} \neq 3, \\ x_3^{(i)} \text{ given as part of } \mathbf{z}^{(i)} & \text{if } \mathbf{e}^{(i)} = 3. \end{cases}$$

$$x_{\mathbf{e}^{(i)}+1}^{(i)} \leftarrow \begin{cases} R_{\mathbf{e}^{(i)}+1}(0) & \text{if } \mathbf{e}^{(i)} \neq 2, \\ x_3^{(i)} \text{ given as part of } \mathbf{z}^{(i)} & \text{if } \mathbf{e}^{(i)} = 2. \end{cases}$$

(b) Obtain $\text{View}_{\mathbf{e}^{(i)}+1}^{(i)}$ from $\mathbf{z}^{(i)}$.

(c) $\text{View}_{\mathbf{e}^{(i)}}^{(i)} \leftarrow \text{Update}(\dots \text{Update}(x_{\mathbf{e}^{(i)}}^{(i)}, x_{\mathbf{e}^{(i)}+1}^{(i)}, k_{\mathbf{e}^{(i)}}^{(i)}, k_{\mathbf{e}^{(i)}+1}^{(i)}) \dots)$

(d) $y_{\mathbf{e}^{(i)}}^{(i)} \leftarrow \text{Output}(\text{View}_{\mathbf{e}^{(i)}}^{(i)})$

(e) $y_{\mathbf{e}^{(i)}+1}^{(i)} \leftarrow \text{Output}(\text{View}_{\mathbf{e}^{(i)}+1}^{(i)})$

(f) $y_{\mathbf{e}^{(i)}+2}^{(i)} \leftarrow y + y_{\mathbf{e}^{(i)}}^{(i)} + y_{\mathbf{e}^{(i)}+1}^{(i)}$

2. Re-compute the commitments for views $\text{View}_{\mathbf{e}^{(i)}}^{(i)}$ and $\text{View}_{\mathbf{e}^{(i)}+1}^{(i)}$. For $j \in \{\mathbf{e}^{(i)}, \mathbf{e}^{(i)} + 1\}$:

$$C_j^{(i)} \leftarrow \text{Com}(k_j^{(i)}, x_j^{(i)}, \text{View}_j^{(i)}, y_j^{(i)})$$

3. Set $\mathbf{a}'^{(i)} \leftarrow (y_1^{(i)}, y_2^{(i)}, y_3^{(i)}, C_1^{(i)}, C_2^{(i)}, C_3^{(i)})$ taking $C_{\mathbf{e}^{(i)}+2}^{(i)}$ from $\mathbf{a}^{(i)}$.

4. If $\mathbf{a}'^{(i)} = \mathbf{a}^{(i)}$ for all $i \in [t]$, output **Accept**, otherwise **Reject**.

Scheme 2: The verifier of the ZKB++ Σ -protocol.

Not Including Input Shares. Since the input shares are generated pseudorandomly using the seed k_i , we do not need to include them in the view when $\mathbf{e} = 1$. However, if $\mathbf{e} = 2$ or $\mathbf{e} = 3$, we still need to send one input share for the third view for which the input share cannot be derived from the seed. Thus we explicitly specify the input share when required and do not include it in $\text{View}_i^{(j)}$.

Prove_H($1^\kappa, y, x$): 1. For each iteration $i \in [1, t]$: Sample random tapes $k_1^{(i)}, k_2^{(i)}, k_3^{(i)}$ and run the decomposition to get an output view $\text{View}_j^{(i)}$ and output share $y_j^{(i)}$. In particular, for each player P_j :

- (a) $(x_1^{(i)}, x_2^{(i)}, x_3^{(i)}) \leftarrow \text{Share}(x, k_1^{(i)}, k_2^{(i)}, k_3^{(i)})$
- (b) $\text{View}_j^{(i)} \leftarrow \text{Update}(\dots \text{Update}(x_j^{(i)}, x_{j+1}^{(i)}, k_j^{(i)}, k_{j+1}^{(i)}) \dots)$
- (c) $y_j^{(i)} \leftarrow \text{Output}(\text{View}_j^{(i)})$
- (d) Commit $C_j^{(i)} \leftarrow \text{Com}(k_j^{(i)}, \text{View}_j^{(i)})$, and let $\mathbf{a}^{(i)} \leftarrow (y_1^{(i)}, y_2^{(i)}, y_3^{(i)}, C_1^{(i)}, C_2^{(i)}, C_3^{(i)})$.

2. Compute the challenge: $\mathbf{e} \leftarrow H(\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(t)})$. Interpret the challenge such that for $i \in [1, t]$, $\mathbf{e}^{(i)} \in \{1, 2, 3\}$

3. For each iteration $i \in [1, t]$, let $\mathbf{z}^{(i)} = (D_{\mathbf{e}^{(i)}}^{(i)}, D_{\mathbf{e}^{(i)}+1}^{(i)})$.

4. Output $\pi \leftarrow [(\mathbf{a}^{(1)}, \mathbf{z}^{(1)}), (\mathbf{a}^{(2)}, \mathbf{z}^{(2)}), \dots, (\mathbf{a}^{(t)}, \mathbf{z}^{(t)})]$

Verify_H($1^\kappa, y, \pi$): 1. Parse π as $[(\mathbf{a}^{(1)}, \mathbf{z}^{(1)}), (\mathbf{a}^{(2)}, \mathbf{z}^{(2)}), \dots, (\mathbf{a}^{(t)}, \mathbf{z}^{(t)})]$.

2. Compute the challenge: $\mathbf{e}' \leftarrow H(\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(t)})$. Interpret the challenge such that for $i \in [1, t]$, $\mathbf{e}'^{(i)} \in \{1, 2, 3\}$.

3. For each iteration $i \in [1, t]$: If there exists $j \in \{\mathbf{e}'^{(i)}, \mathbf{e}'^{(i)} + 1\}$ such that $\text{Open}(C_j^{(i)}, D_j^{(i)}) = \perp$, output Reject. Otherwise, for all $j \in \{\mathbf{e}'^{(i)}, \mathbf{e}'^{(i)} + 1\}$, set $\{k_j^{(i)}, \text{View}_j^{(i)}\} \leftarrow \text{Open}(C_j^{(i)}, D_j^{(i)})$.

4. For each iteration $i \in [1, t]$: If $\text{Reconstruct}(y_1^{(i)}, y_2^{(i)}, y_3^{(i)}) \neq y$, output Reject. If there exists $j \in \{\mathbf{e}'^{(i)}, \mathbf{e}'^{(i)} + 1\}$ such that $y_j^{(i)} \neq \text{Output}(\text{View}_j^{(i)})$, output Reject. For each wire value $w_j^{(\mathbf{e})} \in \text{View}_e$, if $w_j^{(\mathbf{e})} \neq \text{Update}(\text{view}_e^{(j-1)}, \text{view}_{e+1}^{(j-1)}, k_e, k_{e+1})$ output Reject.

5. Output Accept.

Scheme 3: The ZKBOO non-interactive proof system.

No Additional Randomness for Commitments. Since the first input to the commitment is the seed value k_i for the random tape, the protocol input to the commitment doubles as a randomization value, ensuring that commitments are hiding. Further, each view included in the commitment

must be well randomized for the security of the MPC protocol. In the random oracle model the resulting commitments are hiding (the RO model is needed here since k_i is used both as seed for the PRG and as randomness for the commitment. Since one already needs the RO model to make the proofs non-interactive, there is no extra assumption here). Hence the hash-based commitment scheme can be defined as:

$\text{Com}(M)$: Set $C \leftarrow H(M)$ and return (C, M) ;

$\text{Open}(C, D)$: Return M if $H(D) = C$, and return \perp otherwise.

Not Including the Output Shares. The output shares y_i are included in the proof as part of **a**. Moreover, for the two views that are opened, those output shares are included a second time. First, we do not need to send two of the output shares twice. We actually do not need to send any output shares at all as they can be deterministically computed from the rest of the proof as follows:

For the two views that are given as part of the proof, the output share can be recomputed from the remaining parts of the view. Essentially, the output share is just the value on the output wires. Given the random tapes and the communicated bits from the binary multiplication gates, all wires for both views can be recomputed.

For the third view, recall that the **Reconstruct** function simply adds the three output shares to obtain y . But the verifier is given y , and can thus instead recompute the third output share. In particular, given y_i , y_{i+1} and y , the verifier can compute: $y_{i+2} = y - y_i - y_{i+1}$. Thus we explicitly specify the output share when required and do not include it in $\text{View}_i^{(j)}$.

When applying the Fiat-Shamir and Unruh transforms to ZKB++ to obtain a signature scheme, we can also perform the following modifications.

Not Including Commitments. It is unnecessary to send all three commitments to the verifier. Since for the two views that are opened, the verifier can recompute the commitment. Only for the third view that the verifier is not given the commitment needs to be explicitly sent.

Security. One can observe that all optimizations except *"No Additional Randomness for Commitments"* are equivalence transformations, and, therefore, do not impact the security of the overall ZKB++ proof system. In Section 5.1, we formally confirm that using no additional randomness for the commitments does not impact the security of the ZKB++ proof system.

2.8 LowMC

LowMC [ARS⁺15, ARS⁺16] is a very parameterizable symmetric encryption scheme design enabling instantiation with low AND depth and low multiplicative complexity. Given any blocksize, a choice for the number of S-boxes per round, and security expectations in terms of time and data complexity, instantiations can be created minimizing the AND depth, the number of ANDs, or the number of ANDs per encrypted bit. Table 1 lists the choices for the parameters for security levels L1, L3, L5.

The description of LowMC is possible independently of the choice of parameters using a partial specification of the S-box and arithmetic in vector spaces over \mathbb{F}_2 . In particular, let n be the blocksize, m be the number of S-boxes, k the key size, and r the number of rounds, we choose round constants $C_i \leftarrow^R \mathbb{F}_2^n$ for $i \in [1, r]$, full rank matrices $K_i \leftarrow^R \mathbb{F}_2^{n \times k}$ and regular matrices $L_i \leftarrow^R \mathbb{F}_2^{m \times n}$ independently during the instance generation and keep them fixed. Keys for LowMC are generated by sampling from \mathbb{F}_2^k uniformly at random.

LowMC encryption starts with key whitening which is followed by several rounds of encryption. A single round of LowMC is composed of an S-box layer, a linear layer, addition with constants and addition of the round key, i.e.

$$\begin{aligned} \text{LOWMCROUND}(i) &= \text{KEYADDITION}(i) \\ &\quad \circ \text{CONSTANTADDITION}(i) \\ &\quad \circ \text{LINEARLAYER}(i) \circ \text{SBOXLAYER}. \end{aligned}$$

SBOXLAYER is an m -fold parallel application of the same 3-bit S-box on the first $3 \cdot m$ bits of the state. The S-box is defined as $S(a, b, c) = (a \oplus bc, a \oplus b \oplus ac, a \oplus b \oplus c \oplus ab)$.

The other layers only consist of \mathbb{F}_2 -vector space arithmetic. LINEARLAYER(i) multiplies the state with the linear layer matrix L_i , CONSTANTADDITION(i) adds the round constant C_i to the state, and KEYADDITION(i) adds the round key to the state, where the round key is generated by multiplying the master key with the key matrix K_i .

Algorithm 1 gives a full description of the encryption algorithm.

LowMC is very flexible in the choice of parameters: the block size n , the key size k , the number of 3-bit S-boxes m in the substitution layer and the allowed data complexity d of attacks can independently be chosen. To reduce

the multiplicative complexity, the number of S-boxes applied in parallel can be reduced, leaving part of the substitution layer as the identity mapping. The number of rounds r needed to achieve the goals is then determined as a function of all these parameters. We discuss concrete choices of the parameters in Section 4.1.

2.8.1 Comparison to Other Primitives

The signature size depends on constants that are close to the security expectation. The only exceptions are the number of binary multiplication gates, and the size of the ring, which both depend on the choice of the primitive. In this section we compare LowMC to existing standardized primitives and to other primitives with a low number of multiplications.

Standardized Primitives. The smallest known Boolean circuit for AES-128 needs 5440 AND gates, AES-192 needs 6528 AND gates, and AES-256 needs 7616 AND gates [BMP13]. An AES circuit in \mathbb{F}_{2^4} might be more efficient in our setting, as in this case the number of multiplications is lower than 1000 [CGP⁺12]. This results in an impact on the signature size that is equivalent to 4000 AND gates. Even though collision resistance is often not required, hash functions like SHA-256 are a popular choice for proof-of-concept implementations. The number of AND gates of a single call to the SHA-256 compression function is about 25000 and a single call to the permutation underlying SHA-3 is 38400.

Lightweight Ciphers. Most early designs in this domain focused on small area when implemented in hardware where an XOR gate is by a small factor larger than an AND or NAND gate. Notable designs with a low number of AND gates at the 128-bit security level are the block ciphers Noekeon [DPVAR00] (2048 ANDs) and Fantomas [GLSV14] (2112 ANDs). Furthermore, one should mention Prince [BCG⁺12] (1920 ANDs), or the stream cipher Trivium [DP08] (1536 AND gates to compute 128 output bits, with 80-bit security).

Custom Ciphers with a Low Number of Multiplications. Motivated by applications in SHE/FHE schemes, MPC protocols and SNARKs, recently a trend to design symmetric encryption primitives with a low number of multiplications or a low multiplicative depth started to evolve. This is a trend we can take advantage of.

We start with the LowMC [ARS⁺15] block cipher family. In the most

recent version of the design [ARS⁺16], the number of AND gates can be below 500 for 80-bit security, below 800 for 128-bit security, and below 1400 for 256-bit security. The stream cipher Kreyvium [CCF⁺16] needs similarly to Trivium 1536 AND gates to compute 128 output bits, but offers a higher security level of 128 bits. Even though FLIP [MJSC16] was designed to have especially low depth, it needs hundreds of AND gates per bit and is hence not competitive in our setting.

Last but not least there are the block ciphers and hash functions around MiMC [AGR⁺16] which need less than $2 \cdot s$ multiplications for s -bit security in a field of size close to 2^s . Note that MiMC is the only design in this category which aims at minimizing multiplications in a field larger than \mathbb{F}_2 . However, since the size of the signature depends on both the number of multiplications and the size of the field, this leads to a factor $2s^2$ which, for all arguably secure instantiations of MiMC, is already larger than the number of AND gates in the AES circuit.

LOWMC has two important advantages over other designs: It has the lowest number of AND gates for every security level: The closest competitor Kreyvium needs about twice as many AND gates and only exists for the 128-bit security level. The fact that it allows for an easy parameterization of the security level is another advantage. We hence use LOWMC for our concrete proposal.

Keccak is a family of cryptographic primitives including hash functions and extensible output functions (XOF). It was selected as the successor to SHA-2 and was standardized as SHA3 [NIS15]. SHAKE is an XOF constructed from SHA3. Since both SHA3 and SHAKE have a large number of AND gates (as described above), we do not use them for Picnic key generation. However, other parts of the ZKB++ protocol require a hash function. We will use SHAKE in two modes

1. as a hash function with a fixed output length
2. as a key derivation function, where we expand a fixed length seed into a larger pseudorandom value, by requesting larger, variable sized SHAKE outputs.

2.9 Signature Schemes

In the following we recall a standard definition of signature schemes along with two widely used security notions.

Definition 2.6. A signature scheme Σ is a triple $(\text{Gen}, \text{Sign}, \text{Verify})$ of PPT algorithms, which are defined as follows:

$\text{Gen}(1^\kappa)$: This algorithm takes a security parameter κ as input and outputs a secret (signing) key sk and a public (verification) key pk with associated message space \mathcal{M} (we may omit to make the message space \mathcal{M} explicit).

$\text{Sign}(\text{sk}, m)$: This algorithm takes a secret key sk and a message $m \in \mathcal{M}$ as input and outputs a signature σ .

$\text{Verify}(\text{pk}, m, \sigma)$: This algorithm takes a public key pk , a message $m \in \mathcal{M}$ and a signature σ as input and outputs a bit $b \in \{0, 1\}$.

Besides the usual correctness property, Σ needs to provide some unforgeability notion. We consider two notions, namely existential unforgeability under adaptively chosen message attacks (EUF-CMA security) and its strong variant (sEUF-CMA security), which we define below.

Definition 2.7 (EUF-CMA). A signature scheme Σ is EUF-CMA secure, if for all PPT adversaries \mathcal{A} there is a negligible function $\varepsilon(\cdot)$ such that

$$\Pr \left[(\text{sk}, \text{pk}) \leftarrow \text{Gen}(1^\kappa), (m^*, \sigma^*) \leftarrow \mathcal{A}^{\text{Sign}(\text{sk}, \cdot)}(\text{pk}) : \right. \\ \left. \text{Verify}(\text{pk}, m^*, \sigma^*) = 1 \wedge m^* \notin Q^{\text{Sign}} \right] \leq \varepsilon(\kappa),$$

where the environment keeps track of the queries to the signing oracle via Q^{Sign} .

Definition 2.8 (sEUF-CMA). A signature scheme Σ is strongly EUF-CMA (sEUF-CMA) secure, if for all PPT adversaries \mathcal{A} there is a negligible function $\varepsilon(\cdot)$ such that

$$\Pr \left[(\text{sk}, \text{pk}) \leftarrow \text{Gen}(1^\kappa), (m^*, \sigma^*) \leftarrow \mathcal{A}^{\text{Sign}(\text{sk}, \cdot)}(\text{pk}) : \right. \\ \left. \text{Verify}(\text{pk}, m^*, \sigma^*) = 1 \wedge (m^*, \sigma^*) \notin Q^{\text{Sign}} \right] \leq \varepsilon(\kappa),$$

where the environment keeps track of the queries and responses to and from the signing oracle via Q^{Sign} .

3 The picnic-FS and picnic-UR Signature Schemes

Let us recall that the Picnic signature scheme essentially uses a Fiat-Shamir (FS) transformed, or Unruh (UR) transformed, respectively, version of the ZKB++ protocol to prove knowledge of a witness with respect to a language L_R with associated witness relation R of pre-images of a one-way function $f_u : D_\kappa \rightarrow R_\kappa$ sampled uniformly at random from a family of one-way functions $\{f_u\}_{u \in K_\kappa}$:

$$((y, u), x) \in R \iff y = f_u(x).$$

In the so-obtained signature scheme the public verification key \mathbf{pk} contains the image y and a description of f , and the secret signing key \mathbf{sk} is a random key x from D_κ .

In Scheme 5 we provide a general description of our two signature schemes. In both schemes `Prove` is implemented with ZKB++, in `picnic-FS` it is made non-interactive with the FS transform, while in `picnic-UR`, Unruh’s transform is used. For the security analysis of the schemes we refer the reader to Sections 5.2.1 and 5.2.2.

<p><u>$\text{Gen}(1^\kappa)$</u> : Choose $u \xleftarrow{R} K_\kappa$, $x \xleftarrow{R} D_\kappa$, compute $y \leftarrow f_u(x)$, set $\mathbf{pk} \leftarrow (y, u)$ and $\mathbf{sk} \leftarrow (\mathbf{pk}, x)$ and return $(\mathbf{sk}, \mathbf{pk})$.</p>
<p><u>$\text{Sign}(\mathbf{sk}, m)$</u> : Parse \mathbf{sk} as (\mathbf{pk}, x), compute $p = (r, s) \leftarrow \text{Prove}_H((y, u), x)$ and return $\sigma \leftarrow p$, where internally the challenge is computed as $c \leftarrow H(r, \mathbf{pk} m)$.</p>
<p><u>$\text{Verify}(\mathbf{pk}, m, \sigma)$</u> : Parse \mathbf{pk} as (y, u), and σ as $p = (r, s)$. Return 1 if the following holds, and 0 otherwise:</p> $\text{Verify}_H((y, u), p) = 1,$ <p>where internally the challenge is computed as $c \leftarrow H(r, \mathbf{pk} m)$.</p>

Scheme 5: Generic description the picnic-FS and picnic-UR signature schemes.

3.1 Instantiation and Optimizations of Unruh’s Transform

We can apply Unruh’s transform to ZKB++ in a relatively straightforward manner by modifying our protocol. Although ZKB++ has 3-special sound-

ness, whereas Unruh’s transform is only proven secure for Σ -protocols with 2-special soundness, the proof is easily modified to 3-special soundness.

Since ZKB++ has 3-special soundness, we would need at least three responses for each iteration. Moreover, since there are only three possible challenges in ZKB++, we run Unruh’s transform with $\mathbf{E} = \{1, 2, 3\}$ and $M = 3$ – i.e., every possible challenge and response. We can then proceed as follows:

- $\text{Prove}_H(1^\kappa, (x, m), w)$:
1. For $i \in [t]$:
 - (a) Start \mathbf{P} on $(1^\kappa, x, w)$ and obtain first message \mathbf{a}_i .
 - (b) For all $\mathbf{e}_{i,j} = j \in \mathbf{E}$, obtain response $\mathbf{z}_{i,j}$ for challenge $\mathbf{e}_{i,j}$.
 2. For $(i, j) \in [t] \times \mathbf{E}$, set $g_{i,j} \leftarrow G(\mathbf{z}_{i,j})$.
 3. Let $(J_1, \dots, J_t) \leftarrow H(m, (\mathbf{a}_i)_{i \in [t]}, (g_{i,1}, \dots, g_{i,3})_{i \in [t]})$
 4. Return $\pi \leftarrow ((\mathbf{a}_i)_{i \in [t]}, (g_{i,1}, \dots, g_{i,3})_{(i,j) \in [t]}, (\mathbf{z}_{i,J_i})_{i \in [t]})$

As we no longer randomly select the challenges, we can omit them as input to the hash function and do not need to include them in the proof.

To instantiate the function G in the protocol, Unruh shows that one does not need a random oracle that is actually a permutation. Instead, as long as the domain and codomain of G are the same size, it can be used, since it is indistinguishable from a random permutation. So let $G : \{0, 1\}^{|\mathbf{z}_{i,j}|} \rightarrow \{0, 1\}^{|\mathbf{z}_{i,j}|}$ be a hash function modeled as a random oracle. The size of the response changes depending on what the challenge is. If the challenge is 0, the response is slightly smaller as it does not need to include the extra input share. So more precisely, this is actually two hash functions, G_0 used for the 0-challenge response and $G_{1,2}$ used for the other two challenges. In our specification document we define G precisely.

Optimization 1: Making Use of Overlapping Responses. We can make use of the structure of the ZKB++ proofs to achieve a very significant reduction in the proof size. Although we refer to three separate challenges, in the case of the ZKB++ protocol, there is a large overlap between the contents of the responses corresponding to these challenges. In particular, there are only three distinct views in the ZKB++ protocol, two of which are opened for a given challenge.

Instead of computing a permutation of each *response*, $\mathbf{z}_{i,j}$, we can compute a permutation of each *view*, $v_{i,j}$. For each $i \in \{1, \dots, t\}$, and for each $j \in \mathbf{E}$, the prover computes $g_{i,j} = G(v_{i,j})$.

The verifier checks the permuted value for each of the two views in the response. In particular, for challenge $j \in \{1, 2, 3\}$, the verifier will need to check that $g_{i,j} = G(v_{i,j})$ and $g_{i,j+1} = G(v_{i,j+1})$.

Optimization 2: Omit Re-Computable Values. Moreover, since G is a public function, we do not need to include $G(v_{i,j})$ in the transcript if we have included $v_{i,j}$ in the response. Thus for the two views (corresponding to a single challenge) that the prover sends as part of the proof, we do not need to include the permutations of those views. We only need to include $G(v_{i,(j+2)})$, where $v_{i,(j+2)}$ is the view that the prover does not open for the given challenge.

3.2 Seed Generation

We generate seeds for the random tapes using SHAKE with the private key, message and the public key as input and requesting the required number of output bytes from the XOF. Complete details are given in the specification document.

3.3 Random Tapes

We generate random tapes using SHAKE as a KDF by first hashing the seed and then requesting the required number of output bytes from the XOF. Complete details are given in the specification document.

3.4 Challenge Generation

For both the FS and Unruh transform the challenge is computed with a hash function $H : \{0, 1\}^* \rightarrow \{0, 1, 2\}^t$ (implemented using SHAKE) and rejection sampling: we split the output bits in pairs of two bits and reject all pairs with both bits set.

3.5 Function G

As explained in Section 3.1, G may be implemented with a hash function with the same domain and range. We implement $G(x)$ with SHAKE, where the requested number of output bits is $|x|$.

Security level	Blocksize	S-boxes	Keysize	Rounds
	n	m	k	r
L1	128	10	128	20
L3	192	10	192	30
L5	256	10	256	38

Table 1: Parameters for LOWMC targeting security levels L1, L3 and L5. All parameters are computed for data complexity $d = 1$.

4 Choice of Parameters

4.1 LowMC Parameters

To minimize the number of AND gates for a given keylength k and data complexity d , we want to minimize $r \cdot m$ (where r is the number of rounds and m is the number of sboxes). One strategy would be to set m to 1, and to look for an n that minimizes r . Examples of such an approach are already given in the document describing version 2 of the LowMC design [ARS⁺16]. In our setting, this approach may not lead to the best results, as it ignores the impact of the large amount of XOR operations it requires. While Picnic signatures defined with these parameters have minimal length, the large number of XOR gates make signing and verification slow. To find the most suitable parameters, we thus explore a larger range of values for m , looking to balance signing and verification cost with signature size.

Whenever we want to instantiate our signature scheme with LOWMC with κ -bit quantum security, we set $k = n = 2 \cdot \kappa$. This matches AES, and the security levels in the NIST call for proposals.

Furthermore, we observe that for a given key the adversary only ever sees a single plaintext-ciphertext pair (namely, the public key in the Picnic scheme). This is why we can set the data complexity $d = 1^2$.

4.2 Number of Parallel Repetitions

A single repetition of ZKB++ has a soundness error of $2/3$, which means that we need to perform parallel repetitions to achieve the desired soundness error. Hence we need 219 parallel repetitions for 128-bit classical security

² d is given in units of $\log_2(n)$, where n is the number of pairs. Thus setting $d = 1$ corresponds to 2-pairs, which is exactly what we need for our signature schemes.

$((3/2)^{219} \geq 2^{128})$. For 128-bit PQ security, we must set our repetition count to $t := 438$. This is double the repetition count required for classical security due to Grover’s algorithm [Gro96]. The required number of repetitions for the L1, L3 and L5 security levels are given in Table 2.

level	# parallel repetitions
L1	219
L3	329
L5	418

Table 2: Number of parallel repetitions required at each security level.

5 Formal Security Analysis

In this section we first formally confirm that the modifications/optimizations induced by moving from ZKBOO [GMO16b] to ZKB++ [CDG⁺17] as presented in Section 2.7 do not impact the security of the Σ -protocol. Then we move on to formally analyze the security of the Picnic signature scheme. In particular, we separately analyze the security of the two variants picnic-FS and picnic-UR schemes, which differ in the way the underlying Σ -protocol based on ZKB++ is made non-interactive.

5.1 Security Analysis of ZKB++

First, we observe that all modifications except “*No Additional Randomness for Commitments*” are equivalence transformations, and, therefore, do not impact the security of the overall ZKB++ Σ -protocol. Thus, we only have to show that including no additional randomness in the commitments preserves completeness, 3-special soundness, and special honest-verifier zero-knowledge of ZKB++.

First, we observe that completeness is clearly not impacted and the corollary below follows from this and [GMO16b, Proof of Proposition 2].

Corollary 5.1. *The modified version of ZKBOO—where the commitments no longer contain additional randomness—is complete, i.e., ZKB++ is complete.*

Second, under the observation that removing the randomness in the commitments does not impact the binding property of the commitments we can derive the following corollary from [GMO16b, Proof of Proposition 2].

Corollary 5.2. *The modified version of ZKBOO—where the commitments no longer contain additional randomness—is 3-special sound, i.e., ZKB++ is 3-special sound.*

What remains is to prove the following theorem.

Theorem 5.3. *The modified version of ZKBOO—where the commitments no longer contain additional randomness—is special honest-verifier zero-knowledge in the random oracle model, i.e., ZKB++ is special honest-verifier zero-knowledge in the (quantum) random oracle model.*

Before we prove the theorem, we recall that—as the challenge can be determined a priori in the proof for special honest-verifier zero-knowledge—we can use the 2-privacy simulator of the (2,3)-decomposition underlying ZKB++ (cf. Section 2.6 for details) to produce satisfying transcripts for the two views which need to be opened according to the challenge. Now, in the original proof [GMO16b, Proof of Proposition 2], the hiding property of the commitment which is not required to be opened ensures that the simulation works out. We have to argue that this still holds when no additional randomness is included in the commitments. Since we already use the random oracle heuristic for our signature scheme, we also rely on the random oracle heuristic for the subsequent proof.

Proof (Sketch). Essentially, for every commitment to a view we have that all values which depend on this seed are also only contained in this commitment and do not appear in other values. Now if a commitment is not opened, it is immediate that in the random oracle model the probability of breaking the hiding property is upper bounded by $2^{-\kappa}$, with κ being the min-entropy of the seed for the PRG, is negligible. \square

5.2 Security Analysis of Picnic

In the following sections we show that making the ZKB++ protocol non-interactive via the respective transformation yields a simulation-sound extractable non-interactive zero-knowledge proof system. As we will show in Section 5.3, this property directly yields strongly unforgeable signatures, i.e.,

the schemes satisfy strong existential unforgeability under chosen message attacks (sEUF-CMA) security.

5.2.1 Security Analysis of picnic-FS

If we view ZKB++ as a canonical identification scheme that is secure against passive adversaries one just needs to keep in mind that most definitions are tailored to (2-)special soundness, and the 3-special soundness of ZKB++ requires an additional rewind. In particular, an adapted version of the proof of [Kat10, Theorem 8.2], which considers this additional rewind, attests the security of picnic-FS. We obtain the following:

Corollary 5.4. *picnic-FS instantiated with ZKB++ and a secure one-way function yields an EUF-CMA secure signature scheme in the ROM.*

However, we actually aim for a stronger result, i.e., sEUF-CMA, which also excludes malleability of the signatures. To show that picnic-FS also satisfies this property, we need to view ZKB++ as a Σ -protocol which is transformed to its non-interactive counterpart via the FS transform and show that this protocol is actually simulation-extractable. We base our argumentation upon the argumentation of [FKMV12] to confirm simulation-extractability. What we have to do is to show that the FS transformed ZKB++ is zero-knowledge and provides quasi-unique responses. We do so by proving two lemmas. Combining those lemmas with [FKMV12, Theorem 2 and Theorem 3] then yields simulation-extractability as a corollary.

Lemma 5.5. *Let Q_H be the number of queries to the random oracle H , Q_S be the overall queries to the simulator, and let the commitments be instantiated via a RO H' with output space $\{0, 1\}^\rho$ and the committed values having min entropy ν . Then the probability $\epsilon(\kappa)$ for all PPT adversaries \mathcal{A} to break zero-knowledge of κ parallel executions of the FS transformed ZKB++ is bounded by $\epsilon(\kappa) \leq s/2^\nu + (Q_S \cdot Q_H)/2^{3 \cdot \rho}$.*

The subsequent proof is similar to the general results for Σ -protocols from [FKMV12], yet we have to account for the additional challenge that the simulator only outputs transcripts which are statistically close to original transcripts (which is in contrast to the identically distributed transcripts in [FKMV12]). Furthermore, we also provide concrete bounds.

Proof. We bound the probability of any PPT adversary \mathcal{A} to win the zero-knowledge game by showing that the simulation of the proof oracle is statistically close to the real proof oracle. For our proof let the environment maintain a list \mathbb{H} where all entries are initially set to \perp .

Game 0: The zero-knowledge game where the proofs are honestly computed, and the ROs are simulated honestly.

Game 1: As Game 0, but whenever the adversary requests a proof for some tuple (x, w) we choose $\mathbf{e} \leftarrow_{\mathcal{R}} \{0, 1, 2\}^\kappa$ before computing \mathbf{a} and \mathbf{z} . If $\mathbb{H}[(\mathbf{a}, x)] \neq \perp$ we abort and call that event E . Otherwise, we set $\mathbb{H}[(\mathbf{a}, x)] \leftarrow \mathbf{e}$.

Transition - Game 0 \rightarrow Game 1: Game 0 and Game 1 proceed identically unless E happens. The message \mathbf{a} includes 3 RO commitments with respect to H' , i.e., a lower bound for the min-entropy is $3 \cdot \rho$. We have that $|\Pr[S_0] - \Pr[S_1]| \leq (Q_S \cdot Q_H) / 2^{3 \cdot \rho}$.

Game 2: As Game 1, but we compute the commitments in \mathbf{a} so that the commitments which will never be opened according to \mathbf{e} contain random values.

Transition - Game 1 \rightarrow Game 2: The statistical difference between Game 1 and Game 2 can be upper bounded by $|\Pr[S_1] - \Pr[S_2]| \leq \kappa \cdot 1/2^\nu$ (for compactness we collapsed the s game changes into a single game).

Game 3: As Game 2, but we use the HVZK simulator to obtain $(\mathbf{a}, \mathbf{e}, \mathbf{z})$.

Transition - Game 2 \rightarrow Game 3: This change is conceptual, i.e., $\Pr[S_2] = \Pr[S_3]$.

In Game 0, we sample from the first distribution of the zero-knowledge game, whereas we sample from the second one in Game 3; the distinguishing bounds shown above conclude the proof. \square

Lemma 5.6. *Let the commitments be instantiated via a RO H' with output space $\{0, 1\}^\rho$ and let $Q_{H'}$ be the number of queries to H' , then the probability to break quasi-unique responses is bounded by $Q_{H'}^2 / 2^\rho$.*

Proof. To break quasi-unique responses, the adversary would need to come up with two valid proofs $(\mathbf{a}, \mathbf{e}, \mathbf{z})$ and $(\mathbf{a}, \mathbf{e}, \mathbf{z}')$. The last message \mathbf{z} (resp \mathbf{z}')

only contains openings to commitments, meaning that breaking quasi unique responses implies finding a collision for at least one of the commitments. The probability for this to happen is upper bounded by $Q_{H'}^2/2^p$ which concludes the proof. \square

Combining Lemma 5.5 and Lemma 5.6 with [FKMV12, Theorem 2 and Theorem 3] yields the following corollary.

Corollary 5.7. *The FS transformed ZKB++ protocol provides simulation-extractability.*

5.2.2 Security Analysis of picnic-UR

Here we prove that the proof system we get by applying our modified Unruh transform to ZKB++ is both zero knowledge and simulation-extractable in the quantum random oracle model.

Before we begin, we note that the quantum random oracle model is highly non-trivial, and a lot of the techniques used in standard random oracle proofs do not apply. The adversary is a quantum algorithm that may query the oracle on quantum inputs which are a superposition of states and receive superposition of outputs. If we try to measure those states, we change the outcome, so we do not for example have the same ability to view the adversary's input and program the responses that we would in the standard ROM.

Here we rely on lemmas from Unruh's work on quantum-secure Fiat-Shamir like proofs [Unr15]. We follow his proof strategy as closely as possible, modifying it to account for the optimizations we made and the fact that we have only 3-special soundness in our underlying Σ -protocol.

Zero Knowledge This proof very closely follows the proof from [Unr15]. The main difference is that we also use the random oracle to form our commitments, which is addressed in the transition from game 2 to game 3 below.

Consider the simulator described in Figure 6. From this point on we assume for simplicity of notation that View_3 includes x_3 .

We proceed via a series of games.

Game 1: This is the real game in the quantum random oracle model. Let H_{com} be the random oracle used for forming the commitments, H_{chal} be the random oracle used for forming the challenge, and G be the additional random permutation.

Game 2: We change the prover so that it first chooses random $e^* = e^{*(1)}, \dots, e^{*(t)}$, and then on step 2, it programs $H_{chal}(a^{(1)}, \dots, a^{(t)}, h^{(1)}, \dots, h^{(t)}) = e^*$.

Note that each the $a^{(1)}, \dots, a^{(t)}, h^{(1)}, \dots, h^{(t)}$ has sufficient collision-entropy, since it includes $\{h^{(i)} = (g_1^{(i)}, g_2^{(i)}, g_3^{(i)})\}$, the output of a permutation on input whose first k bits are chosen at random (the $k_j^{(i)}$), so we can apply Corollary 11 from [Unr15] (using a hybrid argument) to argue that Game 1 and Game 2 are indistinguishable.

Game 3: We replace the output of each $H_{com}(k_{e^*(i)}, \mathbf{View}_{e^*(i)})$ and $G(k_{e^*(i)}, \mathbf{View}_{e^*(i)})$ with a pair of random values.

First, note that H_{com} and G are always called (by the honest party) on the same inputs, so we will consider them as a single random oracle with a longer output space, which we refer to as H for this proof.

Now, to show that Games 2 and 3 are indistinguishable, we proceed via a series of hybrids, where the i -th hybrid replaces the first i such outputs with random values.

To show that the i -th and $i + 1$ -st hybrid are indistinguishable, we rely on Lemma 9 from [Unr15]. This lemma says the following: For any quantum A_0, A_1 which make q_0, q_1 queries to H respectively and classical A_C , all three of which may share state, let P_C be the probability if we choose a random function H and a random output B , then run A_0^H followed by A_C to generate x , and then run $A_1^H(x, B)$, that for a random j , the j -th query A_1^H makes is measured as $x' = x$. Then as long as the output of A_C has collision-entropy at least k , the advantage with which A_1^H , when run after A_0, A_C as described, distinguishes (x, B) from $(x, H(x))$ is at most $(4 + \sqrt{2})\sqrt{q_0}2^{-k/4} + 2q_1\sqrt{P_C}$.

In other words, if we can divide our game into three such algorithms and argue that the A_1 queries H on something that collapses to x with only negligible probability, then we can conclude that the two games are indistinguishable. Let A_0 run the game up until just before the i th iteration in the proof generation. Let A_C be the process which chooses $k_1^{(i)}, k_2^{(i)}, k_3^{(i)}$ and generates $\mathbf{View}_1^{(i)}, \mathbf{View}_2^{(i)}, \mathbf{View}_3^{(i)}$, and outputs $x = k_{e^*(i)}, \mathbf{View}_{e^*(i)}$. (Note that this has collision entropy $|k_{e^*(i)}|$ which is sufficient.) Let A_1 be the process which runs the rest of the proof, and then runs the adversary on the response.

Now we just have to argue that the probability that we make a measurement of A_1 's j -th query to H and get x is negligible. To do this, we reduce to the security of the PRG used to generate the random tapes (and hence the views). Note that besides the one RO query, $k_{e^*(i)}$ is only used as input to the PRG. So, suppose there exists a quantum adversary A for which the resulting A_1 has non-negligible probability of making an H -query that collapses to x . Then we can construct a quantum attacker for the PRG: we run the above A_0, A_C , but instead of choosing $k_{e^*(i)}$ we use the PRG challenge as the resulting random tape, and return a random value as the RO output. Then we run A_1 , which continues the proof (which should query $k_{e^*(i)}$ only with negligible probability since k s are chosen at random), and then runs the adversary. We pick a random j , and on the adversary's j -th query, we make a measurement and if it gives us a seed consistent with our challenge tape, we output 1, otherwise a random bit. If P_C is non-negligible then we will obtain the correct seed and distinguish with non-negligible probability.

Game 4: For each i instead of choosing random $k_{e^*(i)}$ and expanding it via the PRG to get the random tape used to compute the views, we choose those tapes directly at random.

Note that in Game 3, $k_{e^*(i)}$ are now only used as seeds for the PRG, so this follows from pseudo-randomness via a hybrid argument.

Game 5: We use the simulator to generate the views that will be opened, i.e. $j \neq e^*(i)$ for each i . We note that now the simulator no longer uses the witness.

This is identical by perfect privacy of the circuit decomposition.

Game 6: To allow for extraction in the simulation-extractability game we replace the random oracles with random polynomials whose degree is larger than the number of queries the adversary makes. The argument here is identical to that in [Unr15].

Online Extractability Before we prove online simulation-extractability, we define some notation to simplify the presentation:

For any proof $\pi = e, \{b^{(i)}, g^{(i)}, z^{(i)}\}_{i=1..t}$, let $\text{hash-input}(\pi) = \{a^{(i)}, h^{(i)} = (g_1^{(i)}, g_2^{(i)}, g_3^{(i)})\}$ be the values that the verifier uses as input to H_{chal} in the verification of π as described in Figure 4.

For a proof $\pi = (e, \{b^{(i)}, g^{(i)}, z^{(i)}\}_{i=1\dots t})$, let $\text{open}_0(z^{(i)})$, $\text{open}_1(z^{(i)})$ denote the values derived from $z^{(i)}$ and used to compute $C_{e_i}^{(i)}$ and $C_{e_i+1}^{(i)}$ respectively in the description of Ver in Figure 4.

We say a tuple $(a, j, (o_1, o_2))$ is valid if $a = (y_1, y_2, y_3, C_1, C_2, C_3)$, $C_j = H_{\text{com}}(o_1)$, $C_{j+1} = H_{\text{com}}(o_2)$ and o_1, o_2 consist of k , View pairs for player $j, j+1$ that are consistent according to the circuit decomposition. We say $(a, j, (O_1, O_2))$ is set-valid if there exists $o_1 \in O_1$ and $o_2 \in O_2$ such that $(a, j, (o_1, o_2))$ is valid and set-invalid if not.

We first restate lemma 16 from [Unr15] tailored to our application, in particular the fact that our proofs do not explicitly contain the commitment but rather the information the verifier needs to recompute it.

Lemma 5.8. *Let q_G be the number of queries to G made by the adversary A and the simulator S in the simulation extractability game, and let n be the number of proofs generated by S . Then the probability that A produces $x, \pi^* \notin \text{simproofs}$ where x, π^* is accepted by Ver^H , and $\text{hash-input}(\pi^*) = \text{hash-input}(\pi')$ for a previous proof π' produced by the simulator, is at most $n(n+1)/2(2^{-\kappa})^{3t} + O((q_G+1)^3 2^{-\kappa})$ (Call this event MallSim .)*

Proof. This proof follows almost exactly as in [Unr15].

First, we argue that G is indistinguishable from a random function exactly in [Unr15].

Then, observe that there are only two ways MallSim can occur:

Let e' be the hash value in π' . Then either S reprograms H sometime after π' is generated so that $H(\text{hash-input}(\pi'))$ is no longer e' , or π^* also contains the same e as π , i.e. $e = e'$. S only reprograms H if it chooses the same hash-input in a later proof - and hash-input includes $g_j^{(i)}$, i.e. a random function applied to an input which includes a randomly chosen seed. Thus, the probability that S chooses the same hash-input twice is at most $n(n+1)/2(2^{-\kappa})^{3t} + O((q_G+1)^3 2^{-\kappa})$, where $(2^{-\kappa})^{3t}$ is the probability that two proofs use all the same seeds, and $O((q_G+1)^3 2^{-\kappa})$ is the probability that two different seeds result in a collision in G , where the latter follows from Theorem 8 in [Unr15].

The other possibility is that $\text{hash-input}(\pi^*) = \text{hash-input}(\pi')$, and $e = e'$, but $b^{(i)}, g^{(i)}, z^{(i)} \neq b'^{(i)}, g'^{(i)}, z'^{(i)}$ for some i . First note, that if $e = e'$ and $\text{hash-input}(\pi^*) = \text{hash-input}(\pi')$, then $g^{(i)} = g'^{(i)}$ and $b^{(i)} = b'^{(i)}$ for all i , by definition of hash-input . Thus, the only remaining possibility is that $z^{(i)} \neq z'^{(i)}$ for some i . But since $h^{(i)} = h'^{(i)}$ for all i , this implies a collision

in G , which again by Theorem 8 in [Unr15] occurs with probability at most $O((q_G + 1)^3 2^{-\kappa})$.

We conclude that MallSim occurs with probability at most $n(n+1)/2(2^{-\kappa})^{3t} + O((q_G + 1)^3 2^{-\kappa})$. \square

Here, next we present our variant of lemma 17 from [Unr15]. Note that this is quoted almost directly from Unruh with two modifications to account for the fact that our proofs do not explicitly contain the commitment but rather the information the verifier needs to recompute it, and the fact that our underlying Σ -protocol has only 3 challenges and satisfies 3-special soundness. H_0 in this lemma will correspond in our final proof to the initial state of H_{chal} .

Lemma 5.9. *Let G, H_{com} be arbitrarily distributed functions, and let $H_0 : \{0, 1\}^{\leq \ell} \rightarrow \{0, 1\}^{2t}$ be uniformly random (and independent of G). Then, it is hard to find x and π such that for $\{a^{(i)}, (g_1^{(i)}, g_2^{(i)}, g_3^{(i)})\} = \text{hash-input}(\pi)$ and $J_1 || \dots || J_t := H_0(\text{hash-input}(\pi))$*

(i) $g_{J_i}^{(i)} = G(\text{open}_0(z^{(i)}))$ and $g_{J_{i+1}}^{(i)} = G(\text{open}_1(z^{(i)}))$ for all i .

(ii) $(a^{(i)}, J_i, (\text{open}_0(z^{(i)}), \text{open}_1(z^{(i)})))$ is valid for all i .

(iii) For every i , there exists a j such that $(a^{(i)}, j, G^{-1}(g_{i,j}), G^{-1}(g_{i,j+1}))$ is set-invalid.

More precisely, if A^{G, H_0} makes at most q_H queries to H_0 , it outputs (x, π) with these properties with probability at most $2(q_H + 1)(\frac{2}{3})^{t/2}$

Proof. Without loss of generality, we can assume that G, H_{com} are fixed functions which A knows, so for this lemma we only treat H_0 as a random oracle.

For any given value of H_0 , we call a tuple $c = (x, \{a^{(i)}\}_i, \{g_j^{(i)}\}_{i,j})$ a candidate iff: for each i , among the three transcripts, $(a^{(i)}, 1, G^{-1}(g_1^{(i)}), G^{-1}(g_2^{(i)}))$, $(a^{(i)}, 2, G^{-1}(g_2^{(i)}), G^{-1}(g_3^{(i)}))$, and $(a^{(i)}, 3, G^{-1}(g_3^{(i)}), G^{-1}(g_1^{(i)}))$ at least one is set-valid, and at least one is set-invalid. Let $n_{\text{twovalid}}(c)$ be the number of i 's for which there are 2 set-valid transcripts. Let $\mathbf{E}_{\text{valid}}(c)$ be the set of challenge tuples which correspond to only set-valid conversations. (Note that $|\mathbf{E}_{\text{valid}}(c)| = 2^{n_{\text{twovalid}}(c)}$.) We call a candidate an H_0 -solution if the challenge produced by H_0 only opens set-valid conversations, i.e. in lies in $\mathbf{E}_{\text{valid}}(c)$. We now aim to prove that A^H outputs an H_0 solution with negligible probability.

For any given candidate c , for uniformly random H_0 , the probability that c is an H_0 -solution is $\leq (\frac{2}{3})^t$. In particular, for candidate c the probability is $(\frac{2}{3})^t * 2^{n_{\text{twovalid}}(c)-t}$.

Let \mathbf{Cand} be the set of all candidates. Let $F : \mathbf{Cand} \rightarrow \{0, 1\}$ be a random function such that for each c $F(c)$ is i.i.d. with $Pr[F(c) = 1] = (2/3)^t$.

Given F , we construct $H_F : \{0, 1\}^* \rightarrow \mathbb{Z}_3^t$ as follows:

- For each $c \notin \mathbf{Cand}$, $H_F(c)$ is set to a uniformly random $y \in \mathbb{Z}_3^t$.
- For each $c \in \mathbf{Cand}$ such that $F(c) = 0$, $H_F(c)$ is set to a uniformly random $y \in \mathbb{Z}_3^t \setminus \mathbf{E}_{\text{valid}}(c)$.
- For each $c \in \mathbf{Cand}$ with $F(c) = 1$, with probability $2^{n_{\text{twovalid}}-t}$, choose a random challenge tuple e from $\mathbf{E}_{\text{valid}}(c)$, and set $H_F(c) := e$. Otherwise $H_F(c)$ is set to a uniformly random $y \in \mathbb{Z}_3^t \setminus \mathbf{E}_{\text{valid}}(c)$.

Note that for each c , and e the probability of $H(c)$ being set to e is 3^{-t} . Suppose A_0^H outputs an H_0 -solution with probability μ , then since H_F has the same distribution as H_0 , $A^{H_F}()$ outputs an H_F solution c with probability μ . By our definition of H_F , if c is an H_F solution, then $F(c) = 1$. Thus, $A^{H_F}()$ outputs c such that $F(c) = 1$ with probability at least μ .

As in [Unr15], we can simulate $A^{H_F}()$ with another algorithm which generates H_F on the fly, and thus makes at most the same number of queries to F that A makes to H_F . Thus by applying Lemma 7 from [Unr15], we get

$$\mu \leq 2(q_H + 1)\left(\frac{2}{3}\right)^{t/2}.$$

□

Finally, as the sigma protocol underlying our proofs is only computationally sound (because we use H_{com} for our commitment scheme), we need to argue that an extractor can extract from 3 valid transcripts with all but negligible probability.

Lemma 5.10. *There exists an extractor E_Σ such that for any PPT quantum adversary A , the probability that A can produce $(a, \{(\nu_{1,j}, \nu_{2,j})\}_{j=1,2,3})$ such that $(a, j, (\nu_{1,j}, \nu_{2,j}))$ is a valid transcript for $j = 1, 2, 3$, but $E_\Sigma(a, \{(\nu_{1,j}, \nu_{2,j})\}_{j=1,2,3})$ fails to extract a proof, is negligible.*

Proof. Recall that $a = (y_1, y_2, y_3, C_1, C_2, C_3)$, and if all three transcripts are valid, $C_j = H_{\text{com}}(\nu_{1,j}) = H_{\text{com}}(\nu_{2,j-1})$ for $j = 1, 2, 3$. Thus, either we

have $\nu_{1,j} = \nu_{2,j-1}$ for all j or \mathcal{A} has found a collision in H_{com} . But, Theorem 8 in [Unr15] tells us that the probability of finding a collision in a random function with k -bit output using at most q queries is at most $O((q+1)^3 2^{-k})$, which is negligible. If $\nu_{1,j} = \nu_{2,j-1}$ for all j , then we have $3 k_j || \text{View}_j$ values, all of which are pairwise consistent, so we conclude by the correctness of the circuit decomposition, and the fact that $(x = y, w) \in R$ iff $\phi(w) = y$ that if we sum the input share in $\text{View}_1, \text{View}_2, \text{View}_3$, we get a witness such that $(x, w) \in R$. \square

Theorem 5.11. *Our version of the Unruh protocol satisfies simulation-extractability against a quantum adversary.*

Proof. We define the following extractor:

1. On input π , compute $\text{hash-input}(\pi) = \{a^{(i)}, h^{(i)} = (g_1^{(i)}, g_2^{(i)}, g_3^{(i)})\}$
2. For $i \in 1, \dots, t$: For $j \in 1, 2, 3$, check whether there is a solution $\nu_{1,j} \in G^{-1}(g_j^{(i)}), \nu_{2,j} \in G^{-1}(g_{j+1}^{(i)})$ such that $(a^{(i)}, j, (\nu_{1,j}, \nu_{2,j}))$ is a valid transcript. If there is a valid transcript for all j , output $E_\Sigma(a^{(i)}, \{(\nu_{1,j}, \nu_{2,j})\}_{j=1,2,3})$ as defined by Lemma 5.10 and halt.
3. If no solution is found, output \perp .

First we define some notation, again borrowed heavily from [Unr15]:

Let $\text{Ev}_i, \text{Ev}_{ii}, \text{Ev}_{iii}$ be events denoting that A in the simulation extractability game produces a proof satisfying conditions (i) , (ii) , and (iii) from Lemma 5.9 respectively.

Let SigExtFail be the event that the extractor finds a successful $(a, \{(\nu_{1,j}, \nu_{2,j})\}_{j=1,2,3})$, but E_Σ fails to produce a valid witness.

Let ShouldExt denote the event that A produces x, π such that Ver^H accepts and $(x, \pi) \notin \text{simproofs}$.

Then our goal is to prove that the w produced by the extractor is such that $(x, w) \in R$. I.e., we want to prove that the following probability is negligible.

$$\begin{aligned}
& \Pr[\text{ShouldExt} \wedge (x, w) \notin R] \\
\leq & \Pr[\text{ShouldExt} \wedge (x, w) \notin R \wedge \neg\text{MallSim}] \\
& + \Pr[\text{MallSim}] \\
= & \Pr[\text{ShouldExt} \wedge (x, w) \notin R \wedge \neg\text{MallSim} \wedge \neg\text{Ev}_{iii}] \\
& + \Pr[\text{ShouldExt} \wedge (x, w) \notin R \wedge \neg\text{MallSim} \wedge \text{Ev}_{iii}] \\
& + \Pr[\text{MallSim}] \\
\leq & \Pr[(x, w) \notin R \wedge \neg\text{Ev}_{iii}] \\
& + \Pr[\text{ShouldExt} \wedge (x, w) \notin R \wedge \neg\text{MallSim} \wedge \text{Ev}_{iii}] \\
& + \Pr[\text{MallSim}] \\
= & \Pr[\text{SigExtFail}] \\
& + \Pr[\text{ShouldExt} \wedge (x, w) \notin R \wedge \neg\text{MallSim} \wedge \text{Ev}_{iii}] \\
& + \Pr[\text{MallSim}] \\
= & \Pr[\text{SigExtFail}] \\
& + \Pr[\text{ShouldExt} \wedge (x, w) \notin R \wedge \neg\text{MallSim} \wedge \text{Ev}_i \wedge \text{Ev}_{ii} \wedge \text{Ev}_{iii}] \\
& + \Pr[\text{MallSim}] \\
\leq & \Pr[\text{SigExtFail}] \\
& + \Pr[\text{Ev}_i \wedge \text{Ev}_{ii} \wedge \text{Ev}_{iii}] \\
& + \Pr[\text{MallSim}]
\end{aligned}$$

Here, the second equality follows from the definition of SigExtFail and Ev_{iii} , and the description of the extractor. The third equality follows from the fact that $\neg\text{MallSim}$ means that the hash function on $\text{hash-input}(\pi)$ has not been reprogrammed, and the fact that ShouldExt means verification succeeds, which means that conditions (i) and (ii) are satisfied.

Finally, by Lemmas 5.10, 5.9, and 5.8, we conclude that this probability is negligible.

5.3 Strong Unforgeability of **picnic-FS** and **picnic-UR**

We have shown that **picnic-FS** and **picnic-UR** are a simulation-extractable NIZK proof systems in the classical (resp. quantum) random oracle model

against classical (resp. quantum) adversaries. Strong unforgeability (sEUFCMA security) follows directly: this is a well known result in the classical model, and shown in [Unr15] in the quantum setting. For completeness, we briefly sketch this result:

Suppose there exist an adversary \mathcal{A} who can break the strong unforgeability property (cf. Definition 2.8). Then we can construct an adversary against the ZK property of the NIZK, the simulation-extractability property of the NIZK, or the one-wayness of the one-way function. We proceed through a series of games. In the first transition, we switch the signature algorithm to use the ZK simulator rather than the prover. This is indistinguishable by ZK, so the adversary will still produce forgeries with high probability, or we have a distinguisher which breaks the ZK property. Then, when the adversary produces a valid forgery, we run the extractor to produce a pre-image of the one-way function. If this extractor does not succeed with high probability whenever the adversary produces a forgery, we break simulation-sound extractability. Note here, that our extractor is guaranteed to work as soon as either the statement or the proof is different from what the simulator produced, so we will be able to extract from new signatures on previously signed messages as required in strong unforgeability. Otherwise, we have produced a pre-image given only the output of the one-way function (recall that we use the simulator to sign, so we do not need the pre-image there), so we break the one-wayness property.

6 Analysis with Respect to Known Attacks

In this section we analyze the Picnic signature scheme with respect to known attacks. First, we observe that in case we deal with ideal primitives, Corollary 5.4 already gives us a provable bound for EUFCMA security. Since those primitives are instantiated with concrete building blocks however, we consider concrete attacks on those building blocks. In our scheme, we use the classical approach to turn Σ -protocols into signature schemes in the random oracle model. Based on the fact that, since the introduction of the random oracle model [BR93], no attack which arises from the assumption that a hash function behaves as a random oracle (except for some artificial counterexamples such as [CGH98]) was found [KM15], we claim that the best attacks against our scheme are attacks which also invalidate the claims made for the underlying symmetric primitives.

All cryptographic primitives, except the one-way function LowMC, rely on the SHA3 function SHAKE [NIS15], a well established and standardized primitive, and we use it in a standard way. For those primitives, we have already gained substantial confidence regarding security due to extensive cryptanalysis efforts within the community. We therefore do not see these building blocks as a central attack surface and assume that the bounds given in [NIS15, Table 4] hold. We note that improvements in attacks against those primitives also lead to improvements in the attacks against our signature scheme.

6.1 Usage and security margin of LowMC

Consequently, we henceforth put focus on attacks on the one-way function f . Essentially, the function f could be any one-way function, but since we found block ciphers—and, in particular the LowMC family of block ciphers [ARS+15, ARS+16]—gave the most efficient signatures, we decided to use them in our signature scheme. In particular, we assume that using LowMC as described below yields a suitable family of one-way functions $\{f_u\}_{u \in \mathcal{K}_\kappa}$. We use this function to establish a suitable relation between secret and public keys. In particular, let

$$f_u(x) := E(x, u),$$

and let E denote LowMC encryption with respect to a single block u under key x . The keys in our signature scheme are generated as follows. First, one chooses a LowMC encryption key x , as well as a single block u uniformly at random. Then, the the public verification key \mathbf{pk} , as well as the secret signing key \mathbf{sk} are defined as follows

$$\mathbf{pk} := (y, u) = (f_u(x), u), \quad \mathbf{sk} := (\mathbf{pk}, x).$$

The choice of the number of rounds within LowMC comes with a significant security margin. For security level L1 with the specified 20 rounds, the best attack known is on 12 rounds. For security level L3 with the specified 30 rounds, the best attack known is on 19 rounds. For security level L5 with the specified 38 rounds, the best attack known is on 26 rounds. And even those attacks require an attacker to see two plaintext-ciphertext pairs for the same key, whereas within our signature scheme an attacker only ever sees a single input-output pair for every key.

6.2 Attacks in the Single-User Setting

In the single-user setting, the attacker only ever sees a single key pair for the Picnic signature scheme, i.e., a single plaintext-ciphertext pair $(f_u(x), u)$ of LOWMC with respect to a uniformly random key x and a uniformly random block u . Consequently, in this setting cryptanalytic results for LOWMC also directly apply to our scheme, and also the claimed bounds carry over to the Picnic signature scheme. Note that, one could even globally fix x to further shrink the size of the public verification key pk . However, we chose not to do so, as we also want to consider attacks in the multi-user setting (as discussed below).

6.3 Attacks in the Multi-User Setting

The multi-user setting more accurately models reality, in that there are multiple users, each with a public key, and the adversary is considered successful if he can attack any one of the users.

Multi-user EUF-CMA. Even if single-user EUF-CMA security generically implies multi-user EUF-CMA (MU-EUF-CMA) security under a polynomial loss [GMS02], we put concrete focus on attack scenarios which become applicable by moving to the multi-user setting. Here, the adversary may see many signing key pairs and we need to be cautious with respect to more sophisticated attacks that might apply.

In particular—in contrast to the single-user setting—our decision to choose an independent and uniformly random block u , being the encryption function $E(\cdot, u)$ of LOWMC, per signing key pair turns out to be important. This is because using the same, fixed block u with independent keys x_1, \dots, x_n for each of the n users would allow to apply multi-user key recovery attacks [Bih02] and generic time-memory trade-off attacks like [Hel80] and in particular time/memory/key trade-off attacks [BMS05]. In these attacks one of n block cipher keys may be recovered in less time than the cost of recovering a single key, and the attacks become more efficient for large n . Intuitively, the random block chooses a unique function per user, and work done to attack one user (function) can not be used to simultaneously attack another user (function). In addition, Banegas and Bernstein [BB17] have recently shown that parallel collision search attacks [vOW94] can also be applied in the quantum setting which also supports making a random choice of u per user. Finally, we note that one could choose a smaller value that is unique

per user (with a potential decrease in security) to reduce the size of the public key. However, since public keys in our schemes are small (at most 64 bytes), our design uses a full random block.

Key-substitution attacks. These are attacks where an adversary who is given a signature σ_A on message M under A 's public key \mathbf{pk}_A manages to come up with a public key \mathbf{pk}_E (different from \mathbf{pk}_A) such that σ_A verifies under \mathbf{pk}_E and message M . Menezes and Smart in [MS04] provide a formal model to cover such attacks, which are not covered by EUF-CMA security. Therefore, we explicitly consider such attacks. Security against these types of attacks can be generically achieved. This has been shown in [MS04], and we recall their theorem below.

Theorem 6.1 ([MS04], Theorem 6). *Let $(\text{Gen}, \text{Sign}, \text{Verify})$ be an EUF-CMA secure signature scheme. Then, $(\text{Gen}, \text{Sign}', \text{Verify})$ with $\text{Sign}' := \text{Sign}(\text{sk}, \mathbf{pk}||m)$ and \mathbf{pk} being an unambiguous encoding of the public key is a secure signature scheme in the multi-user setting.*

We stress that the above result in particular holds for sEUF-CMA secure signature schemes.

As discussed in Section 3 (see also Section 4.3 of the specification), the public key is prepended to the message on signing, and the specification provides an unambiguous encoding (since the public key is a pair of bitstrings, the encoding is trivial). Consequently, we have the following corollary.

Corollary 6.2. *picnic-FS and picnic-UR provide security in the multi-user setting in the sense of [MS04, Definition 6].*

7 Expected Security Strength

Since a Picnic keypair is a block cipher key and plaintext/ciphertext pair, we chose to define parameters at the L1/AES-128, L3/AES-192 and L5/AES-256 security levels. By the CFP, this means that L2 is implicitly defined by L3 and L4 is defined by L5.

We expect each parameter set to provide security equivalent to AES. For example, at L1, we expect 128-bits of security against classical attacks, and at least 64-bits against quantum attacks. Like AES, key recovery attacks using Grover's algorithm against LowMC are the best known quantum attacks on Picnic. The estimate of 64-bits comes from an idealized version of

Grover’s algorithm capable of running for a long time. Like AES this may be conservative, even more so in the case of LowMC, since the circuit is much larger than AES, due to the large amount of constant data required to implement it.³ For example, at level L1, the Picnic LowMC instance requires 84KBytes of constant data. These constants must either be encoded into the circuit, or the circuit must be expanded to recompute them on the fly. Thus the LowMC circuit is orders of magnitude larger than AES, making attacking LowMC with Grover’s algorithm at least as difficult as attacking AES.

We’re similarly set the number of parallel iterations to provide 128-bit security against classical attacks, and 64-bits against quantum attacks using an idealized version of Grover’s algorithm. Like with LowMC, the circuit required to break ZKB++ soundness with Grover’s algorithm is orders of magnitude larger than the AES circuit.

In more detail, suppose an attacker is trying to forge a proof as a generic search problem. In particular, if an attacker can find a permutation of a set of transcripts that hash to a challenge chosen in advance, he can forge a proof. Consider a T round protocol (Picnic specifies 219, 329 and 438 rounds at levels L1, L3 and L5, resp.). Then there are 3^T possible challenges that can come from hashing those $3T$ transcripts (since there are 3 challenges).

Now consider an attacker who constructs invalid ZKB++ “proofs” such that for each ZKB++ iteration, he can give a valid response to two of the challenges but not the third. If we model the hash function as a random oracle, the probability of getting a challenge for which he can respond is $(\frac{2}{3})^T$, and thus we expect that if the attacker searches a space of $(\frac{3}{2})^T$ candidates (i.e., permutations of transcripts that are constructed in this manner) he can find one. Grover’s algorithm allows the attacker to search the space in time $(\frac{3}{2})^{T/2}$.

However, the items in the space are larger, and changes in one value (e.g., a seed value) requires re-computing many others (the random tape, the MPC transcript, and the commitments). Clearly this is far more expensive than a single AES evaluation, and so we assume that Grover’s algorithm applied to breaking ZKB++ soundness is at least as costly as AES key recovery.

³Previously, when we compared the LowMC circuit size to AES, we were looking only at AND gates, but here we’re considering all gates.

7.1 LowMC Parameter Selection

The choice of LowMC parameters may seem aggressive in the context of a general-purpose block cipher. The LowMC spec recommends an additional 1.3 times the number of rounds, as a security margin against unknown attacks. Picnic does not use these additional rounds.

The general block cipher security definition gives attackers as much power as possible, to model the worst case scenario. Consider the CPA security game, where attackers may choose plaintexts, query the encryption oracle many times, and must only distinguish encryptions of chosen plaintexts in order to break the cipher (as opposed to recovering plaintext or private keys). This strong security definition is sensible when the primitive will be used in a variety of (potentially unforeseen) applications.

By contrast, for the security of Picnic signatures, attackers are severely restricted. They are given a single plaintext/ciphertext pair, for a randomly chosen block and key, and succeed if they can recover the key. Signatures are zero-knowledge proofs, so by definition provide no additional useful information. The success criteria for a Picnic attacker is key recovery, which is more difficult than indistinguishability, while the capabilities are more limited. In this context, the parameters we have chosen for LowMC are arguably very conservative. We believe that reducing the number of rounds further would maintain security, but chose to use the full number of rounds as a security margin, given LowMC is a relatively new design.

Further, it is difficult to quantitatively support parameter selection in our restricted attack model, since most research focuses on the standard security definition. Our claim is that the complexity of a key recovery attack against LowMC, given only a Picnic public key, is at least as difficult as attacking the CPA security of AES (for equivalent key size, and with Picnic the block size always matches the key size).

As we improve our understanding of LowMC security in this restricted attack model, we expect to be able to justify reducing the number of rounds further. The motivation is the direct impact this has on the size of Picnic signatures. For example, at L1 with 128-bit blocksize and keysize, with 10 s-boxes, the recommended number of rounds is 20 and signatures are about 33KB. Reducing the number of rounds to 10 would make signatures 24KB.

7.2 Hash Function Security

Picnic depends on secure hash functions when computing signatures, for commitments and the challenge. In our security analysis we have modeled these as random oracles. While choosing parameters we also took into account some more specific security properties (all implied by a random oracle).

All hash functions are implemented with SHAKE128 with 256-bit digests at security level L1, and SHAKE256 at levels L3 and L5, with 384 and 512-bit digests, respectively.⁴ We expect the concrete security provided by SHAKE for collision and preimage resistance claimed in [NIS15], extended to quantum attacks.

For preimage resistance, in the classical case it is common to assume $O(2^n)$ operations for standard hash functions. When considering quantum algorithms, Grover’s algorithm can find preimages with $O(2^{n/2})$ operations. Therefore, we assume that our uses of SHAKE128 and SHAKE256 provide this level of preimage resistance.

When considering quantum algorithms, in theory it may be possible to find collisions using a generic algorithm of Brassard et al. [BHT98] with cost $O(2^{n/3})$ (for n -bit digests). A detailed analysis of the costs of the algorithm in [BHT98] by Bernstein [Ber09] found that in practice the quantum algorithm is unlikely to outperform the $O(2^{n/2})$ classical algorithm. Multiple cryptosystems have since made the assumption that standard hash functions with n -bit digests provide $n/2$ bits of collision resistance against quantum attacks (for examples, see papers citing [Ber09]). We make this assumption as well, and in particular, that SHAKE128 with 256-bit digests provides 128 bits of PQ security, SHAKE256 with 384-bits provides 192-bits and SHAKE256 with 512-bit digests provides 256-bits.

⁴We considered using SHAKE256 at all three levels for simplicity, but L1 signing and verify times increased by roughly 10%.

8 Advantages and Limitations

8.1 Compatibility with Existing Protocols

A preliminary version of the Picnic implementation has been integrated with the Open Quantum Safe (OQS) project.⁵ Then, using a version of OpenSSL⁶ modified to use OQS, we were able to create X.509 certificates signed with Picnic and certifying Picnic public keys. These keys and certificates were then used to establish TLS 1.2 connections, where the key exchange algorithm was one of the Microsoft Research version of NewHope, Frodo or SIDH. We then built Apache web server with this version of OpenSSL, and configured it to use our PQ ciphersuites, to serve a simple HTTP request. On the client side we used a command line tool that ships with OpenSSL for testing, as modifying a browser to support these ciphersuites is a much larger effort.

The parameter set for this earlier version of Picnic most resembled `Picnic-L5-FS` and `Picnic-L5-UR`, so signatures were 100KB.

OpenSSL had to be patched in one place to handle larger signature sizes, since the TLS 1.2 standard has a limit of $2^{16} - 1$ bytes. Signatures for the L1 parameter sets are under this limit, but for L3 and L5 we would likely need an extension to support larger signatures. Ideally a future version of TLS would allow larger signatures, but this is not pressing as we expect the L1 parameter set to provide sufficient security in the short and medium term. Otherwise the integration was smooth, and performance seemed acceptable in our limited experiments. In particular the certificate stack (X.509/ASN.1) worked unmodified with signatures this large, something we did not expect.

The final version of this document will describe our experience using the current version of Picnic in this software stack along with benchmarks.

9 Additional Security Properties

9.1 Side-Channel Attacks

Key Generation. Key generation requires generating a random LowMC key and plaintext, and computing the LowMC block cipher. A fast implementation of the LowMC block cipher may use precomputed data, and have

⁵<https://github.com/open-quantum-safe/liboqs>

⁶https://github.com/christianpaquin/openssl/tree/add_oqs_sig_with_picnic

cache-timing side channels, because the access pattern depends on the secret key. However, there are a couple mitigations:

1. Since key generation happens infrequently, a slower LowMC implementation with a constant access pattern can be used.
2. Even if a side channel is present in key generation, since only one encryption with a given secret key is ever computed, and known attacks require observing multiple runs, the feasibility of a successful attack is unlikely.

Signing. Generally speaking, since the three party protocol simulated during signing is circuit-based, the same operations are performed, regardless of the values on the input wires of the circuit.

Signing is not constant time in the absolute sense, but is constant time with respect to the operations that depend on the ephemeral random values (that in turn depend on the signing key). The timing (and signature size) variation is due to the different operations performed depending on the (public) bits of the challenge. The reference implementation is constant time relative to the secret key.

Our fast implementation of the binary matrix multiplication step for LowMC uses precomputed data (the so-called "method of four Russians"). Look-ups to the table of precomputed data are done based on the LowMC state and round key. Therefore, the access pattern to memory is dependent on secret data (the access pattern of one of the three players is sensitive and should remain secret). This is mitigated since the LowMC secret keys used by each of the players is a randomized secret sharing of the actual key. The shares are only ever used once, and the randomization means that access pattern information learned by an attacker in one parallel iteration of ZKB++ can not be combined with information from other iterations. Since cache-timing side-channel attacks overwhelmingly require multiple observations (called traces), we expect this mitigation to be effective.

The circuit decomposition technique is similar to the side channel countermeasure called masking, commonly used to protect block cipher implementations from side channel attacks. An early, well cited paper on the topic is Goubin and Patarin [GP99]. With further study, we may find that the ZKB++ circuit decomposition provides other types of side channel resistance, for example, resistance to differential power analysis.

9.2 Security Impact of Using Weak Ephemeral Values

The specification recommends that the per-signature random values used when computing a signature be derived from the signing key and the message, to simplify testing and to mitigate the security impact of a defective random number generator during signing. The goal is that signatures are secure, provided the random number generator was secure during key generation.

Like (EC)DSA, given the random values used when computing a signature, it is possible to recover the signer’s secret key. Recall that in each parallel iteration of ZKB++, three seed values are generated, one for each party in the MPC protocol. In normal operation, two of these seed values are revealed, and one is kept secret. The MPC protocol remains secure when two of the three parties are corrupted (i.e., have their seed exposed, which exposes their state, input and output shares). Given the entire third seed, it is possible to recover the input share of the third player, and recover the secret key.

Unlike (EC)DSA, slight biases in the random number generator do not allow the secret to be recovered from multiple signatures. This is because the seed values are never used directly; they are always hashed, then expanded with an extendable output functions (XOF) into a random tape, and the random tape values are used.

Regardless of how the seeds are generated, a bias in the XOF output may lead to an attack. For example, if the XOF/PRF used to derandomize (EC)DSA, EdDSA, and other ElGamal-like signatures was biased, the ephemeral value is biased, and the lattice attacks studied in the context of RNG biases apply [Ble00, HS01]. For Picnic, it’s not clear if this could be exploited.

Alternative Parameter Selections. We could reduce the feasibility of Grover key recovery attacks against LowMC by increasing the keysize and keeping the block length fixed. There would still be a chance to use Grover’s algorithm to break the soundness of ZKB++ (unless the number of parallel iterations was increased). However, a quick inspection reveals that the computation of attacking soundness is computationally much more complex than attacking LowMC. For example, checking whether a candidate secret key corresponds to a given public key requires one LowMC evaluation, while checking whether a set of cheating commitments leads to a challenge that does not catch the cheating requires hundreds of SHA3 computations.

9.3 Parameter Integrity

With some cryptographic primitives if the system parameters are changed, security is lost. For example, if an elliptic curve secret key is used on a weak curve, the primitive may still work, but leak the secret key.

In Picnic, the parameters are small integer values like the parallel repetition count that tend to be hard-coded in software, and the LowMC matrices and constants. Using weak parameters for LowMC could weaken the cipher to the point where key recovery attacks are feasible. If weak parameters are used for key generation it is possible to generate a weak keypair, however, it will not produce signatures that verify with respect to the correct parameters. So the common PKI practice of signing a certificate request with the subject key will catch keys generated with invalid parameters.

Creating signatures with invalid parameters can also be a security risk. Suppose a set of weak LowMC parameters were used, so that the signing algorithm proves knowledge of the signing key k , for an new circuit E' , i.e., $E'_k(x) = \text{pk}$, where E' is LowMC with invalid parameters. The signature would be invalid in most cases, unless key generation and verification also used E' . However, the invalid signature contains enough information to recover $E'_k(x)$, from which it may be possible to recover k .

10 Efficiency and Memory Usage

This section gives performance benchmarks of the Picnic signature scheme. A single core/thread was used for all benchmarks. All times are in milliseconds. We benchmark three implementations:

Reference An expository C implementation. Makes no performance optimizations.

Optimized-C A somewhat optimized implementation using C only.

Optimized An optimized implementation that uses processor-specific compiler intrinsics for vector instructions, e.g. SSE2 and AVX2 on Intel x86-64 and NEON on ARM v8.

10.1 Description of the Benchmark Platforms

10.1.1 Platform A

The primary benchmarking platform, **Platform A**, has the following specifications:

CPU Intel(R) Core(TM) i7-4790 CPU @ 3.60GHz

Memory 16 GB

OS Ubuntu 17.04

Compiler GCC 6.3

Intel Turbo Boost (dynamic frequency scaling) was disabled. For comparison, OpenSSL version 1.0.2g reports 0.03 ms for ECDSA signing and 0.08 ms for verification on Platform A⁷.

10.1.2 Platform B

The secondary benchmarking platform, **Platform B** (Raspberry Pi 3 Model B), has the following specifications:

CPU Quad Core 1.2GHz Broadcom BCM2837 64-bit CPU, ARM Cortex A53 (ARMv8)

Memory 1 GB RAM

OS openSUSE Leap 42.2

Compiler GCC 6.2

openSSL: 11ms sign, 40ms verify

For comparison, OpenSSL version 1.0.2j reports 11 ms for ECDSA signing and 40 ms for verification on Platform B.

⁷As reported by the command `openssl speed ecdsap256`

10.1.3 Platform C

A third platform, **Platform C**, is an older x64-based system, has the following specifications:

CPU Intel(R) Xeon(R) CPU E31230 @ 3.20GHz

Memory 8 GB

OS Ubuntu 16.04.3 LTS

Compiler GCC 5.4

For comparison, OpenSSL 1.0.2g reports 0.05 ms for ECDSA signing and 0.12 ms for ECDSA verification on Platform C.

10.2 Description of the benchmarking methodology

Timings results for key generation, signing and signature verification were averaged over 1000 runs.

On Platforms A and C we measured CPU cycles using the `perf_event` performance monitoring subsystem of the Linux kernel. On Platform B CPU cycles were measured using the hardware performance counter available via the `MRS` instruction.

10.3 Benchmark Results: Timings

In this section we describe the time required for various operations, on the three benchmark platforms. In all tables presented in this section we give the timing information as milliseconds and CPU cycles.

In Tables 3, 4 and 5 we present the benchmark results of all implementations on Platform A. On this platform we observe speed improvements of the **optimized** implementation compared to the **optimized-C** implementation of 20% up to 50% percent. Additionally, Tables 6 and 7 also give benchmark results with the GCC link time optimization (LTO) feature enabled. The **optimized-C** implementation shows improved performance figures with LTO enabled. The same holds for the **optimized** implementation except for `Picnic-L5-FS`, which is interestingly now closer to the performance figures of `Picnic-L5-UR`.

Parameters	Keygen	Sign	Verify
Picnic-L1-FS	0.05	36.50	23.91
(cycles)	163850	131390415	86062091
Picnic-L1-UR	0.04	44.12	29.48
(cycles)	146193	158826399	106128443
Picnic-L3-FS	0.10	122.85	81.03
(cycles)	364079	442257404	291723398
Picnic-L3-UR	0.10	144.96	97.40
(cycles)	369536	521842013	350635309
Picnic-L5-FS	0.20	298.11	196.85
(cycles)	722494	1073183185	7086526474
Picnic-L5-UR	0.21	329.86	221.46
(cycles)	740633	1187481996	797249015

Table 3: Benchmarks for the **reference** implementation, on benchmark Platform A.

Parameters	Keygen	Sign	Verify
Picnic-L1-FS	0.02	3.35	2.29
(cycles)	86938	12055646	8245436
Picnic-L1-UR	0.02	4.33	3.06
(cycles)	86631	15590725	11014466
Picnic-L3-FS	0.06	11.86	8.07
(cycles)	208123	42683333	29039354
Picnic-L3-UR	0.06	13.81	9.55
(cycles)	214387	49699928	34373293
Picnic-L5-FS	0.11	38.04	25.81
(cycles)	378326	136942395	92906288
Picnic-L5-UR	0.10	39.70	27.35
(cycles)	375633	142906149	98449354

Table 4: Benchmarks for the **optimized-C** implementation, on benchmark Platform A.

Parameters	Keygen	Sign	Verify
Picnic-L1-FS	0.01	2.10	1.45
(cycles)	41635	7544639	5228793
Picnic-L1-UR	0.01	2.78	2.01
(cycles)	43189	10024950	7223730
Picnic-L3-FS	0.06	7.41	5.23
(cycles)	199127	26669631	18813324
Picnic-L3-UR	0.06	9.47	6.76
(cycles)	205247	34092091	24331615
Picnic-L5-FS	0.06	14.88	10.49
(cycles)	204178	53572733	37749441
Picnic-L5-UR	0.06	19.75	14.33
(cycles)	209861	71097205	51603839

Table 5: Benchmarks for the **optimized** implementation, on benchmark Platform A.

Parameters	Keygen	Sign	Verify
Picnic-L1-FS	0.02	2.73	1.87
(cycles)	79586.35	9820159.24	6741989.53
Picnic-L1-UR	0.02	3.44	2.45
(cycles)	78805.89	12369790.39	8812970.22
Picnic-L3-FS	0.05	10.72	7.23
(cycles)	196082.92	38608875.12	26042494.71
Picnic-L3-UR	0.05	10.98	7.58
(cycles)	196882.47	39510719.80	27284410.02
Picnic-L5-FS	0.10	33.72	22.92
(cycles)	342773.08	121409878.24	82504407.64
Picnic-L5-UR	0.10	37.52	25.99
(cycles)	345184.62	135055232.98	93576995.97

Table 6: Benchmarks for the **optimized-C** implementation with LTO enabled, on benchmark Platform A.

Parameters	Keygen	Sign	Verify
Picnic-L1-FS	0.01	1.99	1.39
(cycles)	40654	7154708	5018443
Picnic-L1-UR	0.01	2.69	1.94
(cycles)	40335	9690701	6995335
Picnic-L3-FS	0.05	7.21	5.06
(cycles)	191327	25945491	18224318
Picnic-L3-UR	0.06	9.26	6.60
(cycles)	199197	33350422	23743318
Picnic-L5-FS	0.06	17.22	12.46
(cycles)	214424	61992252	44849301
Picnic-L5-UR	0.06	17.49	12.70
(cycles)	216542	62963501	45732244

Table 7: Benchmarks for the **optimized** implementation with LTO enabled, on benchmark Platform A.

Next we give the results of the evaluation of our implementations on Platform B in Tables 8, 9 and 10. Again we observe improved performance figures for the **optimized** implementation, but they are not as significant as on Platform A. GCC’s optimizer vectorizes the code on Platform B more aggressively for **optimized-C** implementation and thus the performance gains of the **optimized** implementation are smaller.

10.4 Memory Requirements

In this section we give the memory requirements for our implementations. The memory requirements of an implementation are assumed to be the same for all platforms.

10.4.1 Reference Implementation Detailed Memory usage

Memory usage was benchmarked using the Valgrind⁸ tool Massif. Massif was run on an example program, that generates a key pair, creates a signature, then verifies it, using the API in picnic.h. Then the tool `massifcherrypick`⁹ was used to determine the peak memory usage of specific functions.

⁸<http://valgrind.org/docs/manual/ms-manual.html>

⁹<https://github.com/lnishan/massif-cherrypick>

Parameters	Keygen	Sign	Verify
Picnic-L1-FS (cycles)	2.10 2518975.68	276.70 332042859.49	182.70 219239633.21
Picnic-L1-UR (cycles)	2.09 2508982.43	338.02 405624756.71	228.30 273961444.67
Picnic-L3-FS (cycles)	2.58 3095449.36	954.93 1145919907.20	633.51 760208950.65
Picnic-L3-UR (cycles)	2.60 3117483.80	1135.01 1362017384.32	766.11 919332592.31
Picnic-L5-FS (cycles)	3.43 4121029.50	2392.14 2870568905.80	1596.34 1915611672.97
Picnic-L5-UR (cycles)	3.48 4175474.33	2648.95 3178740376.65	1789.57 2147486636.49

Table 8: Benchmarks for the **reference** implementation, on benchmark Platform B.

Parameters	Keygen	Sign	Verify
Picnic-L1-FS (cycles)	2.01 2411719.08	47.61 57126852.40	32.02 38425485.52
Picnic-L1-UR (cycles)	2.05 2463859.98	50.82 60982466.78	34.79 41751597.25
Picnic-L3-FS (cycles)	2.32 2787111.83	173.52 208218410.35	116.65 139974921.54
Picnic-L3-UR (cycles)	2.32 2778689.69	180.27 216320405.46	121.81 146168964.21
Picnic-L5-FS (cycles)	2.60 3121580.16	376.07 451282793.61	251.26 301510063.31
Picnic-L5-UR (cycles)	2.61 3126934.73	389.13 466958725.13	262.32 314780997.66

Table 9: Benchmarks for the **optimized-C** implementation, on benchmark Platform B.

Parameters	Keygen	Sign	Verify
Picnic-L1-FS	1.99	43.69	30.06
(cycles)	2386054.21	52430033.06	36072091.59
Picnic-L1-UR	2.04	47.58	32.94
(cycles)	2445307.79	57099807.09	39529669.74
Picnic-L3-FS	2.44	186.95	125.55
(cycles)	2925039.83	224339885.99	150664840.51
Picnic-L3-UR	2.34	179.84	121.38
(cycles)	2807774.21	215809662.97	145655170.18
Picnic-L5-FS	2.52	357.11	238.07
(cycles)	3025127.71	428529315.98	285685718.82
Picnic-L5-UR	2.56	369.85	248.75
(cycles)	3068088.66	443818338.13	298501331.24

Table 10: Benchmarks for the **optimized** implementation, on benchmark Platform B.

Parameters	Keygen	Sign	Verify
Picnic-L1-FS	0.07	51.92	33.92
(cycles)	212881	166159669	108529436
Picnic-L1-UR	0.06	62.51	41.71
(cycles)	193569	200039945	133478227
Picnic-L3-FS	0.14	172.82	113.70
(cycles)	462757	553026245	363831127
Picnic-L3-UR	0.15	203.82	136.54
(cycles)	471108	652218966	436915884
Picnic-L5-FS	0.28	415.36	275.93
(cycles)	901355	1329138266	882984854
Picnic-L5-UR	0.29	459.44	308.87
(cycles)	922864	1470196303	988380511

Table 11: Benchmarks for the **reference** implementation, on benchmark Platform C.

Parameters	Keygen	Sign	Verify
Picnic-L1-FS	0.03	4.50	3.12
(cycles)	100469	14403397	9998428
Picnic-L1-UR	0.03	5.84	4.19
(cycles)	100825	18678042	13413739
Picnic-L3-FS	0.07	14.68	10.14
(cycles)	239214	46980834	32456544
Picnic-L3-UR	0.07	19.87	14.05
(cycles)	234956	63576699	44966117
Picnic-L5-FS	0.14	46.41	31.76
(cycles)	435994	148516664	101624906
Picnic-L5-UR	0.14	51.44	36.37
(cycles)	440258	164615202	116384582

Table 12: Benchmarks for the **optimized-C** implementation, on benchmark Platform C.

Parameters	Keygen	Sign	Verify
Picnic-L1-FS	0.02	3.50	2.42
(cycles)	48587	11193948	7745406
Picnic-L1-UR	0.02	4.87	3.48
(cycles)	48760	15594829	11132090
Picnic-L3-FS	0.07	13.46	9.39
(cycles)	226283	43076300	30061948
Picnic-L3-UR	0.07	17.36	12.42
(cycles)	230344	55549079	39737866
Picnic-L5-FS	0.10	32.21	22.22
(cycles)	311304	103056039	71098134
Picnic-L5-UR	0.10	37.77	26.82
(cycles)	314544	120862142	85830001

Table 13: Benchmarks for the **optimized** implementation, on benchmark Platform C.

Parameter set	Sign	Verify
Picnic-L1-FS	227,103	183,154
Picnic-L1-UR	315,579	254,772
Picnic-L3-FS	480,835	379,860
Picnic-L3-UR	645,335	565,034
Picnic-L5-FS	802,197	652,539
Picnic-L5-UR	1,091,277	959,553

Table 14: Peak memory usage (stack and heap combined) of reference implementation, in bytes. This excludes memory used for the LowMC constants, which was stored as static data in program binary.

Massif was invoked with the command:

```
valgrind --tool=massif --stacks=yes ./example
```

When creating a signature, peak memory usage ranged from about 227K to 1091K bytes, as shown in Table 14, while verification ranged from about 183K to 959K bytes.

Massif measures memory usage by sampling so there is some variability in these measurements. The variance was low so these are a reasonable estimate, since we’re only interested in peak usage.

10.4.2 Optimized Implementation Detailed Memory usage

With the same methodology as used for the reference implementation, peak memory usage data was generated with `valgrind`. The data is presented in Table 15. A more detail discussion of the precomputed constants and data follows below.

10.5 Size of precomputed constants and data

Size of LowMC constants The LowMC block cipher uses a large amount of constant data when compared to traditional block ciphers. This data may be computed on-the-fly as needed, or precomputed and stored in advance. All our implementations compile this data into the binary.

We did not investigate the cost of re-computing the LowMC constants at runtime. The output of the Grain LSFR is used as a self-shrinking generator to create the constants.

Parameter set	Precomputed Data	Sign	Verify
Picnic-L1-FS	1,835,744	133,598	79,724
Picnic-L1-UR	1,835,744	196,889	126,590
Picnic-L3-FS	7,078,944	230,300	108,570
Picnic-L3-UR	7,078,944	373,415	214,508
Picnic-L5-FS	11,797,792	398,580	189,216
Picnic-L5-UR	11,797,792	642,546	370,548

Table 15: Peak memory usage (stack and heap combined) of the optimized implementation, in bytes. This excludes memory used for the LowMC constants, which was stored as static data in program binary.

The **optimized-C** and **optimized** implementations include additional precomputed matrices used to compute the round keys, which allow one to split the round key computation into parts that are affected by the Sbox layer and parts which are not. The sizes of those matrices are given in Table 17 in the Reduced Key Matrices column.

Params	Linear Matrices	Round Constants	Key Matrices	Total
L1	40,960	320	43,008	84,288
L3	138,240	720	142,848	281,808
L5	311,296	1,216	319,488	632,000

Table 16: Size of constant data (in bytes) required by LowMC as used by the reference implementation, with the parameters used in Picnic at security levels L1, L3 and L5.

Size of precomputed data In order to reduce the time required to implement the binary matrix multiplication step and the linear layer of LowMC, the optimized implementations pre-compute some data based on the constants. The size of the data is given in Table 18. Once the implementation has this data, it no longer needs the constant data from Table 16 (except the round constants).

References

- [AGR⁺16] Martin R. Albrecht, Lorenzo Grassi, Christian Rechberger, Arnab Roy, and Tyge Tiessen. MiMC: Efficient encryption and

Params	Linear Matrices	Round Constants	Key Matrices	Reduced Key Matrices	Total
L1	41,600	960	43,680	14,400	100,640
L3	185,280	1,920	191,456	30,784	409,440
L5	315,512	2,432	320,736	49,216	684,896

Table 17: Size of constant data (in bytes) required by LowMC as used by the optimized implementations, with the parameters used in Picnic at security levels L1, L3 and L5.

Params	Linear Matrices	Key Matrices
L1	1,311,360	458,816
L3	5,899,200	983,104
L5	9,962,688	1,572,928

Table 18: Size of precomputed data (in bytes) required by LowMC, with the parameters used in Picnic at security levels L1, L3 and L5.

cryptographic hashing with minimal multiplicative complexity. In *ASIACRYPT*, pages 191–219, 2016.

- [ARS⁺15] Martin R. Albrecht, Christian Rechberger, Thomas Schneider, Tyge Tiessen, and Michael Zohner. Ciphers for MPC and FHE. In *EUROCRYPT*, 2015.
- [ARS⁺16] Martin R. Albrecht, Christian Rechberger, Thomas Schneider, Tyge Tiessen, and Michael Zohner. Ciphers for MPC and FHE. *IACR Cryptology ePrint Archive*, 2016:687, 2016.
- [BB17] Gustavo Banegas and Daniel J. Bernstein. Low-communication parallel quantum multi-target preimage search. *Cryptology ePrint Archive*, Report 2017/789, 2017. <http://eprint.iacr.org/2017/789>.
- [BCG⁺12] Julia Borghoff, Anne Canteaut, Tim Güneysu, Elif Bilge Kavun, Miroslav Knezevic, Lars R. Knudsen, Gregor Leander, Ventzislav Nikov, Christof Paar, Christian Rechberger, Peter Rombouts, Søren S. Thomsen, and Tolga Yalçın. PRINCE - a low-latency block cipher for pervasive computing applications - extended abstract. In *ASIACRYPT*, 2012.

- [Ber09] Daniel J. Bernstein. Cost analysis of hash collisions: Will quantum computers make SHARCS obsolete? 2009. <http://cr.yp.to/hash/collisioncost-20090823.pdf>.
- [BHT98] Gilles Brassard, Peter Høyer, and Alain Tapp. Quantum cryptanalysis of hash and claw-free functions. In Claudio L. Lucchesi and Arnaldo V. Moura, editors, *LATIN 1998*, volume 1380 of *LNCS*, pages 163–169. Springer, Heidelberg, April 1998.
- [Bih02] Eli Biham. How to decrypt or even substitute des-encrypted messages in 2^{28} steps. *Inf. Process. Lett.*, 84(3):117–124, 2002.
- [Ble00] Daniel Bleichenbacher. On the generation of one-time keys in dl signature schemes. Presentation at IEEE P1363 Working Group meeting, November 2000. Unpublished., 2000.
- [BMP13] Joan Boyar, Philip Matthews, and René Peralta. Logic minimization techniques with applications to cryptology. *Journal of Cryptology*, 26(2):280–312, 2013.
- [BMS05] Alex Biryukov, Sourav Mukhopadhyay, and Palash Sarkar. Improved time-memory trade-offs with multiple data. In *Selected Areas in Cryptography, 12th International Workshop, SAC 2005, Kingston, ON, Canada, August 11-12, 2005, Revised Selected Papers*, pages 110–127, 2005.
- [BPW12] David Bernhard, Olivier Pereira, and Bogdan Warinschi. How not to prove yourself: Pitfalls of the Fiat-Shamir heuristic and applications to Helios. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 626–643. Springer, Heidelberg, December 2012.
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *ACM CCS*, 1993.
- [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*, 2016.

- [CDG⁺17] Melissa Chase, David Derler, Steven Goldfeder, Claudion Orlandi, Sebastian Ramacher, Christian Rechberger, Daniel Slamanig, and Greg Zaverucha. Post-quantum zero-knowledge and signatures from symmetric-key primitives. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, Dallas, USA, October 30 - November 3, 2017*, 2017. to appear.
- [CDS94] Ronald Cramer, Ivan Damgård, and Berry Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In *CRYPTO*, 1994.
- [CGH98] Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited (preliminary version). In *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 209–218, 1998.
- [CGP⁺12] Claude Carlet, Louis Goubin, Emmanuel Prouff, Michaël Quisquater, and Matthieu Rivain. Higher-order masking schemes for s-boxes. In *FSE*, 2012.
- [Dam10] Ivan Damgård. On Σ -protocols. 2010. <http://www.cs.au.dk/~ivan/Sigma.pdf>.
- [DP08] Christophe De Cannière and Bart Preneel. Trivium. In *New Stream Cipher Designs - The eSTREAM Finalists*. 2008.
- [DPVAR00] Joan Daemen, Michaël Peeters, Gilles Van Assche, and Vincent Rijmen. Nesses proposal: Noekeon. In *First Open NESSIE Workshop*, 2000.
- [FKMV12] Sebastian Faust, Markulf Kohlweiss, Giorgia Azzurra Marson, and Daniele Venturi. On the non-malleability of the fiat-shamir transform. In *INDOCRYPT*, 2012.
- [FS86] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *CRYPTO*, pages 186–194, 1986.

- [GLSV14] Vincent Grosso, Gaëtan Leurent, François-Xavier Standaert, and Kerem Varici. Ls-designs: Bitslice encryption for efficient masked software implementations. In *FSE*, 2014.
- [GMO16a] Irene Giacomelli, Jesper Madsen, and Claudio Orlandi. ZKBoo: Faster zero-knowledge for boolean circuits. Cryptology ePrint Archive, Report 2016/163, 2016. <http://eprint.iacr.org/2016/163>.
- [GMO16b] Irene Giacomelli, Jesper Madsen, and Claudio Orlandi. ZKBoo: Faster zero-knowledge for boolean circuits. In *USENIX Security*, 2016.
- [GMS02] Steven D. Galbraith, John Malone-Lee, and Nigel P. Smart. Public key signatures in the multi-user setting. *Inf. Process. Lett.*, 83(5):263–266, 2002.
- [GP99] Louis Goubin and Jacques Patarin. DES and differential power analysis (the “duplication” method). In Çetin Kaya Koç and Christof Paar, editors, *CHES’99*, volume 1717 of *LNCS*, pages 158–172. Springer, Heidelberg, August 1999.
- [Gro96] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *STOC*, 1996.
- [Hel80] Martin Hellman. A cryptanalytic time-memory trade-off. *IEEE transactions on Information Theory*, 26(4):401–406, 1980.
- [HS01] Nick Howgrave-Graham and Nigel P. Smart. Lattice attacks on digital signature schemes. *Des. Codes Cryptography*, 23(3):283–290, 2001.
- [IKOS09] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge proofs from secure multiparty computation. *SIAM Journal on Computing*, 39(3):1121–1152, 2009.
- [Kat10] Jonathan Katz. *Digital Signatures*. Springer, 2010.
- [KM15] Neal Koblitz and Alfred J. Menezes. The random oracle model: a twenty-year retrospective. *Des. Codes Cryptography*, 77(2-3):587–610, 2015.

- [MJSC16] Pierrick Méaux, Anthony Journault, François-Xavier Standaert, and Claude Carlet. Towards stream ciphers for efficient FHE with low-noise ciphertexts. In *EUROCRYPT*, 2016.
- [MS04] Alfred Menezes and Nigel P. Smart. Security of signature schemes in a multi-user setting. *Des. Codes Cryptography*, 33(3):261–274, 2004.
- [NIS15] NIST. SHA-3 Standard: Permutation-Based Hash and Extendable-Output Functions. National Institute of Standards and Technology (NIST), FIPS PUB 202, U.S. Department of Commerce, 2015.
- [Unr12] Dominique Unruh. Quantum proofs of knowledge. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 135–152. Springer, Heidelberg, April 2012.
- [Unr15] Dominique Unruh. Non-interactive zero-knowledge proofs in the quantum random oracle model. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 755–784. Springer, Heidelberg, April 2015.
- [Unr16] Dominique Unruh. Computationally binding quantum commitments. In *EUROCRYPT*, 2016.
- [vOW94] Paul C. van Oorschot and Michael J. Wiener. Parallel collision search with application to hash functions and discrete logarithms. In *CCS '94, Proceedings of the 2nd ACM Conference on Computer and Communications Security, Fairfax, Virginia, USA, November 2-4, 1994.*, pages 210–218, 1994.

- $\text{Prove}_H(1^\kappa, y, x)$: 1. For each iteration $i \in [t]$: Sample random tapes $k_1^{(i)}, k_2^{(i)}, k_3^{(i)}$ and obtain output view $\text{View}_j^{(i)}$ and output share $y_j^{(i)}$. For each player P_j compute
- (a) $(x_1^{(i)}, x_2^{(i)}, x_3^{(i)}) \leftarrow \text{Share}(x, k_1^{(i)}, k_2^{(i)}, k_3^{(i)})$
 - (b) $\text{View}_j^{(i)} \leftarrow \text{Update}(\dots \text{Update}(x_j^{(i)}, x_{j+1}^{(i)}, k_j^{(i)}, k_{j+1}^{(i)}) \dots)$
 - (c) $y_j^{(i)} \leftarrow \text{Output}(\text{View}_j^{(i)})$
 - (d) Commit $C_j^{(i)} \leftarrow \text{Com}(k_j^{(i)}, x_j^{(i)}, \text{View}_j^{(i)}, y_j^{(i)})$, and let $\mathbf{a}^{(i)} \leftarrow (y_1^{(i)}, y_2^{(i)}, y_3^{(i)}, C_1^{(i)}, C_2^{(i)}, C_3^{(i)})$.
2. Compute the challenge: $\mathbf{e} \leftarrow H(\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(t)})$.
3. For each iteration $i \in [1, t]$ set: $\mathbf{z}^{(i)} \leftarrow (\text{View}_2^{(i)}, k_1^{(i)}, k_2^{(i)})$ if $\mathbf{e}^{(i)} = 1$ and $\mathbf{z}^{(i)} \leftarrow (\text{View}_{\mathbf{e}^{(i)}+1}^{(i)}, k_{\mathbf{e}^{(i)}}^{(i)}, k_{\mathbf{e}^{(i)}+1}^{(i)}, x_3^{(i)})$ otherwise, and return $\pi \leftarrow (\mathbf{e}, \mathbf{z}_{i \in [t]}^{(i)})$.

$\text{Verify}_H(1^\kappa, y, \pi)$: Parse π as $(\mathbf{e}, \mathbf{z}_{i \in [t]}^{(i)})$.

- 1. For each iteration $i \in [t]$ reconstruct the views, input and output shares that were not explicitly given as part of the proof $\mathbf{z}^{(i)}$:
 - (a) Set $x_{\mathbf{e}^{(i)}}^{(i)} \leftarrow R_{\mathbf{e}^{(i)}}(0)$ if $\mathbf{e}^{(i)} \neq 3$, otherwise obtain $x_3^{(i)}$ from $\mathbf{z}^{(i)}$. Set $x_{\mathbf{e}^{(i)}+1}^{(i)} \leftarrow R_{\mathbf{e}^{(i)}+1}(0)$ if $\mathbf{e}^{(i)} \neq 2$, otherwise obtain $x_3^{(i)}$ from $\mathbf{z}^{(i)}$.
 - (b) Obtain $\text{View}_{\mathbf{e}^{(i)}+1}^{(i)}$ from $\mathbf{z}^{(i)}$.
 - (c) $\text{View}_e^{(i)} \leftarrow \text{Update}(\dots \text{Update}(x_{\mathbf{e}^{(i)}}^{(i)}, x_{\mathbf{e}^{(i)}+1}^{(i)}, k_e^{(i)}, k_{\mathbf{e}^{(i)}+1}^{(i)}) \dots)$
 - (d) $y_{\mathbf{e}^{(i)}}^{(i)} \leftarrow \text{Output}(\text{View}_{\mathbf{e}^{(i)}}^{(i)})$, $y_{\mathbf{e}^{(i)}+1}^{(i)} \leftarrow \text{Output}(\text{View}_{\mathbf{e}^{(i)}+1}^{(i)})$
 - (e) $y_{\mathbf{e}^{(i)}+2}^{(i)} \leftarrow y + y_{\mathbf{e}^{(i)}}^{(i)} + y_{\mathbf{e}^{(i)}+1}^{(i)}$
- 2. Re-compute the commitments for views $\text{View}_{\mathbf{e}^{(i)}}^{(i)}$ and $\text{View}_{\mathbf{e}^{(i)}+1}^{(i)}$. For $j \in \{\mathbf{e}^{(i)}, \mathbf{e}^{(i)} + 1\}$: $C_j^{(i)} \leftarrow \text{Com}(k_j^{(i)}, x_j^{(i)}, \text{View}_j^{(i)}, y_j^{(i)})$.
- 3. Set $\mathbf{a}'^{(i)} \leftarrow (y_1^{(i)}, y_2^{(i)}, y_3^{(i)}, C_1^{(i)}, C_2^{(i)}, C_3^{(i)})$ taking $C_{\mathbf{e}^{(i)}+2}^{(i)}$ from $\mathbf{a}^{(i)}$.
- 4. Re-compute the challenge: $\mathbf{e}' \leftarrow H(\mathbf{a}'^{(1)}, \dots, \mathbf{a}'^{(t)})$. If $\mathbf{e} = \mathbf{e}'$ output Accept, otherwise Reject.

Scheme 4: The Fiat-Shamir transformed ZKB++ protocol.

Algorithm 1 LowMC encryption for key matrices $K_i \in \mathbb{F}_2^{n \times k}$ for $i \in [0, r]$, linear layer matrices $L_i \in \mathbb{F}_2^{n \times n}$ and round constants $C_i \in \mathbb{F}_2^n$ for $i \in [1, r]$.

Require: plaintext $p \in \mathbb{F}_2^n$ and key $y \in \mathbb{F}_2^k$

$s \leftarrow K_0 \cdot y + p$

for $i \in [1, r]$ **do**

$s \leftarrow Sbox(s)$

$s \leftarrow L_i \cdot s$

$s \leftarrow C_i + s$

$s \leftarrow K_i \cdot y + s$

end for

return s

$p \leftarrow \text{Sim}(x)$: In the simulator, we follow Unruh, and replace the initial state (before programming) of the random oracles with random polynomials of degree $2q - 1$ where q is an upper bound on the number of queries the adversary makes.

1. For $i \in [1, t]$, choose random $e^{(i)} \leftarrow \{1, 2, 3\}$. Let e be the corresponding binary string.
2. For each iteration $r_i, i \in [1, t]$: Sample random seeds $k_{e^{(i)}}^{(i)}, k_{e^{(i)}+1}^{(i)}$ and run the circuit decomposition simulator to generate $\text{View}_{e^{(i)}}^{(i)}, \text{View}_{e^{(i)}+1}^{(i)}$, output shares $y_1^{(i)}, y_2^{(i)}, y_3^{(i)}$, and if $e^{(i)} = 1$ $x_3^{(i)}$.

For $j = e^{(i)}, e^{(i)} + 1$ commit $[C_j^{(i)}, D_j^{(i)}] \leftarrow [H(k_j^{(i)}, \text{View}_j^{(i)}), k_j^{(i)} || \text{View}_j^{(i)}]$, and compute $g_j^{(i)} = G(k_j^{(i)}, \text{View}_j^{(i)})$.

Choose random $C_{e^{(i)}+2}, g_{e^{(i)}+2}^{(i)}$

Let $a^{(i)} = (y_1^{(i)}, y_2^{(i)}, y_3^{(i)}, C_1^{(i)}, C_2^{(i)}, C_3^{(i)})$. And $h^{(i)} = g_1^{(i)}, g_2^{(i)}, g_3^{(i)}$.

2. Set the challenge: program $H(a^{(1)}, \dots, a^{(t)}) := e$.
3. For each iteration $r_i, i \in [1, t]$: let $b^{(i)} = (y_{e^{(i)}+2}^{(i)}, C_{e^{(i)}+2}^{(i)})$ and set

$$z^{(i)} \leftarrow \begin{cases} (\text{View}_2^{(i)}, k_1^{(i)}, k_2^{(i)}) & \text{if } e^{(i)} = 1, \\ (\text{View}_3^{(i)}, k_2^{(i)}, k_3^{(i)}, x_3^{(i)}) & \text{if } e^{(i)} = 2, \\ (\text{View}_1^{(i)}, k_3^{(i)}, k_1^{(i)}, x_3^{(i)}) & \text{if } e^{(i)} = 3. \end{cases}$$

4. Output $p \leftarrow [e, (b^{(1)}, z^{(1)}), (b^{(2)}, z^{(2)}), \dots, (b^{(t)}, z^{(t)})]$.

Scheme 6: The zero knowledge simulator