



HAL
open science

Biscuit: New MPCitH Signature Scheme from Structured Multivariate Polynomials

Luk Bettale, Delaram Kahrobaei, Ludovic Perret, Javier Verbel

► **To cite this version:**

Luk Bettale, Delaram Kahrobaei, Ludovic Perret, Javier Verbel. Biscuit: New MPCitH Signature Scheme from Structured Multivariate Polynomials. Applied Cryptography and Network Security (ACNS), Mar 2024, Abu Dhabi, United Arab Emirates. pp.457-486, 10.1007/978-3-031-54770-6_18. hal-04293399

HAL Id: hal-04293399

<https://hal.sorbonne-universite.fr/hal-04293399>

Submitted on 18 Nov 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Biscuit: New MPCitH Signature Scheme from Structured Multivariate Polynomials

Luk Bettale¹, Delaram Kahrobaei^{2,3,4,5}, Ludovic Perret⁶, and Javier Verbel⁷

¹ IDEMIA, France

² Departments of Computer Science and Mathematics, Queens College, City University of New York, USA

³ Initiative for the Theoretical Sciences, Graduate Center, City University of New York, USA

⁴ Department of Computer Science, University of York, UK

⁵ Department of Computer Science and Engineering, Tandon School of Engineering, New York University, USA

⁶ Sorbonne University, CNRS, LIP6, PolSys, Paris, France

⁷ Technology Innovation Institute, UAE

Abstract. This paper describes **Biscuit**, a new multivariate-based signature scheme derived using the **MPCitH** approach. The security of **Biscuit** is related to the problem of solving a set of quadratic structured systems of algebraic equations. These equations are highly compact and can be evaluated using very few multiplications. The core of **Biscuit** is a rather simple **MPC** protocol which consists of the parallel execution of a few secure multiplications using standard optimized multiplicative triples. This paper also includes several improvements with respect to **Biscuit** submission to the last **NIST PQC** standardization process for additional signature schemes. Notably, we introduce a new hypercube variant of **Biscuit**, refine the security analysis with recent third-party attacks, and present a new **avx2** implementation of **Biscuit**.

Keywords: Post-Quantum · Digital Signature · MPC-in-the-Head · Multivariate Polynomials

1 Introduction

Biscuit is a new multivariate-based digital signature scheme submitted to the recent **NIST** standardization process for additional post-quantum signature schemes [1]. The security of **Biscuit** is proven assuming the hardness of the so-called **PowAff2** problem (Definition 1), which is a structured version of the well-known Multivariate Quadratic (MQ) problem [16].

Biscuit is in the lineage of the **Picnic** signature scheme [20,35], which was selected as an alternate candidate in the first **NIST** post-quantum cryptography standardization process [6]. The security of **Picnic** relies in the hardness of finding that secret key \mathbf{sk} used to produce plaintext-ciphertext pair using a particular block-cipher. The design of **Picnic** builds over a Multi-Party Computation (**MPC**) protocol to check that shared a triple of elements in a finite field (z, x, y) is multiplicative, i.e., the **MPC** protocol checks whether or not $z = xy$. Then it follows the **MPC-in-the-Head** (**MPCitH**) paradigm (introduced by Ishai, Kushilevitz, Ostrovsky, and Sahai in 2007 [28]) to obtain Zero-Knowledge Proof-of-Knowledge (**ZKPoK**) of the key \mathbf{sk} . Finally, the signature scheme is obtained by applying the Fiat-Shamir transformation [26] to the **ZKPoK** protocol.

As in **Picnic**, the design of **Biscuit** follows the **MPCitH** paradigm, and it uses the same **MPC** protocol to check multiplicative triples. Contrary to **Picnic**, in **Biscuit** the main goal is to build a **ZKPoK** of a pre-image \mathbf{s} of a quadratic system of multivariate polynomial equations \mathbf{f} over a finite field. The private and public keys in **Biscuit** are respectively \mathbf{s} and (\mathbf{f}, \mathbf{t}) , where $\mathbf{t} = \mathbf{f}(\mathbf{s})$.

The performance of **Picnic** is proportional to the number of multiplications required to evaluate the circuit defining the underlying encryption function with the secret key \mathbf{sk} . This fact motivates the use of a set \mathbf{f} of polynomial equations that require a small of multiplication to be evaluated. **Biscuit** considers polynomials of the form $f_i = A_0 + A_1 \cdot A_2$, where each A_i is linear polynomial. These polynomials can be evaluated computing only one multiplication, while a random quadratic polynomial would require as many multiplications as the number of variables.

1.1 Other Submitted MPCitH Signature Schemes

Since `Picnic`, the use of MPCitH for designing post-quantum signature schemes has become extremely popular. This is evidenced in the new NIST standardization process for post-quantum signature schemes⁸, where eight among forty of the submitted schemes are MPCitH-based. These schemes follow the same design methodology but differ on the hard problems considered.

AIMer is based on the hardness of key-recovery of a MPC-friendly block-cipher [31], MIRA and MiRitH are based on the MinRank problem [9,4], MQOM is based on the problem of solving random quadratic equations [23], PERK is based on the Permuted Kernel Problem [3], RYDE is based on the rank syndrome decoding problem [8], and SDith relies on the syndrome decoding problem [32]. All these schemes proposed several parameter sets parameter in order to optimize either the signature size (short variant) or the signing and verification times (fast variant).

Name	Performance (cycles)			Size (bytes)		
	keygen	sign	verify	sk	pk	σ
AIMer-L1PARAM4	54 435	78 022 625	73 813 256	16	32	3 840
MIRA-128s	112 000	46 800 000	43 900 000	16	84	5 640
MiRitH-Ias	108 903	41 220 707	40 976 634	16	129	5 673
MQOM-L1-gf31-short	67 000	44 360 000	41 720 000	78	47	6 352
PERK-I-short5	91 000	36 000 000	25 000 000	16	24	6 006
RYDE128s	33 100	23 400 000	20 100 000	32	86	5 956
SDith-L1-hyp	7 083 000	13 400 000	12 500 000	404	120	8 260
Biscuit-128s (this work)	62 484	27 922 077	28 484 726	16	68	5 748
FAEST-128s	200 000	25 580 000	25 830 000	32	32	5 006

Table 1: Performance of level I short variants of MPCitH-based candidates in the first round of the NIST new call for post-quantum signature schemes.

In Table 1, we overview a performance of the MPCitH NIST candidates with version of Biscuit described in this paper. In the table we also includes FAEST whose security is based on AES, and uses a new zero-knowledge technique, named VOLE-in-the head, that improves the MPCitH approach [13]. For each scheme, we report one short variant of achieving level I of security (i.e. equivalent to the security of AES128). The key-generation (`keygen`), signature generation (`sign`) and verification (`verify`) times are shown in number of clock-cycles (`cycles`). Table 1 also includes secret-key (`sk`), public-key (`pk`) and signature (`σ`) sizes in bytes.

Remark 1. Few days before finalizing this manuscript a new preprint appeared [24] that seems to significantly improve MQOM as well as many MPCitH-based signature schemes (including Biscuit).

1.2 Organization of the Paper and Main Results

After this introduction, the paper is organized as follows. Section 2 introduces basic notations, the new hard problem considered in Biscuit (`PowAff2` problem, Section 2.2), as well as the basic cryptography building blocks underlying its design: Multi-Party Computation (MPC), MPC-in-the-Head approach (MPCitH), Zero-Knowledge Proof of Knowledge (ZKPoK), BN-like proof systems and the hypercube technique for MPCitH-based signature schemes.

Section 3 describes the core sub-protocols underlying Biscuit. Due to the structure of the algebraic systems considered in Biscuit, the evaluation of a `PowAff2` solution requires only one multiplication per equation. This leads to a rather simple MPC protocol (Section 3.1) for `PowAff2` that is based on the parallel execution of secure multiplication using Beaver multiplicative triples [15] with some

⁸ <https://csrc.nist.gov/projects/pqc-dig-sig/round-1-additional-signatures>

optimizations from [14,30]. Then, we derive a new ZKPoK for PowAff2 (Section 3.2) using the MPCitH approach. Note that the protocol presented here (Figure 3) differs from the one described in the initial Biscuit submission [19]. In particular, we use the hypercube technique [33] and also include a security proof (Theorem 1) of the new ZKPoK.

Section 4 presents the Biscuit signature scheme and details the key generation (Figure 7), signature generation (Figure 8) and verification (Figure 9) algorithms. Biscuit is constructed using the traditional Fiat-Shamir transform from the ZKPoK described in Figure 3. We conclude this part with Table 2 that summarizes the secret-key, public-key, and signature sizes for the three security levels of NIST. In particular, Biscuit achieves a signature of 5.7KB for the first security level. This is comparable to other recent MPCitH-based signature schemes (Section 1.1).

Section 5 analyzes the security of the parameters proposed in Table 2. This section revisits the security analysis performed in the initial submission of Biscuit by taking into account new third-party analysis. In Section 5.1, we first explain the connection between the hardness of PowAff2 and the difficulty of solving the Learning With (bounded) Errors (LWE) problem [34]. In Section 5.2, we consider the key-recovery problem where the best attack against is a dedicated hybrid approach (i.e. that combines exhaustive search and Gröbner bases [18,17,12]) for solving PowAff2 equations described by Charles Bouillaguet on the NIST PQC⁹ mailing-list. In Section 5.3, we consider forgery attacks. In particular, we refine the analysis of Kales and Zaverucha [29] for forgery attacks against 5-pass Fiat-Shamir based signature schemes. This leads us to introduce a variant of the PowAff2 problem where the attacker has to solve a sub-system with fewer equations; leading to the introduction of the PowAff2_u problem (Definition 1).

Finally, Section 6 presents an optimized implementation of Biscuit which outperforms the previous implementation. First, we use a new canonical representation of the PowAff2 equations (Lemma 1) which allows us to further simplify their evaluation. Then, we integrate the hypercube framework for even further improvement.

2 Preliminaries

This section presents the preliminary concepts and notations used throughout the paper.

2.1 Notations

All over this paper, we use λ for the security parameter, $[n]$ refers to the set $\{1, \dots, n\}$ for an integer $n \in \mathbb{N}$, \mathbb{F}_q is the finite field of q elements (where q is prime or a prime power), \mathbb{F}_q^m denotes the vector space of dimension m over \mathbb{F}_q and $\mathbb{F}_q[x_1, \dots, x_n]$ the ring of polynomials in the variables x_1, \dots, x_n over the field \mathbb{F}_q .

Bold lower-case letters denote vectors, $\mathbf{x} + \mathbf{y}$ denotes the element-wise addition, $\mathbf{x} \odot \mathbf{y}$ denotes the element-wise product and \parallel the concatenation of two vectors or two-byte strings. $a \leftarrow \mathcal{A}(x)$ indicates that a is the output of an algorithm \mathcal{A} on input x , $a \xleftarrow{\$} \mathcal{S}$ means that a is sampled uniformly at random from a set \mathcal{S} .

Let \mathcal{R} be a ring and $a \in \mathcal{R}$. *additive sharing* of a , denotes $\llbracket a \rrbracket$ is a tuple $\llbracket a \rrbracket := (\llbracket a \rrbracket_1, \dots, \llbracket a \rrbracket_N) \in \mathcal{R}^N$ such that $a = \sum_{i=1}^N \llbracket a \rrbracket_i$. Each component $\llbracket a \rrbracket_i$ of $\llbracket a \rrbracket$ is called a *share* of a . Throughout this paper, we only consider additive sharing and use the word sharing to refer to additive sharing.

A *Multi-Party Computation* (MPC) protocol is an interactive protocol executed by a set of N parties knowing a public function f . Its goal is to compute the image $z = f(x_1, \dots, x_N)$, where the value x_i is only known by the i -th party. An MPC protocol is considered secure and correct if, at the end of the protocol, every party i knows z , and no information about its secret input value x_i is revealed to the other parties.

⁹ https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/sw8NueiNek0/m/2sa_emjABQAJ

2.2 The PowAff2_u Problem

The core problem considered in Biscuit is the one of solving a system of multivariate equations defined as the product of two affine forms. This problem, denoted as PowAff2_u, is parameterized by a tuple of positive integers (n, m, u, q) , where n is the number of variables, m the number of equations, u is a parameter related to forgery (Section 5.3), and q is the finite field size.

Definition 1 (The PowAff2_u problem). Let $A_{1,0}, A_{1,1}, A_{1,2}, \dots, A_{m,0}, A_{m,1}, A_{m,2} \in \mathbb{F}_q[x_1, \dots, x_n]$ be affine forms, i.e.:

$$A_{k,j}(x_1, \dots, x_n) = a_0^{(k,j)} + \sum_{i=1}^n a_i^{(k,j)} x_i, \text{ with } a_0^{(k,j)}, \dots, a_n^{(k,j)} \in \mathbb{F}_q. \quad (1)$$

Input. A vector $\mathbf{t} = (t_1, \dots, t_m) \in \mathbb{F}_q^m$ and multivariate polynomials $\mathbf{f} = (f_1, \dots, f_m) \in \mathbb{F}_q[x_1, \dots, x_n]^m$ defined as:

$$f_k(x_1, \dots, x_n) = A_{k,0}(x_1, \dots, x_n) + \prod_{j=1}^2 A_{k,j}(x_1, \dots, x_n), \forall k \in [m]. \quad (2)$$

Question. Find – if any – a vector $(s_1, \dots, s_n) \in \mathbb{F}_q^n$ and set $J \subseteq [m]$ of size $m - u$ such that:

$$f_j(s_1, \dots, s_n) = t_j, \forall j \in J.$$

Definition 2 (The PowAff2 problem). We use PowAff2 to denote the PowAff2₀ problem. Also, we shall call PowAff2 algebraic system the set of non-linear equations $f_1, \dots, f_m \in \mathbb{F}_q[x_1, \dots, x_n]$ defined as in (2).

PowAff2 is the problem corresponding to key-recovery whilst PowAff2_u, with $u > 0$, is a relaxation that corresponds to signature forgery.

A detailed hardness analysis of PowAff2_u is provided in Section 5. The current best attack against Biscuit has been sketched by Charles Bouillaguet on the NIST PQC mailing-list¹⁰. In particular, it was mentioned that the multivariate equations defined as in Definition 1 can be reduced to a simple, but equivalent, structure.

Lemma 1. Let $\mathbf{f} = (f_1, \dots, f_m) \in \mathbb{F}_q[x_1, \dots, x_n]^m$ be a PowAff2 algebraic system. Then, with high probability, there exists an invertible matrix $\mathbf{L} \in \text{GL}_n(\mathbb{F}_q)$ such that :

$$\begin{aligned} \mathbf{f}(\mathbf{x} \cdot \mathbf{L}) &= (u_1(\mathbf{x}) \cdot (x_1 + c_1) + w_1(\mathbf{x}), \dots, u_n(\mathbf{x}) \cdot (x_n + c_n) + w_n(\mathbf{x}), \\ &A'_{n+1,0}(\mathbf{x}) + \prod_{j=1}^2 A'_{n+1,j}(\mathbf{x}), \dots, A'_{m,0}(x_1, \dots, x_n) + \prod_{j=1}^2 A'_{m,j}(\mathbf{x})) \end{aligned}$$

where $\mathbf{x} = (x_1, \dots, x_n)$, $A_{n+1,0}, A_{n+1,1}, A_{n+1,2}, \dots, A_{m,0}, A_{m,1}, A_{m,2}, u_1, \dots, u_n, v_1, \dots, v_n \in \mathbb{F}_q[x_1, \dots, x_n]$ are affine polynomials and $c_1, \dots, c_n \in \mathbb{F}_q$.

Proof. By construction, we have :

$$f_k(x_1, \dots, x_n) = A_{k,0} + \prod_{j=1}^2 A_{k,j}, \forall k \in [m],$$

with $A_{1,0}, A_{1,1}, A_{1,2}, \dots, A_{m,0}, A_{m,1}, A_{m,2} \in \mathbb{F}_q[x_1, \dots, x_n]$ affine forms as defined in (1).

Thus, we can write $A_{k,2}(x_1, \dots, x_n) = (x_1, \dots, x_n) \cdot \mathbf{b}_k + c_k$, where $\mathbf{b}_k = (a_1^{(k,2)}, \dots, a_n^{(k,2)}) \in \mathbb{F}_q^n$ and $c_k = a_0^{(k,2)} \in \mathbb{F}_q$. Let $\mathbf{C} \in \mathbb{F}_q^{n \times n}$ be the matrix whose rows are $\mathbf{b}_1, \dots, \mathbf{b}_n$. We want to find a non-singular matrix $\mathbf{L} \in \text{GL}_n(\mathbb{F}_q)$ such that $\mathbf{I}_n = \mathbf{C} \cdot \mathbf{L}$, where \mathbf{I}_n is the identity matrix of size n . This reduces to compute, if any, the inverse of \mathbf{C} .

¹⁰ https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/sw8NueiNek0/m/2sa_emjABQAJ

2.3 Digital Signature Scheme

Definition 3. A Digital Signature Scheme (DSS) is a tuple of three probabilistic polynomial-time algorithms (KeyGen , Sign , Verify) verifying:

1. $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^\lambda)$. The key-generation algorithm KeyGen takes as input a security parameter 1^λ and outputs a pair of public/private keys (pk, sk) .
2. $\sigma \leftarrow \text{Sign}(\text{sk}, \text{msg})$. The signing algorithm Sign takes a private key sk and a message $\text{msg} \in \{0, 1\}^*$ and outputs a signature σ .
3. $b \leftarrow \text{Verify}(\text{pk}, \sigma, \text{msg})$. The verification algorithm Verify is deterministic. It takes as input a message $\text{msg} \in \{0, 1\}^*$, a signature σ , and a public key pk . It outputs a bit $b \in \{0, 1\}$: 1 means that it **accepts** σ as a valid signature for msg , otherwise it **rejects** outputting 0.

A signature scheme is correct if for every security parameter $\lambda \in \mathbb{N}$, every $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^\lambda)$, and every message $\text{msg} \in \{0, 1\}^*$, it holds that

$$1 \leftarrow \text{Verify}(\text{pk}, \text{msg}, \text{Sign}(\text{sk}, \text{msg})).$$

The standard security notion for a DSS is Existential Unforgeability under Adaptive Chosen-Message Attacks (EU-CMA). We say that a signature scheme is EU-CMA-secure if for all probabilistic polynomial-time adversaries \mathcal{A} , the probability

$$\Pr \left[1 \leftarrow \text{Verify}(\text{pk}, \text{msg}^*, \sigma^*) \mid \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^\lambda) \\ (\text{msg}^*, \sigma^*) \leftarrow \mathcal{A}^{\mathcal{O}_{\text{Sign}(\text{sk}, \cdot)}}(\text{pk}) \end{array} \right]$$

is a negligible function in λ , where \mathcal{A} is given access to a signing oracle $\mathcal{O}_{\text{Sign}(\text{sk}, \cdot)}$, and msg^* has not been queried to $\mathcal{O}_{\text{Sign}(\text{sk}, \cdot)}$.

Auxiliary Functions. Biscuit also relies on further basic cryptographic building blocks that we do not explicitly introduce such as commitments, collision-resistant hash functions, key-derivation functions, and pseudo-random number generators. As explained in [19], we can use the SHAKE256 [21] extendable-output function (XOF) to instantiate these functions.

During signature, the signer must generate a set of N seeds and reveal $N - 1$ of them to the verifier for each iteration (TreePRG). The verifier then uses these seeds to check that the MPC protocol was correctly simulated. A binary tree structure allows generating the seeds using one root seed from a binary tree. Instead of sending $N - 1$ seeds in the signature, this allows sending only $\lceil \log_2 N \rceil$ seeds that will be used to reconstruct all $N - 1$ seeds required. We refer to [19] for the description of TreePRG.

2.4 5-Pass Identification Schemes

An IDentification Scheme (IDS) is an interactive protocol between a *prover* P and a *verifier* V , where P wants to prove its knowledge of a secret value sk to V , with (pk, sk) satisfying a given relation, for a public value pk .

Definition 4 (5-pass identification scheme). A 5-pass IDS is a tuple of three probabilistic polynomial-time algorithms (KeyGen , P , V) such that

1. $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^\lambda)$. The key-generation algorithm KeyGen takes as input a security parameter 1^λ and outputs a pair of public/private keys (pk, sk) .
2. P and V follow the protocol in Figure 1, and at the end of this, V outputs 1, if it **accepts** that P knows sk , otherwise it **rejects** outputting 0.

A *transcript* of a 5-pass IDS is a tuple $(\text{com}, \text{ch}_1, \text{rsp}_1, \text{ch}_2, \text{rsp}_2)$, as in Figure 1, referring to all the messages exchanged between P and V in one execution of the IDS. We require an IDS to fulfill the following security properties.

Correctness: if for any security parameter $\lambda \in \mathbb{N}$ and $(pk, sk) \leftarrow \text{KeyGen}(1^\lambda)$ it holds, $\Pr[1 \leftarrow V(pk, com, ch_1, rsp_1, ch_2, rsp_2)] = 1$, where $(com, ch_1, rsp_1, ch_2, rsp_2)$ is the transcript of an execution of the protocol between $P(pk, sk)$ and $V(pk)$.

Soundness (with soundness error ε): if, given a key pair (pk, sk) , for every polynomial-time adversary \mathcal{A} such that

Honest-verifier zero-knowledge: if there exists a probabilistic polynomial-time *simulator* $\mathcal{S}(pk)$ that outputs a transcript $(com, ch_1, rsp_1, ch_2, rsp_2)$ from a distribution that is computationally indistinguishable from the distribution of transcripts of an honest execution of the protocol between $P(pk, sk)$ and $V(pk)$.




Fig. 1: Canonical 5-pass IDS.

2.5 MPC-in-the-Head : From MPC to Zero-Knowledge

MPC-in-the-Head (MPCiTH) is a generic technique, introduced as “IKOS” [28]), that allows to build a Zero-Knowledge Proof of Knowledge (ZKPoK) from a secure MPC protocol.

Consider a MPC protocol where N parties $P_1 \dots, P_N$ collaborate to securely evaluate a public function f on a secret input x . Assuming that the protocol is perfectly correct and that the views of $t < N$ parties leak no information on x , then one can construct a ZKPoK from the MPC protocol as follows:

1. Simulation.

- Prover P generates a random sharing $\llbracket x \rrbracket := (\llbracket x \rrbracket_1, \dots, \llbracket x \rrbracket_N)$ of x such that $x = \sum_{i=1}^N \llbracket x \rrbracket_i$ and assign a share $\llbracket x \rrbracket_i$ to each party P_i .
- P emulates “in his/her mind” execution of the MPC protocol with N parties $P_1 \dots, P_N$.
- P commits on the *views* of each P_i , meaning the messages they send/receive during the protocol execution and their internal states. These commitments are sent to the verifier V .

2. Challenges.

- P possibly receives random challenges from V on the MPC, executes local computations accordingly and send the results to V . This step can be repeated several times.
- V challenges P to open a random subset of t parties.
- P returns the requested views.

3. Verification.

- P then checks that the views¹¹ are consistent, and the output of the circuit corresponds to the result expected.

¹¹ If only one party is opened then there are no pairs to check consistency. In this case, the prover does not commit to the views, but actually to the point-to-point channels between the parties.

Since its introduction, the initial approach for MPCitH from [28] has been improved in different ways. In particular, Katz, Kolesnikov and Wang (KKW) extended the MPCitH paradigm to support what is known in MPC as the *preprocessing model*, where MPC protocols are split into an offline phase that is independent of the sensitive inputs, and an online phase, with the former being typically the bottleneck in terms of efficiency. The benefit is that the prover does not need to include the preprocessing as part of the views of the parties, and instead, the preprocessing can be checked. As an application, KKW allowed to significantly decrease the signature size of the Picnic initial version.

Also, [33] described a so-called hypercube variant of MPCitH that allows improving efficiency for a large number of parties in the MPC protocol. Indeed, a large number of parties leads to shorter signatures but increases signature generation and verification times. We detail the approach in the case of Biscuit in Section 3.1.

2.6 Proof Systems for Arbitrary Circuits

In [27], Giacomelli, Madsen and Orlandi demonstrated the efficiency of the MPCitH approach for generating ZKPoK. Doing so, the authors also introduced a new generic proof system, called ZKBoo, which ultimately resulted in the first version of the Picnic signature scheme. In such work, the virtual/emulated parties actually *execute* some MPC protocols, and the verifier checks this execution. In [14], Baum and Nof proposed an improved proof system, called BN, for arithmetic circuits. [14] observed that the prover knows all the wire values in the circuit, and instead of computing a protocol, the prover can distribute sharings for each intermediate wire value, and the virtual parties only need to execute a protocol that checks the correctness of the multiplication gates. This allows batching the checks by taking random linear combinations. In [30], Kales and Zaverucha built on top of BN and introduced BN++ that includes several optimizations, leading to a roughly 2.5× communication improvement. The authors apply their techniques to AES as well as LowMC (Picnic signature).

The BN and BN++ proof systems rely on the concept of multiplicative triple (or Beaver triple [15]). Given $x, y, z \in \mathbb{F}_q$, we say that the triple $(\llbracket x \rrbracket, \llbracket y \rrbracket, \llbracket z \rrbracket) \in \mathbb{F}_q^N \times \mathbb{F}_q^N \times \mathbb{F}_q^N$ is a *multiplicative triple* if it holds that $z = x \cdot y$. The Biscuit MPC protocol will rely on a somewhat standard protocol introduced in [14] (along with the optimization given in [30, Section 2.5]) to check multiplicative triples of sharing (Section 2.6). A multiplication triple $(\llbracket x \rrbracket, \llbracket y \rrbracket, \llbracket z \rrbracket) \in \mathbb{F}_q^N \times \mathbb{F}_q^N \times \mathbb{F}_q^N$ can be checked using a *correlated* triple $(\llbracket a \rrbracket, \llbracket y \rrbracket, \llbracket c \rrbracket) \in \mathbb{F}_q^N \times \mathbb{F}_q^N \times \mathbb{F}_q^N$ with $a \in \mathbb{F}_q$ and $c = a \cdot y \in \mathbb{F}_q$ as follows:

1. The parties get a random element $\varepsilon \xleftarrow{\$} \mathbb{F}_q$.
2. The parties locally set $\llbracket \alpha \rrbracket \leftarrow \llbracket x \rrbracket \cdot \varepsilon + \llbracket a \rrbracket$.
3. The parties open $\llbracket \alpha \rrbracket$ so that they all obtain α .
4. The party locally compute $\llbracket v \rrbracket = \llbracket y \rrbracket \cdot \alpha - \llbracket z \rrbracket \cdot \varepsilon - \llbracket c \rrbracket$.
5. The parties open $\llbracket v \rrbracket$ to obtain v .
6. The parties output **accept** if $v = 0$ and **reject** otherwise.

The security of this simple protocol has been proven in [30]. In particular, the false success probability is given by:

Lemma 2. *Let $x, y, z, a, c \in \mathbb{F}_q$. If the shared multiplicative triple $(\llbracket x \rrbracket, \llbracket y \rrbracket, \llbracket z \rrbracket) \in \mathbb{F}_q^N \times \mathbb{F}_q^N \times \mathbb{F}_q^N$ is incorrect, i.e. $z \neq x \cdot y$, or the helping multiplicative triple $(\llbracket a \rrbracket, \llbracket y \rrbracket, \llbracket c \rrbracket) \in \mathbb{F}_q^N \times \mathbb{F}_q^N \times \mathbb{F}_q^N$ is incorrect, i.e. $c \neq a \cdot y$, then the parties outputs **accept** with probability at most $1/q$.*

3 Interactive Protocols for PowAff2

This section describes the MPC protocol underlying Biscuit (Section 3.1) and the corresponding ZKPoK (Section 3.2) obtained using the MPCitH paradigm (Section 2.5) together with the hypercube technique [5].

3.1 Multi-Party Computation Protocol for PowAff2

In Figure 2, we detail the MPC protocol used in Biscuit to check a solution of a PowAff2 algebraic system. The protocol is executed by N parties sharing a secret vector $\mathbf{s} \in \mathbb{F}_q^n$. Every party knows the target vector

$\mathbf{t} = (t_1, \dots, t_m) \in \mathbb{F}_q^n$, affine forms $A_{1,0}, A_{1,1}, A_{1,2}, \dots, A_{m,0}, A_{m,1}, A_{m,2} \in \mathbb{F}_q[x_1, \dots, x_n]$ as in (1) and the corresponding PowAff2 algebraic equations $\mathbf{f} = (f_1, \dots, f_m) \in \mathbb{F}_q[x_1, \dots, x_n]^m$ defined as:

$$f_k = A_{k,0} + A_{k,1} \cdot A_{k,2}, \forall k \in [m]. \quad (3)$$

The MPC protocol (Figure 2) consists of m iterations of the multiplicative checking protocol described in Section 2.6). At the end of the protocol, the parties output **accept** indicating they are convinced that the shared vector \mathbf{s} satisfies $\mathbf{t} = \mathbf{f}(\mathbf{s})$. Otherwise, they output **reject**.

Public data: $\mathbf{t} = (t_1, \dots, t_m) \in \mathbb{F}_q^m$, affine polynomials $A_{1,0}, \dots, A_{m,2} \in \mathbb{F}_q[x_1, \dots, x_n]$ and $\mathbf{f} = (f_1, \dots, f_m) \in \mathbb{F}_q[x_1, \dots, x_n]^m$ as defined in (3).

Inputs : The i -th party knows $[[\mathbf{s}]]_i \in \mathbb{F}_q^n$, $[[\mathbf{a}]]_i \in \mathbb{F}_q^m$ where $\mathbf{a} = (a_1, \dots, a_m) \xleftarrow{\$} \mathbb{F}_q^m$, and $[[\mathbf{c}]]_i \in \mathbb{F}_q^m$ where $\mathbf{c} = (c_1, \dots, c_m) \in \mathbb{F}_q^m$ such that $c_k = A_{k,2}(\mathbf{s}) \cdot a_k, \forall k \in [m]$.

MPC protocol:

for $k \in [m]$ **do**

- 1 : The parties locally compute $[[z_k]] \leftarrow t_k - A_{k,0}([[s]]), [[x_k]] \leftarrow A_{k,1}([[s]])$, and $[[y_k]] \leftarrow A_{k,2}([[s]])$.
- 2 : The parties get a random element $\varepsilon_k \xleftarrow{\$} \mathbb{F}_q$.
- 3 : The parties locally set $[[\alpha_k]] \leftarrow [[x_k]] \cdot \varepsilon_k + [[a_k]]$.
- 4 : The parties open $[[\alpha_k]]$ so that they all obtain α_k .
- 5 : The party locally compute $[[v_k]] = [[y_k]] \cdot \alpha_k - [[z_k]] \cdot \varepsilon_k - [[c_k]]$.
- 6 : The parties open $[[v_k]]$ to obtain v_k .

The parties output **accept** if $v_k = 0, \forall k \in [m]$ and **reject** otherwise.

Fig. 2: MPC protocol Π to check that $\mathbf{t} = \mathbf{f}(\mathbf{s})$.

The following proposition follows easily from Lemma 2.

Proposition 1. *Suppose that a set of N parties genuinely follow the MPC protocol given in Figure 2 with inputs $\mathbf{t} \in \mathbb{F}_q^m, \mathbf{f} = (f_1, \dots, f_m) \in \mathbb{F}_q[x_1, \dots, x_n]^m$, and $[[\mathbf{s}]] \in (\mathbb{F}_q^n)^N$. Suppose $\mathbf{s} \in \mathbb{F}_q^n$ is a solution to $\text{PowAff2}_u(\mathbf{f}, \mathbf{t})$ but not a solution to the $\text{PowAff2}_{u-1}(\mathbf{f}, \mathbf{t})$. If $u = 0$, i.e., $\mathbf{t} = \mathbf{f}(\mathbf{s})$, then the parties **accept**. Otherwise, the parties **accept** with probability at most $1/q^u$.*

3.2 ZKPoK for PowAff2

In Figure 3, we derive a ZKPoK for the PowAff2 problem using the MPC protocol Π of Figure 2. To do so, we use the traditional MPCitH approach combined with the recent hypercube technique.

In Phase 1, for each $\ell \in [D]$ the prover generates an input set $S_\ell = ([[\mathbf{s}]]_{(\ell,j)}, [[\mathbf{c}]]_{(\ell,j)}, [[\mathbf{a}]]_{j \in [2]})$ for a two parties instance the MPC protocol Π (Figure 2). The set S_ℓ is called the ℓ -th set of *main shares*. The sets of main shares are computed in two steps. First, the prover generates and commits to inputs $([[\mathbf{s}]]_i, [[\mathbf{c}]]_i, [[\mathbf{a}]]_i)$ of one $N = 2^D$ parties instance of Π . Then, for each $(\ell, j) \in [D] \times [2]$, the main share $[[\mathbf{s}]]_{(\ell,j)}$ is computed as the sum of the shares $[[\mathbf{s}]]_i$ for which j equals the ℓ -th bit of i plus 1. Similarly, the main shares $[[\mathbf{c}]]_{(\ell,j)}$ and $[[\mathbf{a}]]_{(\ell,j)}$. In Phase 3, the prover executes the protocol Π for every set of main shares using $\varepsilon_1, \dots, \varepsilon_m \in \mathbb{F}_q$ as the random elements for all D executions. This particular execution of the protocol Π on the set of main shares S_ℓ is shown in Figure 4. The output of each execution the shares $([[\alpha_k]]_{(\ell,j)}, [[v_k]]_{(\ell,j)})_{(k,j) \in [m] \times [2]}$ of the opened values α_k and v_k , and its corresponding hash $H_\ell = \mathbb{H} (([[\alpha_k]]_{(\ell,j)}, [[v_k]]_{(\ell,j)})_{(k,j) \in [m] \times [2]})$ ¹². In Phase 5, the prover sends $((\text{seed}^{(i)}, \rho_i)_{i \neq \bar{i}}, \text{com}^{(\bar{i})}, \Delta \mathbf{s}, \Delta \mathbf{c}, [[\alpha]]_{\bar{i}})$ sends to the verifier, where $[[\alpha]]_{\bar{i}} = ([[\alpha_1]]_{\bar{i}}, \dots, [[\alpha_m]]_{\bar{i}})$, $[[\alpha_k]]_{\bar{i}} = [[x_k]]_{\bar{i}} \cdot \varepsilon_k + [[a_k]]_{\bar{i}}$ and $[[x_k]]_{\bar{i}} = A_{k,1}([[s]]_{\bar{i}})$. We

¹² As noted in [10], the security of proof knowledge protocols using the hypercube technique with additive shares is the same with or without these intermediate hash values H_ℓ , but it might help to reduce the memory demand of the protocol when the implementation of the hash \mathbb{H} is not incremental.

highlight that the prover does not send explicitly instead of sending $N - 1$ strings of the form $(\text{seed}^{(i)}, \rho_i)$ but it sends instead the $\log_2(N)$ nodes of the tree $\text{TreePRG}(\text{root})$ so that the verifier can recompute the values $(\text{seed}^{(i)}, \rho_i)_{i \neq \bar{i}}$. Finally, in the verification phase, the verifier recomputes $(\text{seed}^{(i)}, \rho_i)_{i \neq \bar{i}}$, and uses them to recompute the sets main shares partially. We say partially recompute and not just recompute because for each set S_ℓ one of the main shares triples (either the one corresponding to $j = 1$ or $j = 2$) is missing the addition of the shares corresponding to the \bar{i} -th party. After, the verifier, for every set of main parties, follows the algorithm in Figure 5 to check the execution of the MPC protocol Π . Finally, the verifier recomputes h_0 and h_2 and outputs accept if these two values match the ones the prover sent. Otherwise, the verifier rejects.

Security proofs of the zero-knowledge protocol. The following theorem establishes the security properties of the proof of knowledge protocol described in Figure 3.

Theorem 1. *The PowAff2 affine zero-knowledge protocol described in Figure 3 has the following properties:*

- **Completeness:** *A Prover with the knowledge of a solution to an instance (\mathbf{f}, \mathbf{t}) of the PowAff2 is always accepted by the Verifier.*
- **Soundness:** *Let $\epsilon = \frac{1}{N} + \frac{1}{q^u} \cdot (1 - \frac{1}{N})$, where $p = 1/q^u$. Suppose there exists a prover $\tilde{\mathcal{P}}$ who convinces the verifier to accept with probability $\tilde{\epsilon} > \epsilon$. Then there is an efficient probabilistic extraction algorithm \mathcal{E} , which has rewindable black-box access to $\tilde{\mathcal{P}}$, that, in expectation, with at most*

$$\frac{4}{\tilde{\epsilon} - \epsilon} \cdot \left(1 + \tilde{\epsilon} \cdot \frac{2 \ln(2)}{\tilde{\epsilon} - \epsilon} \right).$$

calls to $\tilde{\mathcal{P}}$ outputs either a solution to an instance (\mathbf{f}, \mathbf{t}) of the PowAff2 $_{u-1}$ problem or a collision to the commitment scheme Com or the hash H.

- **Honest-verifier zero-knowledge:** *If the pseudo-random generator PRG and the commitment scheme com are indistinguishable from the uniform random distribution, then the algorithm in Figure 3 honest-verifier zero-knowledge.*

Proof. (sketch) Here follow the same reasoning as the one in [10, Theorem 1]. We highlight the main parts of the proof. The details follow along the lines with [10, Theorem 1]

- **Completeness:** By following, step by step, the protocol in Figure 3 it is not hard to see that a Prover that follows the protocol with inputs $(\mathbf{f}, \mathbf{t}, \mathbf{s})$ such that $\mathbf{t} = \mathbf{f}(\mathbf{s})$ will always be accepted.
- **Soundness:** The structure of the proof will be the following:
 1. We prove that a prover $\tilde{\mathcal{P}}$ who has not a solution for the PowAff2 $_{u-1}$ problem can cheat with probability at most $\epsilon = \frac{1}{N} + \frac{1}{q^u} \cdot (1 - \frac{1}{N})$.
 2. Assuming that
 - (a) No collisions to Com nor H can be found.
 - (b) There exists a cheater $\tilde{\mathcal{P}}$ who has cheating probability $\tilde{\epsilon} > \epsilon$.

We show how to extract a solution to the PowAff2 $_{u-1}$ problem whenever rewindable black-box access to $\tilde{\mathcal{P}}$ is given.

For part 1, suppose that at step 7 the vector $\mathbf{s} = \llbracket \mathbf{s} \rrbracket_1 + \dots + \llbracket \mathbf{s} \rrbracket_N$ is not a solution of the PowAff2 $_{u-1}$ problem defined by (\mathbf{f}, \mathbf{t}) . With such a vector \mathbf{s} the prover can be accepted by the verifier in only two situations:

- The prover honestly follows the protocol, and for each $k \in [m]$, the value $v_k = y_k \alpha_k - z_k \varepsilon_k - c_k$, which is the value that would be obtained from a genuine executing of the MPC protocol with challenges ε_k (see Figure 2), equals zero, or
- The prover dishonestly deviates from the protocol, yet the verifier believes that all the honest v_k are zero, but in reality, at least one of them is not.

In the first situation, we would have a false positive case of the MPC protocol in Figure 2. By Proposition 1, this happens with probability at most $1/q^u$. In the second situation, the prover cheats during the simulation of at least one party. Since the verifier checks the correct execution of all the parties but one, the prover has to cheat on exactly one party. Otherwise, the verifier rejects. Cheating in one party i' means that the prover uses a set of different shares than an honest party, holding the same input seed $\text{seed}^{(i')}$, would use. Since every party aggregates to exactly one of the main shares for all of the D bi-party protocols. For each one of these bi-party protocols, there is one share that has been dishonestly computed, i.e., not following the MPC protocol. Thus, the prover won't be detected with

<p>PoK(Prover(f, t, s), Verifier(f, t))</p> <hr/> <p>Phase 1: Prover commits to the inputs of the MPC protocol in Figure 4</p> <p>1 : $\text{root} \xleftarrow{\\$} \{0, 1\}^\lambda, \quad (\text{seed}^{(i)}, \rho^{(i)})_{i \in [N]} \leftarrow \text{TreePRG}(\text{root})$</p> <p style="padding-left: 20px;">for $i \in [N]$ do</p> <p>2 : $\llbracket \mathbf{s} \rrbracket_i, \llbracket \mathbf{c} \rrbracket_i, \llbracket \mathbf{a} \rrbracket_i \leftarrow \text{PRG}(\text{seed}^{(i)})$</p> <p>3 : $\text{com}^{(i)} \leftarrow \text{Com}(\text{seed}^{(i)}, \rho_i)$</p> <p>4 : $h_0 \leftarrow \mathbb{H}(\text{com}^{(1)}, \dots, \text{com}^{(N)})$, and send h_0 to Verifier</p> <p>5 : $\mathbf{a} \leftarrow \sum_{i \in [N]} \llbracket \mathbf{a} \rrbracket_i, \quad \mathbf{c} \leftarrow (A_{k,2}(\mathbf{s}) \cdot a_k)_{k \in [m]}$</p> <p>6 : $\Delta \mathbf{s} \leftarrow \mathbf{s} - \sum_{i \in [N]} \llbracket \mathbf{s} \rrbracket_i, \quad \Delta \mathbf{c} \leftarrow \mathbf{c} - \sum_{i \in [N]} \llbracket \mathbf{c} \rrbracket_i$</p> <p>7 : $\llbracket \mathbf{s} \rrbracket_1 \leftarrow \llbracket \mathbf{s} \rrbracket_1 + \Delta \mathbf{s}$ and $\llbracket \mathbf{c} \rrbracket_1 \leftarrow \llbracket \mathbf{c} \rrbracket_1 + \Delta \mathbf{c}$</p> <p>8 : Initialize $\llbracket \mathbf{s} \rrbracket_p, \llbracket \mathbf{c} \rrbracket_p$ and $\llbracket \mathbf{a} \rrbracket_p$ to zero objects for each $p \in [D] \times [2]$</p> <p style="padding-left: 20px;">for $i \in [N]$ do</p> <p>9 : $(i_1, \dots, i_D) \leftarrow i$ // Binary representation of i.</p> <p style="padding-left: 40px;">for $\ell \in [D]$ do</p> <p>10 : $\llbracket \mathbf{s} \rrbracket_{(\ell, i_\ell+1)} \leftarrow \llbracket \mathbf{s} \rrbracket_{(\ell, i_\ell+1)} + \llbracket \mathbf{s} \rrbracket_i, \llbracket \mathbf{c} \rrbracket_{(\ell, i_\ell+1)} \leftarrow \llbracket \mathbf{c} \rrbracket_{(\ell, i_\ell+1)} + \llbracket \mathbf{c} \rrbracket_i$ and</p> <p>11 : $\llbracket \mathbf{a} \rrbracket_{(\ell, i_\ell+1)} \leftarrow \llbracket \mathbf{a} \rrbracket_{(\ell, i_\ell+1)} + \llbracket \mathbf{a} \rrbracket_i$</p> <p>Phase 2: First challenge</p> <p>12 : Verifier samples $\varepsilon_1, \dots, \varepsilon_m \xleftarrow{\\$} \mathbb{F}_q$ and sends them to Prover</p> <p>Phase 3: Prover's first response // Prover executes MPC protocol for every set of main shares.</p> <p style="padding-left: 20px;">for $\ell \in [D]$ do</p> <p>13 : Prover gets H_ℓ and $(\llbracket \alpha_k \rrbracket_{(\ell, j)}, \llbracket v_k \rrbracket_{(\ell, j)})_{(k, j) \in [m] \times [2]}$ from the algorithm in Figure 4</p> <p>14 : $h_1 \leftarrow \mathbb{H}(H_1, \dots, H_D)$ and send h_1 to Verifier</p> <p>Phase 4: Second challenge</p> <p>15 : Verifier samples $\bar{i} \xleftarrow{\\$} [N]$ and sends it to Prover</p> <p>Phase 5: Prover's second response</p> <p>16 : $\llbracket \alpha \rrbracket_{\bar{i}} \leftarrow (\llbracket \alpha_1 \rrbracket_{\bar{i}}, \dots, \llbracket \alpha_m \rrbracket_{\bar{i}})$, where $\llbracket \alpha_k \rrbracket_{\bar{i}} = \llbracket x_k \rrbracket_{\bar{i}} \cdot \varepsilon_k + \llbracket a_k \rrbracket_{\bar{i}}$ and $\llbracket x_k \rrbracket_{\bar{i}} = A_{k,0}(\llbracket \mathbf{s} \rrbracket_{\bar{i}})$</p> <p>17 : $\text{rsp} \leftarrow ((\text{seed}^{(i)}, \rho_i)_{i \neq \bar{i}}, \text{com}^{(\bar{i})}, \Delta \mathbf{s}, \Delta \mathbf{c}, \llbracket \alpha \rrbracket_{\bar{i}})$ and send rsp to Verifier</p> <p>Verification:</p> <p>18 : Verifier partially recomputes $(\llbracket \mathbf{s} \rrbracket_p, \llbracket \mathbf{c} \rrbracket_p, \llbracket \mathbf{a} \rrbracket_p)_{p \in [D] \times [2]}$ from $(\text{seed}^{(i)}, \rho_i)_{i \neq \bar{i}}$ by following Phase 1 but skipping the steps involving a \bar{i}-th share or the seed $\text{seed}^{(\bar{i})}$</p> <p style="padding-left: 20px;">for $\ell \in [D]$ do</p> <p>19 : Verifier gets H_ℓ and $(\llbracket \alpha_k \rrbracket_{(\ell, j)}, \llbracket v_k \rrbracket_{(\ell, j)})_{(k, j) \in [m] \times [2]}$ from the algorithm in Figure 5</p> <p>20 : Verifier accepts if and only if $h_0 = \mathbb{H}(\text{com}^{(1)}, \dots, \text{com}^{(N)})$ and $h_1 = \mathbb{H}(H_1, \dots, H_D)$, where $\text{com}^{(i)} = \text{Com}(\text{seed}^{(i)}, \rho_i)$ for each $i \neq \bar{i}$.</p>

Fig. 3: Proof of Knowledge protocol for PowAff2.

probability $\frac{1}{N}$. Consequently, a prover without a correct solution of the PowAff2_{u-1} problem will be accepted with probability at most $\epsilon = \frac{1}{N} + \frac{1}{q^u} \cdot (1 - \frac{1}{N})$.

Now, for the second part, we assume that no collisions to Com nor \mathbb{H} can be found and there exists a cheater $\tilde{\mathcal{P}}$ who has cheating probability $\tilde{\epsilon} > \epsilon$. First, we prove that a solution \mathbf{s} of the PowAff2_{u-1}

Inputs : A set of main shares $\left(\left(\llbracket \mathbf{s} \rrbracket_{(\ell,j)}, \llbracket \mathbf{c} \rrbracket_{(\ell,j)}, \llbracket \mathbf{a} \rrbracket_{(\ell,j)} \right) \right)_{j \in [2]}$ and the challenges $\varepsilon_1, \dots, \varepsilon_m$

Outputs : H_ℓ and $\left(\llbracket \alpha_k \rrbracket_{(\ell,j)}, \llbracket v_k \rrbracket_{(\ell,j)} \right)_{(k,j) \in [m] \times [2]}$

for $k \in [m]$ **do**

for $j \in [2]$ **do**

1 : $\llbracket x_k \rrbracket_{(\ell,j)} \leftarrow A_{k,1}(\llbracket \mathbf{s} \rrbracket_{(\ell,j)})$

2 : $\llbracket \alpha_k \rrbracket_{(\ell,j)} \leftarrow \llbracket x_k \rrbracket_{(\ell,j)} \cdot \varepsilon_k + \llbracket \mathbf{a} \rrbracket_{(\ell,j)}$

3 : $\alpha_k \leftarrow \llbracket \alpha_k \rrbracket_{(\ell,1)} + \llbracket \alpha_k \rrbracket_{(\ell,2)}$ // The parties open $\llbracket \alpha_k \rrbracket_{(\ell,j)}$ to obtain α_k .

4 : $\llbracket z_k \rrbracket_{(\ell,1)} \leftarrow t_k - A_{k,0}(\llbracket \mathbf{s} \rrbracket_{(\ell,1)})$

5 : $\llbracket y_k \rrbracket_{(\ell,1)} \leftarrow A_{k,2}(\llbracket \mathbf{s} \rrbracket_{(\ell,1)})$

6 : $\llbracket v_k \rrbracket_{(\ell,1)} \leftarrow \llbracket y_k \rrbracket_{(\ell,1)} \cdot \alpha_k - \llbracket z_k \rrbracket_{(\ell,1)} \cdot \varepsilon_k - \llbracket c_k \rrbracket_{(\ell,1)}$

7 : $\llbracket v_k \rrbracket_{(\ell,2)} \leftarrow -\llbracket v_k \rrbracket_{(\ell,1)}$

8 : $H_\ell \leftarrow \mathbb{H} \left(\left(\llbracket \alpha_k \rrbracket_{(\ell,j)}, \llbracket v_k \rrbracket_{(\ell,j)} \right)_{(k,j) \in [m] \times [2]} \right)$

Fig. 4: Simulation of the MPC protocol Π for the ℓ -th set of main shares.

Inputs: Partially computed main shares $\left(\left(\llbracket \mathbf{s} \rrbracket_{(\ell,j)}, \llbracket \mathbf{c} \rrbracket_{(\ell,j)}, \llbracket \mathbf{a} \rrbracket_{(\ell,j)} \right) \right)_{j \in [2]}$, the first challenges $\varepsilon_1, \dots, \varepsilon_m$, the second challenge \bar{i} , and the $\llbracket \alpha \rrbracket_{\bar{i}}$

Outputs : H_ℓ and $\left(\llbracket \alpha_k \rrbracket_{(\ell,j)}, \llbracket v_k \rrbracket_{(\ell,j)} \right)_{(k,j) \in [m] \times [2]}$

1 : $(\bar{i}_1, \dots, \bar{i}_D) \leftarrow \bar{i}$ // Binary representation of \bar{i} .

2 : $\llbracket \alpha_1 \rrbracket_{\bar{i}}, \dots, \llbracket \alpha_m \rrbracket_{\bar{i}} \leftarrow \llbracket \alpha \rrbracket_{\bar{i}}$

for $k \in [m]$ **do**

for $j \in [2]$ **do**

3 : $\llbracket x_k \rrbracket_{(\ell,j)} \leftarrow A_{k,1}(\llbracket \mathbf{s} \rrbracket_{(\ell,j)})$

4 : $\llbracket \alpha_k \rrbracket_{(\ell,j)} \leftarrow \llbracket x_k \rrbracket_{(\ell,j)} \cdot \varepsilon_k + \llbracket \mathbf{a} \rrbracket_{(\ell,j)}$

5 : $\llbracket \alpha_k \rrbracket_{(\ell, i_\ell+1)} \leftarrow \llbracket \alpha_k \rrbracket_{(\ell, i_\ell+1)} + \llbracket \alpha_k \rrbracket_{\bar{i}}$ // Adding missing share of $\llbracket \alpha_k \rrbracket_{(\ell, i_\ell+1)}$.

6 : $\alpha_k \leftarrow \llbracket \alpha_k \rrbracket_{(\ell,1)} + \llbracket \alpha_k \rrbracket_{(\ell,2)}$ // The parties open $\llbracket \alpha_k \rrbracket_{(\ell,j)}$ to obtain α_k .

7 : Set $i^* = 2$ if $\bar{i}_\ell = 0$, otherwise set $i^* = 1$.

8 : $\llbracket y_k \rrbracket_{(\ell, i^*)} \leftarrow A_{k,2}(\llbracket \mathbf{s} \rrbracket_{(\ell, i^*)})$

9 : $\llbracket z_k \rrbracket_{(\ell, i^*)} \leftarrow t_k - A_{k,0}(\llbracket \mathbf{s} \rrbracket_{(\ell, i^*)})$

10 : $\llbracket v_k \rrbracket_{(\ell, i^*)} \leftarrow \llbracket y_k \rrbracket_{(\ell, i^*)} \cdot \alpha_k - \llbracket z_k \rrbracket_{(\ell, i^*)} \cdot \varepsilon_k - \llbracket c_k \rrbracket_{(\ell, i^*)}$

11 : $\llbracket v_k \rrbracket_{(\ell, \bar{i}_\ell+1)} \leftarrow -\llbracket v_k \rrbracket_{(\ell, i^*)}$

12 : $H_\ell \leftarrow \mathbb{H} \left(\left(\llbracket \alpha_k \rrbracket_{(\ell,j)}, \llbracket v_k \rrbracket_{(\ell,j)} \right)_{(k,j) \in [m] \times [2]} \right)$

Fig. 5: Check the simulation of the MPC protocol Π in the ℓ -th set of main shares.

problem can be extracted from two valid transcripts of the form \mathcal{T}_1 and \mathcal{T}_2 produced by $\tilde{\mathcal{P}}$ that have the same initial commitment h_0 and different second challenges \tilde{i}_1 (for \mathcal{T}_1) and i_1 . Finally, we prove that such transcripts \mathcal{T}_1 and \mathcal{T}_2 can be extracted from \tilde{P} (assuming rewindable black-box access to \tilde{P}) with an expected number of calls upper bounded by

$$\frac{4}{\tilde{\epsilon} - \epsilon} \cdot \left(1 + \tilde{\epsilon} \cdot \frac{2 \ln(2)}{\tilde{\epsilon} - \epsilon} \right).$$

Details of how this second part is proven follow in a similar as shown in [10, Theorem 1].

- **Honest-verifier zero-knowledge:** Now we sketch the proof of the honest-verifier zero-knowledge property of the protocol in Figure 3. The goal here is to show that the distribution of the transcripts output by the simulator described in Figure 6 on input (\mathbf{f}, \mathbf{t}) are indistinguishable from those coming from a genuine interaction between a prover and an honest verifier, where the prover input is $(\mathbf{f}, \mathbf{t}, \mathbf{s})$ and $\mathbf{t} = \mathbf{f}(\mathbf{s})$.

Simulator(\mathbf{f}, \mathbf{t})

- 1 : Sample first challenge: $\boldsymbol{\varepsilon} = (\varepsilon_1, \dots, \varepsilon_m) \xleftarrow{\$} \mathbb{F}_q^m$
- 2 : Sample second challenge: $\tilde{i} \xleftarrow{\$} [N]$
- 3 : $\text{root} \xleftarrow{\$} \{0, 1\}^\lambda$
- 4 : $(\text{seed}^{(i)}, \rho^{(i)})_{i \in [N]} \leftarrow \text{TreePRG}(\text{root})$
for $i \in [N]$ **do**
- 5 : $[\mathbf{s}]_i, [\mathbf{c}]_i, [\mathbf{a}]_i \leftarrow \text{PRG}(\text{seed}^{(i)})$
- 6 : $\text{com}^{(i)} \leftarrow \text{Com}(\text{seed}^{(i)}, \rho_i)$
- 7 : $h_0 \leftarrow \text{H}(\text{com}^{(1)}, \dots, \text{com}^{(N)})$
- 8 : $\Delta \mathbf{s} \xleftarrow{\$} \mathbb{F}_q^m, \Delta \mathbf{c} \xleftarrow{\$} \mathbb{F}_q^m$
- 9 : $[\mathbf{s}]_1 \leftarrow [\mathbf{s}]_1 + \Delta \mathbf{s}$ and $[\mathbf{c}]_1 \leftarrow [\mathbf{c}]_1 + \Delta \mathbf{c}$
- 10 : Initialize $[\mathbf{s}]_p, [\mathbf{c}]_p$ and $[\mathbf{a}]_p$ to zero objects for each $p \in [D] \times [2]$
for $i \in [N] \setminus \{\tilde{i}\}$ **do**
- 11 : Simulate the i party to obtain $[\alpha_k]_i$ and $[v_k]_i$ for each $k \in [m]$
- 12 : $[\alpha_k]_{\tilde{i}} \xleftarrow{\$} \mathbb{F}_q$ and $[v_k]_{\tilde{i}} \xleftarrow{\$} \mathbb{F}_q$ for each $k \in [m]$
- 13 : $\text{com}^{(\tilde{i})} \xleftarrow{\$} \{0, 1\}^\lambda$
- 14 : For each $(k, \ell, j) \in [m] \times [D] \times [2]$ compute $[\alpha_k]_{(\ell, j)}$ and $[v_k]_{(\ell, j)}$
- 15 : Set $H_\ell \leftarrow \text{H}\left(\left([\alpha_k]_{(\ell, j)}, [v_k]_{(\ell, j)}\right)_{(k, j) \in [m] \times [2]}\right)$ for each $\ell \in [D]$
- 16 : $h_1 \leftarrow \text{H}(H_1, \dots, H_D)$
- 17 : $\text{rsp} \leftarrow ((\text{seed}^{(i)}, \rho_i)_{i \neq \tilde{i}}, \text{com}^{(\tilde{i})}, \Delta \mathbf{s}, \Delta \mathbf{c}, [\alpha]_{\tilde{i}})$, where $[\alpha]_{\tilde{i}} = ([\alpha_1]_{\tilde{i}}, \dots, [\alpha_m]_{\tilde{i}})$

Output $(h_0, \boldsymbol{\varepsilon}, h_1, \tilde{i}, \text{rsp})$

Fig. 6: Honest-verifier zero-knowledge simulator.

The idea is to create a sequence of simulators that ends with the simulator described in Figure 6. The first simulator of the sequence consists of a legitimate prover, which holds a solution \mathbf{s} and simulates the verifier by randomly sampling the challenges, as an honest verifier would do. These transcripts are indistinguishable from those coming from a legitimate execution of the protocol in proof of knowledge protocol.

Finally, the proof is completed by showing that the transcripts output by any simulator in the sequence are indistinguishable from those in the previous simulator. This implies that the transcripts of the

simulator in Figure 6 are indistinguishable from those in for deduced by the actual protocol. Details of this part follow similarly as shown in [10, Theorem 1]. □

4 Biscuit Signature Scheme

In this part, we describe the Biscuit signature scheme. It is obtained by applying the Fiat-Shamir transformation [25] to the zero-knowledge protocol given in Figure 3. The corresponding key-generation, signing, and verification algorithms are described in figures 7, 8 and 9, respectively.

The secret-key is a random vector $\mathbf{s} \in \mathbb{F}_q^n$ and the public-key is a pair $(\mathbf{f} = (f_1, \dots, f_m), \mathbf{t} = \mathbf{f}(\mathbf{s})) \in \mathbb{F}_q[x_1, \dots, x_n]^m \times \mathbb{F}_q^m$ such that for all $k \in [m]$:

$$f_k(x_1, \dots, x_n) = A_{k,0}(x_1, \dots, x_n) + A_{k,1}(x_1, \dots, x_n) \cdot A_{k,2}(x_1, \dots, x_n), \quad (4)$$

where $A_{1,0}, \dots, A_{m,2} \in \mathbb{F}_q[x_1, \dots, x_n]$ are random affine forms as in (1).

We use two seeds $\text{seed}_f, \text{seed}_s \in \{0, 1\}^\lambda$ that are extended by a pseudo-random generator (PRG) to obtain the public polynomials $\mathbf{f} \in \mathbb{F}_q[x_1, \dots, x_n]^m$ and the secret vector $\mathbf{s} \in \mathbb{F}_q[x_1, \dots, x_n]^m$. Finally, the vector $\mathbf{t} \in \mathbb{F}_q^m$ is computed as $\mathbf{t} = \mathbf{f}(\mathbf{s})$.




Fig. 7: Key-generation algorithm for Biscuit.

Remark 2. The notation $\mathbf{f} \leftarrow \text{PRG}(\text{seed}_f)$ is a shortcut for extending the seed from a PRG and casting the bit string into a set of algebraic equations as in (4). Similarly, $\mathbf{s} \leftarrow \text{PRG}(\text{seed}_{sk})$ stands for extending the seed and interpreting the bit string as a vector in \mathbb{F}_q^n .

The signing procedure `Biscuit.Sign` is given in Figure 8. It takes as input a key-pair (sk, pk) and the message $\text{msg} \in \{0, 1\}^*$ to sign. It is obtained by applying the Fiat-Shamir transform to the ZKPoK for `PowAff2` described in Section 3.2.

The verification process (Figure 9) is very similar to the signature process (Figure 8) as the verifier has to replay the MPC protocol for each of the τ participants except one. The algorithm takes as input a message $\text{msg} \in \{0, 1\}^*$, a signature sig and a public-key pk . It returns a bit $b \in \{0, 1\}$.

Theorem 2 (EU-CMA). *Let PRG be a $(t, \epsilon_{\text{PRG}})$ -secure pseudo-random generator function, and that any adversary running in time t has advantage at most $\epsilon_{\text{PowAff2}}$ against the underlying `PowAff2` _{$u-1$} problem. Suppose that H_0, H_1, H_2, H_4 behave as random oracles that output binary strings of size 2λ . Let \mathcal{A} be an adversary, who has access to a signing oracle, making q_i queries to H_i and q_s queries to the signing oracle. Then, the probability that \mathcal{A} outputs a forgery of for the Biscuit signature scheme is:*

$$\Pr[\text{Forge}] \leq \frac{3(q + \tau N \cdot q_s)^2}{2 \cdot 2^{2\lambda}} + \frac{q_s(q_s + 5q)}{2^{2\lambda}} + \epsilon_{\text{PRG}} + \epsilon_{\text{PowAff2}} + \Pr[X + Y = \tau],$$

where τ is the number of rounds of the signature, $X = \max_{i \in [0, q_2]} \{X_i\}$ with $X_i \sim \mathcal{B}(\tau, \frac{1}{q^u})$, and $Y = \max_{i \in [0, q_4]} \{Y_i\}$ with $Y_i \sim \mathcal{B}(\tau - X, \frac{1}{N})$.

```

Sign(pk, sk, msg)
1 : (seedf, t) ← pk, seedsk ← sk
2 : f ← PRG(seedf), s ← PRG(seedsk)
Step 1: Commit to the inputs of the MPC protocol in Figure 4
3 : salt  $\xleftarrow{\$}$  {0, 1}2λ
   for e ∈ [τ]
4 : seed(e)  $\xleftarrow{\$}$  {0, 1}λ, (seed(e,i))i∈[N] ← TreePRG(salt, root(e))
   for i ∈ [N] do
5 : [s]i(e), [c]i(e), [a]i(e) ← PRG(seed(e,i))
6 : com(e,i) ← H0(salt, e, i, seed(e,i))
7 : h0(e) ← H1(salt, e, com(e,1), ..., com(e,N))
8 : h1 ← H2(salt, msg, h0(1), ..., h0(τ))
9 : a(e) ← ∑i∈[N] [a]i(e), c(e) ← (Ak,2(s) · ak(e))k∈[m]
10 : Δs(e) ← s - ∑i∈[N] [s]i(e), Δc(e) ← c(e) - ∑i∈[N] [c]i(e)
11 : [s]1(e) ← [s]1(e) + Δs(e) and [c]1(e) ← [c]1(e) + Δc(e)
12 : Initialize [s]p(e), [c]p(e) and [a]p(e) to zero objects for each p ∈ [D] × [2]
   for i ∈ [N] do
13 : (i1, ..., iD) ← i // Binary representation of i.
   for ℓ ∈ [D] do
14 : [s](ℓ,iℓ+1)(e) ← [s](ℓ,iℓ+1)(e) + [s]i(e), [c](ℓ,iℓ+1)(e) ← [c](ℓ,iℓ+1)(e) + [c]i(e) and
15 : [a](ℓ,iℓ+1)(e) ← [a](ℓ,iℓ+1)(e) + [a]i(e)
Step 2: First challenge
16 : ((ε1(e), ..., εm(e)))e∈[τ]  $\xleftarrow{\$}$  PRG(h1)
Step 3: First response
   for e ∈ [τ] do
   for ℓ ∈ [D] do
17 : Follow the algorithm in Figure 4 to get Hℓ(e), which is defined instead as
18 : Hℓ(e) = H3(salt, ℓ, [αk](ℓ,j)(e), [vk](ℓ,j)(e))(k,j)∈[m]×[2]
19 : h2 ← H4(salt, msg, pk, h1, H1(e), ..., HD(e))
Step 4: Second challenge
20 : i1, ..., iτ  $\xleftarrow{\$}$  PRG(h2)
Step 5: Second response
   for e ∈ [τ] do
21 : [α]i(e) ← ([α1]i(e), ..., [αm]i(e)), where [αk]i(e) = [xk]i(e) · εk(e) + [ak]i(e), and
   [xk]i(e) = Ak,1([s]i(e))
22 : σ ← (salt, h1, h2, ((seed(e,i))i≠ie, com(e,ie))e∈[τ], (Δs(e), Δc(e), [α]ie(e))e∈[τ])
23 : Output σ

```

Fig. 8: Biscuit signing algorithm.

```

Verify(pk, σ, msg)


---


1 : (f, t) ← ExpandPK(pk) // Expanding the public key pk.
Step 1: Parse signature
2 : (salt, h1, h2, ((seed(e,i))i≠ie, com(e,ie))e∈[τ], (Δs(e), Δc(e), [α]ie(e))e∈[τ]) ← σ
3 : i1, ..., iτ ←$ PRG(h2)
Step 2: Recompute h1 and the inputs of the MPC protocol
  for e ∈ [τ]
    for i ∈ [N] \ {ie} do
4 : [s]i(e), [c]i(e), [a]i(e) ← PRG(seed(e,i))
5 : com(e,i) ← H0(salt, e, i, seed(e,i))
6 : h0(e) ← H1(salt, e, com(e,1), ..., com(e,N))
7 : h1(e) ← H2(salt, msg, h0(1), ..., h0(τ))
8 : a(e) ← ∑i∈[N]\{ie [a]i(e), c(e) ← (Ak,2(s) · ak(e))k∈[m]
9 : Δs(e) ← s - ∑i∈[N]\{ie [s]i(e), Δc(e) ← c(e) - ∑i∈[N]\{ie [c]i(e)
    if ie ≠ 1 then
10 : [s]1(e) ← [s]1(e) + Δs(e) and [c]1(e) ← [c]1(e) + Δc(e)
11 : Initialize [s]p(e), [c]p(e) and [a]p(e) to zero objects for each p ∈ [D] × [2]
    for i ∈ [N] \ {ie} do
12 : (i1, ..., iD) ← i // Binary representation of i.
    for ℓ ∈ [D] do
13 : [s](ℓ,iℓ+1)(e) ← [s](ℓ,iℓ+1)(e) + [s]i(e), [c](ℓ,iℓ+1)(e) ← [c](ℓ,iℓ+1)(e) + [c]i(e) and
14 : [a](ℓ,iℓ+1)(e) ← [a](ℓ,iℓ+1)(e) + [a]i(e)
Step 3: Recompute h2
  for e ∈ [τ] do
    for ℓ ∈ [D] do
15 : Use (ε1(e), ..., εm(e)), [α]ie(e) and the ℓ-th set of main shares as inputs in
16 : the algorithm in Figure 5 to get Hℓ(e), which is defined instead as
17 : Hℓ(e) = H3(salt, ℓ, [αk](ℓ,j)(e), [vk](ℓ,j)(e))(k,j)∈[m]×[2]
18 : h2 ← H4(salt, msg, pk, h1, (H1(e), ..., HD(e))e∈[τ])
Step 4: Verify signature
19 : Output (h1 = h1) ∧ (h2 = h2)

```

Fig. 9: Biscuit verification algorithm.

4.1 Parameters

Table 2 provides the parameter sets **Biscuit**, along with the corresponding size of the keys and signatures. Each parameter set aims to provide a security level either I, III or V according to the NIST guidelines. A more detailed description of the claimed security level of each parameter set is given in Section 5.

Level	Version	λ	q	n	m	N	τ	Bit-Security	sk	pk	σ
I	short	128	256	50	52	256	18	143	16	68	5748
	fast						28	143			7544
III	short	192	256	89	92	256	25	207	24	116	12969
	fast						40	210			17784
V	short	256	256	127	130	256	33	272	32	162	23523
	fast						53	275			32575

Table 2: Parameters of **Biscuit**, bit security, public-key (**pk**), secret-key (**sk**) and signature (σ) sizes in bytes.

The size of the public-key is $\lambda + \log_2(q) \cdot m$ bits, the size of the secret-key is λ bits and the bit-size of the signature is:

$$\underbrace{6\lambda}_{\text{salt}, h_1, h_2} + \tau \left(\underbrace{(n + 2m) \log_2 q}_{\Delta \mathbf{s}^{(e)}, \Delta \mathbf{c}^{(e)}, [\alpha]_{\bar{i}_e}^{(e)}} + \underbrace{\lambda \cdot D}_{(\text{seed}^{(e, i)})_{i \neq \bar{i}_e}} + \underbrace{2\lambda}_{\text{com}^{(e, \bar{i}_e)}} \right).$$

5 Security Analysis

This part is dedicated to the security analysis of **Biscuit** against *key-recovery* (Section 5.2) and *forgery* (Section 5.3) attacks. Before that, Section 5.1 discusses the motivations for using structured systems as **PowAff** and the connection with the Learning With Errors (LWE, [34]) problem.

From now on, let $(\mathbf{f} = (f_1, \dots, f_m), \mathbf{t} = \mathbf{f}(\mathbf{s})) \in \mathbb{F}_q[x_1, \dots, x_n]^m \times \mathbb{F}_q^m$ be a **Biscuit** public-key and $\mathbf{s} \in \mathbb{F}_q^n$ be the corresponding secret-key.

5.1 About the Hardness of PowAff2

A fundamental assumption in the design of **Biscuit** is that solving algebraic systems generated essentially from power of affine forms are not much easier to solve than a random system of quadratic equations. Whilst the complexity of solving structured equations can be difficult to assess in general, the hardness of solving random quadratic equations has been deeply investigated and only exponential algorithms are known, e.g. [18, 17, 12, 16].

We emphasize **PowAff2** algebraic equations already appeared previously in the literature. In particular, the authors of [7, 11] demonstrated that attacking the Learning With Errors (LWE) problem [34] reduces to solve a structured algebraic system similar to **PowAff2**. An instance of LWE is given by a pair $(\mathbf{A} = \{a_{i,j}\}, \mathbf{c} = \mathbf{s}\mathbf{A} + \mathbf{e}) \in \mathbb{F}_q^{n \times m} \times \mathbb{F}_q^m$ where $\mathbf{s} \in \mathbb{F}_q^n$ is a secret and $\mathbf{e} \in \mathbb{F}_q^m$ is an error vector. LWE (search) asks to recover the secret \mathbf{s} . Arora and Ge exhibits in [7, 11] a rather natural algebraic modeling of LWE. More precisely, Arora and Ge show that LWE secrets can be recovered by solving:

$$f_1(x_1, \dots, x_n) = P(c_1 - \sum_{k=1}^n a_{k,1} x_k) = 0, \dots, f_m(x_1, \dots, x_n) = P(c_m - \sum_{k=1}^n a_{k,m} x_k) = 0, \quad (5)$$

where P depends on the error distribution. In particular, $P(X) = X(X-1) \in \mathbb{F}_q[X]$ for binary errors and [7] introduced the assumption that a system such as (5) behaves such as a semi-regular sequence. As a consequence, a new fast algorithm for **PowAff2** will lead to a new fast algebraic algorithm for binary LWE.

5.2 Key recovery attacks

A key-recovery attack against `Biscuit` consists of solving the `PowAff2` problem, i.e. recovering $\mathbf{s} \in \mathbb{F}_q^m$ from the system defined as :

$$\mathbf{t} = \mathbf{f}(\mathbf{x}), \text{ with } \mathbf{x} = (x_1, \dots, x_n). \quad (6)$$

Currently, the best attack against `Biscuit` is a dedicated hybrid approach for solving `PowAff2` equations described by Charles Bouillaguet¹³. The hybrid approach is a classical technique for solving algebraic systems that combines exhaustive search and a Gröbner basis-like computations [18,17,12]. The efficiency of such approaches is related to the choice of a *trade-off*, denoted $k \leq n$, between these two methods.

We sketch below the approach described on the NIST PQC mailing list. Let $\mathbf{g} = (g_1(\mathbf{x}) = u_1(\mathbf{x}) \cdot (x_1 + c_1) + w_1(\mathbf{x}), \dots, g_n(\mathbf{x}) = u_n(\mathbf{x}) \cdot (x_n + c_n) + w_n(\mathbf{x})) \in \mathbb{F}_q[x_1, \dots, x_n]^m$, with $\mathbf{x} = (x_1, \dots, x_n)$, $u_1, \dots, u_n, v_1, \dots, v_n \in \mathbb{F}_q[x_1, \dots, x_n]$ affine polynomials and $c_1, \dots, c_n \in \mathbb{F}_q$. According to Lemma 1, with high probability, there exists $\mathbf{L} \in \text{GL}_n(\mathbb{F}_q)$ such that:

$$\mathbf{f}(\mathbf{x} \cdot \mathbf{L}) = (\mathbf{g}, A'_{n+1,0}(\mathbf{x}) + \prod_{j=1}^2 A'_{n+1,j}(\mathbf{x}), \dots, A'_{m,0}(\mathbf{x}) + \prod_{j=1}^2 A'_{m,j}(\mathbf{x}))$$

where $A_{n+1,0}, A_{n+1,1}, A_{n+1,2}, \dots, A_{m,0}, A_{m,1}, A_{m,2} \in \mathbb{F}_q[x_1, \dots, x_n]$ affine forms.

Then, for every guess $(a_1, \dots, a_k) \in \mathbb{F}_q^k$ of the vector of variables (x_1, \dots, x_k) , we obtain k linear polynomials, namely $g_1(a_1, \dots, a_k, x_{k+1}, \dots, x_n), \dots, g_k(a_1, \dots, a_k, x_{k+1}, \dots, x_n)$. These k linear polynomials are expected to be linearly independent with a probability close to $1 - 1/q$. Hence we can use them to substitute k additional variables in the remaining polynomials. The attack is finalized by solving the resulting quadratic system of $m - k$ equations in $n - 2k$ variables.

Complexity. The cost of the attack is dominated by

$$\min_{0 \leq k < \frac{n}{2}} q^k \cdot \text{MQ}(n - 2k, m - k, q), \quad (7)$$

where $\text{MQ}(n, m, q)$ denotes the complexity of solving a random system of m quadratic equations over in n variables over \mathbb{F}_q . To compute the exact complexity, we can rely on the `MQEstimator` software tool which is part of the more general `CryptographicEstimators`¹⁴ library [22].

5.3 Forgery attacks

In the context of forgery, the attacker has to solve the `PowAff2u` problem (Definition 1) which is a variant of the problem considered before for key-recovery (Section 5.2). In the `PowAff2u` problem, the goal is to find a vector $\mathbf{s}' \in \mathbb{F}_q^n$ that vanishes a subset of size $m - u$ of the system (6). Without loss of generality, we assume \mathbf{s}' vanishes the first $m - u$ polynomials and not the remaining equations. That is, $f_k(\mathbf{s}') = t_k$, for $k \in [m - u]$, and $f_k(\mathbf{s}') \neq t_k$ for $k = m - u + 1, \dots, m$.

By Proposition 1, a set of N parties that follows the MPC protocol in Figure 2 on inputs $\llbracket \mathbf{s}' \rrbracket$ and (\mathbf{f}, \mathbf{t}) will output **accept** with probability $p_1 = 1/q^u$. In the context of `MPCitH`, the value p_1 is referred in the literature as the false positive rate of the MPC protocol.

Thanks to the $\text{}$, it is known that `MPCitH`-based signature scheme that consist of τ repetitions of a MPC protocol with false positive rate p_1 can be forged by computing of average

$$\text{KZ}_\tau(p_1, p_2) = \min_{\{\tau_1, \tau_2 | \tau_1 + \tau_2 = \tau\}} \left\{ \frac{1}{\sum_{i=\tau_1}^{\tau} \binom{\tau}{i} p_1^i (1 - p_1)^{\tau-i}} + \frac{1}{p_2^{\tau_2}} \right\},$$

calls to some hash functions, where p_2 is the probability of guessing some of the views of parties that remain unopened, e.g., $p_2 = 1/N$ for `Biscuit`.

Let $\mathbf{C}_u(q, n, m)$ denote the complexity of finding a preimage to a chosen subset S of the system $\mathbf{t} = \mathbf{f}(\mathbf{x})$ of size $m - u$ and $\mathbf{s}' \in \mathbb{F}_q^n$ be a solution than vanishes the equations of S . Then, \mathbf{s}' might, by chance, be a solution of any equation in S^c , i.e., any equation that is not in S . If there remain $k \in [u]$ equations in S^c for which \mathbf{s}' is not a solution, then an attacker can mount a forgery attack with complexity $\text{KZ}_\tau(q^{-k}, N^{-1})$.

¹³ https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/sw8NueiNek0/m/2sa_emjABQAJ

¹⁴ <https://github.com/Crypto-TII/CryptographicEstimators>

Let (\mathbf{f}, \mathbf{t}) a Biscuit public-key selected uniformly at random, and let S be a subset of the equations $\mathbf{t} = \mathbf{f}(\mathbf{x})$ of size $m - u$ selected uniformly at random. Then, a random solution $\mathbf{s}' \in \mathbb{F}_q^n$ of the equations in S follows a uniform distribution. Hence, $f_\ell(\mathbf{s}')$ is a uniform element in \mathbb{F}_q . Therefore, the probability that \mathbf{s}' is a solution of exactly j equations in S^c is $\binom{u}{j} \cdot (q-1)^{u-j}/q^u$. Consequently, if p_k denotes the probability that \mathbf{s}' is not the solution of at most k equations in S^c , then,

$$p_k = \frac{\sum_{j=u-k+1}^u \binom{u}{j} \cdot (q-1)^{u-j}}{q^u}.$$

In order to secure Biscuit against forgery attacks, we must have for every pair (k, u) , where $0 \leq k \leq u \leq m$:

1. $\text{KZ}_\tau(q^{-k}, N^{-1}) > 2^\lambda$, or
2. $\frac{1}{p_k} \cdot \mathcal{C}_u(q, n, m) > 2^{\lambda+C_\lambda}$,

where $C_\lambda = 15$ if $\lambda = 128$ or 192 and $C_\lambda = 16$ otherwise.

Following these analyses, we propose in Table 2 a set of 3 parameters for 128, 192 and 256 bits of classical security.

6 Implementation

6.1 Canonical Representation Optimization

As seen in Lemma 1, an equivalent system where, for the first n equations, one of the affine form is only composed of one variable. Without loss of generality, we can choose to have this variable in $A_{k,0}$. In other word, we can chose for the algorithm a system f_1, \dots, f_m as

$$f_k(x_1, \dots, x_n) = (x_k + a_k) + A_{k,1}(x_1, \dots, x_n) \cdot A_{k,2}(x_1, \dots, x_n),$$

for $k \leq n$, and

$$f_k(x_1, \dots, x_n) = A_{k,0}(x_1, \dots, x_n) + A_{k,1}(x_1, \dots, x_n) \cdot A_{k,2}(x_1, \dots, x_n),$$

for $n < k \leq m$, where $A_{k,j}$ are affine forms.

The effect is that the evaluation of the polynomial will be much faster as only 2 affine form evaluations have to be performed instead of 3 for most of the equations. In the implementation, we chose to simplify $A_{k,0}$ to save some code, as $A_{k,1}$ and $A_{k,2}$ can be computed in the same way in a loop.

6.2 Hypercube Optimization

The algorithms described in Figure 8 and Figure 9 use the hypercube variant. The simulation of the MPC protocol does not need to compute all the values as in Figure 4. We first compute α_k using directly the opened values \mathbf{s} and \mathbf{a} . Then, need to compute $\llbracket \alpha_k \rrbracket_{(\ell,j)}$ only for $j = 1$. The value for $j = 2$ can be derived from α . Similarly, we can do the same for $\llbracket v_k \rrbracket_{(\ell,j)}$. This can also be applied to the verification. All in all, we usually require to keep only $\log_2(N)$ shares.

6.3 Vectorization

The main data structure in the algorithm is a vector of value in \mathbb{F}_q . We have

- the secret value which is a vector of n elements in \mathbb{F}_q
- the public key which is a vector of m elements in \mathbb{F}_q
- intermediate values which are vectors of m elements in \mathbb{F}_q .

For each of this vector, we need to compute operations component-wise. We can then pack all elements in the largest possible integer handled by the CPU. Typically this could be a 64-bit word that can contain 8 elements in \mathbb{F}_{2^8} for instance.

When vectorized instructions are available (SSE, AVX, ...), even lagrer integer types can be used. For instance, with AVX2 a 256-bit integer can be used to pack vector of \mathbb{F}_q elements. In characteristic 2, the component-wise addition of a vector of elements can be done in one instruction using the VPXOR instruction.

6.4 Performances and Memory Consumption

In this section, we show the performance and memory consumption of our instances. Our implementation is optimized to use `avx2` vectorized instructions on a little-endian 64-bit CPU.

The code is compiled with `GCC` version 12.2.0 on Debian GNU/Linux. Number of cycles was measured by counting `PERF_HW_COUNT_CPU_CYCLES` events on an `11th Gen Intel(R) Core(TM) i7-1185G7 @ 3.00GHz` CPU. Even if frequency modification should not affect this metric, we deactivated Intel’s TurboBoost feature anyway. The number of cycles is averaged over 100 executions.

In Table 3, we give the figures for the implementation strictly following the description in the NIST submission but with the new parameters proposed in Table 2.

In Table 4, we include the canonical representation optimization as described in Section 6.1. This improves the performances by 18 to 28 percent.

Finally, in Table 5, in addition to the previous optimization, we integrated the hypercube variant. With this variant, the memory consumption is greatly improved especially for large values of N . This is because we have to keep track of only $\log_2(N)$ shares instead of N . The performances are improved by 50 to 83 percent for the small variant, and by 41 to 69 percent for the fast variant.

The code is available in [2].

Name	Memory (bytes)			Performance (cycles)		
	keygen	sign	verify	keygen	sign	verify
<code>biscuit128s</code>	512	1 654 288	122 480	88 870	69 849 933	69 009 007
<code>biscuit128f</code>	512	329 904	25 712	88 614	13 762 950	13 150 570
<code>biscuit192s</code>	608	3 438 832	194 544	252 374	191 148 678	191 044 372
<code>biscuit192f</code>	608	708 944	49 392	252 420	38 750 065	37 226 999
<code>biscuit256s</code>	800	7 414 000	335 312	504 271	635 041 660	631 093 761
<code>biscuit256f</code>	800	1 537 904	98 768	506 853	128 301 281	123 744 104

Table 3: Time performance and memory consumption of Biscuit on `avx2` impl.

Name	Memory (bytes)			Performance (cycles)		
	keygen	sign	verify	keygen	sign	verify
<code>biscuit128s</code>	512	1 651 088	122 480	62 413	57 621 312	56 705 942
<code>biscuit128f</code>	512	326 704	25 712	61 947	11 367 980	10 739 162
<code>biscuit192s</code>	608	3 430 288	194 544	173 744	150 259 434	148 692 324
<code>biscuit192f</code>	608	700 400	49 392	175 783	30 356 710	29 144 437
<code>biscuit256s</code>	800	7 393 680	335 312	341 657	456 813 570	453 782 090
<code>biscuit256f</code>	800	1 517 584	98 768	344 467	92 359 826	88 932 332

Table 4: Time performance and memory consumption of Biscuit on `avx2` impl. using canonical optimization.

6.5 Security Against Side-Channel Attacks

In this section, we briefly discuss the possible side-channel vulnerabilities of our proposition, and how to address them.

Name	Memory (bytes)			Performance (cycles)		
	keygen	sign	verify	keygen	sign	verify
biscuit128s	576	814 256	40 144	62 484	27 922 077	28 484 726
biscuit128f	576	201 744	14 096	62 043	6 577 335	6 260 131
biscuit192s	704	1 686 416	67 376	173 536	49 655 357	49 852 658
biscuit192f	704	433 008	28 272	173 460	13 612 332	13 084 840
biscuit256s	960	3 556 624	117 424	342 041	77 132 303	77 300 577
biscuit256f	960	928 368	57 648	344 956	28 316 992	27 347 243

Table 5: Time performance and memory consumption of Biscuit on avx2 impl. using canonical and hypercube optimization.

Timing Attacks. First of all, it is worth noticing that the signature procedure of Biscuit is independent of the secret value as long as the field arithmetic and the hash functions do not leak information on the manipulated data. Indeed, there is no branching that depends on the value of any secret in the algorithm. This allows to make an isochronous implementation by focusing on the field arithmetic.

Side-Channel Attacks. The most popular technique to prevent these attacks is to use *masking*: we compute a sharing of the secret using fresh random value at each execution.

In our scheme, it happens that most of the time, the secret value s is already shared into N shares due to the MPC protocol. Nevertheless, this does not guarantee security at order $N - 1$ because all the shares (except one) will finally become public during the verification process. This sharing is not secure from a side-channel attacker point of view. However, as our construction uses fields of characteristic 2, classical Boolean masking techniques can be applied throughout the scheme.

7 Acknowledgements

The authors would like to thank Charles Bouillaguet and Julia Sauvage for discussions on the hardness of PowAff2 and choice of parameters.

The third author would like to thank Google which partially supported this work thanks to a gift for supporting post-quantum research.

References

1. NIST. Call for Additional Digital Signature Schemes for the Post-Quantum Cryptography Standardization Process . <https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/call-for-proposals-dig-sig-sept-2022.pdf>.
2. Biscuit github repository, 2023. <https://github.com/BiscuitTeam/Biscuit>.
3. Najwa Aaraj, Slim Bettaieb, Loïc Bidoux, Alessandro Budroni, Victor Dyesryn, Andre Esser, Philippe Gaborit, Mukul Kulkarni, Victor Mateu, Marco Palumbi, Lucas Perin, and Jean-Pierre Tillich. PERK specification. 2023. available at https://pqc-perk.org/assets/downloads/PERK_specifications.pdf.
4. Gora Adj, Stefano Barbero, Emanuele Bellini, Andre Esser, Luis Rivera-Zamarripa, Carlo Sanna, Javier Verbel, and Floyd Zweyding. MiRitH specification. 2023. available at https://pqc-mirith.org/assets/downloads/mirith_specifications_v1.0.0.pdf.
5. Carlos Aguilar Melchor, Nicolas Gama, James Howe, Andreas Hülsing, David Joseph, and Dongze Yue. The return of the SDitH. pages 564–596, 2023.
6. Gorjan Alagic, Jacob Alperin-Sheriff, Daniel Apon, David Cooper, Quynh Dang, John Kelsey, Yi-Kai Liu, Carl Miller, Dustin Moody, Rene Peralta, Ray Perlner, and Angela Robinson and Daniel Smith-Tone. Status Report on the Second Round of the NIST Post-Quantum Cryptography Standardization Process. Technical Report NISTIR 8309, NIST, 2022. <https://nvlpubs.nist.gov/nistpubs/ir/2020/NIST.IR.8309.pdf>.

7. Martin R. Albrecht, Carlos Cid, Jean-Charles Faugère, and Ludovic Perret. Algebraic algorithms for lwe. Cryptology ePrint Archive, Paper 2014/1018, 2014. <https://eprint.iacr.org/2014/1018>.
8. Nicolas Aragon, Magali Bardet, Loïc Bidoux, Jesús-Javier Chi-Domínguez, Victor Dyseryn, Thibault Feneuil, Philippe Gaborit, Antoine Joux, Matthieu Rivain, Jean-Pierre Tillich, and Adrien Vinçotte. RYDE specification. 2023. available at https://pqc-ryde.org/assets/downloads/RYDE_Specifications.pdf.
9. Nicolas Aragon, Magali Bardet, Loïc Bidoux, Jesús-Javier Chi-Domínguez, Victor Dyseryn, Thibault Feneuil, Philippe Gaborit, Romaric Neveu, Matthieu Rivain, and Jean-Pierre Tillich. MIRA specification. 2023. available at https://pqc-mira.org/assets/downloads/mira_spec.pdf.
10. Nicolas Aragon, Loïc Bidoux, Jesús-Javier Chi-Domínguez, Thibault Feneuil, Philippe Gaborit, Romaric Neveu, and Matthieu Rivain. Mira: a digital signature scheme based on the minrank problem and the mpc-in-the-head paradigm, 2023.
11. Sanjeev Arora and Rong Ge. New algorithms for learning in presence of errors. In Luca Aceto, Monika Henzinger, and Jirí Sgall, editors, *Automata, Languages and Programming - 38th International Colloquium, ICALP 2011, Zurich, Switzerland, July 4-8, 2011, Proceedings, Part I*, volume 6755 of *Lecture Notes in Computer Science*, pages 403–415. Springer, 2011.
12. Magali Bardet, Jean-Charles Faugère, Bruno Salvy, and Pierre-Jean Spaenlehauer. On the Complexity of Solving Quadratic Boolean Systems. *J. Complex.*, 29(1):53–75, 2013.
13. Carsten Baum, Lennart Braun, Michael Kloöß, Christian Majenz, Shibam Mukherjee, Emmanuela Orsini, Sebastian Ramacher, Christian Rechberger, Lawrence Roy, and Peter Scholl. FAEST specification. 2023. available at <https://faest.info/faest-spec-v1.1.pdf>.
14. Carsten Baum and Ariel Nof. Concretely-efficient zero-knowledge arguments for arithmetic circuits and their application to lattice-based cryptography. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *Public-Key Cryptography – PKC 2020*, pages 495–526, Cham, 2020. Springer International Publishing.
15. Donald Beaver. Efficient multiparty protocols using circuit randomization. In Joan Feigenbaum, editor, *Advances in Cryptology - CRYPTO '91, 11th Annual International Cryptology Conference, Santa Barbara, California, USA, August 11-15, 1991, Proceedings*, volume 576 of *Lecture Notes in Computer Science*, pages 420–432. Springer, 1991.
16. Emanuele Bellini, Rusydi H. Makarim, Carlo Sanna, and Javier A. Verbel. An estimator for the hardness of the MQ problem. pages 323–347, 2022.
17. Luk Bettale, Jean-Charles Faugère, and Ludovic Perret. Hybrid approach for solving multivariate systems over finite fields. *J. Math. Cryptol.*, 3(3):177–197, 2009.
18. Luk Bettale, Jean-Charles Faugère, and Ludovic Perret. Solving polynomial systems over finite fields: improved analysis of the hybrid approach. In Joris van der Hoeven and Mark van Hoeij, editors, *International Symposium on Symbolic and Algebraic Computation, ISSAC'12, Grenoble, France - July 22 - 25, 2012*, pages 67–74. ACM, 2012.
19. Luk Bettale, Ludovic Perret, Delaram Kahrobaei, and Javier Verbel. Biscuit: Shorter MPC-based Signature from PoSSo, June 2023. Specification of NIST post-quantum signature.
20. Melissa Chase, David Derler, Steven Goldfeder, Claudio Orlandi, Sebastian Ramacher, Christian Rechberger, Daniel Slamanig, and Greg Zaverucha. Post-quantum zero-knowledge and signatures from symmetric-key primitives. pages 1825–1842, 2017.
21. NIST Computer Security Division. SHA-3 Standard: Permutation-Based Hash and Extendable-Output Functions. FIPS Publication 202, National Institute of Standards and Technology, U.S. Department of Commerce, May 2014.
22. Andre Esser, Javier Verbel, Floyd Zweyding, and Emanuele Bellini. **CryptographicEstimators**: a software library for cryptographic hardness estimation. Cryptology ePrint Archive, Paper 2023/589, 2023. <https://eprint.iacr.org/2023/589>.
23. Thibault Feneuil and Matthieu Rivain. MQOM specification. 2023. available at <https://mqom.org/docs/mqom-v1.0.pdf>.
24. Thibault Feneuil and Matthieu Rivain. Threshold computation in the head: Improved framework for post-quantum signatures and zero-knowledge arguments. Cryptology ePrint Archive, Paper 2023/1573, 2023. <https://eprint.iacr.org/2023/1573>.
25. A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In A. M. Odlyzko, editor, *CRYPTO 1986*, pages 186–194, Berlin, Heidelberg, 1987. Springer Berlin Heidelberg.
26. Amos Fiat and Adi Shamir. How to Prove Yourself: Practical Solutions to Identification and Signature Problems. In Andrew M. Odlyzko, editor, *Advances in Cryptology - CRYPTO '86, Santa Barbara, California, USA, 1986, Proceedings*, volume 263 of *Lecture Notes in Computer Science*, pages 186–194. Springer, 1986.

27. Irene Giacomelli, Jesper Madsen, and Claudio Orlandi. ZKBoo: Faster Zero-Knowledge for Boolean Circuits. In *25th USENIX Security Symposium (USENIX Security 16)*, pages 1069–1083, 2016.
28. Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from secure multi-party computation. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 21–30, 2007.
29. Daniel Kales and Greg Zaverucha. An attack on some signature schemes constructed from five-pass identification schemes. In Stephan Krenn, Haya Shulman, and Serge Vaudenay, editors, *Cryptology and Network Security*, pages 3–22, Cham, 2020. Springer International Publishing.
30. Daniel Kales and Greg Zaverucha. Efficient lifting for shorter zero-knowledge proofs and post-quantum signatures. Cryptology ePrint Archive, Paper 2022/588, 2022. <https://eprint.iacr.org/2022/588>.
31. Seongkwang Kim, Jihoon Cho, Mingyu Cho, Jincheol Ha, Jihoon Kwon, Byeonghak Lee, Joohee Lee, Jooyoung Lee, Sangyub Lee, Dukjae Moon, Mincheol Son, and Hyojin Yoon. AIMER specification. 2023. available at <https://aimer-signature.org/docs/AIMer-NIST-Document.pdf>.
32. Carlos Aguilar Melchor, Thibault Feneuil, Nicolas Gama, Shay Gueron, James Howe, David Joseph, Antoine Joux, Edoardo Persichetti, Tovahery H., Randrianarisoa, Matthieu Rivain, and Dongze Yue. SDITH specification. 2023. available at <https://sdith.org/docs/sdith-v1.0.pdf>.
33. Carlos Aguilar Melchor, Nicolas Gama, James Howe, Andreas Hülsing, David Joseph, and Dongze Yue. The return of the sdith. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology - EUROCRYPT 2023 - 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, April 23-27, 2023, Proceedings, Part V*, volume 14008 of *Lecture Notes in Computer Science*, pages 564–596. Springer, 2023.
34. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):34:1–34:40, 2009.
35. Greg Zaverucha, Melissa Chase, David Derler, Steven Goldfeder, Claudio Orlandi, Sebastian Ramacher, Christian Rechberger, Daniel Slamanig, Jonathan Katz, Xiao Wang, Vladimir Kolesnikov, and Daniel Kales. *Picnic* : Algorithm specification and design document.