

SOFIA: \mathcal{MQ} -based signatures in the QROM

Ming-Shing Chen¹ and Andreas Hülsing² and Joost Rijneveld³ and
Simona Samardjiska³ and Peter Schwabe³

¹ Department of Electrical Engineering, National Taiwan University,
and Research Center for Information Technology Innovation, Academia Sinica,
Taipei, Taiwan

`mschen@crypto.tw`

² Department of Mathematics and Computer Science,
Technische Universiteit Eindhoven, Eindhoven, The Netherlands

`andreas@huel sing.net`

³ Digital Security Group, Radboud University, Nijmegen, The Netherlands
`joost@joostrijneveld.nl`, `simonas@cs.ru.nl`, `peter@cryptojedi.org`

Abstract. We propose SOFIA, the first \mathcal{MQ} -based signature scheme provably secure in the quantum-accessible random oracle model (QROM). Our construction relies on an extended version of Unruh’s transform for 5-pass identification schemes that we describe and prove secure both in the ROM and QROM.

Based on a detailed security analysis, we provide concrete parameters for SOFIA that achieve 128-bit post-quantum security. The result is SOFIA-4-128 with parameters carefully optimized to minimize signature size and maximize performance. SOFIA-4-128 comes with an implementation targeting recent Intel processors with the AVX2 vector-instruction set; the implementation is fully protected against timing attacks.

Keywords: Post-quantum cryptography, multivariate cryptography, 5-pass identification schemes, QROM, Unruh’s transform, vectorized implementation.

1 Introduction

At Asiacrypt 2016 [11], we presented a post-quantum signature scheme called MQDSS, obtained by applying a generalized Fiat-Shamir transform to a 5-pass identification schemes (IDS) with security based on the hardness of solving a system of multivariate quadratic equations (\mathcal{MQ} problem). Unlike previous \mathcal{MQ} signature schemes, MQDSS comes with a reduction from a random instance of \mathcal{MQ} ; it does not need additional assumptions on the hardness of related problems like the Isomorphism of Polynomials (IP) [29] or MinRank [13,20].

Unfortunately, the security reduction of MQDSS is in the random oracle model and highly non-tight, while our ultimate goal is (as stated in [11]) a scheme

This work was supported by the Netherlands Organization for Scientific Research (NWO) under Veni 2013 project 13114, by the European Commission through the ICT Programme under contract ICT-645622 PQCRYPTO and by the Faculty of Computer Science and Engineering at the “Ss. Cyril and Methodius” University.

with a tight reduction from \mathcal{MQ} in the quantum random oracle model (QROM) or even in the standard model. In this paper, we take a step closer towards such a scheme. More specifically, we propose SOFIA, a digital signature scheme that is provably EU-CMA secure in the QROM if the \mathcal{MQ} problem is hard and allows for a tight reduction in the ROM (albeit not in the QROM).

To achieve this, we start from Unruh’s transform [33] for transforming Σ -protocols to NIZK proofs (and signatures) in the QROM. The reason for a different transform comes from the inherent problems of the Fiat-Shamir transform (and also the generalization to 5-pass schemes) in the QROM. Namely, the proof technique introduced by Pointcheval and Stern [30] requires rewinding of the adversary and adaptively programming the random oracle. Not only does this cause problems in the QROM, but it also produces non-tight proofs in the ROM. Unruh’s transform avoids these problems by adopting and tweaking an idea from Fischlin’s transform [21] that solves the rewinding problem.

Recently, Kiltz, Loss, and Pan [27] considered a generalization of the Fiat-Shamir transform to 5-pass schemes, and provided a tight reduction in the ROM. However, the technique faces similar issues when it comes to the QROM such as adaptive programming of the random oracles. Hence, no proof in the QROM is known. Therefore, it is not applicable for SOFIA. In concurrent work, Kiltz, Lyubashevsky, and Schaffner [28] provide a viable alternative to Unruh’s transform in the QROM. The authors prove security of the Fiat-Shamir transform in the QROM, using the additional assumption of “lossiness” of the IDS. While this requires modifications in the IDS and re-parametrization of MQDSS, it seems promising future work to see whether one can obtain a more efficient scheme with a QROM proof this way.

MQDSS builds on a 5-pass \mathcal{MQ} -based IDS from [31]. While [31] also introduces a 3-pass IDS, we showed in [11] that the 5-pass scheme leads to smaller signatures due to a smaller soundness error. Hence, we do not simply apply the Unruh transform to the 3-pass IDS but extend it such that it applies to any 5-pass IDS with binary second challenge (named $q2$ -IDS in [11]) and thus to the \mathcal{MQ} -based 5-pass IDS from [31]. We prove that the signature scheme resulting from the application of the transform is post-quantum EU-CMA secure (PQ-EU-CMA) in the QROM. This proof follows a two-step approach: We first give a (tight) proof in the ROM, and then discuss the changes necessary to carry over to the QROM. We then instantiate the construction with the \mathcal{MQ} -based 5-pass IDS by Sakumoto, Shirai, and Hiwatari [31] and provide various optimizations particularly suited for this specific IDS. These optimizations almost halve the size of the signature compared to the non-optimized generic transform.

We instantiate SOFIA with carefully optimized parameters aiming at the 128-bit post-quantum security level; we refer to this instance as SOFIA-4-128. A comparison with MQDSS-31-64 from [11], which targets the same security level, shows that the improvements in security guarantees come at a cost: with 123 KiB, SOFIA-4-128 signatures are about a factor of 3 larger than MQDSS-31-64 signatures and our optimized SOFIA-4-128 software takes about a factor of 3 longer for both signing and verification than the optimized one for MQDSS-

31-64. However, like MQDSS, SOFIA features extremely short keys; specifically, SOFIA-4-128 public keys have 64 bytes and the secret keys have 32 bytes.

SOFIA is not the first concrete signature scheme with a proof in the QROM. Notably, TESLA-2 [1] is a lattice-based signature scheme with a reduction in the QROM, while Picnic-10-38 [10] is the result of constructing a signature scheme from a symmetric primitive using the transform by Unruh [33] that was mentioned above. Relying on even more conservative assumptions, the hash-based signature scheme SPHINCS-256 [6] has a tight proof in the standard model. Although SOFIA-4-128 remains faster than SPHINCS-256 (which is, because of its standard model assumptions, arguably the ‘scheme to beat’), we do significantly exceed its 40 KiB signature size. Conversely, but on a similar note, SOFIA-4-128 outperforms Picnic-10-38 both in terms of signing speed and signature size. TESLA-2 remains the ‘odd one out’ with its small signatures but much larger keys; it strongly depends on context whether this is an upside or a problem. See Table 3 for a numeric overview of the comparison.

Organization of the paper. Section 2 gives the necessary background on identification schemes and signature schemes. Section 3 presents the extended Unruh transform to support $q2$ identification schemes. Section 4 revisits the 5-pass identification scheme introduced in [31]. Section 5 introduces the SOFIA signature scheme and finally Section 6 explains our parameter choices for SOFIA-4-128 and gives details of our optimized implementation.

Availability of software. We place all software presented in this paper into the public domain to maximize re-usability of our results. It is available for download at <https://joostrijneveld.nl/papers/sofia>.

2 Preliminaries

In the following we provide basic definitions used throughout this work. We are concerned with post-quantum security, i.e., a setting where honest parties use classical computers but adversaries might have access to a quantum computer. Therefore, we adapt some common security notions accordingly, modeling adversaries as quantum algorithms.

Digital signatures. In this work we are concerned with the construction of digital-signature schemes. Due to space limitations, we omit the standard definitions for digital signatures and their security. They are included in the full version of the paper.

Identification Schemes. An identification scheme (IDS) is a protocol that allows a prover \mathcal{P} to prove its identity to a verifier \mathcal{V} . More formally:

Definition 2.1 (Identification scheme). *An identification scheme with security parameter k , denoted $\text{IDS}(1^k)$, is a triplet of PPT algorithms $\text{IDS} = (\text{KGen}, \mathcal{P}, \mathcal{V})$ such that the key generation algorithm KGen is a probabilistic algorithm that outputs a key pair (sk, pk) , and \mathcal{P} and \mathcal{V} are interactive algorithms, executing a common protocol. The prover \mathcal{P} takes as input a secret key sk and*

the verifier \mathcal{V} takes as input a public key pk . At the conclusion of the protocol, \mathcal{V} outputs a bit b with $b = 1$ indicating “accept” and $b = 0$ indicating “reject”.

For correctness of an IDS, we require that for all $(\text{pk}, \text{sk}) \leftarrow \text{KGen}()$ we have $\Pr[\langle \mathcal{P}(\text{sk}), \mathcal{V}(\text{pk}) \rangle = 1] = 1$, where $\langle \mathcal{P}(\text{sk}), \mathcal{V}(\text{pk}) \rangle$ refers to the common execution of the protocol between \mathcal{P} with input sk and \mathcal{V} on input pk . We denote by $\text{trans}(\langle \mathcal{P}(\text{sk}), \mathcal{V}(\text{pk}) \rangle)$ the transcript of messages exchanged during this execution.

In this work we are concerned with canonical 5-pass IDS, where the prover and the verifier exchange two challenges and replies. More formally:

Definition 2.2 (Canonical 5-pass identification schemes). Consider $\text{IDS} = (\text{KGen}, \mathcal{P}, \mathcal{V})$, a 5-pass identification scheme with two challenge spaces \mathcal{C}_1 and \mathcal{C}_2 . We call IDS a canonical 5-pass identification scheme if the prover can be split into three subroutines $\mathcal{P} = (\mathcal{P}_0, \mathcal{P}_1, \mathcal{P}_2)$ and the verifier into three subroutines $\mathcal{V} = (\text{ChS}_1, \text{ChS}_2, \text{Vf})$ such that:

$\mathcal{P}_0(\text{sk})$ computes the initial commitment com sent as the first message and a state state fed forward to \mathcal{P}_1 . ChS_1 computes the first challenge message $\text{ch}_1 \leftarrow_R \mathcal{C}_1$, sampling at random from the challenge space \mathcal{C}_1 . $\mathcal{P}_1(\text{state}, \text{ch}_1)$ computes the first response resp_1 of the prover (and updates the state state) given access to the state and the first challenge. ChS_2 computes the second challenge message $\text{ch}_2 \leftarrow_R \mathcal{C}_2$. $\mathcal{P}_2(\text{state}, \text{ch}_2)$ computes the second response resp_2 of the prover given access to the state and the second challenge. $\text{Vf}(\text{pk}, \text{com}, \text{ch}_1, \text{resp}_1, \text{ch}_2, \text{resp}_2)$ upon access to the public key and the whole transcript outputs \mathcal{V} 's final decision.

Note that the state forwarded among the prover algorithms can contain all inputs to previous prover algorithms if they are needed later. We also assume that the verifier keeps all sent and received messages to feed them to Vf .

We will consider a particular type of 5-pass identification protocols where the size of the two challenge spaces is restricted to q and 2.

Definition 2.3 ($q2$ -Identification scheme). A $q2$ -Identification scheme IDS with security parameter $k \in \mathbb{N}$ is a canonical 5-pass identification scheme where for the challenge spaces \mathcal{C}_1 and \mathcal{C}_2 it holds that $|\mathcal{C}_1| = q$ and $|\mathcal{C}_2| = 2$. Moreover, the probability that the commitment com takes a given value is $\leq 2^{-k}$, where the probability is taken over the random choice of the input and the used randomness.

Our goal is to construct signature schemes from identification schemes. It is well known that passively secure identification schemes suffice for this. In this setting, security is defined in terms of two properties: special soundness and honest-verifier zero-knowledge (HVZK). To prove security of our signature scheme, we will make use of the existence of so called $q2$ -extractor which is a variant of special soundness.

Definition 2.4 ((computational) PQ-HVZK). Let $k \in \mathbb{N}$, $\text{IDS}(1^k) = (\text{KGen}, \mathcal{P}, \mathcal{V})$ an identification scheme with security parameter k . We say that IDS is computational post-quantum honest-verifier zero-knowledge (PQ-HVZK) if there exists a probabilistic polynomial time algorithm \mathcal{S} , called the simulator, such that for any polynomial time quantum algorithm \mathcal{A} and $(\text{pk}, \text{sk}) \leftarrow \text{KGen}()$:

$$\text{Succ}_{\text{IDS}(1^k)}^{pq-hvzk}(\mathcal{A}) = |\Pr[1 \leftarrow \mathcal{A}(\text{sk}, \text{pk}, \text{trans}(\langle \mathcal{P}(\text{sk}), \mathcal{V}(\text{pk}) \rangle))] - \Pr[1 \leftarrow \mathcal{A}(\text{sk}, \text{pk}, \mathcal{S}(\text{pk}))]| = \text{negl}(k).$$

Intuitively it must be hard for any cryptographic scheme to derive a valid secret key given a public key. To formally capture this intuition, we need to define what valid means. For this we define the notion of a key relation.

Definition 2.5 (Key relation). *Let IDS be a $q2$ -Identification scheme and R a relation. We say IDS has key relation R iff R is the minimal relation such that*

$$\forall (\text{pk}, \text{sk}) \leftarrow \text{KGen}() : (\text{pk}, \text{sk}) \in R$$

Now that we have defined what valid means, we can define key-one-wayness.

Definition 2.6 (PQ-KOW). *Let $k \in \mathbb{N}$ be the security parameter, $\text{IDS}(1^k)$ be a $q2$ -Identification scheme with key relation R . We call IDS post-quantum key-one-way (PQ-KOW) (with respect to key relation R) if for all quantum polynomial time algorithms \mathcal{A} ,*

$$\text{Succ}_{\text{IDS}(1^k)}^{pq-kow}(\mathcal{A}) = \Pr[(\text{pk}, \text{sk}) \leftarrow \text{KGen}(), \text{sk}' \leftarrow \mathcal{A}(\text{pk}) : (\text{pk}, \text{sk}') \in R] = \text{negl}(k)$$

In [11] it was shown that in general, for $q2$ -Identification Schemes, it is not possible to efficiently extract a matching secret key from two related transcripts alone (as in the case of 3-pass schemes fulfilling special soundness). In order to capture the nature of these schemes and provide sufficient conditions for efficient extraction, we proposed the definition of a $q2$ -Extractor. In the following we give a slightly refined definition that uses the notion of key relation to capture what kind of secret key the extractor returns.

Definition 2.7 ($q2$ -Extractor). *Let $\text{IDS}(1^k)$ be a $q2$ -Identification scheme with key relation R . We say that $\text{IDS}(1^k)$ has a $q2$ -Extractor if there exists a polynomial time algorithm \mathcal{K}_{IDS} , the extractor, that, given a public key pk and four valid transcripts with respect to pk*

$$\begin{aligned} \text{trans}^{(1)} &= (\text{com}, \text{ch}_1, \text{resp}_1, \text{ch}_2, \text{resp}_2), & \text{trans}^{(3)} &= (\text{com}, \text{ch}'_1, \text{resp}'_1, \text{ch}_2, \text{resp}_2), \\ \text{trans}^{(2)} &= (\text{com}, \text{ch}_1, \text{resp}_1, \text{ch}'_2, \text{resp}'_2), & \text{trans}^{(4)} &= (\text{com}, \text{ch}'_1, \text{resp}'_1, \text{ch}'_2, \text{resp}'_2), \end{aligned} \quad (1)$$

where $\text{ch}_1 \neq \text{ch}'_1$ and $\text{ch}_2 \neq \text{ch}'_2$, outputs a secret key sk such that $(\text{pk}, \text{sk}) \in R$ with non-negligible success probability in k .

3 From $q2$ -IDS to signatures in the QROM

In [11], we showed that the Fiat-Shamir transform can be generalized to the case of 5-pass IDS whose ChS_2 is bounded to two elements. We showed that the Pointcheval-Stern proof [30] can be extended to this case, and the obtained signature scheme can be shown EU-CMA secure in the random oracle model. This result is further extended to any $2n+1$ round identification scheme that fulfills a

certain kind of special soundness in [15]. However, similar to the standard Fiat-Shamir transform, these proofs rely on the forking lemma, which introduces two serious problems in the post-quantum setting: rewinding of the adversary, and adaptively programming the random oracle. While it is known how to deal with the latter [33], the former seems to become a real show stopper [3]. The only known way (at the time of writing⁴) to fix the Fiat-Shamir transform in the QROM setting is using oblivious commitments [14], which are a certain kind of trapdoor commitments, effectively avoiding rewinding at the cost of introducing the necessity of a trapdoor function. This makes the solution not applicable in our setting as there are no known trapdoor functions with a reduction from the \mathcal{MQ} -problem.

In [33], Unruh proposes a different transform, based on Fischlin’s transform [21], that turns 3-pass zero-knowledge proofs into non-interactive ones in the QROM. In addition, Unruh shows how to use his transform to obtain a signature scheme. The transform essentially works by “unrolling” Fischlin’s transform and then applying a few tweaks. This works, as Fischlin’s transform already avoids rewinding. The basic idea is to let the signer generate several transcripts for a commitment. This is iterated for several initial commitments. Next, the signer “blinds” all responses in the transcripts by applying a length-preserving hash. All the obtained data is hashed together with the public key and the message to obtain a challenge vector. This challenge vector determines one transcript per commitment that has to be unblinded, i.e., for which the response must be included in the signature. The signature consists of all the transcripts with “blinded” responses and the unblinded responses for the transcripts identified by the challenge vector. The reasoning behind the transform is that without knowing the secret key, a forger cannot know sufficiently many valid openings to be able to include all the challenged responses. On the other hand, a security reduction can replace the length-preserving hash (modeled as QRO) by an invertible function (e.g. a QPRP). That way, a reduction can “unblind” the remaining responses in the signature by inverting the function. Now, it can be argued that an adversary with non-negligible success probability must have known several valid transcripts for at least one commitment. The unblinding reveals those transcripts and they can be used to run the extractor.

Here, we show that a similar transform can be applied to 5-pass IDS with a binary second challenge (i.e., $q2$ -IDS). Basically, we treat the second challenge-response round like the first. However, as we have a binary second challenge, we ask that for each first challenge, a transcript for both values of the second challenge is generated. The main difference between the security reduction of Unruh’s transform and our extension to $q2$ -IDS is a more involved argument to show that we get sufficiently many valid transcripts that follow the pattern needed to extract a valid secret key. As this argument is essentially independent of the RO, we first give a proof in the classical ROM. This also allows us to show that the reduction is tight in the ROM. Afterwards we describe how things

⁴ Very recently, Kiltz et al. [28] proposed the use of “lossy” IDS which enabled them to prove security of the Fiat-Shamir transform in the QROM.

```

Sign(sk, M)


---


For  $j \in \{1, \dots, r\}$  do
   $(\text{state}^{(j)}, \text{com}^{(j)}) \leftarrow \mathcal{P}_0(\text{sk})$ 
  For  $i \in \{1, \dots, t\}$  do
     $\text{ch}_1^{(i,j)} \leftarrow_R \text{ChS}_1 \setminus \{\text{ch}_1^{(1,j)}, \dots, \text{ch}_1^{(i-1,j)}\}$ 
     $(\text{state}^{(i,j)}, \text{resp}_1^{(i,j)}) \leftarrow \mathcal{P}_1(\text{state}^{(j)}, \text{ch}_1^{(i,j)})$ 
     $\text{cr}_1^{(i,j)} \leftarrow \text{H}_1(\text{resp}_1^{(i,j)})$ 
     $\text{resp}_2^{(i,j,0)} \leftarrow \mathcal{P}_2(\text{state}^{(i,j)}, \text{ch}_2 = 0), \text{resp}_2^{(i,j,1)} \leftarrow \mathcal{P}_2(\text{state}^{(i,j)}, \text{ch}_2 = 1)$ 
     $\text{cr}_2^{(i,j,0)} \leftarrow \text{H}_2(\text{resp}_2^{(i,j,0)}), \text{cr}_2^{(i,j,1)} \leftarrow \text{H}_2(\text{resp}_2^{(i,j,1)})$ 
   $\text{trans}_{\text{full}}(j) := \text{com}^{(j)}, \{\text{ch}_1^{(i,j)}, \text{cr}_1^{(i,j)}, (\text{cr}_2^{(i,j,0)}, \text{cr}_2^{(i,j,1)})\}_{i=1}^t$ 
 $\text{md} \leftarrow \mathcal{H}(\text{pk}, M, \{\text{trans}_{\text{full}}(j)\}_{j=1}^r)$ 
  Read  $\text{md}$  as vector  $((I_1, B_1), \dots, (I_r, B_r))$ 
   $\text{trans}_{\text{red}}(j) := \text{com}^{(j)}, \{\text{ch}_1^{(i,j)}, \text{cr}_1^{(i,j)}, (\text{cr}_2^{(i,j,0)}, \text{cr}_2^{(i,j,1)})\}_{i \neq I_j, i=1}^t$ 
 $\sigma := \left( \text{md}, \{\text{trans}_{\text{red}}(j), \text{ch}_1^{(I_j,j)}, \text{resp}_1^{(I_j,j)}, \text{resp}_2^{(I_j,j,B_j)}, \text{cr}_2^{(I_j,j,\neg B_j)}\}_{j=1}^r \right)$ 

```

Fig. 1. Signature generation

change in the QROM along the lines of Unruh’s QROM proof. This is where the reduction becomes loose. It remains an interesting open question whether this is a fundamental issue with QROM reductions or the existing techniques are just not sufficiently evolved, yet.

3.1 Extending Unruh’s transform to $q2$ -IDS

Let $\text{IDS} = (\text{KGen}, \mathcal{P}, \mathcal{V})$ be a $q2$ -IDS, with $\mathcal{P} = (\mathcal{P}_0, \mathcal{P}_1, \mathcal{P}_2)$, $\mathcal{V} = (\text{ChS}_1, \text{ChS}_2, \text{Vf})$, and let $r, t \in \mathbb{N}$ be two parameters, where $2 \leq t \leq q$. Moreover let $\text{H}_1 : \{0, 1\}^{|\text{resp}_1|} \rightarrow \{0, 1\}^{|\text{resp}_1|}$, $\text{H}_2 : \{0, 1\}^{|\text{resp}_2|} \rightarrow \{0, 1\}^{|\text{resp}_2|}$, and $\mathcal{H} : \{0, 1\}^* \rightarrow \{0, 1\}^{\lceil \log 2^t \rceil r}$ be hash functions, later modeled as random oracles. We define the following digital signature scheme $(\text{KGen}, \text{Sign}, \text{Vf})$. The key generation algorithm just runs $\text{IDS.KGen}()$. Signature and verification algorithms are given in Figures 1 and 2.

For ease of exposition, we will use the notation $T(j, i, b)$ for a string that has the format of a transcript of the IDS (not necessarily a valid transcript), corresponding to the j -th round of the non-interactive protocol, with i and b being the indices of the corresponding challenges ch_1 and ch_2 , i.e.

$$T(j, i, b) := (\text{com}^{(j)}, \text{ch}_1^{(i,j)}, \text{resp}_1^{(i,j)}, \text{ch}_2 = b, \text{resp}_2^{(i,j,b)}),$$

where $j \in \{1, \dots, r\}$, $i \in \{1, \dots, t\}$, $b \in \{0, 1\}$.

3.2 PQ-EU-CMA-Security in the ROM

In the following, we first establish post-quantum security under key-only attacks (PQ-KOA). More specifically, we will show that a successful KOA-forgery \mathcal{A} can

<p>$\mathbf{Vf}(\mathbf{pk}, \sigma, \mathbf{M})$</p> <hr/> <p>Read md as vector $((I_1, B_1), \dots, (I_r, B_r))$</p> <p>For $j \in \{1, \dots, r\}$ do</p> <p style="padding-left: 20px;">$\mathbf{ct}_1^{(I_j, j)} \leftarrow \mathbf{H}_1(\mathbf{resp}_1^{(I_j, j)}), \quad \mathbf{ct}_2^{(I_j, j, B_j)} \leftarrow \mathbf{H}_2(\mathbf{resp}_2^{(I_j, j, B_j)})$</p> <p>$\mathbf{md}' \leftarrow \mathcal{H}(\mathbf{pk}, M, \{\mathbf{trans}_{\text{full}}(j)\}_{j=1}^r)$</p> <p>Check that $\mathbf{md}' \stackrel{?}{=} \mathbf{md}$</p> <p>For $j \in \{1, \dots, r\}$ do</p> <p style="padding-left: 20px;">Check that $\mathbf{ch}_1^{(1, j)}, \dots, \mathbf{ch}_1^{(t, j)}$ are all distinct</p> <p style="padding-left: 20px;">Check $1 \stackrel{?}{=} b \leftarrow \mathbf{Vf}(\mathbf{pk}, \mathbf{com}^{(j)}, \mathbf{ch}_1^{(I_j, j)}, \mathbf{resp}_1^{(I_j, j)}, B_j, \mathbf{resp}_2^{(I_j, j, B_j)})$</p> <p>If all checks succeed, output success.</p>

Fig. 2. Verification

be used to extract a valid secret key for the underlying IDS. Afterwards, we will extend the result to existential unforgeability under chosen message attacks.

PQ-KOW \Rightarrow PQ-KOA. The following lemma gives an exact relation between the key-one-wayness of the identification scheme and the security of the proposed signature scheme under key-only attacks.

Lemma 3.1. *Let $k, t, r \in \mathbb{N}$ be the parameters of the signature scheme from Figures 1 and 2, using a $q2$ -IDS that has a key relation R , a $q2$ -extractor, and is PQ-KOW secure. Let \mathcal{A} be a quantum algorithm that implements a KOA forger which given only the public key \mathbf{pk} outputs a valid message-signature pair (M, σ) with probability ϵ . Then, in the random oracle model there exists an algorithm $\mathcal{M}^{\mathcal{A}}$ that given oracle access to any such \mathcal{A} breaks the KOW security of IDS in essentially the same running time as the given \mathcal{A} and with success probability*

$$\epsilon' \geq \epsilon - (q_{\mathcal{H}} + 1)2^{-r \log \frac{2t}{t+1}}. \quad (2)$$

Moreover, $\mathcal{M}^{\mathcal{A}}$ only manipulates the random oracles $\mathbf{H}_1, \mathbf{H}_2$ and leaves \mathcal{H} untouched.

Proof. We show how to construct such an algorithm $\mathcal{M}^{\mathcal{A}}$. On input of an IDS public key \mathbf{pk} , $\mathcal{M}^{\mathcal{A}}$ first runs $\mathcal{A}(\mathbf{pk})$. Let $\mathcal{E}_{\mathcal{A}}$ be the event that \mathcal{A} outputs a valid message-signature pair (M, σ) with

$$\sigma = \left(\mathbf{md}, \left\{ \mathbf{trans}_{\text{red}}(j), \mathbf{ch}_1^{(I_j, j)}, \mathbf{resp}_1^{(I_j, j)}, \mathbf{resp}_2^{(I_j, j, B_j)}, \mathbf{ct}_2^{(I_j, j, \neg B_j)} \right\}_{j=1}^r \right).$$

Then $\mathcal{E}_{\mathcal{A}}$ implies that for every $j \in \{1, \dots, r\}$, $T(j, I_j, B_j)$ is a valid transcript of IDS and the Verifier \mathbf{Vf} accepts. Now, our goal is to use the $q2$ -extractor to extract. This means, we need to obtain four valid transcripts $T(j, i_1, 0)$, $T(j, i_1, 1)$, $T(j, i_2, 0)$, $T(j, i_2, 1)$ for some $j \in \{1, \dots, r\}$. To this end, $\mathcal{M}^{\mathcal{A}}$ simulates the random oracles \mathbf{H}_1 and \mathbf{H}_2 for \mathcal{A} in the common way. The important point is that this way $\mathcal{M}^{\mathcal{A}}$ learns all of \mathcal{A} 's queries together with the given responses. Hence, when given \mathcal{A} 's forgery, $\mathcal{M}^{\mathcal{A}}$ can open all blinded responses in the signature.

Now, \mathcal{M}^A will only fail to extract if among all the $2tr$ opened transcripts of the signature, there are no four valid transcripts with the above pattern. Consider the event $\mathcal{E}_{\text{-ext}}$ which describes this case.

$\mathcal{E}_{\text{-ext}}$: $\forall j \in \{1, \dots, r\}$, and $\forall i_1, i_2 \in \{1, \dots, t\}$, $i_1 \neq i_2$, at least one of $T(j, i_1, 0)$, $T(j, i_1, 1)$, $T(j, i_2, 0)$, $T(j, i_2, 1)$ is not a valid transcript of the IDS.

We will upper bound $\Pr[(\mathcal{E}_A \cap \mathcal{E}_{\text{-ext}})]$ and thereby lower bound \mathcal{M}^A 's success probability. Let (M, σ) be \mathcal{A} 's output under the event $(\mathcal{E}_A \cap \mathcal{E}_{\text{-ext}})$. First, (M, σ) must be valid because of \mathcal{E}_A . Now, consider the set $\mathcal{S}_{\text{-ext}}$ of tuples $(\text{pk}, M, \{\text{trans}_{\text{full}}(j)\}_{j=1}^r)$, such that for every $j \in \{1, \dots, r\}$ there is at most one I_j^* with $T(j, I_j^*, 0)$ and $T(j, I_j^*, 1)$ being valid transcripts of IDS. It is clear that \mathcal{A} 's output under the event $\mathcal{E}_A \cap \mathcal{E}_{\text{-ext}}$ must come from $\mathcal{S}_{\text{-ext}}$. Indeed, if a tuple does not satisfy the given condition, then there exist at least two indices I_j^*, I_j^{**} such that $T(j, I_j^*, 0)$, $T(j, I_j^*, 1)$, $T(j, I_j^{**}, 0)$, $T(j, I_j^{**}, 1)$ are valid transcripts of IDS, which is in contradiction to the event $\mathcal{E}_{\text{-ext}}$.

Let $(\text{pk}, M, \{\text{trans}_{\text{full}}(j)\}_{j=1}^r)$ be such a tuple. Then the indexes that define the required openings in σ are obtained as the output of the random oracle \mathcal{H} on input of the tuple, i.e. $((I_1, B_1), \dots, (I_r, B_r)) \leftarrow \mathcal{H}(\text{pk}, M, \{\text{trans}_{\text{full}}(j)\}_{j=1}^r)$. In order for the signature to pass verification, for each $j \in \{1, \dots, r\}$, the transcript $T(j, I_j, B_j)$ must be valid. Given the conditions of $\mathcal{E}_{\text{-ext}}$, for each $j \in \{1, \dots, r\}$, there are at most $t + 1$ valid transcripts per j . Hence for the entire $((I_1, B_1), \dots, (I_r, B_r))$ at most $(t + 1)^r$ possible values. Thus, the probability for the adversary to produce a valid signature from such a tuple is $\frac{(t+1)^r}{(2t)^r} = 2^{-r \log \frac{2t}{t+1}}$.

Now let $q_{\mathcal{H}}$ be the number of queries of the adversary to the oracle \mathcal{H} . Then

$$\Pr(\mathcal{E}_A \cap \mathcal{E}_{\text{-ext}}) \leq (q_{\mathcal{H}} + 1)2^{-r \log \frac{2t}{t+1}},$$

as \mathcal{A} can try at most $q_{\mathcal{H}}$ tuples to obtain a valid signature and output a signature based on a new tuple otherwise. Towards obtaining a bound on \mathcal{M}^A 's success probability, note that \mathcal{M}^A succeeds in the event $(\mathcal{E}_A \cap \neg \mathcal{E}_{\text{-ext}})$, and

$$\Pr(\mathcal{E}_A \cap \neg \mathcal{E}_{\text{-ext}}) = \Pr(\mathcal{E}_A) - \Pr(\mathcal{E}_A \cap \mathcal{E}_{\text{-ext}}) \geq \epsilon - (q_{\mathcal{H}} + 1)2^{-r \log \frac{2t}{t+1}}.$$

This proves the claimed bound. \square

PQ-KOA \Rightarrow PQ-EU-CMA. Given the above lemma, it suffices to reduce PQ-KOA to PQ-EU-CMA security to eventually prove PQ-EU-CMA security of the proposed scheme, i.e. we have to show that we can answer an adversary's signature queries without knowledge of a secret key. This is done in the following lemma. Afterwards we can derive the main theorem of the section.

Lemma 3.2. *Let $k, t, r \in \mathbb{N}$ be the parameters of the signature scheme from Figures 1 and 2 above, using a $q2$ -IDS that is PQ-HVZK. Let \mathcal{A} be a quantum algorithm that breaks the PQ-EU-CMA security of the signature scheme with probability ϵ . Then, in the random oracle model there exists an algorithm \mathcal{M}^A*

that breaks the PQ-KOA security of the signature scheme in essentially the same running time as \mathcal{A} and with success probability

$$\epsilon' \geq \epsilon(1 - q_{\text{Sign}}q_{\mathcal{H}}2^{-rk}). \quad (3)$$

Moreover, $\mathcal{M}^{\mathcal{A}}$ only manipulates the random oracle \mathcal{H} and leaves H_1, H_2 untouched.

Proof. We show how to construct $\mathcal{M}^{\mathcal{A}}$ that on input a public key pk of the signature scheme (which is also a public key for IDS), access to a HVZK-simulator \mathcal{S}_{IDS} for IDS and the random oracles H_1, H_2, \mathcal{H} , breaks the KOA security of the signature scheme. The running time and success probability of $\mathcal{M}^{\mathcal{A}}$ are essentially the same as that of \mathcal{A} up to a negligible difference.

Upon receiving the public key pk , $\mathcal{M}^{\mathcal{A}}$ runs $\mathcal{A}(\text{pk})$, simulating all signature and random oracle queries for \mathcal{A} . Whenever \mathcal{A} queries H_1 or H_2 , $\mathcal{M}^{\mathcal{A}}$ simply forwards the query to his respective RO. For \mathcal{H} , $\mathcal{M}^{\mathcal{A}}$ keeps a local list $\mathcal{L}_{\mathcal{H}}$. Whenever \mathcal{A} queries \mathcal{H} , $\mathcal{M}^{\mathcal{A}}$ first checks $\mathcal{L}_{\mathcal{H}}$ and returns the stored answer if one exists. Otherwise, $\mathcal{M}^{\mathcal{A}}$ forwards the query to his oracle \mathcal{H} and stores the query together with the result in $\mathcal{L}_{\mathcal{H}}$ before returning the response. Whenever \mathcal{A} makes a signature query on a message M , $\mathcal{M}^{\mathcal{A}}$ does the following:

1. Samples $\tilde{\text{md}} \leftarrow_R \{0, 1\}^{\lceil \log 2t \rceil r}$ and interprets it as challenge string, i.e., $((I_1, B_1), \dots, (I_r, B_r)) := \tilde{\text{md}}$.
2. Runs the HVZK-simulator \mathcal{S}_{IDS} r times to obtain r valid transcripts of IDS:

$$\left\{ (\text{com}^{(j)}, \text{ch}_1^{(I_j, j)}, \text{resp}_1^{(I_j, j)}, \text{ch}_2^{(j)}, \text{resp}_2^{(I_j, j, B_j)}) \right\}_{j=1}^r,$$

and uses them as the challenged transcripts $T(j, I_j, B_j)$ for $j \in \{1, \dots, r\}$.

3. Blinds the responses $\text{resp}_1^{(I_j, j)}$ and $\text{resp}_2^{(I_j, j, B_j)}$ for every $j \in \{1, \dots, r\}$:

$$\text{cr}_1^{(I_j, j)} \leftarrow H_1(\text{resp}_1^{(I_j, j)}), \quad \text{cr}_2^{(I_j, j, B_j)} \leftarrow H_2(\text{resp}_2^{(I_j, j, B_j)}).$$

4. For all $j \in \{1, \dots, r\}$, and all $(i, b) \in \{1, \dots, t\} \times \{0, 1\} \setminus \{(I_j, B_j)\}_{j=1}^r$,
 - samples a first challenge $\text{ch}_1^{(i, j)} \leftarrow_R \text{ChS}_1 \setminus \{\text{ch}_1^{(I_j, j)}, \text{ch}_1^{(1, j)}, \dots, \text{ch}_1^{(i-1, j)}\}$,
 - samples fake responses $\text{resp}_1^{(i, j)} \leftarrow_R \text{RespS}_1$, $\text{resp}_2^{(i, j, b)} \leftarrow_R \text{RespS}_2$,
 - blinds the fake responses $\text{cr}_1^{(i, j)} \leftarrow H_1(\text{resp}_1^{(i, j)})$, $\text{cr}_2^{(i, j, b)} \leftarrow H_2(\text{resp}_2^{(i, j, b)})$,
 - sets $\text{trans}_{\text{full}}(j) := \text{com}^{(j)}, \left\{ \text{ch}_1^{(i, j)}, \text{cr}_1^{(i, j)}, (\text{cr}_2^{(i, j, 0)}, \text{cr}_2^{(i, j, 1)}) \right\}_{i=1}^t$.
5. Checks if there is already an entry for $(\text{pk}, M, \{\text{trans}_{\text{full}}(j)\}_{j=1}^r)$ in $\mathcal{L}_{\mathcal{H}}$. If so, $\mathcal{M}^{\mathcal{A}}$ aborts. Otherwise, $\mathcal{M}^{\mathcal{A}}$ stores $((\text{pk}, M, \{\text{trans}_{\text{full}}(j)\}_{j=1}^r), \tilde{\text{md}})$ in $\mathcal{L}_{\mathcal{H}}$.
6. Outputs the signature

$$\sigma = \left(\text{md}, \left\{ \text{trans}_{\text{red}}(j), \text{ch}_1^{(I_j, j)}, \text{resp}_1^{(I_j, j)}, \text{resp}_2^{(I_j, j, B_j)}, \text{cr}_2^{(I_j, j, \neg B_j)} \right\}_{j=1}^r \right),$$

$$\text{where } \text{trans}_{\text{red}}(j) := \text{com}^{(j)}, \left\{ \text{ch}_1^{(i, j)}, \text{cr}_1^{(i, j)}, (\text{cr}_2^{(i, j, 0)}, \text{cr}_2^{(i, j, 1)}) \right\}_{i \neq I_j, i=1}^t.$$

Finally, \mathcal{M}^A outputs whatever \mathcal{A} outputs.

Now, \mathcal{M}^A must succeed with probability at most differing from \mathcal{A} 's success probability by a negligible additive term as long as it does not abort (to be more precise, the term is $rq_{\text{Sign}}\text{Succ}_{\text{IDS}(1^k)}^{pq-hvzk}(\mathcal{A})$). This is the case because otherwise \mathcal{A} could be used to break the PQ-HVZK property of the used IDS. All RO queries follow the correct distribution and so do the signatures. An abort only occurs if \mathcal{A} queried \mathcal{H} before on the value for which \mathcal{M}^A wants to program. The value has the form $(\text{pk}, M, \{\text{trans}_{\text{full}}(j)\}_{j=1}^r)$ with $\text{trans}_{\text{full}}(j) := \text{com}^{(j)}, \left\{ \text{ch}_1^{(i,j)}, \text{cr}_1^{(i,j)}, (\text{cr}_2^{(i,j,0)}, \text{cr}_2^{(i,j,1)}) \right\}_{i=1}^t$. The $\{\text{trans}_{\text{full}}(j)\}_{j=1}^r$ term has at least rk bits of entropy as the commitments have at least k bits of entropy according to the definition of $q2$ -IDS and there is one commitment for each of the r rounds. This is merely a very loose (but more than sufficient) lower bound on the entropy as the blinded responses also add additional entropy. Hence, if \mathcal{A} makes a total of $q_{\mathcal{H}}$ queries for \mathcal{H} and q_{Sign} signature queries, an abort occurs with probability $\Pr[\text{abort}] \leq q_{\text{Sign}}q_{\mathcal{H}}2^{-rk}$.

Hence, \mathcal{M}^A succeeds with probability $\epsilon' \geq \epsilon(1 - q_{\text{Sign}}q_{\mathcal{H}}2^{-rk})$. \square

PQ-KOW \Rightarrow PQ-EU-CMA. Combining the two previous lemmas we obtain the following theorem.

Theorem 3.3. *Let $k, t, r \in \mathbb{N}$ be the parameters of the signature scheme from Figures 1 and 2 above using a $q2$ -IDS IDS that is PQ-HVZK and has a PQ- $q2$ -extractor. Let \mathcal{A} be a PQ-EU-CMA forger that succeeds with probability ϵ . Then, there exists an algorithm \mathcal{M}^A , that in the random oracle model breaks the PQ-KOW security of IDS in essentially the same running time as \mathcal{A} and with success probability*

$$\epsilon' \geq \epsilon - \epsilon q_{\text{Sign}}q_{\mathcal{H}}2^{-rk} - (q_{\mathcal{H}} + 1)2^{-r \log \frac{2t}{t+1}}. \quad (4)$$

Proof. Suppose there exists a PQ-EU-CMA forger \mathcal{A} that succeeds with non-negligible probability ϵ . We construct a PQ-KOW adversary \mathcal{C} for the $q2$ -IDS as follows. \mathcal{C} runs $\mathcal{A}(\text{pk})$, to construct a key-only forger \mathcal{M}^A as in Lemma 3.2, that succeeds with probability (3). Now as in Lemma 3.1, \mathcal{C} can extract a valid secret key sk , in approximately the same time, and with only negligibly smaller probability (see (2)). This concatenation of the two reductions is possible as the reduction from Lemma 3.2 only manipulates random oracle \mathcal{H} , while the one from 3.1 only touches H_1, H_2 . In total the success probability of \mathcal{C} is exactly (4), and the running time of \mathcal{C} is essentially the same as that of \mathcal{A} . \square

3.3 PQ-EU-CMA security in the QROM

We now show that with only slight changes, the two lemmas above also hold in the QROM. We do this in reverse order, starting with the PQ-KOA to PQ-EU-CMA reduction as it is the easier case. As already the QROM proofs in Unruh's work which we build on are non-tight, we only give our arguments in the asymptotic regime.

PQ-KOA \Rightarrow PQ-EU-CMA. We will first revisit the reduction from PQ-KOA to PQ-EU-CMA. We show the following lemma:

Lemma 3.4. *Let $k, t, r \in \mathbb{N}$ be the parameters of the signature scheme from Figures 1 and 2 above, using a q2-IDS that is PQ-HVZK. Let \mathcal{A} be a quantum algorithm that breaks the PQ-EU-CMA security of the signature scheme with probability ϵ . Then, in the quantum-accessible random oracle model there exists a quantum algorithm $\mathcal{M}^{\mathcal{A}}$ that breaks the PQ-KOA security of the signature scheme in essentially the same running time as \mathcal{A} and with success probability*

$$\epsilon' \geq \epsilon(1 - \text{negl}(k)). \quad (5)$$

Moreover, $\mathcal{M}^{\mathcal{A}}$ only manipulates the random oracle \mathcal{H} and leaves H_1, H_2 untouched.

Proof (Sketch). The proof in the ROM above also applies in the QROM with essentially a single change. The queries to H_1 and H_2 are still just forwarded by $\mathcal{M}^{\mathcal{A}}$ without interaction. This works without any issues in the QROM given that $\mathcal{M}^{\mathcal{A}}$ is now a quantum algorithm (which is unavoidable in the QROM). The only issue is the way $\mathcal{M}^{\mathcal{A}}$ handles \mathcal{H} . It is not possible anymore for $\mathcal{M}^{\mathcal{A}}$ to learn \mathcal{A} 's queries to \mathcal{H} and thereby not possible to abort. However, we only added the abort condition above for clarity: in the classical case $\mathcal{M}^{\mathcal{A}}$ could also simply always program \mathcal{H} . Then \mathcal{A} 's success probability might change if $\mathcal{M}^{\mathcal{A}}$ programmed on an input previously queried by \mathcal{A} . However, we still obtain the same bound on the probability. In the QROM, Unruh showed in [33, Corollary 11] that this adaptive programming only negligibly changes \mathcal{A} 's success probability (the exact argument for our specific case is exactly the one made in the first game hop of the proof of Theorem 15 in [33]). From this it follows that $\mathcal{M}^{\mathcal{A}}$'s success probability still only negligibly deviates from that of \mathcal{A} . \square

PQ-KOW \Rightarrow PQ-KOA. Now we revisit the reduction from PQ-KOW to PQ-KOA in the quantum-accessible ROM. While we still do this in the asymptotic regime, we make the parts of the reduction loss explicit which depend on the parameters r, t of the scheme. This is to support parameter selection in later sections.

Lemma 3.5. *Let $k, t, r \in \mathbb{N}$ be the parameters of the signature scheme from Figures 1 and 2 above, using a q2-IDS that has a key relation R , a q2-extractor, and is PQ-KOW secure. Let \mathcal{A} be a quantum algorithm that implements a KOA forger which given only the public key pk outputs a valid message-signature pair (M, σ) with probability ϵ . Then, in the quantum-accessible random oracle model there exists a quantum algorithm $\mathcal{M}^{\mathcal{A}}$ that given oracle access to any such \mathcal{A} breaks the KOW security of IDS in essentially the same running time as the given \mathcal{A} and with success probability*

$$\epsilon' \geq \epsilon - 2(q_{\mathcal{H}} + 1)2^{-(r \log \frac{2t}{t+1})/2}. \quad (6)$$

Moreover, $\mathcal{M}^{\mathcal{A}}$ only manipulates the random oracles H_1, H_2 and leaves \mathcal{H} untouched.

Proof (Sketch). A QROM version of our proof is obtained by essentially following the proof of Lemma 17 in [33]. The changes in the proof above are as follows. First, $\mathcal{M}^{\mathcal{A}}$ cannot learn \mathcal{A} 's RO queries to H_1 and H_2 by simulating these the classical way, anymore. Instead, $\mathcal{M}^{\mathcal{A}}$ simulates these oracles using one quantum PRP (QPRP) per oracle with a random secret key per QPRP. QPRPs exist as shown in [38] and they are quantum indistinguishable from random functions. Now, $\mathcal{M}^{\mathcal{A}}$ can open the blinded responses in the signature by inverting the QPRP using the secret key. Second, the analysis of $\Pr(\mathcal{E}_A \cap \mathcal{E}_{\text{-ext}})$ changes. As we have shown, the probability of a tuple from $\mathcal{E}_{\text{-ext}}$ to lead to a valid signature is $2^{-r \log \frac{2t}{t+1}}$. We can now follow the analysis in [33] that reduces distinguishing the constant zero function from a Bernoulli distributed boolean function to finding a tuple in $\mathcal{E}_{\text{-ext}}$ that leads to a valid signature. Thereby we get the claimed bound:

$$\Pr(\mathcal{E}_A \cap \mathcal{E}_{\text{-ext}}) \leq 2(q_{\mathcal{H}} + 1)2^{-(r \log \frac{2t}{t+1})/2}. \quad \square$$

PQ-KOW \Rightarrow PQ-EU-CMA. Putting the above two lemmas together allows us to state the following theorem.

Theorem 3.6. *Let $k, t, r \in \mathbb{N}$ be the parameters of the signature scheme above using a q_2 -IDS IDS that is computational honest-verifier zero-knowledge and has a q_2 -extractor. Let \mathcal{A} be a PQ-EU-CMA forger that succeeds with probability ϵ . Then, there exists a quantum algorithm $\mathcal{M}^{\mathcal{A}}$, that in the quantum-accessible random oracle model breaks the PQ-KOW security of IDS in essentially the same running time as \mathcal{A} and with success probability*

$$\epsilon' \geq (\epsilon - 2(q_{\mathcal{H}} + 1)2^{-(r \log \frac{2t}{t+1})/2})(1 - \text{negl}(k)).$$

4 The Sakumoto-Shirai-Hiwatari 5-pass IDS scheme

In [31], Sakumoto, Shirai, and Hiwatari proposed two new identification schemes, a 3-pass and a 5-pass IDS, based on the intractability of the \mathcal{MQ} problem. Unlike previous public key schemes, their solution provably relies only on the \mathcal{MQ} problem (and the security of the commitment scheme), and not on other related problems in multivariate cryptography such as the Isomorphism of Polynomials (IP) [29], the related Extended IP [17] and IP with partial knowledge [32] problems or the MinRank problem [13,20]. Let us quickly recall the \mathcal{MQ} problem.

Definition 4.1 (\mathcal{MQ} problem (search version)). *Let $m, n, q \in \mathbb{N}$, $\mathbf{x} = (x_1, \dots, x_n)$ and let $\mathcal{MQ}(n, m, \mathbb{F}_q)$ denote the family of vectorial functions $\mathbf{F} : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^m$ of degree 2 over \mathbb{F}_q :*

$$\mathcal{MQ}(n, m, \mathbb{F}_q) = \{\mathbf{F}(\mathbf{x}) = (f_1(\mathbf{x}), \dots, f_m(\mathbf{x})) \mid f_s(\mathbf{x}) = \sum_{i,j} a_{i,j}^{(s)} x_i x_j + \sum_i b_i^{(s)} x_i \mid_{s=1}^m\}.$$

An instance $\mathcal{MQ}(\mathbf{F}, \mathbf{v})$ of the \mathcal{MQ} (search) problem is defined as:

Given $\mathbf{F} \in \mathcal{MQ}(n, m, \mathbb{F}_q)$, $\mathbf{v} \in \mathbb{F}_q^m$ find, if any, $\mathbf{s} \in \mathbb{F}_q^n$ such that $\mathbf{F}(\mathbf{s}) = \mathbf{v}$.

The decisional version of the \mathcal{MQ} problem is NP-complete [23]. It is widely believed that the \mathcal{MQ} problem is intractable even for quantum computers in the average case, i.e., that there exists no polynomial-time quantum algorithm that given $\mathbf{F} \leftarrow_R \mathcal{MQ}(n, m, \mathbb{F}_q)$ and $\mathbf{v} = \mathbf{F}(\mathbf{s})$ (for random $\mathbf{s} \leftarrow_R \mathbb{F}_q^n$) outputs a solution \mathbf{s}' to the $\mathcal{MQ}(\mathbf{F}, \mathbf{v})$ problem with non-negligible probability.

We will later also need the \mathcal{MQ} relation $R_{\mathcal{MQ}}$ which is the relation of \mathcal{MQ} instances and solutions:

Definition 4.2 (\mathcal{MQ} relation). *The \mathcal{MQ} relation is the binary relation:*
 $R_{\mathcal{MQ}(m,n,q)} \subseteq (\mathcal{MQ}(n, m, \mathbb{F}_q) \times \mathbb{F}_q^m) \times \mathbb{F}_q^n : ((\mathbf{F}, \mathbf{v}), \mathbf{s}) \in R_{\mathcal{MQ}(m,n,q)}$ iff $\mathbf{F}(\mathbf{s}) = \mathbf{v}$.

We will omit m, n, q whenever they are clear from the context.

In [31], Sakumoto, Shirai, and Hiwatari propose a clever splitting technique, using the so-called polar form of the function \mathbf{F} which is the function $\mathbf{G}(\mathbf{x}, \mathbf{y}) = \mathbf{F}(\mathbf{x} + \mathbf{y}) - \mathbf{F}(\mathbf{x}) - \mathbf{F}(\mathbf{y})$. Using the polar form and its bilinearity, it becomes possible to split a secret into two shares, such that none of the shares on its own leaks anything about the secret. From this result, they showed how to construct zero knowledge arguments of knowledge for the \mathcal{MQ} problem, using a statistically hiding and computationally binding commitment scheme. They present a 3- and a 5-pass protocol with differing performance properties. Later, in [11], the security properties of the 5-pass scheme were reexamined to provide the minimal requirements for Fiat-Shamir type signatures from 5-pass IDS. For completeness and better readability we provide the description of the 5-pass IDS, together with the properties that we will use.

Let $(\mathbf{pk}, \mathbf{sk}) = ((\mathbf{F}, \mathbf{v}), \mathbf{s}) \in R_{\mathcal{MQ}}$ be the public and private keys of the prover. Without loss of generality, let the elements from \mathbb{F}_q be $\alpha_1, \dots, \alpha_q$. The 5-pass IDS from [31] is given in Figure 3.

Theorem 4.3. *The 5-pass identification scheme from [31] (see Fig. 3)*

1. *is computationally PQ-HVZK when the commitment scheme Com is computationally hiding,*
2. *has key relation $R_{\mathcal{MQ}(m,n,q)}$,*
3. *is PQ-KOW if the \mathcal{MQ} search problem is hard on average, and*
4. *has a $q2$ -Extractor if the commitment scheme Com is computationally binding against quantum polynomial time algorithms.*

A stronger result of the first statement in the classical case was shown in [31], namely that the 5-pass IDS is statistically honest-verifier zero-knowledge when the commitment scheme Com is statistically hiding. Relaxing the requirements of Com to computationally hiding, weakens the result to computationally HVZK, since now, it is possible to distinguish (albeit only with negligible probability) whether the commitment was produced in a valid run of the protocol. This easily transfers to the post-quantum setting, if Com is computationally hiding against quantum PPT algorithms.

The second statement holds by construction. The third statement follows from the second. The $q2$ -Extractor essentially follows from a proof in [11]. In [11] the existence of a $q2$ -Extractor was proven under the condition that the commitment scheme is computationally binding. The proof shows that there exists

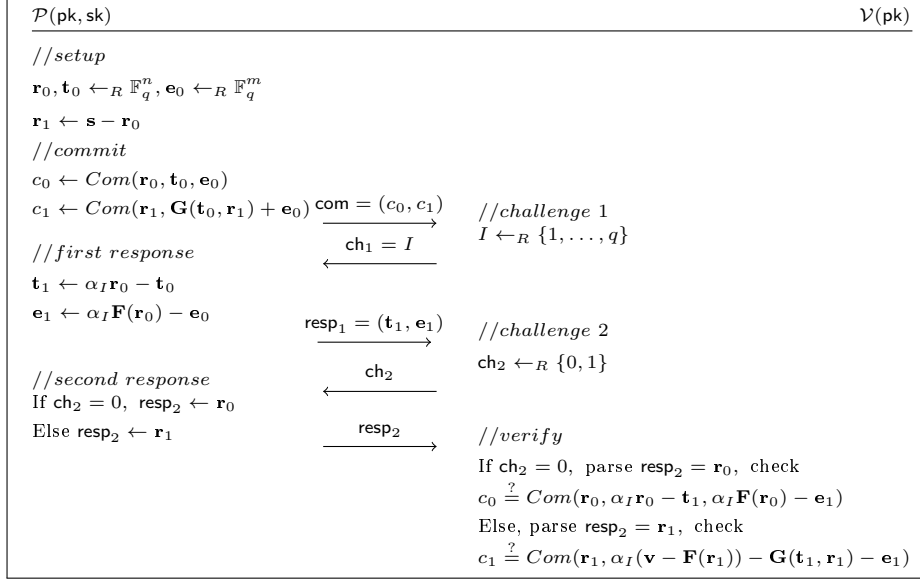


Fig. 3. The 5-pass IDS by Sakumoto, Shirai, and Hiwatari

a PPT algorithm that given four valid transcripts of the IDS with the correct pattern always either extracts a secret key or outputs two valid openings for the commitment. Hence, as long as the used commitment scheme achieves the traditional definition of computationally binding also against quantum polynomial time algorithms, the 5-pass IDS from [31] has a $q2$ -Extractor (as the probability to output two valid openings must be negligible).

5 Instantiation from the Sakumoto-Shirai-Hiwatari 5-pass IDS

In the previous sections, we have defined a signature scheme as the result of a transformed $q2$ -IDS scheme. Here, we define it instantiated with the 5-pass identification scheme proposed in [31].

5.1 SOFIA

We define the signature scheme in generic terms by describing the required parameters and the functions KGen , Sign and Vf , and defer giving concrete parameters m, n, r, t and \mathbb{F}_q for a specific security parameter k to the next section, where we also instantiate the pseudorandom generators (PRGs) and extendable output functions (XOFs). For now, we only need to fix $2 \leq t \leq q$ elements of the field \mathbb{F}_q . Without loss of generality, we denote them by $\alpha_1, \dots, \alpha_t$.

Key generation. The SOFIA key generation algorithm formally just samples a \mathcal{MQ} relation. Practically, the algorithm is realized as shown in Figure 4. The

```

KGen()
-----
sk  $\leftarrow_R \{0, 1\}^k$ 
 $S_{\mathbf{F}}, \mathbf{s}, S_{\text{rte}} \leftarrow \text{PRG}_{\text{sk}}(\text{sk}), \mathbf{F} \leftarrow \text{XOF}_{\mathbf{F}}(S_{\mathbf{F}}), \mathbf{v} \leftarrow \mathbf{F}(\mathbf{s})$ 
pk :=  $(S_{\mathbf{F}}, \mathbf{v})$ 
Return (pk, sk)

```

Fig. 4. SOFIA key generation

secret key is used as a seed to derive the following values: $S_{\mathbf{F}}$, a seed from which the system parameter \mathbf{F} is expanded; \mathbf{s} , the secret input to the \mathcal{MQ} function; S_{rte} , a seed that is used to sample all vectors $\mathbf{r}_0^{(i)}$, $\mathbf{t}_0^{(i)}$ and $\mathbf{e}_0^{(i)}$. Note that S_{rte} is not yet needed during key generation, but is required during signing.

Signature generation. For the signing procedure, we assume as input a message $M \in \{0, 1\}^*$ and a secret key sk . The signing procedure is given in Figure 5. Note that the scheme definition includes several optimizations to reduce the signature size. We discuss these later in this section.

```

Sign(sk, M)
-----
 $S_{\mathbf{F}}, \mathbf{s}, S_{\text{rte}} \leftarrow \text{PRG}_{\text{sk}}(\text{sk}), \mathbf{F} \leftarrow \text{XOF}_{\mathbf{F}}(S_{\mathbf{F}})$ 
pk :=  $(S_{\mathbf{F}}, \mathbf{F}(\mathbf{s}))$ 
 $\mathbf{r}_0^{(1)}, \dots, \mathbf{r}_0^{(r)}, \mathbf{t}_0^{(1)}, \dots, \mathbf{t}_0^{(r)}, \mathbf{e}_0^{(1)}, \dots, \mathbf{e}_0^{(r)} \leftarrow \text{PRG}_{\text{rte}}(S_{\text{rte}}, M)$ 
For  $j \in \{1, \dots, r\}$  do
   $\mathbf{r}_1^{(j)} \leftarrow \mathbf{s}^{(j)} - \mathbf{r}_0^{(j)}$ 
   $c_0^{(j)} \leftarrow \text{Com}(\mathbf{r}_0^{(j)}, \mathbf{t}_0^{(j)}, \mathbf{e}_0^{(j)}), c_1^{(j)} \leftarrow \text{Com}(\mathbf{r}_1^{(j)}, \mathbf{G}(\mathbf{t}_0^{(j)}, \mathbf{r}_1^{(j)}) + \mathbf{e}_0^{(j)})$ 
  com $^{(j)} := (c_0^{(j)}, c_1^{(j)})$ 
  For  $i \in \{1, \dots, t\}$  do
     $\mathbf{t}_1^{(i,j)} \leftarrow \alpha_i \mathbf{r}_0^{(j)} - \mathbf{t}_0^{(j)}, \mathbf{e}_1^{(i,j)} \leftarrow \alpha_i \mathbf{F}(\mathbf{r}_0^{(j)}) - \mathbf{e}_0^{(j)}$ 
    resp $_1^{(i,j)} := (\mathbf{t}_1^{(i,j)}, \mathbf{e}_1^{(i,j)}), \text{cr}_1^{(i,j)} \leftarrow \text{H}_1(\text{resp}_1^{(i,j)})$ 
  resp $_2^{(j,0)} := \mathbf{r}_0^{(j)}, \text{resp}_2^{(j,1)} := \mathbf{r}_1^{(j)}$ 
  cr $_2^{(j,0)} \leftarrow \text{H}_2(\text{resp}_2^{(j,0)}), \text{cr}_2^{(j,1)} \leftarrow \text{H}_2(\text{resp}_2^{(j,1)})$ 
  trans $_{\text{full}}(j) := (\text{com}^{(j)}, \{\text{cr}_1^{(i,j)}\}_{i=1}^t, \text{cr}_2^{(j,0)}, \text{cr}_2^{(j,1)})$ 
md  $\leftarrow \mathcal{H}(\text{pk}, M, \{\text{trans}_{\text{full}}(j)\}_{j=1}^r)$ 
 $((I_1, B_1), \dots, (I_r, B_r)) \leftarrow \text{XOF}_{\text{trans}}(\text{md})$ 
trans $_{\text{red}}(j) := (c_{-B_j}^{(j)}, \{\text{cr}_1^{(i,j)}\}_{i \neq I_j, i=1}^t, \text{cr}_2^{(j, -B_j)})$ 
Return  $(\text{md}, \{\text{trans}_{\text{red}}(j), \alpha_{I_j}, \text{resp}_1^{(I_j, j)}, \text{resp}_2^{(j, B_j)}\}_{j=1}^r)$ 

```

Fig. 5. SOFIA signature generation

The signer begins by effectively performing $\text{KGen}()$ to obtain pk and \mathbf{F} , and then iterates through r rounds of the transformed identification scheme to obtain the transcript. He then uses this as input for $\text{XOF}_{\text{trans}}$ to derive a sequence

of indices $((I_1, B_1), \dots, (I_r, B_r))$, which effectively dictate the responses that should be included unblinded in the signature.

Verification. Upon receiving a message M , a signature σ , and a public key $\text{pk} = (S_{\mathbf{F}}, \mathbf{v})$, the verifier begins by obtaining the system parameter \mathbf{F} and parsing the signature σ as defined by its construction in $\text{Sign}()$. The verification routine that follows is listed in Figure 6.

```

Vf(pk, σ, M)
-----
F ← XOFF(SF)
((I1, B1), ..., (Ir, Br)) ← XOFtrans(md)
For j ∈ {1, ..., r} do
  cr1(Ij, j) ← H1(resp1(Ij, j)), cr2(Ij, Bj) ← H2(resp2(Ij, Bj))
For j ∈ {1, ..., r} do
  If Bj = 0 then
    r0(j) := resp2(Ij, Bj)
    c0(j) ← Com(r0(j), αIjr0(j) - t1(Ij, j), αIjF(r0(j)) - e1(Ij, j))
  Else
    r1(j) := resp2(Ij, Bj)
    c1(j) ← Com(r1(j), αIj(v - F(r1(j))) - G(t1(Ij, j), r1(j)) - e1(Ij, j))
md' ← H(pk, M, {transfull(j)}j=1r)
Return md' = md

```

Fig. 6. SOFIA signature verification

Optimizations. There are several optimizations that can be applied to signatures resulting from a transformed $q2$ -IDS. Some of them are specific for SOFIA and some are more general; similar and related optimizations were suggested in [33], [11] and [10].

Excluding unnecessary blindings. The signature contains blindings of all computed responses, as well as a selection of opened responses $\text{resp}_1^{(I_j, j)}$ and $\text{resp}_2^{(j, B_j)}$. It is redundant to include the values $\text{cr}_1^{(I_j, j)}$ and $\text{cr}_2^{(j, B_j)}$, as these can be recomputed based on the opened responses. This optimization was actually proposed in the generic Unruh transform [33], and applies to any construction similar to Unruh's and ours. However, for the verifier to know which responses were actually opened, they must be able to reproduce the indices $((I_1, B_1), \dots, (I_r, B_r))$, which are derived from the transcript, and without the blinded responses, this transcript is incomplete. To solve this circular dependency, we could include the selected indices in the signature. However, for typical parameters (see Section 6.1), we can do this more efficiently by breaking $\text{XOF}_{\text{trans}}$ into two parts, composing it of a hash function over the transcript \mathcal{H} and an extendable output function XOF_{IB} to derive the indices from the hash output. We then include $\mathcal{H}(\text{pk}, M, \{\text{trans}_{\text{full}}(j)\}_{j=1}^r)$ as part of the signature, so that the verifier can reconstruct the indices, blind the corresponding responses, construct $\text{trans}_{\text{full}}$, and recompute the same hash for comparison.

Fixed challenge-space definition. Following the generic description of the signature, the selected $\alpha^{(i,j)}$ are included in the signature. Depending on the specific choice of t and q , it may be more efficient to include the challenges $\alpha^{(i,j)}$ that were *not* selected. However, there is no reason not to take this a step further and simply fix a challenge space ChS_1 of t elements. That way, all the α 's from ChS_1 will be selected and there is no need to include them in the signature. This not only reduces the signature size, but also simplifies the implementation.

Excluding unnecessary second responses. The underlying IDS from [31] has a specific property, namely that the second responses do not depend on the previous state (that is, on the first challenge and response). Therefore, regardless of the value of α , the second responses are always the same. For this reason, they need to be included only once per commitment (rather than repeating the same value t times). Combined with the previous optimization, this implies that one of the second responses will be opened, and the other will be included blinded.

Omitting commitments. The check that the verifier performs for each round consists of recomputing $c_{B_j}^{(j)}$, and comparing it to one of the commits supplied by the signer. Similar to the above, and as already suggested in [31], the signer can omit the commits that the verifier will recompute. A hash over all commits could be included instead, which the verifier can reconstruct using the commits $c_{B_j}^{(j)}$ he recomputes and the commits $c_{-B_j}^{(j)}$ the signer includes. However, it turns out that this hash is not necessary either: as these commitments are part of the transcript and the verifier is already checking the correctness of the transcript as per the first optimization, the correctness of the recomputed commitments is implicitly checked when comparing the two hashes md and md' .

No need for additional randomness in the commitments. Commitments must be randomized in order for them to be hiding. This is typically done by including a randomization string. In our case, $\mathbf{r}_0^{(j)}$, $\mathbf{t}_0^{(j)}$ and $\mathbf{e}_0^{(j)}$ are all randomly chosen, already providing sufficient randomness in both $c_0^{(j)}$ and $c_1^{(j)}$.

While constructing this scheme, we attempted several other variations. Notably, we explored opening for multiple α -challenges, but that led to no improvement in the number of rounds, and, in some cases, to a contradiction of the zero-knowledge property. Variants that employ a form of internal parallelization by committing to multiple values for \mathbf{t}_0 do reduce the number of rounds, but increase the size of the transcript disproportionately.

Altogether, the above optimizations are crucial: they add up to around 126 KiB, more than halving the signature size of the scheme that results from the transform.

5.2 Security of SOFIA

In Section 3 we described an extension of Unruh's transform to $q2$ -IDS and have proven that it provides PQ-EU-CMA security in the QROM for *any* underlying $q2$ -IDS with a $q2$ -extractor, the HVZK property, and PQ-key-one-wayness. This,

of course, implies that this transform can immediately be applied to the 5-pass \mathcal{MQ} IDS from [31], to give an \mathcal{MQ} signature secure in the QROM.

As discussed in the previous subsection, some optimizations can significantly improve the performance of the scheme. They deviate from the generic construction, however, causing a need for some changes in the security proof. Fortunately, only minor changes are required. We specify the following theorem.

Theorem 5.1. *Let $k \in \mathbb{N}$ be the security parameter. The signature scheme SOFIA is post-quantum existentially unforgeable under adaptive chosen message attacks in the quantum-accessible random oracle model if the following conditions are satisfied:*

- *The search version of the \mathcal{MQ} problem is intractable in the average case,*
- *the hash functions \mathcal{H} , H_1 , H_2 as well as the extendable output functions $XOF_{\mathbf{F}}$ and XOF_{trans} are modeled as quantum-accessible random oracles,*
- *the commitment function Com is computationally binding and computationally hiding against quantum adversaries, and has $O(k)$ bits of output entropy,*
- *the pseudorandom generators, PRG_{rte} , PRG_{sk} have outputs computationally indistinguishable from random for any polynomial-time quantum adversary.*

Proof. First let's consider a signature scheme obtained by applying the optimizations from the previous section on the signature scheme from Figures 1 and 2. We will refer to it as the optimized scheme throughout this proof. We will show that this optimized scheme is PQ-EU-CMA secure, if the underlying $q2$ -IDS satisfies the conditions from Theorem 4.3. We will assume some additional properties of the IDS, that represent a special case of $q2$ -IDS schemes. The optimized scheme is characterized by the following optimizations.

- We fix the challenge space ChS_1 to t elements. Note that this change does not influence the security arguments at all.
- We assume that the underlying IDS of the optimized scheme is such that the second response does not depend on the first challenge and response, but only on the second challenge and the initial output by the prover \mathcal{P}_0 . In this case, in the signature generation, instead of calculating the second response as $\text{resp}_2^{(i,j,\text{ch}_2)} \leftarrow \mathcal{P}_2(\text{state}^{(i,j)}, \text{ch}_2)$ for every $i \in \{1, \dots, t\}$, we calculate it once per round as $\text{resp}_2^{(j,\text{ch}_2)} \leftarrow \mathcal{P}_2(\text{state}^{(j)}, \text{ch}_2)$. The full transcript is now $\{\text{trans}_{\text{full}}(j)\}_{j=1}^r$, with $\text{trans}_{\text{full}}(j) = \text{com}^{(j)}, \left\{ \text{ch}_1^{(i,j)}, \text{cr}_1^{(i,j)} \right\}_{i=1}^t, (\text{cr}_2^{(j,0)}, \text{cr}_2^{(j,1)})$. The reduced transcript $\text{trans}_{\text{red}}(j)$ that is included in the signature is influenced similarly.
- Assuming that the underlying IDS is such that $\text{com} = (c_0, c_1)$, we omit from the signature the commitment c_{ch_2} that the verifier recomputes, depending on the challenge ch_2 . This alters the content of $\text{trans}_{\text{red}}(j)$ but not of $\text{trans}_{\text{full}}(j)$.

It is straight forward to verify that Lemma 3.1 (and in the QROM, Lemma 3.5) still hold for the optimized scheme. We only removed duplicate information from the signature, which the reduction can recompute. The exact probability of abort

in Lemma 3.2 might change as we remove some values from $\{\text{trans}_{\text{full}}(j)\}_{j=1}^r$, maybe reducing its entropy. However, the given bound does not change as it only depends on the amount of entropy coming from the commitments, which remains unchanged. Thus, the claims of Lemma 3.2 remain valid.

Next, recall (cf. Theorem 4.3) that, under the assumption of intractability of the \mathcal{MQ} problem on average, and assuming computationally binding and computationally hiding properties of Com , the 5-pass IDS from [31] is PQ-KOW, is HVZK, and has a $q2$ -Extractor. Furthermore, it satisfies the particular properties that the optimized scheme above requires. Thus applying the optimized transform on the Sakumoto-Shirai-Hiwatari 5-pass IDS scheme, we obtain a PQ-EU-CMA secure signature (cf. Theorems 3.3 and 3.6).

To complete the proof, we note that using a standard game hopping argument, it is straightforward to show that the success probability of a PQ-EU-CMA adversary against SOFIA is negligibly close to the success probability of a PQ-EU-CMA adversary against the optimized scheme from the Sakumoto-Shirai-Hiwatari 5-pass IDS scheme when the outputs of PRG_{rte} and PRG_{sk} are post-quantum computationally indistinguishable from random. \square

6 SOFIA-4-128

Having described the scheme in general terms, we now provide concrete parameters that allow us to specify a specific instance, which we will refer to as SOFIA-4-128. We present an optimized software implementation and list the results, in particular in comparison to MQDSS-31-64. All benchmarks mentioned below were obtained on a single core of an Intel Core i7-4770K (Haswell) CPU, following the standard practice of disabling TurboBoost and hyper-threading. We compiled the code using gcc 4.9.2-10, with `-O3` and `-march=native`.

6.1 Parameters

The previous section assumed a number of parameters and functions. Notably, we must define \mathbb{F}_q , the field in which we perform the arithmetic, and n and m , the number of variables and equations defining the \mathcal{MQ} problem. The number of rounds r is determined by t (i.e. the number of responses $\text{resp}_1^{(i,j)}$, bounded by q in SOFIA) and the targeted security level, using Theorem 3.6.

For MQDSS-31-64, the choice of \mathbb{F}_{31} was motivated by the fact that it brings the soundness error close to $\frac{1}{2}$ while providing convenient characteristics for fast implementation [11]. For SOFIA-4-128, our primary focus is on optimizing for signature size while still maintaining efficiency. To do so, we compute signature sizes for a wide range of candidates, and investigate several in more detail by implementing and measuring the resulting \mathcal{MQ} evaluation functions. In particular, we look at the results of $\mathcal{MQ}(128, 128, \mathbb{F}_4)$, $\mathcal{MQ}(96, 96, \mathbb{F}_7)$ and $\mathcal{MQ}(72, 72, \mathbb{F}_{16})$, and compare to $\mathcal{MQ}(64, 64, \mathbb{F}_{31})$ from [11]. Of these, $\mathcal{MQ}(128, 128, \mathbb{F}_4)$ is the decisive winner, resulting in the smallest signatures while still providing decent performance. This is also the minimum amongst all candidate systems

we looked at – it is not merely beating \mathbb{F}_7 and \mathbb{F}_{16} , but also less common options such as \mathbb{F}_5 and \mathbb{F}_8 . See Table 2 for benchmarks of single evaluation functions and the related signature sizes. Note that, as the number of rounds r does not depend on the choice of \mathbb{F}_q but merely on t , the signing time scales proportionally.

Parameters for $\mathcal{MQ}(m, n, \mathbb{F}_q)$. A straightforward method for solving systems of m quadratic equations in n variables over \mathbb{F}_q is by performing exhaustive search on all possible q^n values for the variables, and testing whether they satisfy the system. Currently, [9] provide the fastest enumeration algorithm for systems over \mathbb{F}_2 , needing $4 \log n \cdot 2^n$ operations. The techniques from [9] can be extended to other fields \mathbb{F}_q with the same expected complexity of $\Theta(\log_q n \cdot q^n)$.

In addition, there exist algebraic techniques that analyze the properties of the ideal generated by the given polynomials. The most important are the algorithms from the F4/F5 family [18,19,4,8], and the variants of the XL algorithm [12,16,36,35]. Although different in description, the two families bear many similarities, which results in similar complexity [37].

In the Boolean case, today’s state of the art algorithms BooleanSolve [5] and FXL [35], provide improvement over exhaustive search, with an asymptotic complexity of $\Theta(2^{0.792n})$ and $\Theta(2^{0.875n})$ for $m = n$, respectively. Practically, the improvement is visible for polynomials with more than 200 variables. A very recent algorithm, the Crossbred algorithm [26] over \mathbb{F}_2 , is likely to further improve the asymptotic complexity, as the authors report that it passes the exhaustive search barrier already for 37 Boolean variables. Unfortunately, at the time of writing, the preprint does not include a detailed complexity analysis that we can use (the authors of [26] confirmed that the complexity analysis is an ongoing work, and will soon be made public).

The current best known algorithms, BooleanSolve [5], FXL [35,36], the Crossbred algorithm [26] and the Hybrid approach [8] all combine algebraic techniques with exhaustive search. This immediately allows for improvement in their quantum version using Grover’s quantum search algorithm [25], provided the cost of running them on a quantum computer does not diminish the gain from Grover. Unfortunately, the current literature lacks analysis of the quantum version of these algorithms. To the best of our knowledge, a detailed analysis has only been done for pure enumeration using Grover’s search [34], showing that a system of n equations in n variables can be solved using $\Theta(n \cdot 2^{n/2})$ operations.

In what follows we will analyze the complexity of the quantum versions of the Hybrid approach and BooleanSolve, and use the results as a reference point in choosing parameters for $\mathcal{MQ}(m, n, \mathbb{F}_q)$ that provide 128 bit post-quantum security. A similar analysis can be made using the algorithms from the XL family.

First of all, we note that $m = n$ is the best choice in terms of hardness of the \mathcal{MQ} problem. Indeed, if there are more equations than variables, they provide more information about the solution, so finding one becomes easier. On the other hand, if there are more variables than equations, we can simply fix $n - m$ variables and reduce the problem to a smaller one, with m variables.

Let $\mathbf{F} = (f_1, \dots, f_m)$, $f_i \in \mathbb{F}_q[x_1, \dots, x_n]$. Without loss of generality, the equation system that we want to solve is $\mathbf{F}(\mathbf{x}) = \mathbf{0}$.

The main complexity in both the Hybrid approach and BooleanSolve comes from performing linear algebra on a Macaulay matrix $Mac_D(\mathbf{F})$ of degree D (with rows formed by the coefficients of monomials of uf_i of maximal degree D). The degree D should be big enough so that a Gröbner basis of the ideal generated by the polynomials can be obtained by performing linear algebra on the Macaulay matrix. The smallest such D is called the degree of regularity D_{reg} , and for semi-regular systems (which is a very plausible assumption for randomly generated polynomials) it is given by $D_{reg}(n, m) = 1 + deg(HS_q(t))$, where

$$HS_q(t) = \left[\frac{(1-t^2)^m}{(1-t)^n} \right]_+, \text{ for } q > 2, \text{ and } HS_2(t) = \left[\frac{(1+t)^n}{(1+t^2)^m} \right]_+,$$

and the $_+$ subscript denotes that the series has been truncated before the first non-positive coefficient. Since D_{reg} determines the size of the matrix, and thus the complexity of the linear algebra performed on it, both algorithms first fix k among the n variables in order to reduce the complexity of the costliest computational step. Now the linear algebra step is instead performed on $Mac_{D_{reg}}(\tilde{\mathbf{F}})$, where $\tilde{\mathbf{F}} = (\tilde{f}_1, \dots, \tilde{f}_m)$ and $\tilde{f}_i(x_1, \dots, x_{n-k}) = f_i(x_1, \dots, x_{n-k}, a_{n-k+1}, \dots, a_n)$, for some $(a_{n-k+1}, \dots, a_n) \in \mathbb{F}_2^k$. The value of k is chosen such that the overall complexity is minimized.

Given the linear algebra constant $2 \leq \omega \leq 3$, the complexity of the Hybrid approach for solving systems of n equations in n variables over \mathbb{F}_q is

$$C_{Hyb}(n, k) = Guess(q, k) \cdot C_{F5}(n - k, n), \quad (7)$$

where $C_{F5}(n, m) = \Theta \left(\left(m \binom{n + D_{reg}(n, m) - 1}{D_{reg}(n, m)} \right)^\omega \right)$, is the complexity of computing a Gröbner basis of a system of m equations in n variables, $m \geq n$, using the F5 algorithm [19], $Guess(q, k) = \log_q(k)q^k$ in the classical case and $Guess(q, k) = \log_q(k)q^{k/2}$ in the quantum case using Grover's algorithm. Here, we assume a rather optimistic factor of $\log_q(k)$ in the quantum case, i.e., it is the same as in the classical case, as opposed to the factor k from [34].

In the case of \mathbb{F}_2 , the BooleanSolve algorithm performs better than the Hybrid approach. It reduces the problem to testing the consistency of a related linear system

$$\mathbf{u} \cdot Mac_{D_{reg}}(\tilde{\mathbf{F}}) = (0, \dots, 0, 1) \quad (8)$$

If the system is consistent, then the original system does not have a solution. This allows for pruning of all the inconsistent branches corresponding to some $a \in \mathbb{F}_2^k$. A simple exhaustive search is then performed on the remaining branches. It can be shown that the running time of the algorithm is dominated by the first part of the algorithm in both the classical and the quantum version, although in the quantum case the difference is not as big, as a consequence of the reduced complexity of the first part. Therefore, for simplicity, we omit the exhaustive search on the remaining branches from our analysis. The complexity of the BooleanSolve algorithm is given by

$$C_{Bool}(n, k) = Guess(2, k) \cdot C_{cons}(Mac_{D_{reg}}(\tilde{\mathbf{F}})), \quad (9)$$

where $Guess(2, k)$ is defined the same as in the Hybrid approach, and

$$C_{cons}(Mac_{D_{reg}}(\tilde{\mathbf{F}})) = \Theta(N^2 \log^2 N \log \log N), \quad N = \sum_{i=0}^{D_{reg}(n-k, n)} \binom{n}{i}$$

is the complexity of testing consistency of the matrix (8), using the sparse linear algebra algorithm from [24].

Table 1 below provides estimates of the minimum requirements for 128 bit post-quantum security of $\mathcal{MQ}(n, n, \mathbb{F}_q)$ with regards to BooleanSolve and the Hybrid Approach using Grover’s search, as well as plain use of Grover’s search. In the estimates we used $\omega = 2.3$, which is smaller than the best known value $\omega = 2.3728639$ [22]. We provide the optimal number of fixed variables in brackets, where actually this number does not equal the number of variables in the initial system. When this is the case, the optimal strategy is to simply use Grover (fix all variables), which we denote with G. Note that since any system of n variables over \mathbb{F}_{2^s} can be efficiently transformed into a system of sn variables over \mathbb{F}_2 , we have scaled the results for BooleanSolve for larger \mathbb{F}_{2^s} accordingly.

	\mathbb{F}_2	\mathbb{F}_3	\mathbb{F}_4	\mathbb{F}_5	\mathbb{F}_7	\mathbb{F}_8
BooleanSolve	221 (200)	/	111	/	/	56
Hybrid	G	G	G	G	G	84 (57)
Grover	251	158	126	108	90	84
	\mathbb{F}_{11}	\mathbb{F}_{13}	\mathbb{F}_{16}	\mathbb{F}_{17}	\mathbb{F}_{31}	\mathbb{F}_{32}
BooleanSolve	/	/	28	/	/	14
Hybrid	77 (51)	73 (43)	69 (40)	69 (40)	61 (30)	60 (21)
Grover	73	68	63	62	51	51

Table 1. Lower bound on number of variables n for 128 bit post quantum security against the quantum versions of Hybrid approach [8] and BooleanSolve [5]. In brackets is the number of fixed variables. G denotes that the best strategy is to fix all variables, i.e. plain Grover search.

As mentioned earlier, a new algebraic method for equations over \mathbb{F}_2 , the Crossbred algorithm, was proposed very recently [26]. The main idea of this approach is to first perform some operations on the Macaulay matrix of degree $D_{reg}(n - k, n)$ of the given system, and fix variables only afterwards. In particular, the algorithm first tries to find enough linearly independent elements in the kernel of a submatrix of $Mac_{D_{reg}(n-k, n)}$, corresponding to monomials of specialized degree in the variables that will later remain in the system (i.e. will not be fixed). These can then be used to form new polynomials in the $n - k$ remaining variables of total small degree d , which added to $Mac_d(\tilde{\mathbf{F}})$ will result in working with a much smaller Macaulay matrix. The advantage here comes from using sparse linear algebra algorithms on $Mac_{D_{reg}(n-k, n)}$ for the first part and dense linear algebra only on the smaller Macaulay matrix in the second part. An external specialization of variables is also possible, but this does not bring any improvement classically, and we have verified for some parameters (including ours) that this is the case also quantumly. Even more, the algorithm can be split

into two distinct parts: thus, the first part, that is more memory demanding can always be performed on a classical computer, and the second part which can make use of Grover’s algorithm can be performed on a quantum computer. Since [26] does not contain a complexity analysis, we refrain from claiming exact security requirements based on the quantum version of the algorithm. Nevertheless, following the description of the algorithm we have estimated the security of our chosen instance $\mathcal{MQ}(128, 128, \mathbb{F}_4)$. We analyzed both the quantum version of the algorithm over \mathbb{F}_2 as described in [26] and the quantum version over \mathbb{F}_4 .

In both cases, as long as the number of the remaining $n - k$ variables is small, the sparse linear algebra part takes much less time, since in this case $D_{reg}(n - k, n)$ is also small. It turns out that actually it is more efficient to work with a Mac_D , with $D > D_{reg}(n - k, n)$, but not too large so that the cost of the first part becomes significant. The complexity thus, is dominated by enumeration of k variables in a system in n variables of degree D over \mathbb{F}_4 , and checking whether the obtained system has a valid solution. Clearly, a quantum version of this part using Grover can quadratically speed up the enumeration, however there will be some additional cost for the Grover oracle.

Our analysis showed that for our parameters, the version over \mathbb{F}_4 is much more efficient. Not counting the evaluation cost of the polynomials and any additional cost of the Grover oracle, the quantum algorithm against $\mathcal{MQ}(128, 128, \mathbb{F}_4)$ takes at least 2^{117} operations for the best found trade-off of parameters of the algorithm. Very likely, the additional cost we did not take into account would be much bigger than 2^{11} operations. In total, we can safely assume that a system of 128 variables over \mathbb{F}_4 provides 128-bit security against the quantum version of Crossbred algorithm. We will include a more detailed analysis for the quantum version once a classical complexity analysis of [26] is available.

Number of rounds r and blinded responses t per round. The choice of t provides a trade-off between size and speed; a larger t implies a smaller error, resulting in less rounds, but more included blinded responses per round (the additional computational cost of which is insignificant). Interestingly, $t = 3$ provides the minimal size, followed by $t = 4$, and, only then, $t = 2$. The decrease in rounds quickly diminishes, making $t = 3$ and $t = 4$ the most attractive choices. Note that t is naturally bounded by q , making these the *only* options for \mathbb{F}_4 .

	cycles ^b	size $t = 3, r = 438$	size $t = 4, r = 378$
$\mathcal{MQ}(128, 128, \mathbb{F}_4)$	21 412	123.22 KiB	129.97 KiB
$\mathcal{MQ}(96, 96, \mathbb{F}_7)$	36 501	129.00 KiB ^a	136.20 KiB ^a
$\mathcal{MQ}(72, 72, \mathbb{F}_{16})$	25 014	136.91 KiB	144.73 KiB
$\mathcal{MQ}(64, 64, \mathbb{F}_{31})$	6 616[11]	149.34 KiB ^a	158.15 KiB ^a

^a Assumes optimally packing the elements of \mathbb{F}_q , which may not be practical.

^b For single evaluation. In practice, batching provides a speedup. See Sec. 6.2.

Table 2. Benchmarks for varying parameter sets

Given the above considerations (and with a prospect of some convenience of implementation), we select the parameters $n = m = 128$, $q = 4$ and $t = 3$. For a security level of 128 bits post-quantum security, it follows from Theorem 3.6 that we must select r such that $2^{-(r \log \frac{2t}{t+1})/2} < 2^{-128}$. This implies $r = 438$.

Required functions. Before being able to implement the scheme, we must still define several of the functions we have assumed to exist. In particular, we need: a string commitment function Com ; pseudorandom generators PRG_{sk} and PRG_{rte} ; extendable output functions $XOF_{\mathbf{F}}$ and XOF_{IB} ; permutation functions H_1 and H_2 ; and a cryptographic hash function \mathcal{H} .

We instantiate the extendable output functions, the string commitment functions, the permutations and the hash function with SHAKE-128 [7]. This applies trivially, except for XOF_{IB} , of which the output domain is a series of ternary and binary indices (as $t = 3$). We resolve this by applying rejection sampling to the output of SHAKE-128 to derive the ternary challenges. Note that this does not enable a timing attack, as the input to SHAKE-128 is public. For $XOF_{\mathbf{F}}$, we achieve a significant speedup by dividing its output in four separate pieces, generating each of them with a domain-separated call to cSHAKE-128 [7]. For the application of \mathcal{H} to the public key, the message and the transcript, collision resilience is achieved by absorbing the transcript into the SHAKE-128 state first, as the included randomness prevents internal collisions.

We also instantiate PRG_{rte} and PRG_{sk} with SHAKE-128, but note that implementations can make different choices without breaking compatibility. In fact, for the optimized Haswell implementation discussed in the next section, we instantiate PRG_{rte} with AES in counter mode, using the AES-NI instruction set.

6.2 Implementation

As part of this work, we provide a C reference implementation and an implementation optimized for AVX2. The focus of this section is the evaluation of the \mathcal{MQ} function, given the abovementioned parameter set $\mathcal{MQ}(128, 128, \mathbb{F}_4)$. The rest of the scheme depends on fairly straight-forward operations (such as multiplying vectors of \mathbb{F}_4 elements by a constant scalar) and applications of existing implementations of AES-CTR and SHAKE-128. The used AES-CTR and SHAKE-128 implementations are in the public domain and run in constant time.

Before discussing the computation, we note that the chosen parameters lend themselves to a very natural data representation. Throughout the entire scheme, we interpret 256 bit vectors as vectors of 128 bitsliced \mathbb{F}_4 elements: the low 128 bits make up the lower bits of the two-bit elements, and the high 128 bits make up the higher bits of each element. This makes operations such as scalar multiplication very convenient in C code, as this can be easily expressed as logical operations on bit sequences, but provides an even more important benefit for AVX2 assembly code. Notably, one vector of \mathbb{F}_4 elements fits exactly into one 256 bit vector register, with the lower bits now fitting into the low lane and the higher bits into the high lane. Whereas other parameter sets could result in having to consider crossing the lanes, in this case the separation is very natural.

When sampling elements in \mathbb{F}_4 from the output of SHAKE-128 or AES-CTR, we can freely interpret the random data to be in bitsliced representation. Similarly, we include the elements in the signature in this representation, as signature verification enjoys precisely the same benefits. Throughout the entire scheme, there is no point at which we need to actually perform a bitslicing operation.

As a side effect of this choice of representation, it is very natural to perform the \mathcal{MQ} evaluation in constant time. While bigger underlying fields might have implied approaches based on lookup tables, for vectors over \mathbb{F}_4 it is much faster to perform the evaluation using bitsliced field arithmetic.

Evaluating \mathcal{MQ} . For a given input \mathbf{x} , we split the evaluation into two phases: computing all quadratic monomial terms $x_i x_j$, and composing them to evaluate the quadratic polynomials.

Computing the quadratic terms. To perform the first step, we use a similar approach as was used in [11]. It can be viewed as a combination of their approach for \mathbb{F}_2 and for \mathbb{F}_{31} , as we now operate on a single register that contains all input elements, but view each lane as 16 separate single-byte registers. Using `vpshufb` instructions, the elements can be easily arranged such that all multiplications can be performed using only a minimal number of rotations. We used the script from [11] as a starting point to generate the arrangement.

A bitsliced multiplication in \mathbb{F}_4 can be efficiently performed using only a few logical operations. The inputs to these multiplications are a register containing \mathbf{x} and a register containing some rotated arrangement of \mathbf{x} . However, some of these operations require the low and high lanes of the vector registers to interact, which is typically costly. As \mathbf{x} is constant, we speed up these multiplications by rewriting them as shown below, and presetting two registers that contain $[\mathbf{x}_{high}|\mathbf{x}_{high}]$ and $[\mathbf{x}_{high} \oplus \mathbf{x}_{low}|\mathbf{x}_{low}]$, respectively. Note that all of these operations are not performed on single bits, but rather on 128 bit vector lanes. The multiplication of 128 elements then requires only two `vpand` instructions, one `vperm` instruction, and a `vpxor` to combine the results.

$$\begin{aligned} c_{high} &= (a_{high} \wedge (b_{low} \oplus b_{high})) \oplus (a_{low} \wedge b_{high}) \\ c_{low} &= (a_{low} \wedge b_{low}) \oplus (a_{high} \wedge b_{high}) \end{aligned}$$

Multiplying, and accumulating results. We focus on two approaches to perform the second and most costly part of the evaluation, in which all of the above monomials need to be multiplied with coefficients from \mathbf{F} and summed into the output vector. They are best described as iterating either ‘horizontally’ or ‘vertically’ through the required multiplications. For the vertical approach, we iterate over all⁵ registers of monomials, broadcasting each of the monomials to

⁵ There are $\frac{n \cdot (n+1)}{2} = 8256$ such monomials, which results in $64 \frac{1}{2}$ 256-bit sequences. We round up to 65 by zeroing out half of the high and half of the low lane. To still get results that are compatible with implementations on other platforms, we create similar gaps in the stream of random values used to construct \mathbf{F} , ensuring that the same random elements are still used for the same coefficients.

each of the 128 possible positions (using rotations), before multiplying with a sequence of coefficients from \mathbf{F} and adding into an accumulator. Alternatively, we iterate over the output elements in the outer-most loop. For each output element, we iterate over all registers of monomials, perform the multiplications and horizontally sum the results by making use of the `popcnt` instruction.

Intuitively, the latter approach may seem like more work (notably because it requires more loads from memory), but in practice it turns out to be faster for our parameters. The main reason for this is that by maintaining multiple separate accumulators, loaded monomials can be re-used while still maintaining chains of logic operations that operate on independent results (as the accumulators are only joined together later), which leads to highly efficient scheduling.

For both cases, delaying part of the multiplication in \mathbb{F}_4 provides a significant speedup. This is done by computing both $[\hat{\mathbf{x}}_{high} \wedge \mathbf{f}_{high} | \hat{\mathbf{x}}_{low} \wedge \mathbf{f}_{low}]$ and $[\hat{\mathbf{x}}_{low} \wedge \mathbf{f}_{high} | \hat{\mathbf{x}}_{high} \wedge \mathbf{f}_{low}]$, with \mathbf{f} from \mathbf{F} and $\hat{\mathbf{x}}$ a sequence of quadratic monomials, and accumulating these results separately. After accumulating, all multiplications and reductions can be completed at once, eliminating the duplicate operations that would otherwise be performed for each of the 65 multiplications.

Evaluating \mathcal{MQ} instances in parallel. As each of the coefficients in \mathbf{F} is used only once, loading these elements from memory causes a considerable burden. Since \mathbf{F} is constant for each evaluation, however, a significant speedup can be achieved by processing multiple instances of the \mathcal{MQ} function in parallel. This applies in particular to the vertical approach, as its critical section leaves some registers unused. Horizontally, there is a trade-off with registers used for parallel accumulators, but there is still considerable gain from parallelizing evaluations.

For SOFIA-4-128, the signer evaluates $r = 438$ instances of \mathbf{F} and its polar form \mathbf{G} on completely independent inputs, which can be trivially batched.

Parallel SHAKE-128 and cSHAKE-128. As will be apparent in the next section, many cycles are spent computing the Keccak permutation (as part of either SHAKE-128 or cSHAKE-128). Some of the main culprits are the commitments, the blinding of responses and the expansion of \mathbf{F} . While the Keccak permutation does not provide internal parallelism, it is straightforward to compute four instances in parallel in a 256 bit vector register. This allows us to seriously speed up the many commitments and blindings, as these are all fully independent and can be grouped together across rounds. Deriving \mathbf{F} can be parallelized by splitting it in four domain-separated cSHAKE-128 calls operating on the same seed, as was alluded to in Section 6.1.

Benchmarks. Evaluating the \mathcal{MQ} function horizontally in batches of three turns out to give the fastest results, measuring in at 17 558 cycles per evaluation. Evaluating vertically costs 18 598 cycles. The cost for evaluating the polar form is not significantly different, differing by approximately a hundred cycles from regular \mathcal{MQ} . Generating the monomial terms $x_i y_j + x_j y_i$ is somewhat more costly, but this is countered by the fact that the linear terms cancel out.

To generate a signature, we spend 21 305 472 cycles. Of this, 15 420 520 cycles can be attributed to evaluating \mathcal{MQ} , and 43 954 to AES-CTR. The remainder is

almost entirely accounted for by the various calls to SHAKE-128 and cSHAKE-128 for the commitments, blindings and randomness expansion. In particular, expanding \mathbf{F} costs 1 120 782 cycles. Note, however, that if many signatures are to be generated, this expansion only needs to be done once and \mathbf{F} can be kept in memory across subsequent signatures. Verification costs 15 492 686 cycles, and key generation costs 1 157 112; key generation is dominated by expansion of \mathbf{F} .

The keys of SOFIA-4-128 are very small by nature, with the secret key consisting of only a single 32 byte seed, and the 64 byte public key being made up of a seed and a single \mathcal{MQ} output.

The natural candidate for comparison is MQDSS-31-64 [11]. While MQDSS has a proof in the ROM, we focus further comparison on post-quantum schemes that have proofs in the QROM or standard model. See Table 3, below; as mentioned in the introduction, we include SPHINCS-256 [6] (standard model), Picnic-10-38 [10] (QROM) and TESLA-2 [2] (QROM). Since [2] does not implement the TESLA-2 parameter set, we include TESLA-1 (ROM) for context.

	$ \sigma $ (bytes)	$ \text{pk} , \text{sk} $ (bytes)	keygen (cycles)	signing (cycles)	verification (cycles)
SOFIA-4-128 ^a	126 176	64 32	1 157 112	21 305 472	15 492 686
MQDSS-31-64 ^a	40 952	72 64	1 826 612	8 510 616	5 752 612
SPHINCS-256 ^b	41 000	1056 1088	3 237 260	51 636 372	1 451 004
Picnic-10-38 ^{c,d}	195 458	64 32	$\approx 36\,000$	$\approx 112\,716k$	$\approx 58\,680\,000$
TESLA-1 ^a	2 444	11 653k 6 769k	? ^e	143 402 231	19 284 672
TESLA-2 ^f	$\geq 4.0k^g$	$\geq 21\,799k^g \geq 7\,700k^g$? ^f	? ^f	? ^f

^a Benchmarked on an Intel Core-i7-4770K (Haswell). ^b Benchmarked on an Intel Xeon E3-1275 (Haswell). ^c Benchmarked on an Intel Core-i7-4790 (Haswell). ^d Converted from milliseconds at 3.6GHz. ^e The benchmarks in [2] omit key generation. In [10], a measurement of approximately 173 billion cycles is reported for the preceding TESLA-768 [1] scheme, which uses a similar key generation operation but is instantiated with smaller parameters. ^f The TESLA-2 parameter set is not implemented in [2]; no benchmarks are available. ^g “Sizes are theoretic sizes for fully compressed keys and signatures” [2].

Table 3. Benchmark overview

References

1. E. Alkim, N. Bindel, J. Buchmann, and O. Dagdelen. TESLA: tightly-secure efficient signatures from standard lattices. Cryptology ePrint Archive, Report 2015/755, 2015.
2. E. Alkim, N. Bindel, J. Buchmann, Ö. Dagdelen, E. Eaton, G. Gutoski, J. Krämer, and F. Pawlega. Revisiting TESLA in the quantum random oracle model. *International Workshop on Post-Quantum Cryptography – PQCRYPTO 2017*, vol. 10346 of LNCS, 143–162. Springer, 2017.

3. A. Ambainis, A. Rosmanis, and D. Unruh. Quantum attacks on classical proof systems: The hardness of quantum rewinding. *FOCS 2014*, 474–483. 2014.
4. M. Bardet, J. Faugère, and B. Salvy. On the complexity of the F5 Gröbner basis algorithm. *Journal of Symbolic Computation*, 70:49–70, 2015.
5. M. Bardet, J. Faugère, B. Salvy, and P. Spaenlehauer. On the complexity of solving quadratic boolean systems. *Journal of Complexity*, 29(1):53–75, 2013.
6. D. J. Bernstein, D. Hopwood, A. Hülsing, T. Lange, R. Niederhagen, L. Papachristodoulou, M. Schneider, P. Schwabe, and Z. Wilcox-O’Hearn. SPHINCS: practical stateless hash-based signatures. *Advances in Cryptology – EUROCRYPT 2015*, vol. 9056 of *LNCS*, 368–397. Springer, 2015.
7. G. Bertoni, J. Daemen, M. Peeters, and G. Van Assche. The KECCAK reference, 2011.
8. L. Bettale, J. Faugère, and L. Perret. Solving polynomial systems over finite fields: improved analysis of the hybrid approach. *Proceedings of the 37th International Symposium on Symbolic and Algebraic Computation – ISSAC ’12*, 67–74. ACM, 2012.
9. C. Bouillaguet, H.-C. Chen, C.-M. Cheng, T. Chou, R. Niederhagen, A. Shamir, and B.-Y. Yang. Fast exhaustive search for polynomial systems in \mathbb{F}_2 . *Cryptographic Hardware and Embedded Systems – CHES 2010*, vol. 6225 of *LNCS*, 203–218. Springer, 2010.
10. M. Chase, D. Derler, S. Goldfeder, C. Orlandi, S. Ramacher, C. Rechberger, D. Slamanig, and G. Zaverucha. Post-quantum zero-knowledge and signatures from symmetric-key primitives. *Cryptology ePrint Archive*, Report 2017/279, 2017.
11. M.-S. Chen, A. Hülsing, J. Rijneveld, S. Samardjiska, and P. Schwabe. From 5-pass \mathcal{MQ} -based identification to \mathcal{MQ} -based signatures. *Advances in Cryptology – ASIACRYPT 2016*, vol. 10032 of *LNCS*, 135–165. Springer, 2016.
12. N. Courtois, E. Klimov, J. Patarin, and A. Shamir. Efficient algorithms for solving overdefined systems of multivariate polynomial equations. *Advances in Cryptology – EUROCRYPT 2000*, vol. 1807 of *LNCS*, 392–407. Springer, 2000.
13. N. T. Courtois. Efficient zero-knowledge authentication based on a linear algebra problem MinRank. *Advances in Cryptology – ASIACRYPT 2001*, vol. 2248 of *LNCS*, 402–421. Springer, 2001.
14. Ö. Dagdelen, M. Fischlin, and T. Gagliardoni. The Fiat–Shamir transformation in a quantum world. *Advances in Cryptology – ASIACRYPT 2013*, vol. 8270 of *LNCS*, 62–81. Springer, 2013.
15. Ö. Dagdelen, D. Galindo, P. Véron, S. M. El Yousfi Alaoui, and P.-L. Cayrel. Extended security arguments for signature schemes. *Designs, Codes and Cryptography*, 78(2):441–461, 2016.
16. C. Diem. The XL-algorithm and a conjecture from commutative algebra. *Advances in Cryptology – ASIACRYPT 2004*, vol. 3329 of *LNCS*, 323–337. Springer, 2004.
17. J. Ding, L. Hu, B.-Y. Yang, and J.-M. Chen. Note on design criteria for rainbow-type multivariates. *Cryptology ePrint Archive*, Report 2006/307, 2006.
18. J.-C. Faugère. A new efficient algorithm for computing Gröbner bases (F4). *Journal of Pure and Applied Algebra*, 139:61–88, 1999.
19. J.-C. Faugère. A new efficient algorithm for computing Gröbner bases without reduction to zero (F5). *Proceedings of the 2002 International Symposium on Symbolic and Algebraic Computation – ISSAC ’02*, 75–83. ACM, 2002.
20. J.-C. Faugère, F. Levy-dit-Vehel, and L. Perret. Cryptanalysis of MinRank. *Advances in Cryptology – CRYPTO 2008*, vol. 5157 of *LNCS*, 280–296. Springer, 2008.

21. M. Fischlin. Communication-efficient non-interactive proofs of knowledge with online extractors. *Advances in Cryptology – CRYPTO 2005*, vol. 3621 of *LNCS*, 152–168. Springer, 2005.
22. F. L. Gall. Powers of tensors and fast matrix multiplication. *Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation – ISSAC '14*, 296–303. ACM, 2014.
23. M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979.
24. M. Giesbrecht, A. Lobo, and B. D. Saunders. Certifying inconsistency of sparse linear systems. *Proceedings of the 1998 International Symposium on Symbolic and Algebraic Computation – ISSAC '98*, 113–119. 1998.
25. L. K. Grover. A fast quantum mechanical algorithm for database search. *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing – STOC '96*, 212–219. ACM, 1996.
26. A. Joux and V. Vitse. A crossbred algorithm for solving boolean polynomial systems. Cryptology ePrint Archive, Report 2017/372, 2017.
27. E. Kiltz, J. Loss, and J. Pan. Tightly-secure signatures from five-move identification protocols. *Advances in Cryptology – ASIACRYPT 2017: 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part III*, 68–94. Springer International Publishing, Cham, 2017.
28. E. Kiltz, V. Lyubashevsky, and C. Schaffner. A concrete treatment of fiat-shamir signatures in the quantum random-oracle model. Cryptology ePrint Archive, Report 2017/916, 2017.
29. J. Patarin. Hidden field equations (HFE) and isomorphisms of polynomials (IP): Two new families of asymmetric algorithms. *Advances in Cryptology – EUROCRYPT '96*, vol. 1070 of *LNCS*, 33–48. Springer, 1996.
30. D. Pointcheval and J. Stern. Security proofs for signature schemes. *Advances in Cryptology – EUROCRYPT '96*, vol. 1070 of *LNCS*, 387–398. Springer, 1996.
31. K. Sakumoto, T. Shirai, and H. Hiwatari. Public-key identification schemes based on multivariate quadratic polynomials. *Advances in Cryptology – CRYPTO 2011*, vol. 6841 of *LNCS*, 706–723. Springer, 2011.
32. E. Thomae. *About the Security of Multivariate Quadratic Public Key Schemes*. Ph.D. thesis, Ruhr-University Bochum, Germany, 2013.
33. D. Unruh. Non-interactive zero-knowledge proofs in the quantum random oracle model. *Advances in Cryptology – EUROCRYPT 2015*, vol. 9056 of *LNCS*, 755–784. Springer, 2015.
34. B. Westerbaan and P. Schwabe. Solving binary \mathcal{MQ} with grover's algorithm. *Security, Privacy, and Advanced Cryptography Engineering*, vol. 10076 of *LNCS*. Springer, 2016.
35. B. Yang and J. Chen. Theoretical analysis of XL over small fields. *Information Security and Privacy*, vol. 3108 of *LNCS*, 277–288. Springer, 2004.
36. B.-Y. Yang and J.-M. Chen. All in the XL family: Theory and practice. *Information Security and Cryptology – ICISC 2004*, 67–86. Springer, 2005.
37. J. Y.-C. Yeh, C.-M. Cheng, and B.-Y. Yang. Operating Degrees for XL vs. F4/F5 for Generic \mathcal{MQ} with Number of Equations Linear in That of Variables. *Number Theory and Cryptography: Papers in Honor of Johannes Buchmann on the Occasion of His 60th Birthday*, 19–33. Springer, 2013.
38. M. Zhandry. A note on quantum-secure PRPs. Cryptology ePrint Archive, Report 2016/1076, 2016.