Sharing the LUOV: Threshold Post-Quantum Signatures

Daniele Cozzo[1][0000−0001−5289−3769] and Nigel P. Smart[1,2][0000−0003−3567−3304]

[1] imec-COSIC, KU Leuven, Leuven, Belgium.
[2] University of Bristol, Bristol, UK.
`daniele.cozzo@kuleuven.be, nigel.smart@kuleuven.be`

Abstract. We examine all of the signature submissions to Round-2 of the NIST PQC "competition" in the context of whether one can transform them into threshold signature schemes in a relatively straight forward manner. We conclude that all schemes, except the ones in the MQ family, have significant issues when one wishes to convert them using relatively generic MPC techniques. The lattice based schemes are hampered by requiring a mix of operations which are suited to both LSSS- and GC-based MPC techniques (thus requiring costly transfers between the two paradigms). The Picnic and SPHINCS+ algorithms are hampered by the need to compute a large number of hash function queries on secret data. Of the nine submissions the two which would appear to be most suitable for using in a threshold like manner are Rainbow and LUOV, with LUOV requiring less rounds and less data storage.

1 Introduction

Ever since the late 1980's there has been interest in threshold cryptography [13]. The interest in threshold signatures being a key concern, as threshold signatures allow one to distribute the signing power to different authorities (using different access structures) which is important in many applications. For example in the 1990's and early 2000's there was work on threshold RSA signatures, see for example [11, 47] amongst other work, and DSA/EC-DSA, again see [16, 36] amongst other works.

The case of distributed signatures for EC-DSA has recently had renewed interest, [15, 31–33], due to applications in block-chain. In addition, general distributed cryptographic solutions for decryption and signature operations are becoming more in vogue, see a the recent NIST workshop in this space[3].

However, these solutions for RSA and EC-DSA do not obviously apply if (or when) a quantum computer is built. Thus if one is to have applications that require threshold signatures in a post-quantum world then one needs to examine how to "thresholdize" post-quantum signatures. However, the tricks to create threshold versions of RSA and EC-DSA make strong use of the number-theoretic structure of such signature schemes; this structure is not available for many of the proposed post-quantum signature algorithms.

The NIST post-quantum cryptography "competition" aims to find replacement public key encryption and signature algorithms for the current number-theoretic standard solutions based on integer factoring and discrete logarithms. There are nine solutions which have made it to the second round of this "competition", and these nine can be divided into four classes, depending on the underlying hard problem on which they are based

- Lattice Based: There are three submissions in this category; Dilithium [35], qTesla [5], and Falconal [42]
- Hash Based: There is one submission in this category, SPHINCS+ [24].
- MPC-in-the-Head (MPC-in-H) Based: Here there is also one submission Picnic [50].
- Multivariate Quadratic Based: Here we have four submissions GeMSS [6], LUOV [4], MQDSS [46] and Rainbow [14].

Generic MPC techniques are so well developed that one can simply apply the MPC sledge hammer to the post-quantum signature nut. But some proposed post-quantum signatures will be more suited to this approach than others. Thus in this work we examine the proposed post-quantum signature schemes submitted to Round-2 of the NIST "competition" in the context of this problem.

---

[3] `https://www.nist.gov/news-events/events/2019/03/nist-threshold-cryptography-workshop-2019`.

Our Contribution: Simply from the underlying assumptions one would suspect Picnic would be the algorithm which lends itself to being converted into an MPC threshold version; after all it is based on MPC-in-the-Head. However, closer examination reveals that this is not the case. Indeed we examine all the post-quantum signature submissions from the point of view of whether one can easily turn them into threshold versions. It turns out that the ones which are most amenable to "thresholdizing" are those based on the MQ family of problems, in particular Rainbow and LUOV, see Table 1 for a summary.

Name	Underlying Assumption	Issues in Obtaining a Threshold Variant
Dilithium	Lattice	A mix of linear operations (suitable for LSSS-based MPC) and non-linear operations (suitable for GC-based MPC) requires costly transferring between the two representations. We expect this to take around 12s to execute.
qTesla	Lattice	A mix of linear operations (suitable for LSSS-based MPC) and non-linear operations (suitable for GC-based MPC) requires costly transferring between the two representations. We expect to take at least 16s to execute.
Falcon	Lattice	A mix of linear operations (suitable for LSSS-based MPC) and non-linear operations (suitable for GC-based MPC) requires costly transferring between the two representations. We expect to take at least 6s to execute.
Picnic	MPC-in-H	Applying SHA-3 to obtain the necessary randomness in the views of the MPC parties.
SPHINCS+	Hash	Applying SHA-3 to obtain the data structures needed.
MQDSS	MQ	Applying SHA-3 to obtain the commitments.
GeMSS	MQ	Potential for threshold implementation, implementation is tricky due to need to extract polynomial roots via Berlekamp algorithm
Rainbow	MQ	Simple LSSS based MPC solution which requires 12 rounds of communication. We expect a signature can be generated in around three seconds
LUOV	MQ	Simple LSSS based MPC solution which requires 6 rounds of communication. We expect a signature can be generated in just over a second

Table 1. Summary of NIST Round-2 Post-Quantum Signature Schemes

The main problems with the lattice based techniques are the need to perform rejection sampling, which means intermediate values need to be kept secret until after the rejection sampling has been accomplished, and they need to be compared to given constants. This results in a number of operations suitable for garbled circuit operation to be performed. However, the rest of the algorithms require operations which are linear. Thus one has both a large number of GC-based operations to perform, as well as conversion to-and-from LSSS based representations so as to help mitigate the number of GC operations needed. This conversion turns out to be a major bottleneck.

Picnic on the other hand requires the signer to privately evaluate a set of PRFs and then reveal the associated keys for a given subset of the PRFs on obtaining the challenge value. This means that the PRFs need to be securely evaluated in a threshold manner. Since the PRFs used in Picnic are not specifically designed to be evaluated in this manner one is left to applying generic MPC techniques which become very expensive due to the gate counts of the underlying PRFs.

The hash-based signature scheme SPHINCS+ has a similar issue in that one needs to securely evaluate the underlying hash functions in a threshold manner. This leads, again to huge gate counts, when one tries to perform this in a threshold manner.

One of the MQ schemes (MQDSS) requires one to also evaluate hash functions on secret data, and so suffers from the same problems as the previous schemes. One MQ scheme (GeMSS) is a plausible candidate to be implemented via MPC, but any implementation will be highly non-trivial due to the need to evaluate Berlekamps' algorithm to extract roots of a univariate polynomial.

This leaves us with the two remaining MQ schemes (LUOV and Rainbow). These are based on the FDH signature construction and hence the main issue is implementing generic arithmetic circuit based MPC over the given finite fields will lead to threshold variants relatively easily. In this case Rainbow requires more rounds of interaction than LUOV, on the other hand Rainbow requires less secure multiplications. In addition, LUOV requires us to store in the shared secret key state less data.

In all of our analysis we try to give a best guess as to the *minimum* amount of time a threshold implementation would take for each of the candidates. This is assuming the current best run-times for evaluating the SHA-3 internal function in an MPC system etc. These estimates are given for the schemes at the security level denoted Level-3 by NIST; when a scheme does not have parameters at Level-3 we pick the set at Level-4. Level-3 corresponds to the difficulty of breaking AES-192 on a quantum computer. This provides less than 192-bits of security quantumly (due to Grover's algorithm), and hence seems a reasonable compromise given current (classical) security levels are usually picked to be equivalent to AES-128.

2   Preliminaries

In this section we define various notations and notions which will be needed in future sections. In particular we describe the underlying MPC systems which we will assume 'as given'. In particular our focus will be on MPC solutions which are actively secure (with abort) against static corruptions.

We assume that all involved parties are probabilistic polynomial time Turing machines. Given a positive integer n, we denote by $[n]$ the set $\{1, \ldots, n\}$. We let $x \leftarrow X$ denote the uniformly random assignment to the variable x from the set X, assuming a uniform distribution over X. We also write $x \leftarrow y$ as shorthand for $x \leftarrow \{y\}$. If $\mathcal{D}$ is a probability distribution over a set X, then we let $x \leftarrow \mathcal{D}$ denote sampling from X with respect to the distribution $\mathcal{D}$. If A is a (probabilistic) algorithm then we denote by $a \leftarrow A$ the assignment of the output of A where the probability distribution is over the random tape of A.

Signature schemes: Digital signature schemes which are defined by

Definition 2.1. *A* digital signature scheme *is given by a tuple of probabilistic algorithms* (KeyGen, Sign, Verify)*:*

- KeyGen (1^λ) *is a randomized algorithm that takes as input the security parameter and returns the public key* pk *and the private key* sk.
- Sign (sk, μ) *is a randomized signing algorithm that takes as inputs the private key and a message and returns a signature on the message.*
- Verify $(\text{pk}, (\sigma, \mu))$ *is a deterministic verification algorithm that takes as inputs the public key and a signature σ on a message μ and outputs a bit which is equal to one if and only if the signature on μ is valid.*

Correctness and security (EU-CMA) are defined in the usual manner, and all signature scheme submitted to NIST in Round-2 meet this security definition.

Multi-Party Computation: As mentioned above we consider actively secure (with abort) MPC for static adversaries in this work. We assume a generic black box for MPC, abstracted in the functionality $\mathcal{F}_{\mathsf{MPC}}$ of Figure 1, which defines MPC over a given finite field $\mathbb{K}$ (or indeed sometimes a finite ring). When instantiating this abstract MPC functionality with state-of-the-art protocols one needs to consider aspects such as the access structure, the field used $\mathbb{K}$, and the computational/communication model. We summarize many of the state-of-the-art of MPC protocols in Table 2.

In terms of access structures the main ones in use are full threshold (for example in the SPDZ protocol family [12]) and Q2-access structures (which includes standard threshold protocols for which $t < n/2$). A Q2-access structure is one in which the union of no two unqualified sets cover the entire set of players. In terms of the field $\mathbb{K}$ for evaluating binary circuits one usually utilizes MPC over $\mathbb{K} = \mathbb{F}_2$. However, for some applications (in particular the MQ signature schemes considered later) it is better to take $\mathbb{K}$ to be a specific

---

The ideal functionality $\mathcal{F}_{\mathsf{MPC}}$ for MPC over $\mathbb{F}_q$

Initialize: On input $(init, \mathbb{K})$ from all parties, the functionality stores $(domain, \mathbb{K})$.

Input: On input $(input, P_i, varid, x)$ from P_i and $(input, P_i, varid, ?)$ from all other parties, with $varid$ a fresh identifier, the functionality stores $(varid, x)$.

Random: On input of $(random, varid)$, if $varid$ is not present in memory then the functionality picks a random value in $\mathbb{K}$ and stores it in $varid$.

Add: On command $(add, varid_1, varid_2, varid_3)$ from all parties, if $varid_1$ and $varid_2$ are present in memory and $varid_3$ is not then the functionality retrieves $(varid_1, x)$ and $(varid_2, y)$ and stores $(varid_3, x + y)$. Otherwise does nothing.

Multiply: On input $(multiply, varid_1, varid_2, varid_3)$ from all parties, if $varid_1$ and $varid_2$ are present in memory and $varid_3$ is not then retrieve $(varid_1, x)$ and $(varid_2, y)$ and store $(varid_3, x \cdot y)$. Otherwise do nothing.

Output: On input $(output, varid, i)$ from all honest parties, if $varid$ is present in memory then retrieve $(varid, y)$ and output it to the environment. It then waits for an input from the environment. If this input is Deliver then y is output to all players if $i = 0$, or y is output to party i if $i \neq 0$. If the input is not equal to Deliver then $\bot$ is output to all parties.

Figure 1. The ideal functionality $\mathcal{F}_{\mathsf{MPC}}$ for MPC over $\mathbb{F}_q$

finite field taylored to the application. Some protocols are tailored to very specific access structures (for example using threshold $(t, n) = (1, 3)$).

The functionality has a command to produce random values in $\mathbb{K}$. This can always be achieved using interaction (via the input command), however, for LSSS based protocols in the Q2 setting (with small numbers of parties) such a command can be executed for free using a PRSS.

To make it simpler to describe MPC protocols, in what follows we use the notation $\langle x \rangle$ for $x \in \mathbb{K}$ to denote a value x stored in the MPC engine (the reader can think of this as $\langle x \rangle$ being the secret sharing of x). The MPC functionality Figure 1 enables one to compute $\langle x \rangle + \langle y \rangle$, $\lambda \cdot \langle x \rangle$ and $\langle x \rangle \cdot \langle y \rangle$. We extend this notation to vectors and matrices of elements in $\mathbb{K}$ in the obvious manner.

In terms of computational model we find the set of practical MPC protocols divided into distinct classes. In some protocols there is a function-independent offline phase, and then a fast offline phase. Other protocols have no offline phase but then pay a small cost in the online phase. In some instances one can choose which class one wants to be in. For example, for Q2 access structures over a general finite field one can use the protocol of [48] if one wishes to utilize an offline phase, but the protocol in [8] if ones wants to avoid the offline phase (but have a slightly slower "online" phase). Although the performance of [8] degrades considerably for small finite fields, whereas that of [48] does not degrade at all (however [48]'s offline phase performance degrades if the finite field is small). Note [8] is expressed in terms of $t < n/2$ threshold adversaries but it can be trivially extended to any Q2 access structure.

The communication model also plays a part with protocols based on Garbled Circuits using a constant number of rounds, whereas protocols based on linear-secret sharing (LSSS) requiring rounds (roughly) proportional to the circuit depth. In all cases the total amount of communication, and computation, is roughly proportional to the number of multiplcition gates with the arithmetic circuit over $\mathbb{K}$ which represents the function to be computed[4]. The LSSS based protocols cost (essentially) one round of communication per each multiplicative depth, and communication cost linear in the number of multplication gates.

It is possible to mix GC and LSSS based MPC in one application, and pass between the two representations. For special access structures one can define special protocols for this purpose, see [28] for example. For general access structures one can apply the technique of doubly-authenticated bits (so called daBits) introduced in [44]. This latter method however comes with a cost. Assuming we are converting ℓ bit numbers, then not only does one need to generate (at least) ℓ daBits, but when transforming from the LSSS to the GC world one requires to evaluate a garbled circuit with roughly $3 \cdot \ell$ AND gates. The more expensive part is actually computing the daBits themselves. The paper [28] claims a cost of 0.163ms per daBit, for fields

---

[4] This is not strictly true as one often does not represent the function as a pure arithmetic circuit. But as a first order approximation this holds

of size 2^{128}. Whilst the fields used in the lattice based post-quantum signature algorithms are much smaller (of the order of 2^{20}) we use the same estimate[5].

Of course one could execute bitwise operations in an LSSS-based MPC for an odd modulus q using the methods described in [7, 10]. But these are generally slower than performing the conversion to a garbled circuit representation and then performing the garbled circuit based operation. Especially when different operations are needed to be performed on the same bit of data.

Protocol Name	Field $\mathbb{K}$	Access Structure	Pre-Proc Model	Rounds	Example Reference
SPDZ family	Large $\mathbb{K}$	Full Threshold	✓	$\approx \mathsf{depth}(C)$	[12]
Tiny-OT family	$\mathbb{F}_2$	Full Threshold	✓	$\approx \mathsf{depth}(C)$	[30, 40]
SPDZ-2k	$(\mathbb{Z}/2^k\mathbb{Z})$	Full Threshold	✓	$\approx \mathsf{depth}(C)$	[9]
n-party GC family	$\mathbb{F}_2$	Full Threshold	✓	constant	[22, 49]
General Q2	Any	Q2	✓	$\approx \mathsf{depth}(C)$	[48]
General Q2	Any	Q2	-	$\approx \mathsf{depth}(C)$	[8]
Special GC	$\mathbb{F}_2$	$(t,n) = (1,3)$	-	constant	[39]

Table 2. Summary of main practical MPC protocols

MPC of Standard Functionalities: A number of the signature schemes submitted to NIST make use of keyed (and unkeyed) symmetric functions which need to be applied in any threshold implementation to secret data. Thus any threshold implementation will need to also enable a threshold variant of these symmetric primitives. Here we recap, from the literature, the best timings and costs one can achieve for such primitives. We will use these estimates to examine potential performance in our discussions which follow.

In [28] the authors give details, also within the context of thresholdizing a NIST PQC submission (this time an encryption algorithm), of an MPC implementation of the SHA-3 round function (within the context of executing the KMAC algorithm). The round function f for SHA-3 requires a total of 38,400 AND gates, and using a variant of the three party honest majority method from [39], the authors were able to achieve a latency of 16ms per execution of f, for a LAN style setting. This equates to around $0.4\mu s$ per AND gate. Any actual application of SHA-3 requires multiple executions of the round function f; depending on how much data is being absorbed and how much is being squeezed.

In [49] give timings for a full-threshold garbled circuit based evaluation of various functions. Concentrating on the case of AES and SHA-256, and three party protocols, the authors obtain a latency of 95ms (13ms online) for the 6800 AND gate AES circuit, and 618ms (111ms online) for the 90825 AND gate SHA-256 circuit, again for a LAN setting. These times correspond to between 1μ and $2\mu s$ per AND gate, thus the three party full threshold setting is slightly slower than the honest majority setting (as is to be expected).

For general arithmetic circuits the estimates in [8] in the honest majority three party setting for a 61-bit prime field give a time of 826ms to evaluate a depth 20 circuit with one million multiplication gates, in a LAN setting. Thus we see that when using arithmetic circuits over such a finite field one can deal we obtain a similar throughput, in terms of multiplications per second, as one has when looking at binary circuits using garbled circuit techniques. However, the fields are of course much larger and we are performing more "bit operations" per second in some sense.

However, the protocol in [8] for 61-bit prime fields assumes a statistical security of 61-bits, i.e. the adversary can pass one of the checks on secure multiplications probability $1/2^{61} = 1/\#\mathbb{K}$. For smaller finite fields the performance degrades as one needs to perform more checks. A back-of-the-envelope calculation reveals one would expect a throughput of roughly 250,000 multiplications per second in the case of $\mathbb{F}_{2^8}$.

---

[5] Arithmetic modulo a prime of size 2^{20} is faster, but on the other hand one then has to perform more work to obtain the same level of active security.

Whilst these times are comparing apples and oranges, they do give an order of magnitude estimate of the time needed to compute these functions. Generally speaking, one is looking for operations which involve a few number of multiplications. Almost all NIST signature submissions make use of SHAKE-256, as a randomness expander, or hash-function. The SHAKE-256 algorithm is based on SHA-3. Recall that an application of SHA-3/SHAKE-256 on an input of ℓ_i bits, to produce an output of ℓ_o bits, will require (ignoring issues when extra padding results in more blocks being processed) a total of

$$\mathsf{rounds}(\ell_i, \ell_o) := \left\lceil \frac{\ell_i}{1088} \right\rceil + \left\lceil \frac{\ell_o}{1088} \right\rceil - 1$$

iterations of the main Keccak round function, since the rate of SHA3-256 is $r = 1088$ bits. In what follows we use the current best MPC evaluation time for this function (of 16ms from [28]) to obtain an estimate of how a specific application of SHAKE-256/SHA-3 will take.

3 Lattice Based Schemes

Lattice based signature schemes have a long history, going back to the early days of lattice based cryptography. Early examples such as NTRUSign [23] were quickly shown to be insecure due to each signature leaking information about the private key [18]. In recent years following the work of Lyubashevsky [34] the standard defence against such problems has been to adopt a methodology of Fiat–Shamir-with-aborts. All of the three lattice based submissions to NIST Round-2 follow this paradigm. However, we shall see that this means that they are all not particularly tuned to turning into threshold variants; for roughly the same reasons; although Falcon is slightly better in this regard. In all our lattice descriptions we will make use of a ring R_q, which one can take to be the cyclotomic ring $\mathbb{Z}[X]/(X^N + 1)$ reduced modulo q.

3.1 Crystals-Dilithium

The Dilithium Signature Algorithm

1. $z \leftarrow \bot$
2. While $(z = \bot)$ do
   (a) Sample a "short" $y \in R_q^l$ with $\|y\|_\infty \le \gamma_1$.
   (b) $v \leftarrow A \cdot y$.
   (c) Let w be the topbits of v.
   (d) $c \leftarrow H(\mu \| w) \in R_q$.
   (e) $z \leftarrow y + c \cdot s_1$.
   (f) If z or the lower bits of $v - c \cdot s_2$ are too big then set $z \leftarrow \bot$.
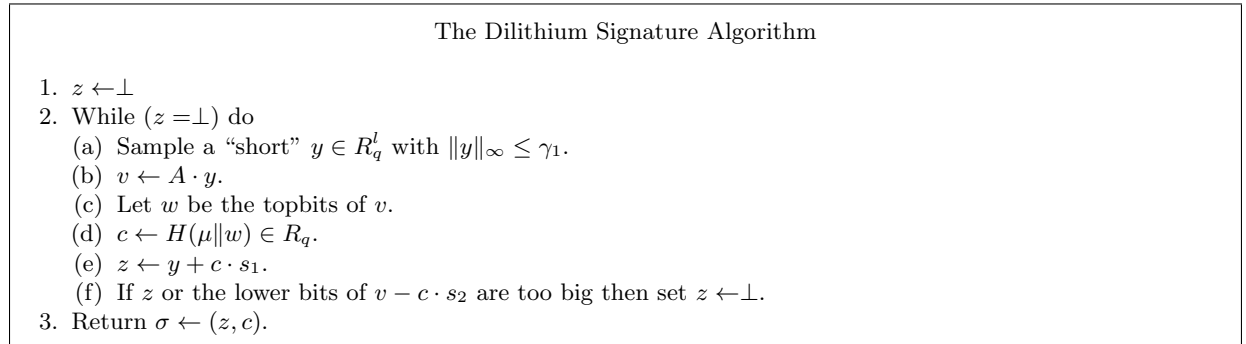3. Return $\sigma \leftarrow (z, c)$.

Figure 2. The Dilithium Signature Algorithm

The Dilithium [35] signature scheme is based on the Module-LWE problem. The secret key is two "short" vectors (s_1, s_2) with $s_1 \in R_q^l$ and $s_2 \in R_q^k$, and the public key is a matrix $A \in R_q^{k \times l}$ and a vector $t \in R_q^k$ such that $t = A \cdot s_1 + s_2$. The high-level view of the signature algorithm for signing a message μ is given in Figure 2, for precise details see the main Dilithium specification. We do not discuss the optimization in the Dilithium specification of the MakeHint function, to incorporate this will involve a few more AND gates in our discussion below. To aid exposition we concentrate on the basic signature scheme above. At the Level-3 security level the main parameters are set to be $N = \deg R_q = 256$, $q = 2^{23} - 2^{13} + 1$ and $(k, l) = (5, 4)$. There are a number of other parameters which are derived from these, in particular $\gamma_1 = (q - 1)/16$ and $\gamma_2 = (q - 1)/32$.

From our point of view we see that Dilithium is a signature scheme in the Fiat-Shamir-with-aborts family. If we did not have the while-loop in the signature algorithm, then the values of z and $v - c \cdot s_2$ would leak

information to the adversary. Thus it is clear that *any* distributed version of Dilithium signatures should maintain the secrecy of these intermediate values. Only values values which pass the size check and are output as a valid signature, can be revealed.

The parameters of the algorithm are selected so that the probability of needing two iterations of the while loop is less than one percent. Thus we can concentrate on the case of only executing one iteration of the loop. We assume that the secret key has been shared in an LSSS scheme over $\mathbb{F}_q$ which supports one of the MPC algorithms for LSSS schemes discussed in the introduction. We now discuss each of the lines of the main while loop in turn:

- **Line 2a:** In the specification the values y are derived in a deterministic manner from a PRF. However, this is not checked in the verification algorithm. Thus any threshold implementation can select y using any randomness process, as long as the values selected are from the same distribution as in the original Dilithium specification; namely elements in R_q^l whose coefficients are selected from a uniform distributed with bounded absolute value $\gamma_1 - 1$. The value $2 \cdot \gamma_1 - 2$ is just less than 2^{20}, thus this operation can be performed by generating 20 bits at random, and then testing whether the resulting number is less than $2 \cdot \gamma_1 - 2$ (which requires 20 AND gates in a garbles circuit). If this passes then the resulting value has γ_1 subtracted from it to obtain a value in $[-\gamma_1, \ldots, \gamma_1]$. This last step requires (approximately) 60 AND gates. In total, this line requires (at least) $20 \cdot N \cdot l = 20480$ random shared bits, and $80 \cdot N \cdot l = 81920$ AND gates.
- **Line 2b:** This line (as well as line 2e) is linear, and hence is much better suited to implementation via a LSSS based MPC operation. Thus we could either convert the output of the previous line to a LSSS based representation using the daBit technique mentioned earlier, or we need to execute this line using a Garbled Circuit. It would seem that the conversion method is better, as we will use the converted values again in line 2e. This conversion will require (at least) $23 \cdot N \cdot l = 23552$ daBits, but no garbled circuit evaluations.
- **Line 2c:** The algorithm to perform this step is defined in Figure 3 of the Dilithium specification. For each coefficient c in the elements in v one takes its value modulo q, and then takes the value $(c - c')/\gamma_1$ where c' is the representative of c modulo γ_1 in the interval $(-\gamma_2, \ldots, \gamma_2]$. This computation is best performed using a Garbled Circuit, which means the output from the previous step needs to be converted via daBits into a Garbled Circuit representation (requiring $23 \cdot N \cdot k = 29440$ daBits and garbled circuits costing roughly $3 \cdot 23 \cdot N \cdot k = 70656$ AND gates). The binary circuit to compute the reduction operation per coefficient requires about 170 AND gates. This means we require another $170 \cdot N \cdot k = 218880$ AND gates for this step.
- **Line 2d:** The Dilithium specification says that H should be a hash function which outputs a polynomial in R_q with exactly 60 non-zero coefficients equal to ± 1. The algorithm used for this is SHAKE-256. The input of μ, along with the $128 \cdot k$ bytes representing w are first absorbed into the sponge. A variable amount of output is then squeezed so as to obtain an output of the correct form. The *minimum* amount of squeezed output needed is 68 bytes (8 for the sign bits and 60 assuming the associated rejection sampling in the function H never rejects).
  In what follows we assume that μ is tiny. For longer messages μ one could of course, compute the initial absorbtion into H in the clear, since the arguments to H are such that μ is passed in first. Thus our estimates for small messages will also apply to large messages. Thus the cost of the evaluation of H will be at least $\mathsf{rounds}(128 \cdot k \cdot 8, 68 \cdot 8) = 5$ of the SHA-3 round function. Given the minimum the execution time of 16ms for a single round of SHA-3 from [28], for a relatively simple MPC instantiation for threshold $(t, n) = (1, 3)$, we expect the time to evaluate Dilithium in a distributed manner will be *at least* $5 \cdot 16 = 80$ms.
  In the context of the implementation of masked implementations of signatures in the GLP family, some authors have proposed solutions which do not mask the inputs to the hash function [2]. However, this comes at the expense of having to assume a non-standard hard problem. If one is willing to assume such a non-standard problem then a threshold implementation could be executed which keeps the values v secrets, but reveals the values w at the end of line 2c. Thus in such a situation line 2d is for free.
- **Line 2e:** This is a linear operation, as c can be held in the clear, and thus given the conversion to the LSSS representation which we performed on line 2b, this line costs nothing in a threshold implementation.

- **Line 2f:** This is much like line 2c. The most efficient way of performing this check is to compute $t = v - c \cdot s_2$ in the LSSS representation and then convert z and t to the GC representation, which requires $(k + l) \cdot 23 \cdot N = 52992$ daBits and $3 \cdot (k + l) \cdot 23 \cdot N = 158976$ AND gates. The execution of the final comparison check requires about $170 + 20$ AND gates per coefficient to verify the size of t and 20 AND gates per coefficient to verify the size of z. Thus the total number of AND gates for this step is $190 \cdot k \cdot N + 20 \cdot l \cdot N = 263680$.

So in total, and bearing in mind we have not described the additions needed for the MakeHint optimization in the specification, and we have underestimated the requirements due to the need to loop at various stages, our threshold variant would require a minimum of 76544 daBits and 794112 AND gates (assuming the SHA-3 operation is performed in the clear). Thus we expect Dilithium to be able to be executed in under a second in a threshold manner, plus the time needed to produce the daBits (which we estimate to be $76544 \cdot 0.163 = 12476$ms, i.e. about 12 seconds).

3.2 qTesla

qTesla is a signature scheme based on the ring-LWE problem, and like the previous one it too uses rejection sampling to avoid information leakage from signatures. The secret key is a pair $s, e \in R_q$, where e is small and R_q has degree N. The public key is a value $a \in R_q$ in R_q along with the value $t = a \cdot s + e$. The high level view of the signature algorithm is given in Figure 3. For the Level-3 security level we have the parameters $N = \deg R_q = 1024$, $B = 2^{21} - 1$, $q = 8404993$, and $d = 22$.
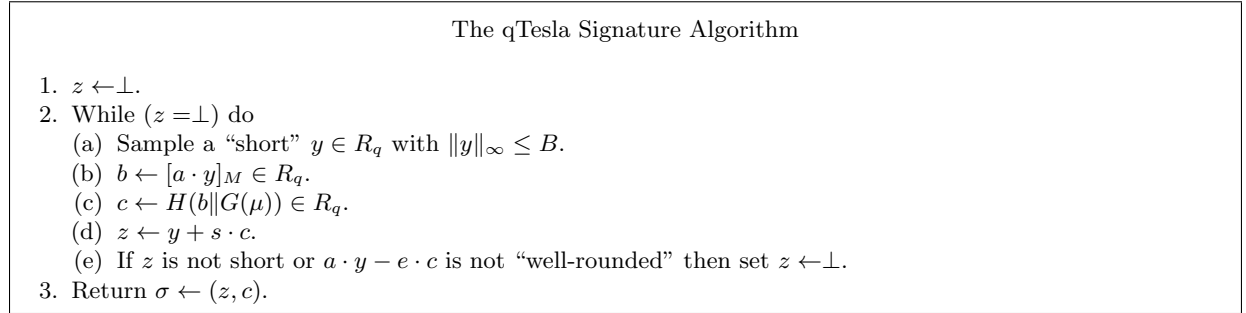

The qTesla Signature Algorithm

1. $z \leftarrow \perp$.
2. While $(z = \perp)$ do
   (a) Sample a "short" $y \in R_q$ with $\|y\|_\infty \leq B$.
   (b) $b \leftarrow [a \cdot y]_M \in R_q$.
   (c) $c \leftarrow H(b\|G(\mu)) \in R_q$.
   (d) $z \leftarrow y + s \cdot c$.
   (e) If z is not short or $a \cdot y - e \cdot c$ is not "well-rounded" then set $z \leftarrow \perp$.
3. Return $\sigma \leftarrow (z, c)$.

Figure 3. The qTesla Signature Algorithm

The operation $[x]_M$ applied to $x \in R_q$ provides a rounding operation akin to taking the top $(\log_2 q - d)$ bits of x in each coefficient. We define $[x]_M = (x \pmod q) - x \pmod{2^d})/2^d$ where the two modular operations perform a centered reduction (i.e. in the range $(-q/2, \ldots, q/2]$. The values of $[x]_M$ are stored in one byte per coefficient.

The function G is a hash function which maps messages to 512 bit values, and H is a hash function which maps elements in $R_q \times \{0,1\}^{512}$ to a 512-bit string c, which is then treated as a trinary polynomial. The functions H and G being variants of SHAKE-256 (or SHAKE-128 for the low security variants). Again much like the Dilithium, due to the rejection sampling the computation of y and the evaluation of H must be done in shared format. And again, we assume that the secret key has been shared in an LSSS scheme over $\mathbb{F}_q$. We again treat each line in turn.

- **Line 2a:** The generation of y is done in the specification via an algorithm which calls cSHAKE-256 (resp. cSHAKE-128) much like the algorithm for Dilithium. However, again much like in Dilithium, this can be generated in a threshold implementation by just generating a random shared y value with the correct distribution (which is again a polynomial with coefficients selected from a small interval in a uniformly random manner). This is simpler than in DiLithium as to produce an element in the range $[-B, \ldots, B]$ requies us to simply sample 22 bits per coefficient. Thus we simply need to sample $22 \cdot N = 22528$ shared random bits. If we sample these in the garbled circuit representation then we can convert the values to the LSSS representation (needed in lines 2b and 2d) using around $24 \cdot N = 24576$ daBits. Thus bar the generation of the daBits and the shared bits this step is for free.

- **Line 2b:** Here we need to compute $a \cdot y \in R_q$, which is a linear operation and best performed using the LSSS representation of y. Then we need to compute the resulting coefficients back to the garbled circuit representation so as to compute the value of $[c]_M$ on each coefficient. This will take another $24 \cdot N = 22528$ daBits and a garbled circuit evaluation of $3 \cdot 24$ AND gates per coefficient, i.e. $3 \cdot 24 \cdot N = 73728$ AND gates in total. However, once we have the shared bit representation of the coefficients of $a \cdot y$, the computation of $[a \cdot y]_M$ is for free.
- **Line 2c:** The execution of H needs to hash approximately $(N \cdot 8) + 256$ bits, which is already a lot so we ignore the required rounds to produce the output c. Thus the evaluation of H will require at least $\mathsf{rounds}((N \cdot 8) + 256, 0)$ secure executions of the SHA-3 round function. Thus we will require a minimum of $\mathsf{rounds}(1024 \cdot 8 + 256, 0) = 8 - 1 = 7$ executions of the SHA-3 round function, equating to a minimum execution time of $7 \cdot 16 = 112$ms. Again this is only an underestimate as we have not estimated the time needed to perform the arithmetic operations modulo q and the rounding operations.

  The same consideration that we had re Dilithium with respect to non-standard assumptions also applies to qTesla. In particular, if one is willing to assume the type of non-standard hard problem introduced in [2], then one can apply the hash function to clear data; thus significantly reducing the cost of performing a threshold implementation.
- **Line 2d:** Given that y and s are presented in LSSS form, and c can be in the clear, the evaluation of this line can be performed in an LSSS based MPC over $\mathbb{F}_q$ for free.
- **Line 2e:** Just as in Dilithium we need to convert the output from the previous line, as well as $w = a \cdot y - e \cdot c$ over to the GC world so as to perform the comparison more efficiently. The conversion requires at least $2 \cdot N \cdot 24 = 49152$ daBits, plus the execution of garbled circuits of size roughly $2 \cdot 3 \cdot N \cdot 24 = 147456$ AND gates. The final comparisons are then relatively easy in the case of qTesla, requiring only $6 \cdot N$ comparisons, i.e. $6 \cdot N \cdot 24 = 6144$ AND gates.

Putting this all together, and assuming the hash function on line 2c is performed in the clear, we find that we require a minimum of 98304 daBits, and 227328 AND gates. Thus we expect qTesla to be faster than Dilithium (purely on an AND gate count), with a rough estimate of $98304 \cdot 0.163 + 227328/1000 = 16250$ms, i.e. about 16 seconds.

3.3 Falcon

Falcon [42] is another lattice based scheme, and the only one to have NTRU-like public keys. It is based on the GPV framework [17]. The private key is a set of four "short" polynomials $f, g, F, G \in R_q$ such that $f \cdot G = g \cdot F$ in the ring R_q. The public key is the polynomial $h \leftarrow g/f$, which will have "large" coefficients in general. Associated to the private key is the private lattice basis in the FFT domain

$$\overline{B} = \begin{pmatrix} \mathsf{FFT}(g) & -\mathsf{FFT}(f) \\ \mathsf{FFT}(G) & -\mathsf{FFT}(F) \end{pmatrix}.$$

There is also a data structure T, called the *Falcon Tree* associated to the private key, which can be thought of as a set of elements in the ring R_q. At the Level-3 security level one has $N = \deg R_q = 768$ and $q = 18435$. A high level view of the signature algorithm is given in Figure 4.

Again, we assume that the secret key has been shared in an LSSS scheme over $\mathbb{F}_q$, and we go through each line in turn.

- **Line 2:** The function H is based on SHAKE-256 and outputs, like all the other lattice based signature schemes, an element in the ring R_q. An interesting aspect of the Falcon design is that the hash function H can be evaluated in the clear in any threshold implementation, since the output value c does not reveal any information about the private key. Thus unlike the previous two designs the evaluation of the hash function can be performed in the clear without the need for additional hardness assumptions.
- **Line 3:** Given a secret shared value of $\overline{B}^{-1}$, the evaluation of t is then a linear operation.

<div style="border:1px solid">

<center>The Falcon Signature Algorithm</center>

1. $r \leftarrow \{0,1\}^{320}$.
2. $c \leftarrow H(r\|\mu)$.
3. $t \leftarrow (\mathsf{FFT}(c), \mathsf{FFT}(0)) \cdot \overline{B}^{-1}$.
4. $z \leftarrow\perp$.
5. While $(z =\perp)$ do
   (a) $z \leftarrow \mathsf{ffSampling}_n(t, T)$.
   (b) $s \leftarrow (t - z) \cdot \overline{B}$.
   (c) If s is not short then set $z =\perp$.
6. $(s_1, s_2) \leftarrow \mathsf{FFT}^{-1}(s)$.
7. $s \leftarrow \mathsf{Compress}(s_2)$.
8. Return $\sigma = (r, s)$.

</div>

<center>**Figure 4.** The Falcon Signature Algorithm</center>

- **Line 5a:** The sub-algorithm $\mathsf{ffSampling}$ applies an FFT like procedure to the input data values, thus it is a highly linear algorithm and lends itself well to being evaluated in a threshold manner. The algorithm essentially provides a randomized rounding on the coefficients of t, with respect to the data structure T However, the randomization within this rounding requires the sampling of a discrete Gaussian distribution.

  The specification document says that how the discrete Gaussian is evaluated is "arbitrary" and "outside the scope of this specification" [42], bar needing to be close in terms of the Rényi divergence. However, a Gaussian sampler is defined within the specification, [42][Section 4.4], for use in the reference implementation. The sampler takes as parameters a centre ν and a standard deviation σ, where we can assume for the following discussion that $\nu \in [0, \ldots, 1)$. The specification of the sampler uses ChaCha20 to define a PRNG, although any PRNG could be used in practice. The outputs from this PRNG are used to sample an integer value z with a half Gaussian with standard deviation $\sigma_0 = 2$ (or sometimes $\sigma_0 = \sqrt{5}$ depending on the precise parameters) and centre zero. Then a random bit b is computed and the value $z' \leftarrow b + (2 \cdot b - 1) \cdot z$ is computed. This value is accepted/rejected on the basis of another bit d chosen with probability

  $$e^{(z-b)^2/(2\cdot\sigma_0^2)-(z-\nu)^2/(2\cdot\sigma^2)}. \tag{1}$$

  To obtain the correct Rényi divergence the sampler needs to operate with a floating point precision of 53-bits.

  Here we see the main issue in terms of a threshold implementation. The output from this sampler needs to be kept secret from the players to ensure security, thus the evaluation of the sampler needs to be performed in some form of threshold manner. Yet we seem to need to compute floating point formulae to high-precision. This is made more complex as the values (ν, σ) change with every call to the sampler, with the precise values depending on the message to be signed. Thus any threshold implementation will need an efficient way of sampling from this distribution in a threshold manner.

  Using the methods from [26] one can construct a Gaussian sampler for the Falcon signature scheme with a precision of 116 bits of security, using a circuit with just over 3700 AND gates[6]. We then take the output of this sampler and convert it back to the LSSS representation, which requires at least 15 daBits per conversion (and no garbled circuit evaluations), plus a garbled circuit with $3 \cdot 15$ AND gates. We need to compute *at least N* of these sampler evaluations followed by conversions, thus we require at least $N \cdot 15 = 11520$ daBits and $3 \cdot N \cdot 15 = 34560$ AND gates.
- **Line 5b:** The evaluation of s in the main while-loop can be evaluated using low depth arithmetic circuit over $\mathbb{F}_q$, since $t - z$ will be secret shared in the LSSS based MPC, as will $\overline{B}$.
- **Line 5c:** The comparison here can be performed just as in the previous lattice based schemes. We convert, using daBits, to into a garbled circuit friendly representation which requires at least $n \cdot \log_2 q = 768 \cdot 15 = 11520$ daBits, plus circuits totalling $3 \cdot 15 \cdot 768 = 34560$ AND gates. Then the comparisons can be performed using $15 \cdot 768 = 11520$ AND gates

---

[6] Private communication with the authors of [26].

- **Line 6:** The computation of (s_1, s_2) via the inverse FFT is also a linear operation.
- **Line 7:** The computation of the compress function is purely for performance reasons and can be computed in the clear if needs be.

Thus in total we require, per Falcon signature, a minimum of 34560 daBits and 80640 AND gates, which would equate to a rough minimum time of $34560 \cdot 0.163 + 80640/1000 = 5713$ms, or nearly six seconds. Again most of the time coming from daBit generation.

4 MPC-in-the-Head Based Scheme

The MPC-in-the-Head paradigm for producing zero-knowledge proofs was developed in [25]. The prover, to prove knowledge of a preimage x of some function $\Phi(x) = y$ (where Φ and y are public), simulates an MPC protocol to compute the functionality Φ, with the input x shared among the simulated parties. The prover executes the protocol (in it's head), then commits to the state and the transcripts of all players. Then it sends the verifier these commitments and randomly opens a (non-qualified) subset of them (the precise subset is chosen by the verifier). The verifier checks that the simulated protocol was correctly executed using the opened values. If everything is consistent, it then accepts the statement that the prover knows x, otherwise it rejects. Typically, the proof has to repeated several times in ordar to achieve high security. Clearly to obtain a signature scheme we apply the Fiat–Shamir transform so that the verifier's choices are obtained by hashing the commitments with the message.

4.1 Picnic

Picnic is a digital signature scheme whose security entirely relies on the security only of symmetric key primitives, in particular the security of SHA-3 and a low-complexity block cipher called Low-MC [1]. The core construction is a zero-knowledge proof of knowledge of a preimage for a one-way function $y = f_k(x)$, where f is the Low-MC block cipher, the values x and y are public and the key k is the value being proved. Using the Fiat–Shamir and MPC-in-the-Head paradigms we obtain a signature scheme with public key (x, y) and private key k.

In this paper we concentrate on Picnic-1, but a similar discussion also applies to the Picnic-2 construction. The specific proof system that implements the MPC-in-the-Head for Picnic-1 is ZKBoo++ [1], which is itself an extension of the original ZKBoo framework from [19]. The simulated MPC protocol is between three parties, and is executed at a high level as in Figure 5
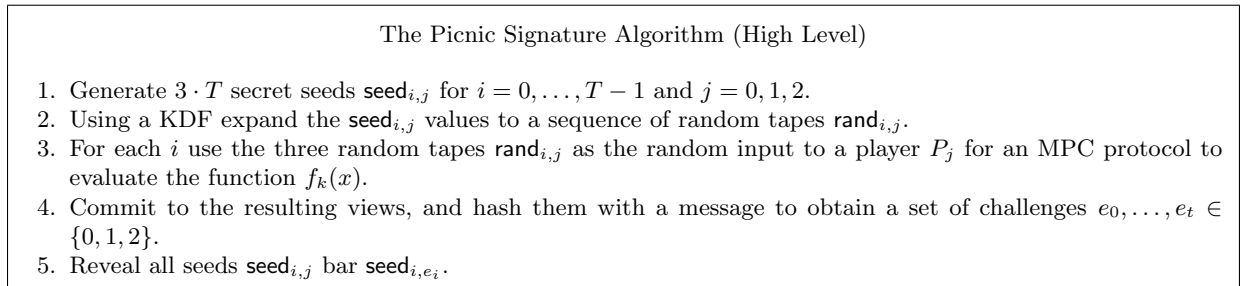

The Picnic Signature Algorithm (High Level)

1. Generate $3 \cdot T$ secret seeds $\mathsf{seed}_{i,j}$ for $i = 0, \ldots, T-1$ and $j = 0, 1, 2$.
2. Using a KDF expand the $\mathsf{seed}_{i,j}$ values to a sequence of random tapes $\mathsf{rand}_{i,j}$.
3. For each i use the three random tapes $\mathsf{rand}_{i,j}$ as the random input to a player P_j for an MPC protocol to evaluate the function $f_k(x)$.
4. Commit to the resulting views, and hash them with a message to obtain a set of challenges $e_0, \ldots, e_t \in \{0, 1, 2\}$.
5. Reveal all seeds $\mathsf{seed}_{i,j}$ bar seed_{i,e_i}.

Figure 5. The Picnic Signature Algorithm (High Level)

In our analysis we will ignore any hashing needed to produce commitments and the challenge, and we will simply examine the operation of the key derivation in step 2 of Figure 5. It is clear that in the MPC-in-the-Head paradigm the seeds need to be kept secret until the final reveal phase, thus the derivation of the random tape from the seed needs to be done in a secure manner in a any threshold implementation.

In Picnic the precise method used to derive the random tape is to use

$$\mathsf{rand}_{i,j} = \mathsf{KDF}\left(H_2\left(\mathsf{seed}_{i,j}\right) \|\mathsf{salt}\|i\|j\|\mathsf{length}\right)$$

where

- The seeds are S bits long.
- The salt is 256 bits long.
- The integers i, j and length are encoded as 16-bit values.
- The output length (length), of the KDF, is $n + 3 \cdot r \cdot s$ when $j = 0, 1$ and $3 \cdot r \cdot s$ when $j = 2$.

We again concentrate on the NIST security Level-3, which is instantiated with the parameters $S = n = 192$, $T = 329$, $s = 10$ and $r = 30$. The hash function H_2 is SHAKE-256 based with an output length of 384 bits. Thus the execution of H_2 requires only two executions of the SHA-3 round function. Each KDF operation is also cheap, requiring either two or three rounds. The problem is we need to execute these operations so many times. The total number of executions of the round function of SHA-3 is given by

$$T \cdot \Big(2 + 2 \cdot \mathsf{rounds}(384 + 256 + 32 + 32, n + 3 \cdot r \cdot s)$$

$$+ \mathsf{rounds}(384 + 256 + 32 + 32, 3 \cdot r \cdot s) \Big)$$

$$= 329 \cdot \Big(2 + 2 \cdot \mathsf{rounds}(704, 1092) + \mathsf{rounds}(704, 900) \Big)$$

$$= 329 \cdot (2 + 2 \cdot (3 - 1) + (2 - 1)) = 2303.$$

Thus given our estimate of a minimum of 16ms for a SHA-3 round execution in MPC we see that even this part of the Picnic algorithm is expected to take $16 \cdot 3290$ ms, i.e. 37 seconds!

5 Hash Based Scheme

Hash based signatures have a long history going back to the initial one-time signature scheme of Lamport [29]. A more efficient variant of the one-time signature attributed to Winternitz is given in [38], where a method is also given to turn the one-time signatures into many-time signatures via so-called Merkle-trees. The problem with these purely Merkle tree based constructions is that they are strictly a statefull signature scheme. The signer needs to maintain a changing state between each signature issued, and the number of signatures able to be issued is bounded as a function of the height of the Merkle tree.

To overcome these issues with state the SPHINCS signature scheme was introduced in 2015 [3], which itself builds upon ideas of Goldreich elaborated in [21], and going back to [20]. In the SPHINCS construction messages are still signed by Winternitz one-time signatures, but the public keys of such signatures are then authenticated via another (similar) structure called a Forest of Random Subsets (which is itself based on earlier work in [43]).

5.1 SPHINCS+

The only hash based signature scheme to make it into the second round of the NIST competition is SPHINCS+ [24]. We refer the reader to the design document [24] for a full description. For our purposes we recall that messages are signed using Winternitz one-time signatures which are then authenticated using a FORS tree. The parameters which are of interest to us are: n the security parameter in bytes, w a parameter related to the underlying Winternitz signature, h the height of the hypertree, d the number of layers in the hypertree, k the number of trees in a FORS, t the number of leaves in a FORS tree. From these two length functions are defined[7]

$$\mathsf{len}_1 = \left\lceil \frac{8 \cdot n}{\log_2 w} \right\rceil, \qquad \mathsf{len}_2 = \left\lfloor \frac{\log(\mathsf{len}_1 \cdot (w - 1))}{\log w} \right\rfloor + 1, \quad \text{and} \quad \mathsf{len} = \mathsf{len}_1 + \mathsf{len}_2.$$

The scheme uses (essentially) four hash functions labelled $\mathbf{F}$, $\mathbf{H}$, $\mathbf{PRF}$ and T_{len}. The function $\mathbf{F}$ is used as the main function in the Winternitz signature scheme, as well as the FORs signature. The underlying expansion

---

[7] Note the definition of len_1 in the specification is wrong and need correcting which we do below

the secret key into secret keys of the trees is done via the function **PRF**. The function **H** is used to construct a root of the associated binary trees, where as T_{len} is used to compress the len Winternitz public key values into a single n-bit value for use as a leaf in the Merkle tree. The evaluation of the **F** and **PRF** calls within a single signature needs to be performed on secret data, even though eventually some of the input/outputs become part of the public signature. The calls to **H** and T_{len} appear to be able to be performed on public data, and will not concern us here.

In what follows we concentrate on the SHAKE-256 based instantiation of SPHINCS+ (to be comparable with other signature schemes in this paper). In the SHAKE instantiation the execution of the function **F** requires two calls to the underlying SHA-3 permutation, where as **H** requires three calls to the underlying SHA-3 permutation, and **PRF** requires one call to the SHA-3 permutation.

To sign a message requires $k \cdot t + d \cdot w \cdot \mathsf{len} \cdot 2^{h/d}$ calls to **F** and $k \cdot t + d \cdot \mathsf{len} \cdot 2^{h/d} + 1$ calls to **PRF**. When instantiated with the parameters at the NIST Level-3 security level (for fast signing) we have $(n, w, h, d, k, t) = (24, 16, 66, 22, 33, 256)$. Leading to $\mathsf{len}_1 = 48$, $\mathsf{len}_2 = 3$ and $\mathsf{len} = 51$. This leads to a grand total of 152064 calls to **F** and 17425 calls to **PRF**. This leads to a total of 321553 calls to the SHA-3 internal permutation which need to be performed securely. With current best garbled circuit implementations this on its own would require 85 minutes to execute. Of course a complete threshold implementation would take longer as we have not looked at other aspects of the signature algorithm.

6 MQ Based Schemes

The history of MQ cryptography, is almost as old as that of hash-based signatures. The first MQ based scheme was presented in 1988 [37]. In terms of signature schemes based on the MQ problem, the original works were due to Patarin and were given the name "Oil and Vinegar" [27, 41]. The basic idea is to define a set of multivariate quadratic equations (hence the name MQ) $P : \mathbb{F}_q^m \longrightarrow \mathbb{F}_q^n$ and the hard problem is to invert this map, where q is a power of two[8]. The intuition being that inverting this map is (for a general quadratic map P) is an instance of the circuit satisfiability problem, which is known to be NP-Complete.

In three of the NIST candidate signature schemes the function P is generated so that there is an efficient trapdoor algorithm which allows the key holder to invert the map P using the secret key. In such situations the secret key is usually chosen to be two affine transforms $S : \mathbb{F}_q^n \longrightarrow \mathbb{F}_q^n$ and $T : \mathbb{F}_q^m \longrightarrow \mathbb{F}_q^m$, plus an easy to invert map $P' : \mathbb{F}_q^m \longrightarrow \mathbb{F}_q^n$ consisting of quadratic functions (note any function can be expressed in terms of quadratic functions by simple term rewriting). Then the public map is defined by $P = S \circ P' \circ T$. Of course the precise definition of this construction implies that one is not using a generic circuit satisfiability problem. However, for specific choices of P', q, n and m the construction is believed to provide a trapdoor one-way function.

Given we have a trapdoor one way function the standard Full Domain Hash construction gives us a signature scheme. Namely to sign a message μ, the signer hashes μ to an element $y \in \mathbb{F}_q^m$ and then exhibits a preimage of y under P as the signature s. To verify the signature the verifier simply checks that $P(s) = y$. Note, that many preimages can exist for y under P, thus every message could have multiple valid signatures. From this basic outline one can define a number of signature scheme depending on the definition of the "central map" P'. All of the Round-2 MQ based signaure schemes, with the exception of MQDSS, follow this general construction method; therefore we deal with MQDSS first.

Inverting Linear Systems in MPC Before proceeding we present a trick which enables us to efficiently solve linear systems in an LSSS based MPC system. We will use this in our analysis of two of the submissions, so we present it here first. Suppose we have a shared $n \times n$ matrix $\langle A \rangle$ over $\mathbb{F}_q$ and an n-dimensional shared vector $\langle \mathbf{b} \rangle$. We would like to determine $\langle \mathbf{x} \rangle$ such that $A \cdot \mathbf{x} = \mathbf{b}$. We do this using the algorithm in Figure 6. This algorithm either returns the secret shared solution or the $\perp$ symbol. This latter either happens because the input matrix has determinant zero, or the random matrix used in the algorithm has determinant zero

---

[8] To enable comparison with the NIST submissions we use the same notation in the sections which follow as used in the submissions. We hope this does not confuse the reader.

(which occures with probability $1/q$). The algorithm requires a total of three rounds of communication and $n^3 + n^2$ secure multiplications.
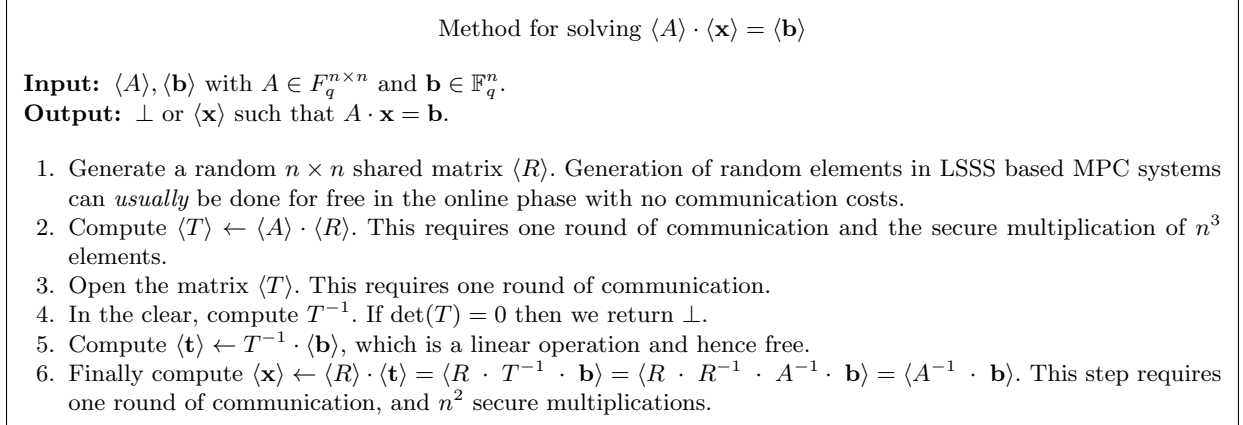

Method for solving $\langle A \rangle \cdot \langle \mathbf{x} \rangle = \langle \mathbf{b} \rangle$

Input: $\langle A \rangle, \langle \mathbf{b} \rangle$ with $A \in F_q^{n \times n}$ and $\mathbf{b} \in \mathbb{F}_q^n$.
Output: $\perp$ or $\langle \mathbf{x} \rangle$ such that $A \cdot \mathbf{x} = \mathbf{b}$.

1. Generate a random $n \times n$ shared matrix $\langle R \rangle$. Generation of random elements in LSSS based MPC systems can *usually* be done for free in the online phase with no communication costs.
2. Compute $\langle T \rangle \leftarrow \langle A \rangle \cdot \langle R \rangle$. This requires one round of communication and the secure multiplication of n^3 elements.
3. Open the matrix $\langle T \rangle$. This requires one round of communication.
4. In the clear, compute T^{-1}. If $\det(T) = 0$ then we return $\perp$.
5. Compute $\langle \mathbf{t} \rangle \leftarrow T^{-1} \cdot \langle \mathbf{b} \rangle$, which is a linear operation and hence free.
6. Finally compute $\langle \mathbf{x} \rangle \leftarrow \langle R \rangle \cdot \langle \mathbf{t} \rangle = \langle R \cdot T^{-1} \cdot \mathbf{b} \rangle = \langle R \cdot R^{-1} \cdot A^{-1} \cdot \mathbf{b} \rangle = \langle A^{-1} \cdot \mathbf{b} \rangle$. This step requires one round of communication, and n^2 secure multiplications.

Figure 6. Method for solving $\langle A \rangle \cdot \langle \mathbf{x} \rangle = \langle \mathbf{b} \rangle$

6.1 MQDSS

The MQDSS scheme takes a general quadratic function $\mathbf{F} : \mathbb{F}_q^n \longrightarrow \mathbb{F}_q^n$, i.e. one with no secret trap door (note the domain and codomain have the same dimension). To describe the scheme the authors make use of the polar form of the function $\mathbf{F}$, given by

$$\mathbf{G}(\mathbf{x}, \mathbf{y}) = \mathbf{F}(\mathbf{x} + \mathbf{y}) - \mathbf{F}(\mathbf{x}) - \mathbf{G}(\mathbf{y}).$$

It is readily verified that $\mathbf{G}$ is a bilinear map.

The public key is an image $\mathbf{v} \in \mathbb{F}_q^n$, with the private key being a preimage $\mathbf{s} \in \mathbb{F}_q^n$ such that $\mathbf{F}(\mathbf{s}) = \mathbf{v}$. Thus unlike the other signature schemes the underlying MQ hard problem is suspected to be harder than an MQ problem for a function with a trapdoor. The signature is then based on an interactive proof of knowledge of a pre-image of this function, which forms the basis of an identification protocol given in [45]. This is then converted into a signature scheme via the Fiat–Shamir heuristic in the standard way. Unlike traditional signature schemes the method in MQDSS uses a five round identification scheme as opposed to a three round scheme; but this is purely for efficiency reasons. The underlying identification protocol has a relatively low soundness error, thus this is amplified by a parameter r which gives the number of rounds executed. At the NIST Level-3 security level the authors recommend $(n, q, r) = (64, 31, 202)$.

In the specification the secret key, map $\mathbf{F}$, and randomness used within the protocol are all derived from hash functions and PRF's. However, for a threshold implementation which is functionally equivalent such expansion is not needed (bar the expansion needed to generate the function $\mathbf{F}$ for the public key). Thus we can ignore the expansion in our discussion and consider the public key as the pair $(\mathbf{F}, \mathbf{v})$ and the private key as $\mathbf{s}$. The signature scheme uses a hash function which generates a random per-signature nonce R, and then combines this with the message and public key to obtain a value $D \in \mathbb{F}_q^r$ which is the value which will be 'signed' via the Fiat-Shamir heuristic. This inner signing algorithm we give in Figure 7.

The values $\mathbf{r}_0^{(1)}, \ldots, \mathbf{r}_0^{(r)}, \mathbf{t}_0^{(1)}, \ldots, \mathbf{t}_0^{(r)}, \ldots, \mathbf{e}_0^{(1)}, \ldots, \mathbf{e}_0^{(r)}$ in the scheme are derived in a deterministic manner in the scheme from a PRG, this is never verified by the verification algorithm. Thus in a threshold implementation we can generate these values via any secure randomization method.

From a perspective of turning the algorithm into a threshold version the evaluation of the hash functions H, H_1 and H_2 can be applied to public data. Thus the key part is that the inputs to the commitments Com_0 and Com_1 need to be kept secret from the threshold players, as we cannot allow an adversary to learn r_0 and r_1 in the clear.

<div style="border:1px solid">

MQDSS Inner Signing Algorithm

1. $\mathbf{r}_0^{(1)}, \ldots, \mathbf{r}_0^{(r)}, \mathbf{t}_0^{(1)}, \ldots, \mathbf{t}_0^{(r)}, \ldots, \mathbf{e}_0^{(1)}, \ldots, \mathbf{e}_0^{(r)} \leftarrow \mathbb{F}_q^n$.
2. For $j \in \{1, \ldots, r\}$
   (a) $\mathbf{r}^{(j)} \leftarrow \mathbf{s} - \mathbf{r}_0^{(j)}$
   (b) $c_0^{(j)} \leftarrow \mathsf{Com}_0\left(\mathbf{r}_0^{(j)}, \mathbf{t}_0^{(j)}, \mathbf{e}_0^{(j)}\right)$
   (c) $c_1^{(j)} \leftarrow \mathsf{Com}_1\left(\mathbf{r}_1^{(j)}, \mathbf{G}\left(\mathbf{t}_0^{(j)}, \mathbf{r}_1^{(j)}\right) + \mathbf{e}_0^{(j)}\right)$
   (d) $\mathsf{com}^{(j)} \leftarrow \left(c_0^{(j)}, c_1^{(j)}\right)$
3. $\sigma_0 \leftarrow H\left(\mathsf{com}^{(1)}\|\ldots\|\mathsf{com}^{(r)}\right)$
4. $\mathsf{ch}_1 \leftarrow H_1(D, \sigma_0) = \left(\alpha^{(1)}, \ldots, \alpha^{(r)}\right) \in \mathbb{F}_q^r$.
5. For $j \in \{1, \ldots, r\}$
   (a) $\mathbf{t}_1^{(j)} \leftarrow \alpha^{(j)} \cdot \mathbf{r}_0^{(j)} - \mathbf{t}_0^{(j)}$
   (b) $\mathbf{e}_1^{(j)} \leftarrow \alpha^{(j)} \cdot \mathbf{F}\left(\mathbf{r}_0^{(j)}\right) - \mathbf{e}_0^{(j)}$
   (c) $\mathsf{resp}_1^{(j)} \leftarrow \left(\mathbf{t}_1^{(j)}, \mathbf{e}_1^{(j)}\right)$
6. $\sigma_1 = \left(\mathsf{resp}_1^{(1)}\|\ldots\|\mathsf{resp}_1^{(r)}\right)$
7. $\mathsf{ch}_2 = H_2(D, \sigma_0, \mathsf{ch}_1, \sigma_1) = \left(b^{(1)}, \ldots, b^{(r)}\right) \in \{0, 1\}^r$.
8. For $j \in \{1, \ldots, r\}$ set $\mathsf{resp}_2^{(j)} \leftarrow \mathbf{r}_{b^{(j)}}^{(j)}$
9. $\sigma_2 \leftarrow \left(\mathsf{resp}_2^{(1)}\|\ldots\|\mathsf{resp}_2^{(r)}\|c_{1-b^{(1)}}^{(1)}\|\ldots\|c_{1-b^{(r)}}^{(r)}\right)$
10. Return $\sigma \leftarrow (\sigma_0, \sigma_1, \sigma_2)$

</div>

Figure 7. MQDSS Inner Signing Algorithm

The specification states that the string commitments are implemented using a modification of SHAKE-256. These takes arrays of F_q elements, absorb into SHAKE-256, and then output things of size $k/8$ bytes, where k is the security parameter (which at Level-3 we have $k = 192$). Thus Com_0 takes as input $3 \cdot n \cdot \log_2 q$ bits of input and outputs k bits of data, whereas Com_1 takes as input $2 \cdot n \cdot \log_2 q$ bits of input and outputs k bits of data. And we need to execute Com_0 and Com_1 a total of r times.

Thus ignore at the Level-3 security level the number of secure evaluations of the SHA-3 inner permutation are given by

$$r \cdot \left(\mathsf{rounds}(3 \cdot n \cdot \log_2 q, k) + \mathsf{rounds}(2 \cdot n \cdot \log_2 q, k)\right)$$
$$= 202 \cdot \left(\mathsf{rounds}(960, 192) + \mathsf{rounds}(640, 192)\right)$$
$$= 202 \cdot ((1 + 1 - 1) + (1 + 1 - 1))$$
$$= 202 \cdot 2 = 404.$$

Thus just this aspect of a threshold implementation will take at least $404 \cdot 16 = 6454$ ms, i.e. about 6.5 seconds. Any full threshold implementation will thus take much longer than this, as we have not factored in the need to actually evaluate the functions $\mathbf{F}$ and $\mathbf{G}$.

6.2 GeMSS

GeMSS is a MQ based signature scheme whose signing algorithm requires the extraction of roots of polynomials over a characteristic two finite field. The secret key is a tuple $(F, \mathbf{S}, \mathbf{T})$ where F is a polynomail in $\mathbb{F}_2[X, v_1, \ldots, v_v]$ of degree D and $\mathbf{S}$ (resp. $\mathbf{T}$) are random square matrices over $\mathbb{F}_2$ of dimension $n + v$ (resp. n). For the Level-3 parameter sets we have $n = 265$, $v = 20$ and $D = 513$. The polynomial F is of the special form

$$F = \sum_{\substack{1 \le j < i < n \\ 2^i + 2^j \le D}} A_{i,j} \cdot X^{2^i + 2^j} + \sum_{\substack{0 \le i < n \\ 2^i \le D}} \beta_i(v_1, \ldots, v_v) \cdot X^{2^i} + \gamma(v_1, \ldots, v_v)$$

15

where β_i is a linear function and γ is a quadratic function. The scheme is based on the FDH paradigm, thus the main issue is whether one can produce a distributed variant is purely a question of whether one can distribute effectively the arithmetic operations.

The key part of the signing algorithm is the need to execute nb_ite times (where at Level-3 we have nb_ite = 4) a procedure which finds a root of the multivariate equation

$$F(Z, z_1, \ldots, z_v) - D = 0$$

for values of D which depend on the message and secret key. Such a root is to lie in the set $\mathbb{F}_{2^n} \times F_2^v$. To do this a random set of values $(z_1, \ldots, z_v) \in \mathbb{F}_2^v$ is selected, and then the resulting univariate polynomial is passed to the classical Berlekamp algorithm for root extaction.

Given $\mathbf{z} \in \mathbb{F}_2^n$ we define the polynomial $G(X) = F(X, z_1, \ldots, z_v)$, but in our threshold version we would need the $\mathbf{z}$ to be kept secret, as well as the constant D (and the values $A_{i,j}$, β_i, γ_i defining the polynomial F). Thus the entire polynomial $G(X)$ needs to be secret shared. We then take the greatest common divisor of $G(X)$ with $X^{2^n} - X$ to obtain a polynomial of degree r which is a product of linear terms. This is then passed to the equal degree factorization algorithm so as to obtain all the roots in $\mathbb{F}_{2^n}$, of which one is selected. We are interested in both the number of multiplications needed in $\mathbb{F}_{2^n}$ to extract a root, as well as the multiplicative depth.

In general we expect the depth for this algorithm to be large, in the worst case the greatest common divisor of $G(X)$ with $X^{2^n} - X$ will have depth D (since we need to apply in the worst case D euclidean divisions to obtain a final result). This would have to be implemented in a data oblivious manner, so the depth is likely to be larger[9]. Thus we expect a multiplicative depth of at least $4 \cdot D$.

If we let $M(D)$ denote the number of multiplications in $\mathbb{F}_{2^n}$ to multiply two polynomials of degree D, then we expect $O(M(D) \cdot \log(D) \cdot \log(D \cdot 2^n)) = \tilde{O}(n \cdot D)$ operations in $\mathbb{F}_{2^n}$ are needed to produce the root of $G(X)$. Thus in comparison to the non-MQ schemes it seems feasible to produce a threshold variant of GeMSS. However, the multiplicative depth and the requires number of finite field multiplications is likely to be quite large. We will see that other MQ schemes are much simpler, and require much less depth and multiplications to implement in a threshold manner.

6.3 Rainbow

The Rainbow signature scheme can be seen as a multilayer version of the original UOV. In its original version, the number of layers is determined by a parameter u. For $u = 1$ this is just the basic UOV scheme, whereas the candidate submission chooses $u = 2$. As described earlier we pick for the secret key two affine transforms $\mathcal{S} : \mathbb{F}_q^m \to \mathbb{F}_q^m$ and $\mathcal{T} : \mathbb{F}_q^n \to \mathbb{F}_q^n$. Along with a function $\mathcal{F}$, called the central map, which can be defined by quadratic functions. The public key is then the map $\mathcal{P} = \mathcal{S} \circ \mathcal{F} \circ \mathcal{T} : \mathbb{F}_q^n \to \mathbb{F}_q^m$.

In the Rainbow specification the affine maps $\mathcal{S}$ and $\mathcal{T}$ are chosen to be given by matrix multiplication by upper triangular matrices S and T. This means that the inverse matrices S^{-1} and T^{-1} are also upper triangular. In particular the inverses are selected to have the following block form

$$S^{-1} = \begin{pmatrix} \mathbf{1}_{o_1} & S_{o_1 \times o_2} \\ \mathbf{0}_{o_2 \times o_1} & \mathbf{1}_{o_2} \end{pmatrix} \quad \text{and} \quad T^{-1} = \begin{pmatrix} \mathbf{1}_{v_1} & T_{v_1 \times o_1} & T'_{v_1 \times o_2} \\ \mathbf{0}_{o_2 \times v_1} & \mathbf{1}_{o_1} & T''_{o_1 \times o_2} \\ \mathbf{0}_{o_2 \times v_1} & \mathbf{0}_{o_2 \times o_1} & \mathbf{1}_{o_2} \end{pmatrix}$$

where $S_{a \times b}$ etc denotes a matrix of dimension $a \times b$, $\mathbf{0}_{a \times b}$ denotes the zero matrix of dimension $a \times b$ and $\mathbf{1}_a$ denotes the identity matrix of dimension a.

To define the central map we define three constants (v_1, o_1, o_2), which at the Level-3 security level are chosen to be $(68, 36, 36)$. From these we define further parameters given by $v_2 = v_1 + o_1, n = v_3 = v_2 + o_2$ and $m = o_1 + o_2$. Note this means that $n = v_1 + m$. We then define the sets $V_i = \{1, \ldots, v_i\}$ and $O_i = \{v_i + 1, \ldots, v_{i+1}\}$, for $i = 1, 2$, which will be referred to as the vinegar (resp. oil) variables of the ith layer.

---

[9] We assume school book gcd algorithms are used here, which are likely to be the most efficient in this case.

The Rainbow central map $\mathcal{F} : \mathbb{F}_q^n \to \mathbb{F}_q^m$ can then be defined by the set of m quadratic polynomials $f^{(v_1+1)}, \ldots, f^{(n)}$ having the form

$$f^{(k)} = \begin{cases} \sum_{i,j\in V_1, i\leq j} \alpha_{ij}^{(k)} x_i \cdot x_j + \sum_{i\in V_1} \sum_{j\in O_1} \beta_{ij}^{(k)} x_i \cdot x_j & k = v_1 + 1, \ldots, v_2, \\ \sum_{i,j\in V_2, i\leq j} \alpha_{ij}^{(k)} x_i \cdot x_j + \sum_{i\in V_2} \sum_{j\in O_2} \beta_{ij}^{(k)} x_i \cdot x_j & k = v_2 + 1, \ldots, n, \end{cases}$$

where the coefficients $\alpha_{i,j}^{(k)}, \beta_{i,j}^{(k)}$ are randomly chosen from $\mathbb{F}_q$.

Signature generation (for the EUF-CMA scheme) is done by the steps

1. Compute the hash value $\mathbf{h} \leftarrow H\left(H\left(\mu\right) \| \mathsf{salt}\right) \in \mathbb{F}_q^m$, where μ is the message, salt is a random l-bit string and $H : \{0,1\} \to \mathbb{F}_q^m$ is an hash function.
2. Compute $\mathbf{x} \leftarrow S^{-1} \cdot \mathbf{h} \in \mathbb{F}_q^m$,
3. Compute a preimage $\mathbf{y} \in \mathbb{F}_q^n$ of $\mathbf{x}$ under the central map $\mathcal{F}$.
4. Compute $\mathbf{z} \leftarrow T^{-1} \cdot \mathbf{y} \in \mathbb{F}_q^n$.
5. Output $(\mathbf{z}, \mathsf{salt})$.

Inversion of the Rainbow central map

Input: The central map $\mathcal{F} = \left(f^{(v_1+1)}, \ldots, f^{(n)}\right)$, a vector $\mathbf{x} \in \mathbb{F}_q^m$
Output: A vector $\mathbf{y} \in \mathbb{F}_q^n$ satisfying $\mathcal{F}(\mathbf{y}) = \mathbf{x}$.

1. Choose random values for the variables $\hat{y}_1, \ldots, \hat{y}_{v_1}$ and substitute these values into the polynomials $f^{(v_1+1)}, \ldots, f^{(v_2)}$.
2. Perform Gaussian elimination on the system

$$f^{(v_1+1)}(\hat{y}_1, \ldots, \hat{y}_{v_1}, y_{v_1+1}, \ldots, y_n) = x_{v_1+1}$$
$$\vdots$$
$$f^{(v_2)}(\hat{y}_1, \ldots, \hat{y}_{v_1}, y_{v_1+1}, \ldots, y_n) = x_{v_2}$$

to obtain the values of the variables $y_{v_1+1}, \ldots, y_{v_2}$, say $\hat{y}_{v_1+1}, \ldots, \hat{y}_{v_2}$.
3. Substitute the values $\hat{y}_{v_1}, \ldots, \hat{y}_{v_2}$ into the polynomials $f^{(v_2+1)}, \ldots, f^{(n)}$.
4. Perform Gaussian elimination on the system

$$f^{(v_2)}(\hat{y}_1, \ldots, \hat{y}_{v_2}, y_{v_2+1}, \ldots, y_n) = x_{v_2+1}$$
$$\vdots$$
$$f^{(n)}(\hat{y}_1, \ldots, \hat{y}_{v_2}, y_{v_2+1}, \ldots, y_n) = x_n$$

to obtain the values of the variables $y_{v_2+1}, \ldots, y_n$, say $\hat{y}_{v_2+1}, \ldots, \hat{y}_n$.
5. Return $\mathbf{y} = (\hat{y}_1, \ldots, \hat{y}_n)$.

Figure 8. Inversion of the Rainbow central map

The main work of the signing algorithm occurs in step 3 which is done using the method described in Figure 8. As all the components $f^{(k)}$ of the central map are homogeneous polynomials of degree two, we can represent them using matrices. Specifically, substituting the first layer of the vinegar variables $\hat{y}_1, \ldots, \hat{y}_{v_1}$ into the first o_1 components of $\mathcal{F}$ is equivalent to computing

$$(\hat{y}_1, \ldots, \hat{y}_{v_1}) \cdot \begin{pmatrix} \alpha_{11}^{(k)} & \cdots & \alpha_{1v_1}^{(k)} \\ & \ddots & \vdots \\ & & \alpha_{v_1 v_1}^{(k)} \end{pmatrix} \cdot \begin{pmatrix} \hat{y}_1 \\ \vdots \\ \hat{y}_{v_1} \end{pmatrix}$$

$$+ (\hat{y}_1, \ldots, \hat{y}_{v_1}) \cdot \begin{pmatrix} \beta_{1v_1+1}^{(k)} & \cdots & \beta_{1v_2}^{(k)} \\ \vdots & & \vdots \\ \beta_{v_1v_1+1}^{(k)} & \cdots & \beta_{v_1v_2}^{(k)} \end{pmatrix} \cdot \begin{pmatrix} y_{v_1+1} \\ \vdots \\ y_{v_2} \end{pmatrix},$$

for $k = v_1 + 1, \ldots, v_2$. With a similar equation occuring for the second layer, namely,

$$(\hat{y}_1, \ldots, \hat{y}_{v_2}) \cdot \begin{pmatrix} \alpha_{11}^{(k)} & \cdots & \alpha_{1v_2}^{(k)} \\ & \ddots & \vdots \\ & & \alpha_{v_1v_2}^{(k)} \end{pmatrix} \cdot \begin{pmatrix} \hat{y}_1 \\ \vdots \\ \hat{y}_{v_2} \end{pmatrix}$$

$$+ (\hat{y}_1, \ldots, \hat{y}_{v_2}) \cdot \begin{pmatrix} \beta_{1v_2+1}^{(k)} & \cdots & \beta_{1n}^{(k)} \\ \vdots & & \vdots \\ \beta_{v_2v_2+1}^{(k)} & \cdots & \beta_{v_2n}^{(k)} \end{pmatrix} \cdot \begin{pmatrix} y_{v_2+1} \\ \vdots \\ y_n \end{pmatrix}$$

for $k = v_2 + 1, \ldots, n$. We call the $2 \cdot (n - v_1)$ matrices in these equations $A^{(k)}, B^{(k)}$. So (abusing notation a bit) we write the equations as $f_k = \hat{\mathbf{y}} \cdot A^{(k)} \cdot \hat{\mathbf{y}}^{\mathsf{T}} + \hat{\mathbf{y}} \cdot A^{(k)} \cdot \mathbf{y}^{\mathsf{T}}$. Recall at any stage we know $\hat{\mathbf{y}}$ and we want to solve the equations for $\mathbf{y}$.

It is clear that signing, given $\mathbf{h} \in \mathbb{F}_q^m$, is a purely algebraic operation over $\mathbb{F}_q$. Thus it can be accomplished in a threshold manner via any LSSS based MPC protocol which evaluates arithmetic circuits over $\mathbb{F}_q$, such as those mentioned earlier. We assume that the private key already exists in secret shared form, i.e. we have sharings $\langle S^{-1} \rangle$, $\langle T^{-1} \rangle$, $\langle \alpha_{i,j}^{(k)} \rangle$ and $\langle \beta_{i,j}^{(k)} \rangle$.

We now look at the signing algorithm's complexity from the point of view of MPC evaluation. We count both the multiplicative depth, as well the number of secure $\mathbb{F}_q$ multiplications needed.

- The first two operations of the signing algorithm come for free, as they are a public hash calculation, followed by the linear opeation $\langle \mathbf{x} \rangle \leftarrow \langle S^{-1} \rangle \cdot \mathbf{h}$.
- We then need to evaluate the map $\mathcal{F}$. This executes in a number of phases.
  - We generate shared values $\langle y_1 \rangle, \ldots, \langle y_{v_1} \rangle$ at random.
  - We then translate the first level of o_2 equations $f_k = \mathbf{x}^{(k)}$ for $k = v_1 + 1, \ldots, v_2$ into a linear system to solve for $\mathbf{y}_1 = (y_{v_1+1}, \ldots, y_{v_2})$. Thus we find an $o_1 \times o_1$ shared matrix $\langle C \rangle$ and a vector $\langle \mathbf{b} \rangle$ such that $C \cdot \mathbf{y}_1 = \mathbf{b}$. To *determine* this system requires two rounds of communication and

$$M_1 = o_1 \cdot \left(\sum_{i=1}^{v_1} i + v_1 + (v_2 - v_1) \cdot v_1 \right)$$
$$= o_1 \cdot (v_1 \cdot (v_1 + 1)/2 + v_1 + o_1 \cdot v_1)$$
$$= o_1 \cdot (o_1 \cdot v_1 + v_1 \cdot (v_1 + 3)/2) = 175032$$

  secure multiplications.
  - Solving our linear system to obtain $\langle \mathbf{y}_1 \rangle$ using our method from Figure 6, which requires three rounds of communication and $M_2 = o_1^3 + o_1^2 = 47952$ secure multiplications.
  - We then repeat with the second layer of the central map, which requires

$$M_3 = o_2 \cdot \left(\sum_{i=1}^{v_2} i + v_2 + (n - v_2) \cdot v_2 \right)$$
$$= o_2 \cdot (v_2 \cdot (v_2 + 1)/2 + v_2 + o_2 \cdot v_2)$$
$$= o_2 \cdot (o_2 \cdot v_2 + v_2 \cdot (v_2 + 3)/2) = 335088.$$

  secure multiplications, and another two rounds of communication.

- We now solve this new linear system to obtain $\langle \mathbf{y}_2 \rangle$ using Figure 6. Again this requires three rounds of communication and $M_4 = M_2$ secure multiplications.
- We then compute $\langle \mathbf{z} \rangle \leftarrow \langle T^{-1} \rangle \cdot \langle \mathbf{y} \rangle$. This requires one round of communication, and due to the special form of T^{-1} it requires $M_5 = v_1 \cdot (o_1 + o_2) + o_1 \cdot o_2 = 6192$ secure multiplications.
- Finally we need to open $\langle \mathbf{z} \rangle$ to obtain the signature in the clear which takes one round of communication.

In summary we require $2+3+2+3+1+1 = 12$ rounds of communication and $M_1+M_2+M_3+M_4+M_5 = 612216$ secure multiplications. Note the last two steps could be computed by opening the last o_2 variables (one round), and then computing $v_1 \cdot o_1 = 2448$ secure multiplications (one round), with another round of communication to open the first $v_1 + o_1$ variables. In practice we expect the extra round to be more costly than the extra multiplications.

If the above algorithm aborts, which can happen if the linear systems have zero determinant, or the random matrices in the trick to solve the linear systems also have zero determinant, then we simply repeat the signing algorithm again. The probability of an abort is bounded by $4/q$. The Rainbow specification uses $q = 2^8$, thus we expect to need to repeat the signing process with probability about 1.5 percent. As mentioned in the introduction a LSSS based MPC protocol can process at least a 250,000 secure multiplications per second over the field $\mathbb{F}_{2^8}$ in the honest majority setting. Thus we expact an implementation of a threshold version of Rainbow to take around three seconds. A *major* disadvantage of this threshold variant of Rainbow is the need to store so much data in secret shared form, namely $\langle S^{-1} \rangle$, $\langle T^{-1} \rangle$, $\langle \alpha_{i,j}^{(k)} \rangle$ and $\langle \beta_{i,j}^{(k)} \rangle$.

6.4 LUOV

Here we present the LUOV signature scheme [4]. As we shall see this is almost entirely made up of low depth algebraic operations, making this scheme a prefect candidate for a threshold variant. The main non-linear component is a map $\mathcal{F} : \mathbb{F}_{2^r}^n \to \mathbb{F}_{2^r}^m$ with components $(f_1, \ldots, f_m)$ where

$$f_k(\mathbf{x}) = \sum_{i=1}^{v} \sum_{j=1}^{n} \alpha_{i,j,k} x_i x_j + \sum_{i=1}^{n} \beta_{i,k} x_i + \gamma_k,$$

with the coefficients $\alpha_{i,j,k}, \beta_{i,k}$ and γ_k being chosen from the field $\mathbb{F}_{2^r}$ by expanding a seed which forms part of the secret key. The integers n, m and v are related by the $v = n - m$. The elements in $\{x_1, \ldots, x_v\}$ are called the *vinegar* variables and that the ones in $\{x_{v+1}, \ldots, x_n\}$ are the *oil* variables. Note that the polynomials $f_1, \ldots, f_m$ contain contain no quadratic terms $x_i \cdot x_j$ with both x_i and x_j oil variables.

The central map $\mathcal{F}$ has to be secret and in order to hide the structue of $\mathcal{F}$ in the public key, one composes $\mathcal{F}$ with an affine map $\mathcal{T} : \mathbb{F}_{2^r}^n \to \mathbb{F}_{2^r}^m$. The public key consisting of composition $\mathcal{P} = \mathcal{F} \circ \mathcal{T} : \mathbb{F}_{2^r}^n \to \mathbb{F}_{2^r}^m$, and the private key being $\mathcal{P}$. At the Level-4 security level (Level-3 is not provided for LUOV) there are two sets of parameters $(r, m, v) = (8, 82, 323)$ and $(64, 61, 302)$.

The LUOV signature generation

Input: The message to be signed μ, and the data $\mathbf{C}, \mathbf{L}, \mathbf{Q}_1$ and $\mathbf{T}$.
Output: A signature $(\mathbf{s}, \mathsf{salt})$ on the message μ.

1. $\mathsf{salt} \leftarrow \{0, 1\}^{16 \cdot 8}$.
2. $\mathbf{h} \leftarrow H(\mu \| \mathsf{0x00} \| \mathsf{salt})$
3. While no solution $\mathbf{s}'$ for the system $\mathcal{F}(\mathbf{s}) = \mathbf{h}$ is found
   (a) $\mathbf{v} \leftarrow \mathbb{F}_{2^r}^v$.
   (b) $\mathbf{RHS} \| \mathbf{LHS} \leftarrow \text{BuildAugmentedMatrix}(\mathbf{C}, \mathbf{L}, \mathbf{Q}_1, \mathbf{T}, \mathbf{h}, \mathbf{v})$
   (c) If $\det(\mathbf{LHS}) \neq 0$ set $\mathbf{o} \leftarrow \mathbf{LHS}^{-1} \cdot \mathbf{RHS}$.
4. $\mathbf{s} \leftarrow \begin{pmatrix} \mathbf{1}_v & -\mathbf{T} \\ \mathbf{0} & \mathbf{1}_m \end{pmatrix} \cdot \begin{pmatrix} \mathbf{v} \\ \mathbf{o} \end{pmatrix}$
5. Return $(\mathbf{s}, \mathsf{salt})$.

Figure 9. The LUOV signature generation

19

The LUOV public and private keys are in practice expanded from a random seed to define the actual data defining the various maps. However, for our threshold variant we assume this expansion has already happened and we have the following data values $\mathbf{C} \in \mathbb{F}_{2^r}^m$, $\mathbf{L} \in \mathbb{F}_{2^r}^{m \times n}$, $\mathbf{Q}_1 \in \mathbb{F}_{2^r}^{m \times (\frac{v(v+1)}{2} + v \cdot m)}$, and $\mathbf{T} \in \mathbb{F}_{2^r}^{v \times m}$, where $\mathbf{C}$, $\mathbf{L}$, and $\mathbf{Q}_1$ are public values and the matrix $\mathbf{T}$ is a secret parameter. In our threshold variant the parameter $\mathbf{T}$ will be held in secret shared form $\langle \mathbf{T} \rangle$. There is another matrix $\mathbf{Q}_2$, but that will not concern us as it is only related to verification. The signing algorithm is given in Figure 9, and makes use of an auxillary algorithm given in Figure 10 and a hash function $H : \{0,1\}^* \longrightarrow \mathbb{F}_{2^r}^m$ The auxiliary algorithm builds a linear system $\mathbf{LHS} \cdot \mathbf{o} = \mathbf{RHS}$ which we solve to obtain the oil variables.
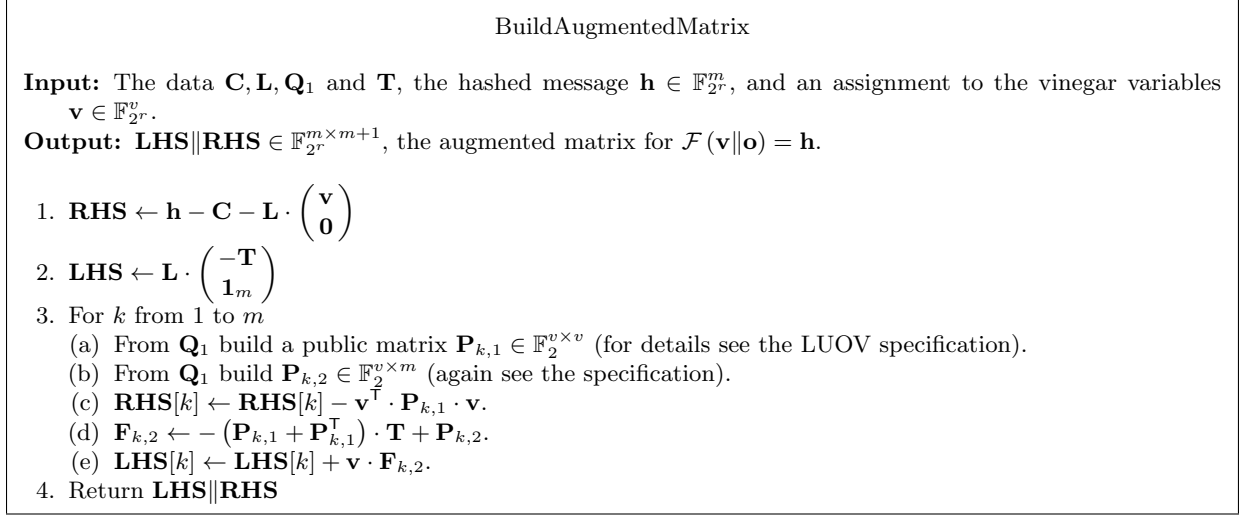

BuildAugmentedMatrix

Input: The data $\mathbf{C}, \mathbf{L}, \mathbf{Q}_1$ and $\mathbf{T}$, the hashed message $\mathbf{h} \in \mathbb{F}_{2^r}^m$, and an assignment to the vinegar variables $\mathbf{v} \in \mathbb{F}_{2^r}^v$.
Output: $\mathbf{LHS}\|\mathbf{RHS} \in \mathbb{F}_{2^r}^{m \times m+1}$, the augmented matrix for $\mathcal{F}(\mathbf{v}\|\mathbf{o}) = \mathbf{h}$.

1. $\mathbf{RHS} \leftarrow \mathbf{h} - \mathbf{C} - \mathbf{L} \cdot \begin{pmatrix} \mathbf{v} \\ \mathbf{0} \end{pmatrix}$
2. $\mathbf{LHS} \leftarrow \mathbf{L} \cdot \begin{pmatrix} -\mathbf{T} \\ \mathbf{1}_m \end{pmatrix}$
3. For k from 1 to m
   (a) From $\mathbf{Q}_1$ build a public matrix $\mathbf{P}_{k,1} \in \mathbb{F}_2^{v \times v}$ (for details see the LUOV specification).
   (b) From $\mathbf{Q}_1$ build $\mathbf{P}_{k,2} \in \mathbb{F}_2^{v \times m}$ (again see the specification).
   (c) $\mathbf{RHS}[k] \leftarrow \mathbf{RHS}[k] - \mathbf{v}^\mathsf{T} \cdot \mathbf{P}_{k,1} \cdot \mathbf{v}$.
   (d) $\mathbf{F}_{k,2} \leftarrow -\left(\mathbf{P}_{k,1} + \mathbf{P}_{k,1}^\mathsf{T}\right) \cdot \mathbf{T} + \mathbf{P}_{k,2}$.
   (e) $\mathbf{LHS}[k] \leftarrow \mathbf{LHS}[k] + \mathbf{v} \cdot \mathbf{F}_{k,2}$.
4. Return $\mathbf{LHS}\|\mathbf{RHS}$

Figure 10. BuildAugmentedMatrix

We now examine the above algorithm from the point of view of how one could implement it in a threshold manner given a generic MPC functionality for arithmetic circuits over $\mathbb{F}_{2^r}$. We assume that the secret matrix $\mathbf{T}$ is presented in secret shared form $\langle \mathbf{T} \rangle$. First note that the hash function is only called to compute the hash digest, which does not involve any shared input.

In the main while loop we assume the vinegar variables are generated in a shared manner in secret shared form $\langle \mathbf{v} \rangle$. Thus the main call to the BuiltAugmentedMatrix routine has two secret shared input $\langle \mathbf{T} \rangle$ and $\langle \mathbf{v} \rangle$, with the other values being public. The key lines in this algorithm then requiring secure multiplications are lines 3c and 3e to compute $\langle \mathbf{v} \rangle^\mathsf{T} \cdot \mathbf{P}_{k,1} \cdot \langle \mathbf{v} \rangle$ and $\langle \mathbf{v} \rangle \cdot \mathbf{F}_{k,2}$ respectively. The first of these takes v secure multiplications, whereas the latter requires $v \cdot m$ secure multiplications. Giving a total of $v \cdot m \cdot (m+1)$ secure multiplications in total, which can be performed in parallel in one round of communication.

Solving the nonlinear system in line 3c is done using the method in Figure 6, which requires three rounds of interaction and $m^3 + m^2$ secure multiplications. Note the probability that this procedure fails is roughly 2^{-r+1}, which can be essentially ignored for the parameter set with $r = 64$ and is under one percent for the parameter set with $r = 8$. But if it does fail, then we simply repeat the signing algorithm with new shared vinegar variables.

We then need to compute the matrix multiplication $\langle \mathbf{T} \rangle \cdot \langle \mathbf{o} \rangle$. However, note that we can save some secure multiplications by opening the oil variables $\mathbf{o}$ after the matrix inversion (since they are going to be released in any case in the clear). This will require anyway a rouond of interaction, but we are then able to save the $v \cdot m$ secure multiplications required to multiply $\mathbf{T}$ by $\mathbf{o}$, since that operation then becomes a linear operation. Finally, we open the resulting shared signature in order to transmit it in the clear. This requires one round of interaction.

Thus the overall cost of LUOV signature algorithm is $1 + 3 + 1 + 1 = 6$ rounds of interaction and $(m^3 + m^2) + v \cdot m \cdot (m+1)$ secure multiplications. Choosing the Level-4 parameter set with $(r, m, v) = (8, 82, 323)$ this gives a total of 2756430 secure multiplications. Whereas for the parameter set $(r, m, v) = (64, 61, 302)$

this gives us 1372866 secure multiplications. In the former case, where arithmetic is over $\mathbb{F}_{2^8}$ and we expect to perform 250,000 secure multiplications per second, signing will take about 10 seconds. In the latter case, where arithmetic is over $\mathbb{F}_{2^{64}}$ and we expect to perform 1,000,000 secure multiplications per second, signing will take about 1.3 seconds. Another advantage of LUOV is that our threshold variant requires less storage of secret key material. We only need to store $\langle \mathbf{T} \rangle$ in secret shared form

Acknowledgments

References

1. Albrecht, M.R., Rechberger, C., Schneider, T., Tiessen, T., Zohner, M.: Ciphers for MPC and FHE. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part I. LNCS, vol. 9056, pp. 430–454. Springer, Heidelberg (Apr 2015)
2. Barthe, G., Belaïd, S., Espitau, T., Fouque, P.A., Grégoire, B., Rossi, M., Tibouchi, M.: Masking the GLP lattice-based signature scheme at any order. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part II. LNCS, vol. 10821, pp. 354–384. Springer, Heidelberg (Apr / May 2018)
3. Bernstein, D.J., Hopwood, D., Hülsing, A., Lange, T., Niederhagen, R., Papachristodoulou, L., Schneider, M., Schwabe, P., Wilcox-O'Hearn, Z.: SPHINCS: Practical stateless hash-based signatures. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part I. LNCS, vol. 9056, pp. 368–397. Springer, Heidelberg (Apr 2015)
4. Beullens, W., Preneel, B., Szepieniec, A., Vercauteren, F.: LUOV (2019), Submission to NIST PQC "competition" Round-2.
5. Bindel, N., Akleylek, S., Alkim, E., Barreto, P.S., Buchmann, J., Eaton, E., Gutoski, G., Kramer, J., Longa, P., Polat, H., Ricardini, J.E., Zanon, G.: Lattice-based digital signature scheme qTESLA (2019), Submission to NIST PQC "competition" Round-2.
6. Casanova, A., Faugère, J.C., Patarin, G.M.R.J., Perret, L., Ryckeghem, J.: GeMSS: A great multivariate short signature (2019), Submission to NIST PQC "competition" Round-2.
7. Catrina, O., de Hoogh, S.: Improved primitives for secure multiparty integer computation. In: Garay, J.A., Prisco, R.D. (eds.) SCN 10. LNCS, vol. 6280, pp. 182–199. Springer, Heidelberg (Sep 2010)
8. Chida, K., Genkin, D., Hamada, K., Ikarashi, D., Kikuchi, R., Lindell, Y., Nof, A.: Fast large-scale honest-majority MPC for malicious adversaries. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part III. LNCS, vol. 10993, pp. 34–64. Springer, Heidelberg (Aug 2018)
9. Cramer, R., Damgård, I., Escudero, D., Scholl, P., Xing, C.: SPD $\mathbb{Z}_{2^k}$: Efficient MPC mod 2^k for dishonest majority. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part II. LNCS, vol. 10992, pp. 769–798. Springer, Heidelberg (Aug 2018)
10. Damgård, I., Fitzi, M., Kiltz, E., Nielsen, J.B., Toft, T.: Unconditionally secure constant-rounds multi-party computation for equality, comparison, bits and exponentiation. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 285–304. Springer, Heidelberg (Mar 2006)
11. Damgård, I., Koprowski, M.: Practical threshold RSA signatures without a trusted dealer. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 152–165. Springer, Heidelberg (May 2001)
12. Damgård, I., Pastro, V., Smart, N.P., Zakarias, S.: Multiparty computation from somewhat homomorphic encryption. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 643–662. Springer, Heidelberg (Aug 2012)
13. Desmedt, Y., Frankel, Y.: Threshold cryptosystems. In: Brassard, G. (ed.) CRYPTO'89. LNCS, vol. 435, pp. 307–315. Springer, Heidelberg (Aug 1990)
14. Ding, J., Chen, M.S., Petzoldt, A., Schmidt, D., Yang, B.Y.: Rainbow (2019), Submission to NIST PQC "competition" Round-2.

15. Doerner, J., Kondi, Y., Lee, E., shelat, a.: Secure two-party threshold ECDSA from ECDSA assumptions. In: 2018 IEEE Symposium on Security and Privacy. pp. 980–997. IEEE Computer Society Press (May 2018)

16. Gennaro, R., Jarecki, S., Krawczyk, H., Rabin, T.: Robust threshold DSS signatures. In: Maurer, U.M. (ed.) EUROCRYPT'96. LNCS, vol. 1070, pp. 354–371. Springer, Heidelberg (May 1996)

17. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Ladner, R.E., Dwork, C. (eds.) 40th ACM STOC. pp. 197–206. ACM Press (May 2008)

18. Gentry, C., Szydlo, M.: Cryptanalysis of the revised NTRU signature scheme. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 299–320. Springer, Heidelberg (Apr / May 2002)

19. Giacomelli, I., Madsen, J., Orlandi, C.: ZKBoo: Faster zero-knowledge for boolean circuits. In: Holz, T., Savage, S. (eds.) USENIX Security 2016. pp. 1069–1083. USENIX Association (Aug 2016)

20. Goldreich, O.: Two remarks concerning the Goldwasser-Micali-Rivest signature scheme. In: Odlyzko, A.M. (ed.) CRYPTO'86. LNCS, vol. 263, pp. 104–110. Springer, Heidelberg (Aug 1987)

21. Goldreich, O.: Foundations of Cryptography: Basic Applications, vol. 2. Cambridge University Press, Cambridge, UK (2004)

22. Hazay, C., Scholl, P., Soria-Vazquez, E.: Low cost constant round MPC combining BMR and oblivious transfer. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017, Part I. LNCS, vol. 10624, pp. 598–628. Springer, Heidelberg (Dec 2017)

23. Hoffstein, J., Howgrave-Graham, N., Pipher, J., Silverman, J.H., Whyte, W.: NTRUSIGN: digital signatures using the NTRU lattice. In: Joye, M. (ed.) Topics in Cryptology - CT-RSA 2003, The Cryptographers' Track at the RSA Conference 2003, San Francisco, CA, USA, April 13-17, 2003, Proceedings. Lecture Notes in Computer Science, vol. 2612, pp. 122–140. Springer (2003), `https://doi.org/10.1007/3-540-36563-X_9`

24. Hulsing, A., Bernstein, D.J., Dobraunig, C., Eichlseder, M., Fluhrer, S., Gazdag, S.L., Kampanakis, P., Kolbl, S., Lange, T., Lauridsen, M.M., Mendel, F., Niederhagen, R., Rechberger, C., Rijneveld, J., Schwabe, P.: SPHINCS+ (2019), Submission to NIST PQC "competition" Round-2.

25. Ishai, Y., Kushilevitz, E., Ostrovsky, R., Sahai, A.: Zero-knowledge from secure multiparty computation. In: Johnson, D.S., Feige, U. (eds.) 39th ACM STOC. pp. 21–30. ACM Press (Jun 2007)

26. Karmakar, A., Roy, S.S., Vercauteren, F., Verbauwhede, I.: Pushing the speed limit of constant-time discrete gaussian sampling. A case study on the falcon signature scheme. In: Proceedings of the 56th Annual Design Automation Conference 2019, DAC 2019, Las Vegas, NV, USA, June 02-06, 2019. pp. 88:1–88:6. ACM (2019), `https://doi.org/10.1145/3316781.3317887`

27. Kipnis, A., Patarin, J., Goubin, L.: Unbalanced oil and vinegar signature schemes. In: Stern, J. (ed.) EUROCRYPT'99. LNCS, vol. 1592, pp. 206–222. Springer, Heidelberg (May 1999)

28. Kraitsberg, M., Lindell, Y., Osheter, V., Smart, N.P., Alaoui, Y.T.: Adding distributed decryption and key generation to a Ring-LWE based CCA encryption scheme. IACR Cryptology ePrint Archive 2018, 1034 (2018), `https://eprint.iacr.org/2018/1034`

29. Lamport, L.: Constructing digital signatures from a one-way function. Technical Report SRI-CSL-98, SRI International Computer Science Laboratory (Oct 1979)

30. Larraia, E., Orsini, E., Smart, N.P.: Dishonest majority multi-party computation for binary circuits. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part II. LNCS, vol. 8617, pp. 495–512. Springer, Heidelberg (Aug 2014)

31. Lindell, Y.: Fast secure two-party ECDSA signing. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part II. LNCS, vol. 10402, pp. 613–644. Springer, Heidelberg (Aug 2017)

32. Lindell, Y., Nof, A.: Fast secure multiparty ECDSA with practical distributed key generation and applications to cryptocurrency custody. In: Lie, D., Mannan, M., Backes, M., Wang, X. (eds.) ACM CCS 2018. pp. 1837–1854. ACM Press (Oct 2018)

33. Lindell, Y., Nof, A., Ranellucci, S.: Fast secure multiparty ECDSA with practical distributed key generation and applications to cryptocurrency custody. IACR Cryptology ePrint Archive 2018, 987 (2018), `https://eprint.iacr.org/2018/987`

34. Lyubashevsky, V.: Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 598–616. Springer, Heidelberg (Dec 2009)

35. Lyubashevsky, V., Ducas, L., Kiltz, E., Lepoint, T., Schwabe, P., Seiler, G., Stehle, D.: Crystals-dilithium (2019), Submission to NIST PQC "competition" Round-2.

36. MacKenzie, P.D., Reiter, M.K.: Two-party generation of DSA signatures. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 137–154. Springer, Heidelberg (Aug 2001)

37. Matsumoto, T., Imai, H.: Public quadratic polynominal-tuples for efficient signature-verification and message-encryption. In: Günther, C.G. (ed.) EUROCRYPT'88. LNCS, vol. 330, pp. 419–453. Springer, Heidelberg (May 1988)

38. Merkle, R.C.: A certified digital signature. In: Brassard, G. (ed.) CRYPTO'89. LNCS, vol. 435, pp. 218–238. Springer, Heidelberg (Aug 1990)
39. Mohassel, P., Rosulek, M., Zhang, Y.: Fast and secure three-party computation: The garbled circuit approach. In: Ray, I., Li, N., Kruegel, C. (eds.) ACM CCS 2015. pp. 591–602. ACM Press (Oct 2015)
40. Nielsen, J.B., Nordholt, P.S., Orlandi, C., Burra, S.S.: A new approach to practical active-secure two-party computation. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 681–700. Springer, Heidelberg (Aug 2012)
41. Patarin, J.: The oil and vinegar signature scheme (1997), presentation at the Dagstuhl Workshop on Cryptography
42. Prest, T., Fouque, P.A., Hoffstein, J., Kirchner, P., Lyubashevsky, V., Pornin, T., Ricosset, T., Seiler, G., Whyte, W., Zhang, Z.: Falcon: Fast-fourier lattice-based compact signatures over NTRU (2019), Submission to NIST PQC "competition" Round-2.
43. Reyzin, L., Reyzin, N.: Better than BiBa: Short one-time signatures with fast signing and verifying. In: Batten, L.M., Seberry, J. (eds.) ACISP 02. LNCS, vol. 2384, pp. 144–153. Springer, Heidelberg (Jul 2002)
44. Rotaru, D., Wood, T.: Marbled circuits: Mixing arithmetic and boolean circuits with active security. IACR Cryptology ePrint Archive 2019, 207 (2019), https://eprint.iacr.org/2019/207
45. Sakumoto, K., Shirai, T., Hiwatari, H.: Public-key identification schemes based on multivariate quadratic polynomials. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 706–723. Springer, Heidelberg (Aug 2011)
46. Samardjiska, S., Chen, M.S., Hulsing, A., Rijneveld, J., Schwabe, P.: MQDSS (2019), Submission to NIST PQC "competition" Round-2.
47. Shoup, V.: Practical threshold signatures. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 207–220. Springer, Heidelberg (May 2000)
48. Smart, N.P., Wood, T.: Error detection in monotone span programs with application to communication-efficient multi-party computation. In: Matsui, M. (ed.) CT-RSA 2019. LNCS, vol. 11405, pp. 210–229. Springer, Heidelberg (Mar 2019)
49. Wang, X., Ranellucci, S., Katz, J.: Global-scale secure multiparty computation. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) ACM CCS 2017. pp. 39–56. ACM Press (Oct / Nov 2017)
50. Zaverucha, G., Chase, M., Derler, D., Goldfeder, S., Orlandi, C., Ramacher, S., Rechberger, C., Slamanig, D.: The picnic signature scheme (2019), Submission to NIST PQC "competition" Round-2.